SYSTEM-DRIVEN RESEARCH: LEGITIMATE EXPERIMENTAL DESIGN FOR BIOLOGICAL/BIOMEDICAL RESEARCH

by

Eve A. Roberts

Submitted in partial fulfillment of the requirements

for the degree of Doctor of Philosophy

at

Dalhousie University

Halifax, Nova Scotia

March 2015

DEDICATION PAGE

Studies serve for delight, for ornament, and for ability. Their chief use for delight, is in privateness and retiring; for ornament, is in discourse; and for ability, is in the judgement and disposition of business. … To spend too much time in studies is sloth; to use them too much for ornament is affectation; to make judgement wholly by their rules is the humour of a scholar. They perfect nature, and are perfected by experience, for natural abilities are like natural plants that need proyning by study; and studies themselves do give forth directions too much at large, except they be bounded in by experience.

—Sir Francis Bacon, *Of Studies*

I dedicate this dissertation, with respect,

to Professor Dame Sheila Sherlock and her husband, Professor Geraint James.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

ABSTRACT

Exciting new high-throughput methodologies permit innovative experimental designs for current biological and biomedical research. Much of this research is "non-hypothesis-driven": its experiments are not designed around a specific testable hypothesis. This challenge to the hegemony of the hypothesis poses important problems. For philosophers, it yet again questions the hypothetico-deductive method (HDM), the established hypothesis-driven strategy for science. For scientists, it raises practical issues about how to evaluate non-hypothesis-driven experimental design and its data, including what research deserves funding. Overall, it asks what qualifies as *worthy* science.

I define hypothesis-driven research as research where a well-articulated hypothesis immediately governs the experimental design. Several different styles of non-hypothesis-driven research exist. With system-driven research, the biological entity under investigation, demarcated as a system, governs the experimental design. My thesis is that system-driven research, which typically involves high-throughput methodologies collectively known as "omics", constitutes scientific research as worthy as hypothesis-driven research. The HDM derives its cogency from its logical structure. I argue that system-driven research is supported by a similarly powerful scientific epistemology, the Omics Experimental Strategy (OES). I aim to show how the OES operates. I anchor this analysis in a subdiscipline of proteomics, metalloproteomics, which I helped to develop in the early 2000s.

Biological research attempting to examine a complex biological system directly reflects a distinctive epistemological orientation. The OES attempts to be fully inclusive. It proceeds without prejudice of expectation related to a hypothesis and may reveal 'surprising' findings, otherwise inaccessible to a chain of hypotheses. With the OES, excellent experimental design, technical adequacy and adherence to standards of expert performance in the field are crucial for ensuring that data are reliable. Pattern-detection is central to data evaluation. For knowledge production, contextualization of the findings in the system itself is critical. With the HDM, the findings are evaluated in terms of the hypothesis; with the OES, the findings are evaluated in the context of the system.

System-driven research constitutes proper scientific research. It operates alongside hypothesis-driven research, not merely preparatory to it. Importantly, system-driven research offers innovative, unique ways to understand biological systems, and very possibly the methodology of science itself.

LIST OF ABBREVIATIONS USED

Acronym or technical term	Stands for:	Comments:	First mention (Section):
2DE	Two-dimensional electrophoresis		2.3
A, T, G, C	Adenine, thymine, guanine, cytosine	Purines (A, G) and pyrimidines (T, C) of DNA	8.2.3
Ah	Aromatic hydrocarbon	Part of designation of the Ah receptor (AhR, also 'AHR') which regulates induction of cytochrome P450 IA1 and certain other drug-metabolizing enzymes	3.3.3
AIPP	ATPase interacting PDZ protein 1		8.1.3
ATOX1	Antioxidant protein 1	Metallochaperone: an intracellular protein which ferries a metal from one intracellular site to another	2.5.1
ATP7A	[gene designation]	Structural gene for Menkes ATPase, a human metal-transporting P-type ATPase	2.4.2
ATP7B	[gene designation]	Structural gene for Wilson ATPase, another human metal-transporting P-type ATPase	2.4.2
ATP7B	[protein designation]	Wilson ATPase	2.4.2
ATPase	Adenosine triphosphatase		2.4.2
Atx1	Antioxidant protein 1 (yeast)	Same as ATOX1 except that it is the version found in yeast	5.3
C. elegans	*Caenorhabditis elegans*	Transparent roundworm serves as a	2, fn 14

		model organism in genetics.	
C3H	[mouse strain designation]	Background strain for tx-j	2.4.3
CAP	Complete, Accurate, Permanent	Meant to describe a proteome but seems to have fallen out of use	6.4.1 fn 132
cDNA	Complementary DNA	Structural gene DNA without the intervening non-coding regions	5.3
CF	Cystic fibrosis		4.1
CFTR	Cystic fibrosis transmembrane conductance regulator	Gene designation for gene whose mutations cause CF	4.1, fn 64
CIHR	Canadian Institutes of Health Research		2.5.2, fn 18
COMMD1	[gene designation]	Structural gene for COMMD1	2.5.1
COMMD1	Copper metabolism Murr1 domain 1 protein	'Murr1' was a previous name for this protein.	2.5.1
CTR1	Copper transporter 1	Human version is called hCTR1.	2.5.2.2
DCF	Dichlorofluorescein		7.2.1
DCFH-DA	2′,7′-dichlorodihydrofluorescein diacetate		7.2.1
DCTN4	Same as dynactin	Intracellular transporter	8.4.3.3
D-N	Deductive-nomological		4.3.2
DNA	Deoxyribonucleic acid		2.2.2
E. coli	*Escherichia coli*	Gram-negative bacterium	5.3
ESI	Electrospray ionization		2.3
EtOH	Ethanol		8.3.1, fn 186
FKBP52	Immunophilin	Previously known as 'FK-binding protein 52' where "FK" refers to "FK506", an immunosuppressant drug	8.1.3
Golgi	[not an acronym]	Cellular organelle (Golgi apparatus) involved in protein manufacture, named after anatomist	6.4

GRP78	Glucose-regulated protein 78 kD		2.5.3.2
GWAS	Genome-wide association study		9.4
HDM	Hypothetico-deductive method		1
HepG2	[cell line designation]	Continuous human liver cell (hepatocyte) line	2.4.3
HIF-1	Hypoxia-inducible factor-1		8.1.3
HPLC	High performance liquid chromatography		2.3
HT-MS/MS	High-throughput tandem mass spectrometry		2.5.3.2
HUGO	Human Genome Organization		1.3.2
HUPO	Human Proteome Organization		1.3.2
ICAT	isotope-coded affinity tag		2.5.2.2
ICP-MS	Inductively coupled plasma mass spectroscopy		2.5.3.1
IMAC	Immobilized metal affinity column		2.3 fn 23
I-S	Inductive-statistical		4.3.2
JPEG	Joint Photographic Experts Group	Standard compression method for digital images	8.2.3
LA	Laser ablation		2.5.3.1
LAN	Local area network		2.2.2
LC	Liquid chromatography		2.3
LC-MS/MS	Liquid chromatography followed by tandem mass spectrometry		6.2.3
LEA	Long-Evans agouti	Rat strain from which LEC is derived (‘agouti’ indicates banded coat-colour).	2.5.2.3
LEC	Long-Evans cinnamon	Rat model for Wilson disease (‘cinnamon’ describes light brown coat-colour)	2.4.3
m/z	Mass-to-charge ratio		2.3
MALDI	Matrix-assisted laser desorption/ionization		2.3

MEDNIK	Mental retardation, Enteropathy, Deafness, Neuropathy, Ichthyosis, Keratodermia	= Erythrokeratodermia Variabilis type 3 (EKV3)	2.4.2 fn 29
MIAME	Minimum Information About a Microarray Experiment		6.3
MIAPE	Minimum Information About a Proteomics Experiment		6.3
MINSEQE	Minimum Information about a high-throughput Sequencing Experiment		6.3
MIRAGE	Metal Isotope native RadioAutography in Gel Electrophoresis		2.5.3.1
MS	Mass spectrometry		2.3 fn 26
Mtch2	Mitochondrial carrier homologue 2 (protein)	One of the proteins of the mitochondrial outer membrane	5.3
MTTP	Microsomal triglyceride transfer protein		7.4.2
MXCXXC	[not an acronym]	Shorthand for a primary sequence of amino acids in a protein M =methionine C = cysteine X = space holder for any other amino acid	5.3
MXSXXC, MXSXXS	[not an acronym]	See above—S = serine	5.3
NF-κB	Nuclear factor-κB		8.1.3
NIDDK	National Institute of Diabetes and Digestive and Kidney Diseases	An institute within the NIH	5.2.5.2, fn 99
NIH	National Institutes of Health (USA)		2.5.2, fn 38
OES	Omics Experimental Strategy		1.3.1
OX-PHOS	Oxidative phosphorylation	Chain of enzymes in the mitochondrion which are responsible for energy production	2.4.3
p97/VCP	Valosin-containing protein		8.1.3, fn 168
PAGE	Polyacrylamide gel electrophoresis		2.4.3

PC12	[cell line designation]	rat adrenal pheochromocytoma-derived dopamine-producing cell line	2.5.3.1
PCR	Polymerase chain reaction		2.5.1
PDI	Protein disulfide isomerase		2.5.1
PINA	Pineal night-associated protein	*ATP7B* gene product by alternate splicing	6.1.1, fn 111
PIXE	Proton-induced X-ray emission		2.5.3.1
PLZF	Zinc finger promyelocytic leukemia protein		8.4.3.3
PRIDE	Proteomics Identifications database		6.3
PSI	Proteomics Standards Initiative		6.3
Qq	Quadripole		2.3 fn 26
RB	Retinoblastoma		4.4
RNA	Ribonucleic acid		1.2
S. cerevisiae	*Saccharomyces cerevisiae*	Brewer's yeast (also used for bread)	5.3
SDS-PAGE	Sodium dodecyl sulfate polyacrylamide gel electrophoresis		6.2.2
SELDI	Surface-enhanced laser desorption/ionization		2.5.2
SERCA1	Sarco(endo)plasmic reticulum Ca^{2+} ATPase 1		2.4.2
SILAC	Stable isotope labelling by amino acids in cell culture		2.3
SLC31A1	= hCTR1	SLC nomenclature is a different naming system for enzymes.	8.4.3.3
SOD1	Superoxide dismutase 1		2.5.1
SRXRF	Synchrotron radiation X-ray fluorescence		2.5.3.1
STRING	Search Tool for the Retrieval of Interacting Genes/Proteins		8.4.3.3
SXRF	Synchrotron X-Ray fluorescence	Another name for SRXRF	2.5.3.1
T/F	True/false		7.3.2
TEM	Transmission electron micrograph	Technical name for what we think of as a	2.4.3

		standard ‘electron micrograph’	
TOF	Time-of-flight		2.3 fn 26
tRNA	Transfer RNA		5.2.3.2
tx	toxic milk	Two spontaneous mouse models for Wilson disease: one arising in the C3H mouse strain and the other in the C57/Bl strain	2.4.3
tx-j	toxic milk mouse, Jackson lab variety	Displays homozygous point mutation of *Atp7b* in C3H strain	2.4.3
V-R	Video-response		8.2.2, fn 170
XANES	X-ray Absorption Near Edge Structure		2.5.3.1
XIAP	X-linked inhibitor of apoptosis		2.5.1

fn = footnote

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the support and encouragement of numerous people. First of all, I would like to thank my supervisors, Letitia Meynell and Gordon McOuat, for seeing merit in this project, offering sage advice all along, challenging my half-baked ideas, and providing solid encouragement whenever I felt faint-hearted. I also want to thank Richmond Campbell who supervised my MA and launched me in this PhD project: I am most grateful for his abiding support and trenchant, engaged critique of my work. I have had the great benefit of a distinguished Thesis Committee—Peter Schotch and Richmond Campbell in the Department of Philosophy, and Ford Doolittle of the Department of Biochemistry & Molecular Biology—and I am very grateful for their oversight and numerous enjoyable discussions. I would also like to thank my mentors in Toronto. Without the energies of my colleague, Bibudhendra Sarkar, to invent metalloproteomics and invite me to collaborate in that adventure, this dissertation would never have happened. He has listened to and critiqued my scientific and philosophical ideas about metalloproteomics throughout this entire philosophical research project. Robert Loewen has provided non-scientific balance throughout my PhD programme; his advice 'to commit to [an idea] and see where it takes you' was the *leitmotif* for this dissertation.

I would like to acknowledge funding sources for the laboratory-based research described in this dissertation. It included the Coady Family Fund for Hepatic Research, Canadian Institutes of Health Research (CIHR), and the Canadian Liver Foundation (CLF). Other sources are detailed in our publications.

Portions of Chapter 2 have appeared in scientific publications in 2012 and 2014; portions of Chapter 3 were presented to the Off-year Workshop of the International Society for the History, Philosophy, and Social Studies of Biology in Seattle in 2012; portions of Chapter 4 were presented to the biennial meeting of the International Society for the History, Philosophy, and Social Studies of Biology in Montpellier in 2013; and various iterations of key ideas from the dissertation were presented to the International Liver Study Group at annual meetings between 2011 and 2014.

Finally I am delighted to take this opportunity to say a huge thank-you to my family. My children, Graham and Reed, have encouraged me about this project from its earliest days as speculation about 'this interesting philosophical problem' arising from my laboratory research, and they have been constantly supportive. My mother, though of advanced age, continues to be positively excited about this project—an excitement which has been extremely compelling. Most importantly, my husband Ian has been my champion throughout: doubting, dissecting, extending, and complicating my ideas as fast and furiously as I could come up with them, but always affirming the worthiness of this intellectual endeavour. I owe him the greatest thanks.

Chapter 1 INTRODUCTION

For most biological or biomedical scientists, the critical question about a proposed research project or the results of an ongoing research project is: "What is your hypothesis?" This is often the first question a trainee or a colleague gets asked. Such an approach is hardly surprising since most biomedical scientists have been schooled in the firm belief that good and worthy research is organized around a hypothesis. Indeed, in both the brief summary of a grant proposal seeking research funding and the lengthy detailed description of the proposed research, biomedical scientists typically put the stated hypothesis in bold type so that the reviewers won't miss it! However, new high-throughput methodologies permit innovative experimental designs. Our rapidly-advancing knowledge about cellular and structural biology plus our constantly expanding capability to store and manipulate very large data sets via sophisticated computer-based technologies is changing how we do biological and biomedical research. Much biological and biomedical research nowadays is "non-hypothesis-driven": it is not designed around a hypothesis.

This challenge to the hegemony of the hypothesis is highly problematic for scientists and philosophers. For philosophers, it yet again upsets standard 20th-century concepts of how science 'works'. Namely, it calls into question yet again the hypothetico-deductive method (HDM), the established hypothesis-driven strategy for experimental science, worked out in exquisite detail by Carl Hempel, though not invented solely by him. The hypothetico-deductive method derives its cogency from the certainty characteristic of deductive inference and the plausibility which is typical of abductive inference. For scientists, this challenge to the 'hegemony of hypothesis' raises practical issues about how to evaluate experimental design and the resulting data, and how to determine what research deserves funding. In other words, it raises the question of what is worthy science. My objective in this dissertation is to examine how 'non-hypothesis-driven' biological/biomedical research functions effectively as scientific research. The purpose of this introductory chapter is two-fold: to review some

background concepts critical to proceeding with my exercise in scientific epistemology and to outline how that analysis proceeds.

1.1 Varieties of non-hypothesis-driven research

1.1.1 System-driven research

On first inspection, two forms of non-hypothesis-driven research are currently attracting a lot of attention. In this age of 'big data', both of these forms of research are generally grouped together under the label of "data-driven research". This is certainly the thrust of a Microsoft Inc. think tank where the fourth paradigm of scientific research is described as "data-intensive": computer-based and embracing comprehensive data capture, curation, analysis and sharing (Bell, 2009, xv). It is the usage typical of philosophers of science (O'Malley and Soyer, 2012, 58; Leonelli, 2012, 1). I want to make a key distinction here. When the research project is designed to interrogate a collection of data already in hand in order to find new information 'hidden' there, I call this effort "data-driven" research. The collected data at hand is literally at the centre of the research effort. Examples of such research include examining a large clinical database for new disease associations extraneous to the initial purpose for assembling the database, working out population health patterns by combining information from various administrative data sets, and studying social media (for example, tweets and on-line comment-postings) to address an emerging social behaviour. This kind of research is commonly referred to as 'mining a database'. Data-driven research, though highly interesting and emerging as an exciting new research strategy for 'big data' on a global scale, is in fact not the subject of this dissertation. My focus is on high-throughput research methods. With these methods, the biological system[1] in its attendant

[1]For the purposes of my discussion, I will define a system as a delimited collection of components whose structures and function are somehow integrated (see Section 5.1 for discussion and development of this formulation). Since I am dealing with biological systems, the components are those ordinarily found in a biological entity. For example, if I chose a cell as the system of interest, its subcellular organelles (nucleus, mitochondria, lysosomes, etc.) could be the components. High-throughput research methods are those which can perform numerous analyses simultaneously or examine a very large number of specimens rapidly and efficiently. The contrast between arduous single analysis and high-throughput methodology is illustrated by the difference between marking an essay exam for a class of 1000 students or giving a multiple-choice test whose answer sheet is scanned and marked by computerized technology.

complexity is addressed directly. Data sets about such a system assembled prospectively are critically important, for example, for systems biology. In other words, the research is system-driven. Data sets about the system under investigation are assembled prospectively.[2]

Not surprisingly, system-driven research is critically important for the emerging research discipline of systems biology. High-throughput analytical techniques and new technologies for manipulating the constituents and machinery of cellular genetic apparatus have made the experimental project characteristic of systems biology possible, as well as permitting research in other broad disciplines. This type of research generates massive amounts of data, not merely because having lots of data conveys a sense of mastery of the problem under investigation but because the system under investigation generally turns out to be very complicated! Bioinformatics manages and facilitates analysis of these data. One desired outcome of systems biology research is to generate quantitative models of biological systems, perhaps tipping biology toward the perceived legitimacy of the physical sciences. How systems biology uses data sets generated by the high-throughput methods is also outside the scope of this dissertation, but the expectation that reliable data characterizing such systems in great detail will be required before systems biology's goals can be achieved is entirely pertinent.

I hold that biological research which attempts to examine a biological system as such, in its complexity, reflects a distinctive epistemological orientation. Instead of being organized around a hypothesis or an existing data set, the research project is organized around a system. This system[3] can be defined as broadly or as narrowly as the

---

[2] When a data set is the product of a research effort, we say it is prospective research and the data set is itself prospective. When the research effort examines a data set already collected, we say it is retrospective research. The contrast is perhaps most easily appreciated with clinical research. Suppose I want to investigate whether a certain blood test is informative for diagnosing a disease: my prospective study is to perform that test in the next 300 patients who may have that disease and see how well the blood test performed diagnostically. I have complete control over the design and performance of the study. Alternatively, I could do a retrospective study and review the hospital records of 300 consecutive patients in my clinic who had that disease and see how the test performed for making the diagnosis. I have less control over critical details in such a study, which is retrospective.

[3] It could be an extensive functional pathway within a cell (for example, how copper is handled in the hepatocyte, the system of interest to my research I describe herein) or even a disease process (such as abnormal handling of a mutated α_1-antitrypsin protein, discussed in Chapter 10). It is worth noting that

researcher wishes: a subcellular organelle, a cell, a tissue or organ. Importantly, given the system as defined, instead of being hypothesis-driven or data-driven, the research project is system-driven.

1.1.2 Other types of non-hypothesis-driven research

It turns out that biological/biomedical research encompasses several different non-hypothesis-driven research strategies besides system-driven and data-driven. Indeed experimental strategies cannot be divided symmetrically between hypothesis-driven or non-hypothesis-driven research. The array of available styles of biological research is complicated (Table 1.1). There is hypothesis-driven research and then there are all the others. In fact, there are numerous non-hypothesis-driven research strategies. No one seriously doubts that a biochemist working out the structure of a protein is "doing science", even though the experimental design is not directed by a hypothesis but, rather, by the nature of the protein under investigation. I call this "structure-driven research" because the point of the research is to produce a detailed and accurate description of the physical structure of the biomolecule. With minor tweaking, structure-driven research could be said to apply equally to cloning a gene since cloning a gene involves determining in minute detail the structure of the gene. Technique-driven research is more problematic than the rest because it can amount to a wholesale indiscriminate application of a technique, new and improved (or not), to whatever problems might be susceptible to some analysis by that technique. Researchers who engage in this sort of research are sometimes accused of doing "one-trick pony" research. They are more interested in the particular application(s) of their technology than in the problem(s) being addressed. A broader view of technique-driven research, however, includes the research aimed at developing a new technique. Figuring

---

the term 'system' has different meanings for different philosophers of science. These usages may differ from mine. One important instance is that Hans-Jörg Rheinberger speaks of experimental systems where that sort of system is a central feature of an experimental design (Rheinberger, 1997, 27). He examines the cell-free system of microsomes plus cytosol developed by Zamecnik for experiments relating to how proteins are produced in cells (Rheinberger, 1997, 55-73). Another example would be an isolated perfused whole organ for studying drug metabolism (for example, in philosophical parlance, a liver in a vat). The 'system' in "systems biology" is also not exactly identical with my concept of system, mainly because the former has a vaguely cybernetic connotation.

out how to do primary culture of liver cells isolated from a normal (or abnormal) liver would be an example. Developing the techniques for creating a transgenic mouse would

Table 1.1 Broad classification of types of research strategies

	Investigates	**Produces evidence?**	**Rise to prominence**	**Hypothesis-driven?**
Hypothesis-driven	Hypothesis	Yes: evidence determines T/F of hypothesis	1880s[4]	Yes
System-driven	System	Yes: evidence characterizes system	(late) 1990s[5]	No
Data-driven	A collection of data already collected	Yes: evidence reveals new information latent in the database	1999→[6]	No
Structure-driven	Protein structure; structure of some other molecule	Yes: produces a stylized picture of the structure, backed up by detailed mathematical description derived from structural evidence	~1950[7]	No
Technique-driven	The technique	Yes, often very valuable evidence, but this strategy often lacks a unified conceptual framework	1940s[8]	No
'Recipe'-driven	Whatever seems of interest	Yes, evidence of a sort, but at risk of completely lacking *any* conceptual framework	1980s[9]	No

[4] I am taking the move from naturalists working outdoors to laboratory scientists working indoors, mainly a German initiative, as the date for this trend.

[5] Genomics emanates from very late 1980s and starts to flourish in the early 1990s, proteomics from mid-1990s, and systems biology from approximately 2000—see text (Section 1.2). I accept that there were other systematists throughout the 20th century, but I am taking the date from the establishment of "Systems Biology" as such.

[6] This dating is based on A-L Barabási's paper "Diameter of the world-wide web" (Albert et al., 1999, 130). Development of large data bases and rise of social media occur mainly in the early 2000s.

[7] Pauling's work on molecular structure relevant to hemoglobin was published in 1951 (Pauling and Corey, 1951, 729).

[8] This date is a guess since it is typical of many aspects of biological/biomedical research to get a new technique and then apply it widely. For example, X-ray crystallography came into its own rather tentatively in the 1930s.

[9] I am dating this to the commercial availability of pre-fabricated kits for doing cross-species molecular biology assays.

be another example. Relevant to system-driven research, omics research in particular runs the risk of degenerating into the version of technique-driven research infatuated with the technique itself. Krohs quite properly inveighs against such research by calling it “convenience research” (the term invoking the metaphor of convenience foods) (Krohs, 2012, 53), but he seems to overlook that not all system-driven research ends up being technique-driven research. Nevertheless, opprobrium attached to technique-driven research has tended to cloud the epistemic issues relating to system-driven research. Krohs’ most vehement criticism is reserved for what I call “recipe-driven” research, which is a corrupted version of technique-driven research. With recipe-driven research, the researcher purchases the kit for an assay or technique and uses it as s/he likes, typically looking at some biological question because the kit makes it so easy. The problem with some technique-driven research and most recipe-driven research is a lack of critical stance to evaluate the data because the experimental design lacks either a hypothesis or a system driving it.

Terminological problems complicate contemporary discussions of innovative techniques in current biological/biomedical research. In this dissertation I will adhere to the inclusive and detailed taxonomy I have provided above.

1.2 Omics

At the present time many biological and biomedical researchers use various types of system-driven research, but comparatively few of these researchers focus on mathematical model-building as the main deliverable of their research enterprise. Such researchers do not see themselves as systems biology researchers per se. At the practical day-to-day level system-driven research resolves itself into some sort of “omics” research. The biologist or biomedical scientist uses high-throughput methods which epitomize the ‘omics’ research domains (genomics, transcriptomics, proteomics, metabolomics and the like). Although the word ‘genome’ was introduced in the 1920s, the various “omics” research disciplines, all ending up within the general scope of systems biology, are recent. The linguistic convention, developed in the 20th century, is that in biology tacking the suffix ‘-ome’ onto an entity conveys the meaning of the

"collectivity of the units [in something]" (Lederberg and McCray, 2001, 8). That seems to have been the intention when Winkler introduced the term 'genome' in 1920, although it can be argued that the current meaning of genome and his intention have diverged.[10] When the '-ome' suffix is changed to '-omics', the new word indicates the endeavour of investigating whatever constitutes that 'ome'. Thus 'genomics' means the investigation of genomes. This term was invented in the late 1980s, apparently (though perhaps unbelievably) because the geneticist Victor McKusick and his cell biologist colleague Frank Ruddle needed a name for a new journal they were launching (McKusick and Ruddle, 1987, 1). Transcriptomics relates to the RNA species in a system: these mediate between a structural gene and the expression of the encoded protein. Proteomics, the study of all the expressed proteins in a system of interest, was first articulated in the mid-1990s. One of the earliest reports examined the proteome of a *Mycoplasma* species in 1995 (Kahn, 1995, 369; Wasinger et al., 1995, 1090). Metabolomics examines small metabolically-active molecules in a biological system.

Of philosophical importance, the contrast of the 'ome', such as a proteome, and the 'omics', such as proteomics, clearly differentiates ontological and epistemological commitments. The 'ome' suffix signals the ontological entity, and the 'omics' suffix points to the epistemological stance. The proliferation of 'omics' disciplines has been exponential. Indeed it seems fair to say that as of 2013 omics has taken on a conceptual life of its own. A few years ago the journal *Nature* established an Omics Gateway, a website-based portal, or information source, for papers/reports about omics. An organization for promoting omics research publications and conferences has been set up (omics.org). Thus it seems appropriate here to speak of omics as a free-standing neologism in common parlance (no inverted commas, no italics). Omics deals with the investigation of the totality of a component in a biological system. Omics can look at one specific dimension of the system's function, providing a kind of tomographic

[10] The linguistic story is further muddled because many of those working in or around molecular biology teach that the –ome suffix comes from σομα, which is the root for chromosome (coloured body). This account is inaccurate, but it is highly persistent in the lore of contemporary molecular biology.

examination of the system under consideration. Thus as a research strategy for examining biological systems, omics research is system-driven.

1.3 Omics Experimental Strategy

1.3.1 Relationship between hypothetico-deductive method and omics experimental strategy

Currently, the contrast in research design which has generated much controversy is between hypothesis-driven and what I have called system-driven research. This contrast is between experimentation focused on a hypothesis and high-throughput methodologies, usually some sort of omics, not featuring a hypothesis. Given the prevailing view of how scientific method functions (namely, exclusively with a hypothesis), system-driven research may be passed off as "pre-scientific", the classic fatal deprecation of the Logical Positivists. I will argue here that system-driven research constitutes proper scientific research. It operates at least equivalently with hypothesis-driven research, not merely as a run-up to it. System-driven research offers innovative and unique ways to understand biological systems. A central feature of this argument is that system-driven research has a discernable epistemic structure. I will call it the omics experimental strategy (OES). OES is radically different from HDM. OES aims directly to be inclusive and operates without the prejudice of expectation related to a hypothesis. It may reveal unforeseen "surprising" findings. These are distinctive characteristics, absent in HDM.

As pointed out previously, within philosophy, the established hypothesis-driven strategy for experimental science is known as the hypothetico-deductive method (HDM). HDM was developed mainly from research in the physical sciences, and admittedly it has been sometimes difficult to apply it to the biological sciences. Nevertheless most recent biological research has been hypothesis-driven. The HDM turns on a symmetrical relationship between 'prediction' and 'explanation'. In formulating a hypothesis specifically to be tested by an experiment, the biological/biomedical researcher stipulates a possible explanation and then examines via experimentation whether the prediction entailed by that explanation actually occurs.

However, with system-driven research, experiments entertain no specific predictions about what results will emerge, and yet these results and the patterns revealed therein may constitute new knowledge. In other words, such experiments can still produce or explanatory information; however, the linkage between explanation and prediction, characteristic of hypothesis-driven research, is not present. The strategy of interrogating a system directly is characteristic of systems biology and the various types of omics which belong within the general firmament of systems biology.

1.3.2 A philosophical project relating to scientific epistemology

The project proposed here is envisaged as a basic epistemological investigation. How indeed do we investigate a system directly as a whole? Collecting a complete set of data is partly—even largely— a technical issue related to how the technology is utilized to accomplish this feat, but these technical aspects are located in the larger problem of how experimental design relates to data evaluation, which leads to acquisition of actual knowledge. Once the data set is acquired and enumerated, how do we justify an interpretation of these data so as to produce real knowledge? In a scientific enterprise whose results may not be strictly propositional, the definition of knowledge as 'justified true belief' may not apply. If this definition does apply, how is justification achieved? If it does not apply, what definition of knowledge is appropriate? Some principles of explanation and confirmation already developed for other kinds of experiments (mainly those according to the HDM) are likely to apply, but the nature of these principles, along with possible adjustments or alterations, needs to be clarified for omics.

Given this need for principles of explanation and confirmation, experimental design and knowledge acquisition serve as the two major focal-points for this philosophical project. In the first place there is direct contrast between research which utilizes a concatenation of hypotheses (that is to say, a chain or series of hypotheses), informed by the HDM, and omics, described by the OES, in terms of production of evidence. Hypothesis-driven research investigates a hypothesis and produces evidence which supports the truth or falsehood of the hypothesis under consideration. Applicable logical method is binary: annotated as 'true' or 'false'. System-driven research investigates a system and produces evidence which characterizes the system. The

relevant logic for system-driven research may not be binary: it may need to be more elaborate, employing descriptors like 'true', 'false' and 'not yet determined'. Both HDM and OES produce evidence, but in fact in the laboratory what scientists produce and what both the HDM and OES relate to is empirical data from experiments. However, there are important differences (Table 2) in how experimental data are generated, manipulated and evaluated. Evaluation itself, as indicated in Table 1.2, has two aspects: technical accuracy and broader validation.

Table 1.2 Specifics relating to how experimental data are handled in HDM vs OES

	GENERATION	**EVALUATION**	
	How are data generated?	**How are data submitted to quality control in the laboratory? (A technical consideration)**	**How are data validated as constituting worthwhile information? (A conceptual issue)**
Hypothesis-driven (HDM)	Focused experiment	Replicate experiments	Success of the 'Explanation/prediction' duo
System-driven (OES)	Examine a system or one of its components (numerous methods depending on what is being studied, including bioinformatics)	Assure that standards set by professional organization are met	Evaluation in the context of the system (contextualization); Evaluation in terms of promoting coherence of knowledge about the system
	Experimental design		***Knowledge* acquisition**

With HDM the experiment has a narrow focus directed at the hypothesis being tested. If the hypothesis is substantiated by the experiment, then it is expected that repeating the experiment will produce essentially the same results. In other words, replicating the experiment is part of the evaluating procedure. With OES the situation is quite different, since a system or one of its components is being studied. The methodology is complex; it typically includes sophisticated bioinformatics. It is not obvious that replicating experiments, as with HDM, is requisite for data evaluation: this will be an important point of analysis in this project. At the current time data evaluation

is organized to a great extent by the quality control standards set by professional organizations such as the Human Genome Organization (HUGO) and the Human Proteome Organization (HUPO). As I will discuss later (briefly in Section 1.3.2 and in Chapter 6), this can be regarded as a highly sophisticated form of social epistemology of science, but it is not the same as our customary peer-review. The process of transforming high quality experimental data into new knowledge is a separate aspect from getting the data in the first place. I believe this requires an innovative approach to data evaluation, an innovative scientific epistemology. Data are evaluated by contextualization in the system under investigation. I describe the detection of patterns in the data as the critical feature of this epistemology. However, assigning so much importance to patterns necessitates an analysis of what patterns are and how they behave—an analysis not already found within the available philosophical literature in relation to biological science, though considered for physical science.

1.4 General outline of the dissertation

In effect, research aimed at examining a system *directly as a whole*—that is, prominently, research utilizing the OES since most of this biological/biomedical research utilizes high-throughput methodology typical of omics—proposes a new way of doing science.[11] The principal features of this new scientific enterprise are inclusive characterization and an integrative stance. Taking system-driven research as a prime example, we can articulate the epistemological issue as follows: how do we investigate a system directly in its entirety and gain new knowledge? The general approach is to examine the OES—sometimes in comparison with the HDM—in terms of how experimental data are generated (column 1 in Table 1.2), how they are submitted to

[11] I recognize that some will object that the OES simply modernizes and re-introduces strategies enshrined in 19th-century projects called "natural history". Other will point out that the most ancient science, astronomy, has operated with a similar strategy for centuries. In my opinion, these objections enhance my philosophical project. Importantly, the situation in the 19th century for naturalists is not the same as that of contemporary researchers doing proteomics: no one questioned that 19th century naturalists were actually engaged in "doing science". Likewise it is probably worthwhile to abandon the navel-gazing posture whereby biological/biomedical researchers (or those who administer biological/biomedical research) completely ignore effective research strategies operating in other disciplines.

quality control (column 2 in Table 1.2) and how they are validated as knowledge (column 3 in Table 1.2).

1.4.1 How the dissertation rolls out

In **this introductory chapter**, I attempt to describe the broad problem to be analyzed and also show how I have delimited it somewhat for easier analysis. I am not attempting a philosophy of systems biology, and I cannot provide an analysis of every type of non-hypothesis-driven research. I limit my analysis to system-driven research. In concentrating on omics research, I will focus mainly on proteomics, the (omics) endeavour to identify and characterize all the expressed proteins in a specified biological system at a given state. Genomics, which has come to function as the poster-child for omics, presents certain conceptual difficulties. The main one is that genomics has become extremely complicated, in part because defining a genome (or a gene) is not straightforward in 2015 and also because there are numerous subgroups of genomics, such as functional genomics. Proteomics, by contrast, is much easier to analyze, since in general proteins are comparatively easy to define, although some will argue that the difference between proteins and peptides can be rather subtle. There are important theoretical reasons for choosing to focus on proteomics. The degree of definition is sufficient to achieve the high degree of definition required by HDM, and thus it permits clear comparisons to HDM. It is not possible to argue that definitional uncertainty as to what proteins are limits or incapacitates comparison with HDM. Proteins play a unique role physiologically: they are responsible for the workings of an organism, specifically its actual versatility within an environment. Thus the protein complement serves as a snapshot of the organism as it is, whereas the genome is comparatively static. As one of the omics research areas present from nearly the beginning of omics itself, proteomics amply illustrates the key conceptual problems to be tackled. Moreover, focusing on proteomics places the discussion squarely in the post-genomic era during which the focus has moved from genes (and genomes) to gene products. Among gene products, proteins and proteomes are of prime interest.

Moreover, I will anchor the theoretical issues of this dissertation in ‘real-world’ science, in fact, mainly in my own experimental work. Most of the experimental work

motivating this philosophical project comes from work relating to a genetic disease of the liver in which the handling of copper in liver cells is abnormal. In **Chapter 2** I provide background information about this disease, called Wilson disease, after the physician who first described it approximately 100 years ago in the venerable journal *Brain* (Wilson, 1912, 295). The problem of how to investigate copper disposition in liver cells led to developing a systems biology approach called 'metalloproteomics', a specialized version of proteomics. I had the privilege of contributing to the invention of this line of research, which indeed has called many of the issues to be addressed in this dissertation to my attention. Because this dissertation is so immediately based on highly technical scientific research, Chapter 2 provides a detailed overview of relevant structural and cell biology related to Wilson disease, as well as an up-to-date review of proteomics itself. These topics do appear to be formidably technical, and my objective is to provide the necessary background information to facilitate the analysis of scientific epistemology which follows.

It may be attractive to argue that the dichotomy of research being governed by a hypothesis or else not so governed is artificial. In **Chapter 3** I examine some objections to my formulating this clear and real differentiation between the two as worthy of philosophical examination. The dichotomy of being governed by a hypothesis or else not so governed currently has a powerful grip on scientific thinking. I believe we must regard it as an actual distinction, even if to some it makes little sense. Likewise I investigate and reject the opinion that I am overstating the problem—that all biological/biomedical research is actually hypothesis-driven. Importantly, I consider the objection that context of discovery versus context of justification distinction 'solves' the relationship between hypothesis-driven and non-hypothesis-driven research. In my opinion, if these dissenting views could not be rebutted, then my project would not get off the ground.

Since this dissertation is fundamentally about hypotheses (or lack thereof) and experiments (and how they are designed), these concepts requires some concerted attention. In **Chapter 4** I consider what a hypothesis is. The problem I identify for

biological/biomedical research is that hypotheses come in three different forms, and this multiplicity complicates any discussion of whether biological/biomedical research is "hypothesis-driven" or not. Given the hierarchy of hypotheses I describe, it possible for a system-driven experiment to be located within an epistemic environment of 'over-arching' and 'hunch' hypotheses and yet not be "hypothesis-driven".

In **Chapter 5** I consider what qualifies as an experiment. Likewise, the nature of 'experiment' involves some subtleties. I argue that an experiment is something performed in real time and that the notion of experiment necessarily involves being a deliberate, well-designed investigation, either a test and or a coherently organized set of observations. While considering several specialized types of experiment, I critically appraise 'exploratory experimentation' as developed in the past decade or so within philosophy of biology (Burian, 1997b, 28; Franklin, 2005, 888; O'Malley, 2007, 338). This has come to be identified with experimental design involving high-throughput methodology and/or omics. I believe that these analyses are flawed by focusing on one particularly challenging form of omics experiment, namely with DNA microchips. The contrast to my analysis involves their focusing on the technology—the epistemic consequences of how the experiment is done—as opposed my epistemological approach of examining what we are trying to achieve by doing an experiment which is system-driven. Importantly, the term 'exploratory' inevitably reinforces the old prejudice delimiting context of discovery versus context of justification, and it attributes higher epistemic worth to generation/investigation of hypotheses. Of interest, Steinle (Steinle, 1997, S70) did not make this misstep when he invented the same term for experiments in the physical sciences, but his notion of exploratory experimentation is also a little different from subsequent interpretations for biological research. O'Malley emphasizes novelty through metagenomics research (O'Malley, 2007, 347), and her analysis is closer to mine, though limited by being rooted conceptually in this notion of 'exploratory'. My analysis of system-driven research is not located in this line of philosophical assessment, although my work is relevant to it if their assertion that exploratory experimentation really is about omics is accepted.

Collecting a complete set of data is partly—even largely—a technical issue related to how the relevant technology is utilized. These technical aspects are located within the larger problem of data evaluation, which leads to acquisition of actual knowledge. In **Chapter 6** I examine the problems relating to quality control for proteomics data. From the technological point of view, data evaluation is currently organized by standards set by professional organizations of researchers in the field, such the Human Proteome Organization (HUPO) for proteomics. These standards have identified vulnerabilities in the processes for obtaining these data and begun to establish required uniformities in data collection, analysis and reporting to ensure excellent data sets. In fact, important progress has taken place to permit broad collaboration among proteomics researchers due to assurance of high-quality data and homogenous language for presenting and depicting the data. Finally, I begin to explore the theoretical issues about the nature of a proteomics data set itself, taking as a basis for this discussion an ideal situation where we suppose that no possibility of technical error exists.

Furthermore, once the data set is acquired and enumerated, how do we justify these data as real knowledge? This philosophical problem is examined in detail in **Chapter 7**. In a scientific enterprise whose results may not be strictly propositional—after all, as I point out subsequently, we are examining patterns within results—the traditional definition of knowledge as 'justified true belief' may not apply. Principles of explanation and confirmation already developed for other kinds of experiments may apply, but the nature of these principles, along with possible adjustments or alterations, needs to be clarified for omics. To address this problem, I compare these two ways of designing biological/biomedical research. There is direct contrast between research which utilizes a concatenation of hypotheses, informed by the HDM, and omics, described by the OES, in terms of production of evidence. Hypothesis-driven research investigates a hypothesis and produces evidence which supports the truth or falsehood of the hypothesis under consideration. Applicable logical method is binary. System-driven research investigates a system and produces evidence which characterizes the

system. The logic applied for evaluating the data may not be binary: for example, it may have categories such as 'works', 'doesn't work', and 'don't know yet' where "works" is a purposely vaguely term capturing the notion of function. This is the kind of assessment in play when you are putting a picture puzzle together. This metaphor is apt because a biological system could be conceptualized as a kind of picture puzzle.

With the HDM, the experiment has a limited focus directed at the hypothesis being tested. This entails having a clearly stated testable hypothesis. If the hypothesis is substantiated by the experiment, then it is expected that repeating the experiment will produce essentially the same results. In other words, replicating the experiment (in your lab and in the labs of expert colleagues) is an important aspect of the evaluating procedure; indeed, it constitutes most of it. With OES the situation is quite different, since a system or a class of its components is being studied. It is not obvious that performing the experiment repeatedly, as with HDM, is requisite for data evaluation.

Given these considerations, I propose in Chapter 7 to show how the OES functions as an epistemological basis for omics (taking proteomics as the concrete example). This analysis seeks to establish the OES as a scientific epistemology. It takes inspiration from process reliabilism as it relates to justification, but it is not attempting to produce a general epistemology. With the OES, a well-described system and technically pristine practice can produce high quality data, which is to say reliable data. We develop standards to ensure the high quality of these data. With respect to knowledge acquisition, the technological story—while fascinating and important—is not the whole story. A key question remains: suppose we had technically and technologically 'perfect' data from a proteomics experiment, how would we evaluate it as new knowledge? I argue that whatever drives experimental design also contributes to the evaluative machinery. Just as the results of the HDM experiment are assessed in relationship to the hypothesis, the results of the OES experiment are assessed in relationship to the system. Do they fit with what is known about the system? If the answer is 'no' or 'not certain', then further investigation is needed. (A problem with the system-driven experiment is that it is hard to exclude results as irrelevant.) Indeed some

prominent experts in proteomics have called on cell biologists to 'get with it' and initiate these evaluations! (Bell et al., 2007, 784). Thus the role played by accurate prediction with HDM is played by 'good fit' with contextualization. Whereas HDM operates in an epistemic environment of synergy between explanation and prediction, OES typically exists in the epistemic context of a system that is coherent—it literally fits together.

Chapter 8 is devoted to working out how the OES actually produces justification of non-propositional empirical data. Specifically, I examine a distinctive epistemic action: the detection of patterns within the data set. I develop a more stringent definition a pattern as 'a salient orderly juxtaposition of component entities, characterized by the relation(s) among them'. Those relations include recurrence and interaction. In proteomics, patterns can be direct (for example, with the colour display of a protein chip) or indirect (as when a list of proteins is grouped according to functional or anatomical classifications). Patterns associated with proteomics can also feature how entities interact in a definite fashion. The important evaluative process is contextualization within the system. Identifying these patterns and seeing how they fit (or do not fit) into the system under investigation provides a testing and justification of data via a strategy which is broadly coherentist. Such an approach is still consistent with a non-dichotomous evaluation, allowing for the possibility of not being able to tell whether some data fit into the system or not based one available information about that system.

In **Chapter 9** I discuss some of the implications of the OES as a scientific epistemology. In addition to some but certainly not all of the standard epistemic virtues, a characteristic epistemic virtue is novelty. The OES can reveal knowledge which is really new knowledge, not necessarily inaccessible to the HDM but unlikely to be revealed by the HDM in any efficient fashion. An epistemic advantage of dealing with patterns vis-à-vis the OES is that patterns may permit a more open-ended assessment of the system. The nascent pattern may suggest various possible overall patterns (I call this 'suggestibility', first introduced in Chapter 7), which are different from what was anticipated. Reverting to the problem of the context of discovery versus context of

justification dichotomy, I argue that thanks to the OES discovery regains its important epistemic role in how we do biological/biomedical research. Finally I discuss the interactions among the different types of research styles in the course of performing actual real-life biological/biomedical research. In a comprehensive research project, different research strategies may be employed for different aspects of the problem under investigation. In the biological sciences there is an easy conversance among all these different strategies. No single approach is *primus inter pares*. One of the researcher's objectives is to choose the best experimental design for the problem under investigation. However, to recognize that hypothesis-driven and non-hypothesis-driven research methods (notably system-driven research such as omics) function as equal partners in current scientific research is to recognize that non-hypothesis-driven research[12] is legitimately a way of "doing science".

Chapter 10 serves as a summary chapter for the key themes and issues in this dissertation. I briefly extend this analysis beyond Wilson disease to another genetic liver disease, α_1-antitrypsin deficiency, to illustrate that the issues considered here are not limited to a relatively rare disease and human biology of trace metals. This example shows that the problems considered here have broad scope. Finally, I suggest that even if at the practical everyday level everybody accepted system-driven research as '*obviously*' real science, the impetus for developing a scientific epistemology for it would remain.

1.4.2 Roads not taken

There are several different aspects of current biological/biomedical research which could motivate this dissertation and serve as the chief point-of-view, but I have chosen not to address them in any detail. First, within the past decade funding agencies have been hesitant to invest grant-monies into system-driven research (O'Malley et al., 2009, 613; Keating and Cambrosio, 2012, 37). Their explanation (or excuse) was that it

[12] The objective of my work here is to demonstrate that this assertion holds for system-driven research by developing a scientific epistemology for it. I believe it also holds for structure-driven research. Other types of non-hypothesis-driven research pose certain complex problems in terms of formulating a scientific epistemology: developing a scientific epistemology for each of them is beyond the scope of my current project.

was really not “good science”—or indeed, not even “real science”. The perception that science must be organized around a hypothesis has pervaded the scientific and policy-generating communities in a way which has important political features and implications. The politics of ‘grantsmanship’ countenances (or even promotes) this epistemic confusion about experimental design. This is an important practical societal aspect which may figure in latter part of this dissertation or else serve as a topic for future work. However, sorting out different kinds of hypotheses (in Chapter 4) serves to disentangle some of the problems here.

Secondly, the degree of operational complexity among research styles (hypothesis- and non-hypothesis-driven) may account for the general perplexity within philosophy of biology about the nature of experimentation and knowledge production in biological/biomedical research. As important discoveries accrue through this to-and-fro of hypothesis-driven and non-hypothesis-driven research, we see that HDM plays an important complementary role with other research styles but does not hold an exclusive title to scientific research. I fully agree with Ian Hacking that “we should not expect something as motley as the growth of knowledge to be strapped to one methodology” (Hacking, 1983, 152). Exciting new vistas for biological/biomedical research extend before us. These research opportunities will require some adjustments in our social epistemology of biological/biomedical science, a topic lying once again beyond the scope of this dissertation. The discussion of establishing technical standards (in Chapter 6) provides glimpses at some of these issues.

1.5 Summary and significance

Biological/biomedical research that attempts to examine a biological system as such in its complexity, reflects a distinctive epistemological orientation. Instead of being organized around a hypothesis or an existing data set, the research project is organized around a system. This system can be larger or smaller (as the researcher wishes: a subcellular organelle, a cell, a tissue or organ; a functional pathway within a cell, or even a disease process), but it requires a clear and detailed description. Importantly, given the system as defined, instead of being hypothesis-driven or data-driven, the research

project is system-driven. Currently, the contrast in research design which has generated much controversy is between hypothesis-driven and omics, which is system-driven research dependent on high-throughput methodologies not featuring a hypothesis. Given the prevailing view of how scientific method functions (namely, exclusively with a hypothesis), many consider system-driven research as somehow not scientific. I argue here that system-driven research constitutes proper scientific research. It possesses an epistemic structure for this high-throughput research design, namely the OES; the OES itself has interesting epistemological implications. System-driven research operates at least equivalently with hypothesis-driven research, not merely as a run-up to it. Importantly, system-driven research offers innovative and unique ways to understand biological systems, and perhaps even the methodology of science itself.

CHAPTER 2 BIOLOGICAL UNDERPINNINGS

Biological research is inherently complicated. The object of study, terrestrial life, is complicated. Even ostensibly simple, unicellular organisms—such as bacteria—display complex mechanisms for metabolism and reproduction. Although it makes sense to dissect organisms into individual component parts and study each part separately in exquisite detail—and this has been a highly productive research strategy[13]—there is a lurking suspicion that these separate components function differently in the context of the organism, as opposed to how they function in isolation. It is worthwhile to decompose the entity into manageable pieces, but that effort, though extremely productive, seems incomplete. There is an undeniable urge to recompose the pieces and study the parts functioning within the whole. Yet even that effort is attended by some worry that the reconstituted biological entity will not be quite the same as the original. If it were possible to examine a system in its entirety as such, it would be of great interest. Omics facilitates this kind of analysis. Different kinds of omics examine a system from different perspectives: genomics, the genes; transcriptomics, RNA; proteomics, proteins.

Since even a so-called 'simple' organism is composed of functional parts working together, it is possible to construe such an organism as a system composed of subsystems. A bacterial cell can be seen this way, and likewise so can a mammalian cell. A tissue can be regarded as a system where the component subsystems are cells. Organs are composed of various tissues. In ecology biologists have analyzed whole communities of biological entities with a 'systems' kind of strategy. So, in a sense, what we are looking for is a scientific strategy for studying a biological system. Systems

[13] The most important contemporary example of this strategy is Bechtel and Richardson's analysis of biological research whereby an entity or phenomenon is analyzed piecemeal in its component parts (Bechtel and Richardson, 2010, 7). This process is called 'decomposition'. Mechanistic accounts of biological processes are well discussed in a recent paper (Bechtel, 2011, 535). More recently, Bechtel has been examining reconstituting the entity by putting these parts back together in a process of reassembling the components or 'integration' advances the epistemic project. This latter comment is based on a lecture by Bechtel at biennial meeting of the International Society for the History, Philosophy and Social Studies of Biology (ISHPSSB), Montpellier, France, 2013.

biology is prepared to take on this challenge: the biological entity is regarded as a system[14], dismantled and analyzed, and reassembled as a schematic model with mainly mathematical descriptions. Systems biology is therefore inherently integrative. It requires the detailed and massive data sets which omics can supply (Aggarwal and Lee, 2003, 176), but systems biology as a whole exceeds the remit of high-throughput methodologies in biology. My interest is not focused on the issues related to mathematical reconstruction as a means to understanding a biological entity; rather, it is for the rather more immediate problems of ensuring that relevant omics data sets are of sufficiently high quality to support systems biology, or any investigational approach to gaining new knowledge about system, in a meaningful fashion. 'Integrative biology' is another candidate for studying biological systems. It includes transdisciplinary research efforts bridging biology to chemistry, physics, engineering and the like. It is too general for my purposes, indeed not focused enough, and the term 'integration' comprises too many meanings and connotations in contemporary science (and philosophy) to be manageable. Thus, in a sense what we really need is a ***system*** biology.[15]

For proteomics especially, the recent notion of 'network biology' may serve effectively as a practical version of system biology. Within cells, proteins are found interacting with other proteins and other molecules. Despite the potential for subtle chemical diversity, including both minor transient chemical alterations and striking permanent chemical modifications after production (called post-translational

[14] It is worthwhile noting that when we take systems biology as the starting point, the concept of 'system' is somewhat different from some notions of 'system' in the philosophical literature. We are not limited to "model organisms", such as *C. elegans* or any specific knockout mouse, as the system under study. Likewise the term 'system' is sometimes taken to refer to the experimental apparatus employed such as tissue slices maintained in aerated physiological buffer or a cell-free system. A detailed discussion of what I mean by 'system' follows in Chapter 6.

[15] Some may object that this terminology is too similar to "systems biology" to be effective, or even workable, but I am hoping that this discussion makes the rationale for wanting a *system* biology clear and facilitates this possibly idiosyncratic usage. I have denoted this concept typographically in this dissertation with the word 'system' italicized. While this typographical convention might seem a little fussy, it may facilitate a clear differentiation from systems biology (with an 's'). Incidentally, this term has been suggested independently by Vincent Collura and Guillaume Boissy, who say that "system biology can be described as the study of the cellular network and cell components organization" (Collura and Boissy, 2007, 173). I am disinclined to such a tight definition, but their conception fits with my notion of *system* biology.

modifications), a solo protein in isolation has little scope for action. The cell as a system of subsystems can be viewed as being composed of organelles or it can be viewed as composed of protein networks. This latter view makes intuitive sense in light of the observation that protein-rich membranes within a cell are interconnected. Some membranes are continuous with one another; some are connected by proteins serving as tracks from one part of a cell to another.

In this chapter I will first develop the narrower concept of *system* biology based on a brief consideration of the general character and objectives of systems biology. Then I will review the basics of proteomics as the omics discipline of interest, a key example of *system* biology. In particular, I will present the history of metalloproteomics and show how it illustrates many of the problems of omics experimentation. I will also argue that our concept of metalloproteomics formulated in the early 2000s includes, at least in a *forme fruste*, the key features of contemporary proteomics and thus forms an effective basis for examining the attendant epistemological problems of proteomics today. Finally I will review some pertinent aspects of the clinical disorder, Wilson disease, and the relevant cell biology because this metabolic disorder of copper disposition was the motivation for inventing metalloproteomics in the first place. These discussions are intended as a kind of tutorial on the relevant science in order to facilitate the philosophical analyses which follow. Wherever possible, I will draw illustrative examples from aspects of copper metabolism or Wilson disease.

2.1 Systems biology

Systems biology considers the organism as a whole and is interested in complex interactions (Ideker et al., 2001, 343). In a general way it has been around for decades and actually predates omics. In my view, systems biology typically addresses the connection between structure and function; however, some descriptions are less specific on this point. Although some argue that systems biology is only vaguely defined (Callebaut, 2012, 72), one standard description of systems biology is as follows:

> Systems biology is a comprehensive quantitative analysis of the manner in which all the components of a biological system interact functionally over time. Such an

> analysis is executed by an interdisciplinary team of investigators that is also capable of developing required technologies and computational tools. In this model, biology dictates what new technology and computational tools should be developed, and, once developed, these tools open new frontiers in biology for exploration. Thus, biology drives technology and computation, and, in turn, technology and computation revolutionize biology. (Aderem, 2005, 511)

Establishing systems biology as a research endeavour is possible nowadays because we can generate and analyze large, highly inclusive data sets.[16] In the past 20 years, many omics disciplines have been developed which relate to the types of processes which take place in hepatocytes (the dominant cell of the liver parenchyma). A few of these disciplines are genomics, transcriptomics, proteomics, and metabolomics. In a general fashion, any omics discipline operates by a systems biology approach.

Recently, a distinction has been made between top-down systems biology and bottom-up systems biology (Bruggeman and Westerhoff, 2007, 46-7). With bottom-up systems biology one formulates in detail each interactive behaviour one already knows and builds up to the whole system from that. With top-down systems biology one examines an intact system globally and expects to find new relationships. In fact, top-down systems biology is what most people consider as the archetype of systems biology or omics research as such. I would argue that the bottom-up systems biology described by Bruggeman is nothing other than conventional reductive experimental design done on a large scale. Further, these descriptors do not necessarily enhance our thinking about systems biology. They have been in general parlance since the 1970s and relate to matters as diverse as methodology for developing a computer program to characterizing types of inference! These terms are employed in areas ranging from business practice to neuroscience. Their uneasy application to systems biology is indicated by the further term “middle-out” (Kohl et al., 2010, 30; O'Malley and Dupre, 2005, 1272-3), which is also applied to systems biology. “Middle-out” means to start somewhere in the middle

[16] This more contemporary version of systems biology is the one I prefer for discussion in this dissertation. A concise description is the following: “Systems biology is a new field of science that develops a system-level understanding by describing quantitatively the interaction among all the individual components of the cell” (Aggarwal and Lee, 2003, 175). This view of systems biology is directly dependent on how omics disciplines can supply large amounts of data to permit the systems analysis.

and work in both upward and downward directions. This terms most accurately depicts how most biological research operates. You start somewhere, but generally toward the bottom.

I would prefer to abandon 'top-down' and 'bottom-up' terminology and use some other metaphor less burdened by prior application. Taking a metaphor from photography, for example, we could talk about 'big-picture' and 'narrow-focus' methodologies in biological research. 'Big-picture" corresponds to top-down, and 'narrow-focus' corresponds to bottom-up. Omics fits most naturally with big-picture methodology. This is not to say that omics necessarily zeroes in on the biggest picture: it might be applied to the big picture of something small (not the whole cell but only an organelle in it; not all of the possible neoplasms of a tissue but just one in particular). Narrow-focus methodology typically starts with something very circumscribed: a single protein or single gene; however, investigating that entity may take the inquiry in various directions within a system. If that inquiry moves from one hypothesis to another, then it tends to be linear irrespective of direction, and we could characterize it as one-dimensional. In contrast, it probably not an exaggeration to say that big-picture methodologies tend to be broader and multi-dimensional. In any case, although I would certainly like to avoid the terms 'top-down' and 'bottom-up' in my analysis,[17] these terms are so entrenched that I will continue to allude to them.

A reason that 'bottom-up' and 'top-down' might be applied to systems biology is that the output goal of systems biology is not so very different from the goal of computer programming: the overall objective is to capture the system as a mathematical model, or perhaps more circumspectly as a model in the kind of

[17] Another reason to avoid these terms is that they are also subject to considerable variability in how they are used within the proteomics literature. A convincing example is found in the review by Bensimon and colleagues where their notion of bottom-up proteomics is a variance with conventional usage (Bensimon et al., 2012, 381). The reader must at least be wary of this slipperiness in terminology. Similarly 'bottom-up' and 'top-down' find technically specialized meanings in the philosophy of science literature. Katherine Brading uses 'bottom-up' to mean "from world to theory" (Brading, 2010, 832) in the sense of "the choice of how to structure the phenomena is not driven *bottom-up, by how the world is*" (Brading, 2010, 831, my emphasis). Elsewhere in philosophical literature we may find bottom-up somehow conflated with inductive reasoning and top-down associated with deductive reasoning.

mathematical language employed for engineering and computer analysis. The point is made that at the present time we have the mathematical tools, specifically those for showing stochastic processes, to accomplish the feat, at least in principle (Bittner and Dougherty, 2012, 188; Dougherty and Shmulevich, 2012, 577-8). Detailed, complete accurate data sets are required. Thus systems biology encompasses omics (which can provide the sorts of data sets required) but goes well beyond the domain of omics.

It may be worth noting that bottom-up (narrow-focus) systems analysis explicitly assumes that there is a network of entities, for example, with proteomics a network of interacting proteins which will permit or even expedite the analysis. Top-down systems analysis also assumes there is a network but somewhat more subtly: one of its objectives is to disclose that network in greater detail. I will examine these issues again in the subsequent section devoted to 'network biology' (Section 2.1.2).

2.2 Analyzing biological systems as such: how to envision *system* biology

2.2.1 Integrative biology

If integrative biology is seen as the 'new biology' of the 21st century, it might be regarded as having a strong claim to being the main approach to studying biological systems. A circumspect version exists which resembles systems biology without the mathematics: it traces its lineage to the interesting work of Robert Rosen on what he called relational biology (Paton, 2002, 64). Creating models to represent and explore biological mechanisms seems to be an important output. However, the version of integrative biology which is adopted as the name of university departments and research institute programmes displays a much more ambitious agenda. It is described by Marvalee Wake as "a way of perceiving and practicing science and of transforming science—its processes and results—to deal with societal issues" (Wake, 2008, 349). Wake argues that integrative biology is not amorphous: it is broadly transdisciplinary, and it inclusive of systems biology (Wake, 2008, 350). The attitude of integrative biology alters how research is done (mainly in teams) and how science is taught through collaboration across disciplines in curricula and even in laboratory geography. Integrative biology also sees an important role for integration into society, what we in

Canada would call "KT" (knowledge translation).[18] Robert Full, Wake's colleague in the Department of Integrative Biology at Berkeley, argues in a sidebar comment to her 2008 *BioScience* paper, that "transformative biological research will require integration with physics, engineering, mathematics, chemistry, and computer science to a degree not yet seen" (Wake, 2008, 351). Without subscribing to Full's overriding interest in engineering or invoking the terms of reference of systems biology, this seems like a place where proteomics might find itself.

What is interesting about Wake's description of integrative biology is the importance it attaches to innovation (Wake, 2008, 350). Integrative biology is faster science because it is more efficient than the solitary researcher[19]; moreover, the cross-pollination of ideas and experimental techniques, taken at their most inclusive, eventuates in greater innovation. Overall, the currently pervading view of integrative biology emphasizes addressing complex biological problems, bringing together diverse research expertise to investigate these problems, being fearless about extending research into non-traditional arenas, and developing new ways to educate people about science and train them to do science (Wake, 2003, 240).

What seems to plague integrative biology is that, though 'integration' is absolutely central to its conception of itself, integration implies different commitments

[18] Unquestionably, the current emphasis in biological/biomedical research is on research done by consortia and teams. The era of the lone researcher in his/her small fiefdom laboratory seems to be over. One way for such teams to work is that they are assembled on a modular basis: the researcher supplements his/her already extensive laboratory resources by deploying certain aspects of the research project to collaborators who already possess the requisite expertise. I am grateful to Stephen James for pointing out the current all-importance of such a collaborative approach (28 October 2014) in part a response to the need to get research done rapidly and cheaply. Nevertheless I see this collaborative style as operating for decades, only on a smaller scale. Rheinberger chronicles its variable implementation and success in connection with Paul Zamecnik's research (Rheinberger, 1997, 84-97)]. Open journals and open databases are other aspects of this collaborative style. KT is already a priority in Canada since it was enshrined in the Act of Parliament which created the Canadian Institutes of Health Research (CIHR).

[19] Though somewhat contradicting herself, Wake pays tribute to Charles Darwin as the "quintessential integrative biologist" (Wake, 2003, 240). She writes: "We recognize with awe that Darwin did systematics, morphology, development, behavior, physiology, and ecology—and out of it all emerged a grand synthesis and a new conception that has become the foundation of much of biology—evolution through natural selection" (Wake, 2003, 240). Apart from the intrinsic interest of these comments about Darwin, it is also of some interest that one criticism of my rendition of omics is that we are just going back to the future: we are regaining the experimental approaches of 19th-century naturalists.

to different people. Thus it is like a vague but formidable assumption. You are not entirely certain what you are assuming. Omics fits into integrative biology in various ways. When we use complex machinery to do proteomics, all highly dependent on central aspects of biophysics, we are also integrative in our biological research, but at a more modest level. When we try to coordinate different data sets from different omics disciplines, we are certainly performing an integrative step. On balance, integrative biology is too broad for my general purposes here, but it may ultimately be relevant.[20]

2.2.2 Networks as systems: Network biology

A very basic concept in cellular/structural biology is that proteins rarely, if ever, function in isolation. Perhaps the most accessible example is that of an enzyme (always a protein) which accelerates a biological process by interacting with the components of the process—the usual figurative language is as a "lock and key"—and then regenerating itself as its role is concluded. Proteins almost always interact with each other. Even structural proteins, which are in effect immobile, interact with each other through cross-binding elements and serve as scaffolding for other dynamic proteins. The molecular details of protein-to-protein and protein-to-DNA/RNA interactions are becoming definable through new analytical techniques of structural biology. Thus we tend to conceive of processes mediated by proteins—and indeed, that means most cellular processes—as involving 'networks'. Certainly there are well-established examples of protein networks such as the cascade of activated proteins which effect blood clotting and the sequential handling of a chemical in drug metabolism. The oxidative phosphorylation "chain", which is made up of a five-enzyme cluster arranged in a series in the mitochondrial membrane, serves as the machinery for energy production in mammalian cells. Complicated diagrams of the interlocking metabolic pathways of cells have been produced for decades. More recently, determining the set of proteins any one protein might interact with has been of interest to cell physiologists. This set is generally called an interactome (Biron et al., 2006, 5591).

[20] High-throughput methods are likely to have societal impact, for example, with so-called 'personalized medicine'.

The notion of 'network biology' is, however, somewhat different from this fundamental idea that proteins and other cellular molecules function by interacting. The motivating idea with network biology has to do with fundamental organizing principles of how networks work. Networks are regarded as having essential characteristics; moreover, laws can be discovered which regulate the characteristics and functioning of networks (Barabasi and Oltvai, 2004, 105). Unquestionably computers organized as a LAN (that is, a local area network) in your house—or more dramatically, a nation full of computers linked via the worldwide web—interact with each other differently from the way cellular proteins, whose primary function is to move copper around the cell, interact with each other. However, when we import network theory into biology to create 'network biology', the expectation is that the same characteristics and regulation attributed to the network of computers can be identified as actuating what happens in the biological setting.

Specifically the version of network biology which has captured the imagination of the scientific community emanates from the work of the physicist/mathematician, Albert-László Barabási, who continues to be its main exponent.[21] Most recently he has described it as "network science" (Barabasi, 2009, 412). Barabási's work challenges the orthodoxy that the organization of networks is random. In the late 1990s, he began examining the organization of the World Wide Web and found that its organization appeared to be non-random: the probability that a web page had exactly *k* links was given by a power law distribution, very different from the Poisson distribution predicted by random network theory (Barabasi, 2009, 412). This formula represents an important part of a mathematical depiction of how the network is organized. Such a network is inhomogeneous, composed of a few nodes with numerous links ("hubs") and then few or many other nodes with few links. Borrowing a term from the statistical physics

[21] Part of the excitement around this concept arises from publication in the lay press. Other early important publications were in *Science* "Emergence of scaling in random networks" (Barabasi and Albert, 1999) and *Nature* "Error and attack tolerance of complex networks" (Albert et al., 2000) and "The large-scale organization of metabolic networks" (Jeong et al., 2000).

literature, he called such networks 'scale-free'[22] (Barabasi and Oltvai, 2004, 104). Moreover, scale-free networks get larger by a process which is skewed. New nodes link preferentially to those nodes which already have numerous links. Barabási calls this the rich-get-richer effect, although it is also known as the Matthew effect (Keller, 2005b, 1060), as a Biblical allusion. Then he found that these scale-free networks were almost ubiquitous. Examples could be found in studies of social networking, cell biology, cancer biology, and information technology relating to computer, among others. There are also hierarchical networks, but these are rare.

For cellular biology the impact is broad. According to Barabási and Oltvai, "an important development in our understanding of the cellular network architecture was the finding that most networks within the cell approximate a scale-free topology" (Barabasi and Oltvai, 2004, 104). This viewpoint has been slightly revised to take into the constraints chemical bonds and molecular shape impose on interactions (Barabasi, 2009, 413). Nevertheless, among three possible types of network models in biology (Barabasi and Oltvai, 2004, 105), the scale-free network seems to be pre-eminent. This analysis of scale-free networks lends itself to interactome analysis, modularity and assessment of robustness in biological systems.

For all its interest, Barabási's notion of scale-free networks poses some problems for biological/biomedical research. The notion that most cellular network architecture is scale-free raises the possibility that the mathematical formulation could be a universal descriptor for cellular function. When Loscalzo and Barabási extend their concept to network medicine, they describe an approach to disease "defined as the science of integrating genetic, genomic, biochemical, cellular, physiological, and clinical data to create a network that can be used to model predictively disease expression (and

[22] A scale-free network is a network characterized by the presence of large hubs. Such a network has a power law distribution. A key concept here is that "networks with power-law distributions a[re] called scale-free because power laws have the same functional form at all scales." (http://mathinsight.org/scale_free_network, accessed 2 November 2014) (Nykamp). (I have corrected the typo on the webpage.) It is difficult to avoid a certain circularity in this definition, namely, that the network is scale-free because it has a power-law distribution. An important contrast here is between the Poisson (normal) distribution and the power-law distribution.

response to therapy)" (Loscalzo and Barabasi, 2011, 621). However, this concept is effectively the same as "integrative medicine", and 'network' is not required. Evelyn Fox Keller has provided a formidably trenchant critique of Barabási's network theory. She criticizes it on the basis of its being much less innovative than Barabási has portrayed it and much less informative, and she is worried about its terminology. Her first main point is that scale-free networks are easy to generate: there are various ways to do it. Her second main point is that "the mere fact that a nodal degree distribution follows a power law actually implies nothing either about the mechanism giving rise to it or about its particular architecture" (Keller, 2005b, 1061). She illustrates the problem with the example of five obviously different network graphs which each have the same number of nodes and links and four of which have the same power law node distribution (the fifth serves as a kind of control) (Keller, 2005b, 1064). She doubts that network theory, or network science as it might now be called, is going to generate a fundamental principle of biology. In summary, Fox Keller allows that scale-free networks exist, but power law distributions are neither new nor special; indeed it is pretty easy to fit data to such distributions, and once you have that characteristic it does not supply mechanistic information (Keller, 2005b, 1066).

The concept of 'network' in terms of protein-function is sufficiently fundamental and entrenched in ordinary scientific practice that proteomics is already being extended to include network biology in the weak sense, as well as in the strong Barabásian sense. The central concept here is that networks of molecules mediate phenotypes from genotypes (Bensimon et al., 2012, 379). Instead of being simply vectorial, the route is complex, reticulated and may vary with context within which the network is operating (Bensimon et al., 2012, 381). As we will see, this insistence on the importance of context is an important feature of our notion of metalloproteomics. Since how/when/where proteins are expressed reflects the dynamics of a biological system, proteomics positions our analysis on relationships between the genetic apparatus and external environment. Applying proteomics data to network biology requires high-quality data (Bensimon et al., 2012, 384). However, current technological improvements for

proteomics permit detailed investigation of protein-protein interactions as well as interactions of proteins with other biomolecules such as metabolically active small molecules and specific RNA and DNA sequences (Cox and Mann, 2011, 275). Affinity-purification of the biological specimen[23] followed by standard mass-spectrometry-based proteomics investigation has become the preferred experimental method for studying such interactions because important 'natural' aspects of the protein interaction can be preserved (Cox and Mann, 2011, 285-8). Thus proteomics is well-positioned to be an important contributor to network biology, largely because it is proving to be an efficient way to investigate interactomes (see Section 8.4.3.3).

The whole idea of 'network biology' is appealing because it seems so intuitive.[24] In the context of cellular copper handling (and metalloproteomics), for example, as I will discuss in greater detail subsequently, we formulated the question of interest to us about copper in hepatocytes as: "what is the network of proteins which is responsible for the disposition of copper with in a hepatocyte?" This formulation betrays an assumption that proteins really do work in networks. In fact, it is difficult to discuss this assumption because it is so engrained. In 2001, although network biology was not yet a really 'hot topic', nevertheless the notion of a protein network was fundamental to this strategy for experimental design. In effect, for the purpose of this project relating to cellular copper handling, the liver cell could be identified as the system within which we were operating or, with the focus narrowed further, the network of proteins responsible for copper disposition in the cell was the system of interest.

The discrepancy between the intuitive, almost axiomatic, grasp of networks in biology and 'network biology' as such is an issue which demands continuing attention while I am addressing philosophical issues associated with proteomics and the OES. Network biology by its very nature enshrines the assumption that a network is present and dominant in any biological system. We have seen that this assumption operates in

[23] That means that you first treat the specimen somehow so that proteins in the specimen capable of binding to the molecular of interest have a chance to do so. The concept here is parallel to that of the IMAC column for capturing metal-binding proteins.

[24] Whether we want to jump on the Barabási bandwagon is a separate issue. I believe it generates serious problems, but it is sufficiently influential at the present time that it cannot be ignored.

both big-picture (top-down) and narrow-focus (bottom-up) systems analysis but in either type of analysis it may be open to criticism. As I will show, I am inclined to be wary of the broad implications of network biology, specifically, that the adequate description of the system as a network can be conveyed in power-law formula. Consequently, I cannot settle on 'network biology' as the descriptor I am seeking, although I am rejecting it for reasons which are very different from my reasons for rejecting 'integrative biology' as the descriptor.

2.2.3 *System* biology

Against the background of these other candidates (systems biology, integrative biology and network biology) the reason for articulating *system* biology is that none of these others demarcates the territory I wish to investigate. Systems biology is too inclusive and mathematical; integrative biology is unmanageably broad, as well as intrinsically tautological by being integrative because it seeks to be integrative; network biology is too restrictive and essentialist in its insistence that every system is a network, describable by laws of 'networks' and oblivious to context. Yet, each of these terms describing a biological research endeavour is on to an important aspect of current biological research. Systems biology seeks large detailed data sets describing individual aspects of the system under investigation, manipulates those data through complex mathematical procedures and attempts to produce a mathematical model of how the system works. Among other aspects of systems biology, the detailed examination of the system and objective of modeling it retain importance for omics. *Pace* integrative biology, integration of methodologies is critically important for investigating a system, as we will see in the need for proteomics, for example, to contextualize its findings in the physiology of the system being studied. A more distal aim is to coordinate the findings of different omics disciplines around the system being studied. It almost goes without saying that the technological suite required for omics research spans numerous capabilities ranging from engineering to molecular and structural biology, and where translational research is desired, to clinical science. Thus the integrative aspect cannot be disregarded.

Network biology looks like a good fit to the subject at hand, and I will return to it regularly, mainly because it is germane to the study of proteins. The existence of a protein network within hepatocytes responsible for copper disposition was a key assumption to our research. However, taking the concept of 'network' effectively as an assumption could be problematic: it could end up being an impediment to the larger analysis in situations where the degree of complexity exceeds what can be described as networks. The limiting versions of *system* biology, as I conceive it, are molecular structures at the diminutive end and ecological systems such as forests at the massive end. In a sense a molecular structure is a system in microcosm, and it would not fit into the conceptual vision of network biology. An ecological system may be a complicated mesh of networks: dealing with ecological systems is well beyond the scope of this dissertation. I want to retain the option of considering how network biology is relevant to my discussion, but I do not want to make it the focus of my discussion.

If I am correct in my assertion that biological research which attempts to examine a biological system as such reflects a distinctive epistemological orientation, then I may further assert that I need a new name for what I am talking about. I submit that this is '*system* biology'. Unlike systems biology, it does not promise an integrated, detailed, mathematical deliverable. It will not necessarily produce a model characterized by intimidating mathematics. *System* biology does describe research which is organized around a system of interest, where that system is defined narrowly or broadly (an intracellular protein network, subcellular organelle, a cell, a tissue or organ, or even a whole organism—but if a whole organism, then it is likely to be at least a compact organism, like a virus or bacterium). So the motivating question here is: how do we investigate a system directly in its entirety? The point is that new and emerging technology, including those complicated technologies which support omics types of experiment, permits organizing a research project this way. Moreover, we conceptualize biological mechanisms as operating as systems, not as unilinear causal chains, and this is where the active metaphor of 'network' may operate. Thus we may speak of '*system* biology' as such. It encompasses the strategy of interrogating a system directly. This

strategy is strikingly different from the conventional strategy of a concatenation of hypotheses. In fact, these new techniques and views of biology demand that we revisit the role of hypotheses in biological/biomedical research.

I will further argue that *system* biology involves a kind of thinking by looking. This epistemic action is not mere description. What distinguishes it from description is that identifying mechanisms is the eventual goal. This mission is not very different from the mission of Causal Role analysis of function. The parts of a complex structure are not described simply to sort out the architecture of the structure: the long-term goal is to understand how those parts function together. To some extent 'mechanism' has supplanted the place of 'function' in contemporary biology. Similar to the arguments made by Amundsen and Lauder (Amundson and Lauder, 1994, 450), with *system* biology we are describing experimental strategies actually in active use.

It may be objected that this term, *system* biology, is itself obvious and tautological. If the research design is governed by the system being studied, then how could you be doing anything but "system biology"? My point is that systems biology does not describe this kind of research adequately because it brings with it the necessity for a more complex experimental project than I wish to consider. The experimental project characteristic of the systems biology with its complicated diagrammatic model of the system completely described mathematically as preferred output is more complex than the style of research a lot of biological/biomedical researchers are actually doing. Since I believe that this omics research involves a new, if not revolutionary, epistemology, I need to give the type of research a distinctive and accurate name. Thus I propose '*system* biology'.

2.3 PDQ proteomics

As recounted in detail in Chapter 1, the first use of the term "proteomics" dates to 1995. In the mid-1990s proteomics, initially defined in relationship to 'genome' as the set of all proteins generated by a genome, began to gain traction among researchers. The early attitude of molecular biologists toward proteomics as a kind of functional genomics is misleading, partly because genomics can in functionalized in a number of

different dimensions but mainly because proteomics is actually quite different from genomics.[25]

Researchers in proteomics frequently describe their technology as 'hyphenated': several different technologies are tacked on to each other in series. The typical proteomics experiment involves three broad phases done *back-to-back-to-back*. In current jargon the phases are known as stages of work-flow. The first is the preparation of the biological specimen (notably separating proteins in the specimen). The next is analysis of the resulting sample(s) by mass spectrometry. The last step is identification of proteins by bioinformatics methodology. Proteins in the biological specimen may be separated by two-dimensional electrophoresis (2DE) where proteins migrate differently in a gel according to their mass and charge or by a more complicated (and less intuitive) process called liquid chromatography (LC, more commonly by "high-performance" LC: HPLC) (Kocher and Superti-Furga, 2007, 811). Then the proteins (by this time maybe broken down into smaller pieces known as peptides) are converted to gas phase ions by either electrospray (ESI: electrospray ionization) or matrix-assisted laser desorption/ionization (MALDI). At the present time the coupling of HPLC to ESI is preferred for most applications, but early on MALDI was prevalent.[26] Mass spectrometry has been used by biochemists since the 1960s, but its application to protein chemistry is largely dependent on the development of effective methods to ionize proteins. As Rüdi

[25] While this is my own considered opinion based on my experience and what I have read, Jürgen Cox and Matthias Mann spell the differences out in plain language: "proteomics is different from genomics in almost every respect" (Cox and Mann, 2011, 274). Genomics measures genotype (proteomics: phenotype); genomics is founded on the glorious history of genetics (proteomics: on the glorious history of biochemistry); genomics has "generic, scalable, constantly refined technologies" aimed at nucleic acids (proteomics: mass spectrometry and allied technologies—initially limiting but now all improving) (Cox and Mann, 2011). It is possible that conflation of proteomics with genomics arises from the initial vision of McKusick and Ruddle when they invented the term genomics, namely that genomics involves mapping genes and determining their biological significance (McKusick and Ruddle, 1987, 1-2), although a close reading of their editorial does not provide convincing evidence that their notion of genomics specifically included the realm of proteins, as well as genes.

[26] Our 'classic' method was IMAC-2DE MALDI Qq TOF MS. This is deciphered as immobilized metal affinity column (to capture the metal-binding proteins of interest), then two-dimensional electrophoresis (to separate proteins, [conversion to peptides by treatment with trypsin], then matrix-assisted laser desorption/ionization (to convert peptides to gas phase ions), then quadripole (=Qq) time-of-flight (TOF) mass spectrometry (MS). All this will be reviewed again later in this chapter in somewhat greater detail.

Äbersold and Matthias Mann point out in their relatively early review of proteomics: "By definition, a mass spectrometer consists of an ion source, a mass analyser that measures the mass-to-charge ratio (m/z) of the ionized analytes, and a detector that registers the number of ions at each m/z value" (Aebersold and Mann, 2003, 198). For the purposes of this dissertation, greater detail on mass spectrometry is probably not needed, except to point out that the m/z value permits identification of the protein or peptide[27], based in large part on comprehensive databases of gene sequences. Mass spectrometry is currently so sophisticated that it provides reproducible, highly sensitive detection of proteins in a biological specimen (Domon and Aebersold, 2010, 710). Computer-based bioinformatics is highly technical and methodological design is critically important (Domon and Aebersold, 2006, 1921). In the past ten years analytical technology for performing proteomics experiments has improved substantially (Cox and Mann, 2011, 275). Current methods facilitate quantification of proteins in a proteome and identification of post-translational modifications of proteins. The important and exciting finding is that post-translational modifications are more common and more varied (less canonical) than previously thought (Cox and Mann, 2011, 284). They play an important role in regulation of intracellular processes.

2.4 Wilson disease

The biological problem which forms the stimulus for investigating the epistemological questions posed in this dissertation arises from attempting to apply omics methodology to a human disease conventionally called "Wilson disease". Originally called 'hepatolenticular degeneration', Wilson disease was first analyzed in detail by the British neurologist Kinnear Wilson in 1912 as a progressive degenerative neurological disorder accompanied by liver cirrhosis (Wilson, 1912, 295). He recognized

[27] In this notation the m stands for mass and the z for the charge number of an ion, which operationally is the number of electrons lost when a molecule is converted to a charged species (that is, an ion: "ionized analyte[s]" in the quotation cited). This is part of the analytical procedure for mass spectrometry. However, since the number of electrons lost is generally only 1, the denominator drops out arithmetically. Thus m/z in most cases is equivalent to molecular mass (see http://www.shimadzu.com/an/gcms/support/faq/mz.html, accessed 7 July 2014) (Anon.: Shimadzu).

that it was familial; however, in 1912 the main consideration was whether it was due to an infectious agent, not whether it was genetic.[28] Wilson provided extensive data against an infectious etiology, but he had no firm concept of the pathogenesis. What is extraordinary is the rapid rate at which basic and clinical research about Wilson disease has progressed (Gitlin, 2003, 1868; Ala et al., 2007, 397; Roberts and Cox, 2012, 1110). It mirrors the rate at which biological/biomedical research has advanced in the 20th century. Over approximately 90 years, Wilson disease has been transformed from an invariably lethal disease to one which, when successfully treated, is compatible with a normal active life-span. The broad stages of this rapid research progress comprise disease description; clarification of the clinical physiology relevant to the disease; development of various life-saving treatments; identification of the gene mutated in Wilson disease, structural studies of its gene product; and investigation of the cellular physiology of copper, including metalloproteomics as research method suitable for characterizing a network of intracellular proteins involved in copper disposition—that is, in how copper is handled in a cell (Roberts, 2012, 634). Moreover, identifying the human genes associated with the two genetic/metabolic disorders of copper disposition enhanced renewed interest in biological aspects of metals broadly, including in microbes and plants. Identification of proteins involved in copper disposition in bacteria was first reported in the early 1990s (Bull and Cox, 1994, 248).

2.4.1 Description of the clinical disease

Wilson disease is an inherited disorder of intracellular copper handling. The pattern of inheritance is autosomal recessive. It is rare: it occurs world-wide with an average prevalence of approximately 30 affected individuals per million population. It can present clinically as strikingly different forms of disease: liver disease, degenerative neurological disorder, or sometimes as purely psychiatric illness. Younger patients tend to have prominent hepatic disease; older patients tend to present with neurological

[28] In fact, the whole notion of "inborn errors of metabolism" described in detail by Garrod, appeared only in 1908, a just a few years before Wilson's paper. We might reasonably conjecture that Wilson was working on his meticulous disease description at the same time as Garrod's ideas were in development; however, the work of the previous three to four decades on infectious agents causing diseases really dominated thinking about disease pathogenesis.

problems. For individuals who develop symptomatic liver or neurological disease, criteria for diagnosis include very low serum ceruloplasmin (a protein produced in the liver and requiring copper incorporated into the protein in order for it to function as an enzyme); deposition of copper in the cornea as "Kayser-Fleischer" rings (named after the two German physicians who first described them in the very early 1900s) best seen on slip-lamp examination of the eyes; and elevated 'basal 24-hour urinary copper excretion', that is, abnormally high amount of copper excreted in the urine each day, in the absence of administering a drug to increase urinary excretion of copper. The concentration of copper in liver tissue can be measured on a specimen of liver obtained by needle biopsy. In Wilson disease it typically exceeds 250 μg/g dry weight of tissue (at least five times the normal liver tissue copper concentration). Clinically-evident Wilson disease is fatal if not treated. With consistent effective treatment, especially if it is begun at the earliest sign of abnormality, the outlook for a normal healthy life is excellent. Several treatment options are currently available, and the vast majority of patients do well on pharmacological treatment. Both hepatic and neurological damage can regress substantially or resolve completely on treatment. With consistent and effective medical treatment, liver transplantation is hardly ever required.

2.4.2 Molecular biology of Wilson disease

The gene *ATP7B* was identified in 1993 (Bull et al., 1993, 333; Tanzi et al., 1993, 344). At the present time it is the only gene whose mutations are known to lead to Wilson disease (other genetic mechanisms may exist).[29] The *ATP7B* gene, located on human chromosome 13 (specifically at 13q14.3), is large (80 kilobases, 21 exons). It is approximately 57% homologous to the gene *ATP7A* mutated in the extremely rare disorder of total-body copper deficit, Menkes disease. The two major proteins encoded

[29] In early 2013 a disorder was described from Quebec and Italy which has some clinical features similar to those of Wilson disease. This condition is known as MEDNIK syndrome (Martinelli et al., 2013, 873). The genetic basis has been determined and is due to mutations in the *AP1S1* (whose gene product participates in the assembly of an important intracellular protein called clathrin). I would regard it currently as inaccurate to say that mutations in this gene can lead to Wilson disease, mainly because the neurological features are very different from those of Wilson disease, but this newly identified disorder does point to the material possibility of more than one genetic basis for Wilson disease or Wilson disease-like disorder(s).

by these two genes are metal-transporting P-type ATPases. Similar genes and their encoded proteins have been identified in mice and rats (providing model organisms for Wilson and Menkes disease), also in yeast (permitting certain specialized *in vitro* experiments) and bacteria. Similar P-type ATPases are also found in plants, notably in the model organism *Arabidopsis thaliana*. The chemical structure of the encoded protein, ATP7B, known as the Wilson ATPase, has been investigated by homology modeling based on the enzyme SERCA1 (Fatemi and Sarkar, 2002, 346) and by nuclear magnetic resonance (Dmitriev et al., 2006, 5303-4; Fatemi et al., 2010, 8468).

To date, more than 500 mutations of *ATP7B* have been identified. Approximately 80% of patients are compound heterozygotes, carrying one copy of each of two different disease-causing mutations. The search for genotype-phenotype correlations in this clinically heterogeneous disease has been disappointing. I conclude that mutations which totally eliminate, or severely decrease, functional Wilson ATPase produce early severe disease, mainly as liver disease (Thomas et al., 1995, 212; Wilson et al., 2000, 719; Merle et al., 2010, 1 of 6; Okada et al., 2010, 1232). Findings in animal models of Wilson disease support this view. The *Atp7b*[30] knockout mouse, in which *Atp7b* is totally non-functional, has more severe, and more rapidly progressive, disease than the spontaneous mouse model ('toxic milk') which is homozygous for a point mutation.

The Wilson ATPase must have at least two physiological functions within hepatocytes (parenchymal liver cells). It somehow contributes to the incorporation of copper into ceruloplasmin. In Wilson disease the level of ceruloplasmin is typically subnormal, mainly because copper has not been incorporated into that protein when it was manufactured in the liver. Secondly, it is part of the mechanism for excreting copper from the hepatocyte into bile. As shown in various *in vitro* cell models, the Wilson ATPase is found near the trans-Golgi network when intracellular copper concentrations are relatively low and then it moves to the region of the bile canalicular membrane when intracellular copper concentrations are elevated. Evidently the Wilson

[30] The convention is that only with mouse genes the designation has the first letter capitalized, rest in lower case. With humans, rats, dogs, etc., the entire gene designation is capitalized.

ATPase has a third function: monitoring hepatocellular copper concentration. The Wilson ATPase might have other functions. For example, it is known that the Wilson ATPase is capable of interacting with platinum when it turns up in a hepatocyte as an anti-neoplastic agent, cisplatin (Safaei et al., 2012, 8). Of course, this action might not be regarded as physiological.

2.4.3 Model organisms for Wilson disease

Studying disease mechanisms in Wilson disease requires animal models[31]. Overloading the continuous human liver cell (well-differentiated hepatoma) line HepG2 is informative about copper handling (Aston et al., 2000, 367; Seth et al., 2004, 501), but HepG2 cells as such are genetically not a model for Wilson disease. The Long-Evans cinnamon (LEC) rat is an important model but its mutation, a large gene deletion, is not typical of human Wilson disease (Wu et al., 1994, 542). The several mouse models of Wilson disease are complementary. The available *Atp7b* knock-out mouse, where the entire *Atp7b* gene has effectively been excluded, has severe liver disease (Buiakova et al., 1999, 1667-8; Huster et al., 2006, 423). Two naturally-occurring mouse models for Wilson disease, both referred to as 'toxic milk' (tx), have been identified. Importantly, their mutations are point mutations and thus resemble the kind of mutations found in most people with Wilson disease (Theophilos et al., 1996, 1621; Coronado et al., 2001, 793). My laboratory showed that the tx-j mouse has slowly progressive liver disease (Roberts et al., 2008, 54). We investigated tx-j hepatocellular mitochondria which exhibit the same structural abnormalities on transmission electron microscopy (TEM) as

[31] It may be opportune to point out that 'model' is a problematic term in this analysis. Most of the time when I talk about models, I am referring to model organisms (such as mouse or rat models of Wilson disease), a concept well-described in recent literature in philosophy of biology (See Section 5.2.3.4 and references therein). Sometimes the model is not an actual organism but an *in vitro* cell culture preparation, usual a continuous cell line. However, we have already seen that the desired output of systems biology is a model of a system, and this model is typically a schematic diagram of the system described mathematically. The discussion gets more complicated when we encounter Hempel's, and subsequently, Deborah Mayo's use of 'model' as a technical term to describe certain entities in a philosophical analysis of scientific reasoning. Finally, a model can be a pictorial representation of the biological thing or process being investigated, and it may be rather general or even provisional. This borrows of Ronald Giere's concept of model in scientific reasoning. Such a notion of 'model' may have practical affinities with certain notions of what a hypothesis is, another issue to be explored further in Chapter 4.

found in human Wilson disease (Sternlieb, 1968, 354). Why these structural abnormalities occurred was unclear. Using non-denatured polyacrylamide gel electrophoresis (PAGE), we showed that there is a selective loss of Complex IV of the oxidative-phosphorylation (OX-PHOS) chain in the mitochondria in tx-j hepatocytes and loss of mitochondrial function compared to normal C3H mouse, as control, over the first 6 months of life (Roberts et al., 2008, 60). Others have shown that the *Atp7b* knockout mouse has various metabolic abnormalities besides defective copper disposition: specifically, these mice have decreased biosynthesis of cholesterol (Huster et al., 2007, 8351).

2.5 Omics applied to Wilson disease: metalloproteomics

Given the progress in understanding disease mechanisms in human patients and in laboratory animal models of Wilson disease, we might reasonably ask what we expect to accomplish by applying an omics approach to Wilson disease. The really broad goal for such a research project is a comprehensive understanding of the pathophysiology of Wilson disease. Integrating basic and clinical science is the best way to proceed. Clinical research seeks to collect and analyze extensive clinical experience, currently through the development of large databases and registries on a multi-centered and/or multi-national basis. Basic science research uses various—indeed, numerous—conceptual and animal models to investigate the pathophysiology of Wilson disease. The physiological context for such disease-orientated research is how metals contribute to cell and tissue function. Metalloproteomics can make important contributions to these research efforts, as can transcriptomics and metabolomics. A practical output of this high-throughput research might be identification of so-called “biomarkers” to diagnose Wilson disease or determine clinical severity.[32]

[32] Terminology around ‘biomarker’ is extremely current. Biomarkers are all the rage. It might be pointed out that we have identified biochemical features as useful handles on the diagnosis and staging of disease for a long time. For example, extremely low serum ceruloplasmin is veritably a biomarker of Wilson disease, all the more so when the assay is based on the enzymatic action of the protein ceruloplasmin, which in turn depends on having copper incorporated into the molecule, which in turn depends on the normal action of the Wilson ATPase, which is the gene product of the gene abnormal in Wilson disease.

2.5.1 Narrow-focus approach to the hepatocellular Cu-metalloproteome

We can take a narrow-focus (bottom-up) approach to this effort. In this example, much of the analysis depends on information gained from another animal model, the Bedlington terrier. This breed has a genetic copper disorder (Haywood et al., 1996, 230) and was long regarded as a candidate model for Wilson disease. However, Bedlington terrier copper toxicosis is not canine Wilson disease (Dagenais et al., 1999, 755). It typically involves a large deletion of the gene *COMMD1* (van De Sluis et al., 2002, 165). The protein COMMD1 has numerous functions (Sarkar and Roberts, 2011, 23): among them is the ability to interact with the Wilson ATPase (Tao et al., 2003, 41593). COMMD1 is negatively regulated[33] by XIAP (Maine et al., 2009, 601) and XIAP binds copper (Mufti et al., 2006, 775). XIAP levels decrease in tissues with copper overload (Mufti et al., 2007, 172). We found that COMMD1 expression in tx-j liver rises over the first six months of life, as shown by real-time PCR (Roberts et al., 2007, 923). At the same time, copper accumulates in the liver so that liver copper concentrations are high, and hepatocellular XIAP concentrations progressively decrease. We can put this all

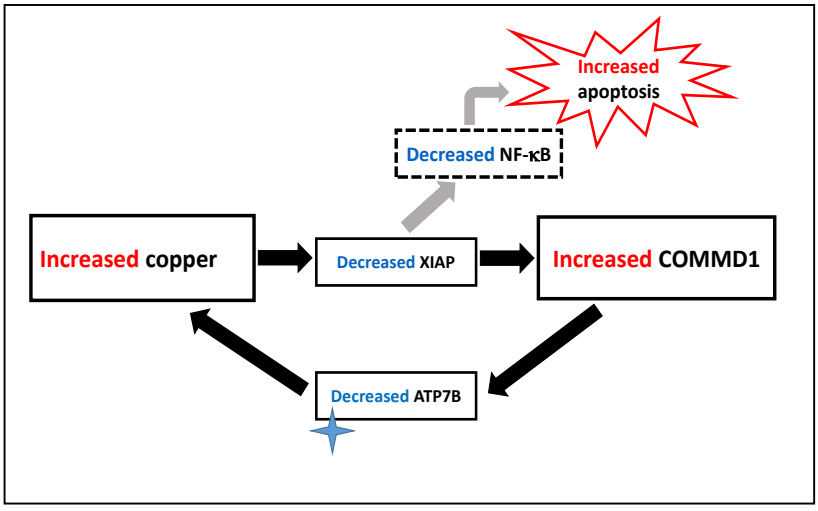

Figure 2.1 Feedback loop of a few hepatocellular proteins involved in disposition of copper

However, when we speak of biomarkers at the present time, the term tends to connote something more complicated, typically a composite signature of genetic or other biochemical features.

[33] 'Negative regulation' is not exactly an intuitive concept, at least not for most people. To say that COMMD1 is negatively regulated by XIAP means that when XIAP levels are low, then COMMD1 levels are high. The negative-regulator relation between X and Y means that if X is high, then Y is low, and similarly if X is low, then Y is high.

together by combining our findings with those of De Bie and colleagues who showed that COMMD1 is a negative-regulator of Wilson ATPase (de Bie et al., 2007, 681). We get a loop, ostensibly a self-amplifying cycle (Figure 2.1), which is enhanced if the Wilson ATPase is itself abnormal ('star' then serves as entry point into the loop). Of note, increased hepatocellular apoptosis is a feature of Wilson disease. Admittedly, this is a very limited system, just a tiny facet of copper disposition in hepatocytes. Data from several different laboratories working on different facets of this problem for several years had to be combined in order to complete this analysis.

2.5.2 Big-picture (top-down) analysis of the hepatocellular Cu-metalloproteome, namely, classic metalloproteomics

Metalloproteomics is a type of proteomics, and we can regard it as a subdiscipline of proteomics.[34] My concept of metalloproteomics, developed in collaboration with Dr. Bibudhendra Sarkar at the University of Toronto, is a big-picture (top-down) systems strategy: it is a way of knowing about a system. Thus the history of our research collaboration has important conceptual implications and is worth recounting briefly, not only to document the achievement but to identify our working assumptions. Proteomics, as we know it today, really started to develop in the late 1990s. Meanwhile, the mid to late 1990s became a period of very active research on the cellular physiology of copper, arising from the cloned genes for copper disposition described in 1993 (human: *ATP7A* and *ATP7B*) and similar discoveries in various model organisms around that time. Sometime between the spring of 1999 and autumn 2000 Bibudhendra Sarkar and I began to discuss whether a proteomics strategy could be devised to examine how copper is handled in the liver. Bibudhendra Sarkar is an eminent bioinorganic chemist in Toronto, best known among clinicians for his work on treatment of Menkes disease and for establishing key aspects of the structural biology of the Wilson ATPase once the *ATP7B* gene had been cloned. I suggested that continuous human liver cell (hepatoma) lines like HepG2 and the tx-j mouse model of

[34] Similarly, another subdiscipline of proteomics, developed since we invented metalloproteomics, is "phosphoproteomics" which is limited proteomes where the constituent proteins have been phosphorylated.

Wilson disease available from the Jackson Laboratories would be suitable for such experiments. Fortunately we had a superb PhD chemist specializing in mass spectroscopy available to work with us, Ye-Min She. At a meeting of the Society of Experimental Biology in Canterbury, UK, in the spring of 2001 I heard a presentation about an emerging 'SELDI' technology from Ciphergen Laboratories, but we decided to use the more complex and versatile technology described above: protein separation by two-dimensional protein electrophoresis followed by classic mass spectroscopy followed by bioinformatics analysis of results. The additional step we introduced was to run the specimen through a column to which the metal of interest was affixed (called an IMAC column[35]) so that the metal-binding proteins would stick to the column. The column could be gently washed off and then the proteins in this 'wash' were subjected to proteome analysis. By November 2001 we had enough data for a poster at the Hospital for Sick Children Research Institute Scientific Retreat, and in peer-reviewed publications in 2002 these data were mentioned as relevant 'unpublished data'. Dr. Sarkar presented our findings to the 3rd International Meeting on Copper Homeostasis and its Disorders: Molecular and Cellular Aspects in Ischia, Italy, in October 2002, and I presented our work at the 1st World Congress of the Human Proteome Organization (HUPO) in Versailles, France, in November 2002. Our abstract introducing this research strategy was thus published in 2002 (Roberts et al., 2002, 662), with the very word 'metalloproteome' leading off the title.[36] Our first formal paper defining metalloproteomics appeared in 2003. Also in 2003 we published a paper operationalizing the concept of metalloproteomics by showing detailed functional characterization of copper-binding by PDI, detected through our metalloproteomics

[35] The immobilized metal affinity column (IMAC) is a well-established technique with numerous applications (Sun et al., 2005, 650-1). It has some predictable weaknesses. Typical metalloproteins (for example, ceruloplasmin) can pass through the IMAC column because constituent metals are strongly bound within the metalloprotein: metal exchange does not occur and consequently no binding to the column. Likewise, proteins with high affinity for metal at their functional binding-site may pass through with an occupied binding-site. Contrariwise, there may be non-specific binding of proteins to the column: these have no functional significance whatsoever but yield false positive results unless disconfirmed by further characterization. IMAC can be modified so that it operates for surface enhanced laser desorption ionization (SELDI).

[36] This is the first published instance of the word 'metalloproteome' in the scientific literature.

work. We next examined the Cu-metalloproteome[37] of HepG2 cells fractionated into cytosol and microsomes and conducted preliminary experiments with tx-j mouse liver. Our further laboratory research in metalloproteomics ended up limited by career changes among several of our team members and the vagaries of research funding.[38] Metalloproteomics[39] has flourished, extending to numerous metals and organisms.

2.5.2.1 Definition of metalloproteomics

When we first started implementing this methodology 10-12 years ago, we knew it was a powerful general concept. We were delighted when, unexpectedly in 2008, metalloproteomes were called the "functional unit[s] of transition metal biology" (Thiele and Gitlin, 2008, 145). One of the noteworthy developments since our initial work has been has been the varied and variable uptake of metalloproteomics: broadly affirmative, though often with a loss of conceptual features to which we attach great importance. This variability within the scientific community is evident from how metalloproteomics is defined. Review of recent literature reveals several definitions of metalloproteomics. They are not identical. One definition is that metalloproteomics is "the structural and functional characterization of metalloproteins on a genome-wide

[37] A metalloproteome could be limited to one metal, or could be determined for several specified metals, or could be determined for every metal present. Our work was limited to one metal at a time. I designate the metalloproteome by prefixing the chemical abbreviation for the metal investigated. Thus: Cu- for copper, Zn-for zinc, Ag- for silver, Pt- for platinum, and so forth. When speaking of a metal in any general way, I avoid the abbreviation.

[38] It is worth noting that we failed to get peer-reviewed funding for this work from either the Canadian Institutes of Health Research (CIHR) or the National Institutes of Health (NIH), and therefore we devoted discretionary and private funding to this research project. Moreover, at its first presentation to an audience of peer-group basic researchers in copper physiology in Ischia in 2002, this work met with scepticism, and even some derision, including from those who authored the *Nature Chemical Biology* paper five years later (Thiele and Gitlin, 2008, 145) (see next paragraph in text).

[39] It is fair enough to say that it has flourished, but not necessarily as we conceived it. In 2010 an impressively large book on metalloproteomics by Permyakov was published (Permyakov, 2009). Scientific conferences on metalloproteomics have taken place. They include, for example: 3rd BioXAS Study Weekend - "Metalloproteomics" sponsored by Synchrotron SOLEIL, Saint-Aubin, France, August 2007, chairman: I. Ascone, Synchrotron SOLEIL and University of Rome "La Sapienza"; 15th International Conference on Biological Inorganic Chemistry (ICBIC 15), Vancouver, August 2011 where 'Metalloproteomics and Genomics' was a conference topic; and on the programmes of three of the prestigious Gordon Research Conferences in 2006, 2012 (discussion leader: Tom O'Halloran) and 2013 (discussion leader: Val Culotta) (http://www.grc.org/search.aspx, accessed 23 April 2014) (Anon.: Gordon Research). I have speculated on the future directions of metalloproteomics (Roberts and Sarkar, 2014, 429).

scale" (Shi and Chance, 2011, 144). It reflects the widely accepted notion that the metalloproteome of a given system is constituted by all the metalloproteins in that system. This definition is attractive because it is linguistically consistent: it establishes a correspondence between metalloprotein (strictly defined as a protein wherein a metal is an integral part of its structure or as a protein whose function is conferred by a metal) and metalloproteome. Such a definition is too narrow. A way to assess the inadequacy of this class of definitions is to consider how definitions of omics disciplines are generated. By convention, an "-omics" discipline examines its "-ome". Genomics examines genomes; proteomics examines proteomes; metabolomics examines metabolomes. Specifying what a metalloproteome is should elucidate what metalloproteomics is. We defined metalloproteome (for any specified system) as "the set of proteins that have metal-binding capacity by being metalloproteins or having metal-binding sites" (She et al., 2003, 1306). Accordingly, a metalloproteome includes metal-binding proteins (that is, proteins which possess metal-binding sites) as well as metalloproteins. Since many metalloproteins have been identified (with the possible exception of microbial/bacterial metalloproteins which continue to be identified anew), a really important focus for metalloproteomics is to identify metal-binding proteins. These may be "dedicated" metal-binding proteins, such as transport proteins and metallochaperones, or proteins which have the metal-binding capability as an incidental or secondary function. The latter metal-binding proteins are equally important as the dedicated metal-binding proteins, and they likely make critical contributions to cell function; however, they may be difficult to identify. Some definitions of metalloproteomics focus exclusively on metal-binding proteins: "structural and functional characterization of metal-binding proteins and their structural metal-binding moieties" (Jakubowski et al., 2004, 1), but again this definition is probably too limited. Another approach is to abandon the term metalloproteome and consider the 'metal-binding proteome', but this concept is vague at best (actually, incoherent since it does not make sense for a proteome to bind metals) and eliminates metalloproteins. Finally, one prominent definition of metalloproteomics is problematic because it adopts a

specialized meaning for 'metalloprotein' (Szpunar, 2005, 445): basically an equivocation on the meaning of 'metalloprotein' whereby that term suddenly includes metal-binding proteins as well. This definition is problematic not merely because it involves a logical fallacy (equivocation),[40] but because it is found in a highly influential, often cited review paper.

From the beginning, when my laboratory started collaborating with the Sarkar laboratory, we regarded metalloproteomics as a dynamic methodology for investigating how copper is handled in cells. It seemed manifest that cellular *patho*physiology of a disease could not be investigated without having a comprehensive grasp of the relevant cell physiology; this in turn was dependent upon the relevant structural biology of the cell. We defined metalloproteomics as follows:

> Metalloproteomics seeks to systematically identify large sets of proteins associated with metals and analyze their regulation, modification, interaction, structural assembly, and function as well as their involvement in physiological processes (including development) and in disease states.

This definition, which represents my own more recent elaboration of a definition we published in 2006 (Kulkarni et al., 2006, 2412), has certain implications. One implication is that it is clearly a systems strategy. A second implication is that, given wording deliberately chosen as 'proteins associated with metals', the set of proteins assembled is not limited to metalloproteins, as has been the preference of numerous investigators subsequently. However, a weakness of our work may be that, while we were clear about the comprehensive nature of our definition, we did not stress the dual nature of metal-associated proteins in our definition of metalloproteome directly and insistently enough. Moreover, our own work focused more on metal-binding proteins, the innovative aspect of the system we were studying. An emerging aspect of

[40] Scientists may not be too worried about logical fallacies (perhaps to their loss), but this definition of metalloprotein ("All the proteins which form complexes with metals which are thermodynamically stable in a given chemical environment and kinetically stable on the time scale of the analytical procedure, are referred to as metalloproteins.") (Szpunar, 2005, 445) actually does not make biochemical sense since on these criteria common protein albumin would be called a metalloprotein, which is contrary to the ordinary received definition of what a metalloprotein is. I am grateful to Bibudhendra Sarkar for pointing out the biochemical subtleties here. In the paper cited the equivocation is not subtle because this redefinition of 'metalloprotein' is placed in a section headed "Metal-binding proteins".

metalloproteomics, indirectly included in our definition but not in others, is that aggregates of proteins, including but not limited to homo- or heterodimers, may function as a unit to bind a metal. None is itself necessarily a metalloprotein nor does any of them bind a metal individually, but the protein aggregate demonstrates the property of metal-binding.

Moreover, evident from the definition above, we view metalloproteomics as an exercise in structural biology in its broadest sense. Structural and functional features of each member of any metalloproteome must be contextualized in cell/tissue physiology and in relevant disease states. Our emphasis on contextualization of results sounds vaguely as if influenced by a feminist epistemology, but it stems from two issues relating to experimental design. One was the worry that we might encounter false-positive identifications through the IMAC technique for capturing metal-binding proteins, and the second—reflecting our commitment to translational research—was the intuition that the metalloproteome as a list of proteins was not really valuable or interesting until you checked how these proteins functioned in the system of interest. According to our concept, a metalloproteome can be defined for a specific metal of interest or—as has become possible with newer technology—the metalloproteome inclusive of all metals in a system might be characterized. Similarly a few selected metals might be investigated in a single metalloproteome. In any case, the critical stipulation is that the characteristics of the system under investigation must be specified in detail: the specific organelle or cell or tissue; the species; the age or stage of development; whether the biological context is normal or abnormal; and so forth.

2.5.2.2 Early work on hepatocellular metalloproteomics of copper

I wish to locate our work briefly in the overall development of metalloproteomics. As I have already pointed out, the rationale for our research was the belief that we need to know what proteins are associated with copper in hepatocytes before we can figure out whether their expression changes in Wilson disease. Our own main interest was for copper-binding proteins, largely because we were starting from our focus on the structure of the Wilson ATPase and the action of the metallochaperone ATOX1, involved in mediating the binding of copper to the Wilson ATPase. We knew the

rudiments of an intracellular protein network for handling copper in hepatocytes: a plasma membrane uptake protein (CTR1), various intracellular interacting proteins (such as the metallochaperone ATOX1, and metallothioneins as storage proteins), and exit routes involving either incorporation into the protein ceruloplasmin or excretion into bile. Thus we could postulate that such a network of proteins existed (since we had identified a few of them), but we were keenly aware that our knowledge of this network was incomplete. In our initial studies we used HepG2 cells and other continuous human liver cell lines and subsequently we used tx-j liver and normal C3H mouse liver as comparator. Using our IMAC-2DE MALDI Qq TOF MS (see Section 2.5.2 to spell out the shorthand) method, we examined the Cu- and Zn-metalloproteomes in various continuous human liver cell lines as well as a specimen of normal human liver by way of non-transformed control (She et al., 2003, 1312). In fact we found broad similarity among the cell lines, although there were some interesting differences. Importantly we established for the first time a post-translational modification dependent on copper. Using the same IMAC-2DE MALDI Qq TOF MS with preparations of cytosol and microsomes from HepG2 cells, we went on to identify 48 proteins in the cytosolic[41] Cu-metalloproteome and 19 proteins in the microsomal Cu-metalloproteome (Smith et al., 2004, 836). Although this was only a first approximation which certainly underestimated the composition of these Cu-metalloproteomes, this work stands as a proof-of-principle experiment. We found extremely interesting proteins in these Cu-metalloproteomes, such as PDI, which we showed conclusively to be capable of binding copper (Narindrasorasak et al., 2003, 408-9). Our HepG2 Cu-metalloproteome included the protein peroxiredoxin, which binds copper by forming a copper-binding site through interaction of homodimers relating to each other via a process of intermolecular tail-swapping (Hirotsu et al., 1999, 12336), thus creating a “virtual” copper-binding site not evident from its primary structure, which is the simple linear chain of amino acids. This is an example of a protein aggregate, albeit a very special kind of protein aggregate,

[41] Cytosol is a laboratory preparation but it is composed of proteins and other chemical components found in the cytoplasm of a cell; likewise, microsomes are a laboratory preparation and constituted mainly of endoplasmic reticulum.

which is capable of binding a metal. We then showed that expression of peroxiredoxin-3 increases markedly over the first seven months of life in the Wilsonian tx-j mouse (Roberts et al., 2008, 59), thus contextualizing the findings from a proteomics experiment in terms of cell biology. In further studies, using isotope-coded affinity tag (ICAT) labeling, we quantified proteins in the hepatic mitochondrial Cu-metalloproteomes in tx-j mice to identify quantitatively important changes associated with the Wilsonian disease state.

2.5.2.3 Other studies relating to metalloproteomics and Wilson disease

Few comparable Cu-metalloproteome studies have been published. A proteomics study performed in young *Atp7b* knock-out mice revealed little difference between the Wilsonian mouse and control (Huster et al., 2007, 3849-50). Specifically, hepatic metallothioneins were induced in the *Atp7b* knock-out mouse, but no change was found in the major copper-containing cytoplasmic metalloprotein SOD1 in the liver. A more recent general proteomics study (not metalloproteomics, also without dietary copper-loading) of LEC rats, the rat model for Wilson disease, showed early progressive changes in protein expression reflecting mitochondrial injury and subsequently apoptosis (Lee et al., 2011, 3701-2). The investigators used a selective approach to proteomics and identified only those proteins apparently increased on 2DE. Most of the work in proteomics, and also with transcriptomics, has looked at copper-responsive proteomes. Simpson and colleagues studied the hepatic proteome in the copper-tolerant Cambridge sheep on a high-copper diet. Although their experimental design differed from ours, they found proteins similar to those we had identified. Of note they also identified PDI in this copper-responsive proteome (Simpson et al., 2004, 528). An investigation of copper-responsive proteins in HepG2 by SELDI technology revealed several proteins, uncharacterized but most likely metallothioneins (Roelofsen et al., 2004, 735). A transcriptomics study of copper-responsive transcripts in HepG2 cells showed that various genes were up-regulated but COMMD1 was down-regulated (Muller et al., 2007, 500). Copper-chaperones did not appear to be transcriptionally regulated. After copper-depletion of the HepG2 cells, the transcriptome looked very similar. A transcriptomics study of the *Atp7b* knock-out mouse revealed decreased

expression of genes relating to lipid metabolism and increased expression of genes relating to cell cycle control (Huster et al., 2007, 8347-8).

2.5.3 Progress with metalloproteomics in the past 7-10 years

2.5.3.1 Methodological advances

In the past 5-7 years, in parallel with advances in the technology for proteomics, there have been important advances in the technology for metalloproteomics so that currently much more comprehensive experiments can be performed. These represent major technical improvements on our now classic method IMAC-2DE MALDI Qq TOF MS. Sample preparation is critical in any case: addition, or not, of antioxidant chemicals to the tissue/cell homogenate, denaturing versus non-denaturing electrophoretic methods (da Silva et al., 2010, 388). A critical issue in evaluating analytical advances for metalloproteomics is to differentiate between the inherent problems of examining metalloproteins and those of characterizing metal-binding proteins. In other words, how we define metalloproteomics has implications beyond denoting the field of experimentation, a perhaps unusual situation where the definition really matters. I am not the only person in this field identifying this problem: it was also pointed out recently by others (Barnett et al., 2012b, 3311-2). For metalloproteins, high-throughput X-Ray absorption spectroscopy is a promising new method (Shi et al., 2011, 898). For metal-binding proteins, technical challenges persist. Although IMAC continues to be employed and regarded as an established standard methodology (Sun et al., 2005, 649-50), methods to capture metal-binding proteins need further development. For example, 2DE followed by the analyses of the same spot both by Inductively Coupled Plasma-Mass Spectrometry (ICP-MS) for metals and by Matrix-Assisted Laser Desorption/Ionization Time of Flight Mass Spectrometry (MALDI TOF-TOF MS) for protein identification may be promising (Gomes et al., 2011, 238).

Some newer analytical methods permit investigation of more than one metal simultaneously. Methods for direct analysis include non-denatured PAGE followed by synchrotron radiation X-ray fluorescence (SRXRF) or proton-induced X-ray emission (PIXE) (Lobinski et al., 2006, 1592; Ortega et al., 2009, S649). Metal Isotope native RadioAutography in Gel Electrophoresis (MIRAGE) involves culturing cells with a stable

isotope to label proteins of interest, non-denaturing 2DE, autoradiography, and then mass spectrometry: it has been used to investigate the Cu-, Zn- and Fe-metalloproteome in *E. coli*. Its advantage is that it permits visualization of all proteins containing the metal of interest in a specimen non-destructively (Sevcenco et al., 2012, 1967).

Another important advance is the development of microanalytical techniques so that a specific single cell can be studied. The cell can be defined as the system under investigation and its tissue environment can be taken into consideration. Among these techniques, laser ablation (LA) of the tissue is the most developed. A single cell in a tissue specimen can be selected and removed for proteomic analysis. The cell can then be analyzed by extremely sensitive techniques such as micro-PIXE. Recent work, relevant to environmental toxicity of manganese and indirectly to Parkinson disease, has combined micro-SXRF (Synchrotron X-Ray Fluorescence) for imaging and micro-XANES (X-ray Absorption Near Edge Structure) to determine the intracellular location and oxidation state of manganese at the single cell level in the rat adrenal pheochromocytoma-derived dopamine-producing PC12 cell line (Carmona et al., 2014, 822). As yet, these methods have not been extensively applied to examining how proteins in the mammalian hepatocellular Cu-metalloproteome function within cells.

A strictly bioinformatics strategy has permitted analysis across the three domains of life: Archea, Bacteria and Eukaryotes (Andreini et al., 2009, 1472-3). Copper-binding proteins were found to constitute approximately 1% of proteomes across phyla studied; five organisms had no copper protein (Andreini et al., 2008, 213). Eukaryotes had a greater array of copper-binding proteins than Archea or Bacteria. A shortcoming of this approach, now fully acknowledged, is that it misses what we might call “virtual binding sites”, those which arise through protein folding. For example, when we identified peroxiredoxin in the hepatocellular Cu-metalloproteome, we picked up exactly this sort of copper-binding site. Of interest, as already mentioned, expression of mitochondrial peroxiredoxin-3 greatly increases over the first 6 months of life in untreated tx-j mice (Roberts et al., 2008, 61). However, the bioinformatics method has various advantages, including the ability to examine various metals. A major advantage

of this method is that studies can be performed across phyla, including eukaryotes and prokaryotes.

2.5.3.2 Recent findings

An important feature of our initial work was that we investigated metalloproteomes in continuous human hepatocyte lines, most prominently the near-normal line, HepG2. Several diverse studies have recently been reported where some specialized HepG2 proteome is described, in relation to neoplasia (Slany et al., 2010, 6) or hepatocellular physiology (Van Summeren et al., 2011, 109; Prot et al., 2011, e21268). More importantly, some studies in HepG2 relate to hepatocellular metalloproteomics. With respect to iron (Fe), HepG2 Fe-metalloproteome was examined after iron overload (Petrak et al., 2006, G1059). A proteomics study in diabetic db/db mouse liver showed down-regulation of glucose-regulated protein 78 kD (GRP78) and PDI A3 (Yamagishi et al., 2012, 367); in HepG2 cells oleic acid-induced down-regulated GRP78 (Yamagishi et al., 2012, 369). We found both GRP78 and PDI(A3) precursor in the cytosolic HepG2 Cu-metalloproteome (Smith et al., 2004, 837). This report may point toward other unappreciated relationships between copper disposition and lipid metabolism. Our classic metalloproteomics strategy was utilized to examine the nickel-metalloproteomes in primary cultures of human (skin) keratinocytes and B-lymphocytes. The objective was to identify proteins involved in nickel allergy. A novel finding was that many of these nickel-binding proteins were heat shock proteins or other chaperonins, proteins that supervise protein-folding during protein synthesis like PDI (Thierse et al., 2008, 141). Interestingly, the Ni-metalloproteome in the cells examined was different from the human hepatic Cu-metalloproteome we had reported.[42]

Proteomics of microbes is an emerging area of research (Malmstrom et al., 2011, 2948), possibly due to growing interest in microbiomes. Microbes have complex

[42] It seemed likely that the metal nickel would be interact with a different set of hepatocellular proteins, compared to copper, and these data support that conjecture. It cannot be taken for granted that such metalloproteomes are absolutely distinct. For example, the Wilson ATPase is presumably in the hepatocellular Cu-metalloproteome (even though we did not find it there in our work, due to technical issues) and it is also in the hepatocellular Pt-metalloproteome.

mechanisms for metabolizing or otherwise coping with metals in their physical environment. As mentioned, the structural genes of the human P-type ATPases, *ATP7A* and *ATP7B*, display important similarities to bacterial metal-resistance genes. Metalloproteomics can play an important role in determining how metals are handled by bacteria. In early work along these lines, Cvetkovic and colleagues characterized cytoplasmic metalloproteomes from *Pyrococcus furiosus*: they found numerous metals present but, more importantly, novel proteins interacting with metals (Cvetkovic et al., 2010, 779). Other recent work includes the Hg-metalloproteome in *E. coli* (Gao et al., 2013, 913) and the metalloproteome(s) of a marine cyanobacterium (Barnett et al., 2012a, 3373-4). The Cu- and Zn-metalloproteomes have been characterized in *S. Pneumoniae* (Sun et al., 2011, 3288). Metalloproteomics has been employed as part of investigating copper disposition in bacteria comprehensively (Arguello et al., 2013, 1 of 14).

2.5.3.3 Commentary by way of wrapping up

As an investigative strategy, metalloproteomics is potentially very powerful. With respect to hepatocellular physiology of copper, it has already produced new information of potential value to understanding disease mechanisms in Wilson disease. The metalloproteomes for other metals—iron, zinc, platinum, silver—will likely be equally informative for various diseases. What is conceptually interesting about our work in metalloproteomics is that it anticipated aspects which have gained importance in current proteomics: the notion of a subdiscipline of proteomics; dynamic capture of the protein subset of interest; identification of atypical post-translational modifications of proteins; possible recognition of protein-protein interactions such as aggregates which maybe players in the intracellular disposition of metal; need for functional characterization of proteins identified as being in the metalloproteome; and, not least of all, need for contextualization of metalloproteomics results in the relevant cell or tissue. Compared to many experimental strategies, metalloproteomics—like proteomics, a process of direct interrogation of a biological system—is powerful and efficient and permits investigation of a system's features with respect to metal disposition and metabolism, which previously had not been assigned conceptual importance.

2.6 Conceptual issues revealed by our metalloproteomics research

In the course of doing the research which constituted the development of metalloproteomics, it became evident to me that there were important conceptual problems pertaining to that research. Even if we now regard this research as "proof of principle" rather than definitive—our technology, though definitely state-of-the-art then, appears somewhat limited at the present time—these issues persist. One problem focused on the cell biology itself: it was easy enough to repeat the preparation of cell homogenates (or the separation into cellular subsets such as microsomes and cytosol) because the methods were very standard and we had great expertise at working with these methods, but would the metalloproteome be equivalent each time? In other words, how many replicate experiments were required for the optimal running of the experiment? If the proteins found the second time were different from those found the next (that is, third) time, which were correct? In other words, how could you identify error? How could you validate the experimental results? Indeed the Cu-metalloproteome appeared different from the HepG2 whole cell proteome reported by other workers a year or two previously. Was this a source of worry? These were questions relating to false positives, false negatives, and the completeness/adequacy of the results. We were well aware in the early 2000s that with available techniques it was extremely difficult to identify proteins embedded in intracellular membranes, and thus we assumed that these were underrepresented in our determination of the Cu-metalloproteome. This was a problem of false negatives. Likewise, we suspected that proteins which were present in very low abundance within the cell might not be detected, another kind of false negative. We were highly suspicious that some of the proteins we identified in the Cu-metalloproteome were not actually copper-binding proteins, but proteins which were loosely bound to another protein which was indeed a copper-binding protein, a kind of 'tag along' detection. Such 'tag-along' proteins would also be false positives. Alternatively, such proteins might be part of a protein complex involved in copper-binding. Because metalloproteomics anticipates current techniques in proteomics, such as functional proteomics, these problems retain their importance.

Although these issues can be portrayed merely as problems of specificity and sensitivity (false positives and false negatives) in proteomics experiments, they could also be viewed as actually having to do with the very nature of what we are studying and how we learn about it. This latter is my principal claim. These issues raised questions about the very nature of this type of experimental design: how does it work and how does it get evaluated. They relate to *system* biology as such and the high-throughput methods, mainly omics, currently in use for investigating a biological entity as a system. Among omics disciplines, proteomics provides an excellent opportunity to examine these issues in detail. These are essentially philosophical problems and serve as the stimulus for the philosophical examination here. Failing to get peer-reviewed funding for this work enhanced a sense of intellectual urgency. The clear message was that these innovative non-hypothesis-driven experiments were 'sketchy' (not just risky): they failed to make the grade as proper scientific investigation. They were not worthy science. Indeed if having a hypothesis was the key requirement for qualifying as excellent research, then metalloproteomics did not qualify as biological/biomedical science.

Arising from this experience with metalloproteomics (which has proved to be an excellent representative of contemporary proteomics), I find that my initial questions seem rather simplistic: what is a hypothesis and what is an experiment? If I can show that system-driven research such as proteomics is governed by an epistemic superstructure, how does that epistemic superstructure differ from the HDM? What would an omics epistemology look like and how would it operate? Finally, since I believe that (metallo)proteomics does constitute a truly scientific endeavour, I would like to demonstrate through cogent argument that my belief is correct.

CHAPTER 3 IMPORTANT PRELIMINARY PHILOSOPHCAL ISSUES

The assertion that some scientific research in the biological and biomedical sciences proceeds without a hypothesis as such—that is, is "non-hypothesis-driven"—frequently elicits strong negative or defensive reactions from biological/biomedical scientists. Clearly, if my contention that some biological/biomedical research is non-hypothesis-driven is wrong, then my philosophical project to examine epistemological issues around being non-hypothesis-driven will falter. Recognizing that there are several different kinds of non-hypothesis-driven research provides strong support for the existence of this basic dichotomy. Nevertheless, there are some other arguments for rejecting my proposal to examine non-hypothesis-driven research. In this chapter I will examine the main lines of argument against my contention that biological research can be hypothesis-driven or non-hypothesis-driven. I will demonstrate, however, that these arguments do not carry the day. In order to focus the discussion, I will take metalloproteomics, our subdiscipline of proteomics, as the test case.

3.1 Does metalloproteomics actually have a hypothesis?

As previously mentioned, it is fashionable to characterize research strategies in systems biology as either 'top-down' or 'bottom-up' (Bruggeman and Westerhoff, 2007, 46). Some point out that there is a further middle ground approach, called "middle-out" where you start somewhere in the middle and work both up and down (Kohl et al., 2010, 30; O'Malley and Dupre, 2005, 1272-3), but I will leave this middle-out strategy aside for the present time. [43] I have argued that this top-down/bottom-up terminology is overused and suggested that we use other terms, specifically, "big-picture" and "narrow-focus" as replacements (Section 2.1), in part because a biological researcher rarely starts at the very bottom or the very top. I will employ my version for considering biological/biomedical research in general although referring to the customary terminology of "top-down" and "bottom-up" because these terms are so completely

[43] It will receive some discussion subsequently in Chapter 7 (Section 7.3.2).

entrenched. With the narrow-focus approach, you start from one feature in the system which you already know in great detail and then proceed step-by-step from that starting-point to investigate and build up the whole system. With the big-picture approach of top-down omics, you start with the whole system and examine it directly as a whole. Importantly, as a result of these different approaches, the big-picture and narrow-focus approaches generate different kinds of data sets. The big-picture (top-down) data set is large, potentially unwieldy, possibly vague in characterization, and it may be incomplete. The narrow-focus data set of the bottom-up approach is small, granular, well-characterized, and very likely incomplete. The big-picture strategy is what most people think of in connection with omics research as such and it is sometimes called "shot-gun" omics, as in shot-gun proteomics.

Metalloproteomics operates in the big-picture mode. It employs an open-ended experimental design and aims to be as broadly inclusive as possible. However, it is not guided by a specifically-formulated hypothesis. Many scientists[44] so adamantly believe that investigating a hypothesis is essential to what constitutes real science that they reject any research project lacking a hypothesis (O'Malley et al., 2009, 611). They call it *merely descriptive*. On this definition that 'real' science must be governed by a hypothesis[45], metalloproteomics—and indeed proteomics in general and most 'omics' methods—may be dismissed as 'not quite science-as-such'. Typically such projects are disparaged as "fishing expeditions".[46] In competitions for research funding, such

[44] I believe that it is fair to say that some philosophers of science would be inclined to share this opinion, mainly those inclined to Logical Positivist orthodoxy; however, they might not use the researchers' jargon of calling non-hypothesis-driven research merely descriptive.

[45] Another way of postulating that "scientific" research as such requires a hypothesis is to say that it necessarily involves having a control group, or in effect, examining a null hypothesis. This is not to say that the null hypothesis is the hypothesis of record. The null hypothesis comes into existence as a foil to the hypothesis articulated for the experiment. This point arose from discussion with Dr. Michael Torbenson, Department of Pathology, Mayo Clinic, 31 May 2014.

[46] The 'fishing-expedition' terminology is a commonplace, but the reference here is to a comment by scientist Janet Rowley talking about "observationally driven research" as a fishing expedition, the kiss of death for funding today, quoted in (Keating and Cambrosio, 2012, 37). This is also my personal experience in Canada with the Canadian Institutes of Health Research circa 2000-2.

projects typically are "triaged"[47]: they are fast-tracked to the wastebasket and get no detailed consideration or review (O'Malley et al., 2009, 612).

One possible—though I would call it *extremely* desperate—rebuttal to hypothesis-traditionalists' rejection of metalloproteomics as real science (on the basis that it is not hypothesis-driven) is to assert that metalloproteomics actually does test a hypothesis. Specifically, it is contended that metalloproteomics works off some unstated hypothesis.[48] Obviously this is an important issue because, if true, it would effectively eliminate the need for my project to examine philosophical issues surrounding non-hypothesis-driven research. The key suggestion is that the belief that the system under consideration is worthy of investigation serves the same function as a hypothesis. One sympathetic formulation of this objection is to say that top-down metalloproteomics may be a "fishing expedition", *but you know that you are fishing in the right pond*. I find this line of argument totally unconvincing. Such a belief in the essential value of the research project undertaken is a feature of all serious research, whether it features a hypothesis or not.

Moreover, if metalloproteomics has a hypothesis like "this is a productive way to examine an important biological system", equivalent to confidence about fishing in the right pond, then such a hypothesis is inert. It does not go anywhere—experimentally or conceptually. Self-evidently, an inert hypothesis cannot drive a research project. It could be objected that such a hypothesis is about method. I would contend that hypotheses of method, if they exist at all, are problematic: their proof involves strikingly circular

[47] Terms like 'descriptive' and 'triage' are part of the jargon of grant application assessment. It is pejorative to call research "descriptive" because it implies that no interventional experiments are proposed. Strictly speaking, triaging means to sort and allocate according to likelihood of a good outcome (survival of battle or disaster victims, or likelihood of success in some respect). See http://www.merriam-webster.com/dictionary/triage (accessed 8 July 2014) (Anon.: Merriam-Webster). In ordinary CIHR jargon, "triaging" means assigning a fatally low grade so that the application requires no further consideration.

[48] This is an interesting comment because it is recurrent in my experience. I have heard it from more than one person, although I have not seen this line of argument published anywhere. I will examine this question in detail in Chapter 4 where I show that biological/biomedical researchers work within a nested hierarchy of hypotheses. Therefore this objection is not at all surprising. However, importantly, a hypothesis is not driving the experimental design of investigating a system in its entirety. As I will discuss, as a matter of successful grantsmanship, something passing for a hypothesis may be formulated so that the grant application actually gets considered. I have done that myself for a non-hypothesis-driven research project. Obviously it involves epistemic compromise.

reasoning. If the experiment provides interesting data, then this "hypothesis" is vindicated. Need the data be interesting? If the experiment produces any data at all, then it is productive. Mere productivity seems inadequate as a criterion of success; however, it is difficult to determine the degree of interest exhibited by the data without introducing bias. Furthermore, the statement "this is a productive way to examine an important biological system" barely qualifies as a hypothesis. It is more nearly an *evaluation* of method, and this evaluation is *a priori*, antecedent to actually using the method. Evaluating is different from testing. This kind of hypothesis is perilously close to the situation where the stated hypothesis functions as the mere semblance of a hypothesis. Such a hypothesis (if it is one) is very weak. It is formulated to meet the expectation that the research is hypothesis-driven. These hypotheses, simply by being formulated, may seem to elevate a research project to the realm of declared, recognizable 'science', but they hold little epistemic interest.

The notion that metalloproteomics (or proteomics, as an index example of research governed by the OES) actually has a hypothesis is not merely the defensive reaction of my biomedical research colleagues. Equally it is a contention of some people writing in the field of philosophy of science. In his 2012 paper "Convenience experimentation", Ulrich Krohs specifies that top-down systems biology has what he calls a general explanatory hypothesis of delocalization (Krohs, 2012, 54), which is generally not subjected to testing and appears not to be directly testable. He articulates this hypothesis as follows: "Metabolic capacities of the cell are brought about by metabolic networks. Single reactions or enzymes do not have identifiable, stable regulatory roles. Regulatory functions are delocalized dispositions of the network" (Krohs, 2012, 54). According to Krohs, top-down systems biology also has a methodological hypothesis: "Any case of metabolic regulation either matches the known motifs of delocalized regulation or can eventually be explained by the dynamics of yet unknown motifs of a discrete network" (Krohs, 2012, 55), also not open to direct testing. This assessment poses several problems. First of all, I find it highly debatable to call either assertion, whether in relation to explanation or method, a hypothesis. Failing

to be testable, a feature of these 'hypotheses' specified by Krohs, is a fatal flaw disqualifying them as hypotheses, at least on my definition. Thus I would argue that neither the 'delocalization hypothesis' nor the 'methodological hypothesis' is the kind of hypothesis that drives an experiment. Each is an assumption about the system being examined. However, if either were accepted as a hypothesis, it would fall into the 'inert hypothesis' category. The 'explanatory' hypothesis does not motivate a specific research project any more than "right pond" belief does. It may locate the research but it does not take the experimental project into new territory. It could be argued that while not motivating any specific research project, such hypotheses rationalize whatever research is done: they provide the rationale for the experimental project undertaken. A weakness of this argument is that it opens itself to a self-victimizing equivocation, namely, that rationalizing a research project may serve to provide an excuse for its short-comings. The 'methodological' hypothesis is only descriptive and establishes at most something about the *modus operandi* of the experimental strategy. I will return to some of these issues when I consider the nature of the hypothesis in biological/biomedical research. What I find really noteworthy, however, about Krohs's analysis of this issue is that someone outside the biological research community might feel compelled to construe systems biology as having a hypothesis. To reiterate my assessment, very different from Kroh's: although metalloproteomics, which is necessarily big-picture (top-down), does not test a hypothesis, it interrogates a system. The OES may possess assumptions of interest, but these are not operational hypotheses.

There is a further problem with Kroh's assessment of top-down systems biology: namely, the ambiguity of what is meant by systems biology[49] obscures differences between different types of omics disciplines. He clearly prefers a narrow-focus approach where metabolic pathways within the cell can be isolated into neat, self-contained entities, fully characterized in terms of energetics and regulation, and then reassembled

[49] Krohs himself speaks of "systems biology" but it is fairly obvious that he is actually talking about analyzing a biological entity as a system. Thus his discussion is entirely relevant to what I am discussing here. This assessment is corroborated by his focus on experiments which involve "chips" (that is, various sorts of gene arrays).

into the cell. Without any doubt, this strategy has been highly productive in cell biology and biochemistry. It assumes a certain linearity and modularity to cellular processes. What if this assumption is not entirely sound? Perhaps for some or all metabolic pathways the entire milieu actually matters, metabolism is not stepwise, regulation occurs across subcellular domains. If so, the delocalization assumption (a.k.a. "hypothesis) might be more fruitful. Some omics disciplines rely more fully on motif-matching than others. As stated, Krohs's methodological hypothesis seems to relate more to genomics and transcriptomics than to proteomics. Generalizing about omics disciplines requires much caution.[50] The point is that not all top-down systems biology is limited by the restrictions associated with preformed gene arrays. Whether all these disciplines only produce a static analysis of a cell (taken as a stand-in for 'system') (Krohs, 2012, 54) and whether this is a defect is open to debate. Such static snap-shots could be assembled into a picture resembling time-lapse photography. Static snap-shots have their useful role in research strategies where changes over time can be mapped; indeed, each data point on a Michaelis-Menten curve analyzing an enzymatic reaction is a kind of snap-shot. Omics may not be the best experimental strategy for enzyme kinetic data. The strategy has to be suited to the problem to be solved. However, omics may reveal features of the biological system whose kinetics are studied which would not otherwise be taken into account by traditional metabolic pathway analysis.

3.2 The hypothesis-driven/non-hypothesis-driven divide is a mirage

Another contention which could vitiate this philosophical project (and thus demands immediate consideration) is that the dichotomy[51] between hypothesis-driven

[50] An important feature of Krohs's attitude toward these issues is his distrust and deprecation of high-throughput methods which he collectively characterizes as *convenience experimentation* (Krohs, 2012, 53). He calls attention to an important problem, but as we will see in Chapter 6, there is little of convenience in performing proteomics. Protein chip analyses may fit his criticism, but most of proteomics does not. Elsewhere he decries modeling based on high-throughput data sets because the data are easy to obtain (that is, convenient) (Krohs, 2010, 160), and he rightly points out the importance of taking epistemic criteria as the basis for designing omics experiments. This latter approach is clearly how we proceeded with metalloproteomics, as was described in detail in Chapter 2 (Section 2.5.2).

[51] It is worth noting that I am employing a less stringent notion of 'dichotomy' than some might conceive. I use it in the sense of strict disjunction: of two possibilities, take one or the other. The difference between

research and non-hypothesis-driven research is highly artificial and thus, in effect, simply does not exist. If that were the case, my philosophical project to examine the epistemology of non-hypothesis-driven research would seem rather superfluous. Unreality of the hypothesis-driven/non-hypothesis-driven divide is the implication of work currently being done by Maureen O'Malley, who is focusing on the role of integration generally in systems biology (O'Malley and Soyer, 2012, 58). Briefly, she describes the research under discussion here as 'data-driven': "research that is characterized by the generation, collection and potential interpretation of large bodies of biological data, from which subsequent analyses can detect new relationships, processes and phenomena" (O'Malley and Soyer, 2012, 58) and allows that in general it is at least assumed that no particular hypothesis is being tested. What she describes here under her terminology of 'data-driven research' is equivalent to my "system-driven research"[52]. Her contention is that the opposition of data-driven and hypothesis-driven research is misleading and non-productive. She believes that the greatest emphasis should be on integration as "a multi-faceted dynamic, in which methods, bodies of data and explanations are synthesized in order to understand and intervene more effectively in biological systems" (O'Malley and Soyer, 2012, 59). This can include the integration of both data-driven and hypothesis-driven approaches.

In a general way, I agree with her. Similarly, Peter Medawar emphasizes integration. The entities he proposes integrating appear different, but I find strong parallels between non-hypothesis-driven and hypothesis-driven research and the entities in his discussion. He points out that scientific reasoning is a dialogue between

---

the two is clear and real, x or not-x. I do not intend the more demanding concept of dichotomy where the two elements are not only clearly and actually different but also are always separate. With this more stringent kind of dichotomy, each element only occurs separately, never with the other. As we will see, this tight notion of dichotomy does not hold for hypothesis-driven versus non-hypothesis-driven research for various reasons, and it would distort the analysis to think it might apply.

[52] It is important to be very clear that O'Malley's term 'data-driven' and my term 'data-driven' are *not at all* equivalent. In my view, data-driven research is organized around a large existing body of data which seems likely to contain some interesting new information. It has a retrospective tilt. In data-driven research, collected data are the primary starting-point. By contrast, system-driven research is essentially prospective in design, and the system is the starting-point.

what he calls two episodes of thought, the imaginative and the critical[53]: he wants the emphasis to be on the unity of the scientific reasoning amalgamated from these two components (Medawar, 1982, 101). However, a problem with this emphasis on this kind of integration is that it tends to divert attention from the issues of exactly what system-driven (to revert to my term) research does and how it operates to produce new knowledge. Since hypothesis-driven research has the benefit of already being well-analyzed and, moreover, accepted as the standard of research methodology, it may eclipse system-driven research when the focus is on integration. By focusing on the integration stage which O'Malley has identified, we may never really examine that system-driven research is all about. I am not suggesting that integration of hypothesis-driven with system-driven research is unimportant (I will discuss it in Section 9.3), only that some epistemological homework about system-driven research remains to be done before we can get around to talking about such integration.

It may be that hypothesis-driven and system-driven research are operationally closely allied, perhaps even in some sort of dynamic equilibrium, but this relationship needs to be demonstrated. I can easily imagine a dynamic interplay between these two research strategies. From the point of view of practicing scientists, integration of data is a key intrinsic feature of systems biology (Lin and Qian, 2007, 107-8). For example, Alan Aderem pointed out in 2005: "To practice systems biology, one must capture and integrate global sets of biological data from as many hierarchical levels of information as possible" (Aderem, 2005, 511). A problem is that this comment moots on the exact role integration is playing. If looking at a biological problem from different angles and then combining those data sets to get a highly textured picture is central to how one does systems biology, it may or may not invoke hypothesis-driven methodology. Sorting out

[53] For 'imaginative' we may read "non-propositional" or "visual"; for 'critical' we may read "propositional" or "logical". This comment by Medawar in the early 1980s seems to me to reflect a lingering commitment to the distinction between context of discovery and context of justification. The imaginative episode seems to be another way of talking about context of discovery, and the critical episode strongly resembles context of justification. I read Medawar's emphasis on the dialogue between these two as signifying some discomfort on his part at making this division. Consideration of context of discovery and context of justification in relation to omics will constitute the next section of this chapter (Section 3.3).

the role of integration both within system-driven research and with respect to hypothesis-driven research is complicated. The analysis requires at least bracketing the two different methods for initial steps of the analysis. One of the objectives of this philosophical project is to accomplish that task for system-driven research.

Finally, there is a gross pragmatic consideration. Even if the divide between hypothesis-driven and system-driven research is artificial, the supposed opposition of these two ways of doing scientific research matches how a lot of people view it. Specifically, those who take having a hypothesis as the *sine qua non* of scientific research are forced into this dichotomy to account for 'wannabe' scientific research which does not display a hypothesis. The opposition of hypothesis-driven versus non-hypothesis-driven research exerts far-reaching, powerful influence. Therefore it demands attention, even if it is only a "perceived" dichotomy.

3.3 Is OES merely residing in the 'context of discovery', an appetizer to the hypothesis-driven scientific banquet?

Philosophers of science may react to epistemological issues posed by high-throughput research methods and omics in particular (thus necessarily, the OES) by pointing out that the OES belongs in the context of discovery, as theorized by the Logical Positivists. Scientists may cope with these same issues by taking the same tactic, except that they do not attach this technical term. This is an important consideration and requires immediate attention.

3.3.1 The context of discovery/context of justification problem

In the 1930s Hans Reichenbach invented the distinction between 'context of discovery' and 'context of justification'. This distinction coordinated well with concepts of how science works, as promulgated by the Logical Positivists.[54] Reichenbach wrote: "The *objective* relation from the given entities to the solution, and the *subjective* way of

[54] I might suggest that we can view this as one of the most destructive residua of the Logical Positivists because it limits the scope of scientific epistemology. This opinion is also developed by Marcel Weber (Weber, 2005, 51-4, 149-50). What impresses me about this distinction of context of justification from context of discovery is its staying-power. Despite various arguments against it, it still manages to persist and demand our attention.

finding it, are clearly separated for problems of a deductive character; we must learn to make the same distinction for the problem of the inductive relation from facts to theories" (Reichenbach, 1938, 36-7, original emphasis). He laid out the two distinct domains: "The *context of discovery* is to be separated from the *context of justification*; the former belongs to the *psychology of scientific discovery*, the latter alone is to be the object of the *logic of science*" (Reichenbach, 1938, 36, original emphasis). Not surprisingly, this distinction ended up figuring prominently in the analyses of Logical Positivism/Logical Empiricism bearing on science. It was still current and warmly endorsed in discussions occurring around the early 1980s (Siegel, 1980, 297-8).

Latterly these terms have undergone some conceptual modification. In contemporary usage, the context of discovery refers to how a hypothesis first arises (many different ways, including serendipity). Some may view the context of discovery as having the connotation of being messy, unsystematic, and disorderly. This is uncharitable: governance of the context of discovery is not left merely to caprice, though not all of its processes fit established forms of inference, as in the context of justification. Whereas the context of discovery serves as the domain for formulating hypotheses, by contrast, the context of justification is where the proper investigation of a hypothesis takes place. In other words, this is where objective "real" science is done! Most importantly, the context of justification is where the classic hypothetico-deductive model operates. The messy product of the context of discovery gets laundered by the application of clear and dependable logic and by objective experimental technique. Worthless explanations (that is, disproven or falsified hypotheses) never survive to the next round of scientific investigation.

In current common parlance among biological scientists, the context of discovery is regarded as the domain where hypotheses are generated. It has become socially acceptable to characterize research governed by the OES as "hypothesis-generating". This endows such research with a certain intellectual respectability: it participates in 'real science' as the preliminary phase. Systems biology is one example of "data-intensive scientific discovery" (Bell, 2009, xv), term which may be an attempt to

legitimize that context (the context of discovery), or else an inadvertent overlap with old philosophical terminology. Most likely, it is enthusiastic adoption of currently fashionable language featuring 'discovery' like that which currently designates research institutes as discovery districts. More pertinently, "data-intensive scientific discovery" echoes my assertion that systems biology interrogates a system[55] and thus constitutes legitimate scientific activity. Given the connotative residua of the Logical Positivists, though, in relation to omics, it risks placing omics outside of real science, in the discovery phase, preparative to real science phase.

My contention is that system-driven research is not merely a mechanism for generating hypotheses. In the first place, if it were only a hypothesis-generator, that activity would not distinguish system-driven research from hypothesis-driven research, which certainly generates a whole string of hypotheses, in fact, string after string of hypotheses. Indeed it might be argued that hypothesis-generation is a defining characteristic of scientific activity. On such a concept, then, system-driven research—including any omics—is clearly legitimate science. However, this is an extremely weak argument. It brings hypothesis-driven and non-hypothesis-driven research under the same conceptual framework without paying attention to the interesting differences between these two classes of research design. In my opinion, such an argument does not attenuate the mandate for my philosophical project because it is too superficial an argument to carry the day.

Moreover, our own definition of metalloproteomics provides empirical evidence of circumventing this problem. It precludes the divide between discovery and justification by stipulating that the scientific activity described by 'metalloproteomics' has two complementary aspects. It systematically identifies large sets of proteins

[55] What Gordon Bell actually describes "data-intensive scientific discovery" as involving is "three basic activities: capture, curation, and analysis" (Bell, 2009, xv). Curation in this context means a well-curated, or well-tended, database. This terminology turns up early in proteomics, at least circa 2001, and I think the allusion (to the parish curate who looks after the congregation, or to a conservator who looks after a museum) is not entirely accidental. The multi-authored book which this foreword introduces presents a comprehensive vision of data-driven research embracing physical and biological research and actuating the ideas (which he called 'eScience') of the eminent computer scientist James (Jim) Gray.

associated with metals and then puts those results into the physiological context, namely, the system under investigation. At least in terms of metalloproteomics, when we contextualize results, we effectively erase the distinction between context of discovery and context of justification. I believe that this approach can be extended to proteomics in general, and in the past 5-7 years, most proteomics scientists have come around to the same point of view.[56]

3.3.2 Other challengers to this dichotomy—previous to metalloproteomics!

Our operational definition of metalloproteomics is by no means the first such challenge to this distinction. It is quite clear that Paul Feyerabend (Feyerabend, 2010, 150-1) and N. R. Hanson (Hanson, 1961, 70-2), among others[57], rejected the distinction between context of discovery and context of justification. Feyerabend argues forcefully on empirical grounds: actual scientific practice is "a complicated mixture of procedures" (Feyerabend, 2010, 150-1). Distinctions between observation and theory, as he also calls them, can be made but they serve no useful purpose in real-life scientific research (Feyerabend, 2010, 152). Hanson in particular presents a spirited rejection of the context of discovery versus context of justification dichotomy on the basis that consigning how hypotheses are generated to irrelevancy misses the experiential fact of the matter: it often takes extensive work (including both data and genius) to formulate a sound hypothesis (Hanson, 1961, 71-2). Furthermore, anticipating other more recent objections, he writes: "If establishing an hypothesis through its predictions has a logic, so has the conceiving of an hypothesis" (Hanson, 1961, 71). Yet it is possible that with most scientific investigation being dominated by hypothesis-driven strategies in the mid-20th century, these objections did not get off the ground.

At the current time, however, attitudes toward this distinction are somewhat more flexible than in the 1960s. At the theoretical, or analytical, level it is possible to

[56] It may be open to some dispute whether this point can be extended to various other 'omics' disciplines because some scientists hold that all an omics discipline has to do is generate a list. Personally, I find this definition of the mission of proteomics and other omics research to be too narrow. My interpretation may reflect my prejudice that we use omics to investigate cellular or some kind of physiology.

[57] "Others" are enumerated by Feigl as quoted by Feyerabend (Feyerabend, 2010, 150); the list is further updated by Bechtel and Richardson (Bechtel and Richardson, 2010, 4), who nevertheless fail to include the feminist epistemologists I have cited.

argue that methodological features which are critical to the context of justification inform how the context of discovery works, and likewise certain key characteristics of the context of discovery affect (and, in a sense, contaminate) the logical function of the context of justification. This leads to some blurring of divide between the two domains. A narrow practical version of this argument is that omics may operate in context of discovery but it employs standards of data production suitable to the context of justification: indeed those standards serve as a kind of justification. The boundary ends up blurry. Another approach is more pragmatic: it contrasts simple ontology with real-life functioning. It concedes that we can talk about the existence of context of discovery and the context of justification as being distinct entities, but operationally we move back and forth between these contexts such that the boundary is indistinct. What gets lost in this contemporary revision of the relationship between context of discovery and context of justification is the dismissive attitude of the Logical Positivists toward the context of discovery.

In terms of a formulating a criticism of context of discovery versus context of justification based on epistemological theory, feminist epistemology poses a direct challenge. Whether evidence confirms or disconfirms a hypothesis is not independent of the context of discovery. Richmond Campbell points out that the context of justification conceived as an isolated intellectual space impervious to outside influences makes internal feminist epistemology a conceptual impossibility (Campbell, 1998, 24). He argues cogently that the context of discovery cannot be differentiated from the context of justification: that "the very logic of confirmation, according to the familiar empiricist norms cited earlier, depends on the context of discovery" (Campbell, 1998, 25). The empiricist norms to which Campbell alludes are predictive success, observer independence (more specifically conceived as lack of circularity in the formulation of the problem under investigation), and explanatory power (Campbell, 1998, 21). Campbell argues that these empiricist norms constitute roughly a method for testing hypotheses and embrace an affirmation that available hypotheses present a "relatively accurate representation of what the world is like" (Campbell, 1998, 32). A certain degree of

evaluation and editing take place before that product gets delivered to the context of justification. Specifically, these empiricist norms themselves cannot be applied without taking into account the resources and diversity inherent in the context of discovery. There are various ways to appraise 'success'; the problem under investigation might be formulated in several different ways to promote elucidating it; explanatory power requires having several different solutions at hand among which to compare and choose. Likewise, the so-called objective evaluative processes of the context of justification operate with norms and assumptions which have also informed the context of discovery. For example, circular reasoning finds no quarter in the context of justification, just as it does not in the context of discovery. Additionally, in operating upon the 'output' of the context of discovery, if we wish to style it that way, the context of justification embraces whatever norms were operative. Recognizing the interdependence of context of discovery and context of justification provides the possibility of seeing that the contexts of discovery and justification for omics, as system-driven research, are overlapping, or at least possessed of an exceedingly smudged boundary. If the division is not strict, then it is difficult to consign some scientific activity to one and some to the other. Then it becomes very difficult to argue that omics is just a "context of discovery activity", that is, as (merely) hypothesis-generating.

There are two conceptual approaches for exploring the possibility of such interdependence between the two contexts in terms of omics research. The first is based on practical considerations to this criticism, which is after all based on theory. The strict division between the contexts of discovery and justification implies that all credible assessment of the product of the context of discovery can take place only in the context of justification. With omics methodologies, we might reasonably impose some standards of technical competence upon the experimental practice before we accept that data set as worthy of further consideration. For proteomics, as we shall see, such standards have been imposed: technical benchmarks which must be met before the data are suitable for consideration for publication in a peer-reviewed scientific journal. Effectively such formal evaluative criteria push some of the objective, 'scientific' action

of the context of justification into the context of discovery, if you believe that proteomics resides firmly and exclusively in the context of discovery. Consequently the strict division between the contexts of discovery and justification breaks down dramatically. It is not merely that the context of justification shares assumptions with the context of discovery: objective evaluative practices take place in the context of discovery. If the demarcation breaks down, then it becomes awkward to consign proteomics to the context of discovery with the implication that proteomics is operating at a level outside of, and preliminary to, real science.

3.3.3 Looking at the context of discovery/context of justification dichotomy from the vantage-point of the laboratory

The second approach to examining interdependence between the two contexts moves into the research laboratory and uses some real-life situations relating to biological research to look at how this distinction plays out functionally in real life. In my opinion, one of the limitations for the Logical Positivists was that they theorized aggressively about how science is done without actually doing much practical science themselves. What contemporary science there was available to them to motivate their theorizing was in the domain of the physical sciences. Even practical experience in chemistry or physics might not be informative about biology! Thus we can turn to a real-life example, though admittedly a rather minor example, from the lab. We wanted to find out whether the Ah receptor, a cellular protein which regulates a certain genes so that the phenomenon of induction of certain enzymes occurs, was present in human liver tissue—a scientifically important question. One experimental model for examining this question was to determine whether it was detectable in a continuous human liver cell line, HepG2. This seemed like a good and practicable experiment because the protein had been detected in a similar continuous mouse liver cell line. This experience not only served as background data for the experiment but also provided technical cues for the experiment with a human liver cell line. Performing the experiment involved making cytosol, the watery stuff in cells, where this protein may be found, from HepG2 cells. That involved growing the cells in 40 to 50 “P-100” Petri dishes (classic 4-inch diameter plastic dishes) in order to get enough cells to make enough cytosol. The

accepted routine was to scrape the cells off the dish and pool the scrapings prior to the next preparatory step. What was obvious was that up to half the cells were left behind on the plate, no matter how skillfully you scraped the plate. I decided to use an enzymatic lifting of the cells instead, just as was done for propagating the cell line by subculture. In effect I formulated the hypothesis that harvesting the cells enzymatically would be as effective as scraping. In fact my revised method gave a 100% harvesting of the cells from the dish and provided the maximum cell pellet for preparing the cytosol. This modification of technique was based on a hunch, supported by relevant prior experience—pretty typical of how scientific investigation works. (It may be, of course, that I simply hit upon what other labs were doing routinely anyway, but it was an innovation in our technique.) The results of my experiment for detecting the gene-regulating protein were positive and congruent with results in other systems. Subsequently, we went back and did the "context of justification" experiment and showed that the results were equivalent whether the cells were harvested by scraping or enzymatic lifting. After all, it could have been that the enzymatic method lifted cells but also modified the protein of interest. The point of this story is that the action between 'context of discovery' and 'context of justification' was not linear (first discovery and then justification) but interdigitated. In actual scientific practice—at the real-life functional level—the division is not strict. You can label 'context of discovery' and 'context of justification' as such, but functionally they are not distinct.

3.3.4 Philosophers of biology weigh in

When we examine the commentaries of other philosophers of biology on this problem, we find similar lines of argument. In my opinion, these real-life considerations inform William Bechtel's and Robert Richardson's argument against the divide between context of discovery and context of justification. They describe their argument as psychologistic. It is broadly similar to Campbell's argument based on theory and my argument above based on diverse practical experience. Their starting-point, sharply divergent from the outlook of the Logical Positivists, is to be interested in the process of

scientific discovery.[58] Bechtel and Richardson describe two major patterns in the arguments against the dichotomy of contexts of discovery and justification. One pattern upon which they insist demonstrates that attributing exclusive clarity to the processes operating in the context of justification is erroneous. This is similar to Campbell's arguments above. Their other pattern of counter-arguments is that the context of discovery is not devoid of orderly procedures: "discovery too is animated by scientific practice" (Bechtel and Richardson, 2010, 4). They argue that what goes on in scientific discovery is a special case of human problem-solving and needs to be analyzed in terms of the same sort of psychological capacities and mechanisms which generally operate when humans try to solve problems. As opposed to being analytical and deductive—that is, strictly logical—they suggest that the process involves adjustment and pattern-matching (Bechtel and Richardson, 2010, 7). For them, the distinction between context of discovery and context of justification breaks down because similar problem-solving operations are going on in both domains. This argument is not unlike what I have illustrated from real-life examples. It risks some overall weakness because Bechtel and Richardson modify the scope of 'discovery' to "scientific discovery".[59] This conceptual modification may depart too far from the original intent of the Logical Positivists. Their argument becomes somewhat less potent against the context of discovery/context of justification distinction because by specifying discovery as scientific they endow it with characteristics peculiar to the context of justification.

William Wimsatt has also commented on the problem of the context of discovery versus context of justification: he holds that we need to attach more

[58] They point out that "Positivistic accounts not only ignore discovery, but also see justification exclusively in empirical terms" (Bechtel and Richardson, 2010, 4). Theory endows empirical data with relevance. Absent a relevant theory, judging empirical data is difficult, if not impossible. It could be objected to Bechtel and Richardson that the Logical Positivists never entirely abandon the emphasis on logical inference for managing empirical data.

[59] Specifically they say: "We seek a realistic dynamic model of scientific discovery. We seek to understand the cognitive strategies, the procedures, constitutive of scientific rationality. These strategies are, from one perspective, the procedures that define *how* humans approach the problem of understanding the world. From another perspective, the procedure human embody constitute assumptions about the structure of the world, or of the part of it to be explained. They define *what* we think about the world" (Bechtel and Richardson, 2010, 11, original emphasis).

importance to discovery but not forget about the important role of justification. Specifically he states that "we should embed our more traditional concerns of justification and discovery in a theory of practice" (Wimsatt, 2007, 340). His views come as a kind of antidote to the potential problem of an overly enthusiastic affirmation of context of discovery plus proportionate (but for Wimsatt, potentially lamentable) de-emphasis of context of justification.

Wimsatt sounds like the conservative voice here, expressing support for context of justification and indirectly for the dichotomy itself. Such an interpretation of his position seems wrong to me. I believe he is merely pointing out that the importance of validation should not underestimated, even while we get re-focused on discovery. He does not act as apologist for the context of discovery versus context of justification dichotomy. Indeed, he characterizes insistence on the division as a pernicious bias (Wimsatt, 2007, 250). The point of course is that Wimsatt demonstrates ambivalence about this issue: the inflexible insistence on a distinction between 'context of discovery' versus 'context of justification' is non-productive but both discovery and justification have their important roles in the scientific experimentation.

Thus the consensus view among these philosophers of biology is broadly similar to mine.[60] I agree with Feyerabend and Hanson that discovery is reasoned and with feminist epistemologists who point out that the two domains are intermingled. However, the important point for my argument is that we need to look to actual practice for sorting out this issue. As I have shown, if we take the actual practice of experiment and investigation in biology as the focus for sorting out the relative roles of

[60] The one puzzling voice in this discussion is that of Kathleen Ohkrulik who comes ever so close to rejecting the dichotomy between context of discovery and context of justification and then unexpectedly backs off at the last minute. What she says is: "The argument here is *not* that we should abolish the distinction between contexts of discovery and justification, but that we must recognize that on a comparative model factors that influence theory development and theory generation must necessarily influence our confirmation practices and hence the very content of science" (Okhrulik, 1994, 37, original emphasis). Given her cogent argument against the dichotomy, that the context of justification does not in fact succeed at eliminating subjective factors found in the context of discovery, I find it somewhat surprising that she does not press the argument to its radical conclusion. In fact, she does not have to deal the actual body blow to the dichotomy: it is fully implied. Moreover, she may share some of the Wimsattian ambivalence that individually discovery and justification are important concepts in the analysis of scientific experimentation, only not quite as simplistic as Logical Positivists argued.

discovery and justification, we find that both are important and both are overlapping. The strict division of discovery versus justification proves to be artificial (and thus irrelevant).

3.3.5 Context of discovery/context of justification dichotomy does not really exist (at least not for contemporary biology)

Thus I have argued on the basis of theoretical considerations and real-life practice in biological research and clinical medicine that the demarcation between context of discovery and context of justification does not hold up. For omics, these domains functionally overlap. Recognizing the legitimacy of the context of discovery, as argued by Campbell as well as by Bechtel and Richardson, at the very least eliminates the dismissive tone of Logical Positivists' verdict: that omics is merely in the context of discovery, a preparative stage to doing "real science". However, on more substantive criteria, it is not viable to relegate OES to the mere role of hypothesis-generating from the context of discovery. Actual contemporary research practice is much more complicated. Relegating omics (or system-driven research) to merely the preparative phase also fails to recognize that there are some problems which are more likely to be solved by not trying to formulate a hypothesis. An obvious example is our problem of trying to work out the network of proteins which handle copper inside liver cells (hepatocytes). This problem is not going to be addressed efficiently by hypothesizing on a protein-by-protein basis. It requires a global, highly inclusive approach. In this situation system-driven research outperforms hypothesis-driven research.

The context of discovery/context of justification dichotomy holds up effectively only in a universe of discourse ruled strictly by the HDM. (Famously, that universe of discourse disallows certain highly valued philosophical projects, like Aesthetics.) Thus, if we recognize alternative ways, apart from the HDM, to organize biological/biomedical research, then we may expect to be forced to reject the context of discovery/context of justification dichotomy or at the very least to reconstruct it, as Campbell and others have. A reasonable next step in exploring this problem is to take a closer look at what a hypothesis is and how the HDM operates so that we have a firm basis for proceeding

with an analysis of system-driven research designs and the OES. That task occupies the next chapter.

CHAPTER 4 HYPOTHESIS

Having defended this philosophical project from the sorts of objections which could render it a non-starter, I now want to consider two key players in this analysis: namely, *hypothesis* (in this chapter) and *experiment* (in the next chapter). The general view of hypotheses and experiments in biology is somewhat different from that of the physical sciences. Notions based on the character of hypotheses in physics or chemistry cannot be imported directly into the biological sciences because the biological sciences deal with complex but highly particular entities. Marcel Weber explains this difficulty in terms of the multi-faceted mechanistic explanations sought: "biological hypotheses are often proposed mechanistic explanations that contain more than just a set of variables and relations of causal dependence between them" (Weber, 2012, Section 3.2). In terms of biological/biomedical experiments dealing with biological organisms, the experiment is part of an effort to attain a mechanistic understanding of the biological problem under consideration—not just derive a summary (typically mathematical or formulaic) statement about it. Critically examining these two concepts is crucial to the examination of system-driven research because the outcome of that effort may well be revisionary: we expect to abandon hypothesis, and experiment (including how experimental results are justified) may end up radically redesigned. With respect to hypothesis, the subject of this chapter, it is important to know what we are giving up when we switch from hypothesis-driven to system-driven research.

We know in a general fashion what hypotheses and experiments are supposed to *do*. This general knowledge constitutes current orthodoxy. Hypotheses convey testable assertions; experiments are designed to test hypotheses. In general, the testable assertion relates to explanatory information about the thing being studied—what it is or what it does. Thus the assertion is part of an explanation-prediction duo—some would say an explanation-prediction identity. For any hypothesis, an experiment is or can be designed to generate data to investigate its predictions. If the data are inconsistent with what is predicted by the hypothesis, the hypothesis is abandoned,

though it is usually replaced with another. If the data confirm the hypothesis, further hypotheses (and experiments) are designed to explore its scope further and eliminate alternative explanations of the data. This general methodological approach generates a chain of hypotheses, and new knowledge may be produced through the progress from one hypothesis to the next.

It is worth recognizing that positive directionality is often appreciated only after the fact. In real time a hypothesis may be shown to be erroneous and thus lead nowhere. Exploring its modifications constitutes a kind of sideways movement.[61] There may be a bunch of competing hypotheses which require assessment in turn, a process which leaves a pile of discarded disproven hypotheses behind it. This set of articulated hypotheses may not be complete. One of the points Kathleen Ohkrulik makes in her trenchant critique of a comparative method of hypothesis evaluation is that hypotheses which never occur to the researcher (for one reason or another) never get investigated (Okhrulik, 1994, 34-5).

Importantly, this general notion of how hypotheses and experiments operate is derived directly from the HDM as the quintessential version of how the Logical Positivists viewed scientific investigation. Scientific knowledge production emanates directly from a logical analysis. A hypothesis is formulated as part of this logical framework. Then it is operationalized into a 'test implication'. The test implication takes the physical experimental set-up into consideration, along with the various assumptions attendant to its actually working (ordinarily called 'auxiliary hypotheses'). The findings

[61] Lakatos is more flexible and realistic (and forgiving) about how investigating a scientific problem is often not linear. He points this out in relation to his rejection of Popperian falsificationism. He describes his own analysis of scientific research as "*more tolerant* in the sense that it allows a research programme to outgrow infantile diseases, such as inconsistent foundations and occasional *ad hoc* moves. Anomalies, inconsistencies, ad hoc stratagems, even alleged negative 'crucial' experiments, can be consistent with the overall progress of a research programme. The old rationalist dream of a mechanical, semi-mechanical or at least fast-acting method for showing up falsehood, unprovenness, meaningless rubbish or even irrational choice has to be given up" (Lakatos, 1974, 319, original emphasis). He portrays his approach as contrasting to Popper's extremely linear version of the scientific process, which in fact is more akin to the process of testing and discarding hypotheses which I have just described above. Lakatos's notion of research programmes is much grander and broader in scope than my vision of investigating a system, but his concepts resonate to some extent with what I will be describing in terms of system-driven research.

from the experiment are then scrutinized in the logical framework, as a process of justification. What is apparent is that the experiment is tightly bound to the logical framework. The hypothesis/test-implication duo spans between the logical framework itself and the practical world of the laboratory bench. This view of scientific investigation excludes right at the outset any alternative ways of carrying out scientific investigation. Any strategy not conforming to this general rubric (for example, system-driven research as I have defined it) is not part of the realm of scientific investigation, not merely because it goes against the key concepts held by Logical Positivists but because, in its failure to conform to the HDM, it lacks a supportive epistemological structure. The manifest implication is that supplying epistemological structure to a research strategy endows it with legitimacy as a specifically scientific endeavour.

Even though 'hypothesis' is in general parlance, it is not entirely straightforward to describe or define what a hypothesis *is*. Part of the problem is slippage in the concept across generations of philosophers and scientists. Both Plato and Aristotle comment on a concept which appears to correspond to our modern term 'hypothesis'. It is at least conceivable that their usage has stuck to the term as we use it nowadays. Importantly, most of the philosophical discussion of hypothesis has been located within consideration of the physical sciences. Another important aspect of this problem is a recent tendency, especially in the modern Philosophy of Science literature, to conflate the terms 'theory' and 'hypothesis' or at least use them interchangeably.

Among ordinary scientists, there is a different problem of ambiguity in the meaning of 'hypothesis': the hypothesis cited for a research project, especially in the domain of biological or biomedical science, is often a statement of the overarching 'big idea' to which the entire project is dedicated. One of my research mentors once described a hypothesis to me as being an idea for which you would be willing to mortgage your house! If we shift the metaphor to today's business-orientated, entrepreneurial world, we might say that a hypothesis is an idea around which you would be willing to launch a small business, complete with the financial and personal

investment that project necessitates.[62] The point is that such broad 'visionary' concepts have a much larger reach than an operational hypothesis directly tied to an experiment. This big "hypothesis" might more accurately be described as the rationale for a research project or as the important broad concept undergirding the enterprise. If it is a plausible and interesting rationale, it might reasonably be expected to generate actual hypotheses as such. It is easy to confound this rationale with the more focused hypotheses it generates. The habit among biological scientists of calling it some sort of grand hypothesis is thus not entirely surprising.

4.1 Diversity of hypotheses

For the purposes of my discussion, I am looking for a very circumspect concept of what a hypothesis is. The challenge is to work out its characteristics exactly and establish that this concept of hypothesis is suitable for designing biological research. One possible simple formulation is prevalent, perhaps grounded in a very simple formulation that reflects the attitude of Pierce and Hempel, structured on logic: namely, that a hypothesis has a logical form, namely as an "if-then" statement, a conditional. Limiting a hypothesis to being a conditional may prove problematic.[63] Nevertheless we can see that a hypothesis is expressed as a proposition (though not necessarily as a conditional). In terms of the content of such a proposition, a hypothesis exceeds the obvious. Specifically, it is not a tautology and it does not merely restate what is regarded as known. Its content is such as to be testable by actual experiment, for example, in the laboratory or in the field. It may suggest the kind of experiment required to test it. As I have already pointed out, hypotheses tend to come in chains: one hypothesis leads to another. On a day-to-day basis, in actual research 'at the bench' hypotheses may be

[62] This comparison is not entirely farfetched because in a sense a research proposal for a major operating grant is a pitch for a small business. The applicant has to prove his/her expertise, indicate that there is a niche/need for his/her research initiative, and provide a justified budget. Running a laboratory is similar to running a small business with the attendant problems of managing personnel and space and also of advertising productivity.

[63] I am grateful to Richmond Campbell for pointing out that a relevant consideration here is that some intend a hypothesis to say that "for all x, if F(x), then G(x)" whereas others have a narrower viewpoint which resolves itself as "there is some x such that if F(x), then G(x)". It is important to note that here the 'if … then' construction is in the present indicative tense: it is not equivalent to the horseshoe (⊃).

very narrow in scope. Such a hypothesis may be quantitatively specific. It may stipulate as the issue to be tested, for example, that in the presence of condition A, the affinity of protein X for copper is 100 times the affinity when condition A is absent. More often than not, however, the exact quantitative difference between intervention and control is not specified. The contrast sought between the intervention arm of the experiment (such as Condition A present) and the control arm is only envisioned as being statistically significantly different. Finally, because hypotheses do tend to have a genealogy, which is to say that they come in a series or chain of hypotheses, they are often limited by what is already known. In this sense, a hypothesis can itself be limiting or even biasing.

In classic philosophical literature certain variant types of hypothesis abound. In Philosophy of Science contexts such variants serve as technical jargon. I want to discuss these 'arrant' hypotheses briefly and get them off the table because otherwise they may clutter the discussion and render it unnecessarily confusing. A **bridging hypothesis**, or principle as Hempel sometimes calls it, is a concept which "connect[s] the 'unobservable' theoretical entities with the subject matter to be explained" (Hempel, 1966, 74). Hempel goes on to point out that without them a theory would have no explanatory power and also it would be incapable of being tested (Hempel, 1966, 74). An **auxiliary hypothesis** is an assumption about a technical or environmental state of affairs which has a direct bearing on the experiment (Hempel, 1966, 23). More specifically, auxiliary hypotheses are needed to deduce what is predicted in the experiment. For example, auxiliary hypotheses may relate to the collection of assumptions relating to how an analytical instrument works and that it is actually working as designed for the duration of the experiment. Another example is that constant control of temperature or atmospheric pressure actually takes place. Alternatively the correction of data to "standard temperature and pressure", a feature more characteristic of physico-chemical experimentation than biological/biomedical experimentation, is an example of a theoretical version of an auxiliary hypothesis: that such correction can be performed and that it is worthwhile. An example of a similar

phenomenon in biological research is that a continuous cell line is best studied when grown just to confluence, after the phase of rapid "logarithmic" growth is over because the metabolic functions have then stabilized. Although these terms feature the word 'hypothesis', they should not be confused with the operational hypothesis under consideration here. Indeed, I really prefer calling these "auxiliary assumptions" or supporting premises, since investigating them is sometimes the key to understanding the failure of experimental data to support the hypothesis under examination. I find the usage "auxiliary hypotheses" generally confusing. (However, by calling them 'hypotheses', Hempel maintains parallelism within his general logic-based approach to scientific reasoning; moreover, assumptions seem much closer to the actual labour of doing experimentation whereas 'hypothesis' maintains a certain distance from the laboratory bench.) The point is that being able to identify bridging or auxiliary hypotheses which are pertinent to experimental design does not endow the experiment with the characteristic of being governed by a hypothesis as such. These terms should not be regarded as directly relevant to the kind of hypothesis considered here, and they should not be permitted to muddy the analysis of non-hypothesis-driven research.

Likewise, the term 'hypothesis' turns up with various confusing usages in scientific circles. The 'big idea' hypothesis is often articulated in order to communicate the rationale for the research project but, equally, the significance of the experimental work envisioned. It then serves as the 'pitch' for the work proposed. A typical historical example is to assert that finding the gene which is abnormal in cystic fibrosis (CF) will lead to a cure for CF. Of course, such a grand hypothesis arches over a lot of experimental territory, and it sounds—and indeed it *is*—compelling. Everyone would agree that curing CF, the most common monogenic disease in North American whose usual outcome is relatively early death from chronic lung or liver disease, is a worthy endeavour. Some would object that, strictly speaking, it does not follow logically that identifying that gene will actually lead directly to curing CF, but it is plausible that without cloning that gene cure will prove more elusive. Moreover, if your goal were gene transfer therapy, then clearly having the gene in hand is the first step toward that

cure.[64] 'Hypothesis' also turns up in informal terminology which can be confusing. It becomes uncertain whether it is a 'hypothesis' or a 'theory', and neither term in this situation is used with the rigour insisted upon by philosophers of science. This rather indefinite hypothesis, operating at some remove from the experimental situation and often decorated with metaphor, just like the 'big idea' hypothesis about CF, also needs to be distinguished from the kind of hypothesis upon which I wish to focus.

4.2 Brief overview of older usages of 'hypothesis'

The concept of hypothesis is open to the same problem of accretion of meanings, some relevant and some obsolete, as encountered with concepts like 'gene' and 'genome', as pointed out by Evelyn Fox Keller (Keller, 2012, 132). Examining the previous philosophical usage of this term may provide insight into how 'hypothesis' might be defined in biology and/or why we have such a muddled concept of 'hypothesis' in biology. Since sorting out the history of the term 'hypothesis' in the Western tradition of scientific research could serve as a research project in itself, I will limit this discussion to a few points. Moreover, without returning to the original texts, interpretation is difficult. For example, with Plato and Aristotle, my discussion turns on word(s) which are translated as "hypothesis": it is difficult to assess accurately the anachronisms inherent in the translation.[65] It is always easy to introduce new anachronisms into this reading. The short summary about 'hypothesis' in the writings of Plato and Aristotle is that neither said much about it. Much of the interest in the more modern accounts of the 17th century is to figure out what later philosophers may have thought these classical texts were telling them.

[64] Very recent developments in CF therapeutics feature drugs which are targeted to correct the physiological defect in individuals who are homozygous for a specific mutation in *CFTR*, the gene whose mutations result in CF. This is a different example where finding the gene was necessary for developing a new treatment.

[65] I am grateful to Eli Diamond, Department of Classics, Dalhousie University, for kindly reviewing my interpretations in an effort to minimize any distortion of interpretation based on relying upon translations of the writings of Plato and Aristotle, and for very insightful discussion.

4.2.1 Classical philosophy: Plato

Plato uses the concept of hypothesis sparingly. In the *Meno* he refers to hypothesis as a kind of heuristic to guide the argument [*Meno* 86-87] (Plato, 1920, Vol I/367). In the *Republic*, Plato clearly assigns 'hypothesis' to one of the lower levels of intellectual endeavour. For example, a hypothesis is an unexamined assumption which dialectic eliminates [*Republic* Book VI 510-511] (Plato, 1920, Vol I/771-2), and subsequently he asserts that the realm of hypothesis is inferior to the realm of principle [*Republic* Book VII 533] (Plato, 1920, Vol I/793). Thus a hypothesis is a handle to an argument, but not much more than that. It serves as epistemological machinery, but it is not as important as the outcome, or indeed as the process of dialectic, for getting to that outcome. These connotations persist in our contemporary usage: a hypothesis is a kind of heuristic, and in an epistemic hierarchy 'hypothesis' is subservient to a 'theory' which, by definition, is more encompassing. Perhaps this modern epistemic hierarchy is related to Plato's relatively lowly position of 'hypothesis' epistemically. Modern usage attaches more importance to 'hypothesis' than we find in Plato's early writings.

In the later work *Parmenides*, Plato makes an interesting observation about how we should employ a hypothesis. The general role of a hypothesis is unchanged: it still serves a heuristic function. However, here [*Parmenides* 136] (Plato, 1920, Vol II/97) Socrates describes a hypothesis as a speculative concept. He argues that the consequences of its being true as well as the consequences of its not being true should be considered. Socrates then goes on to dismantle the argument of his opponent Parmenides, but he does not dismantle this concept of hypothesis in the process. Thus the notion of 'hypothesis' as epistemically significant whether right or wrong appears to be a concept of hypothesis developed by Plato: it persists into our current thinking about how a hypothesis functions.

4.2.2 Classical philosophy: Aristotle

For Aristotle, who follows Plato in this regard, 'hypothesis' comes off as something uncertain and inferior to theory. What I find interesting about Aristotle's utilization of 'hypothesis', however, is its prominent role in logic. In his system of logic Aristotle contrasts hypothetical proofs to probative proofs. The hypothetical proof is a

reductio ad impossibile [*Prior Analytics*; 40b25-26] (Aristotle, 1984, 64): a supposition (the hypothesis) is made, and then it is demonstrated that this supposition leads to a self-contradiction. Thus hypothetical proof takes on the character of arguing by counterfactuals. It is a classic dialectical strategy. This logical manoeuvre seems to me to be very similar to positing a 'null hypothesis' in statistics.[66] Moreover, later in the *Topics*, somewhat like Plato in *Parmenides*, Aristotle describes a process of investigating a hypothesis which seems highly to resemble falsification as such [*Topics*, Book III; 119b34-120a5] (Aristotle, 1984, 201). This may be a formulation of hypothesis that is closest to modern usage. If you accept the notion that Aristotle assigned prime importance to the principle of non-contradiction in his metaphysics and logic (Lear, 2010, 250), the role of hypothetical proof is important: showing that the counterfactual leads to nonsense is sound argument.

Although Aristotle treats hypothesis mainly as a logician/meta-logician, we get an interesting glimpse of a different view of 'hypothesis' when Aristotle writes as a biologist. In *Generation of Animals* Book III, after presenting an account of how bees reproduce (which is wholly unsatisfying to the modern reader), he writes: "… the facts, however, have not yet been sufficiently grasped; if ever they are, then credit must be given rather to observation than to theories, and to theories only if what they affirm agrees with the observed facts" [*Generation of Animals* Book III; 760b30-4] (Aristotle, 1984, 1178). I believe this text provides grounds for inferring that Aristotle recognized that formulation of theories and related hypotheses required sound data. For hypothetical proof to be convincing, the hypothesis had to be of high quality. Moreover, here with Aristotle the biologist, we espy the contrast between regarding hypothesis, or even 'theory' as a logical construct applied to the operation of scientific research—basically, the enterprise of Logical Positivism which still motivates Hempel's work—and the practical realities of actually doing biological research. Even for Aristotle, the strategy from logic is relevant but it requires some operational, not conceptual,

[66] In the *Topics*, Aristotle speaks of "hypothetical deductions" so that we may infer that hypothesis plays this central role in at least one kind of logical demonstration [*Topics* Book I; 108b12-14] (Aristotle, 1984, 180).

adjustment when applied to biology. This may be a nuance of interpretation which would be evident to a contemporary biologist: I have no basis for arguing that it was extremely important within the whole framework of Aristotle's philosophy.

4.2.3 17th-Century "scientific" revolution

Sir Francis Bacon is credited with catalyzing the development of new ways of investigating the natural world. What Bacon was promoting is regarded as an early version of scientific empiricism, seemingly in reaction to the prevailing practice of his day (late 16th to early 17th century), a kind of *ex cathedra* methodology, which was rooted on received theory rather than real data. Indeed Bacon's innovations are generally depicted as a reaction against orthodoxy founded on Aristotelian philosophy. He criticized Aristotelian views deemed relevant to (and problematic for) emerging modern physical science. It is of interest to consider usage of 'hypothesis' at this time of critical development, which we call anachronistically a "scientific" revolution.

In this connection, it is almost impossible not to consider, first of all, one of the most provocative epigrams of that era relating to hypothesis: Newton's "Hypotheses non fingo", generally translated as "I do not feign hypotheses". The connection to Aristotle is interesting and not necessarily negative. In the typical Aristotelian hypothetical proof, a supposition (the hypothesis) is made. It is like saying in algebra "let x = something-or-other" in order to work the algebraic solution. Interestingly, this usage is the same as in 17th-century English for the word 'feign' which was then taken to mean "to assume fictitiously for purposes of calculation"[67]. This is relevant to Newton's famous disclaimer about feigning hypotheses. So we find a persistence of meaning from Aristotle to at least the 17th century. However, Newton's famous declaration is a good deal more complicated than it might appear. It seems completely disingenuous to me to interpret it simply as it appears ("I do not use hypotheses"), and thus I certainly do not see Newton as an early, albeit clandestine, exponent of non-hypothesis-driven research. It is tempting to say that the statement "hypotheses non fingo" is misinterpreted

[67] *Oxford English Dictionary* (on-line version), Definition 4c, (http://www.oed.com.ezproxy.library.dal.ca/view/Entry/69014?rskey=QChIIi&result=2#eid, accessed 6 January 2015 and *circa* 4 July 2013) (Anon.: OED online)

because we translate ‘fingo’ poorly as ‘feign’. I suggest that in this usage ‘fingo’ might mean what we say nowadays informally about hypotheses: namely, I “spin” a hypothesis. This accords with the OED archaic definition above. However, only a few decades earlier Bacon said that the non-Copernicans “feigned eccentrics” and ellipticals to salvage their cosmic theory (Hacking, 1983, 220) (quotation from Bacon’s 1625 essay ‘Superstition’). This comment puts the Latin verb ‘fingo’ into an entirely different light. The problematic word in Newton’s declaration may be “hypotheses”, not “fingo”. Given his rejections of aligning himself with any of the explanatory theories in fashion in his day, his comment may mean that he refused to feign any of these theories. Then the possible translation in current casual English would be: “I do not have truck with theories.” The classical hierarchy of hypotheses and theories seems to inform the meaning here—and confuse it.

The question relating to Francis Bacon then is what he thought about hypotheses and, specifically, what role he attributed to the hypothesis in what his contemporaries called natural philosophy (later in the 19th century identified as a scientific endeavor). A full examination of this problem deserves more than a few paragraphs and exceeds my remit and resources; however, several points deserve attention. First of all, I believe that interpreting Bacon’s concept(s) of hypothesis requires viewing hypothesis not merely in terms of 17th-century diction but more specifically in relation to how the roughly contemporary researchers of natural philosophy used it. Thus, how Newton used the word ‘hypothesis’ is entirely relevant. Secondly, I put “concept(s)” into a potential plural on purpose. Bacon appears to me to have several different concepts in mind. One is the thing postulated in order to get a calculation started: that is, what is assumed fictitiously for purposes of calculation. More importantly, I see clear evidence that Bacon had some notion of a (modern) testable hypothesis. The basis for my interpretation is in the nature of the ‘crucial experiment’, which he invented. Finally, another concept is some sort of broader concept verging on our modern notion of theory—conceptually problematic today as in the 17th century.

Hypothesis as a supposition to get an investigation going is a heuristic use of hypothesis which seems to be a worthwhile residuum of classical theory. It has methodological utility. Perhaps this seems like an uninteresting type of hypothesis, but we still use such hypotheses today. One example is when you are looking at a new but relatively circumscribed problem and can formulate a hypothesis just to govern a well-designed experiment: in such a situation the results are likely to be of importance whether this testable hypothesis is supported or not. The problem with hypothesis as something which is postulated as a basis for calculation is that it can get perverted into a salvage mechanism for incompetent explanations of data. In such a situation, the "let's suppose" hypothesis is not leading to the experiment itself. Instead it is serving more like a Hempelian auxiliary hypothesis to ease the available empirical data into place. If such hypotheses get elevated to being accepted as theory, then they function as a kind of dogma. It seems that this kind of hypothesis verging on theory was what Newton was rejecting and also what Bacon was rejecting in his oblique allusion in his essay "On Superstition".

The 'crucial experiment' depends on a much more sophisticated notion of what a hypothesis is. Indeed I would describe the hypothesis of a crucial experiment as a clearly articulated testable hypothesis that governs excellent experimental design—but this is none other than the modern hypothesis of hypothesis-driven research as I have defined it, or more specifically, what I will call the 'proximate' hypothesis. If we accept that Bacon invented the crucial experiment (Vickers, 1992, 510; Urbach, 1987, 169), then we need to recognize his concurrent invention of something very like the modern testable hypothesis. Bacon's concept of the *instantia crucis*, the experiment like a signpost pointing in one direction or the other (Bacon, 2000, Book II.36), does not work unless the hypothesis (as well as its partner hypothesis pointing in the other direction) is well-crafted. Such a hypothesis is not any old "let's suppose for the sake of argument" hypothesis. It is not surprising that this type of experiment and its type of hypothesis assumed great importance for researchers, even though they were working in the physical sciences, later in the 17th century.

With Bacon, the relationship between hypothesis and theory is as difficult to ascertain as that relationship is today. The classical tradition of a hierarchy whereby hypothesis is inferior to theory seems to hold. Peter Urbach offers a possible clue about Bacon's view of hypothesis in relation to theory: "Bacon welcomed hypotheses (in the sense of theories going beyond what is immediately given in perception) from the very beginning, and that they were always the intended product of the interpretative method" (Urbach, 1987, 34). This comment really turns on how we understand Bacon's empirical method, which anchored investigation of the natural world on the collection of actual data. Bacon's famous analogical argument about scientific method is critically important to this analysis. He contrasts ants and spiders and chooses a separate methodology:

> Those who have treated of the sciences have been either empiricists or dogmatists. Empiricists, like ants, simply accumulate and use; Rationalists, like spiders, spin webs from themselves; the way of the bee is in between: it takes material from the flowers of the garden and the field; but it has the ability to convert and digest them. This is not unlike the true working of philosophy; which does not rely solely or mainly on mental power, and does not store the material provided by natural history and mechanical experiments in its memory untouched but altered and adapted in the intellect. (Bacon, 2000, Book I.95)

Thus, as Hacking has pointed out (Hacking, 1983, 149), randomly collecting data is not the thrust of Bacon's method. In fact, Bacon is describing experimental method, not mere empiricism. However, collecting data is a key aspect of that method. Bacon demands that raw empirical data must be transformed into evidence and shaken down in relation to each other.[68] His analogical argument here fails to specify the product of

[68] Recognizing that philosophers of science have used or abused Bacon's ideas for the purpose of advancing their own ideas (Vickers, 1992, 511) (and that this practice is objectionable), I want to point out nevertheless some similarities between Bacon's notion of hypothesis verging on theory and my concept of system-driven research. The point of the argument about bees is that the researcher collects data and processes these data. Data are not collected merely to have a collection of data. The outcome of the processing could be a theory. It could be a pattern. (I do not wish to invite here the easy equivalence of theories and patterns, although it might seem appropriate for the physical sciences.) Moreover, that outcome could be a testable hypothesis. It could also be a mandate to collect additional data collection, since Bacon specifically eschewed jumping to conclusion based on insufficient data. What I am suggesting is that Bacon's thought is relevant to the system-driven research and its methodology which I am investigating. He is not a mere data-collector.

such conversion/digestion. Possible outputs include something which serves as the interpreted quintessence of the data. Mary Horton states that Bacon would not describe such an output as a theory (Horton, 1973, 247) but perhaps as an axiom of the first vintage, which she suggests is what we nowadays might call hypotheses—"tentative statements, the purpose of which is to suggest experiments or observations that will either 'support' or 'correct' them" (Horton, 1973, 247). The problem here, as Horton soon points out, is that 'hypothesis' as it was used in the 17th century generally had a negative connotation, and she offers 'preconception' as an approximation to what the 17th-century hypothesis was.[69] Her description of the use of 'hypothesis' in Bacon's time (and how Bacon uses it) is quite specific: "denotes either a conjecture that has already been refuted or discarded, or an assumption that is being maintained in spite of evidence or argument to the contrary" (Horton, 1973, 248). This clarification accords with my interpretations above and does not materially interfere with my claim that with the *instantia crucis* Bacon describes something tantamount to a modern 'hypothesis'. However, we end up not much further ahead sorting out how theory and hypothesis relate to each other, except to note that their connotations, if not denotations, differ from modern usage, and also that a hierarchical sense persists. Theory (axiom of a later vintage?) may yet be a somewhat more advanced, or developed, epistemic entity than hypothesis (axiom of the first vintage). We may find in Bacon's lexicon concepts which bear directly on key modern concepts of experimentation.

4.3 Modern conceptions of 'hypothesis'

4.3.1 Victorian contributions: Whewell on 'hypothesis'

Whewell and his intellectual 'set' (namely, John Herschel, Charles Babbage, and Richard Jones) took Sir Francis Bacon's ideas very seriously and implemented them into the 19th century (Snyder, 2011, 37-9). What seems to have appealed to them was that in effect Bacon insisted on basing scientific activity on scientists' gathering abundant

[69] She digresses momentarily to suggest that Newton's 'I feign no hypotheses' referred to these sorts of preconceptions (Horton, 1973, 248, footnote 12).

observations and then putting their imprint of interpretation upon them.[70] If we follow the recent "back-to-Bacon" turn in understanding experimental science, as formulated by Ian Hacking (Hacking, 1983, 150), as biological/biomedical researchers in the age of molecular biology, then it makes sense to pay attention to Whewell as the chief mediator of Baconian thinking into the modern age of science.

What Whewell actually says about hypotheses deserves further elaboration. He sees a hypothesis as a flexible supposition (Whewell, 1968, 36). He points out that scientists typically articulate several hypotheses and then choose one which seems right[71]; while he does not describe the process of hypothesis formulation, he suggests it requires a creative talent (Whewell, 1968, 130) and he attaches value to that talent.[72] He points out that testing a hypothesis involves ascertaining whether relevant facts "have the same relation in the Hypothesis which they have in reality" (Whewell, 1968, 137). Interestingly, despite being the exponent for hypotheses, he points out that hypotheses can get in the way of scientific work by unduly narrowing the field of investigation and end up "prejudicial" (Whewell, 1968, 147-8). He particularly commends hypotheses which are devised for one set or type of data and then proved applicable to another different set. Whewell writes mainly in the context of the physical sciences, and thus the main importance of his views on hypothesis for the discussion here is to see his contribution to the establishment of hypothesis-driven experimentation as the way to do science.

4.3.2 Pervasive 20th-century concept of 'hypothesis': Hempel's hypothetico-deductive method

Hempel is the most influential contemporary theorist about the nature of 'hypothesis'. Therefore a brief survey of Hempel's hypothetico-deductive method

[70] I recognize that this is an anachronistic description of Bacon's position. The very terminology of 'scientist' did not appear until the 19th century when Whewell invented the term. However, it is very awkward to write that the individual adhering to Bacon's concept of the scientific method would gather data and put his/(her) imprint upon the data.

[71] Ohkrulik will pick up on this comparative approach to hypothesis formation and evaluation.

[72] In a charming comment, he says: "A facility in devising hypotheses, therefore, is so far from being a fault in the intellectual character of a discoverer, that is, in truth, a faculty indispensable to his task" (Whewell, 1968, 145). This suggests, albeit obliquely, that some skepticism about what will become the HDM existed in Whewell's time, just as we see with the OES currently.

(HDM) is important. Quite clearly, Hempel rejects the validity of data-collection as an essential component of doing science. He calls this "*narrow inductivist conception of scientific inquiry*" untenable (Hempel, 1966, 11, original italics). His first reason is that such scientific research could never get started: it would take literally forever to collect all the data! Such an objection might be interpreted as facetious. If it were serious, it would require revision given today's computing capabilities which are far beyond what could be imagined in the early 1960s, when a handheld calculator was a novelty. (Proponents of eScience would say this revolution is taking place.) More likely, it simply reflects Hempel's view that relevant data sets will be infinitely large in order to be definitively specific. In any case, Hempel is absolute in his dismissal: "in sum, the maxim that data should be gathered without guidance by antecedent hypotheses about the connections among the facts under study is self-defeating, and it is certainly not followed in scientific inquiry" (Hempel, 1966, 13). Moreover he holds that hypotheses must direct analysis and classification lest such analysis and classification otherwise be blind (Hempel, 1966, 13). Hempel rejects induction as a scientific thought process, but he is perhaps a little vague about the process of hypothesis formulation. He says that "scientific hypotheses and theories are not *derived* from observed facts, but *invented* in order to account for them" (Hempel, 1966, 15, original emphasis). He does not elaborate on the process of hypothesis invention, which operates somewhere in the context of discovery.

Hempel's conception for scientific reasoning takes the form of some kind of proof. Specifically, according to Hempel's hypothetico-deductive conception, "empirical evidence supports or *confirms* a theory[73] exactly if that evidence is *logically entailed* by the evidence conjoined with additional assumptions" (Weber, 2005, 88). Within the HDM, Hempel ends up describing two main models of reasoning about experiments. The D-N model means "deductive-nomological" or deduction based on laws. A model devised by Hempel somewhat later is the "inductive-statistical" (I-S) model, similar to

[73] This is a typical example of where the word 'theory' gets substituted in context for 'hypothesis' since Hempel is clearly talking about hypotheses, not theories. The additional assumptions are auxiliary "hypotheses" about initial and boundary condition, or how/whether the analytical equipment functions.

the D-N model but where the logical certainty of the D-N model is replaced by a quantified uncertainty, more relevant to real-life situations. Statistical assessment provides the quantification for the degree of certainty or uncertainty that pertains to the relationship of the premises in that particular account. With either model, Hempel appears to make the assumption that the only objective of science is explaining and predicting, and quite simply, for Hempel explanations are arguments and these explanations necessarily involve laws, as shown by his deductive-*nomological* model.

Hempel is meticulous about defining the elements of the HDM. He gives a definite definition of what he means by 'hypothesis': "…whatever statement is under test, no matter whether it purports to describe some particular fact or event or to express a general law or some other, more complex, proposition" (Hempel, 1966, 19). Notably a hypothesis is a proposition, but it could be a mathematical formula (Hempel, 1966, 20). He parries the notion that a hypothesis is some sort of conditional by designating what he calls 'test implications of a specific hypothesis' as conditional. A test implication is expressed as an if-then statement and it describes the state of affairs that may be expected to be found, according to the hypothesis, under specific test conditions (Hempel, 1966, 19). The real character of the test implication stands out as an important facet of the Hempelian HDM. It appears to be somewhat different from a hypothesis, partly because it is particular (and departs from the generality of the hypothesis) and partly because it has a canonical formulation, expressed as an if-then statement. This usage makes for neat book-keeping. It meets the criticism of those who think that hypotheses should never be expressed as if-then statements: the hypothesis can be expressed as a proposition and its corresponding test implication is expressed as the if-then statement. However, I believe this terminological manoeuvre is something of a mirage. If a hypothesis is only viable if it can be tested, then the boundary between hypothesis and test-implication becomes very vague in real life when actual research is being performed to test the hypothesis. It seems just as reasonable to say that the test-implication is a version of a hypothesis restated in the context of the experiment devised to test it. This interpretation of Hempel's 'test-implication' demonstrates that it

serves as the connector between the hypothesis and the specific experimental design. Hempel might reject this view strenuously because he holds that there is distinct difference between a hypothesis and a test-implication, namely, the hypothesis must be combined with appropriate auxiliary assumptions (that is, auxiliary hypotheses) in order to yield a test implication (Hempel, 1966, 31). The argument remains unconvincing. Since Hempel is committed to the idea, redolent of classic Logical Positivist orthodoxy, that a proposition can be a scientific hypothesis only if it is "amenable to objective empirical test" at least in principle (Hempel, 1966, 30), then the relationship between hypothesis and test-implication must be very close indeed. I do not interpret Hempel as having test-implications subservient to a hypothesis: rather, he has a hypothesis operationalized as test-implications in order to actuated as testable, and thus be a scientific hypothesis. In effect, the hypothesis qua test-implication relates directly to experimental design.

In summary, Hempel's HDM serves as the standard for contemporary epistemology of experimental science. His theoretical models are diversified to include situations where laws support straightforward deductive reasoning (D-N) and situations where deductive reasoning involves uncertainty quantifiable by statistical methodology (I-S). Historically, this conceptual apparatus has been most relevant to (and basically inspired by) the physical sciences. The strength of the Hempel's propositional and logic-based approach is its clarity and elegance. A possible weakness with HDM is the critical assumption is that the only objective of science is to explain and predict, whereas it could be argued that science is also about accurate description and comprehension (Salmon, 1978, 684). In biological contexts, description and comprehension might serve to delineate a mechanism, but knowledge of mechanism in a complex situation where some variables operate as 'wild cards' does not translate necessarily into accurate prediction of how the entity will actually function. So, with a biological entity, you might have a complete expertly-derived mechanistic explanation, but it does not mean you can predict what the entity will actually *do* in a given situation. Moreover, Hempel seems to end up wide of the mark in terms of addressing real-world problems. He

seems to be operating at some great distance from the laboratory or research space where empirical data are being produced. Hempel's models as logical constructs do not easily allow for identifying confounders or actual error in the components of what is doing the explaining, even when the shortcomings of auxiliary hypotheses are exposed and rectified. It seems fair to point out, however, that these shortcomings are more evident when the models are applied to biological systems, as opposed to physical systems.

4.3.3 Concept of 'hypothesis' as portrayed by contemporary philosophers of biology

Against this background, it makes sense to review the prevailing usage of 'hypothesis' among contemporary philosophers of biology. What is really interesting is how little literature within the domain of philosophy of biology actually addresses the problem of what a hypothesis is in relationship to biological research. It is as if this were a blind spot within philosophical consideration. As already pointed out at the very beginning of this chapter, Weber describes a hypothesis for biological research as a mechanistic explanation, more complicated than an algebraic formulation of causal relationships between variables (Weber, 2012, Section 3.2).

Deborah Mayo does discuss hypothesis specifically and she presents a hierarchical schema: both her notion of 'hypothesis' and her hierarchy are different from mine. In *Error and the Growth of Experimental Knowledge*, Mayo presents a conceptual schema for examining experimentation. She intends this to be "a framework that permits us to delineate the relatively complex steps from raw data to scientific hypotheses" (Mayo, 1996, 129) and she draws heavily on the early work on Patrick Suppes on how statistical analysis structures our notions of how experimental investigation works. Suppes calls for a hierarchy of models relating to data, experiment, and theories (Suppes, 1962, 260-1). Importantly, the conceptions of 'model' in such theorizing are highly technical and specialized. Moreover, Mayo's schema of models is somewhat different from Suppes's conception. She envisages primary models, experimental models, and data models; the hierarchy runs from the primary model downward so that in descending the hierarchy we get closer to the actual data (Mayo, 1996, 133-5). Mayo's primary model is ostensibly complex. It includes a primary

hypothesis, which itself represents an operational, testable version of the theory being investigated. In fact, Mayo sometimes refers to this element of her schema as the "primary scientific hypotheses or questions" and specifically points out that what she calls a "substantive scientific inquiry" gets broken down into at least one or more hypotheses, which she calls local or topical hypotheses. Each of these "corresponds to a distinct *primary question* or *primary problem*" (Mayo, 1996, 129, original emphasis). These primary hypotheses are not only specific and narrow in scope: they tend to be quantitative in their formulation, devised to estimate numerical values arising from a theory. Then we move down her hierarchy, and the subsequent models all relate to the investigation of this primary hypothesis. They are similarly complex. Mayo's experimental models "serve as the key linkage models connecting the primary model to the data, and conversely" (Mayo, 1996, 129). Her data models provide the theoretical and practical machinery for modeling data and thus providing the modeled data that must be linked to the experimental model. Modeled data, in general, are data which have been evaluated statistically. It is noteworthy that experimental design is a consideration in data-modeling: it does not particularly figure into the experimental model (Mayo, 1996, 135). This feature seems peculiar to me: it may reflect her greater commitment epistemically to statistical methods than to experimental methods. At the very bottom of Mayo's hierarchy, lying outside the hierarchical structure, is a collection of *ceteris paribus* assumptions and also a heap of rejected considerations.[74]

Mayo's robust but complex analysis clearly relates specifically to the research style mandated by the HDM. It provides a way to move critically from the hypothesis which operationalizes some aspect of a larger hypothesis through the experimental process including data analysis. Mayo's primary hypothesis, also called a primary question or primary problem, is the hypothesis closest to the experiment itself. Her primary hypothesis drives the experiment. On balance, Mayo is much more interested in

[74] For example, when Mayo discusses the design and analysis of a clinical trial, she points out that the statistical analysis of the characteristics of the two cohorts showed no statistically important differences for the relevant parameters of age, pregnancies, socio-economic status, etc., but she also mentions that astrological sign was not regarded as relevant and thus was relegated to this discard pile.

the process of data evaluation than hypothesis formulation, and for this project she provides a masterly implementation of statistics (and biostatistics) to the philosophical treatment of experimental procedure. While this may inform some later stages of evaluating system-driven research, the concept of hypothesis she presents here is quite narrow. Introducing the notion of models[75] does not enhance the clarity of what is meant by a hypothesis. Moreover, it is difficult to avoid the suspicion that, broadly-speaking, Mayo's concepts here line up with Hempel's. Specifically, the primary hypothesis is similar to Hempel's hypothesis as such; the experimental model incorporates key features of bridging hypotheses; finally, the data model which mediates how analyzed data relate to both experiment and hypothesis has intimations of Hempel's auxiliary hypothesis.

By way of overall summary, the concept of 'hypothesis' has a long and somewhat confusing history. We find intimations of the current usage of 'hypothesis' in classical philosophy, mainly Aristotle. Since this was the intellectual inheritance for the early modern experimentalists, it is perhaps not surprising that the meaning of 'hypothesis' became quite variable, as we see for example with Newton's usage. Whewell's appreciation of the nuances of hypothesis seems to have been less influential than his contribution to the development of the modern hypothesis-driven research methodology. Philosophical assessment of scientific method in the 20th century has been restricted to hypothesis-driven research. The main thrust is Hempel's HDM. When subsequent 20th-century philosophers of science include 'hypothesis' in their analysis of how scientific experiments are designed and evaluated, they defer to Hempel. In general, their focus is on the physical sciences. Hempel's conception of what a

[75] The concept of 'model' in relation to experimental design receives much attention from Mayo and her contemporaries. Giere's notion of models is the best example. He holds that scientists construct models as a way to expedite "the process of figuring out how the world works" (Giere et al., 2006, 20). He then classifies models as scale models, analog models (where the model functions like a kind of metaphor for the concept of interest), map-like models, and finally theoretical models which play a direct role in the process of scientific reasoning as he analyzes it (Giere et al., 2006, 21-5). The simple description of his analysis is that scientific reasoning involves sorting how the adequacy of fit between the model and the real world. In general, Giere's rendition of scientific method, which emphasizes how a hypothesis mediates prediction and evaluation of the results of an experiment, ends up congruent with the HDM.

hypothesis is and how it functions is as a logical construction, and his problem is to develop an accurate and productive connection to the world of the laboratory, or wherever research is performed. I have suggested that he accomplishes this task by postulating the test-implication as an operationalized version of the hypothesis such that it mediates between the logic of hypothesis-testing and practicalities of experimental design.

4.4 Hypothesis assessed

If philosophers of biology have tended to avoid specifying what is meant by hypothesis, the reason is that the character of hypotheses in biological and biomedical research is complicated. There are various kinds of hypotheses. Some of these hypotheses may fit a Hempelian mode and others not. I contend that much of the difficulty here arises because different kinds of hypothesis can be in play simultaneously in any given biological research project. This characteristic becomes apparent when a researcher doing omics research asserts confidently that s/he has a hypothesis for this research.

More specifically, I submit that perplexity around 'hypothesis' in biological and biomedical research arises because the word is used in at least three different ways. First of all, there is the **hypothesis which locates the general research domain within which you are working**. This overarching hypothesis generally specifies which sector or corner of that domain is relevant. This is generally a broad statement expressing concept(s) upon which there is general agreement. In general, there will also be consensus that the broad issue/problem deserves attention: it is significant. This overarching hypothesis will typically affirm and articulate the importance of the research proposed.[76] Next there is a **hypothesis which is really a hunch**. This states a

[76] It is tempting to say that the overarching hypothesis benchmarks the paradigm within which the research is to be conducted, where 'paradigm' has Kuhnian connotations. At the present time, I prefer to resist this classic and very pervasive nomenclature. We may be witnessing a period of transition in paradigm. Proteomics seems to be most appropriately placed in a 'post-genomic' era whereas genomics must be in the 'genomic' era. However, the paradigm here might be based on how we investigate biological/biomedical problems, and then I would argue that all of omics resides in a paradigm characterized by high-throughput, data-intense methodology. I am not arguing that Kuhnian terminology

broad notion which generally indicates what the research is about, usually the research problem as such. Sometimes this is actually the objective of the research or its rationale which activates the project itself. It typically indicates what is novel or innovative about the research effort itself, in other words, the special insight or expertise the researcher is bringing to the project. Finally there are **individual, rather narrow hypotheses related to specific experiments.** A hypothesis of this sort is directly responsible for the design of an individual experiment. Such a hypothesis could be the hypothesis indicated at the beginning of describing a specific aim. Each specific aim should flow logically from the bigger hypotheses and represent a 'logical' attack on proving or at least investigating that broader "hunch" hypothesis. We can call this narrower hypothesis the hypothesis proximate to the experiment. Except for these proximate hypotheses, none falls neatly into the hypothetico-deductive method. The proximate hypothesis bears strong similarities to Hempel's hypothesis as such and to Mayo's primary hypothesis. However, for my proximate hypothesis to be comparable to Hempel's hypothesis, the latter must be operationalized into a testable version: that is to say, it must be converted into the corresponding test-implication. The important point to be made with my nested hierarchy of types of hypotheses is that when people say that system-driven research has a hypothesis, they are thinking of one of these broader ones. Usually they are referring to the 'hunch' variety of hypothesis, but certainly they will invoke the overarching domain-identifier if all else fails.

What distinguishes hypothesis-driven research from system-driven research is the nature of what is actuating the research most proximate to the experimental design. As we have seen, for hypothesis-driven research, a detailed hypothesis governs experimental design. In other words this is what determines the experimental design itself. That hypothesis is the type of hypothesis which fits with the HDM. For system-driven research, the system being investigated directly informs experimental design, not a hypothesis of the HDM variety. The character of 'hypothesis' in biology has been

---

is irrelevant, only that in biology/biomedical research at the present time the conceptual boundaries are not entirely evident.

elusive because multiple layers of hypothesis are at work. In terms of system-driven research what is driving the experimental design is not a proximate hypothesis, but instead the system itself. The presence of hunch or domain-designating hypotheses does not alter the essentially system-driven nature of this research. What really counts is what governs—indeed, what *drives*[77]—the experimental design.

We can see how this works with some concrete examples. In Wilson disease, copper accumulation in the liver damages liver cells and often leads to cirrhosis. The mechanism of damage was unknown. Based on circumstantial evidence, it was thought to involve toxic chemicals, known as "free radicals", generated within the liver cell by the accumulated copper because of it can be in an oxidized or reduced form thus generate free radicals. I examined this problem directly. I *hypothesized* that copper generates free radicals in liver cells. Under standardized conditions, I grew liver cells in Petri dishes, and then I treated cells in one set of dishes with a range of increasing concentrations of copper for 24-48 hours. Then I directly measured free radicals using a highly reliable new assay. As a control, I did the same experiment with zinc instead of copper, because zinc resembles copper chemically (as it also has a valence of 2^+) but effectively cannot generate free radicals. The assay for free radicals was positive with copper (the amount of free radicals generated rising with increasing dose) but with zinc it was low at all doses. Thus my hypothesis was substantiated: exposure to copper generates free radicals in liver cells. I then went on to do the essentially same experiment, but this time I measured the amount of a protein which combats injury due to free radicals. I found that the concentrations of this protein increased only in the copper-treated liver cells. In this second experiment I tested a new hypothesis, namely, that in the presence of increasing amounts of free radicals, capable of causing cell injury, cells would express more of a protective protein. This hypothesis arose from the first.

[77] In this usage 'drive' has two main connotations: guiding and governing. It is not intended to elicit the inference that the 'thing' being driven is necessarily a machine. I mention this latter point because of the pre-eminence of mechanistic (verging on "mechanical") thinking in current biological/biomedical research.

These experiments clearly established oxidative stress as a key aspect of the mechanism of hepatocellular toxicity. Describing the disease mechanism in greater detail required a comprehensive understanding of how copper is handled in liver cells. For this purpose my colleagues and I used an OES approach for addressing this problem, namely metalloproteomics. An organism's actual variability—including its developmental changes but also responses to stress or noxious agents—is indicated in protein expression. For Wilson disease, we focused on the metalloproteomics relating to copper and were interested in proteins which have copper-binding capacity or contain copper. We ground up cultured liver cells to produce a liquid slurry and ran that liquid through a column specifically designed to capture all proteins capable of binding copper. Then we identified the proteins we captured. We found numerous proteins, some not previously thought to bind copper. Using omics here (without a specific hypothesis) extended our knowledge of copper-handling in liver cells dramatically and opened new vistas relating to disease mechanism. This approach is very different from the focused experiments designed to investigate whether copper generates free radicals. As a matter of fact, the protein assayed in the second experiment described above is one that we had identified in our metalloproteomics work.

There are other types of research which do not involve omics but where, strictly speaking, investigation is not clearly governed by a hypothesis. In the years when we were setting out to identify the gene which is abnormal in Wilson disease, I do not recall clearly enunciating a hypothesis. In fact, our second big grant application leading up to identifying the gene specified no hypothesis at all, although we pointed out that Wilson disease was a significant disease.[78] In the early 1990s the point was merely to clone the gene, a labour-intensive process quite different from the relative ease with which genes are identified today. One can imagine the relevant hypotheses. The overarching hypothesis was something like: Wilson disease is a clinically-important monogenic autosomal recessive disease. In other words, mutations in one gene account for a

[78] That grant application was to Physicians' Services Incorporated, a small independent Ontario granting agency for biomedical research. The grant was successful.

liver/brain disorder we call Wilson disease.[79] The 'hunch' hypothesis for the initial work commencing in 1988 or so was that the relevant gene was on chromosome 13 near the locus known as the retinoblastoma (RB) locus: we could narrow down its location by characterizing that length of the chromosome. This was a pretty strong hunch as there was a lot of evidence for the close spatial connection to the RB locus. Once the stretch of the gene where this gene was most likely to be found was narrowed down through complicated studies of numerous patients, bits of the gene were inserted into yeast and the resulting fragments were pieced together linearly, somewhat like the linear chain of tiles in a game of dominoes, based on overlapping sequences. There was nothing much hypothesis-driven about that step, just as there had been nothing hypothesis-driven about the previous clinical research to identify patterns of DNA-linkage. Then, in early 1993, the gene for another human disease which involves copper disposition was cloned, and thereafter it was hypothesized that the gene abnormal in Wilson disease might have similar features. Looking for these specific features formed the basis of the ensuing experiments which led to it identification. So here we have a mixture of HDM and OES experiments[80] to identify a gene.

Another important line of research, besides gene identification, is to work out the structure of a protein or some molecule in a cell. It is an important part of the research portfolio of structural biology, a contemporary iteration of biochemistry. Perhaps the most famous research relating to the structure of a molecule is the 1953 description of the structure of deoxyribonucleic acid (DNA) by Watson and Crick. Reading the brief communication to *Nature* (Watson and Crick, 1953, 737-8) and also the more detailed paper to *Cold Spring Harbor Symposia on Quantitative Biology*

[79] It is of some scientific interest that as of early 2013 this point is at least disputed. One recent paper (Sheffield) asserts that their extensive genetic data argue against anything but a monogenic disease (Coffey et al., 2013, 1479). Meanwhile a handful of patients have been reported who have a disease clinically very similar to Wilson disease but with a different genetic mechanism. Experience with at least one other hereditary liver disease militates against the Sheffield group's opinion. 99% of patients might have one genetic basis but the other 1% might have a different one. However, in the early 1990s this was an effective and acceptable overarching hypothesis.

[80] It is possible that 20 years later the approach would be purely one of omics. The recently developed techniques of 'exomics' permit rapid and direct identification of genes in monogenic disorders (Bamshad et al., 2011, 745; Fuchs et al., 2012, 609-11).

(Watson and Crick, 1953, 123), one finds that no hypothesis is articulated. It seems unlikely that it was left out as an oversight. No hypothesis was specified because there was none. (In fact, in the *Cold Spring Harbor Symposia* paper, they indicate that the structure permits them to propose a hypothesis about the molecular basis of self-reproduction of genetic material (Watson and Crick, 1953, 123)). What Watson and Crick really wanted to do was describe the structure of DNA. This is similar to system-driven research except that in this case the "system" was the molecule. It also seems fairly obvious from that historical record that the masterly X-ray crystallographer, Rosalind Franklin, whose images of the DNA molecule were critical to getting the structure, thought that taking more and better images would permit the structure to disclose itself (Gibbons, 2012, 69). I see no reason for doubting the eventual success of this approach, which is notable for being open-ended and non-hypothesis-driven. The methodology adopted by Watson and Crick seems extremely different. According to Watson, their method was adapted from Linus Pauling's work: "the α-helix had not been found by only staring at X-ray pictures; the essential trick, instead, was to ask which atoms like to sit next to each other" (Watson, 1968, 50). Pauling's other methodological peculiarity was to explore that question via space-filling molecular models, not pencil-and-paper drawings. Physical model-building could be considered as equivalent to *in silico* experimentation in the era before powerful computers were a commonplace. So then the question becomes whether Watson and Crick employed a hypothesis-driven experimental design or not. Watson and Crick may have thought that what they were doing was spinning numerous small hypotheses. (I cannot find an explicit statement to this effect in Watson's reminiscence *The Double Helix*.) By this interpretation, lines of thinking such as "maybe it is a triple helix" are hypothesis-like and tested in a purposely-designed experiment—their scale model depicting the stereo-chemical and electrochemical specifics of interactions within a molecule. However, their method can be interpreted differently, without invoking hypothesis-dependence. Their physical model was a highly disciplined, reality-testing extension of their ability to visualize a structure. It permitted them to say "do I see a molecular structure with

'feature x'?" and then determine objectively whether that a structure exhibiting 'feature x' could work within the known stereo-chemical and electrochemical constraints. Without the physical picture, it was easy to be misled by a mental picture. In other words, the physical model extended their ability to see a structure accurately. Each of the "maybe it's …" speculations was not a hypothesis so much as a different way of looking at the possible structure. Each "maybe it's …" was a separate interrogation of the structure.

Of course, with this historically prominent example, it seems evident that there were some hypotheses in play. When Francis Crick announced over lunch at the Eagle that "we [he and Watson] had found the secret of life" (Watson, 1968, 197), he might have been disclosing an over-arching hypothesis for this research, namely, that the structure of DNA is significant because it contains the secret of life. Given that α-helices had recently been revealed as important in the structure of biomolecules, their hunch hypothesis seems to have been that the structure of DNA was some sort of helix.[81] It is possible to spin some other potential hypotheses either before or after the fact: DNA is a complex molecule with a regular structure; DNA has a structure which permits storage of genetic data for archiving and retrieval; the structure of DNA permits its role as the repository of genetic information. Some of these hypotheses actually were generated after the structure was determined. Like hypothesis-driven and system-driven research, structure-driven research has over-arching and hunch hypotheses, and it can generate further hypotheses. Although I allow that there is an interpretation of structure-driven research design which seems to invoke numerous minor experiments addressing a succession of small hypotheses, I believe that it is more productive to see structure-driven research as non-hypothesis-driven. Experimental data generate a possible structure which is then evaluated against what is known about molecular structure in general and what is already known about the molecule under investigation. The problem with the DNA example is that Watson and Crick tended to propose a

[81] Watson says: "We could thus see no reason why we should not solve DNA in the same way. All we had to do was to construct a set of molecular models and begin to play—with luck, the structure would be a helix" (Watson, 1968, 50-1).

speculative structure ahead of having the experimental data in hand. As I interpret their method, they would come up with a hunch and then see whether the data could bear it out. In fact it would not be far-fetched to call them the masters of the *hunch*, though not necessarily of the hunch hypothesis. This also seems to be consistent with the opinion rendered by Gibbons (Gibbons, 2012, 68-9). Accordingly, I continue to contend that describing a protein structure (or a nucleic acid structure, or possibly the structure of a virus) constitutes looking at a very tiny system directly and in its entirety, a kind of non-hypothesis-driven research which (by analogy relating to size of the object investigated) is informative about system-driven research.

Of course, much biological/biomedical research is hypothesis-driven. A definition for the type of hypothesis which governs experimental design in biological/biomedical research which *is* hypothesis-driven is required. This is the really operational 'proximate' hypothesis, in contradistinction to the big-picture rationale and the motivating hunch. It has close affinities to the concept of hypothesis as such in Hempel's construct of the hypothetico-deductive method and also to what Deborah Mayo calls a primary hypothesis. Admittedly, my proximate hypothesis governing the experimental design is more directly similar to Hempel's test-implication, which is to say the version of Hempel's hypothesis operationalized into a construct testable by a specific experiment. However, the two really go together. Regrettably, Hempel never ventures into the domain of biological research. Mayo does not devise her 'primary hypothesis' specifically for biological/biomedical research either even though she employs examples from that research domain. I propose that a hypothesis is an explanatory proposition which is susceptible to investigation, and for the biologist it is about an entity or state of affairs in the biological world. Thus this description can be expressed more succinctly for biological and biomedical research: a proximate hypothesis is **an explanatory proposition (about an entity or state of affairs in the biological world) which is susceptible to investigation, which is to say, to being tested.** The explanatory character tends to be mechanistic: that is, it relates directly to biological mechanism. The epistemic advantage of this definition is that it is general enough to achieve the

versatility required for the diversity of biological/biomedical research. It is not limited to law-like mathematical formulations of causal relations. It relates to mechanism, which is important in biology, and incorporates the characteristic of testability. It is possible that such a hypothesis could be formulated as a conditional (an 'if-then' statement), especially if it were precisely a Hempelian test-implication, but it is not a requirement within the purview of the more general definition I am offering to relate to hypothesis-driven research as a whole in biological/biomedical research.

4.5 Structured comparison of 'hypothesis' in HDM and by my analysis

Given the detailed consideration of Hempel's concept of 'hypothesis' in the HDM and my analysis of 'hypothesis' as actually utilized in biological/biomedical research, which will form part of the analysis leading to the OES, the question arises as to how the two broad conceptions of 'hypothesis' compare. Pointing out that Hempel did not seriously consider biological research and then pleading incomparability is not an adequate response. At least conceivably, physicists may want to examine a system directly as a whole and employ non-hypothesis-driven research design(s).

First of all, it seems unlikely to me that there is an HDM-equivalent for the over-arching hypothesis. It may be that my 'over-arching hypothesis' would be an anachronism if imputed to Hempel's analysis because the over-arching hypothesis may be a post-Kuhnian feature demarcating the paradigm, within which the normal science is being performed. More practically, it is a kind of advertisement of worthiness, part of the rhetoric for why this line of research is important. So if the over-arching hypothesis is "how hepatocytes communicate with each other is vital for maintaining liver function", it is immediately apparent that the researcher is a cell-signaling researcher and probably highly conversant with liver physiology. If I write as my over-arching hypothesis that copper disposition in hepatocytes is critical to whole-body copper metabolism, then I am declaring the broad topics of my research and seeking consensus that the research is worth the effort. To some extent this is the product of the times in which we do biological/biomedical research, when we have to advocate for our research against competing claims for resources. These issues did not apply in the early-to-mid

20th-century. Moreover, I would argue that the Logical Positivists took the inherent value of scientific investigation as an assumption. Defending a specific line of scientific investigation would be redundant. It seems unlikely they anticipated the proliferation of technology which has enabled the broad array of biological/biomedical research currently possible.

I have already suggested that the hypothesis which is really operation in determining experimental design (that is, the one dong the driving) is similar to Hempel's operationalized hypothesis in the HDM, and also to Mayo's primary hypothesis, which itself has close parallels to hypothesis in the HDM. For this similarity to hold up, the HDM hypothesis must be converted to the test-implication, a move which basically takes it out of the realm of logic and into the realm of the laboratory or wherever such research is actually being performed. The results of the experiment are analyzed via the logical apparatus. For hypothesis-driven research as I conceive it, this remains the basic operation.

That leaves the 'hunch hypothesis' according to my scheme. Whether HDM encompasses anything equivalent to the hunch hypothesis is a difficult question. The difficulty here is equivocation at the practical level as to how the relationship between the hypothesis and the test-implication really works. In theory, I believe they are essentially identical; in practice, the test-implication could take a novel/risky hypothesis (much like a hunch hypothesis) and transform it into something testable. I would argue that the novel/risky version had been reconstructed into something narrower and testable as a working hypothesis and then converted to the test-implication, but an HDM proponent might argue otherwise. For the HDM the originality or novelty of the hunch hypothesis, which is its distinguishing characteristic—"the new 'something' the researcher is bringing to the research effort"—might be too close to the context of discovery to merit any serious attention. Thus my inclination is to deny correspondence between the hunch hypothesis and hypothesis in the HDM.

Both the HDM and my 'hypothesis' that drives an experimental design establish a limited repertoire of experimental designs but, with the HDM, furnish a logic by which

these experiments and their results can be evaluated. The problem for contemporary biological/biomedical research is that the repertoire of experiments under the HDM does not match the repertoire of experiments actually performed nowadays. As I have shown, there are important and highly productive varieties of biological/biomedical research which are not hypothesis-driven. Accordingly, the issue of what constitutes an experiment needs closer examination, the topic which I will consider in the next chapter. The rationale for the OES is to provide an epistemological structure, doing similar work to that of the HDM, for at least some of these experiments, namely, those which are system-driven.

CHAPTER 5 EXPERIMENT

A possible difference between my approach to a scientific epistemology and that of Hempel or of the Logical Positivists as a group is that I am bringing practical experience to the issues under discussion. My attitude toward the experiments in biological/biomedical research has more in common with cooking than with solving a proof in geometry. While this may be a personal attitude, it has much in common with contemporary philosophers of science, known as the New Experimentalists, who regard the nature of experimentation as a central issue in philosophy of science. I recognize that the philosophical course has to be plotted between the Scylla of too much hands-on practicality and the Charybdis of a desiccated overly-theoretical appraisal. Furthermore, 'experiment' is another of those confusing terms which means different things to different scientists. My notion of what an experiment is may be somewhat broader than that of some of my colleagues. Examining what constitutes an experiment will be motivate the analysis of what constitutes a successful or effective experiment in system-driven research. However, it will also permit a move from specific sorts of 'experiment(s)' to 'experimentation' where the latter deals with how we organize or design experiments, in other words, with the epistemological structures which may be utilized when we design and perform biological/biomedical research.

5.1 Experimental method—Whewell's imprint

Just as Whewell made important contributions to a modern conception of what a hypothesis is, he also contributed to the development of the modern concept of experimentation. Regrettably, Whewell's contribution is not at all straightforward, though highly relevant to the discussion of system-driven research. The relevance of Whewell's work requires some re-analysis because there is some slippage in both the terminology and the concepts he employs. For example, it is not entirely obvious that the meaning of 'induction' is constant and unchanging as it transits from Bacon to Whewell to the late 19th century. Moreover, our own intellectual vantage-point about the nature of experimentation has shifted. For example, if we are no longer in the thrall

of the context of discovery/context of justification dichotomy, our interpretation of Whewell's contributions might differ from those previously held. Thus Whewell may have observations highly relevant to OES but overlooked by a mindset dominated by the HDM.

Whewell's main insight with respect to the nature of experimentation was to operationalize a basic approach to experimentation articulated by Bacon, namely, that we need to go back to data and not formulate theories from geometry-like first principles. Whewell makes a clear differentiation between the methods of geometry and those of science based on actual experience. This may not sound so terribly revolutionary now, and I suspect it was not entirely revolutionary in the early 19th century, since Linnaeus clearly had been gathering data on plants in order to formulate his classification and the elder Herschels had gathered large quantities of data about comets. It would be inaccurate to limit Whewell's concept of research to mere enumerative data collection. His scientific methodology focused on how data sets could play off each other and generate broad concepts about the problem under investigation (McOuat, 2009, 219-20). (Figuratively or actually, the thing generated might be a picture.) Nevertheless, we might appreciate Whewell's contribution as helping to establish the importance of seeking/acquiring sound data or evidence as a hallmark of well-designed research. In this *Theory of Scientific Method*, Whewell writes that "experiment can only show what is, not what must be" (Whewell, 1968, 64), a comment which provides a basic insight into his view of experimentation. Writing mainly about physical sciences, he distinguishes between observation and experiment (the latter apparently must involve some perturbation of the entity being studied), but they seem to be equipotent methodologically.[82] Thus Whewell's main contribution to the modern concept of experimentation focuses on interventional experimentation, which we associate with HDM. However, he does not dismiss the value of what we might call

[82] I make this interpretation from where he writes: "Perception of external objects and experience, experiment and observation are needed, not only, as we have said, to supply the objective element of knowledge—to embody, limit, define and modify our ideas; but this intercourse with objects is also requisite to unfold and fix our ideas themselves" (Whewell, 1968, 70).

"observational experimentation". Thus, even at the inception of what we might regard as modern scientific method, we find someone as insightful as Whewell balancing two styles of experimentation, one involving perturbation of the thing being investigated and the other involving detailed observation (or, as I will develop subsequently, pattern-finding).

5.2 Experiment: what is an experiment in biological research?

Insofar as it has considered what an experiment is, up until quite recently Philosophy of Science has promulgated the critical orthodoxy that experimentation is somehow tied to a hypothesis. Thus we might think of Ben Franklin with his kite and iron key in a Pennsylvanian thunderstorm carrying out an experiment to determine whether lightning really was an electrical discharge or not, or Lavoisier showing that combustion actually made a substance heavier and thus defeating the hypothesis that phlogiston was lost from the substance, or Faraday carrying out any number of experiments to clarify the relationship between electricity and magnetic forces. Similarly, in biological research that is actually organized around a hypothesis, an experiment is designed to show whether that hypothesis is apparently true or definitely false: the former reflects abductive inference and the latter power of (negative) deductive inference. Hacking's analysis of experimental science has shattered this narrow view of what an experiment is. Nowadays most philosophers of science allow, following Hacking, that "experimentation has a life of its own" (Hacking, 1983, 150). It has become acceptable, if not downright reasonable, to examine exactly what an experiment is as such[83] and what myriad of forms it might take. The very notion that experimentation can be investigated on its own permits, or even facilitates, my inquiry as to forms of experimentation which function without being driven by a hypothesis.[84]

[83] Interestingly, the scientific community, including biological/biomedical researchers, has tended to lag behind the philosophical community and perpetuate the more traditional attitude that any proper experiment is intimately tied to a hypothesis which it is designed to test.

[84] Rheinberger takes this analysis in a somewhat different direction from mine. Refreshingly, rather than just quoting Hacking's view about experimentation, he demands more of it: "If experiments are said to have a life of their own, both components of the statement have to be explained: what it means for a system of practices to possess a 'life' and what it means for such a system to have its 'own' life. If we are

However, my inquiry moves the focus of analysis from what sorts of experiments get done in laboratories, and other research spaces, to the experimentation as such by which I denote the epistemological structure of these research efforts.

If we accept that there is a multitude of research strategies within biological/biomedical research, then it would be worthwhile to seek a conception of experiment which could relate to them all. Moreover, it is already mooted in the literature that there are different kinds of experiment besides the 'ordinary' straightforward kind, namely *exploratory* experiments and *pilot* experiments, as well as *natural* experiments, to name a few. These need to be sorted out from the 'experiment' under discussion here.

5.2.1 Standard concept of experiment

The orthodox 'proper' experiment has certain characteristics. Its design is well-thought-out and practicable; it is highly-structured, specifically with all the relevant controls in place; its results are likely to be clear-cut and their interpretation unambiguous. Clearly this describes a very circumspect experiment tied closely to an equally well-articulated hypothesis, indeed a rather narrow hypothesis. These features of a well-designed experiment follow directly from Hempel's description of hypothesis and the HDM, and they are more specifically worked out by Mayo. They also relate to the sorts of experiments which would be utilized to investigate what I call proximate hypotheses, which operate in hypothesis-driven research. Perhaps the most definitive experiments are those described as "crucial experiments", a concept (*instantia crucis*) of Francis Bacon's, but important for Robert Hooke and Isaac Newton (Vickers, 1992, 511). As clearly depicted by Ronald Giere, certain circumstances need to hold in order for a crucial experiment to be even envisioned: only two possible but mutually-exclusive hypotheses[85] are in play and an experiment could be performed where the data would

to take Hacking's notion seriously, we have to explore in what specific sense scientific practice engenders bits of knowledge" (Rheinberger, 1997, 139). (Note that Rheinberger uses 'system' differently from how I do.) Rheinberger backs off from elaborating on these issues. In the long run, his injunction about what we need to explore approximates to the general thrust of this dissertation, namely how the practice of system-driven research engenders knowledge.

[85] It may be argued that the Meselson-Stahl experiment regarding the mechanism of replication of DNA is a crucial experiment in biology (Franklin, 2012, Section 2.52). This experiment has three, not two,

be so distinct as to line up cleanly with one explanation or the other (Giere et al., 2006, 41). In other words, such experiments are elegant in design and attempt to be conclusive in their result. The important word here is 'attempt' since just how successful a crucial experiment is may depend on the prejudices of its critical evaluator. The literature abounds with experiments whose success as crucial experiments is debated. Perhaps the most famous of these contentious experiments in recent times is Eddington's 1919 experiment to demonstrate that the theory of general relativity was correct (Coles, 2001, 21; Hudson, 2003, 112-8).

5.2.2 'New Experimentalists' contribute—but not applied to biology

Ideally the work of those philosophers of science who focus on the problem of experimentation (the 'New Experimentalists') should be both relevant and illuminating to this discussion. The problem with their work is that they are much more interested in research in physics and possibly chemistry, than in biological/biomedical research.[86] Therefore it is difficult to find immediate relevance to the problem at hand. However, according to Allan Franklin, the sorts of information required to assess an experiment include the details of how the experiment was designed and set up, details about the statistics employed and the general analysis of the data, information on data excluded and why, information on the actual authors and the actual apparatus employed. It can be helpful to have a schematic drawing of the apparatus (Franklin, 2013, 3-8). Given the technological intensity of omics research, some of this resonates with requirements for assessing and communicating experimental data, for example, with proteomics. Some of this information is regarded as "metadata" specifying technological details. Review papers on proteomics are replete with cartoons, as they are called, of the apparatus in

candidate mechanisms, but they are mutually exclusive and clearly distinguished in the elegant experimental design.

[86] Marcel Weber attempts to summarize what he regards as the main trends of their thought without asserting that these trends are related to biological/biomedical research. His take on the New Experimentalists is as follows: (1) experimentation has objectives apart from testing high-level theories and specifically an important aspect of experimentation is to be explorative and thus discover new phenomena or regularities; (2) experimentation is not necessarily theory-laden; (3) experimentation involves intervening on the system whereas observation does not; (4) theories and experiments co-evolve so that they are geared to each other; (5) details of experimental practice are critically important for sorting out issues relating to scientific reasoning and theory-testing (Weber, 2005, 128-30).

order to make the experimental design and work-flow as clear as possible. Biological/biomedical research projects may also feature time-lines indicating the how a series of related experiments is envisioned to develop or play out (or indeed how the series of experiments did play out).

5.2.3 More recent conceptions of 'experiment'—some applied to biology

Given the excitement and momentum of the 'New Experimentalists', the relative silence among philosophers of biology regarding the nature of experiments in biology/biomedical sciences is somewhat surprising. I make this statement having consigned the issue of whether the microscope is an extension of the investigator's perceptual faculties, which is extensively discussed in Chapter 11 of Ian Hacking's *Representing and Intervening*, to the category of no longer being an issue. One reason for the silence may have been the concentrated focus of philosophers of biology on evolutionary biology, with its attendant problems of speciation and adaptation. Clearly these were big issues holding the attention of biologists in the 1950s onward. Another reason may have been that experimental biology was itself in the throes of technological transition even before the emerging field of molecular biology wrought its revolution on biological experimental design. Discussing the pioneering work of Jean Brachet regarding nucleic acids, Richard Burian points out that the nature of experiments in biology has been largely overlooked, except by Hans-Jörg Rheinberger in a series of papers and subsequently in his 1997 book *Toward a History of Epistemic Things* (Burian, 1997b, 28). These and Marcel Weber's *Philosophy of Experimental Biology* are the main philosophical discussions of experiment as such in biology. Yet what is interesting about the Burian and Rheinberger analyses is that the investigators discussed are ostensibly doing biochemistry. Figuring out the distribution of nucleic acids in sea urchin eggs or the process by which proteins are produced in a cell, by using a cell-free (microsome) preparation, might seem somewhat remote from biology. Rheinberger expressly makes the point that collaboration between classically-trained biochemists and the emerging 'molecular biologists' in the 1960s involved a clash of intellectual cultures (Rheinberger, 1997, 161). At the same time, we see the turn toward calling the domain of biochemists "physiological chemistry" wherein 'physiological'

invokes a broader application, specifically, an interest in mechanism, necessarily located within the intact functioning of the organism. This focus becomes evident scientifically with the delineation of metabolic pathways—what enzymes are in a specific pathway, their kinetics and regulation—thus producing a complex roadmap of adjacent and/or interlocking metabolic pathways in any given cell. Ulrich Krohs describes such research in some detail, with a note of genuine nostalgia for this type of painstaking, step-by-step, hypothesis-driven research (Krohs, 2012, 53-4). Generalized to transport and signaling pathways, this biological research leads to the sort of philosophical analysis found in Bechtel and Richardson's *Discovering Complexity* (Bechtel and Richardson, 2010, 17). However, if we back up and ask what biologists were doing in the mid-20th century, then the response includes, apart from evolutionary biology and systematics, major efforts in microbiology, virology, and embryology. Virology, for example, was immediately relevant to emerging and potentially soluble research problems in the nascent phase of molecular biology. Also in the mid-20th century, fundamental work relating to human physiology and pathophysiology was being performed: in terms of large organ function, endocrine function, and immunology. Much of this work involved detailed anatomical investigation; experiments to figure out actual mechanisms of such functions seem rather primitive by today's lights. One point is that it would have been difficult to generalize to a canonical philosophical description of experiment in biology given the diversity of experimental style. Here, once again, biology differs from the physical sciences. A second point is that experimental method in biology/biomedical science was itself in flux, and thus unsuited for serving effectively as the substrate for philosophical analysis and discussion.[87]

[87] Another contribution within philosophy of science to the philosophical analysis of biological/biomedical experimentation is philosophical analysis eventuating in a social epistemology of science. Helen Longino's and Bruno Latour's examinations of how scientific research is actually done figure prominently here. These analyses feature biological research but are not limited to it. Although I recognize their importance and relevance, I regard a detailed analysis of their work as being outside the scope of my discussion. In this estimate, I echo what Rheinberger writes about much the same issue: "Let me anticipate the objection that the institutional setting, the societal interests, the political power game involved in make science in post-World War II America, and above all the actors have not been given enough voice. To this I answer that my purpose has been different: I have tried to convey a sense of how experimental systems become articulated, boom, and come of age, how they eventually shape a whole laboratory culture; and

5.2.3.1 Burian's early contribution

Richard Burian's 1997 paper about Jean Brachet's research on nucleic acids stands out as an early contribution to a philosophy of biological experimentation. Burian immediately appreciates the expert audacity of Brachet's approach. The biological problem was definitely important: as Burian notes, "problems surrounding the localization and functions of nucleic acids [DNA and RNA] served as a common technical entry into nucleo-cytoplasmic relations, embryogenesis, chemical embryology and biochemistry of development, and molecular genetics while allowing him to remain faithful to embryological and developmental issues that could not be addressed or resolved by use of procaryotes [sic]" (Burian, 1997b, 28), although there is a touch of anachronism here in relation to 'molecular genetics'. Moreover, Brachet's approach was to choose an important problem already characterized by controversy and then solve controversy. Brachet's method was to develops tools for solving these problems as he proceeded, and Brachet's own description of this strategy, cited by Burian, was to develop and refine and redeploy techniques as he went along and, importantly, to understand the limitations of these techniques and validate the findings by cross-checking the results from one methodology against relevant results from other methodologies. This experimental strategy produced new, and highly reliable, knowledge (Burian, 1997b, 33). Burian calls this kind of research "exploratory experimentation" (Burian, 1997b, 28), and in some respects what he was thus describing was to 'throw the (methodological) book' at the problem. What interests Burian is that somewhere along the line in the plethora of experiments and their data going in numerous directions, "one group of findings became more and more integrated in the work" (Burian, 1997b, 34).

'Exploratory experimentation' here (possibly the first published use of this term by Burian) characterizes a general way of doing biological research. It does not touch on

how they become, remain, and finally cease to be generators of epistemic novelty" (Rheinberger, 1997, 229). Although Rheinberger and I attach different meanings to 'system', our focus on developing a scientific epistemology for biology/biomedical research based directly upon laboratory practice/experimental design is very similar.

issues of whether the research is hypothesis-driven or not. Some of what Brachet was doing would fit into my category of technique-driven research; much of it undoubtedly was hypothesis-driven. My point is merely that this is a typical strategy for mid-20th century biological/biomedical research: its critical feature is that it starts with an important problem and then sets about solving it. We see the same situation with Sydney Brenner's setting out to solve key genetic problems in *C. elegans*, a rapidly reproducing 1-mm long worm, but first he had to figure out (and figure out how to investigate) his innovative model organism. I believe we would see the same general character in the immunological research of Peter Medawar and other pioneering immunologists. The OES differs from Brachet's research strategy, and yet there are some intimations of the OES in Brachet's approach, namely the cross-checking (compare with contextualization) and the consolidation of data into a central finding (compare with pattern detection). I will discuss in the reformulation of the meaning of 'exploratory experimentation' in greater detail in a later section (Section 5.2.5.2), and I will argue that the term got co-opted from its original (1997) intentions to a more contemporary application, which in fact proves to be conceptually counter-productive.

5.2.3.2 Rheinberger's view of emerging modern biology

Hans-Jörg Rheinberger examines the work of Paul Zamecnik at the Huntington Laboratory of Harvard University in the 1950s-1960s in order to develop his scientific epistemology for biological/biomedical research. It is, above all else, a naturalized philosophy of biological research. Zamecnik and his colleagues were ultimately responsible for discovering tRNA: this discovery came through his work on the mechanism of protein synthesis. Rheinberger's painstaking historical analysis overturns certain congenial myths which have crept into the history of molecular biology—fascinating but not germane to the issues under consideration here. Rheinberger's philosophical analysis focuses on two main features: one is a specific experimental system (namely, a cell-free system for *in vitro* protein synthesis) and its vicissitudes, and the other is his more enduring concept of 'epistemic things' (put bluntly) or 'epistemic

objects' (put somewhat more technically).[88] The diction is purposely vague because, as Rheinberger is quick to point out, "epistemic objects … present themselves in a characteristic, irreducible vagueness [which is] inevitable because, paradoxically, epistemic things embody what one does not yet know" (Rheinberger, 1997, 28). Thus I think we can regard the epistemic object as what the biological/biomedical researcher "is going after", something which inevitably changes as the researcher closes in on this conceptual quarry. Burian explicates Rheinberger's concept of epistemic object as what straddles the cusp between the material practices/physical systems of the experiments themselves and the linguistic /interpretative practices which permit interpretation of the experimental results and communication with other researchers about those results (Burian, 1997a, S286-7).

In thus formulating a focus for a scientific epistemology of biological/biomedical research, Rheinberger takes as its starting point an attitude which I regard as very typical of the time and still current today, though perhaps no longer the only way in how biological/biomedical research is organized. Namely, Zamecnik was, first and foremost, investigating a problem. Ostensibly this problem was the malignant growth of cells—cancer—but the focus settled back to normal cellular growth and thence to protein synthesis which appeared, reasonably enough, to be critical to cellular growth. So the problem he really addressed was to figure out how protein synthesis occurs in cells. For this purpose he developed an *in vitro* system consisting of microsomes and cytosol from rat liver, essentially the same type of cell preparation we used for examining the Cu-metalloproteome in HepG2 cells. Later, as the research program progressed, this experimental system was replaced by a cell-free *in vitro* system from *E. coli*. Developing and refining the *in vitro* system and thus appreciating how it can be informative and where it is limited with respect to protein synthesis becomes the practical instantiation of the problem under investigation and the path to its solution. Whether or not

[88] I suspect that 'epistemic thing' is a direct translation of 'epistemisches Ding', and while the translation is not euphonious, it maintains the European acceptance of the word 'Ding' in philosophical language. (This surmise is supported by the name of his book in German: *Experimentalsysteme und epistemische Dinge. Eine Geschichte der Proteinsynthese im Reagenzglas*.)

Zamecnik was adopting an epistemology of simplification in making the cell-free *in vitro* system his experimental strategy seems to be a debate within the province of philosophy (Rheinberger, 1997, 57): it seems fairly obvious that he was making a practical choice because the *in vitro* perfused liver slice system was not working. Anxiety about introduced artifacts led to a practice of checking the *in vitro* results against findings in *in vivo* situations (Rheinberger, 1997, 65), a kind of contextualization. Rheinberger holds that an experimental system serves to establish "a space of representation … for engendering things that otherwise cannot be grasped as objects of epistemic action" (Rheinberger, 1997, 108). Whatever problems arise from Rheinberger's describing the function of an experimental system this way, the point I find interesting is that this 'grasping epistemically' inevitably involves depiction. It may be depiction in the raw data of the investigative procedure (the microscopic image; the display of number of radioactive counts in each fraction of density gradient; bands on a gel) or it may be the flow chart summarizing the experimental design or it may be the researcher's doodles as s/he tries to figure out how all these findings make a coherent explanation of what is being investigated.

Moreover, a productive experimental system must be somewhat open-ended and capacious (Rheinberger, 1997, 74), as Rheinberger points out, and furthermore "an experimental system that gradually acquires contours, creates resonances between different representations, and conveys manageable meanings to stabilized signals, must create at the same time a space for the emergence of things unheard of" (Rheinberger, 1997, 80). (Here 'emergence' has no pretensions to the technical philosophical meaning of "emergence" in general epistemology.) Thus Rheinberger alludes to novelty as being one of the outputs of effective biological/biomedical research. Finally Rheinberger names some epistemic actions which are part of this methodology of experiment based on his 'experimental system': conjunction, hybridization, and bifurcation. These are processes by which various researchers confront and modify and thus develop the experimental system of interest. I mention them mainly for the sake of completeness

since they are important ramparts of his scientific epistemology, though perhaps less immediately critical to my purposes.

It is really important to point out that in general Rheinberger uses the notion of 'system' for his scientific epistemology differently from how I am using it in mine. His experimental system is more nearly an experimental set-up; I tend to regard 'system' as a physiological unit. Zamecnik and his colleagues invested much time and energy into getting their experimental system established and then developing techniques for investigating it. Switching to a cell-free system derived from *E. coli* tweaks the experimental system but does not change it radically, whereas switching to a whole-organ *in vitro* perfusion system would be a radical departure, a brand new experimental system. When we said that we wanted to understand how copper is handled in hepatocytes, then the copper-handling apparatus of the hepatocyte was our 'system'. Studies in cell preparations from hepatocyte cell lines or from whole liver tissue (in either case, either human or mouse) were envisioned as complementary. Each permitted the investigation of the hepatocellular copper-handling apparatus, although subtle, possibly informative differences might be found. It is possible to imagine Zamecnik investigating his rat liver microsome preparation as a system in my sense. He might have subjected it any number of proteomics, or other omics, experiments. This would be a very different line of experimentation, and indeed it might not have addressed the problem he was really trying to solve, that is, how proteins are synthesized. However, it might have addressed the problem of "I want to know what makes this microsomal system tick"! That is to say, a system-driven approach might have elucidated how the *in vitro* cell-free system functions as a whole.

Rheinberger engages with numerous important aspects of biological/biomedical research in developing a naturalized philosophy of experimentation based on the recorded experience of dissecting out the mechanism of protein synthesis in a cell-free *in vitro* system. His notion of experiment seems to be whatever it takes to maximize what can be learned from the system. He is not picky about whether this involves hypotheses or not. He suggests it involves a kind of domestication, that is to say, living

with the system, familiarity. He points out that "once the system has become familiar to those who 'inhabit' it, its own momentum may take over. … The more the experimenter learns to manipulate the system, the better the system comes to realize its intrinsic capacities: it starts to manipulate the researcher and to lead him or her in unforeseen directions" (Rheinberger, 1997, 67). This may be Rheinberger's own reply to his challenge to say what is really meant by attributing to experimentation a life of its own. In contrast, I am unconvinced that Rheinberger makes a solid commitment to decomposition as a key feature of experimental design in biological/biomedical research. He speaks of a patchwork view of research (Rheinberger, 1997, 227), and he touches repeatedly on the tension between simplicity and complexity in this research. However, his foundational case study is less of a campaign for methodological decomposition in experimental design than it is for practicality and traction in biological/biomedical research: given a worthy problem to investigate, you attack it as best you can. This may involve breaking it down into smaller problems, or not; it may require a simplified system even at the risk of creating artifacts, or not.

5.2.3.3 Weber's assessment of experiment in contemporary biology

In his 2005 book *Philosophy of Experimental Biology*, Marcel Weber broadly supports the naturalized philosophical strategy of looking at how evidence is produced in the laboratory in order to understand the reasoning of experimental biologist (Weber, 2005, 130). He reviews Rheinberger's conceptual focus on experimental system(s) and adds his own implementation of such ideas by examining research relating to the oxidative phosphorylation (OX-PHOS) chain of enzymes in mitochondria (Weber, 2005, 136-43). What is interesting is that this second, independent analysis works pretty well. Without detailing the analysis, I will merely summarize the features pertaining to the epistemological assessment. How cells manufacture energy was recognized as an important problem, and experimental methods for isolating mitochondria were already in development in a broadly transdisciplinary scientific community. Technological advances were important. The development and refinement of the experimental system did not operate off any particular theory, and it also operated impervious to attendant controversy or consensus. This Rheinberger experimental system evolved: in general, it

became progressively simpler, moving from isolated mitochondria to sub-mitochondrial particles. Weber concludes: “In summary, the ox-phos [sic] case can be viewed as a history of successive branching, hybdridization and other modifications of an experimental system that originated in cancer research in the interwar years. Isolated mitochondria became ‘epistemic things’ because of the experimental possibilities they offered for differentiating the thousands of metabolic processes that occur in living cells” (Weber, 2005, 143). I differ from Weber somewhat in this last pronouncement: mitochondria became epistemic things because everyone suddenly appreciated that mitochondria were involved in many important cellular metabolic process and going after an understanding of this sector of metabolism was extremely valuable, and practicable.

Weber sees value in Rheinberger’s experimental systems idea as it relates to experimentation in biological/biomedical research, but he thinks it ignores the need for norms (Weber, 2005, 149). This objection seems to turn on whether basing an epistemology of biological/biomedical experimentation on the nature of the experimental system itself introduces an intolerable degree of relativism, refractory to epistemic norms. (The countervailing argument is that it involves a constructive pluralism, which mirrors how real-life experimentation is actually done and additionally is susceptible to normative evaluation.) Weber argues that methodology matters and it involves a normative dimension. The purpose for which an experimental system is developed/refined and also shared with other researchers, who will also develop/refine it, is production of knowledge. An experimental system which fails to yield reproducible results or generates too much artifact is dysfunctional and falls into disuse. It gets abandoned. Weber declares that “the good functioning of an experimental system to a large extent consist [sic] in the *reliability* of the knowledge it produces” (Weber, 2005, 147, original emphasis), where reliability is a methodological concept. Reliability inevitably involves sorting out wheat and chaff in the findings from an experiment: the normative component is the proper assignment of error. Rheinberger seems to want to maintain some generative vagueness (see his comment quoted in Section 5.2.3.2 about

the inherent vagueness of epistemic objects), somewhat at variance with this process of harrowing the data. With the OES, I will take Weber's role for reliability further: the experimental design must be sound, and the instrumentation must be appropriate and technically superb, and then there is yet a further normative component which involves fitting the data back into the system under study. This final additional harrowing of the data is critical for producing new scientific knowledge. Likewise, Weber takes Rheinberger to task for not presenting an epistemic mechanism for the formation of consensus among scientists. These issues will resurface in my discussion of the OES: sources of error in experimental design and instrumentation, the problems of sharing data, and mechanics of pattern detection as a kind of consensus-building on a small scale. Weber does not really touch on these problems in connection with omics, mainly because omics is not what he is talking about.

5.2.3.4 Model organisms in biological/biomedical research

As might be appreciated from this discussion of experimental biology and other biological background information I have presented (for example, see Section 2.4.3), model organisms play an important role in biological/biomedical research. It is somewhat problematic that the word 'model' has so many different applications within biology and philosophy of biology. For example, 'model' can refer to the model of a biomolecule, literally with plastic rods resembling tinker-toy or with space-filling plastic balls or in stylized pictures, or to the mathematic model of a biological system as the output of system biology. In the context of experimentation, model organisms are a special kind of model, and they play an important role in contemporary biological/biomedical research. Some of the diverse model organisms have received considerable attention in the Philosophy of Biology literature. Marcel Weber discusses in detail how *Drosophila melanogaster* started out as the organism of choice for pioneering genetic analysis and ended up as a model organism for molecular biology (Weber, 2005, 154-87). The minute worm *C. elegans* is an intact natural model, developed by Brenner for studying genetic problems and then 'borrowed' into research relating to cell physiology and whole organ function (García-Sancho, 2012, 18-22). The plant *A. thaliana* has served similar research purposes in botany (Leonelli, 2008, 510-1;

Leonelli and Ankeny, 2012, 31). Spontaneously-occurring animal models of human diseases have been identified; additionally, there are genetically-modified ("transgenic") mice. The toxic milk mouse and the Atp7b-knockout mouse are examples, respectively, for Wilson disease. Gene expression can be largely inhibited, though not extinguished, in fish, and the zebrafish has proven to be an attractive model organism for creating models of developmental defects. Philosophical issues relating to animal models for biological/biomedical research are well-described in the cited literature, and an extensive discussion here is beyond the scope of this chapter.

5.2.4 Needed for biological/biomedical research: a more comprehensive concept of experiment

My expanded taxonomy of research styles in biological/biomedical research demands a more expansive concept of what an experiment is. I will argue that we need reclaim the former broader scope of experimentation to include both the notion of testing, which is the priority interpretation currently *and* the notion of observing closely, which was perhaps more widely acceptable in the early to mid-19th century. Secondly, I want to anchor this concept on what contemporary biological/biomedical researchers actually do when they portray their experimental findings in order to keep my conception of experiment as current as possible.

For the purpose of devising an enlarged notion of 'experiment' more suited to diversity of research strategies in biological research, it helps to look back at previous meaning of the word in European common parlance. The Latin root for 'experiment' is the verb *experiri*, which means to test or try. Both 'experience' and 'experiment' are nouns derived from that verb. English usage circa the 14th to 16th centuries appears to have included both "knowledge resulting from observation" for 'experience' and "action undertaken to discover or test something" for 'experiment' (Onions, 1966, 337). So here we find linguistically an early collusion—some will say a confusion—of experiment and experience which characterizes much of the linguistic baggage of our contemporary notion of what an experiment is. In modern French the word for 'experiment' remains

'expérience'.[89] Interestingly a relevant German word, Prüfung,[90] which does not have a Latin derivation, carries a similar array of meanings[91], extending from 'test' to 'inspection' and 'consideration', as well as examination. So the point is that we could readily conceive of 'experiment' as including an observational, or even a descriptive, version as well as the structured-test version.

Like the classic notion of experiment, this conception of experiment includes the characteristics of being well-thought-out and feasible. Importantly, an experiment is something one actually performs physically in real time. This distinguishes a biological/biomedical experiment from a thought experiment. However, it need not be physically "running a gel" or plating cells from a continuous cell line into 50 P-100 Petri dishes with a view to subjecting half of those plates to an intervention being studied and the other half to an appropriate control treatment. It could be an *in silico* study, carried out on a computer, where structural data you have gathered for the protein of interest are compared and fitted to the known structure of a highly homologous protein in order to produce a structure for the novel protein under investigation. This physical aspect of actually doing an experiment is important because it guards against dreaming up an experiment plus its results and it also militates against "*ex cathedra*" science which proclaims results based on dogmatic first principles but without data. I believe it was in reaction to that kind of scientific orthodoxy that Francis Bacon came to insist on the importance of experiment. The word 'experimental' came to mean "based on experience" in the late 16th century.[92] Moreover, I want to reclaim the dual connotations of 'test' and 'observe/inspect/consider'. Given my taxonomy, both pertain

[89] Apparently, by the Renaissance 'expérience' had acquired connotations of observation as the source of knowledge. This information is from http://www.etymonline.com (the *Online Etymology Dictionary*, accessed 8 May 2013) (Harper), but I am unable to confirm this interesting assertion in the published reference literature.

[90] According to the *Oxford Universal Dictionary*, Prüfung is not the source of the English word 'proof' which comes from the Latin *proba*, via French (*prueve*) (Onions, 1955, 1598).

[91] These include examination (exam), test, study, scrutiny, consideration, contemplation, inspection, trial, proof (Springer, 1975, 1222). Prüfung is not the only term for 'experiment' in German.

[92] This etymological information is also from http://www.etymonline.com (accessed 8 May 2013) (Harper). No connection is made to Bacon.

to what an experiment is. It could be a test, but it could be a coherently organized set of observations, namely a deliberate, well-designed investigation.[93]

By way of summing up this discussion, I want to maintain that, for biological/biomedical research, to experiment embraces both the notion of testing and observing closely. An argument in favour of this conception of experiment is that the latter notion (observing closely and coherently) sometimes characterizes experiments done in hypothesis-driven research. My example is "knockout mice", as they are generally called. A knockout mouse is a mouse genetically engineered to have a specific gene inactivated, or colloquially, knocked out. In that case, the gene product is not made. This is the sort of experimental design utilized to investigate the following kinds of hypothesis: "if gene *Q* is not present, then the mouse will have QQ disease" or "if gene *W* is inactivated and thus gene product w is not present, the mouse will not be able to do ww, some specific cell/tissue function." This is clearly a hypothesis-driven experiment. The knockout mouse is created, and then it is observed closely and characterized extensively. For Wilson disease, the *Atp7b*-knockout mouse has severe and rapidly progressive liver disease which is just like Wilson disease, but it does not have neurological disease (Buiakova et al., 1999, 1669). If you try to create an Atox1-knockout, the embryo is extensively and severely damaged such that nearly half of the pups die before weaning and survivors are left with numerous congenital anomalies (structural abnormalities of large organs present at birth), of the which the prominent ones do not affect the liver (Hamza et al., 2001, 6848). Early attempts to create a *CFTR*-knockout mouse resulted in a mouse with very severe intestinal disease but little respiratory disease, thus not really resembling human cystic fibrosis (Wilschanski et al., 1996, 753-4). Characterizing these knockout mice, when possible, produces interesting findings but does not necessarily support the hypothesis leading to the experiment.

[93] If we want to reinforce my description of 'experiment' with the views of an acknowledged worthy, we need only look to Francis Bacon, who negotiated the 17th-century turn to experiment in science. According to Ian Hacking, Bacon "taught that not only must we observe nature in the raw, but that we must also 'twist the lion's tail', that is, manipulate our world in order to learn its secrets" (Hacking, 1983, 149). Thus we have the dual options of well-organized observation and well-designed interventional experimentation.

These various examples of knockout mice illustrate the 'observing' type of experiment which can sometimes be found with hypothesis-driven research. Some will object that with the experiment of creating a knock-out mouse and seeing what happens, you are still just testing the hypothesis, except that the outcome data set is much more extensive and convoluted than in the more circumscribed testing experiment. My response is that this difference is the key difference, and the knockout mouse as a complex outcome of a hypothesis-driven experiment has much in common with an omics data set. Thus, since both testing and observing types experiments are found with hypothesis-driven research, I believe we can defend this dual nature of experiments in current biological/biomedical research. Moreover, clearly we do not have any basis for falling into the trap of aligning 'testing' experiments with hypothesis-driven research and 'observing' experiments with non-hypothesis-driven research.

In addressing the question of how to demonstrate what an experiment is in contemporary biology and biomedical research, I suggest we consider in aggregate the structure of research reports found in current biological/biomedical literature. The typical format for reporting is: background information, aim, methods including statistical methods[94], results, and finally discussion. It is generally regarded as a commonsensical logical progression from what was known to what you did to the implications of your findings—a temporal progression. Given this array of information relevant to a specific research effort involving experiment(s), I submit that we find two essential components, if we take any experiment down to its barest bones. One is methodology and the other is results. Indeed I would say that an experiment consists of methodology and results. The 'aim' may be informative, and the discussion may be insightful. Nevertheless, for describing what an experiment is and evaluating it, the key components are methods and results. It may seem surprising to include 'results'. However, some informative limiting versions of the interconnection between methods and results can be identified immediately. An experiment which produces no results at

[94] Biomedical research and some biological research (in animals) will also include bioethics information along with methods.

all is not viable as an experiment, even if its methodology is perfect. An experiment whose methodology is flawed may produce really interesting results, but their value cannot be determined and thus they are rendered useless. Thus any assessment of what an experiment is in biology must take into account this intimate relationship between methodology and results.

Accordingly, in biological or biomedical research, an experiment is an activity performed in real time which involves clearly described methodology and generates actual results. The methodology may be suitable for testing a hypothesis: typically this kind of experiment involves perturbing some aspect of the thing being studied and seeing how that changes it. Such an experiment could be designed as a functional test with a control arm and an intervention arm, or it could be an experiment where a gene is destroyed in an embryo and the resulting newborn is assessed anatomically (if the change is not deadly to the developing fetus). However, the methodology need not involve perturbation: it could be one for producing a descriptive product, such as a molecular structure or a structural gene or the characterization of a system. The methodology is depicted in sufficient detail to allow for critical assessment of its strengths and weaknesses. Where appropriate, the design of the experiment anticipates and provides for appropriate statistical assessment. The methodology may involve equipment suitable for probing entities under investigation: there is no quarrel with utilizing microscopes, X-ray crystallography, nuclear magnetic resonance as ways of 'seeing' what is under investigation[95]. If methodology is poorly described, or never described, then it is impossible to ascertain the limitations of the techniques involved. Random accumulation of data does not qualify as experiment.[96] However, we do talk about "playing with" a biological problem by way of trying to investigate or solve it experimentally. Such experiments generally involve a series of approximations to either

[95] I follow Hacking in believing that this issue has been rendered a non-problem (Hacking, 1983, 187-94).
[96] In ordinary speech we might say, for example: "This cake is an experiment." If this expression means that you substituted 1 cup of plain yogurt for 1 cup of sour cream because you did not have the latter and wanted a less fat-laden cake anyway, it is a reasonable description. If you took all the ingredients generally found in a plain cake and combined them in random measures, then it is not an experiment, and it may be a culinary disaster.

the desired method or the actual solution of the problem at hand. These experiments may start off looking almost random, but since they have a progressively tighter organization, I believe they fit into my conception of experiment. Thus, by way of summary, in biological/biomedical research an experiment is a deliberately organized examination of a biological process or entity (which may be physiological or pathological). If it is driven by a hypothesis, the hypothesis itself imposes the necessary organization of the experimental design; if it is driven by a system, as I have defined it, or by a structure, then that entity imposes the organization.

Equipped with this concept of what an experiment is in biological/biomedical research, we can go beyond the recognition that the design of such an experiment can be organized in various ways—hypothesis-driven, system-driven, or by any of the other approaches itemized in my taxonomy (Section 1.1.2: Table 1). There are also some broad types of experiment which can be organized as hypothesis-driven or system-driven or 'other' driven (see my taxonomy), and I will consider these in Section 5.2.5 as variations or subtypes of 'experiment'. We can also appreciate that singling out what a biological/biomedical experiment provides scope for figuring out what 'experimentation' is. I believe that in common modern philosophical parlance, experimentation is taken as the aggregate of doing experiments: lots of people doing lots of experiments. In my opinion, this is what Hacking means when he says that experimentation has a life of its own (Hacking, 1983, 150). Experimentation becomes the composite activity of doing experiments. Thus we might talk about proteomics experiments and proteomics experimentation where the latter denotes the general activity of doing proteomics experiments. I will employ the term 'experimentation' somewhat differently. From my point of view, experimentation is inclusive of the epistemological structure undergirding the performance of experiments. Thus the terms are not interchangeable. Experiment has to do with the ontology of experiment(s) whereas experimentation has to do with scientific epistemology.

5.2.5 Variations on the experimental theme

Apart from the crucial experiment, there are several other types of experiment which sometimes are specifically singled out. These include the pilot experiments,

exploratory experiments, and natural experiments. Importantly they can be found with either hypothesis-driven or system-driven research designs. I will argue that we can easily appreciate how these fit within the definition of experiment which I have offered. Thought experiments are entirely different, and I will not consider them here, although it can certainly be argued that thought experiments should achieve the same epistemic rigour as actual-in-the-world experiments.

5.2.5.1 Pilot experiments

Pilot experiments typically are a small-scale version of a larger project. The smaller scale can be reflected in choosing a smaller 'n' for the cohort size, or deliberately excluding a heterogeneous array of test subjects while recognizing that the cohort will be less representative of the population to which the research is actually relevant, or selecting a model which is simpler and easier to manipulate than the definitive model. An example of this last feature is to use a continuous cell line for investigating some aspect of cell physiology with innovative techniques instead of using primary cultures of the cells of interest. Setting up the primary culture might introduce a problematic level of complexity to the experiment, which is already ambitious due to the technical challenges of the new assays. The results in the continuous cell line may be of interest in themselves, though less informative because they are subject to the well-known limitations of that sort of cell culture model.[97] Pilot experiments are performed in real time; they involve clearly described methodology (typically straightforward and/or up-and-running in the research space) and they generate actual results. In general they are designed with the larger and more complicated experiment envisioned. If the pilot experiment proves unmanageable, then the larger experiment is likely to be scuttled. Sometimes a project never gets beyond the pilot stage. There may be problems with getting adequate funds, research subjects or reagents. Occasionally the

[97] Those limitations include that the continuous cell line is transformed into something approximating a neoplasm: for example, it may have lost certain cellular functions over time and thus not mimic the normal cell faithfully. HepG2 is an example of continuous cell line for human hepatocytes, and it is established as very similar to normal hepatocytes (thus "well-differentiated"). A primary culture is where you take the tissue of interest, separate its constitutive cells, and establish a cell culture of those cells. This sounds better than the continuous cell line, and sometimes it is. However, getting suitable tissue can be problematic, and the cells often lose functions over the first 24-48 hours in culture.

pilot project is an adequate proof of principle: more elaborate work is not required to advance the research project.

5.2.5.2 Exploratory experiments

Exploratory experiments have been described previously by philosophers of biology. In my opinion, they should be differentiated from pilot experiments. In a sense, both pilot and exploratory experiments relate to feasibility of a line of research; in my opinion, exploratory experiments operate at an even more preliminary stage than pilot experiments. The problem with exploratory experiments is knowing the connotation of 'exploratory'. Does the term merely convey the important notion of "preliminary"? Is it similar to the usage with 'exploratory surgery'—an operation performed to see whether a more extensive operation is required (or likely to be curative)? Working off the metaphor of the explorer or geographer, does it signal a descriptive mapping out of territory? Some of these connotations put exploratory experiments on the perimeter of the real research territory, or even outside it. This potential risks re-energizing the divide between context of discovery and context of justification. Part of the problem here is that the term has been co-opted into the philosophy of biology literature without much attention to the ambiguities of its day-to-day use.

Among scientists, exploratory experiments generally represent much the same thing as preliminary experiments: they are like putting out feelers to get a sense of the potential problems and pay-offs of a certain line of experiments. Such experiments resemble pilot experiments but are perhaps less focused. Exploratory experiments may include a technological repertoire still in development. With philosophers of science the situation is not so straightforward. Specifically, there is a kind of legerdemain whereby 'exploratory experiment(s)' becomes 'exploratory experimentation'. As I will argue, these are not equivalent. For Ulrich Krohs exploratory experimentation is carried out within a research programme. It is "an experimental attempt to find results that satisfy the general explanatory hypothesis of the programme, and at the same time further specify the particular hypotheses that depend on it" (Krohs, 2012, 55). Implicit in accepting this conception of exploratory experimentation is acceptance of his notions of 'general explanatory hypothesis' and 'methodological' hypothesis', neither of which in

my view (see Section 3.1) is actually a hypothesis as such. The former is a rationale and the latter might be called an analytical strategy. Therefore Krohs's view is quite specialized[98], and it is not enhanced by this terminology which further confuses what a hypothesis is.

A more mainstream assessment of exploratory experimentation is found in how Richard Burian describes it, as he draws on other recent analyses. He sums up these views that exploratory experimentation is generally "limited to situations in which experimental outcomes cannot be accurately predicted by available theories together with general background knowledge plus boundary conditions" (Burian, 2007, 287). On a narrow interpretation, this seems to translate into putting out feelers experimentally in numerous directions to see what looks as if it might be the most promising line of inquiry, as I have described the exploratory experiment from the viewpoint of the working scientist. On a broader interpretation, with respect to omics, the implication is that in a sense each channel of a high-throughput machine's automated process is in some sense an experiment. Burian and others writing about exploratory experimentation in this way are clear about its potential to "[provide] access to molecular structures and mechanisms hitherto far beyond reach" (Burian, 2007, 290). He calls this *discovery science*, not necessarily a term of his own invention, but typical of the contemporary turn to discovery.[99] Thus Burian also lapses into a subtle switch from 'experiment' to 'experimentation'. He also points out, as I touched upon in Chapter 1,

[98] He then goes on to argue that top-down systems biology is inherently incapable of doing exploratory experimentation because it lacks epistemic guidance. It merely collects data. He writes: "Exploration is a matter of allowing for the unknown, learning from error, and being ready to change perspective" (Krohs, 2012, 56). This assessment does not square with my vision of system-driven research and the OES, which I propose as a source of epistemic guidance. However, his comments are not necessarily off-target in relation to gene-chip research, the experimental design which bets fits his idea of convenience experimentation. However, his idea of convenience experimentation is closer to my idea of technique-driven research than to system-driven research, and I pointed out the perils of technique-driven research in Chapter 1.

[99] Indeed, as we will see in Section 9.2, the term "discovery science" has really caught on: it would not be an exaggeration that this is the sort of high-powered science most biological/biomedical researchers currently want to do. (Opinions confirming this attitude came up in a discussion with Dr. Stephen James of the NIH-NIDDK on 28 October 2014 and previously in discussion with Dr. Rod McInnis, Scientific Director of the Lady David Research Institute in Montreal, on 9 December 2013.)

that the interaction of hypothesis-driven research and this exploratory experimentation is really productive.

Perhaps the most extreme example of this conflation is found Lucy Franklin's recent assessment. In her 2005 paper entitled "Exploratory experiments" Franklin immediately moves the focus from 'experiment' to 'experimentation' and sets out to examine important issues germane to those I am addressing in my scientific epistemology. Franklin defines exploratory experimentation as "experimentation that is not guided by a hypothesis (or theory: I use these terms interchangeably)" and holds that "[it] has a broader and more systematic role in scientific inquiry than is commonly realized" (Franklin, 2005, 888). Further she defines what she calls 'wide instruments' those instruments "also known as 'high-throughput' … which allow the simultaneous measure of many features of an experimental system" (Franklin, 2005, 888). Thus for her 'exploratory experimentation' denotes non-hypothesis-driven research: in fact, with her allusion to high-throughput instrumentation, it apparently refers to what I call system-driven research.

Lucy Franklin and I both agree that non-hypothesis-driven research designs are increasingly important, but our philosophical analyses of non-hypothesis-driven research differ in important respects. Franklin pursues a line of analysis which puts the main emphasis on new technology, and she argues that "wide instrumentation [also known as high-throughput instrumentation] can increase the efficiency of an exploratory strategy" (Franklin, 2005, 889). Contrasting an investigations of cell cycle regulation using Northern blotting to examine a specific hypothesis about the role of histone proteins and a top-down shotgun approach utilizing a DNA microarray, she indicates that the Northern blot experiment[100] is theory-driven (or indeed, hypothesis-driven) whereas the DNA microarray experiment "did not did not begin with any particular hypothesis or set of mechanisms that were being explored" (Franklin, 2005,

[100] A Northern blot is a technique for detecting a specific mRNA in a biological sample. It was developed in the 1980s. A Southern blot is an analogous technique for detecting a specific DNA in a biological sample, and a Western blot (also known as immunoblotting) is the analogous technique for detecting a specific protein. Southern blotting was developed first and named after its inventor whose surname was 'Southern'.

892). So far, so good. In fact, Franklin adopts from scientists the notion that such non-hypothesis-driven experiments are "exploratory", perhaps a risky move on her part since, especially in the early 2000s, scientists were under pressure to fit their non-hypothesis-driven experiments into a schema favouring hypothesis-driven research. "Exploratory" and "hypothesis-generating" served this purpose. Franklin suggests that to make sense of such research designs, we need to see that some sort of background theory is "crucial to the success of exploratory experiments, even those carried out using wide experimentation" (Franklin, 2005, 893). Certainly—this sounds similar to my idea of the over-arching hypothesis. Where Franklin and I diverge is that she has accepted the ambient attitude that experimental design using wide instrumentation just is exploratory. She likens this scientific effort to mapping, an attractive but potentially problematic metaphor intrinsic to the notion of being exploratory. Interestingly, in terms of the output of the DNA microarray experiment relating to cell-cycle control, she notes that "rather than a hypothesis directing them to a particular part of the complex web of interaction between different cellular constituents, the exploratory experiment served to find interesting patterns of activity from which scientists could later generate a hypothesis" (Franklin, 2005, 894). Finally she points toward a scientific epistemology of non-hypothesis-driven research which emphasizes efficiency, as opposed to other epistemic virtues. The weakness with this general analysis is two-fold: first, it is too credulous of the status quo.[101] Although whether exploratory experimentation is actually scientific is never questioned, the danger of entrenching the context of discovery versus context of justification dichotomy is never recognized. Secondly, her analysis is derailed by focusing on DNA microarrays.

Writing somewhat later (2007), Maureen O'Malley considers exploratory experimentation further, this time in relation to metagenomics.[102] Citing and building

[101] Franklin briefly considers whether wide instrumentation could be put into the service of a hypothesis-driven experiment and suspects that it is unlikely. On the contrary, it is entirely possible with those sorts of experimental design already being utilized by 2005, namely comparative proteomics. I will discuss this issue in some detail in Chapter 9 (Section 9.4).

[102] O'Malley has already shown the tendency to take a broad view. In 2005, with John Dupré, she wrote: "From the systems-theoretic perspective, the demotion of the genome goes further. Genomes in this

upon Friedrich Steinle's 1997 paper (Steinle, 1997, S69-S70) which distinguished theory-driven experimentation from exploratory experimentation, 'exploratory experimentation' acquires the status of a mature technical term in philosophy of biology. Whereas theory-driven experimentation tests specific expectations, exploratory experimentation seeks to identify empirical regularities and generate concepts and classifications (O'Malley, 2007, 339). Metagenomics is the study of microbial DNA in a specified environment, which may be as 'local' as your own intestinal tract or as vast/remote as the ocean. It is an interesting experimental method in this context because it has well-developed methods and it leads to transcriptomics, proteomics and metabolomics studies of the same microbial object ('community'). Such research is susceptible to similar criticisms as proteomics or metalloproteomics, such as the fishing expedition label (O'Malley, 2007, 342). Much of the paper is devoted to recounting the history of discovering proteorhodopsin by metagenomics methodology, a novel but "accidental" discovery (O'Malley, 2007, 349). O'Malley's view of the status of exploratory experimentation is that instead of being contrasted to theory-driven experimentation, the two should be regarded as extremes on a continuum. This view, which is consistent with her emphasis on integration, features among other virtues the refusal to demote exploratory experimentation as something inferior to theory-driven experimentation. The concession exacted for this parity is that important epistemological features of system-driven research may not get due consideration (see my discussion in Section 3.2) One possibility is to create a different term, natural/history experimentation, to describe some of the experimental practice actual researchers in biology/biomedicine actually perform (O'Malley, 2007, 352).

These views are of interest, though we seem to have lost sight of the circumspect real-life entity, the 'exploratory experiment'. Moreover, I doubt that

framework do not explain anything: they 'merely' constitute some of the components on which higher-level system properties depend. The study of genomes becomes simply an exploratory or first-level tool for the analysis of cells and tissues. This conception of genomes and genomics makes it clear that the primary objects of system-theoretical inquiry are higher-level processes and properties" (O'Malley and Dupre, 2005, 1272). This seems to argue against the "preliminary" characteristic of exploratory experimentation.

conflating exploratory experimentation with omics is conceptually valuable. High-throughput technologies which serve a discriminative objective (that is, answering to the motivation "let's see which of these entities of interest can actually do the function which interests us") lend themselves to exploratory experimentation as described here. It seems that much of Burian's 2007 paper, as well as Franklin's 2005 paper, not unlike Krohs's work, is in the thrall of "gene-chip" DNA microarray work. The gene chip is not the best exemplar for this philosophical analysis, and we have seen that it poses problems for others as well. (Moving to the example of metagenomics enriches the discussion because metagenomics really is about interrogating a system, namely, that community of bacteria which interests us and whose particulars we can define in detail, in this case, whichever genes are present.) When we use omics to interrogate a biological system, the design of the experiment is somewhat, albeit subtly, different from utilizing a gene chip because gene-chip research can lapse into technique- or even recipe-driven research. And when we consider biological/biomedical research aimed at 'situations in which experimental outcomes cannot be accurately predicted by available theories' plus other existing resources, we might ask whether there is any other kind of biological/biomedical research worth the effort. Exploratory experiments on Burian's definition are important, mainly because they are poised to capture the interplay between hypothesis-driven and non-hypothesis-driven research. 'Exploratory experimentation' on Franklin's definition definitely refers to omics. Franklin uses the term to address differences in research styles, which she herself characterizes[103] as Popperian (hypothesis-driven) and Baconian (observation-gathering) (Franklin, 2005, 889).

With a plethora of interpretations exactly what is meant by 'exploratory experimentation' remains vague and thus prone to misinterpretation. It does not advance a clear conception of an individual exploratory experiment. It might be regarded as an experiment designed to see what is possible in a line of research—not quite the same as Burian's conception. Likewise it might be regarded as preliminary or

[103] I do not regard this appellation as appropriate (see Section 4.2.3).

hypothesis-generating or 'descriptive'. Some of these latter versions detract from generating the rich scope of system-driven research as I see it, and therefore I do not want to make 'exploratory experimentation' interchangeable with system-driven research. A key difference is that system-driven research focuses on the system, not on the technology or instrumentation. Some exploratory experimentation as described by philosophers of biology describes other sorts of non-hypothesis-driven research. For those writing between 1995-2010 exploratory experimentation also misses out on an important type of non-hypothesis-driven research, specifically, structure-driven research.

5.2.5.3 Natural experiments

Contemporary natural experiments are a kind of specialized experiment whose existence is just being noticed and whose value is debated. I mention it just to acknowledge it as an option, but it plays little role in omics. The typical natural experiment takes advantage of a change in how something is done and then compares the resulting new status with the status before the change or else with a comparable group where the change has not taken place. These experiments often have a certain time-urgency associated with them, and little time is available for developing a comprehensive theoretical superstructure. For example, the introduction of a policy mandating that all wheat flour be supplemented with folic acid creates the opportunity to examine the benefits and disabilities which might thereby accrue in the population, specifically whether folic acid sufficiency increased across the population and what diseases became less common or less severe (and what new health problems appeared). If the population had been studied before the introduction of folate-supplemented flour, then it could act as its own control. Obviously the time-line would be tight for getting the baseline data. Once the policy was introduced this one-off opportunity would be gone. It might be possible to compare the general population against subpopulations which do not consume wheat flour, for example, people with coeliac disease, but clearly that experimental design has inherent flaws. The main objection to such natural experiments is that data are being gathered without the benefit of any hypothesis or, alternatively, with only a vague idea of what is being

studied. I believe this criticism is largely without basis. The natural experiment involves an intervention on the status quo with the expectation that there will be some effect(s) of that intervention. Whether the effect(s) will be beneficial or not is not postulated. Moreover there is inherent time-pressure with this type of experiment, and this time pressure of a one-off experiment may compromise "perfect" experimental design. Thus while this sort of experiment conforms to my definition, it poses problems relating to coupling a nebulous hypothesis with an observation strategy for data collection.

Another type of natural experiment is a somewhat more conventional than the prospective one-off variety. The design of these more conventional natural experiments takes advantage of variations spontaneously occurring in nature and involves comparison of two or more such variants to investigate a biological problem. One of the more famous examples of this kind of natural experiment involves twin studies where a health-related outcome is examined in identical (monozygotic) twins and in fraternal (dizygotic) twins, in order to tease out the relative contribution of genetic make-up to the outcome. Likewise it could said that if we compared the features of Wilsonian liver disease in the two naturally-occurring mouse models for Wilson disease, which in fact have different strain backgrounds, we would performing a natural experiment in order to examine possible contribution of "non-*ATP7B* genes" (modifying genes) on the liver disease. It qualifies as this kind of natural experiment because the experimental design benefits from the spontaneous occurrence of *Atp7b* point mutations in these two different mouse strains.

5.3 How is a hypothesis-driven experiment different from a system-driven experiment?

For a naturalized scientific epistemology of system-driven experimentation, a key question is how an individual experiment driven by a hypothesis differs from one driven by a system. I wish to consider this question by way of summary for this chapter. Some differences are immediately apparent at first examination. The hypothesis-driven experiment appears much more focused than the system-driven experiment. This is largely because the purpose of the hypothesis-driven experiment is to investigate the claim of the hypothesis, which must be specific, and even rather narrow, in order to be

formulated as testable. In contrast, if the hypothesis-driven experiment feels like looking at something through a narrow aperture, the system-driven experiment feels like viewing with a wide-angle lens. In order to make these intuitions more concrete, and (ideally) more precise, it seems appropriate to consider some representative experiments.

The example of hypothesis-driven experimentation has to do with the metal-binding action of a small protein known in mammals (including humans) as ATOX1. In the mammalian liver, for example in humans and mice, copper is transported across the plasma membrane of the hepatocyte from the blood where copper is loosely bound to albumin and certain amino acids. However, once inside the cell the copper is not allowed to remain as a "naked" metallic ion. It is always bound to an intracellular molecule, and in particular to small proteins which deliver the copper to appropriate intracellular targets. Accordingly they are called copper-chaperones, or more general "metallochaperones". ATOX1 is the copper-chaperone which delivers copper to the Wilson ATPase resident in the trans-Golgi network. The primary amino acids structure of ATOX1 was established circa 1998 (Hung et al., 1998, 1750), and it was found to have one amino acid pattern 'MXCXXC' whereas the Wilson ATPase had six of these patterns in the tail at the amino end of the protein. Since this pattern also is found in various other proteins which bind copper, it is investigated as a copper-binding domain and found to be one. In order to investigate the binding and other properties of ATOX1, the cDNA[104] for ATOX1 was inserted in to *E. coli*, expressed and purified. One of the first 'proximate' hypotheses to be investigated was: in the CXXC motif, the cysteines (that is, the 'C's) are crucial. In order to test this hypothesis, a technique called site-directed mutagenesis was employed. This means that the cDNA for ATOX1 was modified in such a way that any cysteine could be predictably changed to a different amino acid, serine, which lacks the chemical properties of cysteine related to metal-binding. Two of three variants of interest could be generated: MXSXXC and MXSXXS. (The third one, MXCXXS,

[104] The convenience of the cDNA is that it represents the structural gene for the protein with all the introns removed and all the splicing done. It is not the gene for the protein, just the form for translating the protein as such.

proved to be unstable and could not be expressed.) When incubated with metals, both of these variants had very different patterns of electrophoresis in polyacrylamide gel compared to the normal “wild type” protein, consistent with a different shape to the protein, and specifically a shape indicative of failing to bind metal (Narindrasorasak et al., 2004, 112). Therefore the proximate hypothesis was effectively substantiated that the cysteines in the MXCXXC motif are crucial to metal-binding. A second proximate hypothesis was investigated utilizing the same technique. The human protein ATOX1 is structurally very similar to a protein in yeast Atx1, which serves a comparable function in *S. cerevisiae*, ordinary baker’s yeast. The speculation was that the two proteins would function similarly. However, compared to Atx1, ATOX1 has a few extra amino acids and, notably, it has an extra cysteine toward the far end of the protein. The new proximate hypothesis was that this extra cysteine might influence the metal-binding function of the ATOX1. To investigate this second proximate hypothesis, site-directed mutagenesis was performed and the extra cysteine was converted to a serine. This alteration had no effect on metal-binding. The gel electrophoresis patterns were the same in the wild-type and in this latter mutant (Narindrasorasak et al., 2004, 112). So that hypothesis was not supported. These findings generated a further proximate hypothesis that metal-binding by ATOX1 and Atx1 might be very similar because eliminating this extra cysteine made the two proteins very alike. Many of the subsequent experiments reported in this paper were devoted to elucidating this possibility. Thus we find hypothesis-driven experiments which are circumscribed, use methods appropriate just to investigating the issue at hand, and provide step-by-step progress. Experiments may be informative if the hypothesis is supported (in effect, proved) or disproved. The latter are perhaps more clear-cut, both in terms of the interpreting the results and the relevant logical structure of the relevant arguments. Either way, the experiment generates more proximate hypotheses and consequently more experiments.

With respect to system-driven research, I would like to consider two examples. One is William Whewell’s extended investigation of the tides, namely his experiments contributing to so-called ‘tidology’, the study of the laws of the tides. I choose this

example to illustrate many features of system-driven research without invoking molecular biology. However, the second example is contemporary: once again from our metalloproteomics research.

The tidology project is strikingly system-driven, even though it was performed before the age of computers and 'big data'. Several practical considerations motivated Whewell's tidology research: as an island nation, Great Britain was dependent on access by ocean and sea for all its trade, and many of its ports were estuary ports; exploration and naval defences were other considerations; finally, shipwrecks were incredibly frequent and disastrous (Snyder, 2011, 170). Moreover, there was no systematic collection of data about tides, in contrast to extensive knowledge of astronomy (Ducheyne, 2010, 28). Interestingly, Francis Bacon had called (in vain) for an international system of tidal observations (Snyder, 2011, 171). Whewell viewed the study of tides as an opportunity to apply his Baconian notions of science to a real-life problem (Snyder, 2011, 172): theories current at the time were not based on empirical data, and what data were being collected were peculiar to specific ports and not assembled in any general way. While the project was started by one of Whewell's former students, John Lubbock, who got interested because he had a financial stake in some newly opened docks on the Thames, my interpretation is that Whewell was responsible for regarding the problem as a system which needed to be investigated broadly.[105] Moreover, Whewell's approach to solving this problem is highly innovative. Lubbock had figured out that harbourmasters in London had kept detailed daily records of tides since 1805 but no tides tables had ever been constructed (Snyder, 2011, 173). He set about having such tables drawn up: perhaps an early example of what we would now call data-mining where the data were collected for some uncertain reason and then interrogated for another purpose. According to Ducheyne, Whewell's opinion was that "only careful observations could yield insight in patterns"; moreover, a great many

[105] This interpretation is based on Ducheyne's account (Ducheyne, 2010, 28-9) where he reports that Whewell felt that tidology lacked systematically arranged data (bottom, p. 28) and systematic treatment of initial conditions and/or contributing causes would permit "a true and complete synthesis of the tides" (top, p. 29).

observations were needed, as well as numerous parameters (Ducheyne, 2010, 31) in order to avoid erroneous attribution and to derive a solution, which could not be predicted at the start.

At this point Whewell's project lost any hint of being data-driven, in my sense, and became pre-eminently system-driven. He got ordinary people to do his observations for him and create uniformity in reporting; he enlisted sailors, dock masters, family, and world-wide missionaries; he convinced Francis Beaufort, the head of the Hydrographic Office to collaborate (Snyder, 2011, 174-5). This eventuated in a two-week period of reporting from the Office's in the British Isles where observations were made round the clock every fifteen minutes. The outcome of all this work was a huge data set (that is, tens of thousands of data points). After much calculation, this extensive experiment produced tables of the tides around Great Britain, but more importantly a "map showing how tides progress through the Atlantic and onto the shores, into ports, inlets, estuaries and rivers of all the major maritime nations and their colonial possessions" (Snyder, 2011, 177). My point is that Whewell's work can be seen as system-driven experimentation where the system under investigation was oceanic tides around the world. What I believe is typical of system-driven research is that it involved numerous observations which had to be gathered in a standardized fashion; the outcome was not evident at the start of the experimental process; a prominent product of the research was a picture, that is, a graphic depiction of the data, once the data were put together as a whole. This type of experiment also generated hypotheses which could be tested separately such as the notion of a point of no tide (Ducheyne, 2010, 38, footnote 63). Certainly Whewell used his huge data sets to attempt to test the prevailing theories of the day, but the experiment was designed governed by the system under investigation. In my view, it was not designed to test a theory as such. Ducheyne, however, holds that Whewell's favourite theory among the available candidate theories seems to have informed the analysis: thus it might have served as a kind of assumption or bias (Ducheyne, 2010, 38).

To return to omics of the 21st century, I will describe in greater detail one of our later proteomics experiments. We were interested in possible changes in expression of proteins in the hepatocellular *mitochondrial* Cu-metalloproteome in the affected Wilsonian tx-j mouse. We had little if any idea what changes might occur, although we expected there would be changes because several copper-associated proteins are found in mitochondria, copper-chaperones bring copper to some mitochondrial proteins, and mitochondria play an important role in both generating and neutralizing oxidative stress, a feature of Wilson disease. Therefore we proposed investigating the hepatocellular mitochondrial Cu-metalloproteome. For this experiment we utilized a newer technique, isotope coded affinity tags (ICAT) where the tx-j hepatic mitochondrial preparation and a mitochondrial preparation from the normal C3H mouse, from which tx-j is derived, were differentially labeled with radioactive isotopes, mixed together and run on through the mass spectrometer. We were able to identify five mitochondrial proteins which were elevated, on average, two times normal. They included a subunit of succinate dehydrogenase, heat shock protein 65, a subunit of cytochrome C oxidase, Mtch2 protein, and a protein related to ubiquinone. None of these was specifically among our speculated solution set; however, cytochrome C oxidase is well-known to be associated with copper. Thus it is a nearly classic system-driven experiment: we defined our system in detail (isolated mitochondria from liver; tx-j and cognate normal C3H mouse; information regarding age, diet, day-night cycle etc.), we used classic metalloproteomics methodology, only with the refinement of being quantitative, and we ended up with a novel solution set of tx-j proteins in higher (greater than 1.8 times) or lower abundance than in the normal mouse hepatic mitochondria. Some critics would say this was a pilot experiment: it needed to be repeated over a range of ages or with more mice. I would accept that appraisal; however, since the techniques for preparing isolated mitochondria, isotopically labeling proteins and performing proteomics experiments were well-established in our laboratories, this experiment was definitely not exploratory.

Given these examples, I believe that certain clear differences can be discerned between hypothesis-driven and system-driven research. One difference is scale: a focused approach as opposed to a broad-based approach. Because of having few or no preconceptions, system-driven experiments are more likely to produce novel or 'surprising' data, but this feature is not exclusive to system-driven experiments. The large scale of the experiment makes system-driven experiments more likely to produce a graphic representation of the interpreted data. Bias can creep into either variety of experiment. For hypothesis-driven research the source of bias may be in the very nature of the hypothesis and the history of hypotheses upon which it rests (Okhrulik, 1994, 34); for system-driven research this may be in the choice of parameters to describe the system or in the predilections of the researcher which influence what data s/he regards as important. As pointed out previously, generating hypotheses does not distinguish these two types of research. Hypothesis-driven research routinely generates more hypotheses. System-driven research may lead to articulation of hypotheses, but it may mandate different or refined ways of looking at the system.

5.4 Summary: contemporary needs

In this chapter I have asserted that contemporary biological/biomedical research requires a definition of 'experiment' which is inclusive of all the sorts of experiments actually found. The typical definition of 'experiment' reflects the bias of experimentation guided by the HDM, a residuum of the Logical Positivists, which has been adopted by the scientific research community. However, what we find on closer inspection is that we have hypothesis-driven research and various sorts of non-hypothesis-driven research. Accordingly I have suggested the following definition of experiment: a deliberately organized examination of a biological process or entity (which may be physiological or pathological). Examination involves testing or meticulously observing. In hypothesis-driven research the proximate hypothesis provides the 'deliberate organization'. In system-driven research, as an example of non-hypothesis-driven research, the system provides that organization. We can find various special types of experiment. These include pilot, exploratory, and natural experiments,

and either hypothesis-driven research or non-hypothesis-driven research may be found as experimental design. When we move from 'experiment' to 'experimentation', we move from an ontology of experiments to an epistemological issue. In particular, I find the notion of "exploratory experimentation" a particularly unproductive conflation of concepts.

Given the theoretical structure for a hierarchy of hypotheses in biological/biomedical research (described in Chapter 4) and an examination of the nature of experiment, which is to say real-life experiments geared to biological/biomedical problems, I next will turn this discussion to consider the particular problems of experimental design relating to proteomics, the kind of system-driven research which is the focus of this dissertation. Admittedly, proteomics is instrumentation intensive! The analysis of system-driven research turns on several important issues: one of these is the nature of the system itself and another is the technical aspects of actually doing a proteomics experiment. As I have pointed out already, my use of the concept of 'system' is different from its use in some other discussions by philosophers of biology. The technology of proteomics, though daunting, has been subjected to extensive critical evaluation. As the first step in developing a well-described scientific epistemology for omics, based on my experience with proteomics, I propose first to examine various aspects of experimental design in proteomics. These issues receive detailed consideration in the next chapter.

CHAPTER 6 EXPERIMENTAL DESIGN AND STANDARDS IN PROTEOMICS

In the previous chapter I argued that an experiment is a deliberately organized examination of a biological process or entity (which may be physiological or pathological), performed in real time, with a clearly described methodology and the generation of actual results. Experiments are often done in laboratories, but various other settings (except for "in your imagination") are possible.[106] I have suggested that system-driven experiments differ from hypothesis-driven experiments: system-driven experiments have broader scope and enjoy freedom from specific expectations relating to findings or outcome. The system-driven research of interest to my project involves research strategies, identified collectively as 'omics', all characterized by high-throughput methodology. With these high-throughput research strategies, an omics discipline addresses the biological system under investigation directly in its complexity. I hold that omics reflects a distinctive epistemological orientation. Instead of being organized around a hypothesis or an existing data set, the research project is actually organized around a system. But what do we mean by a 'system' in biological/biomedical research? After all, if it is driving the experimental design in omics, the system as such assumes great epistemic importance. Likewise, how do we actually examine a system? I will restrict my discussion here to the question of how we do proteomics. Technical issues relating to proteomics are important to this analysis. These technical aspects inform the larger problem of how system-driven experimental design relates to data evaluation.

The scientific epistemology I am developing for system-driven research emphasizes certain key aspects of biological/biomedical research employing high-throughput methodologies. Excellent experimental design and competently produced empirical data are critically important for providing reliable experimental findings. I will argue that meticulous description of the system is a crucial part of excellent

[106] Thus it could include some sort of natural space—a forest or tundra or lake—for examining a system as part of ecological research or it could be within the confines of a computer for an '*in silico*' experiment. However I am excluding '*in cerebro*' (thought) experiments.

experimental design for the system-driven experiment. Technical competence is complicated. I will take the technical aspects (here limited to proteomics) as a basis for considering how data acquisition leads to acquisition of actual knowledge in omics disciplines. Apart from the technology itself, how the technical aspects of the experiment are described is critical to data evaluation, as it relates to both bioinformatics and data-base compilation. The community of proteomics researchers has dealt with some of these issues by developing standards of technical performance. Nevertheless these important issues relating to nuts-and-bolts aspects of doing a proteomics experiment are not the whole story. I will argue that the very nature of a proteome informs what we can expect to learn by proteomics: variability in experimental results may reflect the extreme variability of biological entities as portrayed by protein expression, and not technical failings. These considerations lead to the third component of the scientific epistemology I am developing, namely, that the findings must be validated within the context of the system under investigation. Consequently the description of the system itself assumes even greater importance.

6.1 Pivotal question: what are we talking about when we say 'system'?

Talking about systems is a comparatively recent development within biological science. The two classic perspectives in biology are 'form' and 'function'. Certainly it is easy to see why the primary emphasis might be on form: form has visual immediacy. In classical epochs and especially from the Renaissance onward, biologists trying to classify living things have tended to rely primarily on features of biological form. These features might be gross—having a vertebral column or not—or esoteric, often microscopic, such as the subtle differences of insect anatomy which serve as the basis for distinguishing different species. It might be tempting to invest this classificatory strategy with a Platonic overlay (of "Forms"), but that seems unnecessary. Thus the predominance of form in biology historically seems to be merely epistemological. Form is something you can study quite easily. By contrast, function is not so manifest. First of all, there is protracted philosophical debate as to what is meant by 'function'. In the present

context, function is taken as relating to what something does, as opposed to how it looks (that is, its form). In human biology the long-established tradition is to classify the components of the human body in terms of anatomical, and by extension, physiological systems. It is easy to see how this approach might have developed from studies of the skeleton, but it seems likely to me that dividing the human body into systems originates in the 17th century with William Harvey and his explanation of the circulation of the blood.[107] The implication is that a system involves functionality.[108] A system is not merely a random collection of anatomical parts. The parts work together to perform (a) task(s).[109] In the human body, each system is made up of organs, which function together. Some, like the endocrine system, have chemical actions which extend beyond their anatomical limits whereas the anatomy and physiology of the circulatory system are closely connected.

With the advent of detailed microscopic examination of tissues, followed shortly thereafter by articulation of the cell theory in the 1840s, it must have seemed natural to extend the notion of system from the macroscopic to the microscopic. Thus, on my interpretation, an individual cell is a kind of system on a very small scale. In fact, eukaryotic cells have internal membranes which both demarcate and connect their constituent organelles, as well as structural fibres which provide shape but also can act as tracks or pathways for movement of vesicles or other components within the cell. In cells, the organelles (such as nucleus, mitochondria, and various membrane components such as the 'endoplasmic reticulum') play a role analogous to the role organs play in the larger systems of the body. Prokaryotic organisms, which are classified as either Archea or Bacteria, are single-celled organisms with a somewhat less complicated cellular

[107] Interestingly, Auffray and Noble have come independently to a similar conclusion (Auffray and Noble, 2009, 1660).

[108] With respect to my interpretation of how we think about physiological systems, I find a sympathetic, affirmative voice in Evelyn Fox Keller who elaborates her notion of a self-organizing system (an issue which I do not wish to address in this dissertation) as follows: "a system in which the entire system is shaped by the combined activities of all the individual components—activities that are generated inside individual components, with effects manifested externally to themselves, but all the while remaining inside the composite self that defines the larger system" (Keller, 2005a, 1073).

[109] Here we can see that my view has some commonality with Cummins' causal role theory of function (Amundson and Lauder, 1994, 447-8).

organization. Single-cell organisms, just as much as the eukaryotic cells making up more complex organisms, can be viewed as systems on a microscopic scale. Thus, whether we are thinking of the mammalian circulatory system or an individual cell (or a subcellular organelle or even just a set of proteins which interact to support some specific cellular action), we can regard a system as a delimited collection of components whose structures and function are somehow integrated.[110]

Several objections might arise immediately to this notion that a system is a delimited collection of entities whose structures and functions are somehow integrated. Being delimited permits the inference of an agent or circumstance that does the delimiting. Surely some, or even many, systems are very extensive. Even a cell, microscopic at one level of observation, looks extensive if viewed from a vantage-point sized to that of a molecule or virus. We risk reifying the big system into some sort of mosaic of arbitrarily devised subsystems. In fact, it seems clear that sometimes circumstances impose limits on a system and sometimes agents, including the scientific investigator, impose those limits. (Indeed with system-driven research the investigator *does* impose those limits as s/he defines the system under investigation.) The limits serve a role similar to that of an assumption, and they need to be specified in order to minimize the downside potential of such an assumption. Another problem is that 'collection' suggests the loosest of all possible associations: the groceries in my shopping-cart are a collection delimited by the wired edges of the cart, and unless they all relate uniquely to one recipe or perhaps one meal, they are integrated only by turning up on the receipt produced by the cash register. This hardly would seem to qualify as a system. Indeed it does not; however, the reason for keeping language which

[110] It is probably worth reiterating yet again that my usage of 'system' as spelled out here is quite different from Hans-Jörg Rheinberger's use of system in his terminology of the "experimental system". His example of an experimental system is a cell-free *in vitro* set-up for protein synthesis. Thus it has much to do with an experimental apparatus in the broadest sense. A Rheinberger system need not be an *in vitro* cell-free system. It could be an *in vivo* set-up where the artery to the left lobe of the liver is tied off and the one to the right lobe is not, and then you compare cellular changes in the two lobes, or it could be a whole organism such as the molecularly engineered *C. elegans* which will be described in Chapter 10. In that example, the engineered *C. elegans* is a Rheinberger system but the disease process itself is the [Roberts] 'system' under investigation by high-throughput methodology. The differences between my use and his are important but can be subtle.

is vague and expansive for describing a biological system is that it may not be entirely obvious at the very outset of an omics experiment just what is actually in the system under examination despite every attempt at a clear delineation of that system.

Finally, it may be objected that this description of system has been produced 'backwards'. The proper approach to figuring out the answer to this pivotal question of what are we talking about when we say 'system' is to determine what a system is and then demonstrate how it relates to the biological world. This is the sort of objection developed by O'Malley and Dupré, who point out that whether you are a pragmatic (not intended as a technical philosophical term: 'practical' might be a better term) systems biologist or a systems-theoretic biologist the question of *what a system is* is twinned with the question of *what biological units map onto such systems* (O'Malley and Dupre, 2005, 1273). For the pragmatic systems biologist "'system' is a convenient but vague term that covers a range of detailed interactions with specifiable functions" (O'Malley and Dupre, 2005, 1271); for the systems-theoretic biologist 'system' represents a fundamental ontological category irrespective of being natural or manmade (O'Malley and Dupre, 2005, 1271). Mapping biological phenomena onto a prefab theoretical construct often proves awkward: implementing the concept of 'species' is an example of such awkwardness. However, O'Malley and Dupré do suggest that "the project of localizing and defining biological systems can instead be developed within existing practice" (O'Malley and Dupre, 2005, 1273). This is the approach taken by my formulation of what a system is, with the added proviso that in actual utilization the investigator has the option to delimit the system as s/he wishes so long as the descriptive parameters are specified accurately and exhaustively.

6.1.1 Describing the system under examination

Typically, biological organisms are anatomically complicated and physiologically dynamic. Even bacteria, for all their apparent simplicity, display structural and functional complexities which are capable of changing over time. Proteins play critical roles in both the structure and function of cells/tissues/organs. What proteins are actually produced from a structural gene is influenced by the cellular environment (resulting in selectively

different reading of the genetic information, a process known as alternative splicing[111]) or by the extracellular environment when hormones or exogenous chemical modulate gene transcription or by temporal activators or silencers reflecting age and stage of development. In addition, proteins undergo chemical modification within cells either as part of normal cellular action or as a consequence of cellular or tissue damage. Chemical modifications related to normal cellular action are known as "post-translational modifications". Abnormal, usually destructive (that is, function-inhibiting), modifications are taken as evidence of cytotoxicity, damage to cellular constituents from toxic chemicals. One of the great assets of proteomic analysis is that the array of proteins captures a snapshot of the system that portrays the system with minute specificity. The asset is also a technical defect: the data sets are large and can be difficult to manage. There are numerous sources of potential technical error.

6.1.1.1 Choosing a system to investigate

Given these opportunities/risks associated with studying proteins, it might not be technically feasible or even particularly interesting to try to develop a proteome, for example, of a whole organism. Such faint-hearted anxieties have not inhibited the attempt. This kind of project certainly is possible for a bacterium, and it has been accomplished for yeast in 2003 (Ghaemmaghami et al., 2003, 737) and again in 2008 (de Godoy et al., 2008, 1253). It has been accomplished for the model organism *C. elegans* (Audhya and Desai, 2008, 205; Geillinger et al., 2012, 4594), and may be a realistic objective even for a paramecium, which is more complicated than *C. elegans* but whose ciliary apparatus has been studied via proteomics (Yano et al., 2013, 113). Determining the entire "human proteome" as such seems like an impossibility, both technically and theoretically. While the technical aspects are potentially solvable and actually in the process of being solved, the theoretical problems are complex.[112] Thus, in general,

[111] For example, the structural gene which encode the Wilson ATPase, known as *ATP7B*, is 'read' as producing that copper-transporting ATPase in most cells where it is expressed at all. However, in the pineal gland the gene product is a novel splice variant, called 'pineal night-specific ATPase' (PINA), which displays marked circadian rhythm variation (Borjigin et al., 1999, 1018).

[112] HUPO proposed a "Human Proteome Project" in 2008. Their general strategy is to take a 'representative protein' approach. This moots on the obvious problems inherent in human variation, reflecting both genetic and environmental variations. The proposed general approach is to identify the

proteomics experiments have tended to narrow the focus somewhat, as with the paramecium, and deal with some subset of the whole large system of interest. These include the proteome of a subcellular organelle (Au et al., 2007, 377): for example, all the proteins expressed in the lysosome. Alternatively, the subset can be delimited regionally, such as, the human plasma proteome: the proteins found in the plasma compartment in humans. In fact, this approach is not merely a matter of technical convenience. The more focused system may address pressing scientific questions of practical importance. Preparing a biological specimen which is enriched in one particular cellular component produces a sample more amenable to proteomic analysis: the proteins are in higher concentration and contamination with other proteins is reduced or eliminated or can be measured (Au et al., 2007, 380). Our approach in metalloproteomics of limiting the proteome under examination to a set of proteins described by specific functional characteristics (metal-binding or metalloproteins) is another way of narrowing the focus in order to address specific biological questions.

Given that there is a discretionary component to exactly what qualifies as the system under investigation, articulating that choice as clearly as possible generates a key set of descriptors. It is an extremely important component of the experimental design. So for example, we might specify the continuous hepatocyte line HepG2 grown to confluence, when the dish is covered with a single layer of cells all abutted to each other, in a specific medium of a given composition. Other descriptors might also be pertinent, such as the source of the cell culture material (for example, directly from the American Type Culture Collection or kindly passed along by a colleague), number of

---

major protein for each protein-coding gene, based on the findings of the Human Genome Project. The general strategy was re-assessed and described in 2011 (Legrain et al., 2011, 1 of 5). Of course, this would not be the entire proteome because it fails to take into account splice variants and post-translational modifications. It limits proteomics to a pared-down version of the original (and I believe defective) definition of proteomics: all the proteins expressed by a specific genome. In terms of an operational definition, here the human genome is restricted to genes which encode actual proteins, elsewhere estimated at 2% of the whole genome. One of the prominent theoretical problems is that it might be possible to construct ***a*** human proteome of some sort (for example, a minimal proteome), but it seems like a non-starter on theoretical considerations to construct ***the*** human proteome. It would be inaccurate, of course, to suggest that HUPO is unaware of these difficulties. There are numerous impediments to this project, including funding. If this mega-project cost only $1 billion, it would be a bargain.

subcultures prior to the proteomics experiment, freedom from subtle contamination by an infectious agent such as a *Mycoplasma* species (which could alter cellular metabolism). As already pointed out, once produced within a cell, a protein may undergo various modifications over time because of metabolic changes in the cell/tissue or in the cell/tissue's environment or because of injury or aging.[113] Changes in proteins reflect and flag what is going on in the system. Consequently the proteome can be said to be dynamic; however, the dynamic nature of proteomes leads to methodological problems. It becomes necessary to include the temporal and environmental specifications of the system among the descriptors: what age or stage of development, what environmental factors (such as availability of chemical substrates or certain nutrients, or presence of oxidative or other forms of stress). If the proteome reflects changes in the system over time, in response to developmental or environmental changes, then in effect you may only be able to determine the proteome**s** (in plural) of a system, which are related like time-slices. In a sense you end up with the opportunity to examine frame by frame changes in the system of interest. I suggest that a system such as a cell or even a subcellular organelle, for example, a mitochondrion, actually generates multiple proteomes in the course of ordinary function, and one problem is to pick which one you want to study. Once that decision is made, then the challenge is to figure out how to work with that system consistently.

Thus by way of summary, with the OES it is entirely permissible that the system under investigation can be chosen to suit the aims of the problem being investigated, but in any case the chosen system needs to be described fully. Important parameters include: the specific organelle or cell or tissue; the species; the age or stage of development; whether the biological context is normal or abnormal; and whether the environment of that system has been modified in any particular way, and if so, how it was modified. The requirement for clear and complete description figures prominently

[113] The specifics of protein expression over time give a certain advantage to proteomics over some of the other omics disciplines. At the very least, it appears to be an advantage over genomics: many of the interesting questions arising genomics have to do with the character and action of encoded proteins expressed in cells.

in our definition of metalloproteomics. To take a different example taken from the experimental record: it is one thing to define the hepatocellular Cu-metalloproteome in a normal adult male mouse maintained on a normal mouse-chow diet, but it is an entirely different matter to define that metalloproteome in a normal adult male mouse fed mouse-chow specifically prepared to be highly enriched in copper. This latter experimental protocol might produce a "copper-responsive" metalloproteome. The difficulty here is that unless the system is well-described, inconsistencies between the two proteomes might be ascribed to technical faults when in fact each proteome addresses a different system, one basal and one influenced by excessive dietary copper.

6.1.1.2 Impact of the 'system' on experimental design—practical aspects

In the long run, the features of the system actually constrain the experimental design. Describing experimental design as 'constrained' sounds limiting and problematic (and indeed it may be both), but it merely recognizes a practical feature. It does not introduce a new consideration. To say that that the system is driving experimental design is not different from saying that it constrains experimental design. When we say that a hypothesis is driving an experiment, it is also constraining the experiment.[114] With system-driven research a problem is that what we want to find out about a system may exceed what we are technically capable of finding out about it. For example, within metalloproteomics we have the problem of efficiently and inclusively identifying metal-binding proteins. However, this technical problem pervades proteomics: identifying proteins which are stuck in membranes within the cell, identifying low-abundance proteins, finding all the plasma proteins when the high-abundance ones get in the way of 'seeing' the proteins present only in very low concentrations. We might prefer smaller systems because they are easier to analyze, but investigating such systems (for example, subcellular organelles like mitochondria) is dependent on the effectiveness of

[114] In some situations the proximate hypothesis can exert a strangle-hold on the experimental design. This problematic aspect of the HDM is a kind of bias, which is sometimes characterized metaphorically as the 'straitjacket'.

our methods to isolate pure preparations of the organelle.[115] With hypothesis-driven research, we typically refine and delimit the proximate hypothesis so that it generates an experiment which can actually be performed. With system-driven research, we may attempt the same move, but it is somewhat less effective. If you delimit the system too much, it might become easier to investigate but it might end up not being the system you actually wanted to investigate in terms of addressing the scientific problem of interest. Continuous cell lines provide an excellent model system for proteomics, but they are famously not exactly the same as normal tissue. This deviance from normal can be identified and annotated.

In general, the OES elicits a technological response to the problem of how systems constrain experimental design: the analytical techniques themselves get changed. Better technology gets developed. A good example is the use of laser dissection so that proteomics can be done on a single, (typically) normal cell, well-described in its original tissue. Another example is the fairly recent development of excellent quantitative mass spectrometry. Meticulous description of the system still serves to keep these practical manoeuvres straight and thus help to minimize error.

6.1.2 Addressing the system

It is unsurprising that both big-picture (top-down, see Section 2.1) and narrow-focus (bottom-up) systems analysis make assumptions which inform how they address the system as an object of investigation. Some of these assumptions typical of both big-picture systems analysis and narrow-focus systems analysis relate to the intrinsic worth of the system under consideration and their ability to characterize it. Both attach value to the system under consideration. However, there are also important differences in their assumptions. Narrow-focus (bottom-up) analysis operates on the assumption that exhaustively working piecemeal with a concatenation of hypotheses will describe the system because all its elements are connected and potentially knowable and because known data generate accurate implications. This is the standard approach, closely

[115] In this respect, bacteria are somewhat more tractable for proteomics, and it is really interesting that most of the recent advances with metalloproteomics involve bacterial systems (Cvetkovic et al., 2010, 779; Barnett et al., 2012a, 3373-4; Arguello et al., 2013, 1 of 14).

connected with the HDM. Big-picture analysis assumes that the whole system has been physically conveyed into whatever technological machinery is being employed and the available analytic apparatus is capable of detecting all the elements in the system. In contrast to the assumptions of narrow-focus analysis, the noteworthy feature of these latter assumptions, typical of the OES and implemented in proteomics, is that these assumptions can be assessed and fine-tuned. The fidelity and completeness of getting the system into the analytic machinery can be determined. The accuracy of the biophysical analysis can also be measured. Of course, either assumption can be shown to be unfounded, but both are susceptible to remedy by improving preparative techniques or analytical apparatus. Then the 'improved' situation can be re-assessed.

One of the problems inherent in examining technical validation in omics is that this process of assessment/improvement has been going on in the past 5-10 years as the field has moved rapidly and shortcomings have been exposed. It can be argued that, like the narrow-focus approaches, big-picture approaches also depend upon the assumption that all the elements of the system are connected. Indeed we have discussed this concept of networks at some length in Chapter 2. Technical improvements in proteomics have not addressed this broader assumption. They have addressed the various assumptions which relate to conveying (with great fidelity) a system into analytic equipment and then carrying out the analysis to a high degree of precision. These are the assumptions which can be articulated, fine-tuned and re-assessed. In the next section we turn our attention to these efforts.

6.2 Generating high quality data: designing and carrying out the experiment

6.2.1 For OES, good experimental design = what?

It is regarded as axiomatic that a poorly-designed experiment will yield worthless results.[116] The problem for the omics experimental strategy (OES) is to figure out what

[116] It is not very difficult to think of experiments where bad design or accident yielded "good results": many researchers can recount experiences of such serendipity. However, I believe this must be classified as a form of epistemic luck in biological/biomedical research. Whether this is the one situation where the context of discovery versus context of justification dichotomy might still find a role is open to debate. I will take the position that the dichotomy does not operate here. First of all, the move to a naturalized

counts as good experimental design. The "pertinence of the methodology" is an important component of evaluating this research (Nabel, 2009, 54). In this regard, I suggest that for the OES good experimental design has a 'concept' component and a 'practicality' component. The **concept** component relates to having a clear delineation of the system to be characterized. I have just discussed this aspect of OES experimental design in some detail. It may seem extremely elementary, but we have already found in the available scientific literature a tendency to equate a basal hepatic Cu-metalloproteome in HepG2 cells (She et al., 2003, 1307) and a copper-induced hepatic proteome (Roelofsen et al., 2004, 733). The two can be profitably compared, but they are not identical systems. Recognizing the actual features (and both the advantages and disadvantages) of a system can be a subtle task. In terms of experimental design, however exhaustively it is described, the system chosen may have inherent advantages and drawbacks. Choosing what system to investigate is always a critical decision.[117] Examining a tissue which has been prepared as a liquidized slurry may be very interesting and highly informative, but does it matter to the interpretation of the experiment that more than one cell type is present in the preparation? Given the relative proportions of the various cells in that tissue, the answer may be 'no'. If the system investigated is a homogeneous continuous cell line, then relevant questions include how faithfully that cultured cell reflects the function of a normal cell and whether absolutely faithful reproduction of the physiological situation matters. Investigating a single component organelle of a cell type—such as just the nucleus or just lysosomes—may sacrifice the complete cellular environment of that organelle but generate a cleaner data set, with better identification of all the proteins in the

---

scientific epistemology, as opposed to one dominated by logical analysis, vitiates the dichotomy. More relevantly, even if epistemic luck sometimes contributes to forward movement in the research project, it still requires observation and insight to appreciate the potential benefit arising from the ostensible disaster. This once again represents a smudging of the boundary between context of discovery and context of justification. See relevant discussion in Chapter 3 (Section 3.3).

[117] This aspect of contemporary experimental biology is effectively a fundamental truism of biological/biomedical research. There are numerous examples. Eric Kandel lavished much consideration before choosing *Aplysia* as his model system (Kandel, 2006, 145-8); Sydney Brenner took the risky approach of developing an entirely new experimental model (*C. elegans*) which involved producing all the probes for molecular biology (García-Sancho, 2012, 18).

proteome and freedom from contamination by irrelevant proteins which would serve as 'noise'.

The **practicality** component relates to whether it is actually possible to do the comprehensive analysis envisioned by proteomics. Mainly in the past 10 years, this aspect has attracted much attention. As is especially apparent to the non-scientist, the technical equipment is extremely complicated. One fancy complex analytical machine is linked up to another fancy complex analytical machine in sequence. The opportunities for the experiment to go 'off the rails' are manifold, and the situation is further complicated because different laboratories use different machines, whose specifications and adequacy may vary, to perform the same step of the analytical procedure. The experimental process, sometime called "work-flow", is typically divided into three phases: sample preparation, analysis by mass spectroscopy and associated physical interventions, and protein identification via bioinformatics. Each phase of work-flow has its points of vulnerability. (Perhaps surprisingly, bioinformatics has numerous attendant problems.) The relevant literature is complicated because the field has moved so fast. It is difficult to escape the impression that circa 2002 proteomics promised much scientific excitement but within a few years was regarded as something of a bust, anything but the envisioned engine for rapid progress. Although it has been speculated that proteomics failed to catch the imagination of cell biologists because it lacks the visual immediacy of live-cell video microscopy (Bell et al., 2007, 783), problems such as variability in identified proteomes (that is, failure of repeatability) and failure to identify low-abundance proteins which on theoretical grounds were obliged to belong to the proteome in question (that is, insensitivity) seem more relevant to the failure to win general interest and approval. Put simply: it was possible to produce a lot of data, but none of the data was very sound. Over the same period, the aims and interests of proteomics scientists also matured: quantitative methods were needed. Protein modification and protein interaction became more urgent priorities for technical development. The range of proteomics experiments was also extended.

Evaluation of proteomics experimental results in the early 2000s also revealed inadequacies of the scientific communication strategies available and in place at that time. Domon and Aebersold wrote in 2006: “The meaningful comparison, sharing, and exchange of data or analysis results obtained on different platforms or by different laboratories remain cumbersome mainly due to the lack of standards for data formats, data processing parameters, and data quality assessment” (Domon and Aebersold, 2006, 1921). This lack imposed immediate demands for figuring out how to expedite collaborative communication among scientists doing proteomics. Part of the solution was the development of standards for reporting data and effecting communication. Another part of the solution was working out how to communicate effectively via web-based virtual links. Moreover, the growing recognition that proteomics needs to be capable of being integrated with other omics disciplines (Bell et al., 2007, 784) has also imposed adjustments in technology and experimental design. Consequently the proteomics field itself exhibits features of being in flux for approximately 10 years. A critical problem has been to invent solutions suitable for proteomics, and one of the difficulties has been a lack of previous examples. In particular, as I will show, taking genomics as an example has not always been a winning strategy. These issues with standardization and communication in proteomics are typical of the problems for omics moving from the genomic era to the post-genomic era.

6.2.2 Sample preparation

In general, sample preparation focuses on getting the sample through initial stages of preparation so that it is ready for analysis by mass spectrometry, the process by which proteins present are actually identified. For the cell biologist, however, the issue is pushed back to the stage of actually working with the system to be studied. A critical consideration is to ensure that the biological system is sufficiently intact when it is submitted to physicochemical analysis. For example, the process of cell preparation may liberate endogenous enzymes which destroy or denature proteins under investigation. Accordingly, the tissue need to be treated with chemicals chosen to inactivate such enzymes or chemicals produced in the cell. What is interesting is that the proteomics literature as such, possibly because it is dominated by biochemists, pays

little attention to this aspect of experimental design. One really attractive solution is to analyze the proteins in a single cell. The fairly recent development of special microdissection techniques, called laser ablation, permit proteomic analysis on a cell-by-cell basis (Alkhas et al., 2011, 5265-6). Developing and standardizing this method are in progress (Alkhas et al., 2011, 5269).

The ordinary focus on sample preparation within the actual practice of proteomics is much narrower. The envisaged problem is, first, to get all the proteins in the specimen, and secondly, to get these proteins separated out and chopped up into small pieces, known as peptides, which can actually be analyzed by mass spectroscopy. Various techniques exist to increase extraction of all proteins in the specimen: these include physical treatments such as high-salt wash followed by detergent extraction and then detergent-phase separation (Bell et al., 2007, 783). In the early days of proteomics, methods were not available to liberate proteins embedded in intracellular membranes for proteomic analysis. For this reason we were not perturbed at failing to find the Wilson ATPase, which is a membrane-spanning protein, in our descriptions of the hepatocellular Cu-metalloproteome. Physical treatments like these now largely solve this problem (Gilchrist et al., 2006, 1265). It is also possible to "refine" the specimen. Such refinements are not random: they reflect the biological issue at hand. For example, with metalloproteomics we were interested in proteins which bind copper and therefore we passed our tissue preparation through an IMAC column charged with copper. By submitting what was stuck to the column to proteomic analysis, we refined our preparation to contain mainly copper-binding proteins. We could have charged the columned with iron or silver, if we had been interested in the metalloproteomes for those metals. For those interested in protein complexes, the cell/tissue preparation can be modified so that it contains only a protein of interest and any proteins possibly bound to it. This is one way to study protein-protein interactions via proteomics (Cox and Mann, 2011, 285-8). As with the metal-charged column, it is possible to charge a column with the protein of interest and thus attempt to capture all possible interacting proteins (Kocher and Superti-Furga, 2007, 809). Methodology is also available to

examine large functional protein complexes in their native form (Kocher and Superti-Furga, 2007, 810) but in essence this strategy operates on making a pure preparation of the protein complex, not so very different conceptually from making a pure preparation of a subcellular organelle.

The second aspect of sample preparation has to do with getting the proteins into the state most favourable for identification by mass spectroscopy. Classically, which is to say in the early years of proteomics, the preferred methodology was two-dimensional gel electrophoresis. This two-step electrophoresis first separates the proteins in the sample by isoelectric point (basically, by characteristics relating to electrical charge on a protein), and then a second type of electrophoresis is applied so that proteins are further separated by molecular mass (that is, by size). An alternative method is one-dimensional separation by electrophoresis using the standard procedure of sodium dodecyl sulfate polyacrylamide gel electrophoresis (so-called “SDS-PAGE”, always the preferred simpler term) in a polyacrylamide gel. With either separation strategy the proteins distributed over the gel are visualized with a dye which would not interfere with the mass spectroscopy step. Then each spot is cut out from the gel and mixed with an enzyme (typically, trypsin) which grinds up the proteins in that spot into peptides. A source of error with this methodology is that extraction efficiency of peptides from the gel is much less than 100%, and the actual structure of the peptide influences its extraction efficiency. One way to circumvent this problem is to digest (break down into component peptides) the proteins adherent to the refining column. The more prevalent solution is to use a different separation strategy. Instead of gel-based protein separation, liquid chromatography is currently preferred. High-pressure liquid chromatography (HPLC) is faster and results in improved recovery of peptides; also HPLC permits automation of the proteomics work-flow. Combination of HPLC with other methods of protein separation improves protein or peptide recovery further. Thus some of the variability in results in the early days of proteomics must reflect the fact that the technology for doing proteomics was relatively underdeveloped. This problem has not entirely been solved. As recently as 2009, variation from laboratory to laboratory in

terms of both methodology for and expertise at “protein identification, protein separation, protease digestion, peptide separation and peptide selection” (Bell et al., 2009, 424) was reckoned as an important source of apparent unreliability of proteomics data.

6.2.3 Problems with mass spectrometry

Mass spectrometry is the analytical process which revolutionized proteomics and permitted its current development. It has been in use as an analytical modality in chemistry for decades, and it is extremely accurate (Brewis and Brennan, 2010, 5-8; Cox and Mann, 2011, 277-8). Likewise, it is fairly obvious that proteomics has prompted extensive development of the technology for how mass spectrometry is performed. A range of designs for mass spectrometry has become available, and the instrumentation chosen needs to fit the proteomics experiment. A misfit between instrumentation and experimental design may lead to a faulty or inadequate dataset. For our work in metalloproteomics, we used some of the most sophisticated equipment then available: it would likely be deemed obsolete only 10 years later. Even in 2007 the recommendation for routine work was different from what was preferred in 2001. Köcher and colleagues write: “If high sample complexity is expected and identification of numerous proteins covering a substantial dynamic range is required, then LC-MS/MS is the method now chosen by most proteomics researchers. Given the importance of unambiguous identification of proteins, most proteomic studies published today are based on LC-MS/MS data acquired from peptides ionized via ESI” (Kocher and Superti-Furga, 2007, 812). Clearly the rate of technical improvement here is very fast.

One of the main pressures for technical improvement relating to the design and performance of the mass spectrometry has been the demand for quantitative measurement of proteomes—not just what proteins are present but how much of each protein. Reviewing best practices for performing mass spectrometry in proteomics research, Bruno Domon and Rüdi Aebersold have specified the desired technical

characteristics as the following: data density and effectiveness,[118] selectivity, dynamic range, reproducibility, repeatability and adequate limits of detection (Domon and Aebersold, 2010, 717). They examine three different types of proteomics experiments and score methodology in terms of these six parameters. Their main point is that no available technology maximizes performance on all six parameters simultaneously, and therefore the technology needs to be individualized to the envisioned experiment.[119] The needs demanded by descriptive big-picture (top-down) "shotgun" proteomics, mainly what we are discussing here, are less stringent than for quantitative proteomics, which in turn poses the greatest demands on technological capability when it undertakes to quantify proteins as biomarkers of disease. These needs, however, are not without some importance.

6.2.4 Bioinformatics and its hazards

Bioinformatics is the computer-based methodology for storage and analysis of very large (which is to say, massive: multi-terabyte) data sets pertaining to biological and biomedical investigations, such as gene sequences and protein primary sequence data (Giometti, 2003, 353; Domon and Aebersold, 2006, 1922; Biron et al., 2006, 5584-8). Reliance on complicated bioinformatics methodology to search databases has become so prevalent that it is easy to overlook the problems inherent in these methods. Both sample preparation and mass spectrometry can introduce lack of reproducibility in proteomics experiments. On balance, however, the accuracy of mass spectrometers is sufficient to eliminate this step as a major source of variability. In terms of the

[118] It is not entirely clear what they mean by "effectiveness" in this context. The term 'effectiveness' can have various meanings and most are subtle. For example, effectiveness can be contrasted to efficacy in pharmacological studies: a drug may be demonstrated as having efficacy in attaining the desired therapeutic effective in the clinical research setting but it may lack effectiveness in the actual real-world context because it is not pleasant to take or it is too expensive or the dosing schedule is inconvenient or people can get over the ailment spontaneously without drug therapy. One reason that the term's meaning is vague is that the authors never really comment on it but merely slip it into their figures and list of critical characteristics. Data density, to which effectiveness is allied, has to do with the number of measurements attained in one experiment: a better mass spectrometer is sensitive and fast enough to produce a large amount of data (Domon and Aebersold, 2010, 716). It appears that this high level of performance amounts to greater effectiveness.

[119] Some of the differences are illustrated graphically in Figure 6 (Domon and Aebersold, 2010, 717) where the performance profile for descriptive top-down proteomics is clearly different from those for quantitative experiments.

bioinformatics analysis, two general sources of inaccuracy can be identified: the completeness of the database (including accurate information and comprehensive data), and the actual design of the search engine.

The extent of this problem for producing technically accurate proteomics data was highlighted in a HUPO-sponsored study whose experimental design was similar to the routine quality assurance protocols utilized by clinical biomedical laboratories. A standardized test sample contained specific amounts of 20 highly-purified human proteins was distributed to 27 laboratories around the world. Their instructions were to identify all 20 proteins as well as 22 unique peptides of a specified mass and report these findings; they were also instructed to submit all the raw data and details of methodology to the lead investigator. All labs used the same protein database for identifying the proteins, but otherwise they were permitted to proceed with their customary instrumentation and techniques. The findings were surprising: initially only 7 of 27 labs were able to identify all 20 proteins and only 1 of 27 was able to identify all 22 peptides. A spectrum of types of error was encountered: some involved problems in handling the specimens (the sort of technical problems discussed in Section 5.2.2 above) but many involved the bioinformatics data analysis including data search errors and overly stringent identification criteria (Bell et al., 2009, 426). Detailed review of the raw data revealed that almost every laboratory had generated high quality mass spectrometry data adequate to identify all 20 proteins and most of the peptides. The problem for overall accuracy lay with the bioinformatics phase of the proteomics experiment. Variability in the search engine used by different laboratories was an important shortcoming in carrying out the experiment (Bell et al., 2009, 428). These findings argue strenuously for standardization of bioinformatics software, since bioinformatics is a weaker link in the work-flow than might have been imagined. Without such standardization, centralized analysis for large collaborative projects becomes critically important. Completeness of the database might also pose problems: even with excellent uniform search engines, a spotty data base may produce incomplete interpretations. This fault is a problem mainly for early proteomics research. Given the

acknowledged shortcoming of this experiment that sample was not at all complex,[120] the urgency to address this source of error is even greater. The importance of this paper is that their conclusion that "a major contributing factor to erroneous reporting resides at the level of data base and search engines used and once corrected for, provided an almost perfect score for most participants" (Bell et al., 2009, 428) was unexpected in terms of the severity of the problem. However, it also points to remediation through development of standards for proteomics experiments.

6.3 Technical adequacy: standards for enumerating a proteome

Several issues are connected to standards for performing proteomics. All relate to validation, at least in the narrow technical sense. The immediate goal is to have a complete data set which is accurate.[121] Standards need to be set so that proteomic data are consistently accurate. In relation to the practicalities of knowledge production, some basic expectations need to be established for how to describe a data set worthy of publication. These stipulations may include certain specifics of methodology as well as specific parameters of the data set which constitutes the 'Results' of the research report. In part, this kind of standard deals with the question of what a peer reviewer needs to know in order to assess the work submitted for publication. Again such standards deal with accuracy of data. A further role for standardization is to facilitate collective research. In 2002 those scientists committed to proteomics recognized that

[120] Specifically, 20 proteins plus 22 peptides is a much smaller number of proteins to be identified in the typical proteomics experiment, and furthermore the sample was artificial in that the proteins were highly purified.

[121] The conceptual language of the biological/biomedical researcher is somewhat impoverished compared to that of the epistemologist of science. The biological/biomedical researcher is seeking correct data. For a proteomics experiment, s/he wants the analytical method to disclose accurately all the proteins in the sample and not report proteins which not there. Moreover the technical performance is to be consistent from day to day. (Issues relating to sample preparation and sensitivity of the technology are relevant to these considerations, but for the purpose of explanation here, I will set them aside.) I will speak of data which are technically correct—conforming to a high degree of technical quality-control—as 'accurate' data. The matter of knowledge acquisition is separate. Whether the data are isomorphic to the system under investigation is a key question for the epistemologist of science, though not inconsequential for the researcher. Data which have been evaluated and found to promote coherent knowledge of the system under investigation have gone through a process of validation, which we might view as a justificatory process. I will call such data 'valid' data.

proteomics could force a major stylistic shift in how contemporary experimental biology was performed: open access to the data from each laboratory. I do not plan to deal with issues of intellectual property and authorship here, although those were immediately apparent, but I do want to examine how the need for effective communication has forced attention to standardization of technical language and content. Standards are required to facilitate communication between laboratories so that every laboratory is operationally speaking the same language by using comparable methodology and instrumentation. This permits one laboratory to utilize another laboratory's data as the basis for new experiments or to check reproducibility of the original data set.

Thus standards in proteomics go beyond addressing the technical accuracy of experimental data: they support a social epistemology of system-driven biological science by permitting public usage of data. This requirement for standardization to promote scientific knowledge production through elaborate collaborations, nowadays often broadly international and multi-continental, is an emerging feature of contemporary social epistemology of science. Unquestionably, having well-articulated standards (and standardization) is critical for these collaborative enterprises, which are typical of omics research. As we look forward toward compiling data about a system across different omics disciplines, the need for standards becomes even more apparent. However, the downside of this standardization is the possibility of stifling individualism and creativity among researchers. "Off-the wall" inventiveness and creative maverick behaviours get lost in this required, and generally constructive, uniformity. Contrariwise, heightened appreciation of features which can confound or derail proteomics research is gained. This appreciation is all-important because many of these features are hidden in the background, nearly inapparent to the individual researcher or analyst because of being so routinely taken for granted.

In terms of developing standards, genomics has taken the lead. Actually, to be precise, the problems of standardization and comparability of data were identified with transcriptomics, where the expression of genes is investigated. Genomes have been

regarded as unchanging[122] whereas what RNAs are expressed and what proteins are produced are both highly dynamic processes. In their recent paper reviewing some of these issues, Keating and Cambrodio lament that establishing standards for analytical performance was not as easy as might have been expected (Keating and Cambrosio, 2012, 43). Their discussion focuses on what are conventionally called gene microarrays, which are actually composed of DNA. These were the first high-throughput technology for assaying and comparing gene expression in different biological systems (Brazma, 2009, 420). In the early days of gene microarray experiments, data were reported in terms of up-regulation or down-regulation relative to a control. Thus the denominator (the control in that lab) was likely to be individualistic, if not idiosyncratic, to that laboratory. Moreover, repeatability within a laboratory was problematic or even erratic. Often the computer program for analyzing the data was peculiar to the laboratory, in fact a guarded secret of the laboratory (Keating and Cambrosio, 2012, 43). What was then sought by the research community was two kinds of standards: for technological programs (that is, bioinformatics) and for reporting practices. Consequently a protocol was developed, with substantial review and input from the relevant scientific community, to establish what information had to be provided about any microarray experiments. This was known as the MIAME (Minimum Information About a Microarray Experiment) protocol.[123] Eventually, scholarly journals in the field adopted the MIAME protocol as a requirement for a manuscript to be considered for publication. However, the MIAME protocol revealed a problem. A key unappreciated point was that standards for what information to provide are ineffective—that is, downright unworkable—without the complementary directives as to format and computer software. The first

[122] Consequently these issues did not arise with the Human Genome Project where the genes to be investigated were seen as being the same irrespective of the individual from whom the specimen came. In 2015 we might regard that attitude as simplistic, given recently acquired knowledge of how epigenetic factors alter an individual's genome (his/her complement of somatic DNA) over a lifetime or in specific environmental circumstances.

[123] Information demanded included the supporting raw data plus "what nucleotide sequences were present on the array, what the assayed samples were and how they were treated, which sample was processed on which array and which data file was obtained in the result, and how were the data processed" (Brazma, 2009, 420).

attempt at a format for MIAME-guided data proved too complex and awkward to be popular, but subsequently a simpler format has proved successful in gaining general uptake among researchers.

Further commentary on the development of MIAME reveals another problem inherent in reporting omics data, whether from functional genomics (namely with these microarrays which interrogate gene expression) which is the subject of Quackenbush's observation, or proteomics or various other omics disciplines: the organism, its age, the tissue actually sampled, and physiological state influence the findings, for example, with DNA microarrays the expression pattern (Quackenbush, 2004, 613). In other words, for what I am calling '*system* biology' the system really matters. We have examined this issue in some detail in Section 6.1.1; however, the additional point made here is that figuring out how to build those details into the reporting machinery, which is mainly computer software, is not straightforward, although critically important. Moreover, standards for describing and reporting experimental data, even when operationalized with user-friendly protocols and software for electronic data exchange, are not equivalent to actuating an objective means for assessing the quality of the experiment (Quackenbush, 2004, 614). For omics experiments the quality of the experiment still hinges on experimental design and overall adequacy of technical resources. Even if it does not solve all the problems[124], MIAME has served as an archetype for other omics

[124] An interesting exchange of views took place in 2006 between Shields and Quackenbush/Irizarry over the effectiveness and scope of MIAME. Robert Shields argues that confidence in microarray technology is misplaced. He deprecates the claim that a comprehensive understanding of the characteristics of high-throughput data cannot compensate for deficiencies in the technology producing those data. He voices worries already articulated by others that microarray technology is such that available platforms are not consistent from one platform to another. He suggests that standards will go only partway to address these problems which are endemic in aspects of the experimental design (Shields, 2006, 65). John Quackenbush and Rafael Irizarry reply, in a spirited fashion, that Shields has misinterpreted their work. They argue that shortcomings are not exactly systemic but emanate from problems in probe design (reflecting limited knowledge of the genome) and occasional problems with annotation. Their basic argument is that microarrays work well so long as you understand what they are capable of and ask them to do no more than that (Quackenbush and Irizarry, 2006, 471). Moreover they point out that MIAME protocol "is not a panacea for all potential problems with microarray experiments" (Quackenbush and Irizarry, 2006, 471). In their view, MIAME is a method for describing the experiment, not directive(s) as to how to do or analyze the experiment (Quackenbush and Irizarry, 2006, 471-2). What we have seen, however, is that delimiting the common language also limits the design and analysis options.

disciplines faced with similar challenges. The benefits and restrictions inherent in such standards, similar to MIAME, are apparent with another type of research community, namely researchers who use the same model organism: facilitating research communication and wide-ranging collaboration but also consolidating research around several canonical model organisms (Leonelli and Ankeny, 2012, 35).

In proteomics, HUPO has taken the lead in establishing similar standards for proteomics research. After the need for standards for reporting was publically discussed[125] in 2004 (Carr et al., 2004, 531), efforts got underway to promote data-sharing by creating standardized data formats and establishing minimum requirements for reporting (much like MIAME), and additionally by attempting to figure out how to ensure public accessibility (Martens et al., 2007, 1666). (Thus development of data repositories assumed major importance early on.) One important point was that "standards serve as essential tools to enable both data quality assessment and its subsequent reuse" (Martens et al., 2007, 1666). With proteomics, however, such standards included not only the detailed description of the system under investigation, but also detailed information as to the technology employed. This latter information is generally called "metadata". Another other key point was that standards need to be subjected to regular review and revision (Martens et al., 2007, 1666), given that proteomics was advancing technologically extremely rapidly. Indeed the proteomics scientific community has made a great deal of progress with respect to these issues since 2004, under the auspices of the HUPO-associated Proteomics Standards Initiative (PSI). As has been established for genomics data, minimum requirements[126] for an adequate proteomics experimental design and reporting of results have been

[125] In addition to discussing the need for standards, the authors of this editorial in *Molecular and Cellular Proteomics* instituted guidelines for preparing manuscripts for publication in this journal. They cited solving the problem of false-positives as one of the incentives for these guidelines. On a broader view, they were interested in making high-quality data available to the general research community.

[126] Note that minimum requirements for reporting are not the same as reporting a "minimum list of protein identifications", as urged by Bell et al. (Bell et al., 2009, 428). This latter requirement, or standard of performance, is focal in nature and aims to minimize redundant reports of proteins, even when these are polymorphisms (sequence variants) of that protein. While the need to keep redundant information to a minimum can be easily appreciated, there are situations where sequence variants of proteins might be of functional interest.

recommended (Taylor et al., 2007, 889). Subsequently this so-called "MIAPE" spawned a series of separate papers which provide guidelines, known as modules[127], for its use (Taylor et al., 2008, 860; Binz et al., 2008, 862; Gibson et al., 2008, 863; Jones et al., 2010, 654; Domann et al., 2010, 654; Hoogland et al., 2010, 655; Martinez-Bartolome et al., 2013, 84) and a semantic validator to check compliance of proteomics data (Montecchi-Palazzi et al., 2009, 5112). A larger report, published after the MIAPE paper, dealt with molecular interaction experiments (Orchard et al., 2007, 894).

What becomes apparent from a 2012 report of the HUPO-PSI is that the original MIAPE has been developed in numerous directions to address the ongoing and emerging needs of the proteomics research community, for example, in terms of studying protein interactions or performing quantitative proteomics (Orchard et al., 2012, 352-3). Moreover, as of 2012, the original module for reporting mass spectrometry (Taylor et al., 2008, 887) was being updated to take into accord technical advances (Orchard et al., 2012, 351). The sheer complexity of assembling literally everyone's data from proteomics experiments is impressive. The apparent capability of contemporary bioinformatics to meet this challenge is also impressive. Thus, in addition to articulating a broad vision of establishing standards, the PSI has also supported the development of diverse formats for data exchange (Kenyani et al., 2011, 2 of 6). A third facet of this effort is the development and refinement of public databases. One of the lynch-pins for all this work has been the development of PRIDE, the proteomics identification database (Jones et al., 2006, D659). A problem is uptake—showing the research community that it is easy and convenient and worthwhile to engage in these data-sharing activities. What is interesting about the 2011 'showcase' paper of Kenyani et al. as an example of these standards in action is that these researchers developed/used tools arising out of the MIAPE enterprise and extending the scope of PRIDE (Kenyani et al., 2011, 2 of 6).

[127] In order, these are for reporting use of mass spectrometry; mass spectrometry informatics; gel electrophoresis; column chromatography; capillary electrophoresis; gel imaging informatics; and for reporting quantitative mass spectrometry-based experiments; plus the check of MIAPE compliance.

Interestingly, there is some difference of opinion about the relative benefits of centralized versus more "grass-roots" control of infrastructure and standards. Cassman and colleagues, mainly representing engineering and bio-engineering research groups, expressed the opinion that lack of infrastructure was hindering development of systems biology as such (Cassman, 2005, 1079). Their prescription was to centralize the data-sharing process with easy-to-access, user-friendly software which can integrate different kinds of data; moreover, they urged annotation of the system studied, much as outlined above, and details of technology employed. The reply from Quackenbush and colleagues bristles from the indictment of innumeracy among biologists (rendered by Cassman et al. and highlighted in a text box by *Nature*): they pointed out the risk of standards imposed from above being based on false assumptions about the research community and ignorance of its needs (Quackenbush, 2006, 24). It appears that the international proteomics research community has managed to maintain a fair amount of democracy in the process by which operationalized standards have been developed.[128] What this discussion sidesteps is the issue of having different operational standards depending on the purpose of the proteomics research (Duncan and Hunsucker, 2006, 47). It might be argued that this early criticism has been met by having multiple modules which address different kinds of proteomics experiments. Moreover, standards must evolve as the field matures. In general (and quite appropriately), standards become more stringent. Such evolution requires input from the "grass-roots" contingent: those who are doing experiments, developing techniques, and formulating new questions to investigate.

A different aspect of the evolution of such standards is that newly-emerging technologies may require entirely new and innovative guidelines, different from the original ones. For MIAME, this progression is represented by the set of standards known

[128] Nevertheless one comment—their closing salvo—by these authors deserves further consideration: "Creating a rigid standard before a field has matured can result in a failed and unused standard, in the best of circumstances, and, in the worst, can have the effect of stifling innovation" (Quackenbush, 2006, 24). The intended thrust is that rigid standards developed prematurely do not work. (Hence we see the need for constant review and revisions of such standards.) There is also the potential limitation on thinking imposed by uniformity in data-reporting. Vocabularies get set and serve as a kind of intellectual commitment, possibly at risk of stifling innovative taxonomies.

as 'Minimum INformation about a high-throughput SEQuencing Experiment' (MINSEQE) (Brazma, 2009, 422). MINSEQE faces somewhat different challenges compared to MIAME: for example, with studies of human subjects, the potential for violating the privacy of the individuals furnishing the samples. It is not clear that proteomics faces the same challenge. Proteomics data are so diverse that easy one-to-one correspondence to reveal the identity of the source of the biological sample is not likely. The issues with proteomics might have to deal more with characterizations of groups based on *post hoc* 'big-data' analysis.[129]

If we step back from focusing on proteomics or other omics disciplines individually, we can see that creating standards is critically important if omics disciplines are to contribute to the classic systems biology project of analyzing and modeling a biological system. Taken in aggregate, omics disciplines permit the broad examination of a system from multiple, complementary viewpoints. I have referred to this as a tomographic analysis of a system (see Section 1.2). Having comparable standards, such as MIAME and MIAPE, which ensure commonality among these disciplines, becomes all-important or else this broader vision cannot be realized. Consequently MIAME and MIAPE serve as important components of the mechanism to organize communication and thus promote integration of different omics disciplines. For example, discussions about harmonizing standards for reporting proteomics and metabolomics data (where proteomics looks at proteins and metabolomics looks at small molecules active in metabolic processes) might promote productive collaborations (Orchard et al., 2012, 353). Thus standards in omics have at least three epistemic roles: first, to ensure so far as possible the technical accuracy of experimental data; second, to promote a social epistemology by permitting general access to data among researchers, and thus

[129] This is an example of where system-driven and data-driven research styles are very closely allied, in effect abutting on each other, and yet are different, with different problems relating to the performance of the research. Here we can see that system-driven and data-driven research can be two sides of the same coin; however, data-driven research is retrospective. The data-driven research is conducted after the system-driven research (proteomics) has furnished the data set.

promote further testing and verification[130]; and third, to generate constructive communication among omics disciplines so that a system, as defined for proteomics, can be examined from different vantage-points. However, MIAPE-type standards may still not ask the most fundamental questions, namely, "how will this particular experiment advance your scientific project?" and "what biological problem are you trying to investigate?"; in general, the current MIAPE standards (and their progeny of recommendations, checklists, and such) address issues of uniformity and completeness of reporting to ensure that the data can actually be interpreted and shared.

6.4 Proteomic evidence—protean evidence?

On Domon and Aebersold's list of criteria for choosing the right mass spectroscopy set-up for the envisioned experiment (namely: data density, selectivity, dynamic range, reproducibility, repeatability and adequate limits of detection) (Domon and Aebersold, 2010, 716), "reproducibility" and "repeatability" stand out as being more global characteristics than the others. Data density, selectivity, dynamic range, and limits of detection all relate to characteristics of the instrumentation itself. Repeatability and reproducibility have broader application. Repeatability refers to running the same experiment twice in the same laboratory and getting the same results each time. Reproducibility refers to the situation where others run the experiment in some other expert laboratory and *they* get the same results. In a general way, based on most current research methods (mainly reflecting hypothesis-driven research) these two features define pragmatically what it is for an experimental method to be reliable. One of the frustrating features of experiments in proteomics is that you end up with a

[130]Martens in 2013 (Martens, 2013, 1548) emphasizes the assumption that proteomics works by a social epistemology of science without actually using that conceptual apparatus by name. This line of discussion continues to be very lively with broadening of collaborative computer techniques (Barsnes and Martens, 2013, 1129; Verheggen et al., 2014, 367). Consequently standards serve to support social epistemology of science in contemporary biology. I regard this as an additional way in which standards contribute to data evaluation at least at the technical level and perhaps in terms of validation. Examination of sound data at a collaborative level may contribute to getting informative patterns worked out and thus promote production of new knowledge.

'shopping list' of proteins in the proteome and then you do the experiment again a couple of weeks later and you get a slightly different list. Maybe it is just the 'normal' variability of the analytical procedure taken as a whole. Maybe you made an error in setting up one of those experiments, or maybe the moon was at a different phase, or maybe proteomes are shifty at the best of times. Nevertheless this is different from lack of reproducibility, the problem where different laboratories get different proteomes for the same experiment. As we have seen, lack of reproducibility is addressed mainly by establishing performance norms. It is also different from recognizing that your proteome did not include some exceedingly low-concentration proteins you expected to find there or some membrane-bound proteins which did not get analyzed in your preparation. These are technical aspects open to improvement through improved technology and instrumentation. A technical approach to solving this problem is to prepare samples for determining a proteome as follows: multiple aliquots of the same tissue or cell preparation are combined into one sample for analysis and then that sample is divided into several aliquots which are run essentially in parallel (Vaudel et al., 2012, 520). However, if we step back and examine this issue more theoretically, the abiding issue of repeatability in proteomics is basically a problem of identifying accurate data as such.

6.4.1 Illusory problem

The problem identified here—the variability of results in the same proteomics experiment done a week apart—may be considered illusory. It is uncommon to run any experiment twice and get exactly the same data. Scientists agree upon a benchmark for what is a conscionable dissimilarity of results. The statistical measure 'coefficient of variance' captures this tolerance for lack of exact repeatability: some percentage margin of difference is specified as acceptable. Technical procedures whose coefficient of variance falls below this threshold are accepted, those with higher variability rejected. The likelihood that the data are accurate is routinely quantified by statistical assessment, but the null hypothesis involved in such assessment requires a comparative experimental design, not the open-ended direct design of the big-picture (top-down) proteomics experiment. Nevertheless, Deborah Mayo's work on error may be relevant

to this problem (Mayo, 2000, S196); however, she has worked mainly with physical sciences and she ties her discussion to hypothesis-driven research. Very possibly, her general approach to interpreting how statistics supports scientific epistemology may be susceptible to modification in order to apply to system-driven research.[131]

Alternatively, as we move to system-driven research, we may need to reorientate our entire attitude toward biological systems. Proteomics, as well as other omics disciplines, forces this revision upon us. If we accept "being dynamic" as a basic feature of biological entities, then it may be flatly naïve to expect to get the same results every time we interrogate the system of interest in our proteomics research. The notion that being able to replicate an experiment is an important indicator of the accuracy of results is central to hypothesis-driven research. It involves the assumption that what is being investigated is somehow static. Given the strictly focused nature of the hypothesis-driven experiment, often designed specifically with a control as comparison, of course you would expect that repeatability within the lab and reproducibility between labs is an important aspect of both the quality-control on the accuracy of the data and the validation of findings in terms of consistent explanation/prediction. This expectation simply may not apply to system-driven research because of the dynamic nature of biological entities and the broad scope of a system-driven experiment. Consequently, one of the problems for proteomics, and various other omics disciplines which deal with dynamic systems, is that the prototype of validation from hypothesis-driven research does not pertain and is actively misleading. Moreover, this is yet another situation[132] where genomics falls short as an archetype for working out what

[131] This analysis would be complicated. The general idea would be to take the statistics methodology appropriated to omics, for example to proteomics, as the basis for working out a theory of error assessment in proteomics. In principle, it should be possible to accomplish this philosophical task. However, I expect that this analysis could constitute a separate dissertation and therefore I have put this project 'on hold' for the time-being, although I offer a preliminary approach to the analysis in Chapter 8 (Section 8.3.3).

[132] Other situations include the Human Proteome Project, mentioned previously, and the CAP (complete, accurate and permanent) proteome concept, examined later in this chapter (toward the end of Section 6.4.2). One of the problems with proteomics is that there is the big temptation to look to genomics for inspiration. In general, this does not work. Sometimes, as with MIAME, genomics can provide some guidance. The importance of genomic databases to furnishing the basis for even getting proteomics moving forward effectively in the first place cannot be discounted. As has been mentioned before, the

counts as validation in proteomics because, although we are learning about their variability, genes are static compared to proteins.

6.4.2 Concept of a plurality of proteomes as such in any dynamic system

First of all, it may be important to establish here that biological entities, from which we select and define systems to investigate, are actually dynamic. Typical features of a living thing include growth and movement. Growth is such a commonplace among living organisms that we regard it almost intuitively as a defining feature. Regeneration, as with Planarian flatworms capable of regenerating fully after being cut in half, is a dramatic example of growth. It is easy to visualize the movement of a single fish or even a school of fish. The process by which a mammalian cell divides epitomizes both growth and movement, as the chromosomes duplicate and then line up on a spindle structure which choreographs the division of the nucleus into two duplicates. Of course there are many living things which do not move and/or which grow at a slow, essentially imperceptible rate. Such systems still have dynamic components. Complex biochemical reactions take place imperceptibly in a yeast cell; moreover, we can see the process of generating daughter cells, known as a budding yeast. The steps by which a protein functioning as an enzyme interacts with the substrate, promotes the chemical reaction to be effected and then regenerates itself have been investigated in detail for many such reactions and described mathematically. The processes by which chemicals are taken up in to cells and excreted from cells have been worked out in detail for many such chemicals. Movement of proteins can be appreciated in cells. For example, with respect to copper, based on studies in various cell lines, the Wilson ATPase is found on the outer aspect of the Golgi apparatus when intracellular copper concentrations are at ordinary levels; however, when those concentrations rise, the Wilson ATPase moves from there to the region of the bile canalicular membrane where it expedites the excretion of copper into bile. Cells may undergo structural changes, as well as biochemical changes, as a result of environmental modifications. Taking certain drugs

current consensus is that proteomics on balance is very different from genomics (Cox and Mann, 2011, 274).

chronically can change the structural features of liver cells (more extensive smooth endoplasmic reticulum) and their biochemical features as well (higher expression of certain drug-metabolizing enzymes). In a sense, even maintaining homeostasis is a dynamic process.

It is at least conceivable therefore that a biological entity, regarded as a dynamic system, generates multiple proteomes over even a relatively short time period. Discrepancies between such proteomes do not represent noise or error as such: they reflect the dynamics of the system. Thus the proteome becomes a kind of time-slice of the system under investigation. I would like to suggest that it is somewhat like looking at a crowded dance floor under a strobe light. Each light burst captures another snapshot of the movement on the dance floor.[133] Scientists will reject this idea as unwieldy. It generates incredibly large data sets. It raises problems about how to choose the time interval for each replicate. Philosophers will imagine time-slices that are thinner and thinner, until they are too thin to work with empirically in a laboratory. Both will probably question what it means to be an "intact system".

In order to examine this idea, let's consider a relatively simple experiment, which could actually be performed in a suitably equipped, well-funded laboratory. You want to determine a whole cell proteome for HepG2 liver cells in stable culture. You have a sample of this cell line available, and you know enough about this sample to describe your system (HepG2 cells at confluence) in the required detail. (This aspect of the experiment was discussed in a previous section: see Section 3.3.3.) You grow these cells in 12 dishes just to confluence, and you perform the experiment on the first day the

[133] This could also be called "Duchampian", after the famous painting by Marcel Duchamp *Nude Descending the Staircase No. 2*, currently found in the Philadelphia Museum of Art. However, I think this allusion is somewhat commonplace and therefore will not employ it although it is highly pertinent to this concept and, additionally, has some interesting historical connections to Wilson disease. This painting was done in 1912, the same year that Kinnear Wilson published his landmark paper on what would come to be called 'Wilson disease'. Apparently Duchamp was very influenced by time-lapse photography when he was creating this painting. Wilson was also fascinated by photography and noted the shutter speed in his legends of photos of patients reported in his 1912 paper: this actually gave a quantitative estimate of the severity of tremor since he set up his camera to suppress blur from the Wilsonian tremor. Later Wilson became interested in film-making, influenced by his friend Charlie Chaplin. Films of patients with various neurological diseases, made by Wilson in the 1920s, have recently been found (literally in an old trunk in the attic of the family home) and made available to the medical community.

cells are covering the dish in a neat monolayer. For actually performing the experiment, you remove the medium, replace it with a simple buffered salt solution, and then you harvest the cells appropriately for performing the proteomics investigation at 10-minute intervals over a two-our period. Then you end up with 12 proteomes. I do not know the exact results of this experiment, but I imagine that if you stack these proteomes one on top of the next, you will find an identical set of proteins among all the proteomes. Indeed it will constitute a large percentage of all proteins present. Likewise there will be a smaller percentage of proteins which vary from proteome to proteome, even in this apparently stable system over a relatively brief time-frame. Even though the experiment has been performed with state-of-the-art technology, some of that variation could represent technical error: faulty preparation of the sample, ambiguous identification of some proteins. For the purpose of argument, let's take experimental error out. I believe that the stacked proteomes will still display the following inherent distribution: many proteins found in all the proteomes and some proteins turning up here and there in one proteome or another. I will call the majority the "sector of consensus" and the variable minority the "sector of variability". If we reneged on our stipulation that there was absolutely no experimental error, we could say, of course, that this experiment merely shows the importance of getting enough replicates for any proteomics experiment and optimizing technology. (It is relevant to that real-life situation.) I am suggesting something different and additional: the sector of variability includes key information about the system. In general, that information is excluded from consideration. It may be that the variation is really important. We might find that we are more interested biologically in the variation[134] (a sort of Δproteome) because it provides information about the mechanics of the system or state of the system, where protein variability discloses phenotypic detail. For example, we might find different patterns of post-translational modifications of proteins between the consensus and variability sectors, or we might find that some precursor proteins were more prevalent in one sector.

[134] This idea came up in the course of discussions with Brian Hall in the Vancouver airport *en route* to the Off-year Workshop of the International Society for History, Philosophy, and Social Studies of Biology (ISHPSSB), August 2012.

This notion might be clearer if we return to the image of the dance floor under a strobe light.[135] Suppose it is a popular place, with numerous people dancing at any one time. With each burst of light, you get a different snapshot of the action on the dance floor, with the movement of the dancers frozen in a single frame. With an appropriately positioned camera (or combination of cameras), you could identify and tally each individual on the dance floor over some period of time, perhaps for example a continuous set of dances between the band's first and second break. Each of these would be like an individual proteome as described above. Stacking these 'dancer-omes', you would probably find that the majority of individuals danced the entire time. This would be like the sector of consensus. However, some dancers might dance only some of the time. In other words, the wall flowers or only occasional dancers are also important. These individuals would be in the sector of variability. Getting a complete picture of this "dance floor system" requires assessing both sectors. It might be that changes in the dancing population depend on the kind of music: jitterbug versus "slow" versus cha-cha. There might be changes in the pattern of dancing (mainly couples or mainly solo) which affect who is in the sectors. At a further level of complexity, within the sector of consensus there might be variation in the composition of couples (due to "cutting-in" and "double-cutting"). Of course, there could be technical problems: on a really crowded dance floor it might be difficult to identify all the dancers all the time, a difficulty alleviated by getting more cameras. The point is that for a dynamic system, focusing only on the consensus sector and discarding all the rest as noise or random variation may produce a faulty assessment of the system and miss critical information. Giving up the particularity of the data and cramming the 'sector of variability' findings into categories conceptually resembling standard deviations or standard errors may sacrifice really valuable information about the system. Indeed the degree and character

[135] This example is inspired by a demonstration at the Ontario Science Centre. I imagine the strobe light flashing at fixed intervals, perhaps every 2 minutes. If the strobe light were geared to the sonic output of the music, then the venue would be a discotheque, not quite what I had in mind, although the comparison of the 'dancer-ome' and 'disco-dancer-ome' might be of some interest. It might approximate to 'proteome' and 'metalloproteome', on the assumption that a discotheque exerts some selection over who comes to that more specialized dance floor.

of changes occurring in the system are likely to be highly informative about the actual function of the system.

By way of a brief summary, what I am saying is that the problem for the proteomics researcher is to choose the appropriate system, define that system fully, and ensure that it remains unperturbed for the duration of the experiment(s). Then s/he has to use the best available analytical instrumentation for the biophysical analysis. The bioinformatics phase of data analysis is optimized by using the best available resources, ordinarily the program vetted and endorsed by the community of proteomics researchers and the available databases. In actual proteomics experiments error must be identified. The point I am making is that some of the 'unreliability' of proteomics is technical and it can be sorted out, but some of that 'unreliability'—the variability from proteome to proteome—is ontological due to the dynamism of the system and cannot be avoided.

A question which arises is whether the sector of consensus is the same thing as the "core proteome". It sounds similar. I will argue that the 'sector of consensus' is not the same as the core proteome. According to its definition, the core proteome is the proteome which would be found if all the genes in the genome were expressed. Although this was an early definition of proteomics, nowadays it poses some subtle problems.[136] Presumably the definition intends that all genes capable of expressing proteins are expressed: we know that this constitutes approximately 2% of the typical mammalian genome. Moreover, the diversity of mammalian physiology, given the size of the typical genome, is due in large part to the phenomenon of alternative splicing.[137] So, again presumably, this means that the set of expressed proteins is limited to those

[136] These are additional problems, quite apart from what seems like an unnecessary handicap of basing the definition of proteomics on genomics, a definitional practice I have already deprecated.

[137] It might be worthwhile to recall the example relevant to Wilson disease. The *ATP7B* gene is the structural gene for the Wilson ATPase found in the liver and various other organs including some kidney cells and various parts of the brain. In the pineal gland this gene produces the 'pineal night-associated protein' (PINA) by alternative splicing (Borjigin et al., 1999, 1019). This protein seemingly has more to do with circadian rhythm than it does with copper disposition, but some aspect of copper disposition in the pineal gland may be relevant. PINA can transport copper but by a different molecular mechanism compared to the Wilson ATPase (Borjigin et al., 1999, 1023).

proteins which would be produced when the genome was expressed in that cell or tissue under the control of various mechanisms for tissue-specific expression. This definition renders a proteome nothing more than a collection of gene-products. It eliminates a critical characteristic of proteins: post-translational modifications and other mechanisms which promote protein diversification.[138] In fact, a separate strategy exists for addressing the set of gene-products of the structural genes in a genome: based on exomics. An exome is the subset of the human genome which is comprised of protein-coding genes. Studying an exome, generally by a process known as exome-sequencing (Bamshad et al., 2011, 745-7) can be very informative, but it is certainly not the same as studying a proteome. Moreover, the core proteome concept leads to awkward contradictions. Under this definition, the core mitochondrial proteome would be the set of proteins produced by expressing the structural genes of the mitochondrial genome. This genome is well-known and rather small; the majority of mitochondrial proteins are encoded in the nuclear genome, a result of the symbiosis which led to mitochondria coming into existence in the first place. So the core mitochondrial proteome would not be equivalent to what we recognize as the apparent mitochondrial proteome. For these reasons, I believe we can dispense with trying to find a conceptual equivalence between the 'sector of consensus' of proteomes of a system and the 'core proteome'.

However, it could be argued that the 'sector of consensus' is the *forme fruste* of the CAP proteome. The notion of the CAP proteome is developed from ideas put forward by Sydney Brenner in a lecture given in Montreal in 2004.[139] The "CAP" of CAP

[138] The richness of proteomes seems obvious (to me, and hopefully to the non-technical reader at this point), but it is described clearly by proteomics experts Bruno Domon and Rüdi Aebersold, who draw attention to the same issue as I have: "The proteome is more than the mere translation of the protein-coding regions of a genome. Processes such as alternative splicing, protein processing and post-translational modification are key to providing the full complexity of life" (Domon and Aebersold, 2010, 710).

[139] Sydney Brenner was talking about research on large systems. We have little information as to what he actually said. What was provided is in an acknowledgement: "The designation CAP principle for the complete, accurate and permanent representation of proteomics data was from the elucidation of the principle as it applied to all large-scale efforts in a Gairdner Award ceremony presented by Dr Sidney Brenner in Montreal, 18 October 2004" (Au et al., 2007, 383). This seems to have been the Hughlings Jackson Lecture at the Montreal Neurological Institute. What the authors of the 2007 paper say is: "We also discuss whether a vision of the cell map constructed of complete, accurate and permanent (CAP) proteomes of each organelle is already becoming a reality" (Au et al., 2007, 376). The interesting problem

proteome is an acronym on Complete, Accurate and Permanent. My concept of the 'sector of consensus' rejects permanence and admittedly makes determining completeness difficult, but it may be relevant to accuracy. The proteins invariably found in the 'sector of consensus' likely have important membership in the protein population of the system under investigation. The major problem with the CAP proteome concept is that it systematically ignores the information to be gained from the 'sector of variability'. It enshrines an essentialist approach to defining a proteome. Although it is difficult to judge from the available literature where the CAP proteome concept has been implemented, it is possible that completeness and permanence of the CAP proteome arises specifically because any variability has been excluded as error, noise or outliers. In other words, the findings have been parsed in service of this concept. Moreover, it is hard to see how the CAP proteome is universally workable in the actual experimental situation. Suppose in investigating a proteome, you expertly and consistently find a protein present in the proteome which is not recognized as part of that system's CAP proteome. Do you disbelieve and reject your finding because it is not part of this established lexicon? Do you amend the CAP proteome? If so, the CAP proteome was not what it purported to be. If you reject your finding to preserve the intactness of the established CAP proteome, then the CAP proteome is restrictive[140].

If the important information that we gain about biological systems by studying proteins arises in part from the versatility and variability of proteins reflecting response to environmental factors (defined broadly—within the system itself or outside the

---

here that Brenner was talking about representation and this nuance seems to get lost in the subsequent use of the term. It could be that Brenner was really invoking what amounts to a jargonized use of the term 'representation', namely, getting data into universally accessible databases. Likewise it could be that the "cell map" is seen as at least a shorthand for 'representation' and there is no loss of conceptual content here. However, it could be criticized as a fallacy of equivocation.

[140] The idea of a cell map constructed by adding up the proteomes of the component organelles is ostensibly an example of mereological reductionism: thus the CAP proteome might also be reductionist. I am perhaps being overly picky. A less fanciful criticism has to do with the problems arising from the CAP proteome concept when we see that it has generated comments about "near-CAP proteomes" (Bell et al., 2007, 783). The notion of a near-CAP proteome involves conceptual problems as to how near is 'near' and epistemological problems of accurately perceiving nearness, however it is defined, when it is postulated. The concept of a near-CAP proteome suggests that a CAP proteome could exist but only functions as an asymptote.

system), then how we investigate and report proteomes needs to be capable of taking actual variability into account. This will involve a suite of validation strategies. Statistical methodology may have some utility; however, I suggest that the component proteins of a given proteome need to be validated functionally in terms of how they contribute to the function of the system under investigation. This type of validation would support mechanistic explanation(s) relevant to that system.

6.4.3 Validation in the context of the system

Optimizing instrumentation and bioinformatics technology is obviously important. Adhering to the professional standards for conducting and reporting proteomics experiments introduces needed uniformity of the quality of experimental work and permits critical appraisal of the results. Statistical analysis may support differentiation of a sector of consensus and a sector of variability. Both sectors are likely to be informative; in fact, sophisticated analysis may prefer the sector of variability in order to examine dynamic changes in a dynamic system. However, these strategies (uniformity of technical excellence, standardized reporting of methodology, statistics) may not suffice for system-driven research even though they play prominent roles in hypothesis-driven research.

To address this problem, I suggest we need to look at differences and similarities between these two ways of designing research. With hypothesis-driven research, under the HDM, the experiment has a restricted focus directed at the hypothesis being tested. This entails having a clearly stated testable hypothesis. If the hypothesis is substantiated by the experiment, then it is expected that repeating the experiment will produce essentially the same results. In other words, replicating the experiment (in your lab and in the labs of expert colleagues) is an important aspect of the evaluating procedure; indeed, it constitutes much of it. With proteomics—or any omics directed by the OES—the situation is quite different, since a biological entity or a class of its components, defined as the system under investigation, is being studied. As previously recognized, for the OES to operate, the system must be defined meticulously and in great detail. This is like having a well-formulated hypothesis. Furthermore, the complex methodology must also be described in detail. From the technological point of view, data evaluation is

currently organized according to the quality control standards set by professional organizations of researchers in the field, such HUPO for proteomics.

With respect to knowledge acquisition, the technological story—while fascinating and important—is not the whole story. When we ask, reasonably enough, how we would evaluate perfectly accurate data from a proteomics experiment as new knowledge, we focus a different question. This question is not merely methodological: it is epistemological. Here, I believe we can get a clue from similarities between hypothesis-driven and system-driven research. What drives experimental design also contributes to the evaluative machinery. So the results of the HDM experiment are assessed in relationship to the hypothesis. I submit that the results of the OES experiment are assessed in relationship to the system. Do they fit with what is known about the system? If the answer is 'no' or 'not certain', then further investigation is needed. As we have seen, a problem with the system-driven experiment is that it is hard to exclude results as irrelevant. Indeed, as mentioned earlier, some prominent experts in proteomics have called on cell biologists to 'get with it' and initiate these evaluations! Whereas HDM operates in an epistemic environment of synergy between explanation and prediction, OES typically exists in the epistemic context of a system that is coherent—it literally fits together.

One reason to recommend coherence as the epistemic environment of system-driven research is that a precedent already exists with another kind of non-hypothesis-driven experimental strategy. Coherence operates for research addressing the structure of a biomolecule. As previously stated, I regard structure-driven research as non-hypothesis-driven. A biomolecule can be treated as a system in miniature. It poses on a small scale the problem of investigating a system directly in its entirety. This precedent to large-scale system-driven research can be appreciated more easily with examples. A structure for the Wilson ATPase which did not account for its interaction with copper would be a non-starter. In the important example of DNA, the Watson-Crick structure explained certain previous key findings (the observed ratios of nucleotides, known as the Chargaff ratios, referring to the scientist who described them) and suggested how

faithful copying of DNA for daughter cells might occur. It is fairly clear that the Chargaff ratios influenced the thinking of Watson and Crick. Their intuition that the molecular structure had to account for those observations was correct. With the double helix as they described it, the paired interlocking nucleotides literally fit right in.[141] Moreover, the structure was consistent with a viable mechanism for faithful copying.

When we consider the practicalities of validating omics data in a system, it immediately becomes evident that the size of the system has to be taken into consideration (Bell et al., 2009, 428). It likely matters. With the example of a biomolecule, as just discussed (relevant, even taken from a different type of non-hypothesis-driven research), comparatively few criteria were critically important. Small-ish systems may be easier to validate. It is tempting to say that small systems ought to be easier to validate, but that assumption appears unwarranted to me. The standard example is the attraction of breaking a system into subsystems and then investigating each subset individually. This is the strategy espoused when looking at cellular organelles, or as we did, with the functionally-defined subsystem of the cell such as those proteins involved in copper disposition. The small system may be more tractable in terms of numbers of proteins to be identified, opportunities to show completeness of a small well-delimited set, and possibilities for distinguishing contaminants from proteins actually in the proteome under consideration. Fitting the smaller system back into the larger system may be a challenge and it may impede the sorts of validation related to pattern detection, which I will discuss shortly. Moreover, size may not matter so very much as our capability for handling extremely large data sets continues to improve.

Finally, a further difference in the justification strategies of hypothesis-driven versus non-hypothesis driven research demands our attention. We are aware of several contrasts between research which utilizes a concatenation of hypotheses, informed by

[141] It is worthwhile noting that Watson's description places the causal relation as running from the helical structure to the ratios: "Two irregular sequences of bases could be regularly packed in the center of a helix if a purine always hydrogen-bonded to a pyrimidine. … Chargaff's rules suddenly stood out as a consequence of a double-helical structure for DNA" (Watson, 1968, 196).

the HDM, and omics, described by the OES, in terms how evidence is produced. One contrast we have not considered is the applicable logical structure. Hypothesis-driven research investigates a hypothesis and produces evidence which attempts to determine the truth or falsehood of the hypothesis under consideration. System-driven research investigates a system and produces evidence which attempts to characterize the system. In order to be inclusive given the size and complexity of this inherently dynamic system, descriptors need to be supplemented with something to 'bookmark' uncertainty. It may be a category like "don't know (yet)", the optimistic version of "don't know". This option militates against exclusion of data for the purpose of immediate simplicity. Judgment about specific aspects of the system is reserved until the picture of the system starts to come together more completely.

6.5 Where are we?

My discussion in this chapter began with a brief consideration of what a system is, given that it plays a commanding role in system-driven research. I have described a system as being a delimited collection of components whose structures and function are somehow integrated. In designing system-driven research, the investigator has a great deal of discretionary space for deciding what system to study and how to delimit it. Then, since biological entities are dynamic, we find that the system under investigation is also dynamic. I recognize the formidable risks of equivocation on the term 'system'. In common parlance we speak of a great many different kinds of biological systems, ranging from an ecosystem to a cardiovascular system to a subcellular organelle. Sometimes an entire biological entity (for example a virus) is the system under investigation. Finally, other philosophers of biology use the term 'system' differently from how I do. Nevertheless I find the term too valuable for my purpose to abandon it. I then reviewed in technical detail how proteomics as system-driven research proceeds in the laboratory, with particular attention to sources of technical error and how proteomics researchers have dealt with these challenges through development of standards and technological improvements (including improvements to bioinformatics).

These improvements notwithstanding, repeatability is problematic in proteomics. I have argued that lack of repeatability is endemic to studying proteins in systems: it reflects of the variability of biological entities. Finally I argued that validation of findings with the HDM turns on the synergy between explanation and prediction but validation with the OES involves coherence of the system under investigation.

Thus with proteomics, as typical of system-driven research, the researcher defines a system related to some biological entity, identifies the system's component proteins, and then evaluates whether all those proteins really are part of the system by working out how they may function in the system. The resulting coherent picture of the system has its peculiar elegance. It may include some unexpected components or configurations. Given this process, clearly and specifically defining the system takes on added importance. Producing technically excellent data is obviously important—through excellent experimental design and adequate technological proficiency (as defined by ambient professional standards, which also expedite review of data). The key epistemological issues of (1) justification as knowledge and (2) pattern detection will be explored in greater detail in the next two chapters. The challenge of taking this real-life laboratory experience with metalloproteomics and creating a scientific epistemology for system-driven research from it is the agenda for Chapter 7.

CHAPTER 7 OES AS A SCIENTIFIC EPISTEMOLOGY FOR PROTEOMICS

An important strength of the HDM is that it is a scientific epistemology, as developed in detail by Logical Positivists and subsequent 20th-century philosophers of science. As a scientific epistemology it demarcates science from pseudo-science. The HDM also has limitations, not the least of which is that it excludes research typified by omics and other high-throughput methodologies from being classified as legitimate science. The project in which I am engaged here is an exercise in scientific epistemology. The fundamental motivating question is from real-life contemporary biological/biomedical research, namely, to ask how we investigate a biological system directly in its entirety and actually generate new knowledge. I have suggested (and portrayed in tabular form in Table 2 of Chapter 1) that my project involves two phases. The first phase has to do with experimental design and the production of high-quality experimental data (which could also be called 'empirical data'—see Section 8.2.3). The second phase has to do with evaluating those experimental data in terms of how they may constitute worthwhile trustworthy information: this is a problem of knowledge acquisition or knowledge production. The first phase might be regarded as mainly a technical problem, a matter of instrumentation, and I have examined those issues in Chapter 6. I have shown that the first phase is not limited to merely technical aspects. I have argued that the very nature of a biological system under investigation imposes important constraints and specifications for the production of accurate experimental data. The second phase, knowledge acquisition (as experimental data are somehow transformed into new knowledge), involves critical conceptual issues. In effect, the question is whether or not an experimental strategy that supports high-throughput methodologies like proteomics qualifies (or can qualify) as a viable scientific epistemology. I expect to answer this question in the affirmative. In this chapter I will explore how the OES constitutes a scientific epistemology.

Specifically, I will briefly review what it means to examine a biological system directly in its entirety, and then I will examine the OES as a scientific epistemology. I will

consider in some detail the mechanics of such a scientific epistemology, which I conceive of as operating through production of consistently accurate experimental data and then contextualization of those experimental data within the system under study, placing prime emphasis on detection of patterns within the data. While I draw on the resources of certain existing general epistemologies, my opinion is that they fall short of offering the resources the OES as a scientific epistemology requires. However, I will draw on reliabilism as the best paradigm among available kinds of general epistemology for developing the OES as a scientific epistemology.

7.1 A preliminary consideration: thinking by looking

In an comparatively early discussion of contemporary systems biology, Olaf Wolkenhauer pointed out that the big task in the post-genomic era was the investigation of the organization and control of genetic pathways (Wolkenhauer, 2001, 258), pathways which he described as "dynamic systems[:] non-linear, adaptive and anticipatory systems to be precise" (Wolkenhauer, 2001, 258). Moreover, as a systems biologist, he expressed greatest interest in the development of formal models to capture such analyses. To summarize briefly: in his estimation what is important about systems biology is that its objective is model-building. As we have noted, in general this objective does not apply directly to omics,[142] although omics research can, and often does, provide information critical for model-building. Each of the omics disciplines contributes a different perspective on the system under consideration: combining these various perspectives is regarded as valuable for comprehensive systems analysis.[143]

[142] We can also quibble with—indeed, reject—his characterization of genomics as "the field of biological research taking us from the DNA sequence of a gene to a structure of the product for which it codes (usually a protein) to the activity of that protein and its function within a cell and, ultimately, the organism" (Wolkenhauer, 2001, 258). This view of genomics seems to be a relic of the genomic period. It turns up in papers written around 2000, possibly with recognition of the broadened scope from genomics as such to functional genomics, and it generates conceptual confusion. In our discussion we have assigned different omics disciplines to the different aspects of this spectrum of investigation.

[143] Proteomics may be exceptional in that it can contribute to this "tomographic" project to characterize a system or it can function very effectively in its own right. For example, proteomics can be employed on its own to identify members of a protein network which performs critical functions within a cell. Equally it can contribute to a more global look at that cell as a system in itself—from the standpoint of proteins found therein.

However, as a systems biologist, Wolkenhauer makes a really important point in connection with model-building. He argues that systems biology is a way of thinking[144]: "system theory is not a collection of facts but a way of thinking, which can help biologists to decide which variables to measure and to validate their 'mental models'" (Wolkenhauer, 2001, 267). Of course, there are 'philosophical' problems with this view: it relegates systems biology to an instrumental role, and it equivocates as to exactly what role the model plays—eliciting the inference that perhaps this role is not so different from what a propositional hypothesis does.[145] My objective here is not to examine Wolkenhauer's assessment in detail and refute it, but instead to move our focus back to the level of looking at the system itself. In *system* biology we address the system and gather data through complicated technologies, as described in Chapters 2 and 6, and then we consider the data critically.

In general, omics is a way of looking, but put perhaps more exactly, a way of thinking by looking. The 'looking' I have in mind, however, is not merely superficial surveillance (as in "taking a look") or census-taking. The output of 'thinking by looking' is knowledge. In terms of knowledge production it is strikingly similar to what is going on with hypothesis-driven research. It would not be far-fetched to describe what is going on with the HDM as 'thinking by spinning hypotheses'. If I am working at my desk and suddenly become aware of a noise in the background, I may respond by formulating a rapid stream of questions, each of which is really a hypothesis (according to my strict definition of governing an experiment), each of only a few seconds' duration: is it the phone? is it my cellular phone? is the backdoor bell? is it the stove timer? With the first three answers checked out as negative and aroma of cake permeating the study, the last

[144] Of local interest to the Dalhousie University community, Wolkenhauer devotes much of this paper to reviewing the work of Robert Rosen as it bears on systems biology (or 'relational biology' as Rosen liked to call it). It would be of interest to review Rosen's ideas in the context of my project (perhaps with the additional consideration of potential connections to feminist epistemology), but I regard that project as located beyond the reasonable scope of this current work. Importantly, note that in the quotation cited, Wolkenhauer actually does use the term "system theory", not "*systems* theory".

[145] My own view is that a model functions more like a hunch hypothesis: it is at one remove from whatever determines experimental design, and it serves as a kind of mental or theoretical space, subordinate to the over-arching hypothesis, within which the investigator frames his/her investigative plan.

question/hypothesis carries. 'Thinking by hypothesizing' is so prevalent we might not pay any attention to it. Likewise, 'thinking by looking' is not as extraordinary as a focus on arcane experimental methods like omics and system-driven research might suggest. I will argue that for proteomics and many (or most, or even *all*) other types of omics experiments the central epistemic action with 'thinking by looking' is *pattern detection*. The important epistemic manoeuvre here is appreciation of how the data resolve into a salient pattern[146] (or salient patterns) which can then be contextualized within the system under investigation. Some of the important philosophical issues before us have to do with the nature of pattern detection as an epistemic entity, how patterns function in this epistemic domain, and how objectivity is maintained (that is, how bias is excluded or, perhaps more realistically, minimized). However, before we address these issues, which I will deal with in Chapter 8, we need to backtrack and attempt to put this variety of knowledge production into some sort of orderly epistemological superstructure.

An immediate question is whether any existing epistemological theory intended to address issues relating to knowledge in general—how we know and how we may regard ourselves as knowing appropriately and accurately, that is, holding a belief with justification—can serve at least as an inspiration or broad exemplar for a much narrower theory regarding scientific knowledge. A scientific epistemology such as the OES must necessarily differ in some important ways from general epistemology. Standard competencies of or objections to various specific epistemological solutions may require revision based on some needs of a practical 'scientific epistemology'. An example may advance this argument. If we take the famous example of real *barns* versus *barn facsimiles* in analytic epistemology, the problems and implications of distinguishing the one from the other are of utmost importance to the epistemologist—and reasonably so. For the landscape painter, however, the differentiation may not really matter so long as the vista retains the aesthetic qualities sought. For the cowherd

[146] I will use the term 'salient' sparingly in this discussion because I believe 'salience' as an epistemic concept is vague. We know what it means but have difficulty saying what it means. In fact, as will be evident in Chapter 8, the concept of 'pattern' suffers from similar vagueness. For the purposes of my discussion, a salient pattern is one that achieves enough coherence to be appreciated. The pattern can then be contextualized. A 'salient pattern' is not limited to the pattern after contextualization.

(and the cows!), the differentiation is important for the practical reason of needing to find overnight shelter for the herd. Scientific epistemology exhibits similar practicality. An objective in my analysis is to find an epistemological structure for omics which optimizes the strengths of a pure epistemological argument without neglecting a special needs of an epistemology for a scientific endeavour, which is fundamentally practical, and specifically for biological/biomedical experiment-based research.

The experimental process itself as performed in the laboratory (or research space) constrains certain aspects of this problem. What interests us is the epistemic movement from a collection of empirical data obtained by experiment to the formation of knowledge. For biological/biomedical scientists, who if nothing else are practical and concrete in their attitude toward doing experiments, the empirical data are 'out there'. This recognition points to a special feature of an epistemological project designed to be relevant to some sort of biological/biomedical experimental practice. The problem before us here is to determine how an omics data set (for example, a proteomics data set) obtained under the auspices of the OES proceeds to being accepted as scientific knowledge. For the biological/biomedical scientist and philosopher of biology such knowledge involves a mechanistic or explanatory component, a requirement which will put extra burden(s) on the epistemological apparatus of the OES. The HDM accomplishes this task (that of involving a mechanistic or explanatory component) for hypothesis-driven research. In general the hypothesis articulated to drive an experiment (in my terminology, the proximate hypothesis) is articulated in such a way as to investigate a mechanism directly or else contribute to accrual of data which will serve in formulating the explanation of what is happening in the process or problem under investigation. The OES addresses explanatory issues relating to mechanism indirectly, but it does have the capability to provide important information about mechanism by permitting a broad examination of the system under investigation. The point is that there is an element of interpretation in either process. It is subtle, though: with HDM, interpreting the problem at hand is interjected at the point of formulating the hypothesis, but with OES it participates latterly in the process of examining and

evaluating the data, that is, in pattern detection. Determining whether and how a pattern fits into a system involves interpretation of both the pattern and the workings of the system itself.

7.2 OES as a scientific epistemology (as illustrated by proteomics)

Utilizing extremely sophisticated equipment and complicated biophysical and biochemical techniques, we can currently employ proteomics to examine numerous aspects of proteins in organisms. We can find out which proteins are expressed within a cell or tissue and how those proteins actually function, including quantity present, post-translational modifications, and complex interactions with other proteins. Criteria for assessing the technical capability of any specific ensemble of analytical instruments (for example, variably-complicated mass spectrometers and attendant set-ups for protein preparation and separation) have been developed by scientists who are expert in these techniques. Complicated stipulations designed to ensure accurate data have been assembled by leaders in the field of proteomics in order to assure, so far as possible, excellence in publication of experimental findings made available for collaborative research. It seems as if what really matters here for knowledge production in proteomics is the getting and sharing, critiquing and further developing high quality data. This epistemic process fulfills the criteria of being a social epistemology of science, and as pointed out in Chapter 6, a complex social epistemology of science is at work with proteomics. However, the OES operates prior to this social epistemology, just as the HDM also operates prior to it.

The near-obsession with technical excellence provides a clue to the epistemic orientation of the OES. The empirical data must be consistently accurate. We hear this emphasis in how biological/biomedical researchers informally describe their best omics empirical data: “sound”, “accurate”, and “good”. I have argued in Chapter 6 that producing ‘good’ proteomics data involves excellent experimental design, in addition to excellent (preferably the best available), expertly-used instrumentation and bioinformatics. Of course, the researcher employing a hypothesis-driven experimental design also wants consistently accurate empirical data. Research methodology which

produces consistently accurate data is reliable. I will argue that the HDM places a premium on reliability, just as the OES does. However, the HDM also promulgates repeatability and reproducibility as indices of reliability. Here the OES necessarily differs from the HDM. Thus, as a scientific epistemology, the OES must establish other criteria for *overall* reliability of empirical data, which will include the recognition of valid—not merely accurate—data. One consequence is that the actual accuracy of the empirical data assumes greater significance with the OES than with the HDM because, unlike the HDM, the OES cannot necessarily fall back on repetition/replication to double-check the accuracy of the data. It is also important to note that taking reliability as an important parameter of the OES as scientific epistemology introduces a normative element into the OES[147]; moreover, as we have seen, the system being investigated serves in large part to specify the pertinent normativity.

7.2.1 HDM: reliability behind the scenes

The importance attached to reliability by the HDM may be best appreciated by considering an example. When I examined the possibility that copper-associated liver damage is due to oxidative stress, I tested a hypothesis, namely, that in a liver cell copper generates toxic 'free radicals' (see Section 4.4). This is the proximate hypothesis in that it governs the experimental design, but I can spell it out in detail: copper in HepG2 cells generates detectable "reactive oxygen" chemical species in a dose-responsive manner, that is, higher concentrations of copper produce higher concentrations of these toxic free radicals. Adherents of the classic HDM approach might consider the first version a hypothesis and the second a test-implication, but in my example not much is gained by putting such a fine point on it. For quantification, we used a fairly new assay which involved adding a dye (2′,7′-dichlorodihydrofluorescein diacetate; DCFH-DA) to the culture, adding copper or zinc, and then measuring a fluorescent product (dichlorofluorescein; DCF), proportionate to the reactive oxygen species produced, at specific wavelengths in a microplate cytofluorimeter (Seth et al.,

[147] I take this conception of 'reliability' as being intrinsically normative from Marcel Weber (Weber, 2005, 147), but my refinement with the OES is that the system under investigation helps to set the criteria for judging reliability in each instance.

2004, 503). According to the HDM, finding increased fluorescence after incubation of HepG2 cells with copper over a range of increasing concentrations would substantiate our proximate hypothesis. As a control, we performed the same experiment with zinc, not copper, as the metal tested. Chemically, zinc is similar to copper, but zinc is not found in two valences and thus is not redox active. Iron, which also can be found in two valences (2+ or 3+) and is redox active, would not have been an appropriate control. What I want to emphasize here is the extensive framework of technical skills, technological assumptions, and background knowledge supporting this hypothesis-driven experiment. The technicians and trainees performing the experiment had to know how to work with cells in culture: how to culture HepG2 cells in microtitre plates so that each well had a similar number of cells without contamination, how to calculate the right amount of copper or zinc to be added, how to prepare the fluorescent dye and add the right amount by having already acquired excellent micropipetting technique. In terms of technology, we needed access to the right kind of fluorimeter, we had to know how to use the machine and calculate or manipulate its data, and we had to assume that the machine was working properly on the days we used it. In choosing that DCF technique for detecting and measuring reactive oxygen species, we accepted the reports and evidence in the scientific literature that it was an accurate and reliable method for this assay. These sorts of considerations get subsumed within the HDM as auxiliary hypotheses (Hempel, 1966, 23), or what I prefer to call auxiliary assumptions (see Section 4.1). The auxiliary hypotheses for this experiment also included the stability of the HepG2 cell line as we were using it. We did not assume that the cells were free of contamination by *Mycoplasma* species (which might have changed their metabolism and rendered our experiment useless): we actually tested our HepG2 stocks for contamination regularly. In other experiments we had previously shown that HepG2 cells do take up copper when it is added to the medium. What was also critical to the experiment—and receives rather little recognition in formal accounts of the HDM—is the supportive role of broad background knowledge relating to such issues as the extent of fidelity of the HepG2 cell to a normal hepatocyte in terms of cell physiology, the

chemistry of copper and zinc, and the cellular pathology of Wilson disease.[148] All of these 'behind-the-scenes' factors contributed to the reliability of the experimental results in this hypothesis-driven experiment.

HDM depends heavily on technical and technological reliability for producing scientific knowledge. If the experiment fails to substantiate the test-implication, then auxiliary hypotheses can be scrutinized for error before the hypothesis is abandoned. Perhaps the fluorimeter was malfunctioning or maybe the technician mislabeled the plates; maybe the DCFH-DA was not prepared freshly enough or perhaps it was exposed inadvertently to light. In terms of the logic of the HDM, if a major premise in the deductive inference is false or one of the auxiliary hypotheses is proven faulty, then the test-implication (and accordingly the hypothesis) cannot be accepted as 'true' and might be rejected as false, despite actually supporting the truth of the hypothesis. The inferential reliability of deductive reasoning is a clear strength of the HDM. If the content of any premise is faulty, then the advantage of that logical reliability is lost. With HDM technical reliability is effectively taken for granted. In this respect the HDM and OES share common ground epistemically. With its highly focused experimental design, the HDM can generate secondary tests for reliability of the findings, namely, that the same laboratory can reproduce its own findings (repeatability) and so can other laboratories following the same methodology (reproducibility). This is not possible for the OES because the system governing its experimental design is highly dynamic and frequently includes redundant mechanisms, that is, more than one way to achieve the same outcome in the system's global function. Therefore examining the nature of reliable experimental data in the OES will take us in a different direction from that of the HDM, mainly because compared to the HDM, the OES is much more complicated. When we abandon the hypothesis as what is driving experimental design, we give up the HDM and we also give up the Hempelian mindset that the hypothesis is a little removed from

[148] I believe this extensive background information would be relegated to the context of discovery. However, without this information the hypothesis, and subsequently the test-implication, could not be fashioned intelligently. Consequently I regard this information as doing the same sort of work as the auxiliary hypotheses. This may be yet another example of how the context of discovery versus context of justification dichotomy breaks down in 'real-life' biological/biomedical experimental research.

real-life where experiments are performed (that is where the test-implication operates) and we also give up the orderly logical assessment of experimental process. What we gain is the opportunity for a naturalized epistemology relating to system-driven experimentation. Examining how proteomics (or other omics) experiments are actually done helps sort out this 'different direction' for the OES.

7.2.2 Classic reliabilism as a resource for developing the OES as the scientific epistemology for omics

Given the importance of reliable empirical data from the experiments of biological/biomedical research and the extensive analysis in the preceding chapter to technical standards of data acquisition in omics (mainly proteomics), it can hardly be surprising that I would consider reliabilism as a key contender for the role of 'touchstone' general epistemology for the OES. According to reliabilism, justification of a belief depends how reliable the process was which led to its formation. However, I regard this criterion, if transferred superficially into a scientific epistemology, as being too weak for my purposes. I will argue that with the OES as a scientific epistemology the dependability of the competently-produced data is evaluated in large part by determining how the data fit into the system under investigation.

7.2.2.1 Goldman's process reliabilism in relation to the OES

According to Alvin Goldman, it was actually Frank Ramsey who first proposed in 1931 a kind of reliabilism in asserting that "a belief is knowledge if it is true, certain and obtained by a reliable process" (Goldman, 2012, 69). Ramsey's contention seems to have gotten lost in the shuffle, and development of classic reliabilism later in the 20th century appears to have been independent. As Jonathan Vogel points out, classic reliabilism has several different aspects: reliabilism furnishes accounts of knowledge that can be based either on reliable belief versus reliable process (both turn on being able to specify real or potential extent of error) or reliabilism provides a theory of justification. The latter effort "seeks to explicate justification (rather than knowledge) in terms of reliability" (Vogel, 2000, 602-3). I am interested mainly in 'reliabilism' as styled by Goldman: "narrowly to refer to process reliabilism about justification" (Goldman, 2012, 68). However, this wording seems somewhat obscure. Berit Brogaard's

description of reliabilism as a general theory of knowledge is more straightforward to import into a consideration for a scientific epistemology: “knowledge is true belief produced by reliable belief-forming processes or faculties” (Brogaard, 2006, 335).

Goldman’s main argument is that “the justificational status of a belief must somehow depend on the way the belief is caused or causally sustained” (Goldman, 2012, 72), based on some simple but nevertheless subtle examples (Goldman, 2012, 72-3) which serve to eliminate non-critical or wishful thinking and a previous stipulation against circularity whereby the justifying elements cannot be of justificatory sort themselves or anything epistemically close to justification. (Truth-value and causal-relation do not get eliminated.) For Goldman, the kinds of processes which confer justification are mental processes such as those related to perception, memory, sound inference, and introspection. Reliability is quantified in a commonsensical fashion by determining what proportion of beliefs generated by that process are true. Thus some external standards for determining accuracy versus error appear to be required. Likewise Goldman styles this as a historical theory.[149] Goldman argues that as a theory of knowledge, this historical character is a strength because he holds that the “traditional notion that justifiedness arises exclusively from one's momentary mental states has always been problematic” (Goldman, 2012, 73). The final piece of this argument is that there must be no countervailing outcome of another relevant reliable process (Goldman, 2012, 74).

Of course, this is very different epistemological territory compared to that of high-throughput omics experimentation. Mental processes conferring justification are very different from the sorts of process I would cite as justificatory. However, given this background, I might argue for the OES that we hold data arising from an experiment as being justified if the way they were generated was reliable. In this situation ‘reliable’ involves that the data are of the highest possible quality—accurate, really “good” data. We get really good data by having really good analytical apparatus and methodology.

[149] Put crassly, in effect Goldman subscribes to the rule (championed and frequently quoted by scientists) of ‘garbage in/garbage out’. Historicity signifies both that how the knowledge came to be was important and that knowledge has a developmental story, as opposed to be instantaneous.

(For proteomics that analytical apparatus includes the bioinformatics apparatus as well as the mass spectrometer and all the associated equipment.) These are definitely external producers of knowledge, clearly separate from and outside of the person (researcher) doing the knowing. The analogy generates quite a narrow view of how the empirical data (where 'data' in this analogical analysis are *very roughly* equivalent to beliefs) are caused: they are caused by the methods/equipment which produce them. We know how good the equipment is because we can calibrate it or put it through practical quality-assessment exercises as previously described for proteomics (Bell et al., 2009, 423). Moreover, we attempt to quantify the risk of being in error by statistical analysis. Thus we can implement various sorts of tests or measures of accuracy. To this extent we seem to be developing in parallel to process reliabilism, arguably without doing much violence to the more general epistemological theory.

However, process reliabilism as a general epistemology serves us only as a touchstone or inspiration. Apart from the fundamental difference between Goldman's mental processes conferring justification, compared to the more concrete sorts of processes we would cite as justificatory, where we seem to get into trouble is with the 'causal relation' to which Goldman attaches so much importance. Is it adequate to say that the applied methodology and the pertinent analytical/interpretative equipment cause the data? It seems not. I contend that we need to recognize that experimental design is also part of what 'causes' these data. This contention finds some connection with Goldman's views. Just as "wishful thinking, confused reasoning, guesswork and hasty generalization" (Goldman, 2012, 73) are instances of defective processes for forming beliefs, inadequate experimental design counts as a defective process in terms of a process-reliabilist scientific epistemology. If the experiment was badly conceived and designed, then the competent technician and his/her state-of-the-art analytical machinery produce data of at least indeterminate value, possibly of no value at all. Stipulating that an important preliminary question relates to the adequacy of experimental design identifies a feature which perhaps attains more importance within a scientific epistemology than it might in a general epistemology. Wishful thinking gets

discounted at the outset. Bad experimental design should get similarly discounted, but it can slip through the critical assessment.[150]

However, in relation to proteomics we have already established that when you want to know about all the proteins in a given system, which has been defined specifically according to the interests of the investigator, then big-picture (top-down) proteomics is an excellent experimental design. It can then be argued that the effectiveness of available experimental methods to execute this general experimental design can be scored quantitatively. Thus perhaps this consideration assumes minimal importance for proteomics. I do not find the resolution here quite so trivial. Even if we work from Berit Brogaard's version of what reliabilism is as a general theory of knowledge, that is, "knowledge is true belief produced by reliable belief-forming processes or faculties" (Brogaard, 2006, 335), we still have some unsolved problems with respect to how proteomics produces new scientific knowledge. The key issues have to do with truth-value and causal-relation: namely, working out within the experimental design what data are true and where the causal relations under investigation come into play.

7.2.2.2 Can we learn something about the OES from the 'value problem'?

It may be helpful to take a closer look at Brogaard's objection to process reliabilism. His objection emanates from the "value problem". According to the value problem, for reliabilism to say that reliably-formed beliefs are valuable because they tend to be true is superfluous: beside the point or, perhaps more charitably, redundant. If a belief is true, it is not more valuable for being reliably produced.[151] Knowledge simply is more valuable than true belief and any competent epistemological theory will account for this greater value. The objection here focuses on whether it matters that we

[150] The HDM has a similar feature but it is somewhat hidden—or built-in. If the hypothesis driving the experimental design is not formulated to be testable or if it demands technical impossibilities (more precision than the equipment can muster or a more stable biological sample than is available), then the HDM breaks down. With the OES we need to take a more explicit and purposive look at experimental design than with the HDM because the OES does not contain as automatic safeguards against this design defect whereas the HDM does, via its deductive apparatus.

[151] Another formulation of this assertion is from Zagzebski: "If the belief is true, it makes no difference if it comes from an unreliable belief-producing source" (Zagzebski, 2003, 13). This appears to me to be a veiled restatement of the context of discovery versus context of justification dichotomy.

can (or cannot) demonstrate how we got to knowledge once we actually have that piece of knowledge. Moreover, in terms of scientific experimentation, we take the directly opposite position to what the objection (based on the value problem) asserts as a critical premise: namely, that the value of the mechanism producing the belief does not transfer to and accrue to the value of the produced belief (Brogaard, 2006, 338). I do actually hold the opposite—that the precision of the equipment invests into the data produced. Data produced by suboptimal technology simply are not as valuable as data produced by highly precise equipment. With experimental data, the situation is somewhat narrower than in a general epistemology. Obtaining accurate data through utilizing best available technology is critically important. This entire argument points to another area of divergence from process reliabilism, but it helps to point out where the OES stands as a scientific epistemology.

The line of argument, by way of objection to process reliabilism, is developed in greater detail by Linda Zagzebski. Her argument further highlights the divergence. Her example has to do with making a cup of espresso coffee. She asserts that not only the reliability of the machine but in fact not anything about the espresso machine is able to make the coffee in the cup any better; moreover, giving the machine credit for producing excellent espresso coffee does not make the coffee taste better either (Zagzebski, 2003, 13). Zagzebski's argument about the machine-product model of belief strikes me as completely off-target for a scientific epistemology, but interesting because it is so specifically connected to a machine. It may simply be an unusually weak argument by analogy, especially if it involves something as subjective as how coffee tastes. Nevertheless, if I must choose between an espresso machine which sometimes but unpredictably produces a *truly* great cup of coffee (allowing that the value of that great individual cup of coffee seems to be located in the cup) and an espresso machine which reliably produces great cups of coffee >95% of the time, I will surely choose the more reliable machine. My basis for this choice is that with the reliable machine I can know pretty well (that is, quantifiably) that the cup of coffee will be tasty. The value of the machine transfers to the value associated with the cup of coffee. Besides, there

might be objective measures available, such as a quantifiable difference in the amount of caffeine extracted into the cup of coffee. Zagzebski decries the facile metaphor that knowledge is like a product of a machine, but of course with omics knowledge generally depends on the product(s) of machine(s) in an important practical way.[152] Thus for the OES the value argument carries little weight as an objection, but it helps to elucidate key components of this reliabilist scientific epistemology. An important point of differentiation between the needs of a general epistemological theory and scientific epistemology is that the scientist wants dependability. The scientist wants the data in hand to be reliable, and s/he also wants the data to be produced tomorrow in a similar experiment to be equally reliable, although s/he may not spell out this latter desire. Dependability might be regarded as an extended version of reliability.

In brief, the value problem as an objection to process reliabilism does not diminish the utility of reliabilism as a resource to the OES. However, it draws attention to how the quality of the machinery (literally) producing the data matters. In other words, it highlights key differences between a general and a scientific epistemology. Technology is important. Experimental design is critical. These factors might be characterized as relating to 'truth value'. Yet, even when the experimental design is entirely adequate and the technological instrumentation the best available, there is a further consideration. The data need to be assessed in relation to what motivates that experimental design. In the case of proteomics, that motivator would be the system itself. Thus I am proposing the possibility that the burden of justification is not limited to the technical accuracy of the analytical equipment (and the expert technicians who make that technology function well) and to the adroitly designed experiment, but it extends to the system which is being investigated. In other words, with proteomics, the

[152] Certain aspects Zagzebski's arguments suggest an interesting parallelism to certain features of the OES as a scientific epistemology. For example, when she writes "if we think of a belief as part of the agent, the belief can get evaluative properties from the features of the agent" (Zagzebski, 2003, 15), we might consider that this could hold for system-driven research if you regard the system as the agent. Although attributing agency to the system is an interesting idea, it might be something of a stretch. The contrarian argument is that in making either a hypothesis or a system drive experimental design, I have in fact endowed each with agency. I reject this objection as it risks unwieldy and unproductive personification.

data have to be assessed in terms of the system. This aspect relates to the issue of causal relations.

7.2.2.3 The matter of causal relations

According to Goldman, causal relations can serve as an admissible source of justification within process reliabilism (Goldman, 2012, 72). We have posited that the analytical apparatus and related technical resources have a causal relationship to the data, as discussed above. We might ask whether any other causal relations might exist. This argument turns on what we think a pattern is.[153] If we hold that characteristics of a system have a causal function in determining patterns associated with that system, then we may attach importance to validating the data within the functioning of the system. Thus, in terms of this proposed scientific reliabilism, I argue that justification requires (or at least revolves around) validation within the system under investigation. Such validation becomes stronger if it employs experimental or investigational techniques which are distinct from those which produced the data in the first place. Distinct techniques avoid the problem of circularity in assessing empirical data. Consequently, as we attempt to take classic process reliabilism as a guide for fashioning a naturalized scientific epistemology based on proteomics (and thus establish the OES as a scientific epistemology), we need to consider the ramifications of validation within the system under investigation.

7.3 Validation within the system

Put simply, my contention, first introduced in Chapter 5 (see Section 5.4.3), is that when a system drives experimental design, it necessarily ends up playing an important role in the validation of the findings, just as a hypothesis plays an important role in the validation of findings from hypothesis-driven research.[154] It seems reasonable to ask, first, where this intuition came from and, secondly, how it exerts a formative influence upon the OES as a scientific epistemology.

[153] I discuss what a pattern is in detail in the next chapter (Section 8.3).
[154] Validation of data is closely connected to knowledge production and thus plays a justificatory role.

7.3.1 'System' in the driver's seat for validation

Since the experience and vision of metalloproteomics, held by myself and my key collaborator, were formative for thinking about validation, I want to position this discussion within that experience, even at the risk of being repetitive, because it bears directly on the OES as a scientific epistemology. In our definition of metalloproteomics we specifically insisted that the mere enumeration of proteins in a system by available techniques was not enough to put together a proper metalloproteome. The identified proteins must be investigated to see whether they possess the key characteristics ascribed to the proteome and also to see how they fit into the system under investigation. Our initial reason for this stipulation was practical, not theoretical (and certainly not "philosophical"). Given our IMAC methodology for capturing those proteins with metal-binding capability, it seemed possible that some of the identified proteins in the Cu-metalloproteome thus determined were "hangers-on": they were loosely bound to another protein which was in fact the copper-binding protein of interest (see Sections 2.5.2.1 and 2.6). Taking a cautious interpretation, we were initially worried about false positive results. Some structural and cell biologists might object that this requirement was too taxing, perhaps even that we were high-handed in demanding it. However, our stipulation indicates that we thought that the findings had to be contextualized in the actual system under consideration in order to be of any value. Thus contextualization was envisioned as a kind of validation activity. Clearly, the epistemological element here was to move from reliably-produced data to knowledge as such, by putting these data through some sort of justification step. The actual justification was to examine how (or indeed whether) the protein actually functioned in the system under investigation. For example, did the protein bind copper? Did intracellular expression of the protein—that is, the amount of protein found in the cell—change when copper concentrations in the cell changed? In fact, this sort of physiological or functional evaluation became even more important when our findings were so novel that they were rejected as entirely improbable according to current dogma and therefore wrong. Consequently the validation manoeuvre shifted from the purpose of excluding 'false positives' to that of including 'true positives'. Philosophically, I see this as a move from excluding error to

affirming actual knowledge production. In terms of how we developed metalloproteomics, this notion of contextual evaluation within the system might be criticized as being an assumption, in the sense of being a preconceived notion or a prejudice. We may have had a predilection for attaching value to validation within the system because of the scientific interests we, a structural biologist and a cell biologist/toxicologist, brought to our collaboration. The way our work developed militates against its being a mere assumption: I think it was an appropriate intuition 'headed in the right direction'.[155] Thus it merits further close investigation.

7.3.2 How far does validation in the system take us?

In order to see how the system might serve to validate experimental data, it may be informative to consider how a hypothesis functions to permit validation of experimental data. Working with a hypothesis, the evaluative stance is to determine whether or not the hypothesis was supported. While it is obviously too strong to say that the hypothesis was proven, the general logic here is broadly dichotomous: the hypothesis is judged as being true or false. (Of course, it cannot be judged as absolutely true or absolutely false, but true or false to the level of stringency previously or consensually postulated as adequate, such as in statistics with $p<0.05$, or with the tacit proviso that it could be judged differently in the future.) A well-articulated hypothesis lends itself to this kind of justificatory move by being specific enough to be evaluated through a well-designed highly pertinent experiment. Then the next hypothesis in the chain of hypotheses is fashioned, and the same apparatus for knowledge production applies.[156]

With system-driven experiments the situation is much different. The investigator has much less control, except for how s/he demarcates and defines the system. The

[155] Subsequently others in the proteomics community have picked up on this concept of contextualization or, perhaps somewhat more likely, figured it out independently. Their realization generates an interesting but discomfiting discussion aimed at cell biologists. They berate the cell biologists for failing to get collaborative with the biochemists in an effort to do this contextualizing efficiently! They write: "It is the cell biologist's role to assure that protein assignments are accurate" (Bell et al., 2007, 784).

[156] A good example of this kind of tight hypothesis-driven experimentation is given in the set of experiments reported by Oswald Avery in his classic paper about a chemical factor transforming the coat of *Pneumococci* (Avery et al., 1944, 144-51).

system gets examined—that is to say, characterized in some inclusive fashion (all the proteins or all the copper-associated proteins or, with other subdisciplines of proteomics, all the phosphoproteins or glycoproteins)—and then a validation process is employed. This situation is quite different from the more straightforward situation with hypothesis-driven research and consequently the justificatory moves are different. No aspect of the system or entity in the data set serves as the unique or invariable starting-point for validation. S/he starts somewhere: that starting-point may be the most pertinent to the data at hand or the technically easiest or the most informative aspect of the system.[157] So instead of being narrow-focus (bottom-up) or big-picture (top-down) or even 'middle-out', it is something else: I will call it "infill with zoom". With the new omics data available, such as a proteome, you zoom in on one aspect of the system in some detail and look at that with the proteomics data and then perhaps turn your attention to another facet of that system with the data.[158] For efficiency, you might choose one or two features of the system under investigation which serve as *sine qua non* features in the sense that if the protein(s) of interest lack(s) this functionality then in all likelihood they are not part of the system and require(s) no further

[157] Admittedly an element of real or apparent caprice may operate here. Where the researcher starts the validation of the data set may depend what feature of the data catches his/her fancy. More objectively, however, practical details are important for such decisions: the techniques available to you, your comprehension of the system you are studying, and your priority issues to be investigated. For example, Carmona et al. report what is in effect a contextualizing study of representative compounds in a Mn-metalloproteome (Carmona et al., 2014, 823): they investigate cell toxicity, intracellular chemistry of the compounds, and intracellular disposition of the compounds—partly because these investigations extend their previous work and utilize their technological expertise and partly because the research is relevant to human disease (Parkinson disease) and major environmental toxicity.

[158] It is possible that my notion reflects some subliminal influence of William Wimsatt in his analysis of levels of organization in entities subject to scientific investigation (Wimsatt, 2007, 205-6) which relates to the notion of zooming up and down through those levels. I will acknowledge this although it was not specifically in mind as I was writing. However, my usage of 'zoom' is not quite the same. More to the point, the concept of zooming-in appears in an essay by Robert Richardson and Achim Stephan in relation to philosophical aspects of systems biology, one of the few books available on the topic in the past 8-10 years. They write: "There are typically various grades of resolution possible in the way we describe the system. These can be thought of as focusing on more or less detail, as if we could zoom in or out on the system revealing more or less about what is happening within the system. … This is a matter of redescription rather than reduction" (Richardson and Stephan, 2007, 131). This is similar and yet again not the same as my very practical concept. Likewise, while 'infill with zoom' has strong similarities to "middle-out", it differs because the real starting point for the research endeavor is the system seen from a big-picture (top-down) vantage point. The "middle-out" segment of the strategy is secondary.

investigation.[159] Importantly, with validation in the system, justification is no longer characterized by dichotomous descriptors (T/F). It seems to require three categories so that judgment can be reserved when results are not clear-cut. The descriptors might be 'works'—'doesn't work'—'don't know'. ("Work" is a purposely vague term of functionality by which I mean 'contributing to the proper functioning of the system': I will elaborate on this word choice later in this chapter, Section 7.3.3.) Proteins found through this system-driven proteomics exercise which show functionality within the system 'work', and those which fail to show functionality 'don't work'. Sometimes the outcome is not clear cut, and then instead of discarding such data as irrelevant, these proteins for which the outcome is not clear cut need to be kept in reserve for further evaluation when the system is better elucidated. So what we see with system-driven research is that the justificatory move consists in the data coming together as a coherent pattern of functioning within the system under investigation.

There is one objection to this approach which arises from the scientific, not the philosophical, community. It deserves some comment. A toxic system resists functional analysis. The toxicity is internal to the working(s) of the system. It is not simply that the system is unpleasant or awkward to work with in the laboratory. (Such a system might be not chosen as a good experimental model in the first place.) The internal toxicity is due to a characteristic of the system whereby it alters or functionally dismantles the protein or other entity which is being contextualized. It seems pretty obvious that my validation strategy will be frustrated in such a system, and I have no easy remedy. However, I do not think that the example of the toxic system (which after all may be mainly theoretical[160]) invalidates the strategy.

[159] For example, a protein identified in the hepatocellular Cu-metalloproteome which, on validation, cannot be demonstrated to bind copper is eliminated from further investigation. However, I say "in all likelihood" because we have enough knowledge of proteins in systems now to recognize that sometimes functionality occurs through partnerships or small groupings of proteins. This was perhaps not so well-recognized when we started doing metalloproteomics.

[160] An example of a toxic system might a knockout mouse where the knocking out of a specific gene proves to be embryolethal. No live progeny are available for further study. Experimental designs might nevertheless be developed to gain some data from this situation. (So far as I know, this is my own idea to suggest this example.)

7.3.3 Mechanics of a three-descriptor evaluative approach

This same process of justification described above for *system* biology can operate in more complex aspects of systems biology. As previously mentioned, a more distal objective of omics is to contribute the requisite data to an important project of systems biology, namely, to look at the system from various angles and then combine the observations for a highly textured account of the system. Here the pattern detection is more complicated, but it remains essentially the same process of justification. This broad objective is to develop a comprehensive picture of the system.[161] In this section I will attempt to sketch out some of the main features of how this OES evaluative machinery for validation in the system under examination might operate.

The evaluative process here turns on a few concepts which I have specified repeatedly but nonetheless end up playing an implicit role. A major consideration is that the OES seeks to be inclusive and yet open-ended: therefore, the way fitting into a system is evaluated will differ from that the HDM evaluative machinery in relation to the governing hypothesis. Taking proteomics as the example, we immediately encounter two basic questions. First, does the entity we have identified as being in the proteome of interest actually belong in that proteome? Secondly, can we show that this entity, now confirmed as being in the proteome of interest, constructively participates in the system?

The first question seems picky and boringly commonsensical. However, there is little point in trying to contextualize a protein in the biological system being studied if that protein does not belong in the proteome in the first place. This may be trivial for some general proteomes, but it is very apposite for a specialized subdiscipline of proteomics like metalloproteomics. For example, if we set out to determine the Cu-metalloproteome of some biological system, we needed to show that the proteins we identified were in fact capable of being associated with copper and thus belonged in a Cu-metalloproteome. We could do this by showing that they had structural features

[161] The paper on the classic systems biology analysis of *Mycoplasma genitalum* (despite its scientific shortcomings) is a good example of this kind of effort (Karr et al., 2012, 389).

typical of copper-binding capability or by citing previous scientific literature confirming the copper-binding capability of a particular protein or by actually doing the assays in the laboratory and demonstrating copper-binding. This last approach generally provided the most extensive information and was reserved for contentious (or interesting) members of the Cu-metalloproteome. Thus the first step in validation is to confirm that each protein actually belongs in the proteome under development in the sense that it is a protein and that it has specialized features if those are being sought. However, what is fairly straightforward for proteins might not be quite so simple for small metabolically active chemicals, for example, in a metabolome. It is also not very easy if the functional entity of interest proves to be an aggregate—but these are technical, rather than theoretical, considerations.

Again taking proteomics as the example, the second question relates to showing the proteins identified as being in the proteome of interest actually fit[162] into the system under investigation. The implication is that just producing an inventory, or what I would call a "shopping-list", is not adequate. In system-driven research the ultimate goal is to understand the system better. Thus, contextualization of the resultant proteome in the system is required. If the system is conceived of as interlocking parts, then contextualizing the findings involves seeing how these proteins fit together in the system. If it is conceived of as a network, then the focus is on how they may interact. For some systems it may be adequate just to show that the protein is functional in that system in the most fundamental sense: it is active, *does something*, in the system being studied. With a view to being inclusive and open-ended, preconceived notions of what the protein does might be discounted. Sometimes it might not be entirely evident what the protein under consideration actually does in the system being studied, or data might be incomplete.

[162] The notion of 'fit' in the sense of 'fitting together' or forming a coherent system certainly has conceptual connections with the Cummins causal role concept of function (Amundson and Lauder, 1994, 447-8). However, the examples of Amundson and Lauder tend to be anatomical whereas those under consideration here tend to be physiological.

In order to be inclusive, I have suggested that the evaluative descriptors are three-fold: broadly speaking, 'yes', 'no', and 'don't know'. This formulation is problematic for a number of reasons. First of all, it is a functional evaluation, not a veridical one. We have much experience with veridical evaluations and their descriptors (yes/no; true/false), but that experience, which is generally relevant to assessment of propositions, may not apply here. Using functional descriptors here takes the analysis into non-propositional territory—and probably rightly so for proteomics. Choosing the right descriptors is difficult. However, it is not so very difficult to figure out the eventual positive outcome. Just as a desirable outcome for a proposition is that it be true, a desirable outcome for a system is that it comes together as an integrated whole. I have already proposed (in Section 7.3.2) that the descriptors might be 'works'—'doesn't work'—'don't know'. Proteins identified as being in some proteome for some system are evaluated physiologically as whether they are doing something in the system. Those which show functionality within the system 'work', and those which fail 'don't work'; where it is impossible to tell, the 'don't know' descriptor is applied. The big problem with this approach is that the concept of 'working' is too vague.[163] One response is that 'works' is shorthand for 'usefully or constructively participates in the system'. A virus or parasite functions in a system but it acts to the detriment or destruction of the system. A rogue protein might be identified that has the same effect: it would not qualify for the assessment "works" even though it had an effect (but an *adverse* effect) in the system. But even this kind of restriction is dicey. Cells, for example, have mechanisms for trashing malformed proteins or damaged organelles; programmed cell death (apoptosis) is a feature of some normal cellular functions such as organ development, although of course it results in the death of the particular cell.

In order to explore this kind of functional evaluative machinery, I suggest considering a very simple model of a system: a conventional jigsaw picture puzzle. The analogy has drawbacks: like a word jumble and various other puzzles, it has one unique

[163] I am grateful to Peter Schotch for raising strenuous objections about my evaluative formulations (July 2014).

solution. Biological systems display redundancy, more than one solution to any problem. However, one advantage of the jigsaw puzzle example is that it has a definite positive endpoint in that the picture comes together as a coherent whole, when all the pieces are in place. Moreover, with respect to any piece of the puzzle, we can articulate simple functional descriptors: 'fits', 'doesn't fit', and 'don't know (yet)'. An advantage of the concept of *fit* in relation to a picture puzzle is that we can define exactly what we mean (not so with "works", as we have discussed). To fit means that an internal piece has to interlock with all adjacent pieces and a border piece has to interlock with pieces above and below it, keeping the straight edge in rank with others around it.

I believe that most of us work a picture puzzle the same way: we find all the pieces with one straight edge and arrange them into the rectangular border.[164] Then we fill in the centre. Some parts may be relatively easy: a group of houses or a collection of flowers. We may put the several pieces together for such a part and leave it unattached to the framing border, just placed roughly where we think it will end up. (Note that this strategy is basically *in-fill and zoom*.) For any piece of the puzzle, we try to interlock it with another piece and ask whether it fits or not. Sometimes with very odd knobby shapes the pieces look as if they fit together but not quite because there is abnormal tension in the join or somehow the bits of picture on the two pieces do not appear to belong together: this gets rated as 'don't know (yet)'. Pieces that fit together properly are seen to make a constructive contribution because the whole picture starts to take shape.

This is the sort of inclusive/open-ended evaluative system I have in mind for the contextualization of proteins within the system being studied. If the descriptors were limited to 'yes' and 'no' types of classification, then important findings might be discounted prematurely simply because relevant data were incomplete or the "picture"

[164] Keeners, or those solving a puzzle bought at a rummage sale, may first dump all the pieces out on the table and turn them upside down to ascertain that all pieces have the same back-paper design and thus belong to this puzzle. Extreme keeners will count the pieces to be certain that the puzzle is complete. This is not so different from ascertaining that everything in the candidate proteome belongs there or trying to assess the completeness of the proteomics experiment itself. Also it is clear that the most distinctive pieces in the puzzle are those with *two* straight edges at right angles, namely the four corners.

of the system has not yet come together adequately. With the picture puzzle, we end up with 'fits'—'doesn't fit'—'don't know (yet) whether it fits' as descriptors. (Similarly, for validation which involves putting together the picture of the system, descriptors might be 'works'—'doesn't work'—'don't know (yet)'.) I am not suggesting that this is a 'logic', not even a rudimentary one. It is nothing more than an examination of how we actually think about solving a picture puzzle.

Allowing for the inherent inadequacies of arguments by analogy, I suggest that our ordinary experience of picture puzzles provides some insight into contextualization of proteomics data in a biological system. The system is chosen and defined with some background information available. (Pertinent specific information is captured in the meticulous description of the system.) This serves as the framework within which the candidate proteome is contextualized. Then each protein's function is characterized within the system to see whether it interacts with other components of the system and whether it somehow promotes the proper functioning of the system, or at least what we believe that proper functioning to be. To retain the property of being open-ended, the researcher must be prepared to apply the descriptor 'don't know (yet)' liberally.

There is one more feature of solving a picture puzzle which is of interest to this contextualization process. It is not unique to picture puzzles. This is the phenomenon that as you get toward the end of putting the puzzle together it goes faster. I called this "epistemic festination" where festination is the typical gait abnormality of Parkinson disease, characterized by a slow steps at first and then progressively rapid forward movement. Somewhere along the line, the puzzle really starts to come together, the solution becomes easier, and the picture just falls together. I submit that a similar phenomenon holds with figuring out a biological system. Somewhere along the line the overall character of the system under consideration starts to become more obvious.

7.4 OES as a scientific epistemology for omics

Classic reliabilism, from which I have drawn inspiration, retains the analytical character of dealing with propositions. With system-driven research, we are seeking validation for knowledge gained through detection of patterns. The analytical approach

furnished by reliabilism may be entirely wrong-headed. Indeed, if you are looking at data which form patterns, either as colour designs or networks, then any sort of proposition-based epistemology is at least awkward, if not irrelevant. Unlike the HDM which is anchored on propositional formulations and logic, the OES must deal with non-propositional formulations. What we need is something entirely different from conventional analytic approaches to knowledge production: namely, a thoroughly non-propositional strategy for justification.

7.4.1 A broader view of justification

Certain aspects of justification as analyzed by Alvin Goldman might still be helpful for addressing this problem. Goldman points out that justification is by its very nature a matter of evaluation, a "term of appraisal" (Goldman, 1979, 1). He sticks to his original claim, namely, that the justificational status of a belief is a function of the reliability of the process producing it. However, Goldman's actual language appeals to causal relation: "the reliability of the process or processes that cause it" (Goldman, 1979, 10). It may be that emphasis on causation can be overplayed since some of his examples of reliable processes appear to me to be essentially matters of experimental design. The incidental suggestion from Goldman's work that experimental design could be a factor determining reliability is highly relevant to our project here. Another aspect where my project for the OES is fully congruent with Goldman's project is that how the belief was generated matters: both are historical. However, in Goldman's general epistemology the processes producing beliefs justified by reliability differ radically from those we encounter with omics: examples include reasoning, memory and perceptual processes (Goldman, 1979, 10). For the OES as a scientific epistemology we need to appeal to something more specialized than cognitive belief-forming processes, even if we regard analytic equipment (mass spectrometers and bioinformatics suites) as directly subservient to—or extensions of—such cognitive processes.

We may find some guidance is Helen Longino's examination of scientific methodology, although she is working well within an HDM universe of discourse. She identifies observation and justificatory reasoning as central components of scientific knowledge (as I do), though she argues that they are essentially social (Longino, 1994,

139). She regards observation as a specialized epistemic activity: "an organized sensory encounter that registers what is perceived" (Longino, 1994, 140). It "involve[s] classification and categorization that notes similarity and dissimilarity with other items of interest" (Longino, 1994, 140). Likewise scientific reasoning is not mere calculation. Longino describes it as a "practice of challenge and response: challenge to a claim is met by the offering of reasons to believe it, which reasons can then be challenged on the grounds both of truth and of relevance, provoking additional reasoning" (Longino, 1994, 141). Longino's broad argument is that these epistemic activities of observation and reasoning are socially-programmed at the individual level and then submitted to further testing and refinement through the disseminated collegial research community. The OES, like the HDM, operates at a stage prior to dissemination. I submit that the core epistemic process of the OES with respect to contextualization is similar to what Longino describes: challenge and response. The researcher looking at the entities in a proteome, for example, asks questions like: does this protein work in this system? How? How can I show it does? What characteristics must it have to work in the system? The OES is sensitive to vantage-point of the researcher including his/her previous experience and active ulterior interests. Pattern detection is susceptible to the limitations of prior knowledge and bias (see Section 8.5), but it is more up-front about this problem than the HDM. However, the OES also includes safeguards: criteria for producing reliable empirical data detailed above.

When we consider how the system itself serves to validate experimental data and locate this role in a non-propositional assessment, Longino's comments about models, within this account of scientific knowledge, provide some useful directions.[165]

[165] Immediately, however, I must disengage from the general thrust of her argument, which attempts to discuss theories in light of models. I do not wish to invite inferences along the lines of "systems = theories" or "patterns = models" or "investigated systems = models". (This is an incomplete list: there might be other similar inferences to be drawn.)The critical connection here is that systems and (many) models are non-propositional. It is possible that I would draw a picture of the copper-handling system of a hepatocytes based on my Cu-metalloproteome, but I might not. I might carry it around as a vague and shifting mental image. I would be unlikely to carry it around as a mental or written list of propositions. I do suspect, however, that if Longino bothered to specify what she means by 'system' in this discussion, that her description of a system vis-à-vis a biological entity would be very close to my notion of system in this dissertation.

Longino holds that a model is neither true nor false: its validity is proportional to how isomorphic it is to the thing in the world being modeled. Importantly, "this isomorphism permits the mapping of relations and structures and processes of the model onto some portion of the world" (Longino, 1994, 147). When we take omics data and contextualize it in the system, we are basically asking whether the picture developing out of these data is isomorphic with the system as we understand it. This is perhaps a tighter, more technical definition of 'works'. Longino's specification of relations, structures, and processes captures key considerations in this system-based evaluative process, and as we will see, also resonates with my definition of a pattern.[166] Contextualizing the data in the system serves as way to establish empirical adequacy. Thus the system governs experimental design and subsequently constrains validity of the data. Unlike the hypothesis, which also exerts these functions, the system allows for greater scope, constructive uncertainty, even dalliance with different versions of the emerging picture or patterns, and indeed emergence of entirely unexpected solutions mainly because it is complex and redundant.

7.4.2 How would the OES as a scientific epistemology work?

To keep this discussion relatively uncomplicated, I will limit it to proteomics. The brief outline of how the OES works for proteomics as a scientific epistemology must include the following elements, some of which have already been discussed: (1) given the suitably chosen and defined system, data regarding component proteins are reliably produced, that is, with maximum available technical accuracy (as discussed in detail in Chapter 6); (2) these data are examined closely and pattern(s) detected (to be discussed in Chapter 8); finally, (3) the patterns are then evaluated in the context of the system. An important component of this epistemology is highly competent experimental design, not merely excellent technical analytic resources. Moreover, the process of pattern detection needs clear definition. Finally, to avoid inherent circularity, evaluation in the

[166] See Section 8.3. This conceptual similarity is borne out by her observation later in the same paper: "Much of what we call 'scientific knowledge' has as its object not singular propositions about mid-sized objects, but patterns and relationships among entities beyond our powers of sensory detection and between such entities and the objects and relationships we do detect" (Longino, 1994, 152).

context of the system needs to involve techniques which are separate from the analytical techniques of proteomics. The open-ended scheme which involves characterizing findings by simple functional descriptors: 'works', 'doesn't work', and 'don't know (yet)' operates mainly in the contextual evaluation, although it might operate in the process of detecting nascent patterns.

As just described, this strategy seems rather spare or rudimentary. Let's see how this works for a now familiar example within this dissertation, namely, PDI as a copper-binding protein. We were clear about the system under investigation: the cellular apparatus for handling copper in hepatocytes. We regarded this as the hepatocellular Cu-metalloproteome. For this initial work on determining a hepatocellular Cu-metalloproteome, we made samples from continuous human hepatocyte lines such as HepG2, which I had been using in my laboratory for years. As a key component of the experimental design, it was routine to describe the cell cultures, media and methods of preparation of the 'cell prep'. Then we used available state-of-the-art methods and instruments for the protein separation and mass spectrometry (with the novel initial step of passing the cell prep over the copper-charged IMAC column): all of this could be, and indeed was, described in the sort of detail which later became routine for proteomics. All this comprises step (1) above. For step (2), in this initial effort pattern detection was not on the grand scale of categorizing proteins found. That was a later refinement (for example, in our next paper). In this first work on metalloproteomes, the number of proteins found was comparatively small, and pattern detection was mainly focused on finding patterns within the proteins of the putative metalloproteome to see whether they were likely to belong in the metalloproteome. This amounted to looking for primary sequence motifs rich in the particular amino acids which tend to bind metals. This type of pattern detection does not immediately provide a rationale for investigating PDI. In fact, why we examined PDI in greater detail is never spelled out. It is an example of not being able to specify the exact choice of starting-point for 'in-fill and zoom'. One possibility is that it was surprising to find such a prominent hepatocellular protein on the list. Another possibility, maybe secondary, is that the resources were

available to study it: specifically, a cDNA was available so that the protein could be expressed and functionally characterized. In any case, these experiments were typical of step (3) and formed the basis for a separate paper. Detailed experiments showed that PDI could bind Cu^{+}, the form of copper found in cells. We needed a series of several different kinds of experiments to move confidently from 'don't know (yet)' to 'works' where 'works' means here "belongs to a set of copper-binding proteins found in hepatocytes". The similarities of structure between PDI and another known copper-binding protein, ATOX1, were pointed out.[167]

Finally, there is a further, forward-looking step in this analysis. It might be called 'creative speculation' or a 'how possibly' explanation. In what might be regarded as step (3+), evaluating PDI in the context of the system led to speculating about what it might do in the system. Because of its location within the endoplasmic reticulum and Golgi apparatus, the general vicinity of where the Wilson ATPase is found, where protein manufacture takes place, one possibility was that PDI did something to deliver copper to nascent copper-bearing metalloproteins such as ceruloplasmin (Narindrasorasak et al., 2003, 412). Another possibility from my clinical perspective was that it might have something to do with fat accumulation in hepatocytes in Wilson disease. PDI forms part of the heterodimeric protein MTTP (microsomal triglyceride transfer protein), which is involved in triglyceride production in hepatocytes. I speculated that the surfeit of copper in hepatocytes in Wilson disease might bind to the PDI component of MTTP and thus interfere with proper disposition of triglyceride. We actually never got to test this idea, but it is typical of the 'idea-flow' of proteomics. In this case, the metalloproteomics experiment enriched our understanding of the functional repertoire of a major hepatocellular protein and also generated some new hypotheses. We cannot say that PDI functions either to deliver copper to apoproteins or to account for some of the

[167] PDI was also shown to be capable of binding silver, chemically not surprising because silver also has valence of 1, like Cu^{+}. Interestingly, PDI was shown not to bind zinc in these experiments (Narindrasorasak et al., 2003, 408). We did not comment much on this finding, but it is noteworthy because PDI also showed up in our candidate Zn-metalloproteome but now would be excluded from the hepatocellular Zn-metalloproteome as such these studies. In effect, it may have been accurate to find PDI in the Zn-metalloproteome, but that finding was not validated.

pathophysiology of Wilson disease, but it is significant that we can easily imagine PDI fitting into a functional network relating to hepatocellular copper disposition in either of these ways.

7.4.3 Prominent role for pattern detection

Faced with a plethora of data relating to a biological system, or even a rather small proteome for some system, the researcher has to do something to come to grips with the experimental data. In general, these data either present themselves as colour patterns or can be classified into groups. Sometimes a statistical procedure like hierarchical cluster analysis is employed to identify patterns in experimental data. This stage of seeking patterns takes place after the raw data have been converted from physicochemical measurements: a large degree of data-processing has taken place. Finding patterns is a very different epistemic activity from the deliberate task of comparing what a hypothesis predicted and what was found. Pattern-detection is at the heart of the OES. It is another distinctive feature, distinguishing it from the HDM. I will turn to an extensive consideration of what patterns are and how we think about them—and with them—in the next chapter.

CHAPTER 8 PATTERNS AND PATTERN DETECTION

In the previous chapter I demonstrated how the OES serves as a scientific epistemology. Certain elements of this scientific epistemology are critical. These include the system itself, experimental design and technical/technological aspects. The resulting empirical data are submitted to identify pattern(s) within the data set, and finally, identified patterns are contextualized in the system under investigation. A key aspect of producing scientific knowledge via the OES involves an epistemic manoeuvre which can be described as "thinking by looking". It is permitting, or coaxing, data to resolve into (a) pattern(s), which then can be evaluated by being contextualized within the system under investigation. I call this process of seeing patterns 'pattern detection'. The customary term is 'pattern recognition'; however, I believe this term often involves the assumption that this is only one pattern to recognize. I am modifying the term in order to avoid this potential assumption. I do not want to exclude the possibility that there is more than one pattern to be found in omics data. Epistemically, this commitment follows directly from wanting to keep the evaluation as open-ended as possible mainly because open-endedness is an important characteristic of the OES as a scientific epistemology.

Although I have alluded to pattern detection, I have not yet considered it in any detail. That is the agenda for this chapter. First, I will illustrate the sorts of patterns we find with omics because it may be difficult to think about patterns in omics without having visual images to consult. Then I will review what has previously been said about patterns by philosophers of science. That literature focuses almost exclusively on the physical sciences. Yet no one would deny that pattern detection plays an important role in the biological sciences. Thus after considering some extremely basic questions about patterns as they relate to biological issues, I will explore in greater detail how pattern detection operates in proteomics. Fascination with patterns is characteristic of both science and the creative arts, two domains often regarded as highly distinct. Yet

patterns as such have received comparatively little concerted philosophical attention. One broad objective for this chapter is to begin to redress that omission.

8.1 Real-life patterns relating to metalloproteomics

An immediate difficulty encountered in talking about patterns is that there are numerous kinds of patterns. Furthermore, as we will see, a problem with talking about patterns in relation to omics is that the relationship between omics and patterns is highly convoluted. Patterns are embedded in omics operations: the very performance of an omics experiment involves pattern identification within the analytical machinery and as a read-out of the experimental findings. This is certainly the case with proteomics, but the same complication occurs with most omics disciplines. Moreover, more than one kind of pattern is relevant to this analysis, depending on the particular omics discipline or the experimental design. Therefore my analysis here is step-wise, starting with actual patterns.

Among the various sorts of patterns we might find with omics, three general types immediately stand out. In the first case, we find data sets where the read-out of the experiment is a colour pattern. It may be a checkerboard of colours, or a graded spectrum resembling a weather map. The second broad category is a classification of results. The data, initially just listed by name, are grouped by some parameter: the structural type of protein or the functional characteristic or the location in the cell or tissue from which the preparation was made. Sometimes these findings are further portrayed as a graphic diagram such as a bar graph or pie chart. Finally, the third type of pattern we often see is a picture (or diagram) showing which proteins interact within a system. This type of pattern has strong affinities to network biology. Other types of patterns are possible, although not so importantly in my personal experience with metalloproteomics. For example, we can draw a picture of how multiple proteins bind to part of a gene and to each other to regulate gene expression. This is also a pattern of interactions among biological entities. We can create a highly stylized picture of protein or other biological molecule, a kind of pattern representation which is very important

with structure-driven research but will not be considered here in relation to system-driven research.

8.1.1 Microarray (some sort of 'chip': "gene", protein, proteome)

Patterns found with a microarray chip have visual immediacy. In fact, they are the empirical 'raw data' from the experiment although some data transformation may have taken place. Such processing is done automatically by the machine producing the data. This (Figure 8.1) is part of a figure from Muller et al. showing results of a DNA microarray analysis in HepG2 cells incubated with copper over a 24-hour period (sampled at intervals from 1.5-24 hours) (Muller et al., 2007,500). It reveals that the expression of some genes was up-regulated at an early stage and others were down-regulated. Of note in connection to Chapter 6, the authors state that the complete protocols and data were deposited with MIAME: thus the research meets certain specifications of the OES regarding technological adequacy.

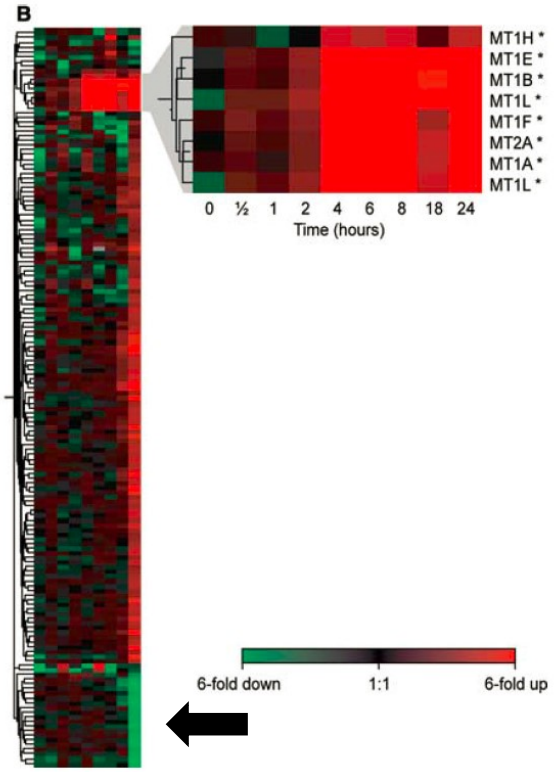

Figure 8.1 DNA microarray analysis of HepG2 cells incubated with copper (adapted from Muller et al., 2007)

The varied colours *are* the data. The one area with numerous up-regulated genes (the red blob in the upper right corner) is shown in detail with the genes identified after the fact. They are all genes for metallothioneins, proteins which sequester copper in the cell. Of interest, there are also some genes which are down regulated (intense green, arrow), and these appear to be COMMD1 and other members of that family.

8.1.2 Classification – 'conceptual' patterns

Some patterns are indirect: they have to be constructed by the investigator who groups or classifies the data. This classification can be displayed as lists or graphically. In our second paper about the Cu-metalloproteome in HepG2 cells, we prepared either cytosol or microsomes from the HepG2 cells (see Section 2.5.2). We identified 67 proteins which were in the Cu-metalloproteome (more precisely called a *candidate* Cu-metalloproteome since they required contextualization); however, 52 of these proteins had putative copper-binding domain on the basis of primary structure patterns. One had a virtual copper-binding site. How proteins appeared (Smith et al., 2004, 836, 838) on the thin layer chromatography plate after 2DE (Figure 8.2) is itself a pattern (cytosol on left, microsomes on right):

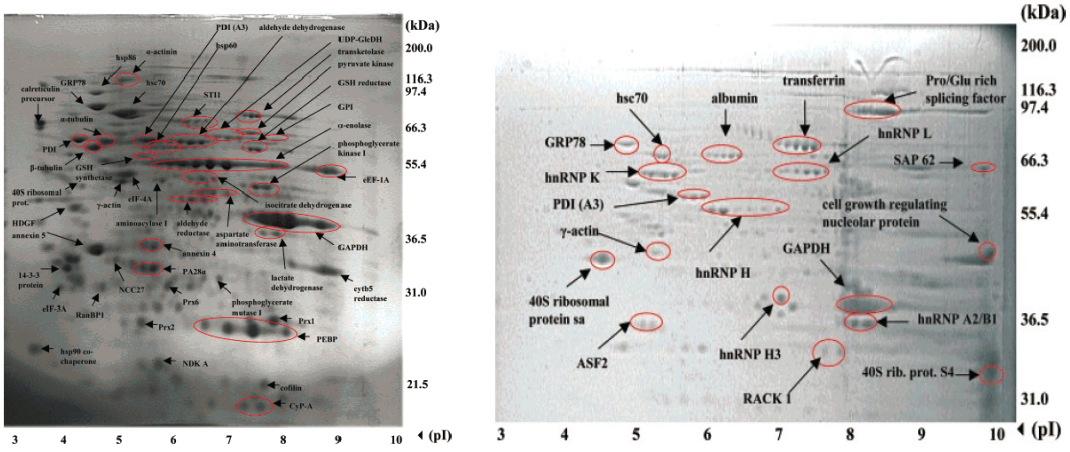

Figure 8.2 2DE maps of HepG2 cytosolic (left) and microsomal (right) proteins showing copper-binding capability (from Smith et al., 2004)

(Of course, the convenient circles and labels were added as part of preparing the figure for publication.) Before mass spectrometry became so critical to proteomics, some biochemists regarded the pattern realized by the distribution of spots on the

chromatography plate as a basis for comparing protein expression in different cells or tissues. Such comparisons were inevitably very general. After identification and classification according to generally accepted categories of types of intracellular proteins, the distribution of proteins found could be depicted (Figure 8.3) graphically (Smith et al., 2004, 838). The actual proteins as identified along with pertinent metadata and specification of metal-binding motifs were listed by name in Tables 1 and 2 in the paper (Smith et al., 2004, 836-7).

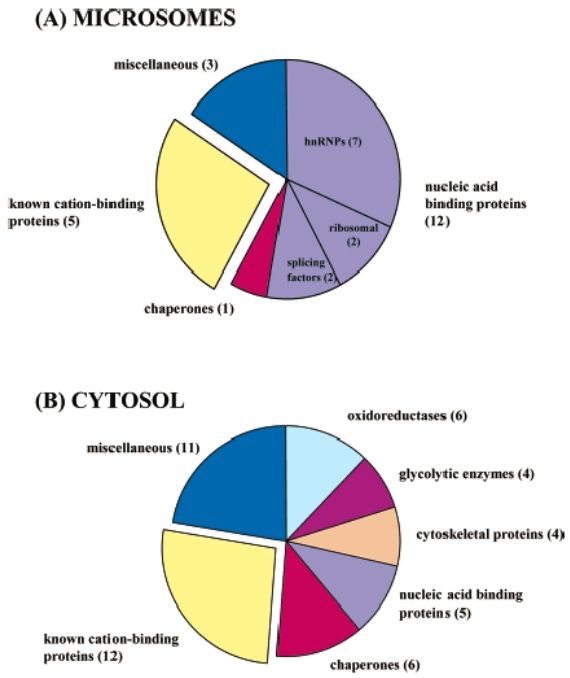

Figure 8.3 Classification of proteins identified in HepG2 microsomes and cytosol (from Smith et al., 2004). This is a depiction of the 'raw' empirical data shown in Figure 8.2 above.

8.1.3 Interactome

Patterns of interaction are important for understanding the action of molecules in organisms and the mechanics of cell physiology. These patterns present a kind of sketch of what interacts with what, although the 'how' of those interactions is rarely depicted.

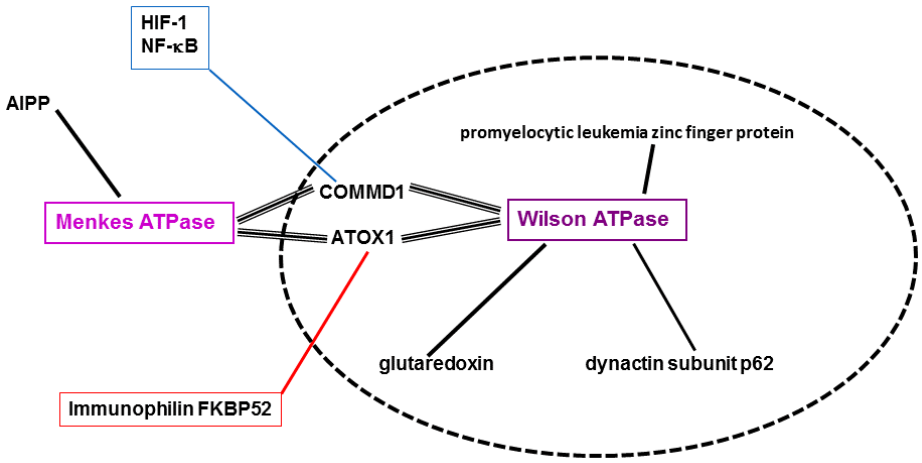

Figure 8.4 Interactomes of the Menkes and Wilson ATPases (modified extensively from de Bie et al., 2007—see text)

This (Figure 8.4) is an early (2008) version of the interactome I will show in a further revised form in Section 8.4.3.3.[168] To construct this figure, I updated a figure proposed by De Bie and colleagues (de Bie et al., 2007, 683) and added further details based on additional published research. It actually shows interactions with the Wilson ATPase and/or the Menkes ATPase as well as proteins which interact with proteins that interact with both ATPases. In fact, the original figure is more than a simple interactome, and my development of that figure departs from the ordinary format of an interactome even further. The 2008 Wilson ATPase interactome is contained within the elliptical dotted line. Once I started adding proteins which interact with those proteins which interact with the ATPase, I started to depict a protein network.

8.1.4 Ontological features of these patterns

If we ask about the relationship between these images which show patterns and the empirical data they depict, we can immediately discern differences among them in terms of that relationship. The microarray pattern grid (Figure 8.1) and the 2DE chromatography plate (Figure 8.2) are the data as such: they are what was found in the system when it was investigated. Moving forward with the metalloproteomics analysis,

[168] In fact, since 2008, one more protein interacting with the Menkes ATPase has been identified: valosin-containing protein, that is, p97/VCP (Yi et al., 2012, 1800). However, I have not included anything about the Menkes ATPase in the figure in section 8.4.3.3. I point this out only for completeness and to indicate that the field continues to move forward.

we use bioinformatics to identify the proteins present. We get a kind of 'shopping list'. One way to deal with these results is to group the proteins identified according to some classification system of interest and then see what patterns are present. The interactome (Figure 8.4) is yet more remote from the raw data: we create a graphic of some sort to organize visually evidence gained from various lines of experimentation with respect to interaction of proteins with some protein or entity of interest.

We might ask where these pictures, which I will consider as real-life examples related to pattern detection, fit into a spectrum of scientific 'pictures' as considered by Linda Perini (Perini, 2005, 273, 275, 272, respectively) and re-assessed by Letitia Meynell (Meynell, 2013, 328-9). Perini's pictures comprise a graph on Cartesian coordinates, a stylized drawing (known to scientists as a "cartoon") of the chemical process of enzyme energetics for bovine heart mitochondrial F_1-ATPase[169], and a transmission electron micrograph (TEM). Clearly, the pictures in Figures 8.1 and 8.2 are most like the TEM because like the TEM they show primary empirical findings. The pie charts in Figure 8.3, like the graph on Cartesian coordinates, summarize a large amount of empirical data efficiently by combining visual icons and terse explanatory text. The graphic of Figure 8.4 most nearly resembles the energetics cartoon: it depicts relationships gained by combining findings from multiple experiments. The accuracy of the graphic can be assessed by referring to the sources they represent. The depictions of both Figure 8.3 and Figure 8.4 amplify comprehension by being pictorial. In contrast, the patterns of Figures 8.1 and 8.2 require development. They may have a syntax but it is not obvious; indeed, arguably the syntax of a TEM is more established than with these images. Like the TEM, as pointed out by Meynell, these visual images are "a type of representation

[169] Though not actually relevant to this discussion, I nevertheless must point out that Perini's choice of the energetics cartoon is most perplexing. The 1994 *Nature* paper provides 17 individual figures which are different sorts of representations of molecular structure, and the energetics cartoon is actually reproduced from a 1981 paper in the *Annual Review of Biochemistry*. The issues of syntax and semantics in how we depict molecular structure are important, but they are not addressed in her paper because of this selection of some other kind of figure from that 1994 paper. If this were a dissertation on structure-driven research, it would be appropriate to consider these problems in great detail because they are important in themselves and also relate to pattern detection, but for the time-being they are outside the scope of my project.

that is both paradigmatically pictorial and plays a crucial epistemic role in science" (Meynell, 2013, 329). Significant forms within these images must be discovered (Meynell, 2013, 341): they are not already stipulated. These pictorial formulations convey what was found when the system of interest was interrogated, and I will argue here that one of the ways for discovering what is significant in these images is through pattern detection. In fact, the epistemic action with the graphs is a kind of pattern-detection by classification—in both examples. With the Cartesian graphic, we have in fact attempted to look at the data in a specific way (after all, we could have chosen semi-log or log-log axes instead of straight Cartesian axes) so that we can find a pattern, which in this case happens to be linear. With the pie diagrams of metalloproteomics data, we have classified the proteins according to some standard criteria and then shown relative proportions of each class, as a way of seeing whether there are recognizable patterns.

With the cartoon of proteins interactions, a somewhat more sophisticated cartoon than the energetics cartoon, it is tempting to regard it as a kind of contextualization vis-à-vis the process of validation presented for the OES (see Section 7.3). Strictly speaking, this is not the case, although there are affinities between the two. This interactome, even if limited exclusively to proteins interacting with the Wilson ATPase, is cobbled together from numerous projects, whose results can be integrated but whose methods may be disparate. Its key assumption is that proteins do actually interact, an assumption for which we have enough empirical data to regard it as reasonable. Putting these proteins into one cartoon based on protein-protein interactions helps to actualize part of the system here, and it relates to the definition of system I have devised (delimited collection of entities whose structures and functions are somehow integrated). Being able to articulate the system, in this case by drawing a cartoon of it, advances the project of contextualization of omics data. As I envision it, those data all come from the same experiment. If the pattern of interactions is similar to or congruent to the pattern of interactions in this cartoon, it would support the validity of the proteomics data. Thus this pattern of interactions, and the very possibility of determining a pattern of

interactions, is critical to being able to contextualize omics data in a system but is not exactly contextualization itself.

8.2 An initial look at patterns and what has been said about them

First of all, a general notion of what a pattern is—both in general and in terms of omics—is required. Descriptions of what a pattern is are uncommon, if not downright rare, in the philosophical literature. Ray Paton, a computer scientist, defines 'pattern' as follows: "A *pattern* (diagram) is a collection of cooperating objects" (Paton, 2002, 70, original emphasis). One problem with this definition is that it conflates 'pattern' and 'diagram'. (Some patterns might not be diagrams as such.) Another problem is that it is not immediately apparent what 'cooperating' is. Nevertheless, Paton's notion of 'pattern' is noteworthy because pattern is connected directly to graphic representation, even if 'diagram' seems too narrow and limiting a concept, and it may be directly connected to diagrams of protein or other biomolecular networks. Nevertheless, in my opinion, it fails to be sufficiently general.

Daniel Dennett explores features of patterns in his classic 1991 paper, but while he asks "What is the pattern a pattern *of*?" (Dennett, 1991, 30, original emphasis), he does not ask what a pattern is. His pronouncement that "a pattern is 'by definition' a candidate for pattern *recognition*" (Dennett, 1991, 32, original emphasis) does not take us very far. In a later paper about patterns, written partly in reaction to Dennett's paper, Norton Nelkin regards patterns as 'relevant to our possessing propositional-attitude concepts, and to our ascribing propositional attitudes to others" (Nelkin, 1994, 70), but he effectively assumes that we all know what patterns are. Indeed, in general, the prevailing tendency is to talk about patterns as if we *just know* what patterns are. Since patterns in omics will do more work than merely subtend propositional-attitude concepts or propositions themselves, we need to take a closer look at exactly what constitutes a pattern (see Section 8.3).

8.2.1 Hanson on patterns

Despite the title of Norwood Hanson's 1961 monograph, he says remarkably little about the nature of patterns as such. He too makes the assumption that we just

know what a pattern is; however, Hanson makes important contributions toward an appreciation of the relationship between patterns and hypotheses. In his view, patterns precede hypotheses ontologically and, by implication, in terms of epistemic potency. Patterns critically connect data and hypotheses (Hanson, 1961, 87-8). He writes: "Perceiving the pattern in phenomena is central to their being 'explicable as a matter of course'" (Hanson, 1961, 87). Hanson is somewhat more interested in 'discovery' than in 'patterns', but he has important things to say about patterns.

Dealing with the physical sciences, Hanson rejects the notion that the laws of physics were arrived at by enumerating particulars—not a Baconian induction—and he also rejects the assertion that they are high-level hypotheses serving as a starting-point. Furthermore he criticizes the HDM for failing to explain how we get the laws of physics in the first place. His attitude about physics resonates with the observations about biological/biomedical research I have described: "Physicists do not start from hypotheses; they start from data. By the time a law has been fixed into an H-D [hypothetico-deductive; = HDM] system, really original physical thinking is over" (Hanson, 1961, 70). Hanson holds that the physicist confronts data with the aim of developing a "conceptual pattern in terms of which his data will fit intelligibly alongside better-known data" (Hanson, 1961, 72). Taking Kepler's determination of the elliptical orbit of Mars as his example, Hanson sketches out the process. It involves taking a data set (which he calls 'surprising' but probably would be better called 'recalcitrant') and playing with it: analyzing it one way or another with different possible patterns in mind and then settling upon the one pattern which engages the data best. Hanson quotes Pierce alluding to Whewell's notion of colligation—bringing together a variety of facts—but Hanson does not attempt a doctrinaire assessment of this process of interpretation of data, although he allows that it is principled, not a whimsical exercise (Hanson, 1961, 71). Thus the process of how patterns are found remains indefinite. In general within Hanson's analysis there seems to be one 'right' or 'adequate' pattern, except that he is interested in equivocating figures. He instances a bird/antelope figure, but he could have cited Wittgenstein's rabbit/duck. In reference to theories, Hanson asserts that

"physical theories provide patterns within which data appear intelligible" (Hanson, 1961, 90) and further he points out that theories as patterns are identified as "picturable" (Hanson, 1961, 91).

Written in terms of physics, at an era well before the rise of contemporary biological/biomedical research, and certainly well before omics was even conceived, Hanson's account may touch rather tangentially on nature of patterns as such, but it addresses the same shortcomings of the HDM that presently plague omics. It would be misleading to conflate patterns in omics data sets with physical theories, since the latter carry considerable conceptual baggage with them. The main point of similarity here is that omics presents the researcher with a data set about the biological system under investigation: it is the starting-point and finding patterns is the activity by which the data set is made intelligible. Hanson will argue that this is the enduring 'method' of physics right through to quantum theory. In terms of omics, I will argue that pattern-detection is how new knowledge is produced. So far as I can tell, Hanson does not stray from physics and mathematics into the realm of biology.

8.2.2 Dennett on patterns

In his 1991 paper entitled "Real patterns", Daniel Dennett examines patterns in order to address the problem of realism (Dennett, 1991, 30) and provides a purpose-built specifically philosophical assessment of patterns. He believes that patterns are a better model system for getting at what is meant by being 'real' than some of the other model systems he has proposed. Moreover, Dennett argues that patterns are important because they help us understand the nature of explanation: patterns constitute a kind of prediction (Dennett, 1991, 30), and that is what explanation is all about. He thus confers a certain law-like character to patterns. Dennett seems to be talking about what I would call a Humean regularity (see Sections 8.3 and 8.3.1). His point is that the pattern typically keeps repeating and thus provides reliable prediction. I see this action as essentially linear, like a hypothesis, or even a scientific law or a mathematical formula. Indeed with his apposition of explanation and prediction, we can immediately grasp that with Dennett we are in HDM intellectual territory. Dennett reinforces this reductive character of a pattern when he writes that "a pattern exists in some data—is

real—if *there is* a description of the data that is more efficient than a bit map, whether or not anyone can concoct it" (Dennett, 1991, 34, original emphasis). Dennett provides an interesting discussion of how noise in a system may obscure emerging patterns. Although his analysis is insightful and entirely relevant to problems of experimental design/interpretation which we are considering,[170] his argument that six patterns shown in Figure 1 (Dennett, 1991, 31) are equivalent because they were generated by the same computer program is unconvincing. The assertion is that having a unique system-generator ensures the identity of apparently multiple systems thus generated (Dennett, 1991, 32). It seems difficult to see why anyone would want to limit the versatility of patterns, and this claim appears totally inaccurate when applied to biological systems. We would never say that different metabolomes (the set of small metabolically active chemicals found in a suitably defined system)—different because they contained some of the same and some different chemicals—were actually all the same because they were generated by the same system, for example, the same cell. His point, however, is probably not directed at diversity of biological systems or redundancy within a single biological system. In his context, he is interested in the effect of noise on how clearly the pattern can be recognized. In relation to a biological system, his point is only that it is a single system even if it looks variable because of different degrees of noisiness in the data.[171] It is worthwhile to be reminded of the problem of noise in data, but Dennett's

[170] Dennett considers the example of six rectangular bit-maps, each of which displays a different pattern within the same general form (size, shape, square components, black/white coloration). He calls them bar codes, but actually they look much more like V-R (video-response) codes. The actual pattern is almost depicted in panel 'D': five large black squares alternating with four large white squares where the signal to noise ratio is only 1:100 (Dennett, 1991, 31). His point is that the information in any of these bar codes can be communicated by transmitting the figure literally bit by bit, verbatim as it were. Alternatively describing each bar codes as a pattern would be an improvement in communication because it would be more efficient and it would be "the description of a real pattern in the data" (Dennett, 1991, 33).

[171] This example (in the figure provided in the paper) is much like a Sudoku puzzle where the same solution can generate apparently different puzzles of varying difficulty depending on which and how many numbers are removed; there is still only one unique solution. Dennett's analysis will assume more importance in relation to the problem of salience where noisy data can obscure recognition of pattern when only a single pattern is present. It fails to speak effectively to the problem of finding patterns in biological/biomedical experiments of the type I am discussing because the redundancy of functional mechanisms in most biological entities being examined implies that there is or maybe multiple patterns present.

analysis fails to speak effectively to the problem of finding patterns in biological/biomedical experiments of the type I am discussing because the redundancy of functional mechanisms in most biological entities examined as systems implies that multiple patterns are or may be present.

8.2.3 McAllister on multiple patterns—physical sciences

James McAllister is keenly interested in the role patterns play in producing scientific knowledge. His analysis is mainly limited to the physical sciences, but it is relevant to the account of patterns I will present for research regarding biological organisms/systems because he holds that any one empirical data set contains, or encompasses, more than one pattern. He writes: "On my account … the world is a complex causal mechanism that produces data in which infinitely many patterns can be discerned" (McAllister, 1997, 224). Like Dennett, he observes the issue of how patterns appear different depending on the acceptable level of noise in the determination, but this is not his main thrust. McAllister proposes that what becomes the 'keeper' pattern (my term, not his) among the numerous patterns available in the data reflects the interests of the investigator, not anything peculiar to the properties of the pattern itself (McAllister, 1997, 224). He illustrates his point with the example of determining the melting-point of lead, an experiment generating different 'keeper' patterns depending on whether you are a physicist or a chemist. These 'keeper' patterns also vary with the scientific era, as better technology permits clearer delineation of a pattern thus generating a different pattern with less noise (McAllister, 1997, 226). While providing support for my contention that there are multiple patterns in an empirical data set, McAllister's assessment of patterns also has important differences. When he focuses on multiple apparent patterns each with its own noise level, the analysis seems very similar to Dennett's argument illustrated by his patterns resembling V-R codes. After all, the melting-point of lead exhibits less redundancy or random variability than a biological organism.

The nature of patterns in sets of empirical data[172] is more completely developed in some of McAllister's subsequent papers. In his 2003 paper on algorithmic randomness in empirical data, he reviews the broad support among 20th-century physical scientists and philosophers of science for the claim "that scientific laws and theories constitute a compression of empirical data sets" (McAllister, 2003, 637) and then demonstrates how the claim that must follow, namely that empirical data sets are algorithmically non-random, is impossible because data points in a data set are inevitably influenced all sorts of intercurrent "perturbations" (for which he produces a long and still incomplete list) which occur entirely by chance and are "local, variable and irreproducible" (McAllister, 2003, 637). This rejection of scientific laws as accurate, in the sense of actually compressing[173] empirical data, seems like old news given the previous analyses by others such as Nancy Cartwright, and few in the field of biology are overly worried about finding biological laws akin to what pass as physical laws.

However, what is relevant to my discussion is that McAllister takes the analysis in the direction of pattern detection. In the first place he attaches importance to the entirety of the empirical data set as the unique and most efficient carrier of information about the physical system being studied: if it could be compressed, important

[172] McAllister defines empirical data as: "[data] … gathered by scientists in observations and experiments, and count as evidence for and against the existence of phenomena and the validity of scientific laws and theories" (McAllister, 2003, 633). This is a helpful definition. However, he elaborates further: "By empirical data, I mean the recorded numerical results of observations and measurements of physical variables. An empirical data set thus consists of a string of digits that is the outcome of a series of observations or measurements. A data point is the recorded result of a single observation or measurement act: it is often represented as a geometrical point on a graphical plot of a data set. Before being appraised for their evidential value, empirical data are usually subjected to various forms of processing, for example to correct or compensate for factors believed to have affected the observations or measurements. What I call 'empirical data' is a set of observational results prior to such processing" (McAllister, 2003, 634). This is possibly the best detailed description of raw data from an experiment, although it may still require some modification for transfer from physics to omics. With certain types of biological/biomedical investigation, including some work in the broad domain of omics, the empirical data need to undergo some degree of processing just to be seen. Examples include ultrasonography and microarray chips. I will refer to empirical data or to experimental data as being roughly equivalent.

[173] McAllister uses this technical term from computer/information theory frequently: compressing data means packaging data in such a way that it requires fewer bits in order to be transmitted or stored. He also has a specialized meaning for 'noise': "not a margin of error or factual inaccuracy in data, but rather the purely mathematical discrepancy between a given pattern and a data set, as in classical information theory" (McAllister, 2003, 643).

information would be lost. Secondly, he allows that the data set can be portrayed as comprised of a pattern with its corresponding noise (where 'noise' is an information concept, not margin of error or technical inaccuracy—just the difference between the pattern and the entire data set).[174] This solves the problem for McAllister of what a scientific law does: it relates just to the pattern element of the pattern/noise duo. Nevertheless, this practical allowance notwithstanding, he sticks by his conclusion of the non-compressibility of empirical data in reply to criticism of this paper (McAllister, 2005, 406-7). McAllister's assessment supports my assertion of the importance of the data set in its particularity and may allay anxieties about variations from one sampling of a system to another. It echoes my insistence on maintaining inclusive open-endedness and attaching importance to the apparent outliers (see 'sector of consensus' versus 'sector of variability' in Section 6.4.2). Comparing his account with mine raises important questions about just what constitutes 'noise'. The variation I find it important and potentially highly informative, perhaps the richest source of information about a biological system which is constantly changing, is not necessarily noise. Noise appears to be whatever obscures the pattern (not error or inaccuracy, but discrepancy between pattern and data set), and it attracts less attention in McAllister's account (McAllister, 2003, 643). Finally, he does not seem to be dealing with systems where more than one pattern can account for the very same cellular function.

Although McAllister prefers to consider experimental data from the physical sciences, there is one instance where he discusses an example taken from molecular biology. The example, DNA sequences as exemplified by digital images, has been thrust upon him by critics, Charles Twardy and colleagues (Twardy et al., 2005, 394). The exchange of ideas is highly informative. He attempts to parry the question of whether a

[174] The question could be raised as to whether McAllister's demand that a noise level be selected and specified level is somehow the same as my stipulation for a meticulous description of the system being investigated. I would deny the equivalence. The role of the system is much more important to the OES than the level of noise to McAllister's version of the research strategy because the system has an important role in justification. Moreover, the noise component—a concept taken from information theory—bears strong similarities to my 'sector of variability', the non-consensus portion of the description of the system, although as we will see, this comparison also need to be refined.

DNA sequence is highly compressible, not entirely successfully, by reducing it to a biophysical read-out which is the "electropherogram", a display of the results of the automated gel electrophoresis which can be 'translated' into the familiar nucleotide code (A, T, G, C) by which we annotate a DNA sequence. Sticking to his definition of empirical data as exactly what comes off the machine, McAllister characterizes both the actual nucleotide sequence and the stylized summary as riddled with hypotheses; moreover, he notes that the machine itself has already generated hypotheses about the nucleotide sequence in sample in the process of generating the electropherogram. So he reformulates the question as to whether the electropherogram is algorithmically compressible, and he asserts that his critics have not adduced any evidence to say it is. What is relevant here, and not cited by McAllister, is that electropherograms have to be parsed for errors, and they are not as straightforward as it would appear.[175] Thus his insistence on empirical data by his definition is not misplaced. However, evidently McAllister and his critics are talking past each other. McAllister concedes that point, but he locates the patterns found in DNA sequences (such as redundancy and repeats) as operating at the level of hypothesis about the information obtained from the electropherogram, not at the level of the empirical data (the electropherogram) itself. This exchange clarifies how McAllister really locates his argument in the physical world, and it also makes it clear that talking about algorithmic compression here really operates on the borderline between how computers actually work (for example, compressing a JPEG file) and how theorists may import that language as metaphor. McAllister retreats to a circumspect, but non-negotiable, position: "If the noise term is algorithmically incompressible, then the data contained in the electropherogram are

[175] It is problematic enough that guides are available for identifying such problems: see, for example, "A Brief Guide to Interpreting the DNA Sequencing Electropherogram Version 3.0", whose version number suggests that the issues are manifold, (https://pmgf.osu.edu/sites/pmgf.osu.edu/files/guide_to_electropherogram_v3.pdf, accessed 6 November 2014) (University) for a description of causes(s) and solutions. McAllister's point about the machine formulating and implementing 'hypotheses' for producing the graphic readout is similar to my point about patterns in computers being involved in discerning patterns in data. This is not to invite the inference that hypotheses and patterns are somehow equivalent: I find McAllister's usage of 'hypothesis' here somewhat idiosyncratic.

also algorithmically incompressible, irrespective of what they suggest about the physical world" (McAllister, 2005, 409). His critics retreat to a broader vantage-point: practically speaking, we compress data all the time, and we can do this in part because data contain patterns. Their parting shot is: "Science is about *explaining* data: finding the patterns and discarding noise (Twardy et al., 2005, 401, original emphasis). In terms of what I am saying about omics and the OES, this dialogue highlights some of the issues which have come up: technical problems as such, the relationships between data and pictures of data, and the problem of noise. Importantly, "noise" which figures in this discussion is not exactly the same as any of the noise in a dataset from a system-driven experiment upon which I have focused. Nevertheless these issues—technical issues, relationships between data and their depictions, and bias—demand attention in a philosophical analysis of patterns with respect to biological/biomedical research. It is also worth pointing out that this discussion about DNA sequencing has little to do with McAllister's claim that there can be more than one pattern in a data set, depending on what interests the investigator.[176]

8.2.4 Mischler and Brandon comment on patterns *en passant*

Finally, in their 1987 paper "Individuality, Pluralism and the Phylogenetic Species Concept", Mischler and Brandon make a somewhat tangential suggestion about patterns in biology. They quietly slip in the notion that patterns constitute a category of criteria by which species individuality is determined (Mischler and Brandon, 1987, 399) and suggest via a subtly placed "i.e." that patterns are equivalent to the effects of biological processes (Mischler and Brandon, 1987, 399). In the context of a discussion of 'species' these patterns are (1) spatial localization and (2) temporal boundaries (Mischler and Brandon, 1987, 399); however, why these should be designated as patterns is not explained. What they do point out is "the lack of correspondence

[176] It is interesting to see the claim as such science is about finding patterns—except that it seems to carry a strong connotation of finding patterns which operate as laws. My analysis for biological/biomedical research operates more nearly at the level of analysis of Twardy and colleagues, but there is the residual problem about what to do with the bits of DNA sequence which do not fall neatly into recognizable patterns. Where I diverge from Twardy and colleagues is that I am not convinced that "discarding it" is the truly productive manoeuvre.

between either these [biological] processes or patterns resulting from them" (Mischler and Brandon, 1987, 400). Thus we are tantalized rather than enlightened by this brief discussion. On balance, Mischler and Brandon waffle about their characterization of pattern. Later, in the context of their preferred species definition (Phylogenetic Species Concept), they do invoke an epistemic manoeuvre which involves pattern identification: "In such cases grouping (estimation of monophyletic groups) will proceed solely by study of patterns of synapomorphy (i.e., shared, derived characters), and a practical ranking concept must be used until something becomes known about biology. This preliminary and pragmatic ranking concept will usually be the size of morphological gaps (i.e., number of synapomorphies along any particular internode of a cladogram) in most cases, a concept in accord with current taxonomic practice" (Mischler and Brandon, 1987, 406). What is noteworthy here is the persisting assumption that we all understand what is meant by 'pattern'. Nor is it entirely clear from their discussion how these patterns are directly equivalent to effects of biological processes.

8.3 After all: what is a pattern?

I have already presented a brief general definition of what I mean by 'pattern' in the Introduction. As a minimal description, I regard a pattern as "consist[ing] of individual entities and the pictorial/integrative relationship of those entities to each other such that a shape[177] is apparent." (If we consider music among our artistic examples or birdsong among the biological ones, then we need to modify this definition to the "… pictorial or acoustic or integrative relationship … .") Despite this expansion to

[177] Invoking 'shape' as a feature of patterns could have some interesting conceptual consequences. Specifically, it invites the notion of a topological epistemology where τοπος = place, spot, passage in a book, **region**, district, space, **locality**, position, rank, **opportunity**, according to a classical Greek dictionary (Feyerabend, 1918, 381). Especially the notions of 'locality' and 'opportunity' may be relevant to the patterns associated with omics. These advantages are not great enough to mandate this definition in which shape plays a central role, but they deserve to be mentioned. We could attach importance to the topology of proteins or, with Barabási, talking about the topology of protein networks. Thus similar ideas are currently in play. Thus it might be extremely interesting to develop a 'topological epistemology', but that is not how I conceive the subject of this dissertation. Thus I will not insist on 'shape' as a feature of patterns—I discard it when I jettison my minimal definition—but I recognize that the concept of 'shape' has some intuitive appeal and I come back to it very briefly in Section 8.3.2.

include ‘pictorial/acoustic/integrative’ relationships, this definition suffers from being extremely permissive. Although I will argue that a really interesting aspect of patterns is the nature of the relationships among entities in a pattern—in other words, that patterns are highly relational in interesting ways—a pattern is something more than simply a relation. I therefore propose describing a pattern as *a salient orderly juxtaposition of component entities, characterized by the relation(s) among them*. The word ‘entities’ is purposely vague. Moreover, implicitly I am regarding a pattern as being more than a straight line, which typically connects only two points. A pattern requires at least three component entities. It is altogether too easy to see lines.[178] One implication of this conception of pattern is that patterns are not inevitably present. However, this definition can accommodate patterns which are *ad hoc*, for example, throwing jacks on the floor. These patterns are “one-off”; we may not encounter any specific one again but we may encounter something very similar which we can recognize to be related because of the patterned nature. This latter situation finds correlates in biological phenomena, and I will explore this matter of repetition of a pattern over time in a later section (Section 8.3.2).

Likewise there are certain usages of ‘pattern’ which I wish to exclude immediately. I am not talking about a pattern such as a sewing or knitting pattern: basically, directions for how to fashion or manufacture something. Thus I will not include a gene as a pattern, partly because of the restrictions of my working definition and partly because I do not entirely accept that a structural gene constitutes the ‘directions’ for producing something. Simplistically, a structural gene encodes a protein, but we have seen that the actual resultant protein depends on more than just the structural gene. Additionally, I am not talking about pattern as archetype, as in the Christmas carol[179] where Jesus is our childhood’s pattern. This might be obsolete diction; however, eliminating the usage of pattern as archetype avoids any possible inference that ‘pattern’ in this context has Platonic overtones whereby the pattern

[178] Edward Tufte notes that the human visual apparatus is specifically developed to pick out straight lines in any random assortment of points (Tufte, 2001, 26).
[179] *Once in royal David’s city*, (Cecil Frances Alexander, 1848), third verse.

stands apart from the biological phenomena and somehow provides an independent reference to what we are seeing or seeking. Nevertheless this is a difficult aspect of the nature of patterns to which we likely will have to return. Dennett has described patterns as efficient descriptions (Dennett, 1991, 32), and the assertion that some patterns may function in a law-like fashion[180] requires further attention in the context of omics experimentation. This sounds to me like an easy-to-understand version of reductionism. I am not enthusiastic about conflating pattern and reduction. Regarding patterns as archetypes or as law-like statements to which data can be reduced or, finally, as the effects of biological processes imposes some conceptual baggage on patterns and pattern-finding which may impede the analysis here.

We might ask why we are interested in patterns at all. Despite the relative paucity of philosophical literature relating to patterns, Hume's commentary on constant conjunctions is likely to be relevant to the discussion. With Hume, we might simply confess to being captivated by observed regularities (Hume, 1978, 111).[181] If we are hard-wired to perceive patterns (a potentially dubious claim which is peripheral to this discussion), perhaps we are at least conditioned by physiological circadian rhythms and the normal electric control of heartbeat, indicated graphically by the electrocardiogram and acoustically by regular heart sounds, to perceive certain broad regularities. Even this claim is problematic because it is at least potentially circular in its biology: organisms (humans, mammals, etc.) have developed physiological features such as circadian rhythms in response to environmental cues of day/night cycles. Likewise, I am disinclined to explore the question of whether we perceive patterns because the facility to organize phenomena according to patterns increases fitness, and thus has definite

[180] Specifically Dennett writes that a pattern exists, as opposed to randomness existing, "if and only if there is some more efficient way of describing it" (Dennett, 1991, 32). Thus a pattern seems to capture some sort of law-like quality, and thus it serves a reductionistic role. Dennett has already claimed that we seek patterns because they facilitate prediction (Dennett, 1991, 30). As will become evident, my utilization of 'pattern' is somewhat different from this conception of pattern.

[181] We might also say that this is compatible with Thomas Nagel, who points out: "Regularities, patterns and functional organization call out for explanation—the more so the more frequent they are" (Nagel, 2012, 47). My impression from the little in the available literature is that patterns (whatever they are—no definition volunteered) really are somehow bound up with explanations.

evolutionary benefit. This is an extremely interesting issue (and in my opinion pattern recognition probably does convey survival benefit), but it is extraneous to this discussion. A related fundamental issue which does require some attention is why we expect to find a pattern or patterns when we confront an omics data set (to be addressed in Sections 8.5.1 and 8.5.2). After all, it might be that there is no pattern.

8.3.1 HDM patterns differ from OES patterns

It can be claimed that in fact scientists are always looking for patterns.[182] Perhaps they are. I believe we need to make some distinctions in the characteristics of patterns which might be sought. Most scientists are always looking for "simple" Humean regularities, where B follows A as it were without exception. It appears that this interest in regularities is tied to the practice among scientists to use hypotheses routinely. On my definition, a regularity described by a hypothesis is at best pattern-like, perhaps describable as a specialized kind of pattern which is very simple and involves two nodes. Another fairly simple regularity is the dissection of the energetics of an enzyme where the enzyme E assumes an intermediate state E* which effects whatever that enzyme is supposed to do and then regenerates itself as E. As previously mentioned, crucially for the ideas explored in this chapter, the patterns operational with omics are not linear. Specifically, unlike the situation with the HDM, a pattern demonstrated by the OES is not linear. The test implication as a conditional, $A \supset B$, generates a linear construct, and in fact the hypothesis itself could be described itself as "linear". The enzyme energetics shown as $E \rightarrow E^* \rightarrow E$ (where E* is the activated form of the enzyme E and further quantitative details may be supplied along the way) is linear or even circular[183] because it really involves only two components, E and E*. It can also

[182] I am grateful to Nathan Brett for raising this point in an unrelated discussion on 8 November 2013. The claim that scientists are always looking for patterns served as a kind of assumption to the actual concept he was advancing in a somewhat different context.

[183] The question can be raised as to whether any circular pattern could exist. In particular the status of a feedback loop can be questioned. On balance, I think that a feedback loop serves as a pattern. Most feedback loops entail the interaction of several entities: X produces Y; when Y exceeds a certain threshold, Z turns off the action of X; should Y drop below a certain threshold then Q causes X to become active again. While this has a closed-loop pattern, it cannot be recast as a linear situation. I am indebted to Eva Boon (7 December 2013) for pointing out this issue. My own view is that it might be more precise to characterize the $E \rightarrow E^* \rightarrow E$ sequence as a Möbius strip, which looks like a circle but in fact is linear.

be written as E⇌E*, a kind of back-and-forth action, with obvious linearity. A comprehensive pattern relevant to system biology, described via the OES, displays at least three nodes (A, B, C) so that there are at least two ways to get from A to C. Thus patterns characteristic of such systems are at least multi-linear and may be regarded as describing an area. This might look like a polygon composed of adjoining triangles. Moreover, the points of this polygon can be single nodes or clusters. Such patterns rapidly become very complex. You do not know exactly which pathway will be taken when: simple regularity does not apply. Finding that the patterns of interest to hypothesis-driven and system-driven research are different discloses another distinction between these two types of research. The idea that a pattern functions for a biological system much as a law does for a physical system then almost immediately suggests itself, and I am inclined to reject this idea outright. [184] Daniel Dennett's analysis of patterns obscures this distinction between characteristic HDM- and OES-patterns. The stochastic nature of how the pattern works within the biological system is inconsistent with having a law-like nature. However, it is difficult to deny that a pattern has at least an organizational role, and that role has potential to be explanatory. However, to summarize, HDM and OES display different kinds of patterned outputs: put simply, linear versus spatial, respectively.

Finally, returning to my claim that the HDM and OES are characterized by different kinds of patterned outputs (regularities for HDM and patterns as such for OES), I wish to point out that, with the OES, looking at data and appreciating a complex, multi-faceted pattern is a unique kind of epistemic activity. It somewhat resembles induction in that from some number of specifics a bigger, more encompassing picture is

[184] Linear patterns (regularities) such as those associated with hypotheses might in some hands look very much like laws, and I can see the basis for asserting that physicists are always looking for patterns (that is, regularities)—mainly so that they can figure out laws at work behind the phenomena they measure. Moreover, McAllister holds that patterns can be law-like (see Section 8.2.3); however, his focus is almost exclusively physical science, not biological science. It will be pointed out that biological systems generally do not have laws. However, we will have to examine the idea put forth by Barabási and colleagues that networks (obviously a kind of pattern qualifying as relevant to *system* biology) can be characterized as having a law-like design, which can be summarized mathematically. This problem will be considered later (Section 8.6.1).

developed, but surely it is not the same as inductive inference. Moreover, and most importantly, it is specifically not the same as generating or testing a hypothesis.[185]

8.3.2 My conception of 'pattern' in greater detail: shape, relations

My definition of a pattern as 'a salient orderly juxtaposition of component entities, characterized by the relation(s) among them' takes its impetus from the necessity imposed by omics of dealing with data which are non-propositional. It seems to me that it is a simple empirical fact of how we do biological/biomedical research that the findings are not composed of words and the expression of these findings is in a form we do not read. We see it. Thus we have individual elements, not words or ciphers. For those elements to be patterned, they must assume some sort of relationship to each other. That relationship often, but not always, has a peculiarly spatial quality to it.[186] This spatial quality is captured in the word 'juxtaposition', which applies equally well to space and time.[187] Perhaps a further example would be helpful. Suppose we hear a birdsong with the following features: three blasts on the first note, then a trill on the note a major third higher, and then four short/sweet tweets on the note a perfect fourth down from that trilled note. The pattern could be displayed as linear, but it has the shape of the intervals above and below the initial note (and—at least to some ears—it can form a tone cluster). If you then hear a "cockeyed" reply: three blasts on a first note, then trill on the next whole note higher, followed by four tweets a tritone down from there—you will recognize that the pattern has the same shape even if the

[185] I am grateful to Richmond Campbell for helping me to see that patterns are not the same as hypotheses, even if the analytical process is construed as somehow involving inductive inference. The question might be asked whether appreciating a pattern is the same as generating a model. My answer is that it is not the same. Models are built, not appreciated. Classifying data is not the same as constructing a model. The reason this is an important distinction is that some will argue that a model is a hypothesis (in 3-D, as often as not). I myself have suggested that a model might function as a 'hunch' hypothesis, and that is just the point: it is not a hypothesis which is capable of driving the experiment.

[186] Biological patterns are usually space-filling or extended over time. For convenience, we depict them on a flat space with various forms of shorthand such as CH_3-CH_2-OH (or 'EtOH', the ethanol we drink in wine or beer) or any metabolic map such as the Krebs cycle. I have purposely specified seen and heard entities: any sensory input might in principle serve as the basis for a pattern but other sensory inputs are less important in scientific research.

[187] I have chosen this word because it is extremely neutral. Alternate possibilities include 'composition' (in the sense of 'being positioned with') and 'arrangement'. The problem with such words is that they subtly introduce a notion of agency—in general, someone creates a composition or makes an arrangement—and references to agency need to be avoided in this definition.

musical intervals differ somewhat. You are more likely to think it is a maladroit singer of the same species than some different bird. What is interesting is the aural pattern consists of a juxtaposition of tones and sound-figures. Effectively it has shape—and shape matters.[188] Shape in this situation is relational, intrinsically involved in the relational nature of a pattern.

The relationship(s) of component entities thus become important for analyzing the broad spectrum of patterns we encounter. I will focus on two general types of relation within patterns: recurrence and interaction. Undoubtedly there are others[189], but these seem to be the important ones for the patterns I will consider here. For example, in connection with a discussion of epistemic virtues Helen Longino identifies other relations which could be imported into this discussion of relations in a pattern. These are "complexity, mutuality, reciprocity" (Longino, 1997, 22). However, I believe these descriptors could apply to both recurrence and interaction within patterns, and thus I do not find that it advances my analysis to include them at this point. When we think about patterns, we notice that for most patterns, being patterned involves exhibiting some kind of repetition. 'Recurrence' is a broader concept than 'repetition'. The relation of 'recurrence' generates repetition in either time or space. Another broad concept of relationship in patterns is 'interaction'. 'Interaction' is a complicated concept, as I will show: it includes 'regulation' but is not limited to it. These relations broaden the scope of what it means for a pattern to be orderly. The general meaning of 'orderly' is to be a regular sequence or arrangement.[190] I believe there are other forms

---

[188] Of course, we are comfortable with the notion of heard patterns being styled as compositions. For both seen and heard entities, composition captures the sense of adjacency. To say that a musical line has "shape" invokes a metaphorical use of the word, at least for some people, although I would argue this metaphor is highly intuitive. Musicians routinely talk about shaping a musical phrase or even a single note.

[189] Some may object that I have omitted the relation which could be named 'mathematical' (inclusive of 'arithmetic'). Such a relation could be given by the formula for series such as 1, 3, 5, 7, … which is $[2n_x +1]$ where $n_1=0$. I have not excluded series outright from my consideration of patterns, but I believe that the pattern relation denoted as 'mathematical' has limited application to the experimental work I am describing here. However, it may be fully relevant to systems biology, which is beyond the scope of my discussion.

[190] Definition is from the on-line version of the Oxford English Dictionary (http://www.oed.com.ezproxy.library.dal.ca/view/Entry/132345?rskey=uFOFfL&result=1#eid, accessed 21 October 2014) (Anon.: OED online).

of orderliness. Entities which have a stable functional interaction, including regulatory actions, are also orderly. Such relations within a pattern may result in a pattern that is unique although it has no internal repetitive elements. Likewise a unique arrangement of features taken all together can be effectively orderly: this generates the unmistakable unique case which can be identified every time it occurs. In effect, the idiosyncratic unique arrangement is the limiting case for 'orderly'. Since these varieties of relation within a pattern give my definition flexibility, they require some further discussion.

8.3.2.1 Recurrence

Recurrence is a feature of many patterns, but repetition generated can be spatial or temporal. Some patterns show actual repeats, a version of internal regularity. This tessellated kind of pattern is probably the easiest to appreciate as a pattern as such. You see it with wallpaper and fabric for draperies—floral chintz and brocade. The rondo form in music (A-B-A-C-A-D-A-E), with its recurring theme A, is also a possible example. However, it can get more complicated. Some patterns are composed of patterns, if the elements are themselves clusters of patterned elements. Acoustic analysis of a bird-song may reveal a monotonous, repetitive concatenation of modular elements which display some internal complexity. The cartoon structure of protein similarly captures a pattern (the structure) composed of patterned elements (such as α-helices and β-sheets).

Another type of repetition can take the more subtle temporal form. A pattern which is some sort of unique configuration has a temporal form of repetition: every time you see it, you recognize it as that unique pattern. For example, every time you see a particular disarrangement of hepatocyte-like cells on a liver biopsy slide, you know it can only be a particular variety of hepatocellular carcinoma (liver cancer): nothing else looks like that. Every time you see the Wicked Witch of the West from *The Wizard of Oz*, you recognize it as that unique personage, even if it is not a photo of Margaret Hamilton, the iconic actress who created this role in the 1939 movie, but another actress wearing the same sort of costume and make-up.

8.3.2.2 Interaction

The relation of 'interacting' (including 'regulating') is perhaps more difficult to explain, but my intuition is that interaction is invariably pattern-producing. Conceptually, interaction is complicated because it can be compositional or functional. When several proteins aggregate to form a unit, this is a compositional interaction. Such an interaction may be automatic given the structural characteristics of those proteins. When proteins in a cell interact to support some function of the cell, then we can distinguish a pattern which supports an intracellular action. We can call such an interaction "functional". It may be instrumental, for example, the intracellular handling of copper. Another important example of a pattern produced by the interacting relation is a protein network which effects intracellular signaling. Interaction patterns include those which regulate. For example, the regulating relation includes feedback loops. Regulating functions have assumed special prominence in biological research since the early days with the *lac* operon. However, feedback loops are not the only interactional pattern capable of achieving regulation. Aggregates of proteins known as transcription factors regulate gene action (transcription), although sometimes single transcription factor perform that function. An interesting philosophical consideration, well outside the scope of this dissertation, is whether any interaction exists in a cell which is not regulating in some respect. So we end up with various kinds of interactions within the concept of the relation of 'interaction' in patterns: compositional or functional, and among those interactions relating to function some which activate processes and some which regulate processes.

By way of summary, not every assemblage of entities is or constitutes a pattern. A pattern has certain features: the component entities are conjoined in an orderly fashion. I have expanded the connotation of 'orderly' to include stable functional interactions and regulatory interactions. I have pointed out that patterns involve relations. I am most interested in recurrence and interaction (the latter inclusive of regulation). Finally, this orderly juxtaposition is salient. It is cohesive enough to jump off the page, whether it gets appreciated or not. For omics research we need a broad

enough concept of what a pattern is in order to explore how pattern detection could be at the centre of the epistemic process of system-driven research.

8.4 Patterns in omics (mainly proteomics)

Now, given this discussion, when we do turn our attention to omics, the immediate problem is that the situation is extremely complicated. Numerous patterns can be identified. The concept of pattern described above is epistemically important for the big picture of the system under investigation; however, patterns often turn up as the immediate output of the experiment itself, and some patterns are embedded in the analytical process. In order to figure out what is going on here, it may help to enumerate the different sorts of patterns generated by different kinds of experiments. In this particular case, it helps to limit our discussion to proteomics, but it does not help very much; however, proteomics remains conveniently representative of omics in general.

8.4.1 Hidden patterns as assumptions

Preliminary to this discussion, it is helpful to recognize that embedded patterns within analytical equipment, the technology of proteomics, function almost as assumptions. What we have not considered so far is that an algorithm is a pattern. Of course it is fundamentally a linear exercise: do this and this and this, and then do it all over again. It can be more complex, essentially branching: do this and this and if you get X do that and that, but if you get Y do something different from that. This analysis is relevant not only because an algorithm is a pattern, but more particularly, because computers run on algorithms. Embedded in the analysis of the mass spectrometry data we find numerous algorithms driving the computer program. It is not merely that the bioinformatics program picks up on peptide patterns; additionally, the computer itself runs on patterns. We have seen that differences from one bioinformatics program to another can introduce error in protein identification. Differences in algorithm affect determination of what proteins are present. Thus it is not an exaggeration to say that omics appears to be patterns stacked one on top of another. It is important to keep in mind that just because you develop the best available consensus bioinformatics

strategy, given all the patterns and pattern-detectors hidden within it, it does not mean that this strategy is free from the bias of hidden assumptions.

8.4.2 Patterns which reveal themselves incrementally

With omics, sometimes the pattern in the data from an experiment is immediately obvious. Sometimes the data require conceptual or classificatory analysis in order for the pattern to be identifiable. In some situations pattern detection seems to progress extremely slowly, as if in slow motion. The pattern becomes apparent gradually, almost by accretion. In the process of appreciating a pattern in a set of omics data, there may be points of non-clarity. The pattern could go in one direction or another. Not uncommonly, these are situations where we assign a temporary epistemic valuation of 'don't know' instead of 'works' or 'doesn't work' (see preceding chapter, Section 7.3.3).

One possible basis for this ambiguity is that biological organisms frequently have multiple mechanisms for achieving a certain physiological objective. This feature is perhaps most evident in the redundancy of metabolic pathways. With hepatic drug metabolism, more than one cytochrome P450 may be involved in the metabolism of any specific drug. For example, the well-known analgesic acetaminophen can be a substrate for three different cytochromes P450 in the human liver (1A2, 2E1, and 3A4), although the product of the action of any of these enzymes seems to be the same toxic intermediate. Thus several different patterns may be associated with the same physiological phenomenon. The metals copper and iron can be found in biological organisms in an oxidized or reduced form (respectively: for copper, Cu^{++} or Cu^{+}; for iron, Fe^{+++} or Fe^{++}). This characteristic endows these two metals with extraordinary versatility as metabolic co-factors; however, we might anticipate that different sets of proteins are responsible for the disposition of Cu^{+} versus Cu^{++}. The question can be raised whether multiple patterns actually reside simultaneously in an omics data set. My interpretation is that this indeed is the case. My rationale for this interpretation is the redundancy of physiological mechanisms in a single organism. If the system under investigation comprises more than one mechanism for accomplishing some function, it may generate more than one pattern of proteins when we examine the relevant proteome. A

supporting argument is that this is very similar to the Rubin vase optical illusion which can be seen as a vase or two human facial profiles. (Wittgenstein's famous duck/rabbit example belongs to the same genre of optical illusion.) While the vase may be salient for some and the profiles are salient for others, nevertheless I would argue that both patterns—the vase and the profiles—are present in the depiction being assessed. In this respect my position is similar to Dennett's regarding the assessment[191] that the pattern—for me, 'pattern(s)'—is/(are) actually there in the data (Dennett, 1991, 49). It is advantageous for some to see one pattern (vase) and others to see the other (profiles), and possibly most advantageous if you can see both. This variability of what is seen might be one way of conceiving what I call 'suggestibility', as characteristic of pattern appreciation. Especially with proteomics, we may identify more than one pattern in a dataset. Additionally, I will show that there is more than one way to identify patterns with proteomics.

In some of these situations the pattern as it becomes evident may itself push our analysis in an unexpected direction. This flexibility of a developing pattern is typical of the OES. It may be regarded as another aspect of a pattern's "suggestibility": a secondary or incidental feature of patterns whereby a pattern can elicit a kind of lateral thinking about a group of entities or phenomena or data. An incomplete pattern may "suggest" possible ways to configure the available entities to complete it. A pattern which looks vaguely like another pattern may direct attention to the similarities and differences so that new concepts arise. Pointing to the possibility of new or alternative patterns is a kind of predictive posture. It serves as a positive value. This is underdetermination in a constructive mode. It is not quite the same as finding holes in the pattern and predicting what entities might fill those holes. It relates to the overall

[191] Where it becomes difficult to track the similarity of our views has to do with the different inherent interests and tensions of a general epistemological problem (with Dennett) versus a scientific epistemology (mine). It is worth pointing out that Dennett tolerates the notion of more than one pattern in a dataset: "There could be two different, but equally real, patterns discernable in a noisy world. … The choice of pattern would indeed be up to the observer, a matter to be decided on idiosyncratic pragmatic grounds" (Dennett, 1991, 49). His conception of such multiple patterns dependent on the interests of the observer, as well as his allusion to information theory—and specifically to compression of data (Dennett, 1991, 32)—seems remarkably congruent to McAllister's views.

description of the pattern itself. The epistemic vigor arises because some patterns are analogous to others,[192] and we can appreciate the similarities between patterns, in addition to pattern-formation of a single pattern (or multiple patterns) within a data set. At some point in the experimental work, it may not be apparent which of several patterns may pertain. The flexibility permits a broader assessment of the issue under investigation. As such, it permits several different patterns to be considered until accumulating data constrain the results to one pattern or indicate the parameters within which one pattern or another might operate. Importantly, such flexibility is consistent with the open-endedness which is a distinctive feature of the OES and assumes considerable epistemic importance in omics. Dennett's emphasis on predictability with respect to patterns obscures this feature of patterns, namely suggestibility. Patterns encountered with omics retain a certain openness, which is not characteristic of a law, the sort of reductive summary of a pattern Dennett would like to identify in any pattern. This characteristic 'suggestibility' might be styled as a heuristic action. In effect, the characteristic of suggestibility casts absolute correctness into doubt: what might be regarded as error could be another way of being correct.[193]

[192] I am grateful to Letitia Meynell for pointing out the existence of 'analogous patterns' as a variety of patterns. When analogous patterns support lateral thinking, they are extremely interesting. However analogous patterns may serve a more mundane purpose more routine: for example, when it describes a general class of appearances, like a neoplastic cell. James McAllister describes some other specific types of relationships between patterns: 'projectability' where a pattern found in one data set is also found in another and 'consilience' where a single pattern found in one data set is reproduced in a data set of a different kind (McAllister, 2010, 808). While he grants that these characteristics and perhaps others unspecified endow patterns with the property of being robust, he argues that such features do not distinguish physically significant patterns from those which are not physically significant. While the latter point may be true, they are characteristics relating to patterns as such and thus of interest here. McAllister's commentary may be influenced by his focus on the physical sciences. Katherine Brading argues there may be a multiplicity of patterns in a given data set but only the robust ones count (Brading, 2010, 832) where robustness is a function of (and evidence for) the restriction of the real world on the pattern (Brading, 2010, 831).

[193] In addition to acknowledging the well-recognized mechanistic redundancies of biological entities, this concept has some other implications. It accords with the view in reliabilism that generally reliable processes will not be absolutely reliable: sometimes the justified belief is wrong after all. Finally it is superficially similar to James McAllister's thesis that multiple patterns can be seen in one data set (McAllister, 2007, 887), which I will consider subsequently (Section 8.6.2).

8.4.3 OES: appreciating patterns directly or indirectly

Pattern detection with the OES is a complicated process mainly because of the complexity of biological organisms. In order to manage this difficulty, at least for proteomics, I suggest that there are two main types of processes available for pattern detection. One is direct (immediate) appreciation of pattern(s) in the data because the graphic is the output of the experiment. The other is indirect and involves constructing pattern(s) by imposing some further analysis or classification. Generally, 'indirect' methodology promotes finding interactions among the entities in the data set. Furthermore, identifying a pattern by the OES is not the same as asserting its relevance. Pattern detection displays a prospective quality whereas relevance is always determined after the fact. Determining relevance necessarily involves evaluation and justification of the pattern, that is, contextualization. Importantly, whether a pattern is direct or indirect has to do with how we perceive the pattern, not necessarily with the character of the pattern itself. Direct patterns display immediacy, but indirect patterns have to be constructed.[194]

8.4.3.1 Direct patterns

With direct patterns, visual patterning is the read-out of an experiment. Instead of a quantitative result, the result is a picture with variable colours or colour intensity to indicate variability among findings. Relevant to proteomics, we can discuss this type of experiment and its resultant data set in terms of using a protein microarray. The protein microarray story is complicated by formidable technical problems, not all of which have been solved in the 10-15 years since protein microarrays were first introduced. Protein microarrays, also known as protein chips[195], are similar to the nearly iconic omics

[194] There is a tendency to call a pattern which is appreciated directly a "direct pattern" and one appreciated indirectly an "indirect pattern", but of course this is in addition to whatever kind of pattern it happens to be. Detecting a pattern is epistemological and has to do with methodology. The urge to invest a pattern with a descriptor reflecting the mode of perceiving it is sufficiently strong that both usages will be found in my discussion.

[195] The technology developed by the American company Ciphergen where a slide has protein binders affixed to it, to which a cell lysate or other cell preparation is applied and then mass spectrometry serves to identify the adherent proteins, is sometimes called a protein chip (perhaps more precisely ProteinChip), but I have referred to it as the SELDI method for proteomics. It is not what I am speaking of as a 'protein chip', although strictly speaking it does qualify as a kind of protein chip. Of note, this Ciphergen

application of this sort, the DNA microarray. These devices, also known as a DNA chips or "gene chips" (Affymetrix Inc. holds the trademark 'GeneChip'), hold a set of DNA fragments representing a subset of the genome or the entire genome of a particular organism.[196] Both varieties of microarray technology have been hailed as a mainstay of contemporary high-throughput technology for biological research (Zhu and Snyder, 2003, 55). With either type of microarray, the read-out is typically a colour pattern. It may be grid with coloured dots or squares of differing intensities of a single colour or of multiple colours, or it may be what is known as a heat map, where the visual colour spectrum is used to indicate something quantitative, much as the temperature map on the weather report.

Protein chips (including *proteome* chips) remain, broadly speaking, in a developmental phase.[197] Gene chips illustrate direct patterns very well. Because DNA

---

technology can be set up to do metalloproteomics, but published practical experience indicated that it could be inconclusive (Roelofsen et al., 2004, 739).

[196] This technology can be designed to investigate presence of alternations in genes of interest or to determine how genes are expressed in a system. What it is really investigating is the transcriptome of a system. Within cells, DNA is transcribed to RNA, and the entire set of these RNA transcripts is known as a 'transcriptome'. The methodology is somewhat convoluted and may seem confusing initially. (Consequently I am describing it in some detail, despite not being directly related to proteomics.) This microtized technology dates from the mid-1990s (Schena et al., 1995, 467).

[197] The protein microarray is another example of classic high throughput technology. The advantage of a protein microarray is that it permits efficient examination of protein expression (not RNA expression, which does not always correlate with protein expression) and protein interactions (Seong and Choi, 2003, 2177). A protein microarray is similar in format to a DNA microarray except that proteins are attached to the slide, not DNA. The physical format is microtized and the detection methods are colorimetric so that the same type of plate-reader and image analysis software can be employed. Numerous technical problem must be solved: how to produce all the proteins, what surface to put on the plate and how to attach these proteins so that they retain tertiary structure and function, how to design and produce specific enough antibodies for immunoassays, how to design functional assay to test enzyme function or protein-protein interaction. Many of these issues were reviewed in detail in a 2003 paper which also stressed the need for a standardized data structure (Seong and Choi, 2003, 2187). The outlook in 2001 was favourable although the obstacles were well-recognized. Protein microarrays were regarded as "second-generation proteomics" complementing proteomics based on mass spectrometry (Albala, 2001, 145). The outlook was optimistic in 2003 in a comprehensive review (Zhu and Snyder, 2003, 62). By 2006 it was possible to describe 'proteome chips' where all the proteins arising from a completely determined genome, for example from a yeast, are put into the microarray (Kung and Snyder, 2006, 617). Proteome chips are interesting because they are really just ordinary protein microarrays except that the organizing principle for how the proteins were chosen for inclusion is clearly specified. They also illustrate well the challenges and advantages of this technology (Kung and Snyder, 2006, 618-9). More recent experience still emphasizes both the versatility and limitations of protein microarrays (Gahoi et al., 2015, 226-7) which are finding broader applications within cell biology. The Ciphergen ProteinChip is a more narrowly-

microarrays have caught the attention of contemporary philosophers of biology, it is worth commenting on how they work. Physically, it is a glass or metal slide with numerous DNA fragments (uniquely representing genes) affixed to it. Each dot/square on the slide has multiple copies of that 'gene' attached to it. The experimental methodology is clearly system-driven. The actual experimental technique involves isolating the RNA from the system (cell(s) or tissue or organism) of interest and then manipulating that RNA pool to isolate the mRNA. The isolated mRNA is then converted into cDNA, which is nothing other than the DNA version of the mRNA present, and in the process it is labeled with a fluorescent dye. Then this colour-tagged cDNA is applied to the chip. If that cDNA finds the DNA fragment complementary to it, then it attaches: the amount of attachment is proportional to the amount of the respective mRNA originally in the preparation. The chip is then washed off and what is measured is the cDNA sticking to the chip. If the cDNA sticks to a DNA fragment on the chip, it indicates that this gene is expressed because the applied cDNA is a stand-in for the mRNA originally found in the system. With DNA microarrays the question often deals with up- or down-regulation of gene expression, not merely whether the gene is expressed at all. Such a question is intrinsically comparative. In such a case, mRNA is prepared from the system of interest and a reference system, such as a newborn tx-j mouse's liver and a 2 month-old tx-j mouse's liver. The labeling is more complicated: a red fluorescent tag for one and a green fluorescent tag for the other. Then the red and green colours on the read-out relate to high expression in newborn liver or the other. The peculiarities of these tags is such that if expression is roughly equivalent in the two system the fluorescence reader sees the colour as yellow. The resulting are patterns are composed of red—green—yellow—black (no binding at all). It is important to emphasize that despite involving a comparison, no hypothesis is operative here. The experiment addresses a "I wonder what happens as the mouse ages …" question, not a "if copper accumulation at

---

focused technology, representative of various similar (competitor) technologies using proteins as biomarkers for complex diseases (Albala, 2001, 149); this remains its main use at the present time.

2 months-old is 10 times normal, then gene *SPQR* will be over-expressed" hypothesis. (A different kind of experiment could investigate this hypothesis more efficiently.) Raw data from the DNA microarray experiment are image files. The extent of expression, where expression is roughly equivalent to the amount present, in some experiments can be determined from the intensity of the colour. These files are typically transformed into numerical data files where the image as such is reported as a quantitative measure, the fluorescence intensity, for each gene (Malone and Oliver, 2011, 3 of 9). Complex bioinformatics is required to analyze these voluminous files of quantitative data. Much of that computerized effort is an exercise in seeing the data, although latterly the data may be reported as grouped or classified.

Practically speaking, what is literally striking about the raw data read-out is the pattern: increased or decreased expression/detection of DNA on the chip (or proteins on a protein microarray), or whole regions of the chip where numerous ones are highly expressed or hardly expressed. The broad questions here include the adequacy of the analytical approach which involves some sort of calculation of the degree of expression. For annotating quantitatively the increased or decreased expression, there is debate as to what qualifies as an increase or decrease. Some biological/biomedical researchers argue in favour of a pre-set fold increase: for example, if the expression is not at least twice base-line it does not count. Others would say that any increase at all is of biological interest. These arguments may miss the point. Specifically, at least in some cases, the human eye (meaning, of course, the eye-brain unit) grasps the results of the experiment better than a quantitative analysis.[198] In such situations the pattern itself may be the unit of comparison. Striking or subtle or predefined differences may not be taken separately: rather the pattern itself as a whole serves as the informative finding.

With these chips, blinded grouping of a visual image, which seems like an operational impossibility, is possible. Here is how it works: when we look at a protein-chip or a gene-chip, we do not actually know what is on each spot. (Even if you have

[198] I am grateful to Joseph Pezacki of the National Research Council of Canada, Ottawa, for pointing this out to me some years ago (probably in 2010) in connection with his own research.

manufactured the chip yourself, it is difficult to keep thousands of spots well sorted out in your own memory.) We may rearrange the colour display output so that the very positive are separated from the very negative. Then we look to see what entities having fallen into one category or another. Likewise we might chose some intensities of expression which we are willing to discount, or we might decide to keep all data for analysis. Patterns may thus emerge. Here the pattern depends on the arrangement of colours whose connection to actual proteins is anonymous to us at the time of pattern detection. With bioinformatics, the analytic approach remains similarly objective.

A microarray can be subjected to some classification prior to the emergence of patterns, in which case it can be argued that the pattern was constructed (or indirect), not direct. For example, in the DNA microarray shown in the Figure 8.1, it looks as if the results were subjected to some computerized reorganization to permit emergence of patterns (see tree-like configuration at left of array, consistent with a hierarchical cluster analysis having been performed). However, I would argue that the key step is here in terms of knowledge production is simply seeing where red or green is prominent.

8.4.3.2 Indirect patterns

The situation is different with a pattern whose detection is necessarily indirect. The process of pattern-construction involves classification: the known findings of an omics experiment are classified according to pre-determined criteria of interest. When we confront a large data set of the proteins identified by mass spectrometry-based top-down proteomics, we have a different situation. As mentioned, we frequently classify the data along some conventional lines, perhaps structural or functional aspects. Some classifications are based on functional or structural classification. This is a commonsensical convention for reporting proteomic data. It permits developing patterns in proteomic data, according to features which may be of interest. Like others, we chose to group our data of relating to the HepG2 Cu-metalloproteome in terms of categories such as *known cation-binding*[199] *proteins*, *chaperones*, *nucleic acid binding*

[199] A cation ("cat-ion") is a positively-charged chemical moiety.

proteins, and *redox proteins*. Additionally, we portrayed our data diagrammatically in pie charts according to these functional classifications so that the findings also had visual impact (see Figure 8.3). One point about this strategy of pattern-creation through classification (as opposed to pattern detection/discovery by examining data blindly) is that it can introduce bias because the categories tend to be conventional or standardized.[200]

This strategy of pattern-creation by grouping data is so prevalent within omics subdisciplines, not just proteomics, that it requires further examination to try to sort out the subtleties of constructed pattern detection. Let's consider the following situation: you purchase at auction, sight unseen, several hundred books from the library of a scholarly polymath. Numerous boxes (perhaps 20-30 of them) of books duly arrive and you empty the boxes haphazardly onto the awaiting shelves of your own capacious library. Now what do you do? On the assumption that you did not buy the books merely for decoration and also that you are not one of those lazy people who unpacks and then procrastinates, it seems likely that you will examine what books you have acquired. Then it seems likely that you will arrange them on the shelves. The question is: how to do this? You could arrange the books based on the physical appearance of the book, for example, according to the height of the spine, tallest at the left and then gradually to the shortest at the right end of the shelf. Similarly you could group all the books with the same colour of binding or dust jacket together. It seems more likely that you will group the books by subject matter: fiction versus non-fiction, poetry within fiction (?—or separately), and non-fiction divided into further categories. But how to do this classification of non-fiction? For example, you could put all the biographies together or you could group biographies with other books on the same general topic. With the latter

[200] The protein-chip is not without bias, however: as with the hidden bias of computer algorithms, the design of the chip predetermines what will be examined. Making the chip involves the choice of what peptides are applied to the chip. There may be 40,000 peptides on the chip, but this is still a selected collection of peptides. The same consideration applies to the gene-chip. In fact you can purchase gene-chips devoted to one aspect of cellular physiology or another. Likewise the metalloproteomics data can be viewed as a pattern directly, by using the electrophoretic pattern as basis for comparison (see Figure 8-2). A limitation is that the pattern of dots permits comparison, but without the benefit of protein identification.

approach the biographies of Gandhi would be shelved alongside of the books about the history of India and instructional manuals about yoga. At the end of this exercise, you become thoroughly acquainted with the books you have acquired—possibly the most important aspect of this exercise. You may find some totally unexpected treasures: a precious first edition, an early obscure monograph, an autographed pamphlet better identified as 'ephemera' than as a book. You also will have learned something about the intellectual moorings of the scholarly polymath. This is a kind of constructed pattern. Importantly, you may find some books which do not fit into any of your preconceived categories. You might sidestep the entire challenge of arranging these books 'somehow' by adopting a preconceived system, such as the Dewey decimal system, to group this collection of books. What I find interesting about this exercise is that the act of classifying, of pattern detection, is this: it not only produces constructed patterns but it also reveals the shortcomings of the classificatory effort. It might reveal bias, especially if you try to force books into categories where they do not quite fit. Alternatively, you may find that the book collection permits several different groupings, each of which has utility. Some of these alternate groupings may only become apparent as you handle the books and move them around on the shelves. Thus more than one pattern may characterize the book collection. A computerized database of the collection with numerous keywords may circumvent the problem of perceiving multiple possible patterns (or arrangements) in the book collection.

Applying these ideas to constructed pattern detection with proteomics data, we can see that pattern detection can be complicated. It is not simply a question of partitioning the data according to preconceived categories. Engaging with the data is necessary. It is entirely legitimate to arrange data differently and find different patterns, each qualifying as a constructed pattern. This again is largely what I mean by the suggestibility built into experiments designed according to the OES. The findings can support different patterns. Sources of bias include a Procrustean approach to classification and the failure to have contingency plans for what to do with the 'miscellaneous' category (the leftovers, "unclassifiable"). However, there is an

advantage to a standardized set of categories, namely to permit communication from one researcher to another. The limitations of standard terms for databases, as we have seen in Chapter 6 so important for validation of technical data, include eliminating data of interest on the basis of not fitting into any preset category.

8.4.3.3 Interaction patterns in proteomics

Interaction patterns are almost always indirect since they are generally developed out from amassed data of numerous experiments; however, current advances in proteomics technology may permit experimental designs generating direct patterns with respect to protein interactions. Interaction patterns deserve separate consideration because 'interaction' is such an important relation pertaining to patterns, and because protein networks are important in cell biology.

An interactome is the set of entities which interact with the entity of interest. In general, the pattern display is restricted to entities of the same class, but in theory it need not be restricted in this fashion. However, typically these are protein networks, and indeed one definition of an interactome is "full set of protein family interactions within a proteome" (Park et al., 2005, 3234). The exercise in determining such a network is a kind of mapping exercise. Methods are very diverse: structural studies examining the interaction of two proteins of interest, methods which document co-precipitation of interacting proteins, and bioinformatics (Collura and Boissy, 2007, 141). Interactomes may help elucidate the diversity of phenotypes arising from one gene (Vidal et al., 2011, 986-7). Techniques for validation of the data and its interpretation are being developed (Lievens et al., 2010, 679).

We can take the Wilson disease as an example for showing what an interactome is. In addition to the self-amplifying feedback loop demonstrated from my own work (see Figure 2.1, Section 2.5.1), we currently have a some idea of quite a few proteins which might interact with the Wilson ATPase in hepatocytes based on recent reports in the literature of individual experiments[201]: ATOX1, COMMD1, glutaredoxin 1, dynactin

[201] These reports span approximately five years of work from many different laboratories (Ko et al., 2006, 719; Lim et al., 2006a, 14006; Lim et al., 2006b, 428; Singleton et al., 2010, 27111; Materia et al., 2011, 10073; Pilankatta et al., 2011, 7389; Inesi et al., 2014, 167; Yanagimoto et

subunit 62 (also known as dynactin p62), PLZF (which stands for 'zinc finger promyelocytic leukemia protein'), clusterin (also known as apolipoprotein j), and protein kinase D. A few other proteins have been proposed as possibly interacting: these include Niemann-Pick protein C1, and possibly other COMMD family members structurally similar to COMMD1. Obviously we do not yet have a comprehensive picture of these interactions. I have constructed a candidate interactome for the Wilson ATPase (ATP7B), shown in Figure 8.5, by critically reviewing and amalgamating experimental data from various experiments in order to update a proposed interactome from 2007 (de Bie et al., 2007, 683). A distinctive and somewhat unusual feature of my candidate interactome is that it conveys functional information. It shows that different proteins interact with the Wilson ATPase in different ways, that is, by accessing different parts of the Wilson ATPase molecule. Some of the proteins interact with the 'tail' comprised of six copper-binding units at the amino end (blue squares) and some on the opposite end of the protein (carboxy end, red square). Proteins depicted may also have different functional effects including participating in some aspect of the action of the Wilson ATPase (yellow circle) or control of the Wilson ATPase protein (purple circle). Moreover, an interactome constructed similarly for the Menkes ATPase, ATP7A, (see Figure 8.3 in Section 8.1.3) reveals that the interactomes of these two closely-related metal-transporting ATPases are similar but not identical. There are other ways to construct an interactome. For example, an interactome for the Wilson ATPase can be developed by a purpose-built bioinformatics program. It seems that the outcome may differ from one mode of assessment to another. For example, the Wilson ATPase interactome developed by STRING 9.0 strategy (http://string-db.org, accessed August 2011) is not identical to the

---

al., 2011, 484) and, as noted in Section 8.1.3, a simpler depiction of this interactome was first presented by di Bie et al. (de Bie et al., 2007, 683) in 2007. As previously mentioned, an important area of current research involves examining how cisplatin interacts with the Wilson ATPase (and possibly the Menkes ATPase), but cisplatin is not a protein. It is fundamentally the metal platinum. Of course it is very interesting that a copper-transporter also serves as a platinum-transporter. Both Cu and Pt are not shown in my interactome, which is limited to proteins and peptides (such as glutathione).

one developed manually. Thus, this type of analysis poses some interesting epistemic challenges possibly reflecting my biases or the program's insensitivities.

Considering this specific problem of the discrepancy of two candidate Wilson ATPase interactomes may elucidate some of epistemic issues associated with these interaction patterns. A reasonable initial reaction might be that these interactomes are not as different as they look. At least three proteins appear in both interactomes: COMMD1, ATOX1, and dynactin subunit p62 (which is within DCTN4). However, it is

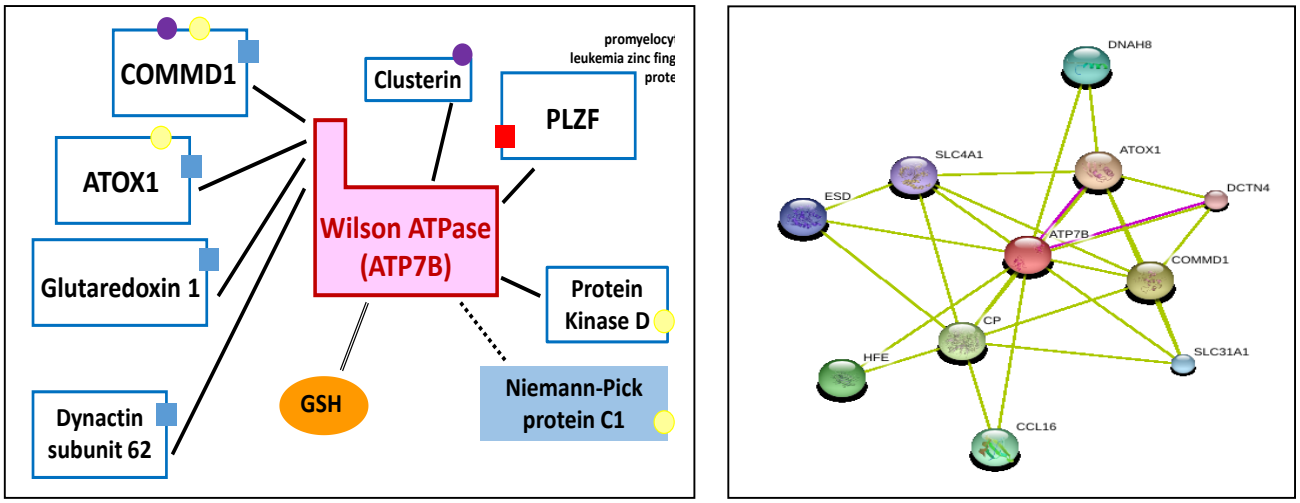

Figure 8.5 Comparison of Wilson ATPase interactome deduced from literature (left) and by the STRING 9.0 (right)

difficult to sustain this explanation. First of all, certain proteins for which published experimental evidence of interaction exists are not included in the STRING interactome but are present in the manually-constructed one. Secondly, some of the proteins included in the STRING interactome have rather circumstantial connections to the Wilson ATPase. For example, esterase D is the gene product of a gene situated on chromosome 13 not far from the *ATP7B* gene; ceruloplasmin is very important in copper utilization but not definitely proven to be an intracellular interactor with the Wilson ATPase (much the same can be said for SLC31A1, which is hCTR1, the copper-transporter in the plasma membrane). So perhaps the STRING interactome is constructed in a faulty fashion, or at least by different rules, such that there is discrepancy. In my opinion, both considerations seem to apply. Interestingly with a more recent STRING interactome for the Wilson ATPase (STRING 9.1, accessed 22

February 2014), the newer version is significantly different from String 9.0 version.[202] Some of the changes, or revisions, meet my objections about the 2011-12 version: namely, proteins whose inclusion I have questioned (like esterase D and ceruloplasmin) do not appear in this new version.

8.5 Some objections about patterns vis-à-vis the OES

8.5.1 *Always* some pattern?

This fundamental question relates to the ontology of patterns. Irrespective of why or how we see pattern(s), there remains this more fundamental ontological question, namely, whether there is always some pattern to be found. This question appears to be similar to saying that given any number of points, they can be joined by a line. Certainly, in principle, any two points can be joined by a line (or whole families of lines), and it is difficult to look at two points and suppress the impulse to connect them by a line. If a pattern necessarily involves entities in relation to each other, it becomes difficult to imagine a collection of entities within a single universe of discourse not in relation to each other.

It is worthwhile noting that numerous modern theorists have struggled with this problem, and I am not going to resolve it here. Chaos theory has investigated whether anything orderly can be said about an apparently chaotic system. Likewise, complexity theory has also sought to find some sort of order in populous and apparently disordered systems. Richard Lewontin has commented in the *Triple Helix* that both attempts have come up short (Lewontin, 2000, 113). Daniel Dennett might regard this question as a non-problem. In his view, an essential characteristic, by definition, of a pattern is to be a candidate for pattern detection (Dennett, 1991, 32). My definition of 'pattern' is not an

[202] STRING (Search Tool for the Retrieval of Interacting Genes/Proteins) is described on its website as follows: "STRING is a database of known and predicted protein interactions. The interactions include direct (physical) and indirect (functional) associations; they are derived from four sources: genomic context, high-throughput experiments, (conserved) co-expression, previous knowledge. STRING quantitatively integrates interaction data from these sources for a large number of organisms, and transfers information between these organisms where applicable. The database currently covers 5'214'234 proteins from 1133 organisms" (www.string-db.org—accessed 22 February 2014) (Anon.: STRING database). The database includes human proteins.

"anything goes" definition. It precludes the automaticity of pattern formation by requiring the characteristic of being orderly and by specifying the sorts of the relations which characterize patterns. Thus I would argue that not every assemblage of entities can be regarded as a pattern. Juxtaposition by itself does not suffice for making something patterned.

8.5.2 Influence of expectations about patterns

Whether or how our expectations may influence what patterns we detect is an epistemological question relevant to biological/biomedical research, and thus it is different from the ontological question about patterns in general which I have just considered. This interesting circularity in how we as researchers see patterns is a very fundamental built-in bias. If we honestly believe we are investigating a system and consider a system to be somehow organized (or else it would not be a system), then we examine our data derived from interrogating that system with the expectation of finding (a) pattern(s). The pattern will initially be detected by examining the data; after contextualization in the system, the pattern corroborated as isomorphic to the system provides important information about the nature of the system. Thus patterns end up with a complex relationship to systems. They also have a very primitive relationship to 'system-ship' due to the almost unavoidable expectation that a well-organized system will display patterns. Likewise, if we truly believe that proteins only function by interacting with other proteins or other biological entities, then we are pre-disposed to discern networks among our proteins identified, a similar type of circularity. Not only do we have to collect the dots before we can join them, but once the dots are collected we feel an irresistible urge to join them into a pattern. Maybe this activity is nothing more than data management.

However, to assert that *system* biology is about pattern detection requires exploring the difference between pattern detection and pattern imposition. The inherent problem with detecting indirect patterns is that the activity of classifying risks forcing the data into categories where they do not really belong. It may become a Procrustean exercise. I have suggested that the activity of perceiving a pattern this way elicits scrutiny of the action itself of classifying; however, I would accept the criticism

that actually paying attention to the need for that critical stance about classifying often gets ignored. We have seen with the interactome example (Section 8.4.3.3) that different methodologies for pattern detection may produce different patterns, some of which seem less accurate than others. As with bias as such (Section 8.5.3), It may be necessary to employ multiple viewing strategies, as well as multiple viewers, in order to keep as many options as possible open and minimize pattern imposition. The actual action of pattern detection may contain a built-in defence against pattern imposition since the fit of the imposed pattern to the data may not be quite adequate, a situation well illustrated by Kepler's ellipses being a more comfortable fit to all the data than the preceding mathematical formulations which had been regarded as the preferred solution.

8.5.3 Inhibiting effects of prior knowledge/interests on pattern-detection

We have all looked at clouds and seen extremely different entities portrayed in the very same cloud formation. Ice floes in the northern Atlantic looked like a cirrhotic liver to me but could have looked like a stone patio to someone else and tectonic plates to yet another. This is different from expecting to find a pattern: it is more nearly the penchant for finding a particularly congenial pattern. The issue here is bias in pattern detection.

The problem of bias, which is an important consideration in feminist epistemology[203], inevitably becomes an important issue in this scientific epistemology which emphasizes pattern detection. Loretta Kopelman provides a useful description of bias in terms of epistemic activities: "*unwarranted inclinations*, or those judgments, dispositions, or belief systems that interfere with any form of careful reasoning because they are untested or unreasonable" (Kopelman, 1994, 24), and she points out that bias, albeit unintentional, can arise from "scientific investigators' enthusiasms about

[203] Much of the impetus for the feminist consideration of bias in scientific research arises from numerous examples where androcentric bias adversely affected experimental design or interpretation of results or both. Such aspects of feminist epistemology of science are adroitly summarized by Jennifer Mather Saul (Saul, 2003, 232-60).

particular hypotheses" (Kopelman, 1994, 24). We deplore bias introduced through conflict of interest (usually financial, sometimes egocentric). We also need to worry about situations where the potential for bias is built into the epistemic activity itself. With respect to the OES, my focus is on sources of bias built into the OES itself, mainly in terms of difficulties associated with pattern detection.[204] Louise Antony notes that bias is inevitable in any acquisition of knowledge—some biases may be productive whereas other instances of bias are counter-productive—and evaluating such biases is itself an empirical question (Antony, 1993, 215). Thus, subjecting such bias or the research wherein it resides to critical review can serve as the best way to deal with bias.

An important part of the problem is that we tend to see what we know. As Margaret Gilbert has said, salience may be strictly psychological "in the sense that what is salient depends on who is involved" (Gilbert, 1989, 64). However, in relation to the OES the issue is not merely psychological. It is more complicated than that, largely because I have included salience as a characteristic of patterns ('a salient orderly juxtaposition of component entities, characterized by the relation(s) among them'—see Section 8.3). Thus, the issue of bias involves certain assumptions about salient patterns. For example, we may say that the salience of patterns somehow involves scientific utility. This statement turns on how we construe 'utility'. It may simply state that patterns that matter biologically are positively informative about a physiological mechanism. It could involve regarding 'utility' as some sort of epistemic virtue where 'utility' equals fecundity or maybe 'lasting significance'. My intention in making salience part of the definition of a pattern is different from these interpretations which attach great importance to salience. My intention is to call attention to the characteristic of sufficient internal relational coherence which makes the pattern capable of jumping off the page, whether it is actually appreciated or not.

[204] I have previously pointed out (see Section 8.4.3.1) that omics technology itself involves potential for selection bias: for one example, even if many thousands of proteins or DNA fragments are on a chip, they still may represent selection of those available and for another example, if your technology cannot pick up membrane-bound proteins, then the proteome will be deficient. Technical sources of bias in proteomics have received attention and remedies have been suggested (Prakash et al., 2007, 1741). In general these efforts are directed toward improving reproducibility of experiments.

There is potential for deleterious bias within the operation of the OES since bias is important with respect to experimental design and more than one type of bias may be in play. However, this is another situation where the peculiarities of a scientific epistemology diverge from the challenges associated with a general epistemology. With system-driven research (and likely with structure-driven research) an impartial arbiter of the experimental findings as organized into pattern exists: it is the system, or structure, itself which is being investigated. Admittedly, getting an adequate experimental design for the contextualization itself is critically important. Likewise, if inclusive open-endedness is disregarded, certain patterns or findings may not be considered. However, I will argue that resources are available for dealing with important forms of systemic, or inadvertent, bias with the OES. Interestingly, critical features of the OES itself appear to fulfill the components of Deborah Heikes' proffered solution to the bias paradox,[205] formulated by Louise Antony and others, which has received much attention in feminist epistemology. In her sketch of a way around this paradox, Heikes calls for implementing a concept of rationality which bridges "theoretical understanding and purposive activity"; she continues: "To be rational one must form and maintain beliefs in a reliable way" (Heikes, 2004, 330). Her elaboration of the rationality she envisages involves forming beliefs according to standards (one aspect of reliability) and such that they entail decisions/action which "respond to the world around one" (Heikes, 2004, 331). This account resonates with reliable production of experimental findings by employing optimal experimental design and adhering to standards of performance and then contextualization of those findings in the system under investigation, as set out by the OES.

An important defence against this type of bias is not making exclusive choices, that is, by keeping all the options in play until some objective reason presents itself for eliminating any specific option or it becomes evident that the solution set is actually not

[205] There are numerous formulations of this paradox, including two versions by Richmond Campbell (Campbell, 1998, 10-1, 56-9). The general form of the paradox is first to reject impartiality, thereby give up any mechanism for evaluating worthiness of competing claims, and end up accepting all accounts as epistemically equally viable, an unacceptable relativism. See extensive discussion in Deborah Heikes' paper (Heikes, 2004, 318-20).

unitary. This attitude is built into the OES as open-endedness.[206] The practicalities of a non-unitary solution set have been faced in genetics and taken in stride: this occurs with alternative splicing whereby one gene can be transcribed in several different ways to yield more than one protein, a phenomenon with the *ATP7B* gene. Similarly it occurs with viruses where the same stretch of DNA produces several different proteins because the transcription starts at different places (so-called overlapping open reading frames). With proteins, an individual protein may have multiple actions: for example, the Wilson ATPase participates in the production of copper-bearing ceruloplasmin and in the biliary excretion of copper from the hepatocytes, and in addition to those two actions it also promotes excretion of platinum from the hepatocyte. Proteins may undergo structural modifications after production to amplify their functions and they interact with different entities in the cell or plasma compartment to amplify that repertoire further. The open-endedness of the OES accords with the physiological realities of multiple roles and redundant mechanisms. The goal is to see whether patterns are there at all, with what some would regard as a rather permissive attitude toward interpreting the findings. Having multiple observers doing the analysis may enhance the analysis by utilizing in the positive way their various 'biases', that is, their different orientations from which they the data set. I am thus arguing that bias potentially associated with the OES can be managed by recognizing it, and sometimes bias can be co-opted into a constructive role in the analysis of data from a system-driven experiment.

Moreover, our experience with patterns in proteomics argues against the intellectual passivity of preferring a favourite pattern. With a proteomics experiment the data can, as it were, speak for themselves. We have discussed way to find patterns within the data, some methods more structured than others. Identification of novel

[206] Really difficult biases plague the OES but are antecedent to the sorts of bias I am discussing here. These biases arise directly from the HDM and the traditional sort of experimental design it subtends. Such biases include the unitary solution and the expectation of utter repeatability/reproducibility. These biases are typically flagged by comments like "well, of course" and diagnosed by the rolling of the eyes or the raising of eyebrows over the "crazy" ideas of the OES. The open-endedness of the OES is a direct challenge to this kind of intrinsic bias inherited from the HDM. Attacking those biases is central to the OES and to the arguments of this dissertation, but perhaps this action needs to be emphasized by my spelling it out again.

findings serves as direct evidence that you are not seeing only what you want to see. Of course, pruning the data set to favour expected findings is problematic. That becomes a clear instance of interpretive bias (not systemic), and it is not a permitted action. In addition to having multiple viewers, the widespread implementation of shared public databases for proteomics findings are at least a potentially effective defence against bias.

8.5.4 Skepticism about statistical analysis

When we analyze a possible pattern by employing some statistical instrument, we think we are sidestepping personal bias affecting interpretation of the data set. It is not obvious that statistical analysis is necessarily objective. For example, when you input data for hierarchical cluster analysis (which is often preferred for finding patterns in a large data set), there are opportunities to make decisions about how the data should be viewed by the computer: what labels might be applied, what small bits of data here and there might be grouped, what degree of stringency to apply. The procedure need not be merely the blind application of an algorithm. I believe that this sort of question has tended to be neglected in the discussion (Keating and Cambrosio, 2012, 38) of the advantages of statistics over bioinformatics, or the reverse, and it is legitimate question to investigate in the context of pattern detection. However, it is complex question deserving a separate, detailed analysis.

Here is one possibility for organizing this discussion. The important challenge with the OES is that when a dynamic system, which is relatively large and complex, is driving the experimental design, it becomes the fact of the matter is that we simply do not expect to get the very same results each time. What this recognition amounts to is the recognition that a proteomics experiment is a "one-off" experiment[207], even though we try to keep everything uniform each time we repeat that experiment. The kind of statistical treatment which might elucidate these findings is akin to meta-analysis, the

[207] This use of the term "one-off experiment" arises from a lecture (20 March 2014) by James Robert Brown at Dalhousie University. While I disagree somewhat with his characterization of all randomized controlled trials examining new drug therapies as one-off experiments, it was immediately obvious to me that this term might be relevant to proteomics experiments.

complex analysis of multiple similar but non-identical research studies. Vaudel and colleagues have the same intuition that some sort of meta-analysis tool might be effective for assigning proper importance to outliers otherwise destined to be scored as noise; however, they point out that this sort of tool does not yet exist (Vaudel et al., 2012, 527). Another possible statistics approach is analysis of variance (ANOVA) whose application is standardized and assumptions declared. However, ANOVA might filter too much data. I doubt that a network biology strategy would be applicable because of its tendency to discount poorly connected nodes, that is, small bit players, precisely the ones sought for further evaluation.

8.6 Patterns in a biological perspective

The notion that an important component of how the OES functions as a scientific epistemology is pattern detection raises questions about the nature of patterns in biological systems. These questions have previously gotten rather little attention. It may be that patterns are the effects of biological processes, as Mischler and Brandon assert (Mischler and Brandon, 1987, 399), but I do not regard myself as having enough data to make this assertion. Nevertheless we do find patterns in biological organisms, and such patterns can play an important role in the investigation of certain biological problems. When we utilize the OES, the act of looking at data from a system-driven experiment and appreciating a complex, multi-faceted pattern is a unique kind of epistemic activity. It is not the same as testing a hypothesis. Some patterns in omics literally take shape before your eyes (direct pattern detection), and other become apparent by grouping findings into relevant categories (constructed pattern-detection). A further kind of pattern appreciation involves working out the interactions between entities.

What might be called a secondary or incidental feature of patterns in biological systems is that they exhibit ‘suggestibility’. A pattern can elicit a kind of lateral thinking about a group of entities or phenomena or data. If it is incomplete, it may suggest possible ways to configure the available entities to complete it. If it looks vaguely like another pattern, it may direct attention to the similarities and differences. I regard this open quality as a kind of constructive underdetermination: it accords with the open-

endedness of the OES and also with the inherent mechanistic redundancy of biological organisms. This is an important point. In this section I will argue against an important but problematic (because reductionistic) account applied to biological entities by Barabási and then very briefly review McAllister's line of analysis in terms of multiplicity of patterns in the physical sciences, which broadly supports my analysis.

8.6.1 Reductionist patterning—network biology

Albert-Lázló Barabási and his team focus on networks, a particular kind of pattern where the relation characterizing it is 'interaction'. No biological/biomedical researcher is likely to deny the importance of networks in contemporary thinking about biological entities. Most would likely agree with his statement that "all systems perceived to be complex, from the cell to the Internet and from social to economic systems, consist of an extraordinarily large number of components that interact via intricate networks" (Barabasi, 2009, 413). Although some might be interested in taking up his call to exploit network topology, many might have some reservations about putting cells, the Internet and economic systems all into the same conceptual basket.[208] As previously recounted (Section 2.2.2), Barabási's main point is that many networks—in fact, almost all of those of interest to us—are characterized by having hubs, as opposed to a random or a hierarchical distribution of interactions. Networks with hubs are scale-free and they have non-random features; in fact, they can all be characterized by a power law distribution $P(k) \sim k^{-\gamma}$ (Barabasi, 2009, 412). What Barabási and his colleagues are proposing is that the high-quality data gained from the sort of system-driven experiments I have been discussing here (proteomics and other high-throughput technologies) be sorted into functional interacting modules according to the theoretic structure of network science. His claim is that laws which he has shown govern the formation and function of physical networks will also describe the function of biological networks.[209] An important set of observations he advances as support for this statement

[208] Others might question his assertion in the abstract for this brief paper (Barabasi, 2009, 412) that the discovery of scale-free networks was only a decade-old in 2009 (see discussion of Evelyn Fox Keller's critique in Section 2.1.2).

[209] Specifically he writes: "It [the rapidly developing theory of complex networks] has led to the realization that the architectural features of molecular interaction networks within a cell are shared to a large degree

is the analysis of metabolic networks in 43 diverse organisms (Jeong et al., 2000, 653) where it turns out that a few substrates participate in many reactions and most substrates participate in only a few: the hubs turn out to be molecules like coenzyme A or the brokers of cellular energetics such as ATP and GTP. Moreover, Barabási call attention to the modularity of networks within a cell and what he calls "small-world effect" such that a small change in one metabolite can affect others in that network very rapidly (Barabasi and Oltvai, 2004, 106), and finally he argues that these scale-free networks account for the robustness of cellular systems (Barabasi and Oltvai, 2004, 110). Although he admits that not all networks in a cell are scale-free (Barabasi and Oltvai, 2004, 106), he evidently believes that the most dominant cellular networks are.

This line of research has interesting implications for cell biology and for structural biology. A relatively early (2004) paper examined data from four protein interaction databases and found that the topology of the networks developed from each data base were scale-free (Yook et al., 2004, 930); they investigated the possibility that these networks were composed of modules themselves comprised of functionally congruent proteins.[210] This aspect of the project is a kind of pattern-finding writ large. However, the theoretical problem posed for detection of patterns in biological entities by this methodology is that ultimately the endeavour gets reduced to finding effectively only one pattern, that described by $P(k) \sim k^{-\gamma}$. This kind of reductionist view of biological patterns certainly contradicts the pluralistic view of patterns I have been developing. It poses a direct challenge to those who reject the possibility of biological laws. Certainly some patterns can be denoted by a mathematical formula as for a Fourier series, but

by other complex systems, such as the Internet, computer chips and society. This unexpected universality indicates that similar laws may govern most complex networks in nature, which allows the expertise from large and well-mapped non-biological systems to be used to characterize the intricate interwoven relationships that govern cellular functions" (Barabasi and Oltvai, 2004, 101).

[210] Their summary is: "In summary, regarding the large-scale topology of protein interaction networks all four databases display the same generic properties: they are all scale-free networks forming a giant cluster accompanied by many small disconnected clusters of proteins; they display a high degree of modularity with a hierarchical organization; and the giant cluster has a small diameter, an indication of its small world property. As these properties are derived from all four databases, they appear to be generic features of the yeast protein interaction network" (Yook et al., 2004, 932). We might speculate retrospectively that in 2003 the databases were not as complete as possible.

some must remain graphic, such as a William Morris wallpaper, in order to have impact as a pattern. Reducing cellular networks to a bunch of power laws loses a lot in the process. Moreover, the results of this exercise may not always to helpful. It comes as no surprise that ATP or coenzyme A are hubs, although it might be of great interest if they failed to be hubs in this analysis. We might try to imagine the topology of the intracellular network for copper: I suspect that CTR1, which is mainly responsible for permitting copper to enter the hepatocyte, would leading candidate as a hub. After that, what? Perhaps one of the intracellular copper-chaperones would win out or maybe an endosomal protein involved in the biliary excretion of copper, or maybe a metallothionein as a storage depot. Analyzing networks according to Barabási's lights may reveal interesting features, even interesting specific patterns. My point, however, is to express distrust for an analytical scheme which apparently aspires to a monopolizing explanation of biological systems (everything is a scale-free network). This is not the epistemic position of the OES.

8.6.2 Commonalities with McAllister's view of patterns

McAllister resides at an opposite pole to Barabási. Although McAllister is focused on the physical sciences, his analysis of patterns has much in common with mine. An important point of commonality is that we both attach importance to patterns in their diversity, and McAllister makes the important point that patterns are a highly efficient way of conveying information about a system—typically a physical system—in detail. McAllister's point is that a pattern cannot be reduced to a proposition: it is the most efficient way to depict the state of affairs. My interpretation of McAllister goes well beyond mere efficiency. A pattern is more richly textured than a proposition and thus conveys more information than a proposition. McAllister might argue that this informational richness is due to a pattern's being pictorial. Without depreciating the importance of being pictorial, I would argue that it is because a pattern is essentially relational and, moreover, it depicts, rather than describes, those relations.

McAllister's practicality in recognizing that an investigator makes choice about patterns and chooses those which play a role in his/her thinking or theorizing in some respects articulates an aspect of bias I have mentioned earlier (Sections 8.5.2 and 8.5.3).

The difference is that I regard such bias as a potential impediment to effective validation of the pattern(s) in the omics data set, for which the happy compensation is to include multiple viewpoints. In the long run, this manoeuvre may promote knowledge production. McAllister sees the difference of viewpoint as determining which pattern among many possible ones ends up being the one of interest to the investigator. There is a spectrum of possible patterns and the researcher chooses the one most useful to his/her endeavor. Thus our interpretations on this point differ, and both differ again from the viewpoint that a family or many families of lines can be drawn connecting two points and some optimal line should be chosen. Another difference in our assessment is that he makes the investigator the arbiter of importance attached to a pattern, whereas I assign that role to the system under investigation. In some situations it might be possible to recast McAllister's experimental scenario into one constructed along my terms, and thus perhaps our positions are not so far apart.

8.6.3 Implications of the OES

I have argued that system-driven research—that is, biological/biomedical research where the experimental design is governed by the system, however it is defined for that research effort, under investigation and generally featuring high-throughput technologies—operates with a distinctive scientific epistemology, the OES. The OES requires clear and meticulous definition of the system under investigation, adherence to the best consensus standards for technical performance, and analysis of data through pattern detection where those findings/patterns are contextualized in the system under investigation. I regard a pattern as a salient orderly juxtaposition of component entities, characterized by the relation(s) among them. The key features here are the juxtaposition of entities and the relations among them which endow the arrangement with orderliness: sufficient unity is achieved that the shape (the pattern) can actually be appreciated. It is possible that more than one pattern exists in any data set. The OES is tolerant of slow accrual of findings about a system, and it entertains an evaluative descriptor of "don't know (yet)" in addition to "works [in the system]" and "doesn't work [in the system]". Thus we might anticipate that the OES will generate

some interesting implications for system-driven and hypothesis-driven research. These implications are the subject of the next chapter.

CHAPTER 9 OMICS EPISTEMOLOGY (OES) GOING FORWARD

The motivation for developing a scientific epistemology for omics, many of whose disciplines are currently an important type of high-throughput biological/biomedical research, was to provide a basis for demonstrating that system-driven research could actually produce new knowledge in a strict philosophical sense. I have limited the discussion to proteomics because it is highly representative of the epistemological issues of interest and also directly related to my own research experience. As presented in Chapter 7, I have shown that the OES functions as a scientific epistemology, finding some conceptual direction in process reliabilism. I have posited that proteomics requires an epistemology where accurate, reliably-produced empirical data from a well-designed experiment constitute the initial phase of knowledge production but where the actual validation process involves the detection of a coherent pattern validated by contextualization in the system under investigation. Specifically, confirmation of a something's (in the case of proteomics, a protein's) being in this pattern relevant to the system under investigation requires some independent line(s) of evidence. In this way I can describe this process of knowledge production with the OES as thinking by looking, although 'looking' has very special characteristics. I must emphasize that the importance of pattern detection to the OES can hardly be underestimated. Pattern detection is a distinctive epistemic activity, and in the previous chapter I have analyzed it in detail with respect to the OES. Thus the OES attempts to describe for proteomics, for example, the basis by which findings, or empirical data, from a proteomics experiment can undergo a process of validation thereby being confirmed, or disconfirmed, as new knowledge, justified true belief.[211] The OES focuses on how these data thus produce scientific knowledge.

[211] It might be worthwhile to note in passing the gentle equivocation on the 'true' here. We take empirical data from a well-designed experiment produced by methods, complicated apparatus, and technicians all not given to egregious technical error, and this constitutes a "reliable" data set. So we can call it a data set which we truly believe to be "true". Yet once we harrow that data set by the justification methods of pattern-finding and contextualizing, we may end up with a smaller but more informative data set, and this one we really truly believe to constitute knowledge as 'justified true belief' because it has gone through

Alvin Goldman has pointed out that even in the context of naturalized epistemology justification is importantly normative (Goldman, 1999, 2). That assertion seems almost self-evident if we understand justification not merely as checking accuracy of data but actually assigning some sort of truth value to the empirical data taken altogether. Science itself is highly normative because it expects, which is to say it 'demands', accurate interpretation of data and accurate explanations/predictions arising from those interpretations. We might expect that the OES as a scientific epistemology would generate features we could style as 'epistemic virtues' even if we do not wish to style the OES as a virtue epistemology as such. In the first section of this chapter I will examine what might be considered an epistemic virtue uniquely relevant to the OES: novelty. It is the propensity for the OES to uncover new and highly valuable knowledge which would not have been found at all, or at least not efficiently, by the HDM. I will then explore the contention that if we attach a great deal of importance to novelty as an epistemic virtue, then we need to attach value to discovery in relation to the OES. In other words, I will revisit the analysis of context of discovery and context of justification and show that the dichotomy simply cannot hold up in relation to contemporary biology. It might be regarded as a logical organizational structure of the HDM, but it is not so with the OES. In fact, the OES eliminates the distinction. Having dismantled the context of discovery versus context of justification dichotomy, I will then show how we need to abandon another dichotomy in terms of day-to-day research, namely the strict and unbending separation of hypothesis-driven and non-hypothesis-driven research. This is a subtle issue, all the more so because this dichotomy was what launched this philosophical project in the first place. One facet of this argument is to examine how the OES encompasses an explanation/prediction duo as is typical of the

---

the validation process. My interpretation is that in this latter situation "true" has taken on a slightly different and even more favourable meaning. I might suggest that if there is a so-called *value problem* in in reliabilism, it might have to do with this subtle transformation of what it is to be true when the beliefs have been evaluated and validated as legitimate knowledge. It is worth noting that this discussion is separate from the other relevant feature of scientific 'truth': namely that it is constantly susceptible to review and revision, thus has a peculiar built-in instability, a feature some philosophers will regard as incompatible with real truth as such. Put briefly, a scientific epistemology might never actually produce truth.

HDM. Another, possibly more practical, aspect of this argument is to show that in real-life biological/biomedical research, the researcher intermingles hypothesis-driven and non-hypothesis-driven experimental designs routinely. Finally I will argue that we have abundant grounds for considering system-driven research a legitimate variety of scientific research. It does actually and unreservedly constitute "doing science".

9.1 Epistemic virtues relating to the OES

Epistemic virtues and vices constitute characteristics of explanations, which recommend them or weigh against them.

9.1.1 What are the epistemic virtues?

The application of epistemic virtues need not be limited to the realm of scientific knowledge. Nevertheless, scientific knowledge seems to be pre-eminent when we consider what an epistemic virtue (or its opposite category, an epistemic vice) is. The apparent monopoly may be due to the pervasive influence of Thomas Kuhn's masterful[212] examination of epistemic virtues in relationship to scientific theories. He regards epistemic virtues as the criteria for identifying an adequate scientific theory, and on the reasonable assumption that we want the best or most adequate scientific theory, they become criteria for theory choice (Kuhn, 1977, 322). More recently, Helen Longino has described epistemic virtues in terms of their being standards: "The so-called epistemic virtues, then, are really, at best, standards around which a cognitive community can coalesce, standards that its members adopt as theirs, but not standards that hold universally" (Longino, 2002, 185). She argues (along with van Fraassen) that no single theory or model can satisfy all the epistemic virtues maximally and you necessarily end up with a situation of trade-offs or compromise.

What is interesting, but given Longino's argument unsurprising, is how little consensus is available within philosophy of science about what is on the lists of these virtues and vices. A secondary problem is whether these epistemic virtues can be transferred without modification into biological/biomedical research since the methods

[212] It may also be masterly, but since it has been so influential I prefer to call it *masterful*.

are different from those of the physical sciences; however, in general this is not a major problem. Kuhn lists accuracy, consistency, scope, simplicity, and fruitfulness as being "standard" (Kuhn, 1977, 321-2). Elliott Sober provides an incomplete list which includes parsimony, generality, fecundity, and familiarity; in his discussion he implies that 'explanatoriness' might also be an epistemic virtue (Sober, 1990, 76). Other candidate virtues include 'non-(ad hocness)', unifying power (Priest, 2014, 217) and non-triviality.[213]

9.1.2 Novelty as a prominent epistemic value of the OES

Novelty is not on most philosophers' lists of epistemic virtues, and yet it seems to me to be one of the foremost positive epistemic features of the OES, almost important enough to be regarded as a defining epistemic virtue of the OES. Cynical critics of research grant proposals and announcements of real/impending technological progress may consider 'novel' as pure exaggeration or 'hype': a fancy, attention-getting way of saying "new". Since I demur from this interpretation, I need to spell out what I mean by 'novel' in an epistemic role. Competent research routinely turns up new information: some of this may be extension of what is already known, additional (possibly critical) data confirming held notions, or application of established facts into areas not yet studied. On my interpretation, 'novel' exceeds 'new' connotatively. It is not just a convenient synonym for breaking the monotony of repeated assertions of being 'new'. In my view, novelty as an epistemic virtue may be limited to the establishing new facts, but typically these new facts challenge received ways of thinking about the matter under consideration or open new avenues of investigation. Importantly, where 'novelty' does turn up among candidate epistemic virtues is with feminist epistemology as described by Helen Longino. She asserts that "feminists endorse the novelty of theoretical or explanatory principle as protection against

[213] Although Graham Priest specified non-(ad hocness) and unifying power, as well as fruitfulness, in his paper cited, which represents his Austin-Hempel Lecture at Dalhousie University on 14 March 2014, he did not mention non-triviality there. I have this added criterion from my recall of his lecture, and I have his written permission to include what I recall of the lecture and attribute it to him (correspondence 7-8 January 2015). I find it important to distinguish between non-triviality and generality.

unconscious perpetuation of the sexism and androcentrism of traditional theorizing or of *theorizing constrained by a desire for consistency with accepted explanatory model*" (Longino, 1997, 21, my emphasis). We differ in our concepts in that she does not accept discovery of new entities as an instance of *bona fide* epistemic novelty. We agree in the rejection of automatic submission to the epistemic virtue of consistency, although I would not reject consistency as an epistemic virtue wholesale.

To say that something is 'novel' is to say that it is new and different.[214] It might prove to be new and revolutionary, but such extreme difference is not essential to being novel. However, the actual nature of the 'difference' is important for assessing novelty as an epistemic virtue. In some situations the degree of difference is such that it challenges or overturns received knowledge. In other situations it gets at data, or knowledge, which was otherwise inaccessible.[215] Such knowledge generally falls into the category of evidence or ideas introduced by the disclaimer: "Never in a millions years would I have imagined …"! This is exactly what happened with our work on protein disulfide isomerase (PDI). The finding that PDI could bind copper was novel because it was a truly new finding, a characteristic of PDI never previously identified in the received lexicon of knowledge of that enzyme. Interestingly, it was sufficiently novel that no one was inclined to believe it.[216]

It is instructive to look at an entirely different example of novelty. This example is taken from the published clinical genetics literature (Worthey et al., 2011, 255), and, despite being complicated to recount, the clinical story is of interest partly because it

---

[214] When a biological/biomedical research assesses the significance of his/her proposed research, s/he may describe it as "novel" meaning "new and important". Of course, this begs the question as to whether novelty is an epistemic virtue. The researcher assumes that novel research is epistemically virtuous. Obviously my interpretation is different from this one.

[215] This more focused usage of 'novel' is certainly found among scientists. Dorothy Hodgkin said in relation to what I am calling structure-driven research, another type of non-hypothesis-driven research: "A great advantage of X-ray analysis as a method of chemical structure analysis is its power to show totally unexpected and surprising structure with, at the same time, complete certainty" (Hodgkin, 1965, 985).

[216] A similar situation developed when we published the case of a 3-year-old child with advanced hepatic Wilson disease. Clinical dogma at the time was that clinically apparent Wilson disease did not happen before 5-7 years of age. One reviewer doubted our diagnosis—to the point of rejecting the paper—even though we supplied genotype data indicating that the child was a homozygous for a very severe mutation of *ATP7B*, in other words, met the key definitional criterion of having Wilson disease (Wilson et al., 2000, 720).

provides some quantification of the power of novelty. A 15-month-old infant presented for medical evaluation of poor growth and malnutrition along with inflammatory bowel disease. The inflammatory bowel disease was diagnosed as Crohn disease (rare but found increasingly in very young children). However, despite very aggressive appropriate treatment, he did not improve clinically. Further diagnostic studies revealed that the entire colon was involved, with inflammation affecting the full thickness of the bowel wall, typical of Crohn disease. Over the next year, stable response to combined medical and surgical treatment could not be achieved. The treating physicians investigated the possibility (what we might term a hypothesis) that the bowel disorder was not Crohn disease but due to some sort of congenital immune deficiency. Tests for known congenital immune deficiencies, though not particularly easy to perform or interpret, revealed nothing diagnostic. With some sense of desperation, the medical team turned to a non-hypothesis-driven approach, namely exome sequencing. The exome is a subset of the genome: it consists of all the DNA segments which encode proteins, approximately 2% of the human genomic DNA. Exome sequencing attempts to sequence all of these segments (known as "exons", hence the term 'exome') and permit identification of segments which differ from the expected sequence, hence potential mutations. These variants are then scrutinized and excluded as false positives or retained as possibly of interest. Given the rarity and clinical severity, the investigators hypothesized that the disorder was recessive. Implementing this premise, they identified 70 candidate variants of which eight were both novel, that is, not found in any publically available gene data sets, and also predicted to be damaging, based on computer analysis for assessing severity of protein dysfunction due to mutation. Two of these eight were highly conserved, and one of them proved to be the abnormal gene which accounted for this child's severe bowel disorder. That gene was the gene encoding XIAP, a protein we recognize as being part of the hepatocellular copper-handling network.[217] However, XIAP has other functions and these include immune

[217] Interestingly, although the mutation seems to have disrupted the copper-binding site of the XIAP molecule, the relationship to copper metabolism was not mentioned by these authors.

regulation. The role of XIAP was confirmed in this child by confirming presence of this mutation in him and showing that his mother was a heterozygote for that mutation,[218] and by *in vitro* functional studies confirming immune dysfunction due to the resulting abnormal XIAP protein.

To recap, in this clinical example a clinical hypothesis of "the diagnosis is Crohn disease" seemed wrong (poor therapeutic response) and a second clinical diagnosis "the diagnosis is immune deficiency" could not be substantiated. Moving to a system-driven approach, exomics, 70 gene variants of interest were identified, of which eight were novel.[219] In this clinical example 'novelty' plays two roles. The one I want to emphasize relates to the epistemic power of the omics approach to problem-solving. By their own description, the clinicians had generated a list of 2006 candidate genes based on published literature, and XIAP was not on the list (Worthey et al., 2011, 261). Thus, hypothesis as the driver of experimental design repeatedly proved inadequate to the task of making a diagnosis. Moreover, the finding was consistent with known data in that patients with XIAP deficiency have been described as having "colitis", though the significance of that observation was underestimated and the colitis previously reported was far milder than in this patient. So, for all these reasons, the conclusion was that the diagnosis would not have been made without omics (in our parlance, without the OES). It can be argued that 'novelty' also functioned as an epistemic virtue or norm within the analysis because they were seeking a novel mutation. I do not think this aspect invalidates the clinical example as indicating how the OES incorporates novelty as an epistemic virtue, although it may be that our attention is drawn to the epistemic virtue novelty because geneticists attach a lot of value to it. Finally we need to appreciate the broader impact of this novel diagnosis: it raises the several specific new possibilities for

[218] Since this is an X-linked gene (recall that XIAP = X-linked inhibitor of apoptosis), her genotype is what counts; variable inactivation of the X chromosome seems to have played a role.

[219] Not to be missed, but relevant to a later section of this chapter, is the back-and-forth interaction of hypothesis- and system-driven research designs here. The exomics data were narrowed to a data set of 70 by *hypothesizing* an autosomal recessive disorder; looking for novel variants and *hypothesizing* that the resulting defect would be severe narrowed that set further. However, looking for novelty could be characterized as a kind of pattern-finding. The family genetic studies and the *in vitro* functional studies serve as contextualization of the exomics (system-driven) results.

understanding further aspects of the disease mechanism(s) of Crohn disease, including potentially a role for copper metabolism in the gut epithelium.

If novelty as an epistemic virtue means 'new and different' where *different* can range from 'generally unexpected' to 'radically at loggerheads with received wisdom'[220], then we might speculate that the extent of its being virtuous is somehow proportional to the degree of deviation from prevailing dogma. Without specifically suggesting a proportionality to degree of deviation, Longino makes the same general observations in her description of novelty as an epistemic virtue: "novelty may be conjoined with a hope of ultimately seeing or engineering an overturning of the theories with which a new view is inconsistent" (Longino, 1997, 22). My interpretation of novelty has another precedent, from perhaps an unexpected source. It connects with Karl Popper's notion of the "riskiness" or "severe testing" with respect to theory evaluation (Popper, 1968, 121). Within the operation of falsifiability, Popper tests the predictions of theory. If the prediction proves inconsistent with empirical findings, then that theory is falsified; if not, then the theory is not yet disproven. What is relevant to the discussion of novelty here is Popper's penchant or preference for examining predictions which seem counterintuitive, or even preposterous, because the theory surviving that more demanding, 'severe' test is a stronger theory (though of course not immune to subsequent falsification). The mathematical relationships here are a little convoluted:

[220] On this definition, novelty as an epistemic virtue need not be limited to the OES, even though I seem to be implying this unique relationship. A good example of novelty in a situation devoid of omics is the early 20th-century example of Archibald Garrod with alkaptonuria. (I have mentioned Garrod before, but his work bears further consideration here.) Garrod was trained at a biochemist, and he was interested in a biochemical basis for human disease. At the time, as we have seen with Wilson's own investigations of "Wilson disease", microbiology had captured everyone's imagination as the prevailing basis for human disease. Garrod focused on several rare diseases we would now term "genetic-metabolic" diseases, such as alkaptonuria. Alkaptonuria is an autosomal recessive disorder approximately 10 times rarer than Wilson disease. It is due to mutation in the gene *HGD* which encodes homogentistate oxidase, an enzyme in the metabolic breakdown pathway for the amino acids phenylalanine and tyrosine. The striking clinical finding is that the affected person's urine turns black when exposed to air. Garrod's research was greeted with a kind of benign tolerance reserved for arcane clinical research; however, he is now celebrated for his work on 'inborn errors of metabolism' and generally regarded as a pioneer in genetic medicine. What Garrod's experience has in common with omics is that the novel findings were made possible by application of innovative strategies for experimental design: namely, looking at chemistry, not microbiology.

the assertion gains epistemic strength by not being defeated by a particularly tough test. (The inverse relationship to a negative cancels out, leaving an overall positive.) I regard Popper's concept, which could be likened to the 'degree of difficulty' multiplier of a dive, gymnastic element, or aria, as similar to my notion of novelty as an epistemic virtue which confers enhanced value in direct proportion to the degree of novelty exhibited. Generating greater novelty is like passing the more severe test: it confers greater value. Of course, conferring value is not equivalent epistemically to conferring justification, which is the epistemic function of Popper's severe test.

9.1.3 Non-virtues of the OES

For the purposes of this discussion, I am regarding a non-virtue as not being the same as a vice. Simplicity does not appear to be an epistemic virtue for the OES. The OES addresses complexity: it takes complexity on, and thus it is not likely to settle for a simple depiction of a solution. In fact it might not be possible to settle for anything simple and 'elegant'. The very complexity of the system being studied militates against simplicity.[221] Feminist epistemologists again lend a support: according to Longino, they endorse 'heterogeneity' as an epistemic virtue which is specifically opposed to the "ontological simplicity and to the associated explanatory virtue of unification" (Longino, 1997, 22). The dissenting opinion from this assessment is that when a solution is no more complex than it need be, then it is simple.[222] To the extent that the final appraisal of a system is a tight coherent story, it is in fact simple, despite the apparent complexity. This line of reasoning may be unconvincing, but I would rate it is as "possible". Thus, after all, in this sense simplicity may be an epistemic virtue of the OES.[223]

---

[221] My point has affinities to Helen Longino's insightful objection to regarding simplicity as an epistemic virtue that it inevitably involves the assumption that the world is simple (Longino, 2002, 185).

[222] This is my own preferred interpretation of Ockham's razor: not that simplicity is always to be sought but that the minimum complexity is always to be sought. This interpretation permits complexity to be a feature of the preferred solution or explanation.

[223] The other dissenting voice will be from the scale-free network theorists who will argue that even the most complex system can be expressed in terms of a straightforward law-like expression which is mathematical formulae describing a power law distribution. This objection is cogent only if you believe that most complex systems are scale-free networks. I have reviewed reasons to think this is not the case

Contrariwise, it is hard to rate 'triviality' as an epistemic vice of the OES.[224] This discussion turns on what is meant by being 'trivial'.[225] One biologist's trivial issue is another biologist's *raison d'être*. The function of one obscure protein or another is not an instance of trivia. It is more a matter of misplaced or unrecognized significance. The OES tends to gravitate toward and generate such problems. Since it is nearly impossible to adjudicate any research endeavour or finding as trivial, it is awkward to rate triviality as a vice with respect to the OES. Indeed the argument could be made that for the OES triviality is an epistemic virtue. Finding crucial 'trivialities' which otherwise get ignored or overlooked may be one of the important contributions of omics.

9.2 Rehabilitation of 'discovery'

Among the initial issues considered (see Section 3.3) was the assertion that the OES resides within the context of discovery as opposed to being part of the context of justification where doing 'real science' takes place. I rejected this notion vigorously and recounted with some disdain my colleagues' convention of referring to omics experiments, effectively guided by the OES (although of course it would be anachronistic to say so), as 'hypothesis-generating". This description was overtly self-defensive: it kept any reviewer or granting-agency bureaucrat committed to an irrevocable connection between hypotheses and real science effectively at bay. I went on to argue that the demarcation between context of discovery and context of justification is extremely indistinct in contemporary biological/biomedical research. That discussion served to create a certain parity between context of discovery and context of justification, effectively eliminating that distinction. The issue I want to discuss here deals with 'discovery' itself. If novelty is an epistemic virtue, can we disregard the legitimate role of discovery in research directed by the OES? In other words, I want to

in Chapter 2 (Section 2.2.2). A-L Barabási has used this sort of mathematical argument to show coherence of theories arising from his concept of network science (Barabasi, 2009, 413).

[224] Triviality was identified as an epistemic vice by Graham Priest in his Austin-Hempel Lecture at Dalhousie University on 14 March 2014. See footnote 213 for details.

[225] Peter Schotch pointed this matter out to me specifically on 18 March 14, but I think my comments here about the shiftiness of what qualifies as being trivial already address this point.

formulate and investigate the claim that 'discovery' is the very business of scientific research. The implied discrediting of *discovery* by the Logical Positivists' dismissive 'context of discovery' is not merely misleading but downright inaccurate, at least with respect to current biological/biomedical research. It may simply be aggrandizement to talk about research establishments as "discovery districts", but that marketing ploy may actually have its basis in an important aspect of contemporary biological/biomedical research. Discovery has assumed newfound importance.[226] I will review my arguments briefly about why the context of discovery versus context of justification dichotomy does not hold up, and then I will examine what we think about discovery nowadays, as opposed to 80 years ago. An associated question is how omics epistemology may actually reshape (and/or dispose of) the context of discovery versus context of justification dichotomy.

9.2.1 Context of discovery as opposed to context of justification

The distinction of context of discovery and context of justification dates from the 1930s when Reichenbach asserted that the psychology of scientific discovery is entirely different from the logic of science, whether that logical inference involved deductive or inductive reasoning (Reichenbach, 1938, 36-7). As I pointed out in Chapter 3, what is impressive about these concepts is their longevity and persisting power to capture philosophers' attention and endorsement. In particular their respective connotations endure: 'context of discovery' being preliminary to real science and somehow unregulated and 'context of justification' being the domain of serious science. Early

[226] Indeed, many *avant garde* biological/biomedical researchers currently see themselves as engaged in "discovery science". Using a service of the National Medical Library via PubMed, I find that the keyword phrase **discovery science** finds 2588 papers in 2014, as opposed to 483 papers in 2004, as opposed to 176 in 1999 when proteomics was truly in its infancy, as were other omics disciplines (http://www.ncbi.nlm.nih.gov/pubmed, accessed 11 November 2014) (Anon.: PubMed). It should not be lost on us that the major group pushing Creationism/Intelligent Design in the United States calls itself the "Discovery Institute", a very slick operation as shown by its extremely impressive and well-designed website. To be engaged in discovery is now a favourable activity. Their adoption of this term is a tribute to its current cachet. Likewise, but completely separately, a popular current style of teaching is styled as 'discovery' where the child works out arithmetic (discovery arithmetic) or narrative writing (discovery spelling) on his own. This educational programme has numerous detractors, but my point here has to do with the general emphasis on discovery, not whether discovery arithmetic is an efficient way to learn how to multiply numbers.

rejections of this dichotomy came from Hanson and Feyerabend on philosophical grounds akin to those of the Logical Positivists; more recent rejections stem from feminist epistemologists of science who hold that norms which operate in the context of justification also operate in the context of discovery. By way of evaluation and revision, the product of the context of discovery has already been subjected to the same norms and assumptions, which operate in the context of justification. The distinction between context of discovery and context of justification must be really strict and sharp for this dichotomy to have any real utility. Showing that the dichotomy is not strict weakens it considerably and, as I have argued previously, opens up the possibility that the context of discovery versus context of justification dichotomy might be irrelevant for omics.

9.2.2 Context of discovery: "dream a little dream of me"

The important point to be made is that to a great extent the dichotomy between context of discovery and context of justification follows from the logical form of the HDM as it was originally conceived. It is impossible even to conjecture that Logical Positivists in the mid-20th century would seriously have considered the possibility of carrying out scientific research without designing it around a hypothesis. Thus, with their highly structured world-view, distinguishing between context of discovery and context of justification makes sense. Truly scientific activity involves working through a hypothesis logically. The 'context of discovery' is the answer to the question "where did you dream up that hypothesis?"—an interesting but objectively irrelevant question. Relinquishing the hypothesis-driven format entails giving up this dichotomy. Choosing to examine a certain system, as in system-driven research, or the structure of a specific molecule, as in structure-driven research, does not locate the researcher in a context of discovery because systems and structures are not dreamed up. While making that choice may emanate from raw curiosity, often its rationale has to do with understanding or solving some other problem, as with our work where sorting out the network of proteins involved in the normal disposition of copper in hepatocytes would be a necessary first step to understanding the disruption of that network in disease. However, once we engage in non-hypothesis-driven research we leave the province of context of discovery and context of justification.

9.2.3 Discovery as a 'deliverable'

Moreover, if the OES regulates looking at a system to understand it more comprehensively and the outcome of that epistemically orderly activity is to identify new and possibly unexpected features of that system, then 'discovery' becomes a deliverable of this research, not something which is preliminary to it. Consequently the dichotomy becomes nonsensical. The deliverable, which is basically new knowledge, is possible through the justificatory mechanics of the epistemology, namely the OES. An implication is that the OES undergirds exploratory experimentation, but here the meaning of "exploratory" has been altered, both from common parlance among scientists and from the technical meaning recently imported into philosophy of biology. Instead of being preliminary as with an exploratory surgical operation or merely hypothesis-generating, my version of *exploratory experimentation* is a proper style of research within the broad domain of biological/biomedical research, encompassing a scientific epistemology conceptualized in the OES.[227]

Another reason for finding renewed positive interest in discovery as such is that it is the important deliverable with "data-driven research" where existing databases are trolled for new and significant information previously overlooked or not sought. Given the tendency to terminological confusion—where some use the term 'data-driven' for what I call 'system-driven'—it is important to note this similarity. Another common use for the term 'discovery' in current biomedical literature has to do with drug discovery, which depends heavily on system-driven research. Some might argue that drug discovery is a hybrid between system-driven research and data-driven research. In any case, discovery of the sort associated with the OES is the output in the domain of pharmacology.

---

[227] My version of exploratory experimentation is more specialized than Steinle's and thus obviously not equivalent to his conception of exploratory experimentation. However, as I have previously argued (Section 5.2.5.2), I believe that Steinle's conception of exploratory experimentation has been misinterpreted when imported into omics as 'hypothesis-generating'. Steinle's specifically states that his concept is independent of considerations relating to context of discovery and context of justification (Steinle, 1997, S71). Steinle appears to be talking about research moving in new directions both empirically and conceptually. Galileo's experiments regarding free fall have been cited as an example of this sort of exploratory experimentation (Van Dyck, 2005, 873).

In summary, then, discovery has lost the pejorative connotation invested in it by the strict dichotomy between context of discovery and context of justification. We might allow the Logical Positivists some slack in terms of how their invented dichotomy might make sense in relationship to their view of scientific activity as it was in their era; however, effective arguments exist for why the distinction is faulty even in connection with the HDM. With the OES, the dichotomy simply cannot exist because justification and discovery are intertwined. We may harbour some minor persisting resentment that the Logical Positivists distorted the meaning of 'discovery' when they appropriated to their technical jargon, but as a contemporary buzz-word 'discovery' it still attracts distorted meanings avidly. The more important point is that making discoveries is one of the reasonable outputs of an epistemically well-founded and perspicacious scientific endeavour.

9.3 Explanation/prediction typical of HDM is possible with OES

Biology may not deal in laws, as physics attempts to do, but it seeks mechanistic explanations. With the HDM the search for an explanation is direct. The hypothesis is postulated as a possible explanation; the hypothesis is then reformulated as a test-implication and respective the experiment is designed to see whether the predicted outcome relative to that explanation holds. Thus explanatory knowledge is acquired through what I have called the 'explanation-prediction duo'. The duo arises because explanation and prediction seem to be the same process viewed from differing temporal vantage-points.[228] In the physical sciences a theory, which is composed of relevant facts, circumstances and laws, has attained some success if it explains the facts at hand. The researcher then applies it prospectively and sees whether it explains similar facts gathered in the future: this is a predictive action. For the biological/biomedical research something similar operates with a biological mechanism. I formulate that mechanistic explanation based on available data and then apply it with other similar instances to see whether it still holds. Failing the possibility of gathering such prospective data, I may

[228] For this reason, some would insist that it is an explanation-prediction *unity*.

turn to a data set assembled previously, though separate from my original set perhaps because it was amassed by someone else, and see whether it works with these data. This situation poses interesting questions: am I engaged in an explanatory exercise because the data are "old" or a predictive one because the data are new to me? I would argue that it constitutes a predictive action on diverse 'old' data sets gathered independently. In biological/biomedical research these old or independent data sets are often called "validation sets". Incidentally this nomenclature accords with the claim that explanation and prediction are closely related aspects of validation.

As I have argued throughout this dissertation, with hypothesis-driven research the complementary relationship of explanation and prediction is fairly straightforward. Both explanation and prediction relate to how the hypothesis functions and therefore to the status of what the hypothesis purports to establish. A hypothesis which was not predictive would not conform to either of Hempel's models. Given the logical structure of those models, especially the D-N model, it would be difficult for a hypothesis to be consistently predictive without being explanatory, although it could happen. Phlogiston might be a good example of such a disjuncture in the physical sciences.

The situation is more complicated with the OES. In terms of system-driven research it is more difficult to identify coordination between explanation and prediction. With the OES, an experiment governed by the system under investigation entertains no specific predictions as to what will be found; however, once data are obtained, they need to be evaluated within the context of the system. Do the findings fit into the system as it is understood? Do they satisfy the critical features of entities belonging in that system? Viewed as a solution set, we may find that this set includes entities not previously suspected to be there. Recognizing these components of the system may lead to new appreciation(s) of how that system works. Thus the OES has potential for generating explanations and predictions by disclosing features of a system not previously envisioned either because a hypothesis can be limiting or biasing or because contextualizing the data may reveal previously unperceived 'holes' in the system, areas where predicted elements still have to be identified. The coherence of the system may

point toward, or indeed predict, the nature of those elements. It is admittedly a complicated scenario: unpredicted results relating to components of the system generate new possible explanations about how the system works. Thus I regard the OES as having an indirect relationship to explanation.

In terms of explanation/prediction the really interesting epistemic action of the OES is to furnish unexpected (or unpredicted) data which advance the research project because they are indeed new knowledge. As we have seen, proteomics can identify elements of a system which were unexpected. For example, it was not expected that the enzyme PDI (protein disulfide isomerase) would be a member of the set of hepatocellular proteins which can bind copper. (In retrospect, it could or should have been obvious from its structure.) Exactly how PDI fits into the network of proteins responsible for the disposition of copper in a liver cell remains undetermined—no one has picked up this line of research—but one possibility is that it is intimately involved in the production of copper-containing ceruloplasmin. So the potential exists for predicting new functional aspects of that system. It could be said that filling in a pattern constitutes a predictive action. It makes manifest connections which were only virtual. I accept the anticipated criticism that this interpretation seems like making a virtue out of necessity: if the experimental design entertained no predictions, then in a sense anything found is unpredicted. However, the comparison to expectation relates to what is already known about the system in terms of its components or how it functions, and thus the finding relates to background data rather than anything postulated, or not postulated, in the design of the experiment.

Finally the action I have described as 'suggestibility' lies within the domain of predictive functions.[229] Pointing to the possibility of new or alternative patterns is a kind of predictive posture. This gets at the positive value of what I have called constructive underdetermination. This is not quite the same as finding holes in the pattern and

[229] Suggestibility was described in Section 8.2.2 as a secondary or incidental feature of patterns whereby a pattern can elicit a kind of lateral thinking about a group of entities or phenomena or data. An incomplete pattern may "suggest" possible ways to configure the available entities to complete it. A pattern which looks vaguely like another pattern may direct attention to the similarities and differences so that new concepts arise.

predicting what entities might fill those holes. It relates to the description of the pattern itself. At some point in the experimental work, it may not be apparent which of several patterns may pertain. This flexibility permits a broader assessment of the issue under investigation. As such, it permits several different patterns to be considered until accumulating data constrain the results to one pattern or indicate the parameters within which one pattern or another might operate. I regard this open-endedness is an important feature of the OES. It is possible that a pattern thus arrived at might serve as a mechanistic explanation, but given the variability and redundancy of most biological systems, it is important to see the output as possibly being multiple patterns, each with a recognized context in which it operates.

9.4 Interplay of OES and HDM

For many reasons, intuitive and conventional and—not least of all—logical, we regard the contrast of hypothesis-driven to non-hypothesis-driven research as the juxtaposition of two mutually exclusive categories. This assessment is supported by the Logical Positivists' attitude toward science. I believe it starts to break down when we find that the categories are not equally balanced: there are many more types of non-hypothesis-driven research than there are hypothesis-driven. In fact, the problem here is that we conflate "mutually exclusive" with "mutually excluding". In the real world of every day biological/biomedical research, taking place 'at the bench', this partition between hypothesis-driven to non-hypothesis-driven research styles does not hold up and may be pernicious. Biological/biomedical researchers are comfortable moving from one type of research to another depending on either the problem they are investigating or the specific aspect of the problem they are investigating.[230] Perhaps surprisingly (and potentially confusingly), methodology associated with system-driven research can be imported for performing an experiment which is hypothesis-driven. Consequently I want

[230] This is different from my observation that researchers engaged in system-driven biological/biomedical research typically work within a conceptual superstructure of overarching hypothesis and hunch hypothesis, as discussed in Chapter 4. As will be shown in some detail, the situation here is the utilization of methods classically associated with system-driven, or even structure-driven, research in the service of a hypothesis-driven experimental design.

to examine this complex relationship between hypothesis-driven and non-hypothesis-driven research.

One reason that the strict differentiation between hypothesis-driven and non-hypothesis-driven research does not work well as the way to characterize the research done by HDM, on the one hand, and the OES, on the other, is that the differentiation is not clear-cut in actual practice. It seems clear-cut if you style HDM as narrow-focus (bottom-up) and OES as big-picture (top-down). This exclusive classification obscures an important feature of scientific research. It could be a kind of category error. If a research methodology is defined as being hypothesis-driven, then clearly all that kind of research utilizes a hypothesis as a central feature of its experimental method. But, for system-driven research, the exclusivity is not so obvious: in fact, I believe it is not there. While hypothesis is not an intrinsic feature of all research methodology, there can be situations where methodologies typically associated with the OES can serve a hypothesis-driven experimental design. I will illustrate my argument with two examples, one relating to genomics and the other to proteomics. In both cases the interactions of HDM and OES involves subtleties which have to be teased out.

If I want to try to find the abnormal gene in a disorder which by all clinical characterization appears genetic, then a good way to proceed to find a very large multi-generational family affected by the disease. Affected members of the kindred must be accurately diagnosed and clearly described. The researcher characterizes the genome of each affected individual in the kindred (by determining haplotype profiles based on expression of minor silent variations in the genetic DNA, called single-nucleotide polymorphisms, or SNPs) and then s/he looks for commonalities among those numerous variations. In fact, the most fruitful comparison is between members of the large kindred spread as far apart as possible, for example, 3-4 generations apart such as newborn infant and great-great-grandmother. This experimental design uses genomics (ostensibly an OES approach), which happens to be the best way to investigate the proximate hypothesis that a specific genetic defect is present in this kindred.

A similar procedure is called a 'genome-wide association study' (GWAS). With GWAS, the genome of every subject in a carefully-defined group of genetically unrelated individuals, who typically all have the same disease, is characterized for the purpose of finding genetic commonality. By its nature, this group is not as strongly connected as the multi-generational family exhibiting exactly the same disorder through several generations. The multi-generational family situation suggests that an experimental design associated with system-driven research can be employed in the service of a proximate hypothesis. With GWAS we need to be cautious: there may be a subtle difference compared to the kindred investigation. This expectation of genetic commonality seems to be some sort of hypothesis, but the question is whether it is a proximate hypothesis. I regard GWAS as taking a different approach. The researcher asserts that there may be a definable genetic component to this disease but then characterizes the system (the cohort of biologically unrelated individuals who happen to have the same disease) to see whether any coherent genetic patterns emerge. Thus it serves as a hunch hypothesis, on my nomenclature. What governs the experimental design is the system. This becomes evident in the ramifications of how the research is actually performed and evaluated: the system must be completely and meticulous described, the techniques are high-throughput, and patterns are sought which must then be contextualized. The nuances of GWAS do not detract from how a similar experimental design can be employed for the more circumspect kindred-based experiment with a proximate hypothesis. However, with the collection of unrelated individuals affected by a disease, having a proximate hypothesis is not a feature of the research design.

A roughly similar situation exists with "comparative" proteomics. Comparing proteomes has become an increasingly attractive way to design proteomics experiments. Thus, for example, I could propose investigating the Cu-metalloproteome in normal mouse liver because it would advance my understanding of normal physiology. If I then proposed repeating the experiment in the liver of a Wilsonian tx-j mouse, I might postulate that I would compare it to the findings in the normal mouse

liver where the Cu-metalloproteome was determined under similar conditions and in temporal proximity. At this point I have set up a controlled experiment. That qualitative experiment actually has a definite proximate hypothesis: the hepatic Cu-metalloproteome in a 4-month-old Wilsonian tx-j mouse differs from that in a 4-month-old normal C3H mouse raised under identical conditions.[231]

The subtlety of experimental design is somewhat easier to see if we imagine a comparative *quantitative* proteomics experiment. We want to interrogate the system as to which proteins in the 4-month-old tx-j mouse hepatic Cu-metalloproteome are elevated more than 10x normal, that is, an order of magnitude over normal. Being quantitative, this experiment involves comparison to the normal C3H control, but that in itself is not a hypothesis-driven experiment. This is the same issue we found with the gene chip and the quantitative proteomics experiment: achieving quantification inherently involves comparison. However, we could design a hypothesis-driven quantitative proteomics experiment. It would be rather complicated but it might look like the following. We have observed that the 2-week-old tx-j mouse with a defective Wilson ATPase behaves as if it were globally copper-deficient and thus broadly resembles the 'brindled' mouse with a defective Menkes ATPase, which is indeed globally copper-deficient, due to its inherited defect in the gene *Atp7a*. Our proximate hypothesis is that the pattern of hepatic proteins elevated an order of magnitude at three weeks of age is extremely similar (in effect, *identical*) in both mouse strains. So we determine a quantitative hepatic Cu-metalloproteome for the 2-week-old tx-j mouse (this necessarily involves determining the same for a 2 week-old C3H mouse) and then we also determine a quantitative hepatic Cu-metalloproteome for the 2-week-old brindled mouse (this involves determining the same for a 2 week-old normal mouse of

[231] With the increasing availability and sensitivity of quantitative methods for proteomics, a quantitative comparison might be sought. Such a quantitative experiment poses some theoretical design problems. Quantifying the proteins in a proteome inherently involves comparison to a 'control'. Thus it is not exactly a controlled experiment without being circular—not a preferred experimental design. Thus the quantitative experiment becomes a matter of characterizing a proteome quantitatively. It remains system-driven. Of course it is worthwhile to quantify the effect of an intervention on a system where such an intervention could be environmental or developmental or disease-related. These experiments relate to a hunch hypothesis—that the intervention appreciably (that is, quantifiably) alters the system.

the same background strain), and then with these quantitative results from each mouse strain we compare the which proteins are elevated at least 10x to test our proximate hypothesis. With this latter quantitative experiment we actually do numerous proteomics experiments in order to test a hypothesis.

These examples may seem arcane, if not far-fetched. However, as we saw with the actual example of the child with congenital XIAP deficiency resulting in failure to thrive and severe colitis resembling Crohn disease, we may employ a non-hypothesis-driven strategy like exome-sequencing and then implement hypotheses into the analysis (in this example, first, if extremely severe and rare, then necessarily autosomal recessive; secondly, highly conserved). This seems like ordinary operating procedure. Choice of experimental design depends upon the problem requiring solution. HDM did not work for the initial stages of solving this diagnostic problem, but OES (as exome-sequencing) did. Additionally, the findings did generate other hypotheses about the actual occurrence of an immune defect demonstrated *in vitro*. Finally, a hypothesis relating to consistency of observations with previous instances was generated and substantiated. Taking this consideration further, we can see that that contextualizing data from an omics experiment may display some features of taking on a hypothesis or generating one. But, as is clear from the clinical example, research performed with the OES is not 'non-scientific' while the HDM experiments are. The two approaches are complementary. If we consider biological/biomedical research from this broader viewpoint, I submit that a hypothesis becomes one instrument of how this research can be done, not the exclusive instrument. This notion would seem quite commonsensical to many biological/biomedical research scientists working at the present time.

9.5 What is 'doing (real) science'?

The original (and generally considered lethal) criticism of omics, specifically including proteomics, has been that it is not real science: only preparatory toward doing

real science.[232] Proteomics, for example, involved sorting out the system, lining up the *dramatis personae*, and generally clearing away the chaff so that real (that is, hypothesis-driven) scientific research could be performed. I have attempted to defeat this criticism by various gambits, most importantly by showing that omics has an epistemological structure (the OES) supporting it, just as hypothesis-driven research does (the HDM). If we are going to evaluate whether omics constitutes 'doing (real) science', then we need to have some idea of what 'doing (real) science' may actually be. On the basis of my foregoing discussion and examples, I propose the following description of "real science": an investigative activity which seeks to generate explanatory and predictive information about the natural world by means of methods which are both transparent and capable of standardization. Methods are transparent when they can be articulated in detail. They are thereby rendered capable of standardization, but standardization also requires the critical review and consensus of the relevant scientific community. Coupled with technological sophistication and excellence, such standardization permits evaluation of data. Experimental data can then be assessed for technical accuracy. This is not the whole story, however, since the OES requires validation of the empirical data within the system under investigation. This composite validation serves as a justificatory move. Thus I would argue that production of new knowledge occurs through the collusion of an epistemological apparatus with excellent technology itself in service of excellent experimental design. I have developed this apparatus as the OES. The information produced by omics disciplines such as proteomics can qualify as new knowledge after being harrowed by process of justification. Thus I support my claim that for system-driven research the OES indeed produces actual scientific knowledge and conclude that for system-driven research is indeed "real science".

[232] I realize that in this very sentence I use 'science' ambiguously: first as a process (epistemological entity) and then as a product (ontological entity). However, a detailed analysis of this flexible usage, though extremely interesting, does not seem really apposite to the argument of this section. Science can be a process and it can be something we produce when we engage in that process, and either way it is open to critical epistemological analysis. My definition of science conflates these two aspects by calling it an activity.

CHAPTER 10 CONCLUSION

In this dissertation I have asserted that we are in the midst of an important and far-reaching revolution in how biological/biomedical research is designed and performed. Thanks to astounding technical and conceptual advances, we have gained the ability to investigate biological entities directly, in their entirety, as complex systems. These investigations are highly detailed. They benefit from technically sophisticated high-throughput methodology combined with our ever-expanding capability to store and analyze huge data sets through complex bioinformatics techniques. Broadly speaking, these research strategies support the development of contemporary systems biology. For many biological/biomedical researchers, however, the enterprise is somewhat more circumscribed than doing formal systems biology. Instead it involves employing in their research some sort of omics, such as genomics, proteomics, transcriptomics, metabolomics, and exomics. Omics critically enables much contemporary biological/biomedical research, which researchers themselves regard favourably as 'discovery science'. In this dissertation I have focused on proteomics. Typical of these high-throughput omics disciplines, proteomics rejects the received notion that scientific research is exclusively hypothesis-driven. Omics experiments typically feature no hypothesis, and neither does the typical top-down "shot-gun" proteomics experiment. Such an experiment entertains no predictions about what findings will emerge, and yet the patterns revealed in its data may constitute—and indeed frequently do constitute—new knowledge.

The resulting situation is highly problematic for both scientists and philosophers. Omics formidably challenges the assumption that hypotheses are required for credible experimental design in biological/biomedical research. Scientists schooled in the necessity of hypothesis to experimental design are left feeling disorientated. Should they happen to care about theoretical aspects of biological/biomedical research design, they also feel unsupported at the conceptual level. Philosophers raised on the hypothetico-deductive method (HDM) are in a quandary, having no alternative

epistemological strategy to contribute. Philosophers of science who have already questioned the omnipotence of the hypothesis for credible research design may nevertheless be surprised at the broad reach of this philosophical problem with the advent of omics.

10.1 Omics Experimental Strategy

My starting-point has been that biological/biomedical research can be either hypothesis-driven or non-hypothesis-driven. What counts currently as the conventional approach to experimental science is hypothesis-driven: examining a well-articulated hypothesis by a purposely-designed, expertly-performed experiment. I have shown that the category of non-hypothesis-driven research is comprised of several subcategories including system-driven research, data-driven research, and structure-driven research. Omics research is within the system-driven category. Being system-driven, omics identifies relevant features of a biological system, otherwise not accessible via HDM. The epistemic strength of omics, and in fact all of system-driven research, arises from providing novel, detailed, inclusive accounts of natural biological systems. Based on my experience with proteomics, I have developed the Omics Experimental Strategy (OES) as a scientific epistemology for system-driven research. I have proposed that the OES involves a well-designed experiment specifically with exhaustive description of the system under investigation, data produced competently by reliable technology, and then contextualization of the findings back within the system. The manoeuvre describing such contextualization is neither “big-picture” (“top-down”) nor “narrow-focus” (“bottom-up”); I have called in “in-fill + zoom”. The findings, when contextualized within the system, provide new insight into how that system functions. This epistemological structure takes its inspiration from process reliabilism but augments that general epistemology with a detailed description of validation. For the purposes of system-driven research, the validation takes place within the system of interest. This is not an outlandish epistemological approach: for hypothesis-driven research, validation takes place in relation to the hypothesis driving the experiment. As demonstrated by specific examples from recent research (see previous chapters and also the example

below), with omics—certainly with proteomics—some of the findings prove to be entirely novel and unexpected. For this experimental strategy (OES), actual novelty serves as an epistemic virtue.

Pattern detection plays a critical role in the validation process of the OES. The production of scientific knowledge via the OES involves an epistemic manoeuvre which can be described as "thinking by looking". This manoeuvre involves letting or getting data to resolve into a pattern or patterns, prior to contextualization. I have defined pattern as generally as possible: "a salient orderly juxtaposition of component entities, characterized by the relation(s) among them", without any connotation of imputing archetypal or law-like behaviour to a pattern. Being in relation to something else elicits the intuition that there are relation-makers: for patterns, according to my conception, these include prominently recurrence and integration. In omics, particularly proteomics, pattern-detection can be direct—as with a visual colour read-out—or indirect, for example, constructed via classification of results.

Finally, an important point which emerged in this discussion is that with high-throughput research methodologies the dichotomy of hypothesis-driven or non-hypothesis-driven is operationally not as iron-clad as it first appeared (or indeed as it appears logically). Some omics experiments can be designed which are the best way to serve the purposes of a hypothesis-driven strategy as I have defined it. Contextualization in a system often takes advantage of a body of knowledge gained through extensive hypothesis-driven experimentation. The take-home message is that both hypothesis-driven and system-driven (such as omics) research are legitimate ways of doing scientific research. Both can be shown to have epistemological undergirding. Scientists and philosophers need to feel comfortable with both the hypothetico-deductive method (HDM) and the omics experimental strategy (OES).

10.2 Another illustrative example

I have illustrated these scientific and philosophical issues with examples from my own research experience relating to Wilson disease and metalloproteomics. Here I would like to recount a different example recently brought to my attention. It illustrates

many of the points which are critical to my arguments in this dissertation. It shows that the issues I have been examining are not limited to a highly specialized subdiscipline of proteomics (that is, metalloproteomics) in service of a genetic disease found world-wide but with sufficiently low incidence to qualify as "rare" (namely, Wilson disease).

In November 2013 the paediatric hepatologist David Perlmutter[233] gave a keynote lecture at the annual meeting of the American Association for the Study of Liver Diseases. His topic was the genetic disorder α_1-antitrypsin deficiency and its treatment. The occasion for the lecture was the 50th anniversary of the identification of the protein α_1-antitrypsin and its relationship to human disease.[234] This protein is manufactured almost exclusively in the liver. It is released into the bloodstream where it serves to neutralize toxic chemicals produced by white blood cells, namely, their chemical artillery against infecting organisms. In this disorder a gene mutation leads to production of an abnormal α_1-antitrypsin molecule which polymerizes and gets trapped in the liver. The consequent lack of α_1-antitrypsin in the blood leads to severe early emphysema, which can be avoided by regular infusions of biologically engineered α_1-antitrypsin. The retention of α_1-antitrypsin in the liver can lead to chronic liver disease and liver cancer. Infusions of biologically engineered α_1-antitrypsin do not abrogate the liver disease. David Perlmutter's research has addressed the cellular pathophysiology of α_1-antitrypsin deficiency of the liver. He introduced his lecture, which would focus on new treatments for the liver disease, with an air of discomfiture bordering on embarrassment. He explained that he would discuss two lines of experiments: the "old way" which constituted his research for the past 20+ years, and the "new way" which employed high-throughput techniques. Through years of beautifully designed

[233] Dr. Perlmutter was the recipient of the Andrew-Sass-Kortsak Award of the Canadian Association for the Study of the Liver/Canadian Liver Foundation in 2006. This award is given in recognition of major achievement advancing paediatric hepatology. Andrew Sass-Kortsak was a paediatrician at the Hospital for Sick Children in Toronto. Arguably one of Canada's first clinician-scientists, he was certainly an early example of an academic physician actively engaged in and promoting translational research. He was an internationally-recognized expert on Wilson disease.

[234] Diane Cox, the Canadian geneticist who would go on to identify *ATP7B*, the gene abnormal in Wilson disease in 1993, first made her international reputation as an expert on human chromosome 14, and specifically the genetics of α_1-antitrypsin deficiency. Until fairly recently, her laboratory functioned as a key international reference laboratory for characterization of serum α_1-antitrypsin (known as PI-typing).

experiments on a mouse model of α_1-antitrypsin deficiency, which Perlmutter and those working in his laboratory had created and validated, he had found that the mechanism of the liver disease involved a defect in autophagy. Autophagy is not an easy concept to grasp, but basically autophagy is the process by which a cell puts its defective or outmoded parts into a recycle bin. These hypothesis-driven experiments eventually generated a hypothesis about treatment: if you could enhance autophagy, then you might ameliorate the liver disease. Accordingly they looked around for a drug with that effect and eventually tried carbamazepine, an anti-epileptic drug which happened to be known to have an enhancing effect on cellular autophagy. It decreased the retention of α_1-antitrypsin in the liver and reduced liver damage in the mouse model of this disease (Hidvegi et al., 2010, 231). A clinical trial in humans has recently been initiated. So much for the "old way" of doing research: ostensibly this hypothesis-driven research has been very successful. The "new way" of doing research arose through a (chance?) collaboration with a research institute colleague. The general idea was to create a model of α_1-antitrypsin deficiency suitable for screening a collection of 12,000 drugs to see whether any of them worked to decrease the intracellular retention of the abnormal α_1-antitrypsin. Accordingly, a model of α_1-antitrypsin deficiency was created in *C. elegans*. The worm lacks a liver. The gut-lining cells could be used as surrogates, and when the abnormal α_1-antitrypsin accumulated in them, it was clearly evident as big globs (Gosai et al., 2010, 6 of 17). These engineered *C. elegans* worms still proliferate at a rapid rate. It was then possible to expose these α_1-antitrypsin-deficient worms to all of the drugs in the collection and measure decrease in the size of the globs. Approximately 6000 compounds could be screened per day (Gosai et al., 2010, 10 of 17). As recounted in his lecture, a few drugs (five) were found which had the desired effect: one of them was carbamazepine but the others were "surprises". Some of these drugs had structural similarities. Appreciating this structural pattern generated the hypothesis of screening for drugs with a similar structure and testing them on the *C. elegans* model. This high-throughput strategy, taken altogether, resulted in identification of numerous potential therapeutic modalities, some of which are undergoing further testing (Li et al., 2014, 1

of 12). The first critical experiment was typical non-hypothesis driven, high-throughput research. Perlmutter's opinion was that we should get used to the new way of doing science.

The story of this research relating to α_1-antitrypsin deficiency illustrates many of the contrasts between hypothesis-driven and non-hypothesis-driven biological/biomedical research. The conventional hypothesis-driven research was elegant and highly successful, and it was very transparent in terms of its logical progression from one experiment to the next. The non-hypothesis-driven strategy required some conceptual flexibility.[235] One problem was finding a model in which to create the system under investigation, namely, the liver disease process itself in α_1-antitrypsin deficiency which involves disposition of abnormal α_1-antitrypsin protein. Clearly it was critically important to know as much about the biology of *C. elegans* as we do. The general thrust of the experimental design was not to ask how the model system functions[236] but how a large collection of diverse drugs operate in the system (hepatic α_1-antitrypsin deficiency thus modeled). This experimental approach was not governed by a proximate hypothesis. There was, as we might expect, a hunch hypothesis: some drug among the 12,000 tested probably has a beneficial effect. The results were quantitatively more productive than with the conventional "old way" of doing research.[237] However, I want to point out two aspects which were not mentioned, perhaps not really appreciated, by David Perlmutter. First of all, there was a superficial validation of the high-throughput findings by observing that one of the drugs was

[235] Critics might say that it demands conceptual flexibility in terms of my definition of system-driven research. It is certainly non-hypothesis-driven. The high-throughput methodology puts this line of research broadly into a system-driven design. However, it is specifically system-driven research. The system here is the disease process itself. It is modeled in *C. elegans* but *C. elegans* itself is not a model organism for the disorder. It would be difficult to style this research as data-driven, since available data about the drugs under consideration were not scrutinized: in fact, new data about those drugs had to be produced.

[236] The *C. elegans* engineered to have a vermiform version of hepatic α_1-antitrypsin deficiency qualifies as a Rheinberger system, and its development could be regarded as a kind of technique-driven research.

[237] Equally, it could be argued that since <0.05% of the drug collection turned out to be "of interest", the likelihood of this hunch hypothesis was very slender indeed. However, the high-throughput methodology identified effective drugs which would not necessarily have been identified by a hypothesis-driven approach. This is a typical example of 'drug discovery' (see Section 9.2.3).

carbamazepine, as identified previously on a hypothesis-driven mechanistic account. Animal models (whether mouse liver or worm gut) had been developed for taking those novel drugs identified and assessing how they fit in with the mechanics of the system. This does not mean that all of them necessarily enhanced autophagy: in fact, only some of these drugs enhanced autophagy and other mechanisms operative with other drugs remain to be discovered. It does provide an example of the importance of contextualizing in the system as a mode of validation for high-throughput data. Secondly, the "new way" of doing research shows the interdependence of hypothesis-driven and non-hypothesis-driven research. Finding a pattern in the high-throughput results generated a hypothesis about the structure of relevant drugs. In fact, the best way to investigate this hypothesis was to utilize another kind of high-throughput methodology. I agree with David Perlmutter that we have to get comfortable with the "new way" of doing research.[238] Omics and high-throughput methodology are here to stay. Moreover, we need to see that hypothesis-driven and non-hypothesis-driven research methodologies are complementary. This entails accepting non-hypothesis-driven research (and perhaps specifically any omics discipline) as a legitimate way of doing science, on any equal footing with the now conventional hypothesis-driven research.

10.3 Significance

The significance of my philosophical project is that it addresses a new type of biological research, which though less than 20 years-old is likely to dominate research methods in the future. Indeed the results of biological/biomedical research utilizing high-throughput methodologies, basically any of the omics disciplines, are truly spectacular. The OES challenges a time-honoured defining feature of scientific research,

[238] A fascinating side issue is that in the summary of this lecture prepared for a 'throw-away' daily newspaper for the conference ("*The Liver Meeting Today*"), there was no mention of the contrast between hypothesis-driven and non-hypothesis-driven research strategies. Indeed there was no mention of the non-hypothesis-driven research at all. Perhaps the science writer thought that it was too complicated for average attendee to understand; however, many of these attendees are highly sophisticated scientists. Alternatively, it may indicate the general resistance to "new ways" of doing biological and biomedical research.

namely to be hypothesis-driven, or at least it appears to do so. In displacing hypotheses from prime importance, what I have called '*system* biology' substantially changes the established practice of modern scientific research. Additionally, we can see that the epistemological considerations I have investigated in relation to omics, and specifically to proteomics, extend well beyond the scope of this project, into other categories of high-throughput research methodologies. Conservative voices from the research community dwell on the limitations, and even the epistemic perils, of omics as high-throughput methodology. A recent review of metalloproteomics emphasizes the analytical difficulties: problems of capturing and characterizing proteins which either incorporate metals in their structure or contribute to the disposition of metals and the unmanageable insouciance of omics in trying to address the entirety of a system (Barnett et al., 2012b, 3319). The authors of this thoughtful review paper emphasize the sources of error, uncertainty, and misinformation. Of course, being skeptical and circumspect is always a safeguard in scientific research. Some of those difficulties are long-recognized and come already equipped with technical coping strategies. As we have shown, other methodological weaknesses are being identified and potentially solved. To shy away from addressing a system directly and comprehensively simply because such an endeavour is fraught with difficulties seems counter to the spirit of scientific research. The OES involves epistemic virtues, as well as epistemic minefields. Proteomics, like other omics disciplines, is proving highly informative. It produces novel findings. That these omics disciplines are so productive has been what has gained their favour among scientists. The productivity has generated sufficient support to counterbalance any qualms about the characteristic lack of a hypothesis driving the experiment design.[239] Once it became obvious that omics was advancing scientific knowledge, it suddenly became less urgent, or even less attractive, to dismiss it as not really science simply because it did not feature a hypothesis, as with hypothesis-driven research. Practicality silenced the detractors. Developing the OES as a scientific

[239] This is based on a conversation with Dr. Roderick McInnes, on 17 December 2013 by phone. However, at the same time we begin to encounter public debate over emerging issues relating to high-throughput methodology: "unease over big science and small science approaches to biology" (Eddy, 2013, R261).

epistemology for system-driven research featuring high-throughput methodology may not be regarded as a practical necessity by biological/biomedical researchers; nevertheless, it can still be defended as a theoretical necessity, certainly at least to philosophers of biology and to philosophers of science more generally.

In fact, that the OES endows system-driven biological/biomedical research with scientific legitimacy has very practical implications for scientists and others assessing research proposals, whether in terms of their scientific value or their suitability for funding. A system-driven biological/biomedical experiment by definition will not feature a proximate hypothesis governing the experimental design. Thus it will not engage with the logic of the HDM. However, it is not free of normative expectations. It must engage with the logic of the OES. Thus the reviewer must be able to confirm that the system under investigation is described appropriately and meticulously, that the high-throughput methodology is suitable and that its technology adheres to the current standards of best performance, and finally that strategies for contextualization of the results have been identified in advance and are in fact the right kinds of contextualization strategies. If the researcher cannot produce a comprehensive description of experimental design to satisfy the key justificatory manoeuvres of the OES, then the proposal deserves to be rejected—but not because the system-driven experiment lacks a hypothesis. The OES can provide direction for evaluation of system-driven research and thus fulfill an important need of both the scientific community and granting agencies. Accordingly it removes the invidious need for the biological/biomedical researcher to make a system-driven experiment appear hypothesis-driven or denigrate excellent scientific research to the status of being only preliminary, that is, hypothesis-generating.

In conclusion, then, I have argued that omics as a high-throughput experimental methodology is a kind of system-driven research and, as such, is susceptible to interpretation and validation in an organized fashion. Further I have argued that the OES functions as a scientific epistemology for system-driven research, and I have investigated these issues in terms of proteomics. Finally, given my definition of science

as “an investigative activity which seeks to generate explanatory and predictive information about the natural world by means of methods which are both transparent and capable of standardization”, I have demonstrated that research carried out according to the OES constitutes actual scientific research, not a preliminary effort or run-up to real science. The HDM is time-honoured and highly effective for much biological/biomedical research. Nevertheless, for effective research about selected (but numerous) problems in contemporary biology and medicine, the OES is the best—possibly the only—scientific epistemology to employ.

BIBLIOGRAPHY

Aderem A. (2005). "Systems biology: its practice and challenges." *Cell* **121**(4): 511-513.

Aebersold R. and Mann M. (2003). "Mass spectrometry-based proteomics." *Nature* **422**(6928): 198-207.

Aggarwal K. and Lee K. H. (2003). "Functional genomics and proteomics as a foundation for systems biology." *Briefings in Functional Genomics and Proteomics* **2**(3): 175-184.

Ala A., Walker A. P., Ashkan K., Dooley J. S. and Schilsky M. L. (2007). "Wilson's disease." *Lancet* **369**(9559): 397-408.

Albala J. S. (2001). "Array-based proteomics: the latest chip challenge." *Expert Reviews of Molecular Diagnostics* **1**(2): 145-152.

Albert R., Jeong H. and Barabasi A. L. (1999). "Diameter of the world-wide web." *Nature* **401**(6749): 130-131.

Albert R., Jeong H. and Barabasi A. L. (2000). "Error and attack tolerance of complex networks." *Nature* **406**(6794): 378-382.

Alkhas A., Hood B. L., Oliver K., Teng P. N., Oliver J., Mitchell D., Hamilton C. A., Maxwell G. L. and Conrads T. P. (2011). "Standardization of a sample preparation and analytical workflow for proteomics of archival endometrial cancer tissue." *Journal of Proteome Research* **10**(11): 5264-5271.

Amundson R. and Lauder G. V. (1994). "Function without purpose: the uses of causal role function in evolutionary biology." *Biology & Philosophy* **9**(4): 443-469.

Andreini C., Banci L., Bertini I. and Rosato A. (2008). "Occurrence of copper proteins through the three domains of life: a bioinformatic approach." *Journal of Proteome Research* **7**(1): 209-216.

Andreini C., Bertini C. and Rosato A. (2009). "Metalloproteomes: a bioinformatic appraoch." *Accounts of Chemical Research* **42**(10): 1471-1479.

Anon. "Gordon Research Conferences." Retrieved 23 April 2014, from http://www.grc.org/search.aspx.

Anon. "Merriam-Webster dictonary." Retrieved 8 July 2014, from http://www.merriam-webster.com/dictionary/triage.

Anon. "OED online." Retrieved 4 July 2013, 21 October 2014, 6 January 2015, from http://www.oed.com.ezproxy.library.dal.ca/view/Entry/69014?rskey=QChIIi&result=2#eid.

Anon. "PubMed." Retrieved 11 November 2014, from http://www.ncbi.nlm.nih.gov/pubmed.

Anon. "Shimadzu analytical and measuring instruments." Retrieved 7 July 2014, from http://www.shimadzu.com/an/gcms/support/faq/mz.html.

Anon. "STRING database." Retrieved August 2011; 22 February 2014, from www.string.db.org.

Antony L. M. (1993). "Quine as feminist: the radical import of naturalized epistemology". *A Mind of One's Own. Feminist Essays on Reason and Objectivity*. L. M. Antony and C. Witt. Boulder, CO, Westview Press**:** 185-225.

Arguello J. M., Raimunda D. and Padilla-Benavides T. (2013). "Mechanisms of copper homeostasis in bacteria." *Frontiers in Cellular and Infection Microbiology* **3**: 73.

Aristotle (1984). *The Complete Works of Aristotle. The Revised Oxford TRanslation*. Princeton, NJ, Princeton University Press.

Aston N. S., Watt N., Morton I. E., Tanner M. S. and Evans G. S. (2000). "Copper toxicity affects proliferation and viability of human hepatoma cells (HepG2 line)." *Human Experimental Toxicology* **19**(6): 367-376.

Au C. E., Bell A. W., Gilchrist A., Hiding J., Nilsson T. and Bergeron J. J. (2007). "Organellar proteomics to create the cell map." *Current Opinion in Cell Biology* **19**(4): 376-385.

Audhya A. and Desai A. (2008). "Proteomics in Caenorhabditis elegans." *Briefings in Functional Genomics and Proteomics* **7**(3): 205-210.

Auffray C. and Noble D. (2009). "Origins of systems biology in William Harvey's masterpiece on the movement of the heart and the blood in animals." *International Journal of Molecular Science* **10**(4): 1658-1669.

Avery O. T., Macleod C. M. and McCarty M. (1944). "Studies on the chemical nature of the substance inducing transformation of Pneumococcal types : induction of transformation by a desoxyribonucleic acid fraction isolated from Pneumococcus type iii." *Journal of Experimental Medicine* **79**(2): 137-158.

Bacon F. (2000). *Francis Bacon: The New Organon. Cambridge Texts in the History of Philosophy*, Cambridge University Press.

Bamshad M. J., Ng S. B., Bigham A. W., Tabor H. K., Emond M. J., Nickerson D. A. and Shendure J. (2011). "Exome sequencing as a tool for Mendelian disease gene discovery." *Nature Reviews Genetics* **12**(11): 745-755.

Barabasi A. L. (2009). "Scale-free networks: a decade and beyond." *Science* **325**(5939): 412-413.

Barabasi A. L. and Albert R. (1999). "Emergence of scaling in random networks." *Science* **286**(5439): 509-512.

Barabasi A. L. and Oltvai Z. N. (2004). "Network biology: understanding the cell's functional organization." *Nature Reviews Genetics* **5**(2): 101-113.

Barnett J. P., Scanlan D. J. and Blindauer C. A. (2012a). "Fractionation and identification of metalloproteins from a marine cyanobacterium." *Analytic and Bioanalytic Chemistry* **402**(10): 3371-3377.

Barnett J. P., Scanlan D. J. and Blindauer C. A. (2012b). "Protein fractionation and detection for metalloproteomics: challenges and approaches." *Analytic and Bioanalytic Chemistry* **402**(10): 3311-3322.

Barsnes H. and Martens L. (2013). "Crowdsourcing in proteomics: public resources lead to better experiments." *Amino Acids* **44**(4): 1129-1137.

Bechtel W. (2011). "Mechanism and biological explanation." *Philosophy of Science* **78**(4): 533-557.

Bechtel W. and Richardson R. C. (2010). *Discovering Complexity. Decomposition and localization as strategies in scientific research*. Cambridge, MA, MIT Press.

Bell A. W., Deutsch E. W., Au C. E., Kearney R. E., Beavis R., Sechi S., Nilsson T. and Bergeron J. J. (2009). "A HUPO test sample study reveals common problems in mass spectrometry-based proteomics." *Nature Methods* **6**(6): 423-430.

Bell A. W., Nilsson T., Kearney R. E. and Bergeron J. J. (2007). "The protein microscope: incorporating mass spectrometry into cell biology." *Nature Methods* **4**(10): 783-784.

Bell G., Ed. (2009). *Foreword*. The Fourth Paradigm. Data-intensive scientific discovery. Redmond, WA, Microsoft Research.

Bensimon A., Heck A. J. and Aebersold R. (2012). "Mass spectrometry-based proteomics and network biology." *Annual Review of Biochemistry* **81**: 379-405.

Binz P. A., Barkovich R., Beavis R. C., Creasy D., Horn D. M., Julian R. K., Jr., Seymour S. L., Taylor C. F. and Vandenbrouck Y. (2008). "Guidelines for reporting the use of mass spectrometry informatics in proteomics." *Nature Biotechnology* **26**(8): 862.

Biron D. G., Brun C., Lefevre T., Lebarbenchon C., Loxdale H. D., Chevenet F., Brizard J. P. and Thomas F. (2006). "The pitfalls of proteomics experiments without the correct use of bioinformatics tools." *Proteomics* **6**(20): 5577-5596.

Bittner M. L. and Dougherty E. R. (2012). "Newton, Laplace, and the epistemology of systems biology." *Cancer Informatics* **11**: 185-190.

Borjigin J., Payne A. S., Deng J., Li X., Wang M. M., Ovodenko B., Gitlin J. D. and Snyder S. H. (1999). "A novel pineal night-specific ATPase encoded by the Wilson disease gene." *Journal of Neuroscience* **19**(3): 1018-1026.

Brading K. (2010). "Autonomous patterns and scientific realism." *Philosophy of Science* **77**(5): 827-839.

Brazma A. (2009). "Minimum Information About a Microarray Experiment (MIAME)--successes, failures, challenges." *Scientific World Journal* **9**: 420-423.

Brewis I. A. and Brennan P. (2010). "Proteomics technologies for the global identification and quantification of proteins." *Advances in Protein Chemistry and Structural Biology* **80**: 1-44.

Brogaard B. (2006). "Can virtue reliabilism explain the value of knowledge?" *Canadian Journal of Philosophy* **36**(3): 335-354.

Bruggeman F. J. and Westerhoff H. V. (2007). "The nature of systems biology." *Trends in Microbiology* **15**(1): 45-50.

Buiakova O. I., Xu J., Lutsenko S., Zeitlin S., Das K., Das S., Ross B. M., Mekios C., Scheinberg I. H. and Gilliam T. C. (1999). "Null mutation of the murine ATP7B (Wilson disease) gene results in intracellular copper accumulation and late-onset hepatic nodular transformation." *Human Molecular Genetics* **8**(9): 1665-1671.

Bull P. C. and Cox D. W. (1994). "Wilson disease and Menkes disease: new handles on heavy-metal transport." *Trends in Genetics* **10**(7): 246-252.

Bull P. C., Thomas G. R., Rommens J. M., Forbes J. R. and Cox D. W. (1993). "The Wilson disease gene is a putative copper transporting P-type ATPase similar to the Menkes gene." *Nature Genetics* **5**(4): 327-337.

Burian R. M. (1997a). "Comments on complexity and experimentation in biology." *Philosophy of Science* **64**(4): S279-S291.

Burian R. M. (1997b). "Exploratory experimentation and the role of histochemical techniques in the work of Jean Brachet, 1938-1952." *History and Philosophy of the Life Sciences* **19**(1): 27-45.

Burian R. M. (2007). "On microRNA and the need for exploratory experimentation in post-genomic molecular biology." *History and Philosophy of Life Sciences* **29**(3): 285-311.

Callebaut W. (2012). "Scientific perspectivism: a philosopher of science's response to the challenge of big data biology." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 69–80.

Campbell R. (1998). *Illusions of Paradox. A Feminist Epistemology Naturalized.* New York, Rowman and Littlefield.

Carmona A., Roudeau S., Perrin L., Veronesi G. and Ortega R. (2014). "Environmental manganese compounds accumulate as Mn(ii) within the Golgi apparatus of dopamine cells: relationship between speciation, subcellular distribution, and cytotoxicity." *Metallomics* **6**(4): 822-832.

Carr S., Aebersold R., Baldwin M., Burlingame A., Clauser K., Nesvizhskii A., Working Group on Publication Guidelines for P. and Protein Identification D. (2004). "The need for guidelines in publication of peptide and protein identification data: Working Group on Publication Guidelines for Peptide and Protein Identification Data." *Molecular and Cellular Proteomics* **3**(6): 531-533.

Cassman M. (2005). "Barriers to progress in systems biology." *Nature* **438**(7071): 1079.

Coffey A. J., Durkie M., Hague S., McLay K., Emmerson J., Lo C., Klaffke S., Joyce C. J., Dhawan A., Hadzic N., Mieli-Vergani G., Kirk R., Elizabeth Allen K., Nicholl D., Wong S., Griffiths W., Smithson S., Giffin N., Taha A., Connolly S., Gillett G. T., Tanner S., Bonham J., Sharrack B., Palotie A., Rattray M., Dalton A. and Bandmann O. (2013). "A genetic study of Wilson's disease in the United Kingdom." *Brain* **136**(Pt 5): 1476-1487.

Coles P. (2001). "Einstein, Eddington and the 1919 eclipse." *Astronomical Society of the Pacific Conference Series* **252**: 21-42.

Collura V. and Boissy G. (2007). "From protein-protein complexes to interactomics". *Subcellular Proteomics* E. Bertrand and M. Faupel. **43:** 135-183.

Coronado V., Nanji M. and Cox D. W. (2001). "The Jackson toxic milk mouse as a model for copper loading." *Mammalian Genome* **12**(10): 793-795.

Cox J. and Mann M. (2011). "Quantitative, high-resolution proteomics for data-driven systems biology." *Annual Review of Biochemistry* **80**: 273-299.

Cvetkovic A., Menon A. L., Thorgersen M. P., Scott J. W., Poole F. L., 2nd, Jenney F. E., Jr., Lancaster W. A., Praissman J. L., Shanmukh S., Vaccaro B. J., Trauger S. A., Kalisiak E., Apon J. V., Siuzdak G., Yannone S. M., Tainer J. A. and Adams M. W. (2010). "Microbial metalloproteomes are largely uncharacterized." *Nature* **466**(7307): 779-782.

da Silva M. A., Sussulini A. and Arruda M. A. (2010). "Metalloproteomics as an interdisciplinary area involving proteins and metals." *Expert Revews in Proteomics* **7**(3): 387-400.

Dagenais S. L., Guevara-Fujita M., Loechel R., Burgess A. C., Miller D. E., Yuzbasiyan-Gurkan V., Brewer G. J. and Glover T. W. (1999). "The canine copper toxicosis locus is not syntenic with ATP7B or ATX1 and maps to a region showing homology to human 2p21." *Mammalian Genome* **10**(7): 753-756.

de Bie P., Muller P., Wijmenga C. and Klomp L. W. (2007). "Molecular pathogenesis of Wilson and Menkes disease: correlation of mutations with molecular defects and disease phenotypes." *Journal of Medical Genetics* **44**(11): 673-688.

de Godoy L. M., Olsen J. V., Cox J., Nielsen M. L., Hubner N. C., Frohlich F., Walther T. C. and Mann M. (2008). "Comprehensive mass-spectrometry-based proteome quantification of haploid versus diploid yeast." *Nature* **455**(7217): 1251-1254.

Dennett D. C. (1991). "Real patterns." *The Journal of Philosophy* **88**(1): 27-51.

Dmitriev O., Tsivkovskii R., Abildgaard F., Morgan C. T., Markley J. L. and Lutsenko S. (2006). "Solution structure of the N-domain of Wilson disease protein: distinct nucleotide-binding environment and effects of disease mutations." *Proceedings of the National Academy of Sciences (USA)* **103**(14): 5302-5307.

Domann P. J., Akashi S., Barbas C., Huang L., Lau W., Legido-Quigley C., McClean S., Neususs C., Perrett D., Quaglia M., Rapp E., Smallshaw L., Smith N. W., Smyth W. F., Taylor C. F. and Minimum Information About a Proteomics E. (2010). "Guidelines for reporting the use of capillary electrophoresis in proteomics." *Nature Biotechnology* **28**(7): 654-655.

Domon B. and Aebersold R. (2006). "Challenges and opportunities in proteomics data analysis." *Molecular and Cellular Proteomics* **5**(10): 1921-1926.

Domon B. and Aebersold R. (2010). "Options and considerations when selecting a quantitative proteomics strategy." *Nature Biotechnology* **28**(7): 710-721.

Dougherty E. R. and Shmulevich I. (2012). "On the limitations of biological knowledge." *Current Genomics* **13**(7): 574-587.

Ducheyne S. (2010). "Whewell's tidal researches: scientific practice and philosophical methodology." *Studies in History and Philosophy of Science* **41**(1): 26-40.

Duncan M. W. and Hunsucker S. W. (2006). "Comments on standards in proteomics and the concept of fitness-for-purpose." *Proteomics* **6 Suppl 2**: 45-47.

Eddy S. R. (2013). "The ENCODE project: missteps overshadowing a success." *Current Biology* **23**(7): R259-261.

Fatemi N., Korzhnev D. M., Velyvis A., Sarkar B. and Forman-Kay J. D. (2010). "NMR characterization of copper-binding domains 4-6 of ATP7B." *Biochemistry* **49**(39): 8468-8477.

Fatemi N. and Sarkar B. (2002). "Structural and functional insights of Wilson disease copper-transporting ATPase." *Journal of Bioenergetics and Biomembranes* **34**(5): 339-349.

Feyerabend K. (1918). *A Pocket-dictionary of the Greek and English Languages*. Philadelphia, PA, David McKay, Publisher.

Feyerabend P. (2010). *Against Method (Fourth Edition)*. London, Verso.

Franklin A. (2012). "Experiment in physics." The Stanford Encyclopedia of Philosophy Winter 2012. Retrieved 20 September 2014, from <http://plato.stanford.edu/archives/win2012/entries/physics-experiment/>

Franklin A. (2013). *Shifting standards : experiments in particle physics in the twentieth century*. Pittsburgh, PA, University of Pittsburgh Press.

Franklin L. R. (2005). "Exploratory experiments." *Philosophy of Science* **72**(5): 888-899.

Fuchs S. A., Harakalova M., van Haaften G., van Hasselt P. M., Cuppen E. and Houwen R. H. (2012). "Application of exome sequencing in the search for genetic causes of rare disorders of copper metabolism." *Metallomics* **4**(7): 606-613.

Gahoi N., Ray S. and Srivastava S. (2015). "Array-based proteomic approaches to study signal transduction pathways: Prospects, merits and challenges." *Proteomics* **15**(2-3): 218-231.

Gao Y., Peng X., Zhang J., Zhao J., Li Y., Li Y., Li B., Hu Y. and Chai Z. (2013). "Cellular response of E. coli upon Hg2+ exposure--a case study of advanced nuclear analytical approach to metalloproteomics." *Metallomics* **5**(7): 913-919.

García-Sancho M. (2012). "From the genetic to the computer program: the historicity of 'data' and 'computation' in the investigations on the nematode worm C. elegans (1963–1998)." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 16-28.

Geillinger K. E., Kuhlmann K., Eisenacher M., Meyer H. E., Daniel H. and Spanier B. (2012). "Dynamic changes of the Caenorhabditis elegans proteome during ontogenesis assessed by quantitative analysis with 15N metabolic labeling." *Journal of Proteome Research* **11**(9): 4594-4604.

Ghaemmaghami S., Huh W. K., Bower K., Howson R. W., Belle A., Dephoure N., O'Shea E. K. and Weissman J. S. (2003). "Global analysis of protein expression in yeast." *Nature* **425**(6959): 737-741.

Gibbons M. G. (2012). "Reassessing discovery: Rosalind Franklin, scientific visualization, and the structure of DNA." *Philosophy of Science* **79**(1): 63-80.

Gibson F., Anderson L., Babnigg G., Baker M., Berth M., Binz P. A., Borthwick A., Cash P., Day B. W., Friedman D. B., Garland D., Gutstein H. B., Hoogland C., Jones N. A., Khan A., Klose J., Lamond A. I., Lemkin P. F., Lilley K. S., Minden J., Morris N. J., Paton N. W., Pisano M. R., Prime J. E., Rabilloud T., Stead D. A., Taylor C. F., Voshol H., Wipat A. and Jones A. R. (2008). "Guidelines for reporting the use of gel electrophoresis in proteomics." *Nature Biotechnology* **26**(8): 863-864.

Giere R. N., Bickle J. and Maudlin R. F. (2006). *Understanding Scientific Reasoning*. Toronto, Thomson Wadsworth.

Gilbert M. (1989). "Rationality and salience." *Philosophical Studies* **57**(1): 61-77.

Gilchrist A., Au C. E., Hiding J., Bell A. W., Fernandez-Rodriguez J., Lesimple S., Nagaya H., Roy L., Gosline S. J., Hallett M., Paiement J., Kearney R. E., Nilsson T. and Bergeron J. J.

(2006). "Quantitative proteomics analysis of the secretory pathway." *Cell* **127**(6): 1265-1281.

Giometti C. S. (2003). "Proteomics and bioinformatics." *Advances in Protein Chemistry* **65**: 353-369.

Gitlin J. D. (2003). "Wilson disease." *Gastroenterology* **125**(6): 1868-1877.

Goldman A. (2012). "Reliabilism and Contemporary Epistemology". Oxford, UK, Oxford University Press (Oxford Scholarship Online).

Goldman A. I. (1979). "What is Justified Belief?". *Justification and Knowledge. New Studies in Epistemology*. G. S. Pappas. Dordecht, The Netherlands, D. Riedel**:** 1-23.

Goldman A. I. (1999). "A priori warrant and naturalistic epistemology." *Nous* **33**(s13): 1-28.

Gomes C., Freire M. S. and Franco O. L. (2011). "Trace element analysis of proteins directly from 2D-PAGE: an efficient strategy for metalloproteomics." *Preparative Biochemistry and Biotechnology* **41**(3): 236-242.

Gosai S. J., Kwak J. H., Luke C. J., Long O. S., King D. E., Kovatch K. J., Johnston P. A., Shun T. Y., Lazo J. S., Perlmutter D. H., Silverman G. A. and Pak S. C. (2010). "Automated high-content live animal drug screening using C. elegans expressing the aggregation prone serpin alpha1-antitrypsin Z." *PLoS One* **5**(11): e15460.

Hacking I. (1983). *Representing and Intervening. Introductory topics in the philosophy of natural science*. Cambridge, UK, Cambridge University Press.

Hamza I., Faisst A., Prohaska J., Chen J., Gruss P. and Gitlin J. D. (2001). "The metallochaperone Atox1 plays a critical role in perinatal copper homeostasis." *Proceedings of the National Academy of Sciences (USA)* **98**(12): 6848-6852.

Hanson N. R. (1961). *Patterns of Discovery*. London, Cambridge University Press.

Harper D. "Online Etymology Dictionary." Retrieved 8 May 2013, from http://www.etymonline.com

Haywood S., Fuentealba I. C., Foster J. and Ross G. (1996). "Pathobiology of copper-induced injury in Bedlington terriers: ultrastructural and microanalytical studies." *Analytical Cellular Pathology* **10**(3): 229-241.

Heikes D. K. (2004). "The bias paradox: why it's not just for feminists anymore." *Synthese* **138**(3): 315-335.

Hempel C. (1966). *Philosophy of Natural Science*. Englewood Cliffs, NJ, Prentice-Hall.

Hidvegi T., Ewing M., Hale P., Dippold C., Beckett C., Kemp C., Maurice N., Mukherjee A., Goldbach C., Watkins S., Michalopoulos G. and Perlmutter D. H. (2010). "An autophagy-enhancing drug promotes degradation of mutant alpha1-antitrypsin Z and reduces hepatic fibrosis." *Science* **329**(5988): 229-232.

Hirotsu S., Abe Y., Okada K., Nagahara N., Hori H., Nishino T. and Hakoshima T. (1999). "Crystal structure of a multifunctional 2-Cys peroxiredoxin heme-binding protein 23 kDa/proliferation-associated gene product." *Proceedings of the National Academy of Sciences (USA)* **96**(22): 12333-12338.

Hodgkin D. C. (1965). "The X-ray analysis of complicated molecules." *Science* **150**(699): 979-988.

Hoogland C., O'Gorman M., Bogard P., Gibson F., Berth M., Cockell S. J., Ekefjard A., Forsstrom-Olsson O., Kapferer A., Nilsson M., Martinez-Bartolome S., Albar J. P., Echevarria-Zomeno S., Martinez-Gomariz M., Joets J., Binz P. A., Taylor C. F., Dowsey A., Jones A. R. and Minimum Information About a Proteomics E. (2010). "Guidelines for reporting the use of gel image informatics in proteomics." *Nature Biotechnology* **28**(7): 655-656.

Horton M. (1973). "In defence of Francis Bacon. A citicism of the critics of the inductive method." *Studies in the History and Philosophy of Science* **4**(3): 241-278.

Hudson R. G. (2003). "Novelty and the 1919 eclipse experiments." *Studes in the History and Philosophy of Modern Physics* **34**(xx): 107-129.

Hume D. (1978). *A Treatise of Human Nature*. Oxford, Clarendon Press.

Hung I. H., Casareno R. L., Labesse G., Mathews F. S. and Gitlin J. D. (1998). "HAH1 is a copper-binding protein with distinct amino acid residues mediating copper homeostasis and antioxidant defense." *Journal of Biological Chemistry* **273**(3): 1749-1754.

Huster D., Finegold M. J., Morgan C. T., Burkhead J. L., Nixon R., Vanderwerf S. M., Gilliam C. T. and Lutsenko S. (2006). "Consequences of copper accumulation in the livers of the Atp7b-/- (Wilson disease gene) knockout mice." *American Journal of Pathology* **168**(2): 423-434.

Huster D., Purnat T. D., Burkhead J. L., Ralle M., Fiehn O., Stuckert F., Olson N. E., Teupser D. and Lutsenko S. (2007). "High copper selectively alters lipid metabolism and cell cycle machinery in the mouse model of Wilson disease." *Journal of Biological Chemistry* **282**(11): 8343-8355.

Ideker T., Galitski T. and Hood L. (2001). "A new approach to decoding life: systems biology." *Annu Rev Genom Hum Genet* **2**: 343-372.

Inesi G., Pilankatta R. and Tadini-Buoninsegni F. (2014). "Biochemical characterization of P-type copper ATPases." *Biochemical Journal* **463**(2): 167-176.

Jakubowski N., Lobinski R. and Moens L. (2004). "Metallobiomolecules. The basis of life, the challenge of atomic spectroscopy." *Journal of Analytical Atomic Spectrometry* **19**(1): 1-4.

Jeong H., Tombor B., Albert R., Oltvai Z. N. and Barabasi A. L. (2000). "The large-scale organization of metabolic networks." *Nature* **407**(6804): 651-654.

Jones A. R., Carroll K., Knight D., Maclellan K., Domann P. J., Legido-Quigley C., Huang L., Smallshaw L., Mirzaei H., Shofstahl J., Paton N. W. and Minimum Information About a Proteomics E. (2010). "Guidelines for reporting the use of column chromatography in proteomics." *Nature Biotechnology* **28**(7): 654.

Jones P., Cote R. G., Martens L., Quinn A. F., Taylor C. F., Derache W., Hermjakob H. and Apweiler R. (2006). "PRIDE: a public repository of protein and peptide identifications for the proteomics community." *Nucleic Acids Research* **34**(Database issue): D659-663.

Kahn P. (1995). "From genome to proteome: looking at a cell's proteins." *Science* **270**(5235): 369-370.

Kandel E. R. (2006). *In Search of Memory. The Emergence of a New Science of Mind* New York, W. W. Norton.

Karr J. R., Sanghvi J. C., Macklin D. N., Gutschow M. V., Jacobs J. M., Bolival B., Jr., Assad-Garcia N., Glass J. I. and Covert M. W. (2012). "A whole-cell computational model predicts phenotype from genotype." *Cell* **150**(2): 389-401.

Keating P. and Cambrosio A. (2012). "Too many numbers: microarrays in clinical cancer research." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 37-51.

Keller E. F. (2005a). "Ecosystem, organisms, and machines." *BioScience* **55**(12): 1069-1074.

Keller E. F. (2005b). "Revisiting "scale-free" networks." *Bioessays* **27**(10): 1060-1068.

Keller E. F. (2012). "Genes, genomes, and genomics." *Biological Theory* **6**(2): 132-140.

Kenyani J., Medina-Aunon J. A., Martinez-Bartolome S., Albar J. P., Wastling J. M. and Jones A. R. (2011). "A DIGE study on the effects of salbutamol on the rat muscle proteome - an exemplar of best practice for data sharing in proteomics." *BMC Research Notes* **4**: 86.

Ko J. H., Son W., Bae G. Y., Kang J. H., Oh W. and Yoo O. J. (2006). "A new hepatocytic isoform of PLZF lacking the BTB domain interacts with ATP7B, the Wilson disease protein, and positively regulates ERK signal transduction." *Journal of Cellular Biochemistry* **99**(3): 719-734.

Kocher T. and Superti-Furga G. (2007). "Mass spectrometry-based functional proteomics: from molecular machines to protein networks." *Nature Methods* **4**(10): 807-815.

Kohl P., Crampin E. J., Quinn T. A. and Noble D. (2010). "Systems biology: an approach." *Clinical Pharmacology and Therapeutics* **88**(1): 25-33.

Kopelman L. M. (1994). "Case method and casuistry: the problem of bias." *Theoretical Medicine* **15**(1): 21-37.

Krohs U., Ed. (2010). *Epistemic consequences of two different strategies for decomposing biological networks*. EPSA Philosophical Issues in the Sciences: Launch of the European Philosophy of Science Association. [Netherlands], Springer Science+Business.

Krohs U. (2012). "Convenience experimentation." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 52-57.

Kuhn T. S. (1977). "Objectivity, value judgment, and theory choice". *The Essential Tension. Selected Studies in Scientific Tradiion and Change*. T. S. Kuhn. Chicago, University of Chicago Press**:** 320-339.

Kulkarni P. P., She Y. M., Smith S. D., Roberts E. A. and Sarkar B. (2006). "Proteomics of metal transport and metal-associated diseases." *Chemistry A European Journal* **12**(9): 2410-2422.

Kung L. A. and Snyder M. (2006). "Proteome chips for whole-organism assays." *Nature Reviews Molecular Cell Biology* **7**(8): 617-622.

Lakatos I. (1974). "The role of crucial experiments in science." *Studies in History and Philosophy of Science* **4**(4): 309-325.

Lear J. (2010). *Aristotle: the desire to understand*. Cambrudge, UK, Cambridge University Press.

Lederberg J. and McCray A. T. (2001). "'Ome Sweet 'Omics--A genealogical treasury of words." *The Scientist* **15**(1): 8.

Lee B. H., Kim J. M., Heo S. H., Mun J. H., Kim J., Kim J. H., Jin H. Y., Kim G. H., Choi J. H. and Yoo H. W. (2011). "Proteomic analysis of the hepatic tissue of Long-Evans Cinnamon (LEC) rats according to the natural course of Wilson disease." *Proteomics* **11**(18): 3698-3705.

Legrain P., Aebersold R., Archakov A., Bairoch A., Bala K., Beretta L., Bergeron J., Borchers C. H., Corthals G. L., Costello C. E., Deutsch E. W., Domon B., Hancock W., He F., Hochstrasser D., Marko-Varga G., Salekdeh G. H., Sechi S., Snyder M., Srivastava S., Uhlen M., Wu C. H., Yamamoto T., Paik Y. K. and Omenn G. S. (2011). "The human

proteome project: current state and future direction." *Molecular and Cellular Proteomics* **10**(7): M111 009993.

Leonelli S. (2008). "Performing abstraction: two ways of modelling *Arabidopsis thaliana*." *Biology and Philosophy* **23**(4): 509-528.

Leonelli S. (2012). "Introduction: Making sense of data-driven research in the biological and biomedical sciences." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 1-3.

Leonelli S. and Ankeny R. A. (2012). "Re-thinking organisms: The impact of databases on model organism biology." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 29-35.

Lewontin R. (2000). *The Triple Helix. Gene, Organism, and Environment*. Cambridge, MA, Harvard university Press.

Li J., Pak S. C., O'Reilly L. P., Benson J. A., Wang Y., Hidvegi T., Hale P., Dippold C., Ewing M., Silverman G. A. and Perlmutter D. H. (2014). "Fluphenazine reduces proteotoxicity in C. elegans and mammalian models of alpha-1-antitrypsin deficiency." *PLoS One* **9**(1): e87260.

Lievens S., Eyckerman S., Lemmens I. and Tavernier J. (2010). "Large-scale protein interactome mapping: strategies and opportunities." *Expert Reviews in Proteomics* **7**(5): 679-690.

Lim C. M., Cater M. A., Mercer J. F. and La Fontaine S. (2006a). "Copper-dependent interaction of dynactin subunit p62 with the N terminus of ATP7B but not ATP7A." *Journal of Biological Chemistry* **281**(20): 14006-14014.

Lim C. M., Cater M. A., Mercer J. F. and La Fontaine S. (2006b). "Copper-dependent interaction of glutaredoxin with the N termini of the copper-ATPases (ATP7A and ATP7B) defective in Menkes and Wilson diseases." *Biochemical and Biophysical Research Communications* **348**(2): 428-436.

Lin J. and Qian J. (2007). "Systems biology approach to integrative comparative genomics." *Expert Reviews in Proteomics* **4**(1): 107-119.

Lobinski R., Moulin C. and Ortega R. (2006). "Imaging and speciation of trace elements in biological environment." *Biochimie* **88**(11): 1591-1604.

Longino H. E. (1994). "The fate of knowledge in social theories of science". *Socializing Epistemology: the Social Dimensions of Knowledge*. F. F. Schnitt. Lanham, MD, Rowman & Littlefield**:** 135-157.

Longino H. E. (1997). "Feminist epistemology as a local epistemology." *Proceedings of the Aristotelian Society, Supplementary Volumes* **71**(Suppl): 19-35.

Longino H. E. (2002). *The Fate of Knowledge*. Princeton, NJ, Princeton University Press.

Loscalzo J. and Barabasi A.-L. (2011). "Systems biology and the future of medicine." *Wiley Interdisciplinary Reviews in Systems Biology and Medicine* **3**(6): 619-627.

Maine G. N., Mao X., Muller P. A., Komarck C. M., Klomp L. W. and Burstein E. (2009). "COMMD1 expression is controlled by critical residues that determine XIAP binding." *Biochemical Journal* **417**(2): 601-609.

Malmstrom L., Malmstrom J. and Aebersold R. (2011). "Quantitative proteomics of microbes: Principles and applications to virulence." *Proteomics* **11**(15): 2947-2956.

Malone J. H. and Oliver B. (2011). "Microarrays, deep sequencing and the true measure of the transcriptome." *BMC Biology* **9**: 34.

Martens L. (2013). "Resilience in the proteomics data ecosystem: how the field cares for its data." *Proteomics* **13**(10-11): 1548-1550.

Martens L., Orchard S., Apweiler R. and Hermjakob H. (2007). "Human Proteome Organization Proteomics Standards Initiative: data standardization, a view on developments and policy." *Molecular and Cellular Proteomics* **6**(9): 1666-1667.

Martinelli D., Travaglini L., Drouin C. A., Ceballos-Picot I., Rizza T., Bertini E., Carrozzo R., Petrini S., de Lonlay P., El Hachem M., Hubert L., Montpetit A., Torre G. and Dionisi-Vici C. (2013). "MEDNIK syndrome: a novel defect of copper metabolism treatable by zinc acetate therapy." *Brain* **136**(Pt 3): 872-881.

Martinez-Bartolome S., Deutsch E. W., Binz P. A., Jones A. R., Eisenacher M., Mayer G., Campos A., Canals F., Bech-Serra J. J., Carrascal M., Gay M., Paradela A., Navajas R., Marcilla M., Hernaez M. L., Gutierrez-Blazquez M. D., Velarde L. F., Aloria K., Beaskoetxea J., Medina-Aunon J. A. and Albar J. P. (2013). "Guidelines for reporting quantitative mass spectrometry based experiments in proteomics." *Journal of Proteomics* **95**: 84-88.

Materia S., Cater M. A., Klomp L. W., Mercer J. F. and La Fontaine S. (2011). "Clusterin (apolipoprotein J), a molecular chaperone that facilitates degradation of the copper-ATPases ATP7A and ATP7B." *Journal of Biological Chemistry* **286**(12): 10073-10083.

Mayo D. (1996). *Error and the Growth of Experimental Knowledge* Chicago, University of Chicago Press.

Mayo D. G. (2000). "Experimental practice and an error statistical account of evidence." *Philosophy of Science* **67**(3): S193-S207.

McAllister J. W. (1997). "Phenomena and patterns in data sets." *Erkenntnis* **47**(2): 217-228.

McAllister J. W. (2003). "Algorithmic randomness in empirical data." *Studies in History and Philosophy of Science* **34**(3): 633-646.

McAllister J. W. (2005). "Algorithmic compression of empirical data: reply to Twardy, Gardner, and Dowe." *Studies in History and Philosophy of Science* **36**(2): 403-410.

McAllister J. W. (2007). "Model selection and the multiplicity of patterns in empirical data." *Philosophy of Science* **74**(5): 884-894.

McAllister J. W. (2010). "The ontology of patterns in empirical data." *Philosophy of Science* **77**(5): 804-814.

McKusick V. A. and Ruddle F. H. (1987). "A new discipline, a new name, a new journal." *Genomics* **1**(1-2): 1-2.

McOuat G. (2009). "The origins of 'natural kinds': keeping 'essentialism' at bay in the Age of Reform." *Intellectual History Review* **19**(2): 211-230.

Medawar P. (1982). *Pluto's Republic*. Oxford, UK, Oxford University Press.

Merle U., Weiss K. H., Eisenbach C., Tuma S., Ferenci P. and Stremmel W. (2010). "Truncating mutations in the Wilson disease gene ATP7B are associated with very low serum ceruloplasmin oxidase activity and an early onset of Wilson disease." *BMC Gastroenterology* **Jan 18;10:8**.

Meynell L. (2013). "Parsing picture: on analzing the content of images in science." *The Knowledge Engineering Review* **28**(3): 327-345.

Mischler B. D. and Brandon R. N. (1987). "Individuality, pluralism, and the Phylogenetic Species Concept." *Biology and Phiosophy* **2**(4): 397-414.

Montecchi-Palazzi L., Kerrien S., Reisinger F., Aranda B., Jones A. R., Martens L. and Hermjakob H. (2009). "The PSI semantic validator: a framework to check MIAPE compliance of proteomics data." *Proteomics* **9**(22): 5112-5119.

Mufti A. R., Burstein E., Csomos R. A., Graf P. C., Wilkinson J. C., Dick R. D., Challa M., Son J. K., Bratton S. B., Su G. L., Brewer G. J., Jakob U. and Duckett C. S. (2006). "XIAP Is a copper binding protein deregulated in Wilson's disease and other copper toxicosis disorders." *Molecular Cell* **21**(6): 775-785.

Mufti A. R., Burstein E. and Duckett C. S. (2007). "XIAP: cell death regulation meets copper homeostasis." *Archives of Biochemistry and Biophysics* **463**(2): 168-174.

Muller P., van Bakel H., van de Sluis B., Holstege F., Wijmenga C. and Klomp L. W. (2007). "Gene expression profiling of liver cells after copper overload in vivo and in vitro reveals new copper-regulated genes." *Journal of Biological Inorganic Chemistry* **12**(4): 495-507.

Nabel G. J. (2009). "Philosophy of science. The coordinates of truth." *Science* **326**(5949): 53-54.

Nagel T. (2012). *Mind and Cosmos*. Oxford, Oxford Univiersity Press.

Narindrasorasak S., Yao P. and Sarkar B. (2003). "Protein disulfide isomerase, a multifunctional protein chaperone, shows copper-binding activity." *Biochemical and Biophysical Research Communications* **311**(2): 405-414.

Narindrasorasak S., Zhang X.-F., Roberts E. and Sarkar B. (2004). "Comparative analysis of metal binding characteristics of copper chaperone proteins, Atx1 and Atox1." *Bioinorganic Chemistry and Applications* **2**(1-2): 105-123.

Nelkin N. (1994). "Patterns." *Mind & Language* **9**(1): 56-87.

Nykamp D. Q. "Scale-free networks." Retrieved 2 November 2014, from http://mathinsight.org/scale_free_network.

O'Malley M. A. (2007). "Exploratory experimentation and scientific practice: metagenomics and the proteorhodopsin case." *History and Philosophy of the Life Sciences* **29**(3): 335-358.

O'Malley M. A. and Dupre J. (2005). "Fundamental issues in systems biology." *BioEssays* **27**(12): 1270-1276.

O'Malley M. A., Elliott K. C., Haufe C. and Burian R. M. (2009). "Philosophies of funding." *Cell* **138**(4): 611-615.

O'Malley M. A. and Soyer O. S. (2012). "The roles of integration in molecular systems biology." *Studies in History and Philosophy of Biological and Biomedical Sciences* **43**(1): 58-68.

Okada T., Shiono Y., Kaneko Y., Miwa K., Hasatani K., Hayashi Y., Mibayashi H., Aoyagi H., Tsuji S., Yoshimitsu M., Hayashi H. and Yamagishi M. (2010). "High prevalence of fulminant hepatic failure among patients with mutant alleles for truncation of ATP7B in Wilson's disease." *Scandinavian Journal of Gastroenterology* **45**(10): 1232-1237.

Okhrulik K. (1994). "Gender and the biological sciences." *Canadian Journal of Philosophy* **Suppl 20**: 21-42.

Onions C. T., Ed. (1955). *The Oxford Universal Dictionary on Historical Principles*. London UK, Clarendon Press.

Onions C. T. (1966). *The Oxford Dictionary of English Etymology*. Oxford, UK, Clarendon Press.

Orchard S., Albar J. P., Deutsch E. W., Eisenacher M., Binz P. A., Martinez-Bartolome S., Vizcaino J. A. and Hermjakob H. (2012). "From proteomics data representation to public data flow: a report on the HUPO-PSI workshop September 2011, Geneva, Switzerland." *Proteomics* **12**(3): 351-355.

Orchard S., Salwinski L., Kerrien S., Montecchi-Palazzi L., Oesterheld M., Stumpflen V., Ceol A., Chatr-aryamontri A., Armstrong J., Woollard P., Salama J. J., Moore S., Wojcik J., Bader G. D., Vidal M., Cusick M. E., Gerstein M., Gavin A. C., Superti-Furga G., Greenblatt J., Bader J., Uetz P., Tyers M., Legrain P., Fields S., Mulder N., Gilson M., Niepmann M., Burgoon L., De Las Rivas J., Prieto C., Perreau V. M., Hogue C., Mewes H. W., Apweiler R., Xenarios I., Eisenberg D., Cesareni G. and Hermjakob H. (2007). "The minimum information required for reporting a molecular interaction experiment (MIMIx)." *Nature Biotechnology* **25**(8): 894-898.

Ortega R., Deves G. and Carmona A. (2009). "Bio-metals imaging and speciation in cells using proton and synchrotron radiation X-ray microspectroscopy." *Journal of the Royal Society Interface* **6 Suppl 5**: S649-658.

Park D., Lee S., Bolser D., Schroeder M., Lappe M., Oh D. and Bhak J. (2005). "Comparative interactomics analysis of protein family interaction networks using PSIMAP (protein structural interactome map)." *Bioinformatics* **21**(15): 3234-3240.

Paton R. (2002). "Process, structure and context in relation to integrative biology." *Biosystems* **64**(1-3): 63-72.

Pauling L. and Corey R. B. (1951). "Configurations of Polypeptide Chains With Favored Orientations Around Single Bonds: Two New Pleated Sheets." *Proceedings of the National Academy of Sciences (USA)* **37**(11): 729-740.

Perini L. (2005). "The truth in pictures." *Philosphy of Science* **72**(1): 262-285.

Permyakov E. A. (2009). *Metalloproteomics*. Hoboken, NJ, John Wiley and Sons.

Petrak J., Myslivcova D., Man P., Cmejla R., Cmejlova J. and Vyoral D. (2006). "Proteomic analysis of iron overload in human hepatoma cells." *American Journal of Physiology Gastrointestinal and Liver Physiology* **290**(5): G1059-1066.

Pilankatta R., Lewis D. and Inesi G. (2011). "Involvement of protein kinase D in expression and trafficking of ATP7B (copper ATPase)." *Journal of Biological Chemistry* **286**(9): 7389-7396.

Plato (1920). *The Dialogues of Plato*. New York, Random House.

Popper K. R. (1968). *The Logic of Scientific Discovery*. London, UK, Hutchinson.

Prakash A., Piening B., Whiteaker J., Zhang H., Shaffer S. A., Martin D., Hohmann L., Cooke K., Olson J. M., Hansen S., Flory M. R., Lee H., Watts J., Goodlett D. R., Aebersold R., Paulovich A. and Schwikowski B. (2007). "Assessing bias in experiment design for large scale mass spectrometry-based quantitative proteomics." *Molecular and Cellular Proteomics* **6**(10): 1741-1748.

Priest G. (2014). "Revising logic". *The Metaphysics of Logic*. P. Rush, Cambridge University Press**:** 211-223.

Prot J. M., Briffaut A. S., Letourneur F., Chafey P., Merlier F., Grandvalet Y., Legallais C. and Leclerc E. (2011). "Integrated proteomic and transcriptomic investigation of the acetaminophen toxicity in liver microfluidic biochip." *PLoS One* **6**(8): e21268.

Quackenbush J. (2004). "Data standards for 'omic' science." *Nature Biotechnology* **22**(5): 613-614.

Quackenbush J. (2006). "Top-down standards will not serve systems biology." *Nature* **440**(7080): 24.

Quackenbush J. and Irizarry R. A. (2006). "Response to Shields: 'MIAME, we have a problem'." *Trends in Genetics* **22**(9): 471-472.

Reichenbach H. (1938). "On probability and induction." *Philosophy of Science* **5**(1): 21.

Rheinberger H.-J. (1997). *Toward a History of Epistemic Things. Synthesizing Proteins in the Test Tube*. Stanford, CA, Stanford University Press.

Richardson R. C. and Stephan A. (2007). "Mechanism and mechanical explanation in systems biology". *Systems Biology. Philosophical Foundations*. F. C. Boogerd, F. J. Bruggeman, J.-H. S. Hofmeyr and H. V. Westerhoff. Oxford, UK, Elsevier**:** 123-144.

Roberts E. A. (2012). "Using metalloproteomics to investigate the cellular physiology of copper in hepatocytes." *Metallomics* **4**(7): 633-640.

Roberts E. A. and Cox D. W. (2012). "Wilson's disease". *Hepatology*. T. Boyer, M. Manns and T. Wright. Philadelphia, Elsevier**:** 1110-1126.

Roberts E. A., Lau C. H., da Silveira T. R. and Yang S. (2007). "Developmental expression of Commd1 in the liver of the Jackson toxic milk mouse." *Biochemical and Biophysical Research Communications* **363**(4): 921-925.

Roberts E. A., Robinson B. H. and Yang S. (2008). "Mitochondrial structure and function in the untreated Jackson toxic milk (tx-j) mouse, a model for Wilson disease." *Molecular Genetics and Metabolism* **93**(1): 54-65.

Roberts E. A. and Sarkar B. (2014). "Metalloproteomics: focus on metabolic issues relating to metals." *Current Opinion in Clinical Nutrition and Metabolic Care* **17**: 425-430.

Roberts E. A., Sarkar B., She Y.-M., Narindrasorasak S. and Yang S. (2002). "Metalloproteome of normal and copper-depleted Hep G2 cells (abstr.)." *Molecular and Cellular Proteomics* **1**: 662.

Roelofsen H., Balgobind R. and Vonk R. J. (2004). "Proteomic analyzes of copper metabolism in an in vitro model of Wilson disease using surface enhanced laser desorption/ionization-time of flight-mass spectrometry." *Journal of Cellular Biochemistry* **93**(4): 732-740.

Safaei R., Adams P. L., Maktabi M. H., Mathews R. A. and Howell S. B. (2012). "The CXXC motifs in the metal binding domains are required for ATP7B to mediate resistance to cisplatin." *Journal of Inorganic Biochemistry* **110**: 8-17.

Salmon W. C. (1978). "Why ask, 'Why?'? an inquiry concerning scientific explanation." *Proceedings and Addresses of the American Philosophical Association* **51**(6): 683-705.

Sarkar B. and Roberts E. A. (2011). "The puzzle posed by COMMD1, a newly discovered protein binding Cu(II)." *Metallomics* **3**(1): 20-27.

Saul J. M. (2003). *Feminism. Issues and Arguments*. Oxford, Oxford University Press.

Schena M., Shalon D., Davis R. W. and Brown P. O. (1995). "Quantitative monitoring of gene expression patterns with a complementary DNA microarray." *Science* **270**(5235): 467-470.

Seong S. Y. and Choi C. Y. (2003). "Current status of protein chip development in terms of fabrication and application." *Proteomics* **3**(11): 2176-2189.

Seth R., Yang S., Choi S., Sabean M. and Roberts E. A. (2004). "In vitro assessment of copper-induced toxicity in the human hepatoma line, Hep G2." *Toxicology In Vitro* **18**(4): 501-509.

Sevcenco A. M., Hagen W. R. and Hagedoorn P. L. (2012). "Microbial metalloproteomes explored using MIRAGE." *Chemistry and Biodiversity* **9**(9): 1967-1980.

She Y. M., Narindrasorasak S., Yang S., Spitale N., Roberts E. A. and Sarkar B. (2003). "Identification of metal-binding proteins in human hepatoma lines by immobilized metal affinity chromatography and mass spectrometry." *Molecular and Cellular Proteomics* **2**(12): 1306-1318.

Shi W. and Chance M. R. (2011). "Metalloproteomics: forward and reverse approaches in metalloprotein structural and functional characterization." *Current Opinion in Chemical Biology* **15**(1): 144-148.

Shi W., Punta M., Bohon J., Sauder J. M., D'Mello R., Sullivan M., Toomey J., Abel D., Lippi M., Passerini A., Frasconi P., Burley S. K., Rost B. and Chance M. R. (2011). "Characterization of metalloproteins by high-throughput X-ray absorption spectroscopy." *Genome Research* **21**(6): 898-907.

Shields R. (2006). "MIAME, we have a problem." *Trends in Genetics* **22**(2): 65-66.

Siegel H. (1980). "Justification, discovery and the naturalizing of epistemology." *Philosophy of Science* **47**(2): 297-321.

Simpson D. M., Beynon R. J., Robertson D. H., Loughran M. J. and Haywood S. (2004). "Copper-associated liver disease: a proteomics study of copper challenge in a sheep model." *Proteomics* **4**(2): 524-536.

Singleton W. C., McInnes K. T., Cater M. A., Winnall W. R., McKirdy R., Yu Y., Taylor P. E., Ke B. X., Richardson D. R., Mercer J. F. and La Fontaine S. (2010). "Role of glutaredoxin1 and glutathione in regulating the activity of the copper-transporting P-type ATPases, ATP7A and ATP7B." *Journal of Biological Chemistry* **285**(35): 27111-27121.

Slany A., Haudek V. J., Zwickl H., Gundacker N. C., Grusch M., Weiss T. S., Seir K., Rodgarkia-Dara C., Hellerbrand C. and Gerner C. (2010). "Cell characterization by proteome profiling applied to primary hepatocytes and hepatocyte cell lines Hep-G2 and Hep-3B." *Journal of Proteome Research* **9**(1): 6-21.

Smith S. D., She Y. M., Roberts E. A. and Sarkar B. (2004). "Using immobilized metal affinity chromatography, two-dimensional electrophoresis and mass spectrometry to identify hepatocellular proteins with copper-binding ability." *Journal of Proteome Research* **3**(4): 834-840.

Snyder L. J. (2011). *The Philosophical Breakfast Club*. New York, Broadway Paperbacks.

Sober E. (1990). "Let's razor Ockham's razor". *Explanation and its Limits*. D. Knowles. Cambridge, Cambridge University Press**:** 73-94.

Springer O., Ed. (1975). *Langenscheidt's Encyclopaedic Dictionary of the English and German Languages*. London, UK, Hodder and Toughton.

Steinle F. (1997). "Entering new fields: exploratory uses of experimentation." *Philosophy of Science* **64**(Suppl): S65-S74.

Sternlieb I. (1968). "Mitochondrial and fatty changes in hepatocytes of patients with Wilson's disease." *Gastroenterology* **55**(3): 354-367.

Sun X., Chiu J. F. and He Q. Y. (2005). "Application of immobilized metal affinity chromatography in proteomics." *Expert Reviews in Proteomics* **2**(5): 649-657.

Sun X., Xiao C. L., Ge R., Yin X., Li H., Li N., Yang X., Zhu Y., He X. and He Q. Y. (2011). "Putative copper- and zinc-binding motifs in Streptococcus pneumoniae identified by immobilized metal affinity chromatography and mass spectrometry." *Proteomics* **11**(16): 3288-3298.

Suppes P., Ed. (1962). *Models of data*. Logic, Methodology, and Philosophy of Science: Proceedings of the 1960 International Congress. Stanford, Stanford University Press.

Szpunar J. (2005). "Advances in analytical methodology for bioinorganic speciation analysis: metallomics, metalloproteomics and heteroatom-tagged proteomics and metabolomics." *Analyst* **130**(4): 442-465.

Tanzi R. E., Petrukhin K., Chernov I., Pellequer J. L., Wasco W., Ross B., Romano D. M., Parano E., Pavone L., Brzustowicz L. M. and et al. (1993). "The Wilson disease gene is a copper transporting ATPase with homology to the Menkes disease gene." *Nature Genetics* **5**(4): 344-350.

Tao T. Y., Liu F., Klomp L., Wijmenga C. and Gitlin J. D. (2003). "The copper toxicosis gene product Murr1 directly interacts with the Wilson disease protein." *Journal of Biological Chemistry* **278**(43): 41593-41596.

Taylor C. F., Binz P. A., Aebersold R., Affolter M., Barkovich R., Deutsch E. W., Horn D. M., Huhmer A., Kussmann M., Lilley K., Macht M., Mann M., Muller D., Neubert T. A., Nickson J., Patterson S. D., Raso R., Resing K., Seymour S. L., Tsugita A., Xenarios I., Zeng R. and Julian R. K., Jr. (2008). "Guidelines for reporting the use of mass spectrometry in proteomics." *Nature Biotechnology* **26**(8): 860-861.

Taylor C. F., Paton N. W., Lilley K. S., Binz P. A., Julian R. K., Jr., Jones A. R., Zhu W., Apweiler R., Aebersold R., Deutsch E. W., Dunn M. J., Heck A. J., Leitner A., Macht M., Mann M., Martens L., Neubert T. A., Patterson S. D., Ping P., Seymour S. L., Souda P., Tsugita A., Vandekerckhove J., Vondriska T. M., Whitelegge J. P., Wilkins M. R., Xenarios I., Yates J. R., 3rd and Hermjakob H. (2007). "The minimum information about a proteomics experiment (MIAPE)." *Nature Biotechnology* **25**(8): 887-893.

Theophilos M. B., Cox D. W. and Mercer J. F. (1996). "The toxic milk mouse is a murine model of Wilson disease." *Human Molecular Genetics* **5**(10): 1619-1624.

Thiele D. J. and Gitlin J. D. (2008). "Assembling the pieces." *Nature Chemical Biology* **4**(3): 145-147.

Thierse H. J., Helm S. and Pankert P. (2008). "Metalloproteomics in the molecular study of cell physiology and disease." *Methods in Molecular Biology* **425**: 139-147.

Thomas G. R., Forbes J. R., Roberts E. A., Walshe J. M. and Cox D. W. (1995). "The Wilson disease gene: spectrum of mutations and their consequences." *Nature Genetics* **9**(2): 210-217.

Tufte E. R. (2001). *The Visual Display of Quantitative Information*. Cheshire, CT, Graphics Press.

Twardy C., Gardner S. and Dowe D. L. (2005). "Empirical data sets are algorithmically compressible: reply to McAllister?" *Studies in History and Philosophy of Science* **36**(2): 391-402.

University O. S. "A Brief Guide to Interpreting the DNA Sequencing Electropherogram Version 3.0." Retrieved 6 November 2014, from https://pmgf.osu.edu/sites/pmgf.osu.edu/files/guide_to_electropherogram_v3.pdf.

Urbach P. (1987). *Francis Bacon's Philosopy of Science: an Account and Reappraisal*. Lasalle, IL, Open Court.

van De Sluis B., Rothuizen J., Pearson P. L., van Oost B. A. and Wijmenga C. (2002). "Identification of a new copper metabolism gene by positional cloning in a purebred dog population." *Human Molecular Genetics* **11**(2): 165-173.

Van Dyck M. (2005). "The paradox of conceptual novelty and Galileo's use of experiments." *Philosophy of Science* **72**(5): 864-873.

Van Summeren A., Renes J., Bouwman F. G., Noben J. P., van Delft J. H., Kleinjans J. C. and Mariman E. C. (2011). "Proteomics investigations of drug-induced hepatotoxicity in HepG2 cells." *Toxicological Sciences* **120**(1): 109-122.

Vaudel M., Sickmann A. and Martens L. (2012). "Current methods for global proteome identification." *Expert Reviews in Proteomics* **9**(5): 519-532.

Verheggen K., Barsnes H. and Martens L. (2014). "Distributed computing and data storage in proteomics: many hands make light work, and a stronger memory." *Proteomics* **14**(4-5): 367-377.

Vickers B. (1992). "Francis Bacon and the progress of knowledge." *Journal of the History of Idea* **53**(3): 495-518.

Vidal M., Cusick M. E. and Barabasi A. L. (2011). "Interactome networks and human disease." *Cell* **144**(6): 986-998.

Vogel J. (2000). "Reliabilism leveled." *The Journal of Philosophy* **97**(11): 602-623.

Wake M. H. (2003). "What is "Integrative Biology"?" *Integrative Comparative Biology* **43**(2): 239-241.

Wake M. H. (2008). "Integrative biology: science for the 21st century." *Bioscience* **58**(4): 349-353.

Wasinger V. C., Cordwell S. J., Cerpa-Poljak A., Yan J. X., Gooley A. A., Wilkins M. R., Duncan M. W., Harris R., Williams K. L. and Humphery-Smith I. (1995). "Progress with gene-product mapping of the Mollicutes: Mycoplasma genitalium." *Electrophoresis* **16**(1): 1090-1094.

Watson J. D. (1968). *The Double Helix. A Personal Account of the Discovery of the Structure of DNA*. New York, Atheneum.

Watson J. D. and Crick F. H. C. (1953). "Molecular structure of nucleic acids. A structure for deoxyribose nucleic acid." *Nature* **171**(4356): 737-738.

Watson J. D. and Crick F. H. C. (1953). "The structure of DNA." *Cold Spring Harbor Symposia on Quantitative Biology* **18**: 123-131.

Weber M. (2005). *Philosophy of Experimental Biology*. New York, Cambridge University Press.

Weber M. (2012). "Experiment in Biology". *The Stanford Encyclopedia of Philosophy*. E. N. Zalta

Whewell W., Ed. (1968). *William Whewell's Theory of Scientific Method*. Pittsburgh, University of Pittsburgh Press.

Wilschanski M. A., Rozmahel R., Beharry S., Kent G., Li C., Tsui L. C., Durie P. and Bear C. E. (1996). "In vivo measurements of ion transport in long-living CF mice." *Biochemical and Biophysical Research Communications* **219**(3): 753-759.

Wilson D. C., Phillips M. J., Cox D. W. and Roberts E. A. (2000). "Severe hepatic Wilson's disease in preschool-aged children." *The Journal of Pediatrics* **137**(5): 719-722.

Wilson S. A. K. (1912). "Progressive lenticular degeneration: a familial nervous disease associated with cirrhosis of the liver." *Brain* **34**(Pt 4): 295-507.

Wimsatt W. C. (2007). *Re-engineering Philosophy for Limited Beings*. Cambridge, MA, Harvard University Press.

Wolkenhauer O. (2001). "Systems biology: the reincarnation of systems theory applied in biology?" *Briefings in Bioinformatics* **2**(3): 258-270.

Worthey E. A., Mayer A. N., Syverson G. D., Helbling D., Bonacci B. B., Decker B., Serpe J. M., Dasu T., Tschannen M. R., Veith R. L., Basehore M. J., Broeckel U., Tomita-Mitchell A., Arca M. J., Casper J. T., Margolis D. A., Bick D. P., Hessner M. J., Routes J. M., Verbsky J. W., Jacob H. J. and Dimmock D. P. (2011). "Making a definitive diagnosis: successful clinical application of whole exome sequencing in a child with intractable inflammatory bowel disease." *Genetics in Medicine* **13**(3): 255-262.

Wu J., Forbes J. R., Chen H. S. and Cox D. W. (1994). "The LEC rat has a deletion in the copper transporting ATPase gene homologous to the Wilson disease gene." *Nature Genetics* **7**(4): 541-545.

Yamagishi N., Ueda T., Mori A., Saito Y. and Hatayama T. (2012). "Decreased expression of endoplasmic reticulum chaperone GRP78 in liver of diabetic mice." *Biochemical and Biophysical Research Communications* **417**(1): 364-370.

Yanagimoto C., Masaru Harada M., Kumemura H., Abe M., Koga H., Sakata M., Kawaguchi T., Terada K., Hanada S., Taniguchi E., Ninomiya H., Ueno T., Sugiyama T. and Sata M. (2011). "Copper incorporation into ceruloplasmin is regulated by Niemann–Pick C1 protein." *Hepatology Research* **41**(5): 484-491.

Yano J., Rajendran A., Valentine M. S., Saha M., Ballif B. A. and Van Houten J. L. (2013). "Proteomic analysis of the cilia membrane of Paramecium tetraurelia." *Journal of Proteomics* **78**: 113-122.

Yi L., Donsante A., Kennerson M. L., Mercer J. F., Garbern J. Y. and Kaler S. G. (2012). "Altered intracellular localization and valosin-containing protein (p97 VCP) interaction

underlie ATP7A-related distal motor neuropathy." *Human Molecular Genetics* **21**(8): 1794-1807.

Yook S. H., Oltvai Z. N. and Barabasi A. L. (2004). "Functional and topological characterization of protein interaction networks." *Proteomics* **4**(4): 928-942.

Zagzebski L. (2003). "The search for the source of epistemic good." *Metaphilosophy* **34**(1-2): 12-28.

Zhu H. and Snyder M. (2003). "Protein chip technology." *Current Opinion in Chemical Biology* **7**(1): 55-63.

APPENDIX A

No chapter in this dissertation was previously published as a stand-alone paper; however, I hereby acknowledge the following:

Some sections of Chapter 2 in essence were in the following 2012 publication:

Using metalloproteomics to investigate the cellular physiology of copper in hepatocytes
E. A. Roberts, *Metallomics*, 2012, **4**, 633
DOI: 10.1039/C2MT20019H

Authors contributing to RSC publications (journal articles, books or book chapters) do not need to formally request permission to reproduce material contained in this article provided that the correct acknowledgement is given with the reproduced material.

If you are the author of this article you do not need to formally request permission to reproduce figures, diagrams etc. contained in this article in third party publications or in a thesis or dissertation provided that the correct acknowledgement is given with the reproduced material.

I interpret their policy (above) as not to require a copyright transfer letter.

Some other sections of Chapter 2 appeared in the following 2014 publication:

Accession Number	00075197-201409000-00009.
Author	Roberts, Eve A. a,b,d,e,f; Sarkar, Bibudhendra c,g
Institution	(a)Division of Gastroenterology, Hepatology and Nutrition, The Hospital for Sick Children (b)Genetics and Genome Biology Program (c)Molecular Structure and Function Program, The Hospital for Sick Children Research Institute (d)Departments of Paediatrics (e)Medicine (f)Pharmacology (g)Biochemistry, University of Toronto, Toronto, Ontario, Canada
Title	Metalloproteomics: focus on metabolic issues relating to metals.[Miscellaneous Article]
Source	Current Opinion in Clinical Nutrition & Metabolic Care. 17(5):425-430, September 2014.

Daniel Hyde of Wolters Kluwer confirmed in an email note dated 26 February 2015 that no permission letter is needed from Wolters Kluwer Health, Lippincott Williams &

Wilkins. Full acknowledgement is required, but that is all. I believe that I have conformed to that policy.