14 The distribution of income, wealth, and economic security: the impact of unemployment insurance reforms in Canada

Lars Osberg, Sadettin Erksoy and Shelley Phipps

Introduction

> The development of the modern welfare states can only be explained by some general aversion to risk . . . if a society, by free choice and democratic procedure, introduces a welfare state, risk aversion must in some way dominate the decision process. (Borch, 1968, p. 74)[1]

Most of the debate on reform of the welfare state has focused on the size of income transfers to and from various groups in society and the possible incentive effects of social welfare programs on labor market behavior. However, very little attention has been paid to the role which potential social transfers also play in reducing economic insecurity for all members of society, although reducing economic insecurity is one of the fundamental reasons for the existence of the welfare state. Insecurity and economic inequality are linked, since the value of the increased security produced by social insurance programs is very different at different points in the income distribution..Changes to the social insurance role of the welfare state therefore affect the overall distribution of economic well-being both directly, by changing the expected transfer income of different groups in society, and indirectly, by altering the economic security and the labor market behavior of individuals.

Since unemployment insurance is the largest single Canadian social insurance program,[2] this chapter focuses on the change in the distribution of economic well-being in Canada implied by changes in unemployment insurance. Over the last two decades, benefit/wage replacement rates have been reduced, entrance requirements have been tightened, and benefit durations reduced. For some workers (e.g. a tenured professor), the impact of these changes has been very small. People who have (1) nil risk of unemployment; or (2) relatively high income (hence relatively small marginal utility of income changes); or (3) sufficient assets and/or access to credit to maintain consumption during spells of low earnings, do not need unemployment insurance to provide income security – hence they lose little by its reduction.

However, low income workers (e.g. research assistants) typically do not have these characteristics. In general, the value of the security which unemployment insurance provides depends on the probability of an individual becoming unemployed, their ability to smooth consumption by borrowing or dissaving, and the change in utility implied by changes in consumption. Since low income individuals have a tendency to (1) face a relatively high probability of unemployment, (2) have a high marginal utility of income, and (3) have low assets and/or poor access to credit, the determinants of the security value of unemployment insurance are correlated. Individuals also differ in the degree to which they can change their labor market behavior in response to changes in unemployment insurance parameters. Given the correlation of these factors, the objective of this chapter is to assess the impact on the distribution of economic well-being of changes in the level of unemployment insurance, as an example of the more general issue of the impact of the welfare state on the distribution of economic well-being.

The next section outlines our theoretical framework for discussion of income risk, and the reasons why we believe it is appropriate to measure the income value of the change in expected utility associated with changes in unemployment insurance regimes, rather than the change in expected income associated with social policy reform. We then discuss the empirical calculations, which have two major steps. We first use a micro-simulation model to estimate, for a representative sample of individuals, the fluctuations in labor earnings and unemployment insurance receipts associated with alternative Canadian unemployment insurance regimes (those in place in 1971 and 1994). The model incorporates individual behavioral responses, such as changes in the incidence and duration of unemployment, which may result from changes in unemployment insurance legislation. Given estimated changes in labor earnings and unemployment insurance transfers, we then use a Stone–Geary utility function to calculate the present discounted values of individual utility associated with each unemployment insurance regime. These are used to calculate the implied certainty equivalent income value of the change in expected utility associated with decreased protection from income loss due to unemployment. We then discuss the distribution of these changes in economic well-being within the Canadian income distribution, and finally draw some conclusions.

Theoretical perspective

In the market for private insurance against such risks as fire, automobile theft, or loss of life, individuals can be observed voluntarily paying premiums which exceed the expected value of their insurance claims (due to administration costs and profits within the insurance industry). Although the marginal individual may be indifferent between purchasing or not purchasing insurance, infra-marginal individuals with greater risk aversion receive net gains in their utility from the existence of the option to decrease risk. Although the net expected income of insurance buyers is on average lower than in the absence of insurance (by the amount of the administration costs and profits of the insurance industry), the average utility of insurance buyers is increased by their ability to avoid risk. In private insurance markets, one can use the demand for insurance at different premium levels to estimate willingness to pay for insurance protection and the implied social benefits of the existence of the option to insure against particular risks.

Unemployment insurance (UI) also provides insurance against the risk of income loss, but evaluation of UI presents a range of difficulties not present in the private insurance example. Private markets for UI are not feasible, partly because private insurers would face huge losses in recessionary times, due to the correlation of unemployment outcomes across individuals. Individuals with private knowledge of their greater risk of personal unemployment could also be expected to self-select for the purchase of UI, at any given premium rate.[3] UI is not, therefore, voluntarily purchased.[4]

Furthermore, since UI, like other social insurance programs, embodies ideas of *social* fairness, not *actuarial* fairness, one can expect to observe systematic differences across individuals in the ratio of expected UI benefits to UI premiums.[5] Nonetheless, UI retains an insurance function – since UI reduces the risk of income loss due to unemployment, risk averse individuals can be expected to be willing to pay some amount, which differs according to individual circumstance, for the reduction in income risk which UI provides.

To fix ideas, figure 14.1 outlines the position of a risk averse individual (decreasing marginal utility of income) who faces some risk of income loss. It contrasts the situation of the individual in a regime with high UI benefits for high premiums (aa') and a regime with low UI benefits and low premiums (bb'). If the individual remains fully employed, a net income of Y'_b is received, when the low benefit UI scheme is in operation. However, since higher UI benefits require higher UI premiums, Y'_a is the

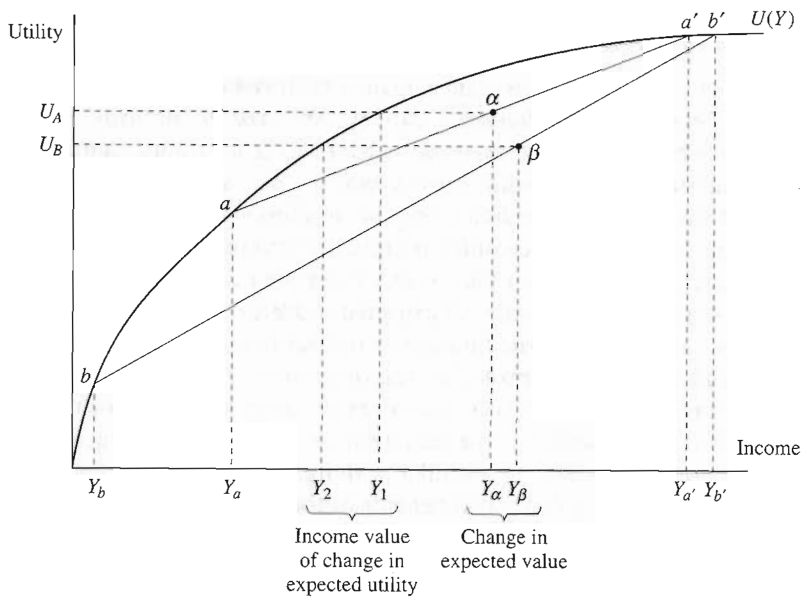

Figure 14.1 The impact on utility of alternative unemployment insurance regimes

individual's income in the absence of unemployment when the high benefit UI scheme is in operation. If the individual experiences unemployment, total income from labor earnings and UI benefits is Y_a if the high benefit UI scheme is in operation, and Y_b under the low benefit UI regime.

The expected income of each individual is determined by their relative probability of unemployment. In figure 14.1 Y_α is the expected income under regime A, where point α defined by $Y_\alpha Y_{a'}/Y_a Y_{a'}$ is the probability of unemployment. Similarly, $Y_\beta Y_{b'}/Y_b Y_{b'}$ is the probability of unemployment under regime B. Usually, $Y_\beta Y_{b'}/Y_b Y_{b'} \neq Y_\alpha Y_{a'}/Y_a Y_{a'}$. Although it is this change in the probability of unemployment that is the focus of most of the literature on UI, one can argue that the focus *should* be the impact of UI on individual well-being $(U_A - U_B)$.

There is no reason to believe that individuals face the same probability of unemployment under different UI regimes – indeed most of the literature on UI is about the possible impacts of the implicit incentives of UI on the probability of unemployment (for example, see Atkinson and Micklewright, 1991). The simulation model which we discuss on p. 337 is built up from a series of estimated behavioral equations which embodies the responses of individuals to changes in the specific parameters of UI in

Canada – hence changes in UI regimes affect individual probabilities of unemployment.

Given these behavioral responses, the relative probability of unemployment corresponding to each UI scheme implies that the expected value of income under the low benefit UI scheme is Y_β and under the high benefit scheme Y_α. The levels of utility associated with these uncertain income streams are graphed on the vertical axis as U_A and U_B. Individuals who are averse to risk get the same level of utility from a certain income as they would get from an uncertain lottery with higher expected value. In figure 14.1, Y_1 is the certainty equivalent income which produces the same level of utility as the expected value of income (Y_α) which the individual would receive under the higher benefit UI scheme. In figure 14.1, Y_α *minus* Y_1 represents the risk premium – the amount which the individual would be willing to pay to rid themselves entirely of the income risk of unemployment (i.e. receive a certain income rather than the uncertain prospect of income Y'_a if no unemployment and Y_a if unemployed). Similarly, Y_2 is the certain income which would generate the utility level U_B, the same level of utility as generated by the uncertain prospect of Y'_b if not unemployed and Y_b if unemployed under the less generous UI scheme. The change in utility associated with the change in UI regimes is $U_A - U_B$ and the money equivalent of that loss in utility (the change in certainty equivalent income) is $Y_1 - Y_2$.

Note that the change in certainty equivalent income has arisen from changes in *both* the level *and* the riskiness of income flows. The change in level of income arises both from changed behavior (hence changes in earnings) and changes in UI transfers. The impact of changes in unemployment insurance legislation on the distribution of economic well-being may not be captured by simple calculation of changes in the expected value of income. In figure 14.1, shifting from a high benefit/high premium to a low benefit/low premium UI scheme *increases* the expected value of income (from Y_α to Y_β). However, the decrease in income security which this entails produces a *decrease* in net utility (from U_A to U_B), the income value of which is given by $Y_1 - Y_2$.[6]

The contrast in figure 14.1 between social insurance revisions which produce a *gain* in expected income and a *loss* in certainty equivalent income deserves some emphasis. Much of the literature on social policy reform (including some of our own work – Erksoy, Osberg and Phipps, 1994a, 1994c) has calculated the change in realized income flows associated with social policy changes and ignored the value of income security. The lesson of figure 14.1 is that even those who realize increases in expected income may experience losses in utility.

Empirical calculations

Micro-simulation of income and labor market outcomes under the 1971 and 1994 UI regimes

The rationale underlying micro-simulation is the idea that if we are to understand how the economy would have functioned under a different set of incentives (such as those embodied in different UI regimes), we have to take full account of the heterogeneity of individual characteristics, the interdependence of economic processes, and the endogeneity of individual characteristics over time. The origins of micro-simulation lie in the work of Guy Orcutt and his colleagues (Orcutt, 1957; Orcutt, Merz and Quinke, 1986). The basic idea is to take microdata on a representative panel sample of individuals and simulate the impact over time of alternative policy scenarios on *each* individual member of the panel. The aggregate implications of changes in policy are, therefore, built up from explicit calculation of its impacts on each and every individual. One of the major advantages of micro-simulation modeling is that it allows the simultaneous consideration of policy impacts on a number of different behaviors (e.g. the impacts of UI on both unemployment and labor force participation) which imposing consistency of predicted impacts (e.g. weeks unemployed *plus* weeks employed *plus* weeks not in the labor force must sum to 52).

In a series of papers and reports (Erksoy, Osberg and Phipps, 1994a, 1994b, 1994c), we have outlined the methodology of our micro-simulation model, presented explicitly the estimated behavioral equations which drive the model and demonstrated its sensitivity to alternative assumptions (e.g. the importance of past labor market outcomes). We present an abbreviated description of the model in appendix B (p. 337) and refer interested readers to our other work for additional detail.

For present purposes, the key thing about our microsimulation model is that it generates, for each of the 19,488 respondents to Statistics Canada's 1984 Survey of Assets and Debts, a predicted vector of labor earnings, UI receipts, weeks unemployed, weeks employed, and weeks not in the labor force for each year of the business cycle. We run the model with two alternative policy frameworks – Canadian 1971 UI legislation and 1944 UI legislation – and compare the results.

The impact of revisions to UI on well-being

If changes to social insurance have only small impacts on income riskiness, and produce the same expected value of income flows, then the

Arrow–Pratt risk premium can be estimated directly from data on absolute risk aversion, without explicit representation of the underlying utility function of individuals. However, the Canadian revisions to UI between 1971 and 1994 are not like that. When policy revisions have relatively large impacts on both income riskiness *and* the expected value of some individuals' incomes, calculation of the impact of those policy changes on certainty equivalent income requires the specification of some underlying utility function.

Economists typically use the idea of utility maximization to drive their theories of individual behavior, and most find it easy to agree on general characteristics – e.g. individual utility is a positive function of income, but there is declining marginal utility of income and wealth. Beyond that, there is little consensus. Indeed, although empirical economists have estimated many models of individual behavior (such as labor supply) whose empirical specification also implies an underlying individual utility function, the characteristics of these implicit utility functions have often not been carefully examined (as Stern, 1986, has emphasized).

In this chapter, we use estimates of Canadian labor supply behavior based on the Stone–Geary utility function to estimate the utility value of income flows associated with 1971 and 1994 UI regimes, over the 1981–89 business cycle. In each case, we calculate the present discounted value of yearly flows of utility over the entire period, using a discount rate of $5\frac{1}{2}$ percent per annum (the average real interest rate on home mortgage indebtedness in Canada from 1981 to 1989). The change in well-being attributable to UI revisions is the difference between the present discounted value of utility under the alternative UI regimes. To put a money value on this change in utility, we calculate the certainty equivalent income which corresponds to the average annual level of utility under each UI regime, at average annual levels of labor supply.

In calculating the change in certainty equivalent income for individuals which correspond to a change in policy regimes we are, in essence, asking each individual: "*If you knew in advance how you would fare under each policy regime, what sum of money (positive or negative) would make you equally well off in one policy regime, compared to the other?*" This chapter asks specifically: "*If you knew in advance how you would have fared, over a business cycle such as that of 1981 to 1989, under the 1994 rules of Canadian UI legislation compared to the 1971 rules, what sum of money per year would make you equally well off with 1994, compared to 1971 UI legislation?*" This can be seen as performing an *ex ante* evaluation of the realized (*ex post*) outcomes of policy reform – in a sense we are asking what fully informed individuals would think would be the implications of policy reform (which may perhaps be a somewhat idealistic type of policy evaluation).

In table 14.1, the Y^*_{71} is the certainty equivalent annual income corresponding to the 1971 UI regime and Y^*_{94} is the certainty equivalent annual income corresponding to the 1994 UI regime. More specifically, define:

Y_{it} = income from earnings and UI under 1971 legislation

Y'_{it} = income from earnings and UI under 1994 legislation

L_{it} = weeks of non-labour time under 1971 legislation

L'_{it} = weeks of non-labour time under 1994 legislation

i = individual

t = year (1981–9)

$$\bar{L}_i = \sum_{t=0}^{8} L_{it}/9$$

$$U_{it} = U(Y_{it}, L_{it})$$

$$U'_{it} = U(Y'_{it}, L'_{it})$$

$\hat{U}_i = \sum_{t=0}^{8} \frac{U_{it}}{(1+0.055)^t}$ = present value of utility of individual i under 1971 legislation

$\hat{U}'_i = \sum_{t=0}^{8} \frac{U'_{it}}{(1+0.055)^t}$ = present value of utility of individual i in period t under 1994 legislation

$$Y^*_{71} = U^{-1}(\hat{U}_i/9 \mid L_{it} = \bar{L}_i)$$

$$Y^*_{94} = U^{-1}(\hat{U}'_i/9 \mid L_{it} = \bar{L}_i)$$

In evaluating the labor supply responses of Canadian households to UI reforms, Phipps (1990, 1991a, 1991b) used the Stone–Geary functional form

$$U_{it} = B[Y_{it} - \gamma_0]^{1-\theta}[L_{it} - \gamma_1]^{\theta}$$

where Y_{it} is earnings *plus* UI transfers received by individual i in period t and L_{it} is weeks of non-labor time.

The Stone–Geary functional form has the convenient feature that Phipps (1990, 1991a, 1991b) has already provided estimates of its

Table 14.1 *Impact of 1971–94 UI revisions on average annual "income," by deciles, ages 16 to 65**

	Males			Females		
	Certainty equivalent income			Certainty equivalent income		
Decile	Y^*_{71}	Y^*_{94}	$\Delta Y^*\%$	Y^*_{71}	Y^*_{94}	$\Delta Y^*\%$
1	619	280	54.8	598	565	5.5
2	2,917	1,874	35.8	1,751	1,679	4.1
3	7,015	6,059	13.6	2,893	2,767	4.4
4	10,435	9,602	8.0	3,534	3,404	3.7
5	14,226	13,251	6.9	4,083	3,967	2.8
6	17,759	16,899	4.8	5,772	5,479	5.1
7	21,232	20,317	4.3	8,557	8,140	4.9
8	25,193	24,266	3.7	12,024	11,515	4.2
9	30,612	29,708	3.0	16,468	15,915	3.4
10	45,390	44,487	2.0	26,510	25,932	2.2
Overall	17,539	16,674	13.7	8,219	7,936	4.0
Gini	0.42003	0.43923	−4.6	0.494	0.498	−0.8
Theil	0.30919	0.34073	−10.6	0.411	0.417	−1.5
CV	0.7850	0.8210	−4.6	0.996	1.010	−1.4
Atkinson						
$r =$ 0.5	0.17377	0.19662	−13.1	0.205	0.207	−0.9
$r = -0.5$	0.70162	0.77475	−10.4	0.633	0.633	0.0
$r = -1.0$	0.91471	0.93762	−2.5	0.859	0.862	−0.3

Note: Labor earnings and UI receipts only, 1981$, for 1981 labor market participants.

parameters in the explicit context of the influence of UI on labor supply.[7] Its disadvantages include the fact that all goods must be substitutes, Engel curves are constrained to be linear and the labor supply function is constrained to be monotonic with respect to wages.[8] Another feature of the Stone–Geary functional form (which we dislike) is that is assumes that all non-labor time has positive utility value. An alternative perspective (dominant in social psychology and sociology) sees unemployment in a more negative light, as something which is psychologically harmful, which people typically do not like (see Hayes and Nutman, 1981; Jahoda, 1979 and Kelvin and Jarrett, 1985). In the economics literature, Narendranathan and Nickell (1985) found a negative valuation on unemployment time, while Clark and Oswald (1994, p. 658) use survey evidence to argue strongly that "unemployed people in Great Britain in 1991 have much lower levels of mental well-being than those in work." In future research, we hope to devise a way of distinguishing the relative utility of time spent unemployed, time spent in leisure, and time engaged in unpaid work. However, since our focus in this chapter is on *income* security, we proceed with the Stone–Geary specification.

Distributional impact of social insurance revisions

Table 14.1 presents the average annual certainty equivalent income[9] for men and women and for each decile of the income distribution under the 1971 and 1994 regimes, as evaluated over the business cycle from 1981 to 1989. Although the dollar value of the average income change associated with the 1971 to 1994 revisions to UI is roughly comparable across many of the deciles of the male distribution, this income loss is a much larger fraction of the annual income of poorer deciles than of richer. For males, there is thus a very significant decrease in income associated with UI cutbacks, which is unambiguously much larger, in percentage terms, for the poorer deciles of the income distribution. For females, it is clear that the 1971–94 revisions to UI in Canada increase the inequality of the distribution of economic well-being – but the distributional impacts are smaller and less clearly concentrated on poorer deciles of the income distribution.[10]

Several caveats are, however, in order. First, since the income concept used in this chapter is that of annual labor earnings *plus* annual receipts of unemployment insurance payments, capital income and (more importantly) pension income and social assistance receipts are excluded. The relatively low annual earnings and unemployment insurance receipts of poorer deciles reflects in part a tendency of those with long duration unemployment spells to withdraw entirely from labor force participation – a tendency which is particularly important for older cohorts and for women. Since our objective is to model the distributional impacts of a particular social insurance program (UI), we do not build in any assumptions of automatic receipt of social assistance by individuals with low annual income or automatic receipt of pension income by retirees. Clearly, however, UI is part of the larger welfare state of public and private social transfers, and the extent to which other transfers will kick in to offset cutbacks in UI is a crucially important issue.

Furthermore, one must emphasize that table 14.1 refers to the population of individuals who participated in the labor market in 1981 (some of whom may have withdrawn from the labor force from 1982–9). Table 14.1 contains no consideration of household income or household size.

Table 14.2 *Males, 16 to 65: Stone–Geary utility functions, percentage gains/losses from UI revisions, by decile of income under 1971 scheme*

	Loss			Nil	Gain		
	More than 50% loss (1)	26%– 50% loss (2)	6%– 25% loss (3)	No change ± 5% (4)	6%– 25% gain (5)	26%– 50% gain (6)	More than 50% gain (7)
Decile							
1	47.02	14.14	5.90	21.65	1.48	1.40	8.41
2	34.53	11.59	14.20	28.58	2.10	0.99	8.02
3	10.45	8.64	18.21	56.60	2.66	1.40	2.04
4	4.95	6.53	18.28	67.70	1.08	0.82	0.64
5	3.82	4.74	14.80	75.35	0.53	0.21	0.56
6	1.84	3.91	15.27	78.27	0.47	0.00	0.23
7	2.11	1.39	14.68	81.55	0.27	0.00	0.00
8	1.18	1.60	11.28	85.87	0.00	0.00	0.07
9	1.31	1.96	12.39	84.31	0.00	0.03	0.00
10	0.43	1.18	6.85	91.35	0.12	0.00	0.07

The Stone–Geary utility function underlying the calculations of certainty equivalent annual income is that for single males and females. Evaluating the distribution of utility across households raises complex issues of equivalence of incomes (see Phipps, 1997) and of the division of resources within households (see Phipps and Burton, 1995) which we hope to address more fully in future work.

Table 14.1 is based on deciles of incomes, as ordered by 1971 and 1994 income, respectively. Since UI revisions imply that some individuals experience gains, while others experience losses, the individuals in each decile of 1971 and 1994 income are not all the same. To illustrate the dispersion in impacts within income deciles, tables 14.2 and 14.3 illustrate the distribution of the percentage change in income, by decile of original income under the 1971 UI regime. All these tables tell essentially the same story (although much stronger for males than for females). The percentage of individuals which is essentially unaffected by revisions to unemployment insurance rises with income – over 90 percent of the top income decile are essentially unaffected by unemployment insurance revisions, while only 20 percent of the males (60 percent of the females) in the bottom income decile are similarly unaffected. Within the lower deciles of the income distribution, there is a very significant minority who experience very large percentage losses in annual income, and a very much smaller proportion who experience gains.

Table 14.3 *Females, 16 to 65: Stone–Geary utility functions, percentage gains/losses from UI revisions, by decile of income under 1971 scheme*

	Loss			Nil	Gain		
	More than 50% loss (1)	26%– 50% loss (2)	6%– 25% loss (3)	No change ± 5% (4)	6%– 25% gain (5)	26%– 50% gain (6)	More than 50% gain (7)
1	3.50	4.32	15.07	62.66	9.57	1.50	3.38
2	1.47	4.41	16.66	70.00	5.52	0.93	1.00
3	1.86	2.02	14.19	79.16	1.67	0.55	0.56
4	0.59	0.39	12.34	86.32	0.37	0.00	0.00
5	0.67	3.76	14.85	80.57	0.14	0.00	0.00
6	0.36	1.48	17.96	79.09	0.98	0.09	0.04
7	0.64	2.54	19.87	75.95	0.77	0.23	0.00
8	0.19	1.69	14.41	83.64	0.00	0.00	0.00
9	0.53	1.07	11.57	86.51	0.32	0.00	0.00
10	0.10	0.65	6.03	93.03	0.19	0.00	0.00

So far, this chapter has focused on the utility gains and losses associated with the risk of *income* fluctuations, but economists normally think of utility as being derived not from *income*, but from the *consumption* which income permits. However, in order to go from a discussion of variations in income flows to an analysis of consumption flows, one needs to discuss saving and dissaving behavior. Up to this point, intertemporal utility maximization has been based on a time separable utility function and a discount rate of 5.5 percent, but savings behavior depends on more than that. The issue we are concerned with in this chapter is income insecurity, and to the degree that income flows are more uncertain, choices between more and less liquid forms of saving (bank accounts and term deposits versus home ownership, automobiles, durables, pensions) will shift toward the liquid assets available to smooth consumption during a spell of unemployment, rather than the illiquid assets which are not. Changes in social insurance legislation are therefore likely to alter both the level and the composition of personal savings.

A full resolution of these issues lies beyond the scope of this chapter. Furthermore, Canada does not now have[11] longitudinal panel data which would enable one to assess directly the impact of income fluctuations over time on individual consumption and savings patterns. Table 14.4 therefore simply asks what percentage of the population in fact has enough assets to smooth consumption to any significant degree – i.e. we use cross-sectional evidence on net household assets to assess the relative importance of liquidity constraints (i.e. the percentage of individuals who are members of households whose assets or access to credit are insufficient to smooth consumption, so that $Y_{it} \simeq C_{it}$ necessarily). Columns (1) and (3) of table 14.4 present the percentage of persons within each 1981 income decile living in households where liquid assets are greater than $3,020[12] while columns (2) and (4) add to those liquid assets the credit limit on available home equity loans.[13]

Table 14.4 *Income deciles and net liquid assets* [a]

	Males		Females	
Deciles	Net liquid assets (1)	Available home equity (2)	Net liquid assets (3)	Available home equity (4)
Poorest 10%	46.43	76.84	33.80	65.92
11–20	32.15	61.86	31.89	58.51
21–30	21.21	44.70	27.11	56.20
31–40	21.74	49.88	30.58	64.08
41–50	24.04	54.75	32.90	60.56
51–60	26.57	58.52	30.07	60.95
61–70	29.17	65.15	31.14	57.30
71–80	34.88	70.01	28.71	59.22
81–90	38.00	76.00	32.89	61.46
Top 10%	47.00	81.88	38.35	67.15
Overall average	*32.12*	*63.96*	*31.74*	*61.14*

Note:
[a] Percentage with greater than $3,020, 16 to 64 years; $3,020 = 1981 single person poverty line for 21.6 weeks (the average uninterrupted duration of a spell of unemployment in 1984).

Conclusions

The innovation of this chapter is its attempt to evaluate the impacts on the distribution of economic well-being of revisions to the social insurance role of the welfare state by combining micro-simulation of behavioral responses to policy changes with an explicit specification of the utility functions of individuals. We have argued that reforms to the welfare state can be expected to alter individual labor market behaviors and that individuals care about both the expected value of their income and the riskiness of that income – hence it is necessary to evaluate jointly the expected value of future income, the uncertainty of income flows, and the behavioral responses of individuals.

Since our micro-simulation model incorporates behavioral response to changes in UI parameters, it explicitly takes account of "moral hazard" in UI induced changes in labor market behavior. Moreover, incorporating behavioral response in a logically consistent model structure implies that we consider the implications of changes in *both* labor force participation and unemployment. Part of the "moral hazard" story about Canadian UI is the argument that easier entrance requirements after 1971 increased the labor supply of some "marginal" workers, who entered the labor force in order to qualify for UI benefits (see Card and Riddell, 1993) – the implication is that when UI qualification becomes more difficult, as in 1994, such workers will decrease their labor supply, thereby decreasing total earnings. Even though the unemployment rate falls when there is a change to a more restrictive UI program, this produces a drop in labor force participation which means that average total incomes also fall.

In doing all this, we are well aware that calculation of the value of a Stone–Geary utility function under alternative UI and unemployment scenarios is a rather bloodless exercise. In the real world, people stay up late at night worrying about how they are going to make ends meet and what they will do if they lose their jobs. Economic anxiety makes real people tense and irritable and they may lash out at their families or try to drown their anxieties in alcohol or drugs. Our model does not capture the sense of anxiety and powerlessness that comes from the increasing withdrawal of the Canadian state from its social insurance role and its denial of responsibility for macroeconomic outcomes.[14]

Much more needs to be done to improve our analysis. We are in the process of revising our micro-simulation model to incorporate demographic transitions and some consideration of household influences on labor market behavior. We need to know more about savings/dissavings and the ability of households to smooth consumption flows over time, despite variations in income flows. We want to experiment further with alternative specifications of the functional form of utility functions and the appropriate specification of utility for individuals within households, and we want to investigate the sensitivity of the importance of income risk to alternative parameter values for individual utility functions.

We have two major concerns with the Stone–Geary specification – that it may not exhibit enough risk aversion to income change and that it imposes a common (positive) valuation on all non-labor time. For the Stone–Geary utility function, evaluated for males at an income of $30,000 and two weeks' leisure, the Arrow–Pratt coefficient of relative

risk aversion is 0.28 – far below other estimates of risk aversion. The logarithmic utility function is popular in the financial literature (see Rubinstein, 1977) and is somewhat more consistent than the Stone–Geary specification with observed risk aversion behavior in financial portfolios, but the Arrow–Pratt coefficient of relative risk aversion is still only 1 – Friend and Blume (1975) concluded that the coefficient of relative risk aversion is at least 2.

Even when unemployment is relatively high (e.g. 10 percent), the vast majority of the labor force (90 percent) are employed – hence the majority of people have relatively little non-labor time, and it seems reasonable that they derive positive utility from it. However, the pleasures of another day with nothing to do decline rapidly as the days add up. Unemployment, particularly long duration unemployment, is very different from "leisure," and we are convinced by such authors as Clark and Oswald (1994) that unemployment generally reduces utility. In future work we hope to experiment with functional forms which allow the utility value of non-labor time to vary with its duration and nature.

Meanwhile, since the current Stone–Geary specification implies that increased non-labor time partly "balances" lower money income and since we suspect that our current results under-state the importance of income risk, a more satisfactory treatment of these two issues would very probably strengthen the conclusions of this chapter.

Although there are many technically complex steps to our argument, the basic conclusion is fairly direct. The 1971–94 revisions to UI in Canada substantially decreased the level of benefits paid by the Canadian UI system, and produced decreases in economic well-being for all deciles of the income distribution. Since declines in economic well-being were especially large, in percentage terms, for the poorest deciles of the distribution, UI revisions increased the inequality of the distribution of economic well-being, in addition to decreasing average economic well-being. And, as we have noted in other work (Erksoy, Osberg and Phipps, 1994c), UI revisions have the same qualitative impact on the distribution of income, whether measured by annual income or the present value of income over the business cycle.

Furthermore, since we get essentially the same result from examination of the distribution of certainty equivalent income, the present value of income flows over the business cycle, or annual money incomes, our hunch is that our results are fairly robust. And we also think that the Canadian public knows that reducing the social insurance role of the modern welfare state increases inequality and decreases economic well-being – when asked in March 1994 whether they approved or disapproved of reducing the level of UI benefits, 61 percent of Canadians replied that they disapproved – and the percentage was significantly higher among low income Canadians than among those with higher incomes.[15] Nevertheless, a major revision to UI in Canada was legislated in 1996, altering many plan parameters and changing the name to "Employment Insurance." When fully phased in, benefits payable are expected to decline by 11 percent nationally and up to 19 percent in provinces with high unemployment (HRDC, 1996).

Appendix A: key features of the 1971 and 1994 UI regimes

	Simulation 1981–9, based on 1971 UI system	Simulation 1981–9, based on 1994 UI system
Benefit wage ratio	(a) 75% of insured earnings for claimants with less than or equal to 1/3 maximum insurable earnings and with dependents (b) 66% for all other claimants	(a) 60% of insured earnings for claimants with less than or equal to 1/2 of maximum insurable earnings and with dependents (b) 55% for all other claimants
Maximum insurable earnings	From $315/week in 1981 to $460/week in 1989	From $315/week in 1981 to $460/week in 1989
Minimum insurable earnings	20% of maximum insurable earnings	20% of maximum insurable earnings
Minimum employment weeks to qualify	8 weeks	from 12 weeks to 20 weeks depending on regional unemployment rates
Maximum annual benefit period	50 weeks	50 weeks
Benefit period determination	(1) Up to 15 weeks of benefits for the first 20 weeks of work (2) Additional 10 weeks for the unemployed at the end of phase 1 (3) Up to 18 weeks of extended benefits for insurable employment beyond 20 weeks (4) If national unemployment rate is 4–5% then an additional 4 weeks of benefits; if the national unemployment rate is above 5% then an additional 8 weeks of benefits	(1) Up to 20 weeks of benefits, based on 1 week of benefits for every 2 weeks of work for the first 40 insured weeks of work (2) Up to 12 weeks of benefits, based on 1 week of benefits for each week of work beyond the first 40 weeks (3) Up to 26 weeks, based on 2 weeks of benefits for every percentage point by which the regional unemployment rate is above 4%

Appendix B: model structure

All equations are estimated using microdata from the 1986/87 Statistics Canada Labor Market Activities Survey.

Complete labor force withdrawal

In any simulation year over the cycle the model starts by computing for each individual the probability of being outside of the labor force for all 52 weeks of the year. The underlying regression is a logit model. Individuals are then ordered in descending order of the probability of being out of the labor force for 52 weeks. Those with highest probability of complete labor force withdrawal are assigned 52 weeks of non-participation weeks until the total number of such individuals reaches the exact proportion of the population with complete labor force withdrawal in that year. For those who are assigned 52 weeks of not being in the labor force, no further calculations of labor market behavior are made for that year. They may still collect UI benefits as a continuation of a claim whose duration has not yet expired from the previous year but they cannot establish a new UI claim in that year. However, they are retained in the model, since they may re-enter the labor force in a subsequent year.

Some weeks of labor force withdrawal

Given the remaining proportion of the population with some labor force participation, the model assigns each individual some weeks out of the labor force based on a Tobit model of weeks of non-labor force participation. Given that each individual has been assigned an estimate of their desired labor supply, the next step is to determine whether or not they can get employment for the weeks in which they are willing to participate in the labor market.

Weeks of unemployment

In this stage the model first computes for each individual a probability of unemployment experience using a logit specification. All individuals are then ordered in descending probability of unemployment. For each individual, unemployment weeks are calculated based on an accelerated failure time model. If an individual's unemployment experience this year is predicted to be greater than that of last year, the person faces no constraint in increasing the weeks of unemployment. If, however, the individual's unemployment weeks are less than that of last year, this person must locate additional employment weeks. Therefore, for all individuals with an expected decrease in unemployment the model computes the probability of being constrained in getting an additional week of work. The individual is assigned one more week of employment (one less week of unemployment) if a random number drawn from a uniform distribution exceeds the estimated probability of constraint. Those who want to increase their labor supply by more than one week of work, given that they have been successful in obtaining one additional week of employment, face a certain probability of being able to get the second additional week of work, etc. We proceed in this way until the individual has either reached the expected additional employment or encountered a constraint in obtaining an additional week of work. Together the duration model and the under-employment model determine for each person the expected weeks of unemployment, if they experience any unemployment.

Since individuals are ordered in descending order of probability of unemployment in a given year, unemployment is assigned to those with the highest probability of experiencing unemployment up to the point where the total number of assigned unemployment weeks equals aggregate unemployment weeks for the year. Aggregate unemployment weeks are given by the product of the total number of labor force weeks and the aggregate unemployment rate in that year.

Alternative scenarios

In each situation run, two scenarios are compared: the "base" and the "shock" scenario. The base and shock scenario differ in the assumed

structure of unemployment insurance legislation, but have identical estimated behavioral response to unemployment insurance parameters, and the influence of personal characteristics.

Stochastic component

In order to preserve the underlying stochastic element in labor market behavior, in each estimated equation random error terms drawn from a distribution with variance consistent with the observed unexplained variance are added to the conditional expectations. Furthermore, we assume that in the real world the underlying stochastic element consists of permanent and temporary features. The former is regarded as the persistent part of the total unobserved heterogeneity which remains constant in each year and represents 30 percent of the individual random error term in each equation. It is generated once and kept constant in each year. The temporary component which corresponds to the remaining 70 percent of the error term is generated separately in each year. Therefore, in a given year the sum total of the permanent and temporary components gives the total stochastic element in each behavioral equation. Further note that random error terms are initially generated for each individual in all behavioral equations and in all years. These random error terms are then retained and used in alternative policy simulations which conveniently allows us to compare alternative policy scenarios directly since we have the same distribution of the stochastic element in all simulations.

The effect of alternative UI systems on the aggregate unemployment rate

The model is embedded within a changing macroeconomic environment by allowing the aggregate unemployment rate to change over time and calculating the associated aggregate weeks of unemployment. The impact of unemployment insurance on aggregate unemployment is a hotly contested empirical issue in Canada. Myatt (1996) presents a summary of fourteen published studies on the impact of the liberalization of UI in 1971. As he notes: "Of these studies, seven found a significant positive effect [of UI on aggregate unemployment], five found no significant effect and two found no significant effect in seven out of ten provinces (it is worth noting that these latter studies disagree on which three provinces have the significant positive effect . . .). A more evenly divided result could not be imagined."

In this chapter, we assume that reductions in the generosity of UI coincide with reductions in aggregate unemployment. These adjustments are based on the presumption that there was a 0.6 percent increase in the unemployment rate due to the introduction of more generous UI regulations by the Unemployment Insurance Act in 1971 (see Grubel, Maki and Sax, 1975). The 1971 Act reduced the minimum employment weeks to qualify for UI benefits from 30 weeks over the previous 104-week qualifying period to 8 weeks in the previous 52-week period, increased the waiting period from one week to two weeks, and raised the replacement rate from 50 percent to 66 percent. In order to be able to extrapolate the effects of changes in the UI system on the aggregate unemployment rate we look at the behavior of an hypothetical individual who follows a repeated cycle of working the minimum required weeks in order to collect the maximum benefits under alternative UI systems. This individual could obtain 2.4 (=(51–15)/15) benefit weeks per week of employment in the pre-1971 period (assuming the 30-week entrance requirement is satisfied in the qualifying period), and 5.25 (=(50–8)/8) benefit weeks per week of employment in the post-1971 period. Therefore, the change in the UI generosity in the post-1971 period compared to the pre-1971 period is about 192 percent (=(3.5–1.2)/1.2). If then one assumes that 192 percent increase in the benefits leads to 0.6 percent increase in the unemployment rate due to the response of our hypothetical individual, one can calculate the effects of UI policy changes on the unemployment rate (U) as follows:

$$\%\text{ change in } U = (a/b \times 0.6\%)$$

where a is the \$ change in UI benefits for a given change in the regulations, and b is the \$ increase in UI benefits in the post-1971 period, which is 192 percent.

Overall, the changes in the UI system in place by 1994 represented a -44.5 percent change in generosity (compared to 1971) for the marginal individual, implying a -0.14 percent change in the aggregate unemployment rate.

Appendix C: Model equations[a]

$$(\Pr(WKSNLF_{it} < 52) = \hat{F}_1(X_{it}, WKSUN_{i,t-1}, LM_t) + \varepsilon_{2i} \qquad (1)$$

$$WK\hat{S}NLF_{it} = \hat{F}_1(X_{it}, WKSUN_{i,t-1}, LM_t, UI_t) + \varepsilon_{2i} \qquad (2)$$

$$\Pr(U_{it} \mid 0 \leq WKSNLF < 52) = F_3(X_{it}, WKSUN_{i,t-1}, LM_t) + \varepsilon_{3i} \qquad (3)$$

$$(WKS\hat{U}N_{it} \mid \hat{U}_{it} = 1) = \hat{F}_4(X_{it}, WKSUN_{i,t-1}, LM_t, UI_t) + \varepsilon_{4i} \qquad (4)$$

$$WKS\hat{E}MP_{it} = 52 - WKS\hat{N}LF_{it} - WK\hat{S}UN_{it} \qquad (5)$$

$$DESIRED_{it}[WKSUNEMP_i \mid UI'_t] \qquad (6)$$

$$IF\ DESIRED_{it} < WKS_{it}THEN \text{ Prob } [WKSUN'_{it} = WKS\hat{U}N_{it} - 1] \qquad (7)$$
$$= 1 - \Pr(CONS)$$
$$\Pr(CONS) = \hat{F}_5(X_{it}, WKSUN'_{it}, LM_t, UI'_t) + \varepsilon_{5i} \text{ model iterates}$$

$\hat{F}_1, \hat{F}_2, \hat{F}_3, \hat{F}_4, \hat{F}_5$	estimated structural relationships (probit, tobit, logit, Weibull, and logit respectively) using the 1986/87 Statistics Canada Labor Market Activities Survey
$\varepsilon_{1i}, \varepsilon_{2i}, \varepsilon_{3i}, \varepsilon_{4i}, \varepsilon_{5i}$	random error term from corresponding structural equation
	i = individual
	$t, t-1$ = period
X_{it}	personal characteristics
$WKSUN_{it}$	weeks unemployed
$WKSNLK_{it}$	weeks not in labor force
U_{it}	experience of some unemployment during year
$CONS$	constrained in obtaining additional weeks work
UI_t	parameters of UI system (UI'_t denotes changed UI parameters)
LM_t	local labor market condition
$WKSEMP_{it}$	weeks employed

Note:
[a] See Erksoy, Osberg and Phipps (1994a) for details.

Notes

We would like to thank the Social Sciences and Humanities Research Council of Canada, the Natural Sciences and Engineering Research Council of Canada, and Human Resource Development Canada for their financial support, and Audri Mukhopadhyay, Lynn Lethbridge, and Tom McGuire for their excellent work as research assistants on this project. Brian MacLean and an anonymous reviewer provided useful comments, but any errors remaining are our sole responsibility.

1 Borch goes on to argue that there may be a minority of risk lovers who would prefer a society with greater risks, and that this may explain some of the dissatisfaction with the welfare state.
2 In 1993, UI benefits of \$18.3 billion were paid to 3.4 million Canadians. Appendix A summarizes the key features of the 1971 and 1994 UI regimes.
3 The problem of self-selection has been recognized for a long time. Atkinson (1991, p. 117) quotes Winston Churchill speaking in the British House of Commons in 1911: "Voluntary schemes of unemployment insurance . . . have always failed because those men likely to be unemployed resorted to

them, and, consequently, there was a preponderance of bad risks . . . which must be fatal to the success of the scheme" (*Hansard*, 1911, vol. 26, col. 495).

4 Given that governments make decisions on macroeconomic policy which may imply an increase in unemployment (e.g. a restrictive monetary policy to combat inflation), it has often been argued that society as a whole should compensate those individuals who are affected. In the 1970s, Canadian UI legislation embodied a recognition of social responsibility for unemployment by requiring the government to pay from general revenues the costs of claims in excess of those associated with 4 percent unemployment (amended in 1976 to unemployment exceeding an eight-year moving average). In the 1980s, denial of macroeconomic responsibility produced legislative amendments which shifted UI funding entirely to premiums.

5 Due to the difference in claims experience, it is normal for insurance companies to charge young male drivers substantially higher premiums for automobile insurance, but although there are similar differences by age and gender in UI claims' experience, such discrimination in UI premiums or claims eligibility is socially unacceptable.

6 Similarly, a legislative change which deprived individuals of protection against financial loss due to home burglary by prohibiting the sale of such insurance could be expected to decrease the incidence of home burglary (as individuals purchased burglar alarms, etc. to decrease the risk of burglary), and would increase the net cash income of individuals, in aggregate (since premiums paid to burglary insurers exceed claims paid out by the amount of administration expenses and industry profits), *but* such a change would decrease the utility of all those who previously purchased insurance. (Figure 14.1 is really a generic diagram of the implications, in general, of higher or lower levels of insurance coverage.)

7 In Phipps (1990, 1991a, 1991b) individual utility is dependent on consumption of goods *plus* non-labor time. These estimates allow for the possibility that time spent unemployed was not a utility maximizing choice – that some unemployed individuals faced demand-side constraints (see also Osberg, 1988). Implicitly, writing utility as a function of income assumes that individuals are liquidity constrained (i.e. consumption equals income) (see table 14.3, p. 332).

8 In practice, annual hours of female labor supply in Canada are a positive function of the wage for low wage workers but a negative function of the wage (i.e. the "income effect" dominates the "substitution effect") for above average wages – see Osberg and Phipps (1993). The calculations reported in tables 14.1 and 14.2 use the parameter values of the Stone–Geary specification reported in Phipps (1991b, p. 202).

9 We also carried out the same analysis using expected incomes. Results were qualitatively very similar, as might be expected given the choice of the fairly restrictive Stone–Geary utility function.

10 Notice that to be affected by changes in UI, one must have been covered by UI in the first place. We conjecture that many women in the lower deciles, while having had positive earnings and UI in 1981, did not maintain eligibility for UI (which requires achievement of threshold weeks per year as well as hours or earnings per week) throughout the 1981–9 period.

11 Statistics Canada is now pre-testing the Survey of Labor and Income Dynamics, which will, when completed and available, provide a six-year panel of income and labor dynamics. The Labor Market Activity Survey of 1986/87 and 1988/90 is now available but contains no consumption information.

12 The Labor Force Survey (LFS) of 1984 recorded the average (uninterrupted) duration of an unemployment spell at 21.6 weeks, and the single person (1978 Base) Statistics Canada, Low Income Cutoff in 1981 dollars was \$7,268 per year – hence \$3,020 is the cash required to finance an average duration unemployment spell, at a poverty line standard of living.

13 A survey of local banks provided the following rule-of-thumb formula for second mortgage credit limits: credit limit = 75 percent of market value of house – balance outstanding on first mortgage.

14 Economic anxiety in Canada is, in fact, much more widespread than the issue of unemployment–unemployment insurance. In October 1993, 34 percent of Canadians agreed with the statement: "It is quite possible that I or someone in my household will seek benefits under the Unemployment Insurance Program over the next year or two." (Another 10 percent were uncertain.) A larger number (52 percent of respondents) agreed with the statement: "I feel I have lost all control over my economic future" (see EKOS Research Associates, 1993).

15 73 percent of Canadians with a household income less than \$15,000 disapproved of reducing the level of UI benefits, while only 53 percent of those with household incomes in excess of \$50,000 disapproved (see Environics Research Group, 1994).

References

Atkinson, A.B., 1991. "Social insurance: the 15th annual lecture of the Geneva Association," *The Geneva Papers on Risk and Insurance Theory*, 16, 113–32, Boston: Kluwer Academic

Atkinson, A.B. and J. Micklewright, 1991. "Unemployment compensation and labour market transitions: a critical review," *Journal of Economic Literature*, 29, 1679–1727

Betson, D. and J. van der Gaag, 1985. "Measuring the benefits of income maintenance programs," chapter 7 in M. David and T. Smeeding (eds.), *Horizontal Equity, Uncertainty and Economic Well-being*, NBER, *Studies in Income and Wealth, vol. 50*, Chicago: University of Chicago Press, 215–38

Bird, E.J., 1993. "An exploratory comparison of income risk in Germany and the United States," Department of Economics, University of Rochester, mimeo

Borch, K.H., 1968. *The Economics of Uncertainty*, Princeton, NJ: Princeton University Press

Card, D. and C.W. Riddell, 1993. "A comparative analysis of unemployment in Canada and the United States," in D. Card and R. Freeman (eds.), *Small Differences that Matter: Labor Market and Income Maintenance in the United States and Canada*, Chicago: University of Chicago Press, 149–89

Clark, A.E. and A.J. Oswald, 1994. "Unhappiness and unemployment," *Economic Journal*, 104, 648–59

EKOS Research Associates, 1993. "A mid-campaign report: the EKOS election 1993 analysis," Ottawa: EKOS Research Associates

Environics Research Group Limited, 1994. *The Focus Canada Report 1994–1*, Toronto: Environics Research Group Limited (March)

Erksoy, S., L. Osberg and S. Phipps, 1994a. "The distributional implications of unemployment insurance – a microsimulation analysis," report prepared for Human Resource Development and Labor, Department of Economics, Dalhousie University, mimeo

1994b. "Panel data and policy analysis," paper presented to the 1994 Meetings of the Canadian Economics Association, Calgary, Department of Economics, Dalhousie University, mimeo

1994c. "The distributional implications of unemployment insurance – revisions," paper presented to the 1994 Meetings of the Canadian Economics Association, Calgary, Department of Economics, Dalhousie University, mimeo

Friend, I. and M.E. Blume, 1975. "The demand for risky assets," *American Economic Review*, 65, 900–22

Grubel, H.G., D. Maki and S. Sax, 1975. "Real and insurance induced unemployment in Canada," *Canadian Journal of Economics*, 8, 174–91

Hansard (1911), *Parliamentary Debates*, 26, col. 995, London: Reuter's

Hayes, J. and P. Nutman, 1981. *Understanding the Unemployed: The Psychological Effects of Unemployment*, London and New York: Tavistock Publications

HRDC, 1996. "Employment insurance: impacts of reform," Human Resources Development Canada, 1, Ottawa (January)

Jahoda, M., 1979. "The psychological meanings of unemployment," *New Society*, 492–5

Kelvin, P. and J.E. Jarrett, 1985. *Unemployment: Its Social Psychological Effects*, Cambridge: Cambridge University Press

Myatt, T., 1996. "Why do we know so little about unemployment determination and the effects of unemployment insurance?" in L. Osberg and B. MacLean (eds.), *The Unemployment Crisis: All for Naught*, Kingston and Montreal: McGill and Queen's University Press, 107–28

Narendranathan, W. and S. Nickell, 1985. "Modeling the process of job search," *Journal of Econometrics*, 28, 29–49

Orcutt, G., 1957. "A new type of socio-economic system," *Review of Economics and Statistics*, 58, 773–97

Orcutt, G., H.J. Merz and H. Quinke (eds.), 1986. *Microanalytic Simulation Models to Support Social and Financial Policy*, Amsterdam: North-Holland

Osberg, L., 1988. "The 'disappearance' of involuntary unemployment," *Journal of Economic Issues*, 22, 707–27

Osberg, L. and S. Phipps, 1993. "Labour supply with quantity constraints: estimates from a large sample of Canadian workers," *Oxford Economic Papers*, 45, 269–91

Phipps, S., 1990. "Quantity constrained household responses to unemployment insurance reform," *Economic Journal*, 100, 124–40

1991a. "Behavioural response to UI reform in constrained and unconstrained models of labour supply," *Canadian Journal of Economics*, 14, 34–54

1991b. "Equity and efficiency consequences of unemployment insurance reform in Canada: the importance of sensitivity analyses," *Economica*, 58, 199–214

(1997). "What is the income cost of a child? Exact equivalence scales for Canadian two parent families," *Review of Economics and Statistics*, forthcoming

Phipps, S. and P. Burton, 1995. "Sharing within families: implications for the measurement of poverty," *Canadian Journal of Economics*, 28, 177–204

Rubinstein, M., 1977. "The strong care for the generalized logarithmic utility model as the premier model of financial markets," in H. Levy and M. Sarrot (eds.), *Financial Decision Making under Uncertainty*, New York: Academic Press, 11–65

Stern, N., 1986. "On the specification of labour supply functions," in R. Blundell and I. Walker (eds.), *Unemployment, Search and Labour Supply*, Cambridge: Cambridge University Press, 143–89