



WP 04_09

Giovanni Gallipoli

University of British Columbia, Canada
The Rimini Centre for Economic Analysis, Rimini, Italy

Laura Turner

University of British Columbia, Canada
The Rimini Centre for Economic Analysis, Rimini, Italy

HOUSEHOLD RESPONSES TO INDIVIDUAL SHOCKS: DISABILITY AND LABOR SUPPLY

Copyright belongs to the author. Small sections of the text, not exceeding three paragraphs, can be used provided proper acknowledgement is given.

The *Rimini Centre for Economic Analysis* (RCEA) was established in March 2007. RCEA is a private, non-profit organization dedicated to independent research in Applied and Theoretical Economics and related fields. RCEA organizes seminars and workshops, sponsors a general interest journal *The Review of Economic Analysis*, and organizes a biennial conference: *Small Open Economies in the Globalized World* (SOEGW). Scientific work contributed by the RCEA Scholars is published in the RCEA Working Papers series.

The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Rimini Centre for Economic Analysis.

Household Responses to Individual Shocks: Disability and Labor Supply

Giovanni Gallipoli, UBC

Laura Turner, UBC *

Draft: Jan. 27, 2008

Abstract

What are idiosyncratic shocks and how do people respond to them? This paper starts from the observation that idiosyncratic shocks are experienced at the individual level, but responses to shocks can encompass the whole household. Understanding and accurately modeling these responses is essential to the analysis of intra-household allocations, especially labor supply. Using longitudinal data from the Canadian Survey of Labour and Income Dynamics (SLID) we exploit information about disability and health status to develop a life-cycle framework which rationalizes observed responses of household members to idiosyncratic shocks. Two puzzling findings associated to disability onset motivate our work: (1) the almost complete absence of ‘added worker’ effects within households and, (2) the fact that single agents’ labor supply responses to disability shocks are larger and more persistent than those of married agents. We show that a first-pass, basic model of the household has predictions about dynamic labor supply responses which are at odds with these facts; despite such failure, we argue that these facts are consistent with optimal household behavior when we account for two simple mechanisms: the first mechanism relates to selection into and out of marriage, while the second hinges on insurance transfers taking place within households. We show that these mechanisms arise naturally when we allow for three features: a linkage between human capital accumulation and life-cycle labor supply, endogenous marriage contracts and the possibility of time transfers between partners. We also report evidence that the extended model with endogenous marriage contracts can fit divorce patterns observed in Canadian data, as well as correlations between disability prevalence and marital status, providing an ideal framework to study intra-household risk-sharing with limited commitment.

*We gratefully acknowledge financial support for this project from the Canadian Labour Market and Skills Researcher Network (CLSRN). We thank participants at seminars/workshops at UBC, CLSRN, NBER, EIEF, the Minneapolis Fed, IFS/UCL, Toronto, Edinburgh, Cambridge, Queen Mary for helpful discussions. John Kennan, Sevi Mora, Victor Rios-Rull, Jonathan Thomas and Aloysius Siow provided many insightful comments. All errors are ours.

Idiosyncratic risk is a common feature of both economic life and economic models. The measurement and effects of idiosyncratic risk have been the object of extensive research, see among others Abowd and Card (1989), Attanasio and Davis (1996), Storesletten, Telmer, and Yaron (2004), Meghir and Pistaferri (2006), Blundell and Preston (1998). The extent to which individuals can insure against idiosyncratic risk, and the way in which insurance can be provided, have also been the object of much interest.¹

In the incomplete markets literature idiosyncratic shocks are often modeled as persistent perturbations to the wage process: this approach simplifies the statistical analysis of income risk but it also presents several shortcomings. For one, it fails to distinguish between the direct effects of a shock and the effects that are due to the optimal responses of agents, understating the economic agency problem. It becomes difficult to analyze the relevant margins through which agents optimally respond to shocks when all idiosyncratic risk is characterized as an unexplained wage residual. This limitations of the standard approach become even more apparent when we look at the optimal behavior of households or couples: shocks are experienced at the individual level, however, the responses to, and consequences of, shocks are often determined at the household level: the significant effects that idiosyncratic shocks have on individual income might as well be the result of optimal household level decisions which are not directly observed. In this paper we provide an analysis of the optimal responses of households to idiosyncratic shocks and characterize the optimal intra-household allocations which are consistent with the observed changes in wages and labor supply associated to shocks. In order to do this, we exploit information from changes in agents' behavior induced by health and disability shocks.

A large economic literature on disability has focused on describing and quantifying the costs, in terms of income and labor time loss, experienced by households with a sick or disabled member (Meyer and Mok (2006), Charles (2003), Stephens (2001), Spector (2006)).²

In this paper we report some comparable evidence on the effects of disability in Canada: we then show how the observed responses of individuals and households are not consistent with a basic household decision problem. Finally we propose a household decision model that can account for single individuals' and couples' responses to disability shocks experienced by the household head. Our primary focus is on labor supply responses of heads and spouses over a multi-year window following disability onset. We examine how total labor supply, participation, and hours of work for workers vary by marital status, and by 'type' of disability

¹The literature on risk-sharing is very large, see for example, Kocherlakota (1996), Ligon, Thomas, and Worrall (2002), Attanasio and Ríos-Rull (2000), Pistaferri, Blundell, and Preston (2004), Attanasio and Pavoni (2007) and Heathcote, Violante, and Storesletten (2008).

²A related literature focuses on individual workers' responses to declining health or disability onset and to the incentives provided by disability insurance programs. Recent contributions based on American data include Burkhauser, Butler, and Gumus (2004), Kreider (1999), Autor and Duggan (2003), Low, Meghir, and Pistaferri (2008), while some of the major contributions based on the Canadian experience include Campolieti and Lavis (2000), Gruber (2000), Campolieti and Krashinsky (2006) and Bolduc, Fortin, Labrecque, and Lanoie (2002). Other studies have focused on the possible added worker effect and spousal responses to individual disability shocks. Coyle (2004), using the HRS, and Charles (1999), using the PSID, provide two recent examples.

shock - specifically, whether or not a reported disability directly affects an individual's ability to engage in market work. We also consider the potential role of disability in marriage formation and dissolution, which are observed, and for intra-household allocations, which are typically unobserved.

We design, estimate and numerically implement a model using observations from Canadian household longitudinal data available in the 1999 and 2002 panels of the Survey of Labor and Income Dynamics (SLID). The SLID resembles large U.S. panel household surveys such as the Panel Study of Income Dynamics (PSID) and the Health and Retirement Study (HRS). Relative to the PSID, which like the SLID follows households at all stages of the life-cycle, the SLID has three advantages and one disadvantage. The first advantage is that income data is subject to less measurement error since about 70% comes from tax records. A second advantage is that the SLID has a cross-sectional dimension about twice as large as that of the PSID. The third advantage is that the questions posed in the disability module are sufficiently detailed for us to classify subjective disability by type of limitation: specifically, disability can be latent, directly labor-limiting or directly leisure-limiting. These data offer a view into the economic effects of disability which is much finer than what can be typically gleaned from income-based panel studies. The one drawback of the SLID is its relatively short six-year panel dimension, which reduces our ability to follow households over a long horizon and requires an unbalanced panel approach.³

The layout of the paper is as follows. In section 1 we introduce a 'basic' life-cycle model of the household. In this 'basic' model, households are subject to a single shock to productivity that encompasses all idiosyncratic risk. Married (or common-law) households are comprised of two members who maximize a joint utility objective which is an explicit function of both members' preferences. From this simple model we draw analytical predictions about 'own' and 'spousal' labor supply responses to idiosyncratic disability shocks. The model provides three predictions: first, married men should in general experience larger swings in labor supply in response to disability shocks than single men, because of the insurance value of marriage; second, if disability shocks are unanticipated and large enough to affect the household's permanent income, we should observe a spousal labor supply response to disability onset in the household's head, a so-called 'added worker effect'; third, the effects of disability on labor supply should not outlast the disability itself and may even lead to higher labor supply in the long run. All these predictions are at odds with the data evidence we have gathered: this interesting puzzle is carefully documented in section 2, where we discuss our data sources and methodology.⁴ First, we observe generally larger swings in response to disability onset for singles than for married agents. Second, in keeping with recent findings

³Disability is self-reported and its severity is subjectively assessed. 'Justification' bias, where individuals with lower incomes or worse labor market prospects report more disability, is a problem with all subjective measures. However two recent studies by Au, Crossley, and Schellhorn (2005) (for health measures) and Campolieti (2002), both using Canadian data from the National Health and Population Study, find that justification bias is small enough to be of less concern than standard measurement error or attenuation bias.

⁴Detailed empirical results from the SLID are provided in a companion empirical paper, Gallipoli and Turner (2008).

using U.S. data, spousal labor supply responses are mostly negligible. Third, also in keeping with U.S. findings, the labor supply responses are quite persistent and in many cases average hours worked continue to fall well after the disability spell has ended. An additional puzzling feature of the data, that cannot be explained by a ‘basic’ household model, is that singles, both those who never marry and divorcees, report significantly higher incidence and chronicity of disability than do married individuals.

One potential explanation for the poor empirical performance of the ‘basic’ model is its lack of flexibility in modeling idiosyncratic risk: in Section 3 we show that simply adding an explicit disability shock process to the basic AR(1) model does very little to rationalize the observed dynamic responses we observe. However, in the same section we argue that three simple extensions to the ‘basic’ model can go a long way in reconciling theory and data. First, the wage process is dynamically linked to labor supply by allowing human capital accumulation through work experience: we use a framework similar to that of Shaw (1989) and Michelacci and Pijoan-Mas (2007), which introduces dynamic considerations linking disability onset, life-cycle labor supply and wages, and makes human capital an important vehicle of consumption smoothing over the life-cycle; moreover, we lift two additional restrictions of the ‘basic’ model: (1) we make it possible for agents to self-select into and out of marriage, and (2) we allow households to provide insurance through transfers of time, as well as consumption. Substantial data evidence appears to support the empirical relevance of both these mechanisms.

Our formulation of the marriage problem follows recent contributions to the dynamic theory of the household, specifically the idea that marriage contracts are subject to renegotiation whenever one spouse’s participation constraint becomes binding. The introduction of intra-marital transfers of time can be seen as a form of ‘home production’, which households allocate optimally to increase total labor market income and leisure. A by-product of optimal home production is the ability of spouses to ‘care’ for each other during periods of disability, through what are effectively intra-household transfers of time. Importantly, this additional channel of insurance interacts with the dynamic human capital accumulation motive to affect labor supply responses of both spouses to disability onset. We conclude the section by describing estimation and calibration of the numerical counterpart of the model.

Section 4 presents our results. We sequentially compare results from our basic/one-shock model, the two-shock extension of this basic model, and, finally, a two shock-model which allows for endogenous marital status and intra-household insurance through different channels (what we call the unrestricted, or full, model). Our unrestricted model does a remarkably good job of predicting intra-household dynamic labor supply responses to disability shocks across marital status and severity of the disability, over a ten-year window following disability onset, while the ‘basic’ model, as our analytical results suggest, performs very poorly. Additionally, the unrestricted model endogenously generates correlations between marital status and disability report rates that are extremely close to what we measure in the data. Finally, the unrestricted model can rationalize the puzzling lack of spousal added worker effect following a string of disability shocks to the household’s head, as well as reconcile this effect with the empirical evidence that spouses of disabled men increase the amount time spent caring

for their husbands. Based on these results we argue that our dynamic model of the household can be useful to evaluate the changing risk-sharing value of marriage over the life-cycle and, more generally, can be used to address questions about the extent of risk-sharing in a limited-commitment environment. We conclude the section, and the paper, by discussing some interesting properties of risk-sharing within marriage.

1 A Simple Model and Two Puzzles

In this section, we introduce a very simple household model and use it to derive some predictions about intra-household allocations in an economy in which all individual-level risk is summarized (and estimated) as idiosyncratic wage shocks. We argue that such predictions are counterfactual in at least two dimensions and that this simple model fails to capture both the timing and nature of married workers' dynamic responses to an important class of idiosyncratic shocks, like the ones that refer to disability and health.

1.1 A Basic Model of the Household

In what follows we introduce the simplest version of our model which has the following characteristics: matching and marriage are exogenous and no self-selection and changes in marital status are allowed. The lack of divorce implies that there is full-commitment between married agents. In this 'basic' model there is no home production: couples can only share consumption goods, but there is no technology to transfer time among them.

Individuals are indexed by their gender g . For simplicity, all couples have the same age (j). Each individual has current wage state w_g , which depends on a set of observable characteristics as well as on a random idiosyncratic shock, evolving over the life-cycle according to an AR(1) process which we later directly estimate from SLID wage data.

Household formation is (for practical purposes) exogenous: in the first period of life a fixed share of individuals are matched and all matches result in marriage. Matching is assortative by education (ed) - low (high) educated women are with low (high) educated men. On all other dimensions, matching is random. Given their respective state variables, agents choose a marriage contract (utility weight of the wife) λ which remains the same over the life-cycle as no renegotiation nor divorce is available.

Optimization occurs at both the individual and household levels. Individual and household-level optimization is the same thing for singles. In this restricted 'basic' model with no divorce, individual outcomes within marriage are fully determined by the solution of the household-level optimization problem under a given, and invariant, marriage contract λ , which is chosen through an individual maximization problem at the start of the couple's married life. Choice variables are leisure l , consumption c , carry-forward household assets a' and marital status. Individuals have gender-specific discount rates β_g ; and age-, gender- and health-specific survival probabilities $\varsigma_g(j, X_g)$, which we shorten to ς_g for simplicity. The set of model state variables is summarized in table 1.

Table 1: Basic Model State Variables

State variable	Description
Household level: fixed	
j	age
a	household asset holdings
ed	education of household members
λ	marriage contract (utility weight of the wife)
Individual level: fixed	
g	gender
Individual level f : variable	
w_f	wage of individual f
Individual level m : variable	
w_m	wage of individual m

We start by describing the dynamic problem for a single individual. He or she solves the following optimization problem at age j

$$V_{j,g}^S(X_g, a) = \max_{\{c,l,a'\}} u(c, l) + \beta \varsigma_g E_j[V_{j+1,g}^S(X'_g, a') | X_g] \quad (1)$$

For married individuals we first define a household-level value function U_j^H as follows

$$U_j^H(X_f, X_m, a; \lambda, \theta^M) = \max_{\{c_f, l_f, c_m, l_m, a'\}} (1 - \lambda)V_{j,m}^M + \lambda V_{j,f}^M \quad (2)$$

$U_j^H(\cdot)$ is the household value function for a husband and wife with state vectors X_f and X_m , marriage contract λ and exogenous, non-pecuniary gain (or cost) to being married, θ^M . The individual value of marriage, for partner g (with spouse ‘ $-g$ ’ and household-level vector of members’ characteristics $X = \{X_g, X_{-g}\}$) is

$$\begin{aligned} V_{j,g}^M(X, a) = u(c_g, l_g) &+ \varsigma_g \varsigma_{-g} \beta_g E[V_{j+1,g}^M(X', a'; \lambda', \theta^M) | X] \\ &+ \varsigma_g (1 - \varsigma_{-g}) \beta_g E[V_{j,g}^S(X'_g, a') | X_g] \end{aligned} \quad (3)$$

Notice that in equation (4) the individual value of marriage is imputed, rather than independently solved through a maximum operator, because the optimal levels of consumption, leisure and savings are jointly determined at the household level by solving the household planner’s maximization problem subject to contract λ . We initialize the utility-weighting parameter λ , given the couple’s age j and state $\{X_f, X_m, a_m, a_f\}$, through cooperative Nash bargaining by solving the following problem

$$\lambda^* = \arg \max_{\lambda} S(j, a_f, a_m, X_f, X_m; \lambda, \theta^M) \quad (4)$$

s.t.

$$S(\cdot) = [V_{j,m}^M(a_m + a_f, X; \lambda, \theta^M) - V_{j,m}^C(a_m, a_f, X)] [V_{j,f}^M(a_m + a_f, X; \lambda, \theta^M) - V_{j,m}^C(a_m, a_f, X)]$$

$$V_m^M(j, a_f + a_m, X; \lambda, \theta^M) - V_m^C(j, a_m, a_f, X) \geq 0$$

Here, $S(\cdot)$ is the product of partners' individual surplus from marriage, conditional on at least one partner's surplus (the male's) being positive. V^C (C for "courting") is the value of remaining single, but still able to marry for another period. In the simplest case (and in the basic model), courtship lasts only one period, so $V_{j,g}^C(a_g, a_{-g}, X) = V_{j,g}^S(a_g, X_g)$; i.e. if no marriage occurs right away, the couple splits. We define the courting problem in greater detail in section 3. Household-level assets at the time of marriage are equal to the sum of both members' privately accumulated assets.

The single endogenous aspect of marriage in the 'basic' model is bargaining over the initial marriage contract. Marriage occurs only if both members of the potential couple find it optimal to marry and commit to the allocations given by λ^* . In the basic model we set θ^M in such a way that individuals who are matched always prefer to marry. After this, solving the married couple's joint optimization problem becomes simply a matter of solving the household-level problem, $U_j^H(\cdot)$ for a given λ .

1.2 The Optimization Problem for Single and Married Households

A single individual of gender g and age j solves (1) subject to

$$\xi_1 : (\tilde{T}_m - l_m)w_m + (1+r)a + b(\cdot) = c_m + a' \quad (5)$$

$$\xi_2 : l < \tilde{T}_m \quad (6)$$

$$\xi_3 : a' \geq \underline{a} \quad (7)$$

$$\xi_4 : c \geq \underline{c} \quad (8)$$

and where the wage process is given by

$$w'_m = f(Y'_m) + \nu' \quad (9)$$

$$\nu' = \rho\nu + \epsilon \quad (10)$$

$$\epsilon \sim (0, \sigma_\epsilon^2) \quad (11)$$

The vector Y_m summarizes the set of deterministic states which determine wages. The function $b(\cdot)$ in the budget constraint captures all the benefits and entitlements to which the individual has access, conditional on his age, wealth, productivity and current participation.

A married couple solves (21), subject to the constraint set

$$\xi_1 : (\tilde{T}_m - l_m)w_m + (\tilde{T}_f - l_f)w_f + (1+r)a + b(\cdot) = c_f + c_m + a' \quad (12)$$

$$\xi_{2a} : l_m < \tilde{T}_m \quad (13)$$

$$\xi_{2b} : l_f < \tilde{T}_f \quad (14)$$

$$\xi_3 : a' \geq \underline{a} \quad (15)$$

$$\xi_4 : c_m + c_f \geq \underline{c} \quad (16)$$

where the wages of both agents follow independent AR(1) processes like the one described in equation (11). We assume that $u_j > 0$, $u_{jj} < 0$, and $u_j(0, k) > A$ for A very large, where $j, k \in \{c, l\}$, $k \neq j$. Similar assumptions hold for period utility $v(\cdot, \cdot)$. For analytical tractability we also focus on the class of preferences for which $u_{jk} = 0$ – that is, individual preferences are separable in consumption and leisure. For the results that follow we consider only interior solutions for labor and for debt holding: that is, solutions to the Lagrangian in which, from complementary slackness, $\xi_2 = \xi_3 = 0$. The first order conditions for own and spousal leisure are:

$$\begin{aligned} u_l &= \xi_1 w_m \\ v_l &= \xi_1 w_f \end{aligned} \tag{17}$$

From the literature on inter-temporal labor supply we know that household responses to disability shocks depend primarily on the degree to which these shocks are foreseeable – that is, the extent to which they are shocks to ξ_1 as opposed to movements along a deterministic life-cycle path. Of course, disability shocks are not likely to be perfectly foreseen by individuals or their families, so it is unlikely that $\frac{d\xi_1}{dw} = 0$. In this case, the uncompensated own (18) and spouse (19) elasticities μ to wage shocks are

$$\mu_{w_m}^{l_m} = +\frac{u_l}{u_{ll}l_m} + \frac{d\xi_1}{dw_m} \frac{w_m^2}{l_m u_{ll}} \geq 0 \quad \mu_{w_m}^{n_m} = -\frac{u_l}{u_{ll}n_m} - \frac{d\xi_1}{dw_m} \frac{w_m^2}{u_{ll}n_m} \leq 0 \tag{18}$$

$$\mu_{w_m}^{l_f} = +\frac{d\xi_1}{dw_m} \frac{w_m w_f}{\dot{\lambda} l_f v_{ll}} < 0 \quad \mu_{w_m}^{n_f} = -\mu_{w_m}^{l_f} \frac{l_f}{n_f} > 0 \tag{19}$$

where $\dot{\lambda} = \frac{\lambda}{1-\lambda}$ is the normalized utility weight of the second earner (with the first earner having normalized weight $\frac{\lambda}{\lambda} = 1$). The magnitude of $\frac{d\xi_1}{dw_m}$ depends on many factors, among them the extent to which the household expected the shock to w_m and the expected duration of this shock (chronicity). For a given level of ‘surprise’ contained in a disability shock and given expectations about its chronicity, the size of $\frac{d\xi_1}{dw_m}$ will depend on the disabled individual’s contribution to household’s resources. The higher the wife’s wages and optimal labor force activity, the smaller should $\frac{d\xi_1}{dw_m}$ be, and the smaller should be husbands’ and wives’ reductions in leisure (increases in labor) in response to disability onset. We therefore make the following predictions for the case in which disability shocks are not perfectly foreseen:

1. for a given level of per-capita resources within the household, married men should experience larger reductions (or smaller increases) in labor supply. In other words, they should give up less leisure than single men in response to both types of disability shocks. Furthermore, the larger the wife’s contribution to married households’ total resources, the more men should reduce labor supply in response to disability onset.
2. there is an ‘added worker effect’ which should be proportionally larger for wives with smaller contributions to total household resources;

3. changes in labor supply due to disability shocks may outlast the limitation itself because the household's lifetime wealth has changed. Individuals should provide relatively more labor in all states of health because less leisure is affordable.⁵

Results in the next section, and empirical evidence provided in Gallipoli and Turner (2008), show that none of the above mentioned predictions fit very well with what we observe in data. The elasticity of single men's labor supply to disability shocks appears quite a bit larger (more negative) than the elasticity of married men's labor supply. Added worker effects, especially for wives, are small and basically non-existent except for very chronic disability. While this might be consistent with well-forecast health shocks and separable preferences, the effects of disability are extremely persistent (by more than can be accounted for by the persistence of the disability itself) and imply large negative wealth effects: both these facts are inconsistent with the basic model with perfectly forecast shocks.

2 Data Sources and Empirical Findings

Our main data source for this study is the Canadian Survey of Labour and Income Dynamics (SLID), a longitudinal survey of Canadian Households maintained by Statistics Canada. For our purposes, the SLID has several advantages, and one noticeable disadvantage, relative to comparable U.S. income-based panel studies and in particular to the Panel Study of Income Dynamics (PSID), which also tracks households over the entire adult life-cycle. The first advantage is that the SLID contains relatively accurate income data since a majority of respondent households - about 70% in each given year - allow Statistics Canada to gather this information from their tax records. Second, the SLID has a cross-sectional dimension about twice as large as the PSID, measuring roughly 17,000 individuals. This large cross-section allows us to maintain reasonable sample sizes even when we disaggregate individuals in the sample by gender, marital status or other demographic dimensions. Finally, compared to other income panel studies, the SLID contains relatively detailed information about the type of disability and its direct consequences on economic life. The relative richness of these data allows us to develop a more sophisticated model of the effects of disability than would be possible using less-detailed work-limitation measures available in the PSID.

The major disadvantage of the SLID for our purposes is that it has a relatively short panel dimension of only six years. The six-year window over which households are observed reduces our ability to observe long-run effects of disability and forces us to use an unbalanced panel approach. We briefly discuss these complications, and our methods of dealing with them, below. More details are available in Gallipoli and Turner (2008).

⁵The predictions derived here refer to the case of separable preferences between own consumption and leisure. We can show that our predictions are unchanged for the case of preferences in which consumption and leisure are complements in individual utility.

Table 2: Incidence of reported disability by type in 2004

age	no disability	latent	<i>l</i> -limiting	<i>n</i> -limiting	<i>l</i> and <i>n</i> -limiting
19-44	85.5%	3.0%	2.6%	1.1%	7.7%
44-69	70.9%	5.8%	5.1%	1.5%	16.6%

2.1 Disability Measures in the SLID

Our measures of disability are constructed from responses to a series of disability-based questions in the 1999 and 2002 panels of the SLID. These are described in detail in Appendix 2. From this information, we classify disability into three broad types. Disability is denoted as ‘latent’ if it limits physical activity but does not directly limit activity at work, home or in other life activities. Disability is ‘work-limiting’ if limits the respondent at work or in other work - or human-capital based jobs such as school or job-search. It is worth noting that disability can have indirect effects on work even if no direct effect is present: for instance, some respondents report being limited in their ability to change jobs or to work their optimal number of hours even though they did not report a direct work-limitation in the current period. Finally, disability is ‘leisure-limiting’ if the individual reports being limited in home-based or in other types of non-work activities such as transportation or leisure. Table 2 reports the incidence of disability by type in 2004 for individuals aged 19-44 and 45-69 (the ages for which all types of limitation are reported).⁶ As can be seen, more than half of all reported disabilities in fact are both work- and leisure-limiting, while relatively few are exclusively work-limiting. As well, the incidence of all types of disability is quite high and increases sharply with age. Nearly 30% of working-age respondents over 45 report at least some dimension of disability.

2.2 Long-run Responses to Disability: Methodology and Sample Correction

Our methodology follows closely the approach of Meyer and Mok (2006), who in turn draw on work by Charles (2003) and Stephens (2001). The estimating equation is as follows

$$y_{it} = \alpha_i + \gamma_t + X_{it}\beta + \sum_h \sum_k \widehat{\delta}_k^h A_{kit}^h + e_{it} \quad (20)$$

where X contains observation-specific demographic and life-cycle information including a cubic in age; years of education; household size; number of children; a dummy for living in a city of at least 50,000; and regional dummies⁷. We control for individual or family fixed effects α_i through the inclusion of sample averages of the covariates, plus a measure of the affected individual’s average self-assessed health (as distinct from disability) over the sample.

⁶Incidence statistics are weighted using the cross-sectional weights provided by Statistics Canada so to represent the Canadian population as of 2004.

⁷The regions are the Atlantic provinces, Ontario, Western provinces and B.C., Quebec.

Index k , ranging from -1 to +10, denotes the number of periods from the initial disability onset, whereas h denotes the severity of disability, or more specifically, the specific sample of disabled people on which the estimation is run (see discussion below). A represents a dummy variable indicating the (h, k) combination under consideration. The coefficients on A , $\widehat{\delta}_k^h$, are what we plot in the graphs in the following subsections.

Our approach differs slightly from previous studies due to the relatively shorter panel dimension of the SLID, which requires some sample adjustment before (20) can be estimated consistently. In contrast to the PSID, we observe individuals in our sample for a maximum of six years. To assess a twelve-year time horizon around disability onset, we therefore rely on individuals' reports of the duration of their current disabling condition - the number of years they had the condition before reporting it in the SLID. The adjustments required to create samples comparable to those used for earlier U.S. studies are described in Appendix 2 and in Gallipoli and Turner (2008). Through this adjustment process, we create a series of samples of the disabled, two of which are considered in this paper: the "disabled" and "chronic disabled" samples. The mean chronicity within the "disabled" sample (the fraction of years, at or after onset, in which the individual reports a *current* latent or limiting disability) is 0.4 (about two periods out of six in the survey), which is the average chronicity of the aggregate post-onset SLID population. This 'chronic disabled' sample consists only of individuals who report being disabled at least 1 out of four periods (chronicity larger than 0.25), and the resulting mean chronicity is 0.62 or nearly four periods out of six.⁸

2.3 Labor Supply Responses to Disability

The primary objective of this paper is to understand and measure the labor supply responses of different types of households to idiosyncratic health shocks. In this section we also report results on the effects of disability onset on wages and family income and examine the relationship between disability reporting and marital status. Additional results on the long-term effects of disability on households' earnings, government and private support, and savings behavior, are presented in Gallipoli and Turner (2008). We focus only on the experience for men and the wives of men experiencing disability. Comparable results for women can be found in Gallipoli and Turner (2008).

Figure 1 reports our main results for men in, respectively, the 'disabled' and 'chronic disabled' samples described above. We plot the estimated $\widehat{\delta}_k^h$ s from (20) in levels, from one year prior to ten years after reported onset of the disability.⁹

The top panels show the changes in total annual hours worked, divided by 50 to give

⁸This procedure necessarily omits individuals reporting their first ever limitation after the final two years of the panel: for these individuals we lack sufficient post-onset information to determine chronicity.

⁹The omitted category (our 'control' group) is described in Appendix 2, are the omitted category. Nodes denote that the estimate is significant at the 95% confidence level. Note that most of the effects of disability typically begin to appear one year after onset, where "onset" is the year in which a current disability is first reported. This is because individuals interviewed in the SLID are asked about a *current* disability, but hours of work are reported for the *previous* year. In section 4, therefore, when comparing our simulated economies to data, we assume that disability shocks hit at time one rather than time zero.

an average weekly level. The middle panels show the changes in average weekly hours for workers (those who work positive hours during the year). The bottom panels show the average percentage point changes in participation, where (20) is estimated as a linear probability model with the binary dependent variable taking a value of 1 if the disabled individual worked positive hours during the previous year and zero otherwise. Figure 2 shows the long-run responses of the female spouse to a partner’s disability shock, along the total and extensive margins.¹⁰

Figures 1-2 do not provide much support for any of the major predictions of the ‘basic’ model described in section 1. We focus on the effects for men. To begin, single men experience larger and more persistent drops in labor supply than their married counterparts. On the total hours margin, the difference for both the disabled and chronic disabled groups is about 80% after ten years from onset, and close to 100% when accounting for the lower labor supply (about 6 hours per week) of the single males’ control group.¹¹ The difference is slightly larger, around 90%, or 110% in percentage terms, on the extensive margin. Gender-onset interaction terms (not shown), when added to the regressions, are significant at 5% in nearly all cases shown, except those for the intensive hours margin of the average-disabled sample.

Second, the labor supply effects of disability shocks appear to be very persistent. There is no evidence that individuals with long term disabilities recover or adapt enough to return to normal working, even for the sample that reports an average recurrence of only 45% of post-onset periods. This finding is consistent with previous U.S. studies, and especially with Meyer and Mok (2006) who also find very persistent effects of disability on labor supply among male household heads who report multiple years of limitation.

Third, we observe relatively small and belated added worker effects for the wives of disabled men. For the average-chronicity sample (results not reported) there is apparently no added worker effect at all, and the total hour effects for the chronic sample are typically insignificant even at 10%. The finding of a negligible added worker effect for wives is consistent with recent U.S. findings by Coyle (2004) using the HRS and Charles (1999) using the PSID.

2.3.1 Wage and Income Effects of Disability Onset

We next turn to the effects of disability onset on ln of hourly wages and on weekly family income, where the family unit is taken to be all relatives by blood or marriage living in a single residence.¹² We again report results for males, disaggregated by marital status. For the wage regressions, we use the panel data selection estimator proposed by Wooldridge (1995), which has the advantage of yielding consistent estimates in the presence of an individual fixed effect. Our selection restrictions are yearly provincial unemployment rate, and deviation of

¹⁰Tables corresponding to each figure will be made available in Gallipoli and Turner (2008) and online at the authors’ websites.

¹¹The estimated coefficients corresponding to year 10 are -3.09 for marrieds vs. -5.56 for singles in the disabled group, with the difference significant at 5%. For the chronic group the coefficients for year 10 are -5.45 for marrieds vs. -9.20 for singles, but the difference is actually larger in percentage terms in periods closer to onset.

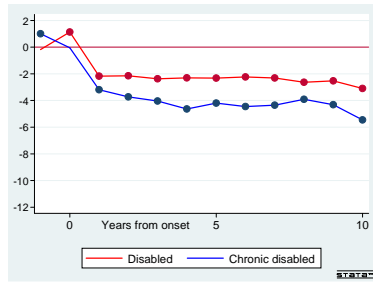
¹²This is consistent with Statistics Canada’s definition of the “economic family”.

this rate from its ten-year average; flow income from household wealth holdings; mother's and father's education; and the presence of pre-school age children in the house. We plot results for the first eight years after onset after which the sample sizes become quite small and the estimates (while still negative and significant) become very noisy.

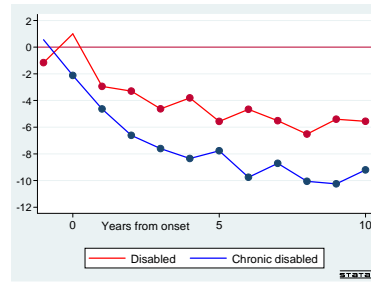
Figure 3 shows results of the effects of disability onset on wages. For married men, the long-run drop in average wages post-onset reaches 9% for the average disabled sample and 12% for the chronic sample, very close to the estimates reported by Meyer and Mok (2006) for PSID male household heads. For singles, the drop is correspondingly larger, dropping below 15% of initial wage for the average-disabled sample and 20% for the chronic disabled sample eight years after onset. This large, consistent drop in wages relative to similar-aged healthy individuals is consistent with the human-capital or learning-by-doing story of wage growth that we develop in subsequent section.

Finally, figure 4 predicts the average change in weekly family income, after public and private transfers are received, over the ten years following disability onset for SLID males, single and married. The top panel plots the drops in levels experienced by married and single men respectively. The bottom panel (dashed lines) show the drops in income as proportions of the average income of the future- and never-disabled control group. The story remains much the same. Single males experience larger drops in family income than married men once public and private transfers are received, and they experience even larger *proportional* drops in income, on the order of 13% vs. 5% for the larger (and more precisely estimated) "disabled" sample. Married men's larger initial family incomes are due in part to the fact that they tend to have much higher wages (by roughly six dollars per hour at the mean) and savings and in part because of their wives' earnings. The patterns are an empirical representation of the prediction of the "basic" model of the previous section that per-capita wealth effects of disability onset should be larger for single men than for married men.

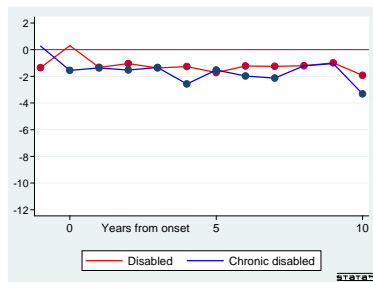
Figure 1: Labor supply responses to disability shocks by marital status: men



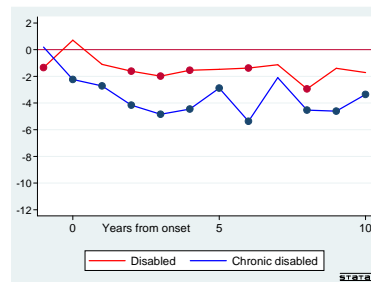
(a) Total hours: Married men



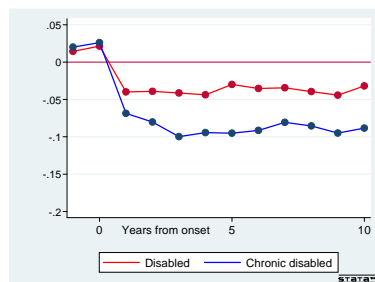
(b) Total hours: Single men



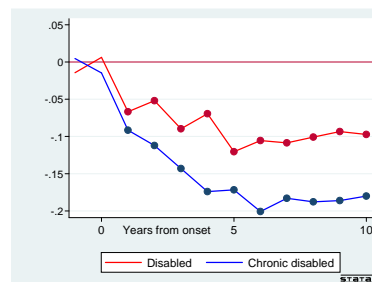
(c) Hours of workers: Married men



(d) Hours of workers: Single men

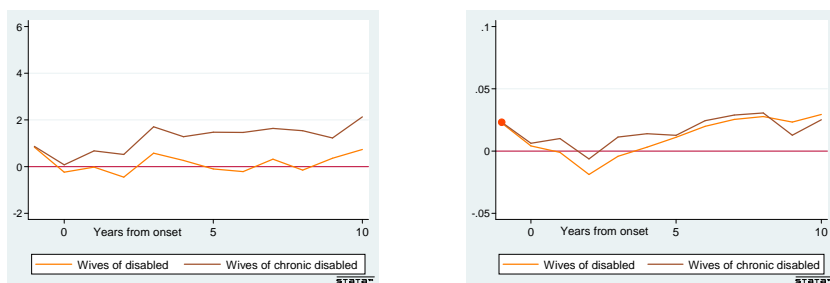


(e) Participation: Married Men



(f) Participation: Single men

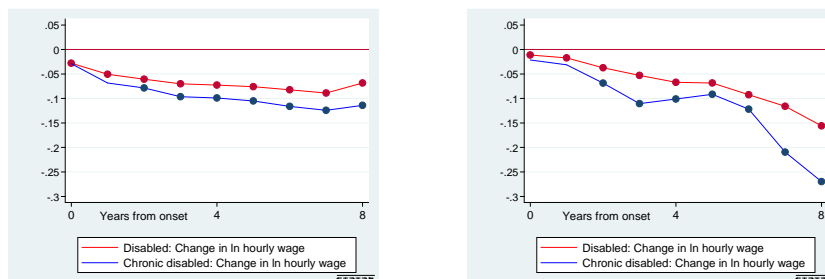
Figure 2: Spousal labor supply responses to disability shocks



(a) Total hours: wives of disabled men

(b) Participation: wives of disabled men

Figure 3: Wage effects of disability onset: men



(a) Hourly wage: married men

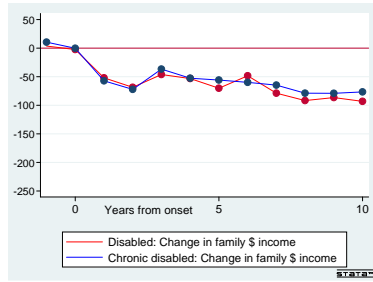
(b) Hourly wage: single men

2.3.2 Disability, Demographics and Marital Status

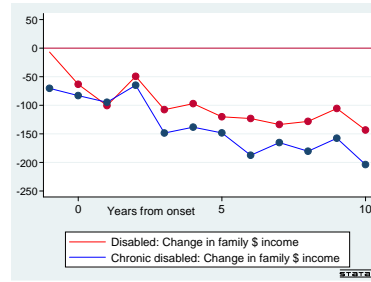
To conclude this section we turn to a brief examination of the incidence and chronicity of disability observed in the 1999 SLID panel, disaggregated by gender, education and marital status. Table 3 shows results, weighted using the 2004 cross-sectional weights provided by Statistics Canada. We report age-adjusted predicted probabilities from an ordered probit regression of the frequency of disability reports (from zero to six) on ten-year age category and dummies for gender (columns 1 and 2), education (columns 3 and 4) and marital status (columns 5-8) respectively. The predicted values are for individuals aged 40-49.

The table shows that women are slightly more prone than men - by about two percentage points for 40-49 year olds - to report a disability at least once. The estimated probit coefficient for the female dummy used to construct the predictions in table 3 is insignificant at 5%. Much

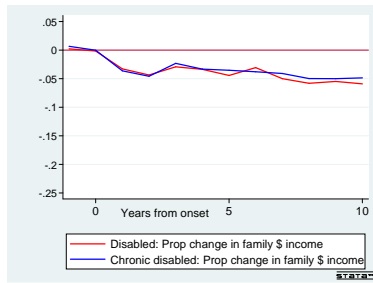
Figure 4: Disability shocks and family income



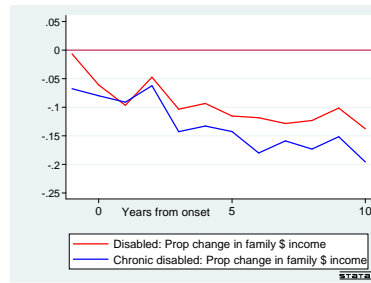
(a) Weekly family income: married men



(b) Weekly family income: single men



(c) Proportional drop in income: married men



(d) Proportional drop in income: single men

larger and more significant differences exist by education category. Low educated individuals are about 8.5 percentage points more likely than high-educated individuals to report at least one disability. In terms of chronicity, a predicted 13.6% of low educated individuals report a disability in four or more of the six years, while the corresponding number for high educated individuals is only 8.9%.

Not surprisingly, given the results in the previous subsection, the age-adjusted results also show strong differences by marital status. Married (M) individuals are 8.9 percentage points more likely to never report a disability and 4.5 percentage points less likely to report four or more times than never-marrieds (S). Of the three unmarried groups, divorcees and separated spouses (D) fare the worst in terms of both incidence and chronicity. Widows (W) fall between marrieds and singles. Wald tests of the estimated coefficients on the three unmarried categories from the ordered probits show that the difference in terms of chronicity between divorcees and widows is significant at 5%, but the difference between singles and divorcees is insignificant even at 10%. The negative effect of being currently married on the

Table 3: Age-adjusted observed and simulated disability by marital status: 40-49 year olds

Frequency of reports (z)	m (%)	f (%)	high ed (%)	low ed (%)	NM (%)	M (%)	S/D (%)	W (%)
0	0.617	0.597	0.553	0.648	0.546	0.634	0.501	0.583
1	0.151	0.155	0.163	0.145	0.164	0.148	0.169	0.158
2	0.076	0.08	0.087	0.071	0.088	0.074	0.095	0.082
3	0.051	0.054	0.06	0.046	0.062	0.048	0.068	0.056
4	0.037	0.039	0.045	0.033	0.046	0.034	0.053	0.041
5	0.035	0.038	0.045	0.03	0.047	0.032	0.055	0.04
6	0.033	0.037	0.046	0.027	0.048	0.029	0.06	0.039

number of disability reports is of course strongly significant, and is robust across different definitions of disability.

3 Generalizing the Model: Marriage, Selection and Insurance

We have established in section 1, and documented in section 2, that a basic household model generates wrong predictions about the dynamic labor supply responses to disability and health shocks. The questions we try to answer in this section are:

- Can we find a way to reconcile the model’s predictions with the basic stylized facts about disability?
- Are there any interesting mechanisms at work within the households which can be identified by exploiting information about dynamic responses to disability shocks?
- Is there a relationship between observed optimal responses to shocks and changes in the allocation of resources within households?
- To what extent can we support alternative mechanisms using micro data at the household level?

In what follows we sequentially discuss three extensions to the basic model that help us rationalize the labor supply puzzles described above by placing them in the context of a life cycle family insurance problem with limited commitment:

1. we replace the standard, single-shock model with a finer two-shock representation: individuals are subject to time-stealing disability shocks as well as (orthogonal) idiosyncratic wage shocks. The two types of shock are related (and separately identified) within a human capital function in which prior hours of work, current disability status, and current wage shock determine current and expected wage

2. we allow for intra-household transfers, not only in goods (income pooling) but also in time
3. we allow for fully endogenous marital formation and divorce, including the ability of members of a couple to renegotiate the power-sharing arrangement within the household

We describe each extension and its contribution separately, and whenever possible we discuss its empirical relevance.

3.1 Disability Shocks and Human Capital Dynamics

An open question in the disability literature relates to the actual effects disability and health shocks have on individuals. There are two alternative ways of modeling the effects of these shocks: in one case, health/disability shocks are interpreted as changes in marginal utility of leisure or consumption; another view is that disability and health shocks affect endowments (especially time endowments) and therefore should be accounted for in the individual budget constraint. Given that time endowments and time use are often measurable, we choose to model disability as an endowment shock.

This subsection also introduces a model of human capital accumulation which links life-cycle labor supply and wages: this approach is motivated by the evidence regarding wage dynamics after disability onset; moreover, it has the advantage of providing a simple way to identify separate sources of idiosyncratic risk (health versus general wage risk). Estimates of the human capital accumulation function from the SLID suggest that a relationship between labor supply patterns and human capital does in fact exist and does a reasonable job of capturing wage dynamics.

3.1.1 Modeling Disability

Disability shocks in our model arrive in two forms: labor-limiting (δ_n) and leisure-limiting (δ_l). Different combinations of $\{\delta_n, \delta_l\}$ constitute the different ds states, including the “healthy” state ds 1, in which $\delta_n = \delta_l = 1$. The δ shocks are modeled as multiplicative factors that ‘steal’ time from individuals by increasing the amount of total time required in order to, respectively, complete a given amount of market work or enjoy a given amount of leisure.

As described in detail in section 2 and Appendix 1, the disability module in the SLID allows us to distinguish between these two types of limitation (and between limiting disabilities and ‘latent’, non-limiting disabilities).¹³ We will see that the presence of both types of shock has implications for the life cycle insurance problem and for optimal responses.

¹³Burkhauser, Daly, Houtenville, and Nargis (2002), drawing on work by Nagi (1965, 1991), describe disability as an economic phenomenon. This is also the way we conceive disability in this paper. Disability is understood as the third stage of a three-stage process. The first stage is the “pathology”, an actual physical limitation such as blindness. The second stage is “impairment”, the ways in which the pathology limits the individual in everyday life, such as the inability to run or walk quickly. The third stage is “disability”: the limitations on economic life or productivity created by the impairment; e.g. blindness would be a severe disability for a policeman. Individuals who report ‘latent’ disabilities in our data are those who report one or more physical impairments but no corresponding direct limitation in work or leisure. As we observe in our

To reduce complexity we assume that only main-earners (males in the model¹⁴) are subject to health and disability shocks: therefore men have two additional individual state variables to those given in table 1: ds , which indexes current disability status, and $risk$, which captures underlying health, or equivalently a man's risk of becoming, and remaining, disabled. To clarify concepts, we restate the maximization problem for a married couple in the presence of disability shocks:

$$U_j^H(X_f, X_m, a; \lambda, \theta^M) = \max_{\{c_f, l_f, c_m, l_m, a'\}} (1 - \lambda)V_{j,m}^M + \lambda V_{j,f}^M \quad (21)$$

subject to the constraint set

$$\xi_1 : \frac{(\tilde{T}_m - \delta_l l_m)}{\delta_n} w_m + (\tilde{T}_f - l_f) w_f + (1 + r)a + b(\cdot) = c_f + c_m + a' \quad (22)$$

$$\xi_{2a} : l_m < \frac{\tilde{T}_m}{\delta_l} \quad (23)$$

$$\xi_{2b} : l_f < \tilde{T}_f \quad (24)$$

$$\xi_3 : a' \geq \underline{a} \quad (25)$$

Disability affects both the household period budget constraint (ξ_1) and (in contrast to a wage shock) also the husband or male head's feasible time constraint (ξ_{2a}). The effects are depicted graphically in figure 5 for the case of a δ_n shock. So long as an interior solution for leisure is optimal, a positive δ_n shock operates exactly like a negative wage shock in rotating the household budget constraint inward along the x (consumption) axis. (The effect of a δ_l shock would rotate the budget line inward along the y- (leisure) axis.) However, unlike a wage shock, a δ_n shock has the additional effect of rotating the feasible time allocation between labour and leisure inward along the y (labour) axis.

The first order conditions for own and spousal leisure would then be:

$$u_l = \xi_1 \frac{\delta_l w_m}{\delta_n} \quad (26)$$

$$v_l = \xi_1 w_f$$

and the uncompensated own (27) and spouse (28) elasticities μ are

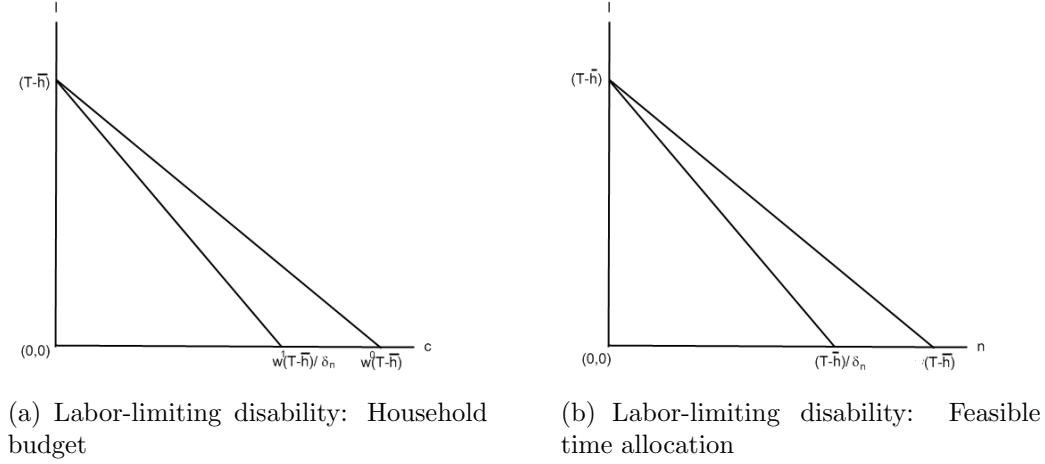
$$\mu_{\delta_n}^{l_m} = -\frac{u_l}{u_{ll} l_m} + \frac{d\xi_1}{d\delta_n} \frac{w_m \delta_l}{l_m u_{ll}} \geq 0 \quad \mu_{\delta_n}^{n_m} = -1 + \frac{u_l}{u_{ll} n_m} \frac{\delta_l}{\delta_n} - \frac{d\xi_1}{d\delta_n} \frac{\delta_l^2 w_m}{\delta_n u_{ll} n_m} \leq 0 \quad (27)$$

$$\mu_{\delta_l}^{l_m} = \frac{u_l}{u_{ll} l_m} + \frac{d\xi_1}{d\delta_l} \frac{w_m \delta_l^2}{l_m u_{ll} \delta_n} < 0 \quad \mu_{\delta_l}^{n_m} = \frac{\delta_l}{\delta_n n_m} \left(-\frac{u_l}{u_{ll}} - l_m - \frac{d\xi_1}{d\delta_l} \frac{w_m \delta_l^2}{u_{ll} \delta_n} \right) \geq 0$$

data, for those with chronic conditions, impairments can develop into disabilities if, for instance, an individual has to change employers or occupations. Similarly, chronic disabilities can be downgraded to impairments if the impaired individual substitutes toward alternative leisure activities for which his impairment is irrelevant. As well, impairments can alternate between being work-limiting or leisure-limiting based on changes in both the condition and other aspects of an individual's situation.

¹⁴Some recent evidence about the higher labor supply and wages of men can be found in Knowles (2007).

Figure 5: Disability and Household Optimality Constraints



$$\begin{aligned}
 \mu_{\delta_n}^{l_f} &= \frac{\xi_1}{\delta_l} \frac{w_f \delta_n}{\lambda l_f v_{ll}} < 0 & \mu_{\delta_n}^{n_f} &= -\mu_{\delta_n}^{l_f} \frac{l_f}{n_f} > 0 \\
 \mu_{\delta_l}^{l_f} &= \frac{\xi_1}{\delta_l} \frac{w_f \delta_l}{\lambda l_f v_{ll}} < 0 & \mu_{\delta_l}^{n_f} &= -\mu_{\delta_l}^{l_f} \frac{l_f}{n_f} > 0
 \end{aligned} \tag{28}$$

where $\hat{\lambda} = \frac{\lambda}{1-\lambda}$ is the normalized utility weight of the second earner (with the first earner having normalized weight $\frac{\lambda}{\lambda} = 1$). These elasticities are similar to the ones given in (18-19) and have similar implications for the two ‘‘puzzles’’ introduced in section 1 (smaller wealth effects and larger drops in labor supply for marrieds; an added worker effect proportional to the wealth effect of the shock). The key difference is the inclusion of a third term: in addition to the traditional substitution and income effects of idiosyncratic shocks, there is also an additional ‘‘time loss’’ effect, which limits the physical possibilities available to the individuals in terms of completed labour and leisure, as depicted in the second panel of figure 5). As we will see in section 4 below, this simple extension improves the fit of the model in terms of dynamic responses dampening the positive wealth effect on labor supply due to a shock, but it is not sufficient to solve the puzzles we documented in section 1.

3.1.2 Measuring Disability Shocks and Health Risk

We calibrate the process for disability as follows: during the working life (before age 66), disability follows a six-state Markov transition process that varies by age, sex and disability risk (*rsk*). The $ds=1$ is the healthy or, equivalently, non-disabled state whereas $ds=2$ is a state of latent disability in which the individual does not know if his condition will progress to a direct limitation. In $ds=3$ individuals are leisure-limited only and in $ds=4$, individuals are labor-limited only. In $ds=5$ or 6 individuals are both leisure and labor-limited, with the combined severity of the limitations greater in $ds=6$. The definitions of ‘latent’, ‘labor-limiting’ etc. are exactly as defined in sections 1 and 2. Each ds state is associated with

unique values of δ_n and δ_l , which we calibrate by using the average labor hours supplied by prime-age singles in each ds state.

Individuals older than 65 are subject to a separate six-state transition matrix based on the severity of leisure limitation. The leisure-limitation indicator is based on two limitation questions (limited ‘at home’ or limited ‘in other activities’, including leisure), each with a two-level answer (limited ‘sometimes’ or limited ‘often’). We construct an index from these responses running from zero to four, and allow an additional ds state for latent limitations, giving us six states. Because we do not observe labor supplies of individuals over 70, we simulate a special group of single men who transition to the retired transition matrix at 55, and match their simulated labor supplies to observed labor supplies of individuals 55 and over in the SLID who have the corresponding leisure limitation index value but no current labor limitation. The effect of this exercise is to generate a life-cycle profile of disability in which both the frequency and magnitude of leisure shocks increases as individuals get older.

For both processes, we estimate Markov transition matrices between ds states by a multi-variate dummy variable regression of ds on lagged ds . The matrices are estimated separately by gender, rsk and age, using the eleven-year age interval around the age in question. To calibrate the process for rsk , we run a probit regression in which the dependent variable is an indicator for having a disability during the course of the panel and the regressors are similar to those used in the propensity matching exercise described in Appendix 2, including age terms. We cut the data at the median so that half of the population is “high” rsk and half “low” rsk . We assume that individuals enter the model in rsk state 1. Over the life cycle, they face an increasing probability of switching permanently to rsk 2 which we take directly from our regression; by age 70, over 95% of the simulated population, and about 98% of the SLID population of men, is in rsk level 2. As well, a small number of individuals in the model enter a special high- rsk state 3, which is a state of permanent disability. Individuals enter rsk 3 only from ds 5 or 6. We calibrate the switching risk for rsk 3 in order to match equate the degree of chronicity of disability in the model with that observed in data.

3.1.3 Human Capital Accumulation and Wage Dynamics

The second component of our two-shock process comprises idiosyncratic shocks to human capital and wages. We propose a process for wages that assumes their evolution as a by-product of past labor market decisions. In particular, we assume that past levels of labor supply have an effect on the current stock of human capital: in effect this amounts to a specific type of ‘on-the-job-learning’. To assess the empirical relevance of this assumption we follow the strategy originally developed by Shaw (1989). In what follows, hours of work are denoted as n and stock of human capital as H . Human capital is paid at rate R_t in year t , where R is taken to be an aggregate wage rate that can vary over time. Individuals accumulate human capital through market work, and their existing stock at a point in time is subject to decay. Human capital also evolves according to an aggregate shock, ζ (captured through the inclusion of year dummies as regressors) and to individual i.i.d. shock v . Finally, education (as completed years of schooling) and disability status affect the human capital stock, both

directly and interactively with the existing stock of human capital and hours worked. The resulting process for wages, $\Theta(\ln(H_{it}), n_{it}, \delta_n, v_{it}, ed)$, is

$$\begin{aligned}
 w_{it} &= R_t H_{it} & (29) \\
 \ln(H_{it+1}) &= \beta_1 \ln(H_{it}) + \beta_2 \ln(H_{it})^2 + \beta_3 \ln(H_{it})n_{it} + \beta_4 n_{it} + \beta_5 n_{it}^2 + \alpha_1 I_{it+1}^{\delta_n > 1} + \zeta_{t+1} + v_{it+1} \\
 \beta_j &= \beta_{j1} + \beta_{j2} \times ed_{it} + \beta_{j3} \times ds_{it} \\
 v_{it} &\sim N(0, \sigma_\eta^{2,H})
 \end{aligned}$$

To Shaw’s basic framework we add (work-limiting) disability status and its interactions with current human capital stock as additional regressors. Our specification assumes the rental rate of human capital R to be constant across the years in our sample, which allows us to estimate the system in logs rather than levels by replacing (unobserved) \ln human capital with (observed) \ln wage rates in the human capital equation.¹⁵ We also allow for endogeneity of lagged hours and wage rates, including the possibility that v_{it} is auto-correlated, though we find evidence that it is not.¹⁶ Details of our estimation strategy are given in the first section of Appendix 2.

Our best specification for the human capital process, for both men and women, is reported in table 4, along with summary statistics that are generated in our simulated economy. The male and female human capital processes have the same functional form. The estimated and simulated variance of the shock is slightly larger for women at .050, compared to .048 for men. In our numerical model, women are not subject to disability shocks; however, we include work-limiting shocks as regressors in the empirical human capital equations for both genders. From our estimations, women in the SLID do not appear to suffer a direct human capital loss from their disability shocks: the coefficients on work-limiting disability and its interactions are insignificant.

For men, several additional results are worth noting. First, disability does appear to have a direct negative effect on male human capital (wages), and this effect is non-monotonic: figure 6 shows that the effect of work-limiting disability on human capital decreases then increases in the current human capital stock. Around the median wage of \$15/hr, these direct effects are quite small: just over 2% of the total wage. They are much higher for low-income workers and moderately higher for workers at the upper tail of the wage distribution.

Second, the effects of hours worked on human capital are increasing in the current stock of human capital, due to the inclusion of the interaction term $n \times \ln(H)$. Moreover, this effect is much stronger for men than for women. The effects of annual (weekly average) hours worked on men’s and women’s next-year’s human capital are shown on the left hand

¹⁵In this assumption, we follow Imai and Keane (2004) who posit a broadly similar model to assess the intertemporal substitution in labor when there is human capital. Since we are estimating human capital using only three years of data on the right hand side, this seems a relatively low-cost simplification.

¹⁶Even if v_{it} is not autocorrelated and we can control for individual fixed effects, the lag of annual hours worked is likely to be correlated with the error if, as it seems possible, hours worked in one year affect the human capital stock and wage rate *in that year*. The annual time dimension of the panel makes this type of endogeneity very plausible. For example, Miller and Sanders (1997) estimate a human capital equation with monthly lags and find that the contribution of past hours to current human capital diminishes very rapidly.

side of figure 7 at wages of \$12 and \$18 per hour. The direct returns to hours worked for women are negligible compared to those for men. Note that the figures do not account for the depreciation of existing human capital, which is higher for men than for women. However, they imply that men have more control than women over the course of their human capital and wages, especially at relatively low wage levels.

Third, the persistence of the human capital stock (that is, one minus the rate of decay) is increasing, in percentage terms, in the size of the of the existing human capital stock.¹⁷ The net effect is shown in the second panel of figure 7 for men with the median level of education in the economy. The figure shows the change in $\ln(H)$ for men with a high draw of v_{it} (at the 70th percentile of the v distribution). The solid line shows effects for individuals with no disability, working 40 hours a week, over a range of H comprising 95% of the male population.¹⁸ The dashed line shows effects for individuals with a work-limiting disability, working 30 hours a week, over the same range of H . The figure gives some idea of how human capital accumulation develops for men, and the role of disability shocks in its development: healthy men who receive positive v shocks can always increase their human capital stock by working full-time, and the rate of increase is higher as more human capital is accumulated. By contrast, low-wage traps in the model occur for the chronically disabled. Even when experiencing good draws of v_{it} , these individuals are not able to accumulate human capital beyond a wage of about \$12 an hour while working a number of weekly hours consistent with what we observe empirically for the work-limited population.

3.2 Intrahousehold Transfers of Time and ‘Caring’

The second component of our augmented model is the ability of spouses to engage in optimal time management through task-sharing¹⁹. Analytically, this amounts to a form of home production that allows the couple to reap economies of scale in total disposable time, similar to economies of scale in consumption typically assumed in household models. In the absence of disability, this feature of the model is simply a way of modeling home production as optimal time management. In the presence of disability, these time transfers allow the healthy spouse to alleviate the disutility experienced by the affected spouse by increasing the amount of

¹⁷Since many of the time-averages of the covariates are omitted from the reported and simulated human capital process, the constant term reported in the table is not the same as from the regression; it is adjusted to produce the mean wages in the male and female SLID sample in the simulation.

¹⁸As can be gleaned from figure 7(b), our quadratic specification for wages implies that returns to lagged human capital in terms of current human capital are monotonically increasing over the distribution of H . The effect is to generate an unreasonable process for human capital accumulation at high end of the wage distribution, where the SLID provides us with relatively few observations to achieve a fit. In this draft, we circumvent this problem by having wages plateau at values above \$40 per hour; that is, once an individual achieves a wage at or above this value, it remains there unless or until a series of negative shocks pushes it back down. This assumption affects about 2% of men and 1% of women in the simulation. The issue is revisited in forthcoming work.

¹⁹Our model is a variant on home production model originally introduced by Gronau (1973) and Gronau (1977). A very interesting treatment of home production versus labor market participation can be found in Rios-Rull (1993).

Figure 6: Effects of δ_{nt} on $\ln(H_{t+1})$

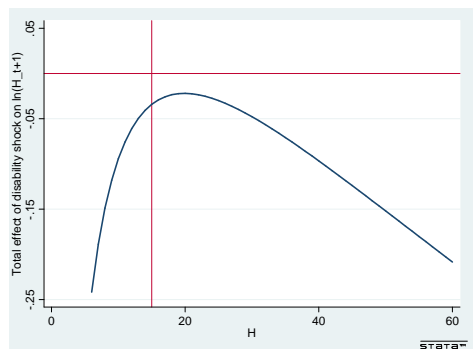
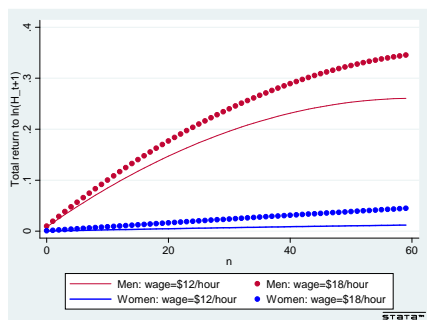
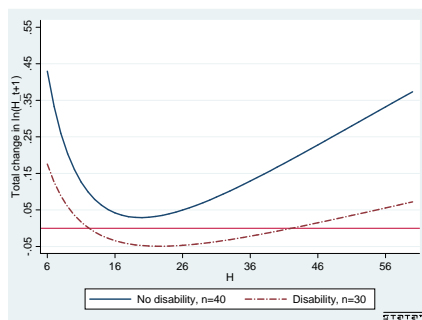


Figure 7: H_{t+1} : returns to n_t and H_t



(a) Returns to n by gender



(b) Returns to H

Table 4: Estimated productivity process parameters: Human capital specification

$\ln(H_{it+1})$	Male	Female
n_{it}	-0.00205 (.0030)	-0.00311 (.0030)
n_{it}^2	-0.0000684 (.000024)	-0.0000768 (.000027)
$\ln(H_{it})$	-1.03 (.362)	-0.714 (.319)
$\ln(H_{it})^2$.340 (.062)	.352 (.059)
$\ln(H_{it}) \times n_{it}$.00373 (.0013)	-0.000249 (.0013)
$I_{\delta_n > 1}$	-1.39 (.49)	-0.89 (.38)
$\ln(H_{it}) \times I_{\delta_n > 1}$.913 (.35)	.432 (.29)
$\ln(H_{it})^2 \times I_{\delta_n > 1}$	-0.153 (.061)	-0.289 (.54)
ed	.00565 (.034)	-0.00403 (.0287)
$\ln(H)_{it} \times ed$.0177 (.024)	.0303 (.021)
$\ln(H)_{it}^2 \times ed$	-0.0049 (.0040)	-0.00967 (.0038)
$n_{it} \times \ln(H)_{it} \times ed$	-0.0000205 (.000005)	.000105 (.0000414)
$cons$	2.54	1.86
σ_η^2	.048 (.003)	.050 (.003)
$Adj.R - sq$.625	.613
n	16040	16004
Summary statistics: human capital		
Mean H	\$17.75	\$15.25
Var $\ln H$.245	.265

leisure he can enjoy, a process we refer to as “caring”. Our estimation and calibration process for intrahousehold time management turns on the assumption that the technology for time transfer when one partner is disabled does not differ fundamentally from the technology when both partners are healthy.

The basic assumption underlying this section is that all individuals are endowed with a given number of waking hours in a week. Out of this natural endowment, every individual must devote some number of hours of non-labor, non-leisure tasks, denoted nll_i , and these tasks must be completed before labor and leisure decisions can be made. The time required for individual i to complete the tasks in his own nll bundle, in the absence of task-sharing, is \bar{h}_i . We assume that neither nll nor \bar{h} varies directly with marital status, but they do vary with age and with demographic factors correlated with marriage, including the number of children in the household. Total hours spent in all nll activities is denoted \bar{h}_i for individual i . Since intrahousehold time transfers are only feasible in multi-member (i.e. married) households, \tilde{h}_i is simply equal to \bar{h}_i for singles. Married partners, however, have the option of allocating the tasks in their respective nll task bundles optimally between them so that \tilde{h}_i differs in general from \bar{h}_i for marrieds.

We denote the spousal transfer of time for individual i as $h_i; i = \{m; f\}$, and the technology for time transfer is ϕ_i , with $\phi'_i \geq 0$ and $\phi''_i \leq 0$. When the wife devotes h_f hours to completing some amount of her husband’s nll tasks, she “returns” $\phi_f(h_f)$ hours to him, which he can then devote to labor and/or leisure. A similar relation, with technology $\phi_m(h_m)$ exists for the husband, capturing his ability to “return” time to his wife by completing tasks in her nll bundle. For the general case, it is assumed that each partner specializes in the tasks within couple’s combined nll activities in which he/she the greatest comparative advantage. As a result, it may be optimal for both partners to simultaneously “transfer” time to each other in order to increase the couple’s aggregate wealth, measured as post- nll time endowments multiplied by wages.

We denote the wage ratio (in levels) of the couple, $\frac{w_f}{w_m} = \bar{w}$. The couple’s complete problem can now be written:²⁰

²⁰Since the leisure-limitation questions in our data source include questions about the ability to perform tasks in the home, we assume that time transfers from the husband to the wife are also subject to time loss from δ_l . However, as can be seen, we also assume that the time transfers *received* from a spouse are *not* reduced by a factor δ_l . This assumption slightly strengthens the results reported in the subsequent sections but is not indispensable.

$$\begin{aligned}
V_j^M(X, a; \lambda, \theta^M) &= \max_{\{c_m, c_f, l_m, l_f, h_m, h_f, a'\}} (1 - \lambda)u(c_m, l_m) + \lambda v(c_f, l_f) + \beta E_j, \varsigma_f \varsigma_m [V_{j+1}^M(X', a')|X] + \\
&\quad \beta_f E_j \varsigma_f (1 - \varsigma_m) [V_{j+1, m}^S(X'_f, a')|X] + \beta_m E_j, \varsigma_m (1 - \varsigma_f) [V_{j+1, f}^S(X'_m, a')|X] \\
&\quad \text{s.t.} \\
\text{budget} &: (\tilde{T} - \delta_l l_m) \frac{w_m}{\delta_n} + (1 + r)a + b + (\tilde{T} - l_f)w_f - c_m - c_f = 0 \\
l_m &: l_m = \frac{\tilde{T}}{\delta_l} \geq 0 \\
l_f &: l_f = \tilde{T} \geq 0 \\
\tilde{T}_m &: \tilde{T}_m = T - \bar{h}_m - \delta_l h_m + \phi_f(h_f) \leq T \\
\tilde{T}_f &: \tilde{T}_f = T - \bar{h}_f - h_f + \phi_m(h_m) \leq T \\
H_f &: w'_f = \Theta_f(n_f, w_f) \\
H_m &: w'_m = \Theta_m(n_m, w_m)
\end{aligned} \tag{30}$$

At an interior solution, abstracting from human capital considerations, the first-order conditions for the h_i s are:

$$\begin{aligned}
\phi'_f(h_f) &= \delta_n \bar{w} \\
\phi'_m(h_m) &= \frac{\delta_l}{\bar{w} \delta_n}
\end{aligned} \tag{32}$$

Focusing on the model in which only husbands are subject to disability shocks and there is no human capital component to wages, it is easy to show that, so long as an interior solution for both members' labor supply continues to hold, $\frac{dh_f}{d\delta_n} \leq 0$ ($\frac{dh_f}{d\delta_n} < 0$ if ϕ_f is strictly concave) and $\frac{dh_f}{d\delta_l} = 0$. In a forthcoming extension, we show that these predictions fail to hold (1) when the husband's participation constraint becomes binding, in which case $\frac{dh_f}{d\delta_l} > 0$ and (2) in the presence of human capital accumulation, in which case the result that $\frac{dh_f}{d\delta_n} < 0$ is weakened, and in fact it is possible though unlikely that $\frac{dh_f}{d\delta_n} \geq 0$.

Next, we define Δh_i as the difference between total hours spent on *nll* activities and the amount of time the individual would spend on *nll* hours if single: $\tilde{h}_i - \bar{h}_i$, with $\Delta h_i = 0$ for singles. Then the equations for total amount of time devoted to (own and spousal) *nll* activities for husbands and wives respectively can be written:

$$\begin{aligned}
\Delta h_m &= h_m - \phi_f(h_f) \\
\Delta h_f &= h_f - \phi_m(h_m)
\end{aligned} \tag{33}$$

We require functional forms for the time-transfer technology that allow us to find closed-form solutions for h_m and h_f and that satisfy $\phi'_i \geq 0$ and $\phi''_i \leq 0$ if $i = \{m, f\}$. A simple specification that meets our needs is:

Table 5: Estimated time-transfer technology

α	.455 (.0117)
a	1.82 (.0630)
β	.001 (.0089)
b	.737 (629.5)
n	4057

$$\begin{aligned}\phi_m(h_m) &= bh_m^\beta \\ \phi_f(h_f) &= ah_f^\alpha\end{aligned}\tag{34}$$

with $a, b > 0$ and $\alpha, \beta \in (0, 1)$. We then combine (32) with (33) to obtain the following system:

$$\begin{aligned}\Delta h_m &= \left(\frac{a\alpha}{w}\right)^{\frac{1}{1-\alpha}} - b(b\beta\bar{w})^{\frac{\beta}{1-\beta}} \\ \Delta h_f &= (b\beta\bar{w})^{\frac{1}{1-\beta}} - a\left(\frac{a\alpha}{w}\right)^{\frac{\alpha}{1-\alpha}}\end{aligned}\tag{35}$$

In theory, this system can be estimated on data using information on wage rates and the *nll* activities for single and married men and women. Unfortunately, we are not aware of any Canadian data source that contains information both on *nll* activities and on wages. Instead, we use U.S. data taken from the 1990-2005 panels of the Panel Study of Income Dynamics (PSID) and the 2004 and 2005 panels of the American Time Use Study (ATUS). The PSID gives us some information on disability experienced by household members both in cross-section and over time. It also provides a measure of housework performed by heads of households and their spouses. The ATUS provides more detailed data on daily time use for a sample representing the current U.S. population. From this data, we are able to calculate the average share of *nll* activity devoted to housework in the ATUS and impute total *nll* activity to individuals in the PSID. Details on the estimation strategy are available in data Appendix 1. Details on exactly how we determine *nll* activity from the ATUS is available in Turner (2008). Table 5 gives our estimates. As can be seen, there is very little evidence that men make time transfers to their wives; but it does appear that women can and do make time transfers to their husbands.

An important question is how exactly *nll* time, and specifically Δh of one partner, changes when the other partner, in the context of our model the husband, receives disability shocks.

Within the model, a δ_n shock should lower Δh while a δ_l shock should have no effect (at interior solutions), unless human capital considerations change these results. In section 4, we reports some evidence that intrahousehold time transfers from husband to wife increase when the husband experiences disability due to the wife’s desire to keep him in the labor force, and that our model captures these effects.

3.3 Household Formation and Dissolution

In the unrestricted framework households can be of four types: courting (C), single (S), divorced (D), or married (M). Time-varying individual-level state variables are again summarized by vector $X_m = \{w_m, ds, rsk\}$ for men and $X_f = \{w_f\}$ for women, and define a couple-level state space is $X = \{X_m, X_f\}$ for married and courting households. Household formation evolves in the following way: in the first period of life individuals are matched exogenously with a member of the opposite sex. Like in the basic model, matching is assortative by education but random on all other dimensions. Couples who are matched but not yet married (or cohabiting) are defined as ‘courting’. Courtship lasts for a maximum of $N = 6$ (compared to $N = 1$ for the basic model) years. If at the end of this period the members of the couple do not get married, they split and remain single for life. Once married, couples remain together until one (or both) prefers to exit the marriage, transiting into the single state (divorce) or dies. Once a couple splits, both partners remain single for the remainder of their lives.²¹

The evolution of assets in the model depends on marital states: courting couples who choose to marry combine their assets so that $a = a_m + a_f$ in the first period of marriage. When couples divorce, assets, net of divorce costs, are divided between the two single households according to $a_g^D = a_g^D(a)$, where a is the married household’s stock of assets. Finally, when a spouse is widowed, he or she inherits all the household’s assets: $a_g^W = a_g^W(a) = a$.

The ability to commit to a fixed marriage contract (as in our ‘basic’ model) is important because commitment allows the household to make ex-ante efficient decisions that are not necessarily ex-post time-consistent for both partners under all realizations of uncertainty. Removing commitment has two effects: first, by adding a set of marriage participation constraints to the household optimization problem, it effectively reduces the feasible choice set relative to the full-commitment case. Second, it reduces the insurance value of marriage because partners who receive a series of bad shocks may be divorced or see their power over household decisions reduced.

Below we sketch out the basic elements of household formation and composition. Our choice of equal-weighted Nash bargaining over the marriage surplus, already described in section 1 is a convenient simplification in the dynamic collective literature and was originally proposed by Manser and Brown (1980). The renegotiation and divorce mechanism described here is drawn from Mazzocco and Yamaguchi (2007). The remaining aspects of the model are our own innovations.

²¹Mazzocco (2004) refers to models of this kind as dynamic collective models without commitment. A clear categorization of this model is difficult, as discussed by Browning, Chiappori, and Lechene (2006).

After a couple is married and choose their optimal starting λ , (which we denote λ^* here), they continue in married life under this contract until one partner finds it suboptimal to continue in the marriage. The conditions for a sustainable marriage are

$$V_{j,m}^M(a, X; \lambda^*, \theta^M) \geq V_{j,m}^D(a_{j,m}^D, X_m; \theta^D) \quad (36)$$

$$V_{j,f}^M(a, X; \lambda^*, \theta^M) \geq V_{j,f}^D(a_{j,f}^D, X_f; \theta^D) \quad (37)$$

where V^D is the value function associated with divorce and θ^D is an exogenous non-pecuniary utility weight associated with the divorced state. There is a one-time financial cost of divorce, borne equally by both partners, of $d_c(ed)$ which differs by the education level (ed) of the husband. We set $d_c(l)=7,500$ for low-educated husbands and $d_c(h)=12,500$ for high-educated husbands, based roughly on the average cost of divorce in Canada in 2000 and the average increase of the cost of divorce with household income. Divorce settlements divide existing household assets, net of d_c , equally between the partners, so that $a_m^D = a_f^D = \frac{a-d_c(ed)}{2}$. This is a standard assumption in the collective literature given current divorce law in Canada and the U.S. – see footnote 2 in Mazzocco and Yamaguchi (2007) for a discussion.

Not all failed marital contracts need end in divorce. It is possible that when the marriage participation constraint binds under the current λ^* , a λ'^* can be ‘renegotiated’ that induces both partners to stay in the marriage and thereby produces a Pareto improvement relative to the no-renegotiation benchmark. In choosing this λ'^* , we directly follow Kocherlakota (1996) and Ligon, Thomas, and Worrall (2002), who show that the most (ex-ante) efficient re-bargaining process (i.e. the one minimizing the insurance loss from non-commitment) selects $\hat{\lambda}$ to be the λ closest to λ^* for which both partners’ participation constraints once more hold with inequality. Formally:

$$\hat{\lambda} = \arg \min_{\lambda} |\lambda - \lambda^*| \quad (38)$$

$$\begin{aligned} s.t. \quad & V_{j,m}^M(a, X; \hat{\lambda}, \theta^M) \geq V_{j,m}^D(a_m^D, X_m; \theta^D) \\ & V_{j,f}^M(a, X; \hat{\lambda}, \theta^M) \geq V_{j,f}^D(a_f^D, X_f; \theta^D) \end{aligned}$$

If a solution to (38) does not exist, then the couple divorces. Couples can renegotiate and/or divorce at any point during the life-cycle. For reasons of tractability, we do not allow remarriage for divorcees or widows. However, even without remarriage, we will see that the ability to renegotiate and exit marriage has large effects on intrahousehold responses to shocks.

Given this structure, we can finally define the individual value functions in detail. In order to make exposition of the maximization problems easier, we define a ‘generalized’ value function denoted as $\Upsilon_{j,g}$. By ‘generalized’ we mean that $\Upsilon_{j,g}$ subsumes the value when married, single, courting or divorced, and can take different numbers of arguments depending on g ’s current marital status. We start by describing the dynamic problem for the single/divorced

individual. He or she solves the following optimization problem at age j

$$\begin{aligned} V_{j,g}^S(X_g, a) &= \max_{\{c,l,a'\}} u(c, l) + \beta_g \varsigma_g E_j[\Upsilon_{j+1,g}(X'_g, a')|X_g] \\ \Upsilon_{j,g} &= V_{j,g}^S \quad \forall j \end{aligned} \quad (39)$$

The divorcee's value function is nearly identical, except for exogenous utility weight θ^D (θ^S is normalized to 1):

$$\begin{aligned} V_{j,g}^D(X_g, a, \theta^D) &= \max_{\{c,l,a'\}} u(c, l) + \beta_g \varsigma_g E_j[\Upsilon_{j+1,g}(X'_g, a')|X_g] \\ \Upsilon_{j,g} &= V_{j,g}^D \quad \forall j \end{aligned} \quad (40)$$

For married individuals, the household-level value function U_j^H is still:

$$U_j^H(X, a; \lambda, \theta^M) = \max_{\{c_f, l_f, c_m, l_m, a'\}} (1 - \lambda)V_{j,m}^M + \lambda V_{j,f}^M \quad (41)$$

We now also define an individual value function, for married partner g (with spouse ' $-g$ ')

$$\begin{aligned} V_{j,g}^M(X, a; \lambda, \theta^M) &= u(c_g, l_g) + \varsigma_g \varsigma_{-g} \beta_g E[\Upsilon_{j+1,g}(X', a', \lambda')|X] \\ &+ \varsigma_g (1 - \varsigma_{-g}) \beta_g E[V_{j+1,g}^W(X', a')|X] \end{aligned} \quad (42)$$

where, $\forall j$:

$$\begin{aligned} \Upsilon_{j,g}(X, a, \lambda) &= \\ = \begin{cases} \max \{V_{j,g}^M(X, a; \hat{\lambda}, \theta^M), V_{j,g}^D(X_g, a_g^D(a); \theta^D)\} & \text{if } V_{j,-g}^M(X, a; \hat{\lambda}, \theta^M) \geq V_{j,-g}^D(X_{-g}, a_{-g}^D(a); \theta^D) \\ V_{j,g}^D(X, a^D(a); \theta^D) & \text{if } V_{j,-g}^M(X, a; \hat{\lambda}, \theta^M) < V_{j,-g}^D(X_{-g}, a_{-g}^D(a); \theta^D) \end{cases} \end{aligned} \quad (43)$$

The previous conditions state that an individual's generalized value function during marriage, conditional on renegotiation of λ , is always equal to the greater between value of being single and value of remaining married, so long as his/her partner also wants to remain married. Because both partners must commit to continuing the marriage, if either spouse prefers divorce, then divorce occurs and both members' value functions are set to $V_{j,g}^D$.

Finally, we define the following value function for a courting individual (i.e. an individual who will not marry this period):

$$V_{j,g}^C(X, a_g, a_{-g}) = \max_{\{c,l,a'\}} u(c, l) + \beta_g \varsigma_g E_j[\Upsilon_{j+1,g}(X', a', \lambda^*)|X] \quad (44)$$

where

$$\begin{aligned} \Upsilon_{j,g}(X, a, \lambda^*) &= \\ = \begin{cases} \max \{V_{j,g}^M(X, a, \lambda^*), V_{j,g}^C(X, a_g, a_{-g})\} & \text{if } V_{j,-g}^M(X, a, \lambda^*) \geq V_{j,-g}^C(X, a_g, a_{-g}) \\ V_{j,g}^C(X, a_g, a_{-g}) & \text{if } V_{j,-g}^M(X, a, \lambda^*) < V_{j,-g}^C(X, a_g, a_{-g}) \end{cases} \end{aligned} \quad (45)$$

where $a = a_m + a_f$. This problem resembles the one for married agents insofar both members of the courting couple must wish to marry before the marriage can take place, and each member’s ability to optimally choose his marital status is dependent on the partner’s willingness to marry.

Note that, in the N th year of courtship $V_{j+1,g}^C = V_{j+1,g}^S$ for $g = \{f, m\}$. As described above, following the initial determination of λ^* for the couple, future values of λ evolve according to:

$$\hat{\lambda}' = \Lambda \left(\lambda^*, X'_f, X'_m, a' \right) \quad (46)$$

3.4 Numerical Counterpart of the Model

In the remainder of this section, we outline additional salient features of our numerical model. These features apply to both the ‘basic’ and ‘augmented’ versions of the model.

3.4.1 Policy Environment and State Space

The simulated economy is populated by one-member (single) and two-member (married) households, consisting of individuals who differ permanently by education (ed) and gender (g). The state space of individuals is $X_i = \{H, ds, rsk\}$, where H is the individual’s current stock of human capital (or current wage), ds is his or her disability state, rsk is the individual’s disability risk – that is, the set of probabilities associated with moving between disability states in future periods – and $i \in \{m, f\}$. For married households, $X_M = \{H_f, H_m, ds_m, rsk_m\}$ due to our simplifying assumption that only males are subject to disability risk. Households, whether one-member or two-member, differ additionally by the age of their head, j (which is also the age of the spouse where applicable), and by their asset holdings, denoted a . For all individuals, the maximum lifespan is 90 years, with age-, gender-, disability-conditional survival probabilities $\zeta_j^{g,ds}$ taken from Canadian vital statistics.²² Men and women in the model may work until age 65. All values in the model are expressed in 2002 Canadian dollars. A model period is one year, but all flow values and results are expressed as if individuals lived for one representative week out of the year.

Policy in the model is based roughly on the existing patchwork of Canadian disability, retirement and anti-poverty programs. Individuals older than 65 receive a benefit of \$115 per week, which approximates the universal Canadian Old Age Security Benefit. Individuals 60 or older additionally receive a retirement benefit in any year they do not work (automatically after 65). The retirement benefit equals 25% of the potential full-time earnings of the individual based on his wage at age 60. Disabled individuals are eligible for a benefit equal to 19% of their current accumulated retirement benefit (based on their current wage) plus a lump sum of \$95, also conditional on nonparticipation. Benefit determination and eligibility in the model approximate the rules governing the retirement and disability benefits available to individuals under the Canada/Quebec Pension Plan (CQPP).²³ As with both CQPP

²²We thank Kevin Milligan for providing us with the Canadian mortality data.

²³To simplify the computations, the ‘‘CPP’’ benefit individuals receive is a function of their human capital

and workers' compensation, the disability benefit is only available to individuals in disability states with a 'qualifying severity' of δ_n , specifically $ds=6$ or who are chronically disabled (*rsk* state 3).²⁴ To fund retirement and disability benefits, individuals face a payroll tax of 9.9% (the Canadian CPP payroll tax rate). There is also a progressive income tax with brackets of {21.2; 31.8; 42.1; 46.4} on income above {\$9,600; \$47,485; \$84,320; \$132,784}, which are approximately equal to 2008 rates and brackets for the median Canadian taxpayer.²⁵ Tax revenues not spent on social security are wasted. Since Canada, unlike the U.S., has a universal set of provincial means-tested welfare programs, individuals may receive welfare payments conditional on: (i) holding less than \$60 (annual equivalent \$3000) in household assets and, (ii), all household members not working, regardless of disability status or age.

Our model is a partial equilibrium economy with no aggregate shocks and no exogenous growth. The interest rate is fixed at 4.2% and the market wage rate is normalized to 1.

3.4.2 Preferences

Period utility of men and women are:

$$\begin{aligned}
 u_{m,t} &= \sum_{ms=\{s,m,d\}} I(ms = ms) \theta(ms) \left[\gamma_c^m \frac{\left(\frac{c_t}{\tilde{n}_t^{j,ed,ms}}\right)^{1-\omega^m}}{1-\omega^m} + \gamma_l^m \frac{l_t^{1-\psi^m}}{1-\psi^m} \right] \\
 u_{f,t} &= \sum_{ms=\{s,m,d\}} I(ms = ms) \theta(ms) \left[\gamma_c^f \frac{\left(\frac{c_t}{\tilde{n}_t^{j,ed,ms}}\right)^{1-\omega^f}}{1-\omega^f} + \gamma_l^f \frac{l_t^{1-\psi^f}}{1-\psi^f} \right]
 \end{aligned} \tag{47}$$

where $l_t = \frac{\tilde{T}_t^{g,j,ms} - \delta_n n_t}{\delta_l}$; \tilde{T} is post-*nll* disposable time; n (again) denotes hours of market work; $\frac{1}{\omega}$ is the intertemporal elasticity of substitution in consumption; $\frac{1}{\psi}$ is the intertemporal (Frisch or marginal-utility-of-wealth-compensated) elasticity of substitution in leisure; \tilde{n} is a age-education-marital status-specific consumption weight; γ_c, γ_l capture the relative preference weightings of consumption and of leisure in overall period utility; and $\theta(ms), ms \in \{s, m, d\}$ are the multiplicative scale factors capturing the relative enjoyment of consumption and leisure in each marital state.²⁶

at age 60 rather than their eligible earnings over the lifetime. Since the stock of human capital in the model is fairly constant over time, we believe this is a reasonable approximation that involves relatively little computational cost. This study began with a detailed study of Canadian disability policy since 2000 which is available from the authors on request.

²⁴It could be argued that eligibility rules for the disability benefits actually approximate more closely those that govern provincial workers' compensation (WC) programs, which explicitly allow payments for temporary disabilities. WC programs are more generous but less universal than CQPP and usually require and/or directly encourage an eventual return to work unless disability can be shown to be permanent.

²⁵The median taxpayer by rates on income other than capital gains and dividends lives in Ontario. We calculate these rates based on the federal standard exemption and the smoothed combined provincial and federal rates.

²⁶Widows are assigned $\theta(d)$.

We set $\omega^f = \omega^m = 2.0$, which are in the middle of the range of inter-temporal elasticities found in the empirical consumption literature. The γ s and ψ s are calibrated to match, respectively, the average labor supplies of prime-age non-disabled men and women (those in $ds=1$ – see below) and the average Frisch or (compensated) inter-temporal elasticity of labor supply substitution of prime-age men and women. For the latter, we use target estimates from the empirical life-cycle literature on labor supply using a MacCurdy (1985) life-cycle framework with uncertainty. Given our specification of preferences, the time-varying Frisch elasticity of labor ϵ_t^n is $(\frac{L_t}{n_t} \frac{1}{\psi})$, which we calculate by averaging over the relevant simulated sample of single males and single females between 25 and 54.

The \tilde{n} consumption factors are calculated from information from the SLID and the traditional OECD equivalence scale, assigning a weight of 1 to the first adult in a household, .7 to every additional adult and 0.5 to every child. Conditional on the initial choice of marital status (whether or not to marry), these weights are fixed. Consumption by married individuals is therefore increased by a factor of $(1/0.85)$ in all childless years and by a factor greater than one in years in which children are present. Married couples have two children, the first at age 26 and the second at age 32, and both children are supported for 18 years. Single men and women have no children. Divorcees share the consumption costs of their children. If divorce occurs at or before the dates of conception of one or both children, then the child is or children is not born; divorcees support only the children who are already alive when the divorce occurs. Disposable time, \tilde{T} – before intrahousehold time transfers take place – is assigned using information from the regressions described in 3.2. The baseline *nll* costs of children within a marriage are the values that single and divorced parents devote to their offspring. Within married couples, of course, they values may be different due to optimal choices over intracouple time allocation.

Finally, individuals discount the future at rate β_g , which varies by gender. From Statistics Canada summary data, we observe that the asset to income ratio is about 2.2 for married couples and that households with male members hold roughly 1.5 times as many assets as households with female members. We use these figures as targets for our β_g s.

3.4.3 Calibration

From the above discussion, we have a total of fourteen parameters to calibrate: $\{\gamma_g\}^{27}$, $\{\delta_n(k)\}$, $\{\delta_l(3)\}$; $\{\psi^g\}$, $\{\beta_g\}$, $\{a, \alpha\}$ and $\{\theta(m), \theta(d)\}$ for : $g \in \{m, f\}$, $k = \{4, 5, 6\}$. The corresponding targets are the mean weekly labor hours of prime-age single women, and of prime-age single men in each *ds* state; male and female Frisch-compensated elasticities of intertemporal labor substitution; the household wealth-to-income ratio; the relative stocks of assets held by single men and women; the responses of married women to changes in the household wage ratio and the distribution of *nll* activities among the married female population; and the (Canadian) national marriage and divorce rates.

The Frisch targets are set to 0.5 for men, following Domeij and Floden (2005) who use real and simulated data to argue that this is the appropriate value for men in the absence

²⁷Although there are two γ s per gender, γ_l for women and γ_c for men are normalized to 1.

of credit constraints. There is some agreement in the literature that Frisch elasticity for females is around twice as large as that for male. We therefore set our female Frisch target to 1.0. These empirical estimates are not strictly correct in the context of our model for two reasons: first, because they are based on a transformation of the first-order condition for labor supply that does not include an intertemporal component (see MacCurdy (1985) or Pencavel (2002) for overviews of the methodology underlying this literature); and second because the life-cycle labor literature typically uses as disutility-of-labor utility function rather than a utility-of-leisure, though the two are so closely related that, in the absence of human capital considerations, the labor elasticity of one implies the leisure elasticity of the other.

Results from the calibration exercise are reported in the first part of Section 4 below.

4 Results

Table 6 gives the calibrated parameters and associated targets in the full augmented model. The choice of the targets is explained in previous sections. The exact value of some of the targets (marked \sim in table 6) are temporarily withheld in deference to Statistics Canada’s disclosure policy. All results in this section are from simulations of 16,000 potential couples (32,000 individuals), 50% male and 50% female.

In the full, unrestricted model, all individuals are matched in the first period of life with a member of the opposite sex. Of these, 88.5% choose to get married within the courting spell, which we set to $N = 5$. There is no remarriage: this assumption simplifies computation substantially and partly understates the value of divorcing; we correct for this by appropriately calibrating a multiplicative utility weight for divorcees which matches the aggregate divorce rate²⁸. When a divorce occurs, individual utility is weighted by a multiplicative weight for divorcees, $\theta(d)$. Former partners set up their own households and revert to the optimization problem for (divorced) single agents.

4.1 The One-Shock versus the Two-Shock Model: Explicit Health Risk

We begin this section by focusing on the mechanisms at work in the single-shock AR(1) model vis-a-vis a two-shock model with disability and human capital accumulation. Our goal is to assess how far a finer shock representation, together with endogenous human capital accumulation, can go in reconciling model and data. For this exercise, we shut down intrahousehold transfers of time and most endogenous aspects of marriage. We assume that individuals are matched exogenously (assortatively by education) so that 85% of the population of each gender receives a match and the rest remain single.²⁹ Moreover we set the courting period N to one year, so that marital pairs form instantaneously. Individual pairs still bargain over their

²⁸The endogenous divorce rate (21%) is roughly the divorce rate of *first* marriages in Canada.

²⁹We choose a θ^M high enough so that those who are matched always choose to marry and remain married until death.

Table 6: Calibrated parameters and targets: augmented model

Parameter	Value	Target	Value: data	Value: model
γ_c^f	1.4	Ufl^a	26.4	25.2
γ_l^{mb}	24.2	Uml^c	35.5	35.7
ψ^f	2.4	ϵ_w^{nd}	.50	.49
ψ^m	2.9	ϵ_w^{ne}	1.0	1.10
ω^f	2.0	taken from literature		
ω^m	2.0	taken from literature		
$\delta_{l,ds=3}$	1.13	$Uml ds 3$	~ 32.0	32.0
$\delta_{n,ds=4}$	1.36	$Uml ds 4$	~ 29.0	30.1
$\delta_{n,ds=5}, \delta_{l,ds5}$	1.11, 1.09	$Uml ds 5$	~ 23.0	26.0
$\delta_{n,ds=6}, \delta_{l,ds6}$	1.36, 1.37	$Uml ds 6$	~ 6.0	8.0
$\delta_{l,ds(r)=3}^f$	1.11	$Uml ds(r) 3$	~ 32.0	32.0
$\delta_{l,ds(r)=4}$	1.37	$Uml ds(r) 4$	~ 29.0	30.1
$\delta_{l,ds(r)=5}$	1.69	$Uml ds(r) 5$	~ 23.0	26.0
$\delta_{l,ds(r)=6}$	2.05	$Uml ds(r) 6$	~ 6.0	8.0
β_m	1.001	$\bar{a}_m / \bar{inc}_m^g$	2.1	2.0
β_f	0.997	\bar{a}_f / \bar{a}_m^h	1.5	1.2
$\theta(m)$	1.017	$P(marry)^i$	90%	88.5%
$\theta(d)$	1.011	$P(divorce)^j$	21%	19.3%

^aPrime-age unmarried female labor supply

^b γ_c^m is normalized to 1.0; γ_c^f is normalized to .8.

^cPrime-age unmarried healthy male labor supply

^dMale Frisch elasticity

^eFemale Frisch elasticity

^f ds states for retirees' leisure limitation transition matrix

^gHousehold wealth to income ratio for hhs with male members

^hRatio of asset holdings for hhs with female and male members

ⁱProportion of individuals who marry

^jProportion of marriages that end in divorce

household sharing rule; however, since individuals differ only by education at the start of life, all education-specific pairs choose the same bargaining parameter.

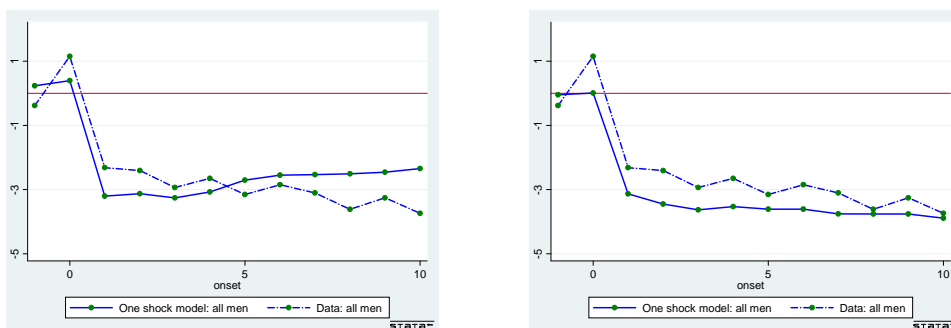
In order to assess the ability of a single shock model to capture responses to disability-like shocks, we need to generate a sample of individuals whose experiences make them comparable to the ‘disabled’ sample in the SLID and in the data simulated from the two-shock model. This is how we do it: we first estimate a standard AR(1) shock process, like the one described in (9) from SLID data.³⁰ The persistence parameters of the AR1 shock, ρ , are 0.944 for men and 0.957 for women, and the variances of the shocks are 0.021 for men and 0.018 for women. We then discretize the estimated shock range into seven levels, with 1/3 of individuals receiving the mean shock of zero in any period. The simulated “disabled” sample includes those who happen to receive negative shocks (of any magnitude) in at least 2 out of every 5 sample periods, over a ten-year window. “Onset” refers to the year corresponding to the first shock in this unusually high sequence of perturbations. The resulting mean age of “onset” is 33, compared to 34 for the disabled population in the simulated two-shock economy.

Figure 8 plots the predicted total hours response of men in the one-shock and two-shock simulations. The simulation based on the two-shock process does a marginally better job of matching the chronicity of the labor supply response to disability, but the differences are negligible. Things change dramatically, however, when we disaggregate responses by marital status, as shown in figure 9. Married men (the blue line) reduce their labor supply in response to a shock, then gradually increase their labor supply over time, either because they receive better shocks or because the wealth effect of the bad shocks begins to dominate the substitution effect. Single men (the green line), on the other hand, actually *increase* their labor supply in response to a sequence of negative wage shocks. These effects are consistent with life cycle theory and with analytical predictions described in section 1 for the case in which the negative wealth effects outweigh substitution effects. However, these patterns are at odds with observed responses to *disability* shocks, especially for single men.

The ability of the two-shock model to provide a better approximation to the dynamic labor supply of men is due to the nature of disability shocks, as discussed in section 3: the reduction in the individual’s total disposable time makes him less likely to respond to negative shocks by increasing his labor supply. Nevertheless, the basic two-shock model is still unable to solve either of the “puzzles” described in Section 1. In figure 9 we observe that married men reduce their labor supply more than singles in response to a disability shock. This is expected, because in a model with no divorce and very little initial sorting, marriage provides a lot of insurance: the presence of a second earner in the household reduces the negative wealth effect of a disability shock and allows the disabled spouse to reduce his labour supply more in response to the shock. As well, given that transfers of time are not available, an ‘added worker’ effect emerges as the optimal response for the spouse. These effects are evident in figure 10, which plots the simulated spousal responses or ‘added worker’ effects. The one-shock and two-shock models produce large, nearly identical, added worker effects.

³⁰Results of this estimation are presented in table 5, Appendix 2.

Figure 8: All men



(a) One shock model

(b) Two shock model

4.2 Unrestricted Model Results: Labor Supply Responses and Marital Patterns

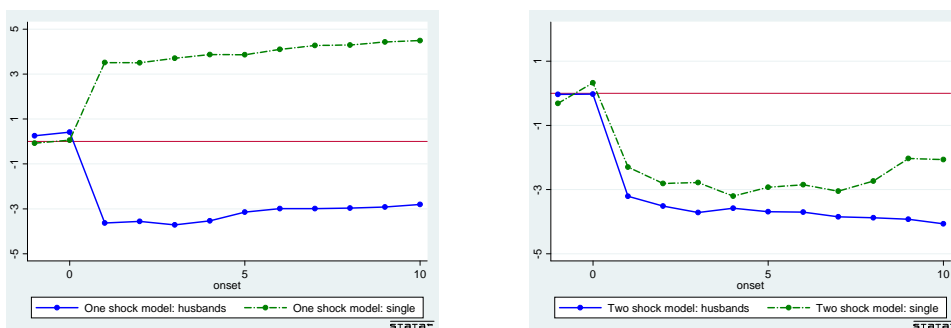
Given the failure of the basic two-shock model to resolve the puzzles, we move on to show the results for the two-shock model in which there are no restrictions on intrahousehold time-transfers and on marriage choices. Figures 11-14 show simulated dynamic labor supply responses to shocks for the ‘disabled’ and ‘chronic-disabled’ samples, corresponding to similarly defined groups in the SLID. We consider the model’s performance along the total hour, intensive hour and extensive hour margins.

The model captures patterns of total labor supply for married and single men, both for the ‘disabled’ and ‘chronic disabled’ populations defined in section 2. Specifically, we are able to rationalize the larger and more gradual drop for single men relative to married men: as in the data, by the tenth year after onset, single men in the ‘disabled’ group have reduced their hours more than twice as much than married men. The model also captures the same broad patterns for hours of workers and for participation.

Figure 15 shows corresponding effects for the post-onset path of wages for the disabled sample. Wages fall slightly at onset, then continue to fall over the next ten years. Unsurprisingly, the effects are larger for single men, who experience larger declines in average work hours, than for married men. This patterns are, again, consistent with what we observe in the data.

What drives these results? It turns out that the primary mechanism by which the restricted and unrestricted two-shock models differ is the endogeneity of marriage, and in particular the selection into and out of marriage. Marital sorting creates a pool of singles with different characteristics from marrieds and with lower average returns to labor. For this group, withdrawing from the labor force in response to shocks, and claiming benefits, is relatively more attractive. A secondary mechanism driving the differences between marrieds

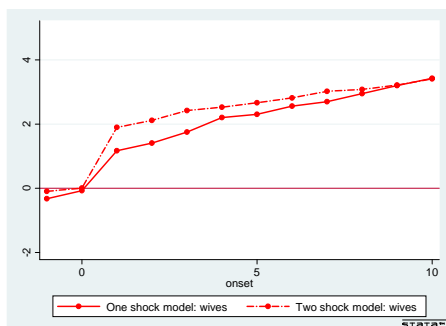
Figure 9: Married and single men



(a) One shock model

(b) Two shock model

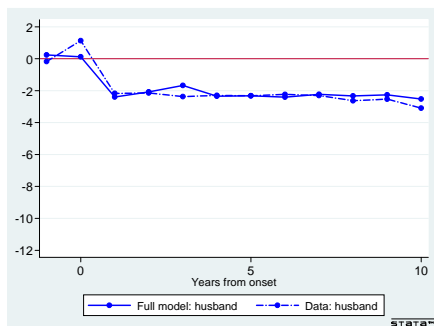
Figure 10: Wives of disabled men



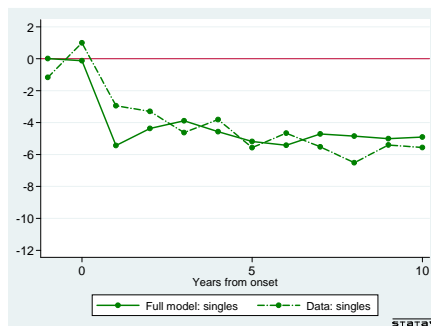
and singles is the set of changes induced by health and disability shocks on married men's bargaining power *within* their marriage. Falling into ill health can reduce a husband's relative utility weighting within the household, so that household "pressures" effectively limit the drop in hours worked. This second effect is not directly observable in data, while the role of sorting is observable and therefore testable.

To give some insight into how and to what extent marital sorting works in the model, table 7 (similar to table 3 in section 2) reports the observed joint distribution of marital status and disability status in the adult SLID population and its counterpart in the unrestricted, simulated economy. In both data and model, married men are much less likely to report a disability over a six-year window than single or divorced men, and are much less likely relative to their unmarried counterparts to report a disability in five or more years out of six. Since, conditional on exogenously varying health (rsk), all men are subject to the same process for disability, these effects are driven entirely by sorting. Men who are perceived as high health risks (those who transition into rsk state 2 or 3 early in life) marry much less in the model,

Figure 11: Labor supply responses: Total hours, Disabled sample

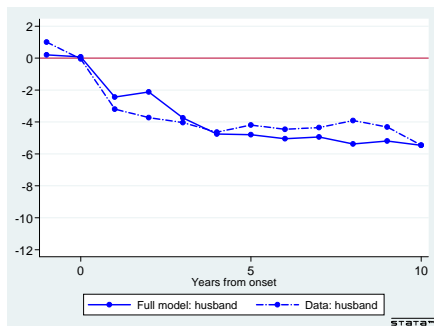


(a) Husbands

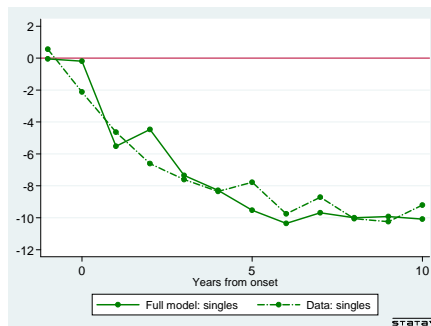


(b) Single Men

Figure 12: Labor supply responses: Total hours, Chronic disabled sample



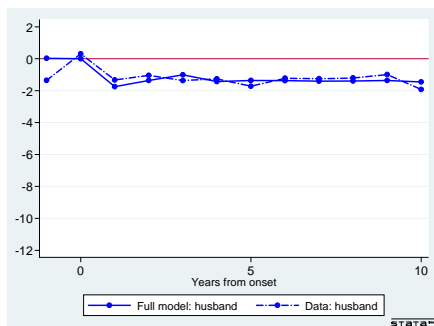
(a) Husbands



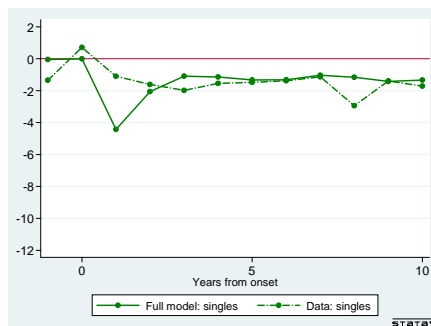
(b) Single men

while married men who receive bad health shocks are more likely to divorce over the long-run. Figure 16 shows the patterns of divorce, by duration of marriage, generated by the model and compares it to the corresponding measure from data on Canadian marriages beginning between 1973 and 1986. As in the real world, divorce in the model is concentrated early in the couple's life cycle, after an initial series of shocks is realized. Early in the life cycle, and the marriage, shocks convey a lot of information about a partner's ability to provide future income and insurance. These early shocks are therefore more likely to trigger a divorce than are similar shocks experienced at later stages of the life-cycle.

Figure 13: Labor supply responses: Hours of workers, Disabled sample

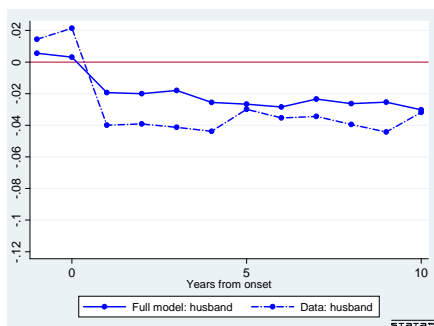


(a) Husbands

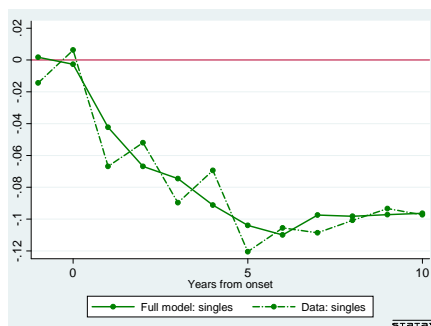


(b) Single Men

Figure 14: Labor supply responses: Participation, Disabled sample



(a) Husbands



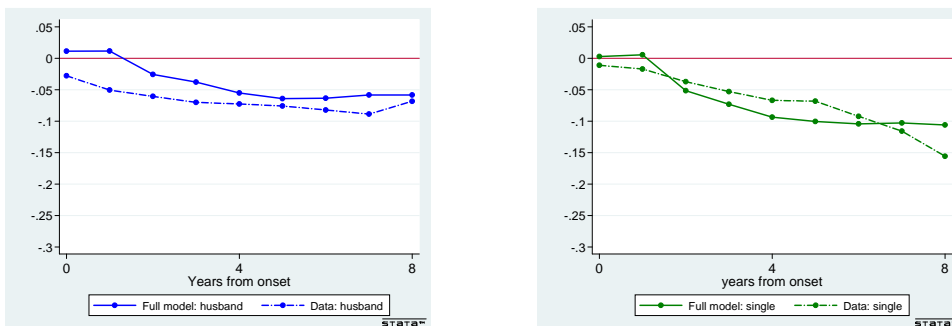
(b) Single men

4.3 ‘Added Worker’ Effects and Intrahousehold Time Transfers

The second puzzle described in section 1 is the near absence of added worker effects, particularly for wives, in Canadian and American data, in spite of the apparently large wealth costs associated with disability onset in the household head. If marriage provides insurance through consumption transfers, we should observe large swings in married individuals’ labor supplies in response to a spouse’s disability onset, like we do in both the AR1 and in the ‘basic’ two-shock model.

The unrestricted model, on the other hand, generates almost no added worker effects, like in the data. The solid red line in figure 17 plots spousal labor supply responses for the full, unrestricted model. In the same figure we also report (gold line) the spousal labor responses

Figure 15: Wage effects of disability, Disabled sample



(a) Husbands

(b) Single men

Table 7: Age-adjusted observed and simulated disability by marital status: 40-49 year olds

Data					Simulation				
Frequency of reports (z)	NM (%)	M (%)	S/D (%)	W (%)	Frequency of reports (z)	NM (%)	M (%)	S/D (%)	W (%)
0	0.516	0.634	0.501	0.583	0	0.488	0.586	0.500	0.563
1-2	0.252	0.222	0.264	0.240	1-2	0.323	0.300	0.310	0.308
3-4	0.108	0.082	0.121	0.097	3-4	0.101	0.076	0.109	0.095
5-6	0.095	0.061	0.115	0.079	5-6	0.087	0.038	0.081	0.035

when, everything else equal, we do not allow any time transfers. It becomes apparent that shutting down the time transfers technology is enough, by itself, to reverse the added worker effect result and generate counterfactual responses.

Why does the full, unrestricted model generate such a small added worker effect? Two simple effects, which naturally arise when we don't restrict intrahousehold time-transfers, explain this result. First, the wealth effects of disability are typically lower for married couples in the full model than in the basic two-shock model, or in a model with time transfers shut down. In part, this is driven by marriage market selection: women tend to marry men with relatively higher earnings, so that households can more easily self-insure against the risk of early retirement. For the most part, however, the wealth effect is generated directly through additional specialization within marriage, which allows husbands to devote more time to labor activities.

To give an idea of the magnitudes involved, table 8 reports some results based on the comparisons of two simulations: one for the full, unrestricted model; the other for same economy, but with all time transfers exogenously and ex-post (after marriage decisions have been taken) set to zero. By construction, matching, marriage and divorce are identical in the two economies: the only difference is that in the second economy all time transfers are

Figure 16: Canadian Divorce rate by Duration of Marriage

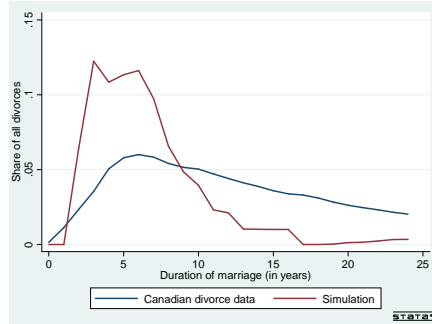
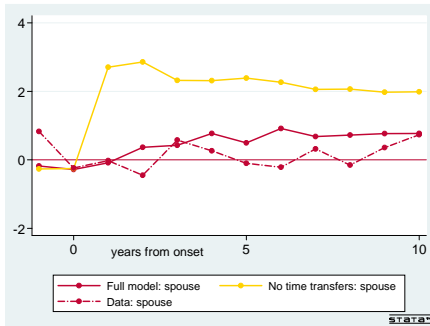


Figure 17: Added worker effects and intrahousehold time transfers



(a) Spouse effects: wives of disabled husbands



(b) Spouse effects: wives of chronic disabled husbands

magically reversed.

The first row of table 8 shows the change in the average marriage surplus across the two economies. Marriage surplus is defined here as the difference between the sum of wife's and husband's value functions in marriage and the sum of wife's and husband's value functions in the event of divorce. It is roughly 15% lower in the model without time transfers than in the parallel unrestricted economy. The next three rows show that eliminating time transfers leads to a modest reduction in the average married male wage of close to 5%; to a negligible corresponding increase in average wife's wage; and to a proportionally larger drop of close to 7% in married households' earned income.

These results are consistent with two facts:

1. from subsection 4.1 we know that labor supply responses to negative wealth (perma-

ment income) changes can be strong. The negative wealth effect from the loss of time transfers, stimulates labor supply by husbands. The greater resulting share of disposable time devoted to labor compensates to some extent for the labor hours lost due to the reduction in husbands' total disposable time in the absence of time transfers, and prevents large drops in husbands' human capital;

2. the low wage returns to women are fully consistent with the low returns to work for women estimated in section 3.1.3. This effect is exacerbated by the fact that women who invest the most in making time transfers to their husbands are those with very low wages and poor prospects for accumulating human capital in the first place.

The slightly larger effect on household earned income relative to wages arises in part due to the fact that couples are poorer in effective resources, with less potential for the husband to rise up through the wage distribution. As a result, a greater share of marrieds (about 1.5% in any period) find it optimal to withdraw from the labor market and go on welfare or disability insurance. The last row of table 8 shows the effect on marriage rates if we allow individuals to choose their marital status in the no-time-transfer economy, *ceteris paribus*. The total amount of marriage that takes place before age 25, the last year of courtship, falls from 88.5% to 77.1%. which is consistent with the drop in marriage surplus.

Table 8: Quantifying the Relevance of Time Transfers

	Unrestricted model	Model w/o time transfers
Average marriage surplus	.0713	.0613
Average male wage	\$17.70	\$16.90
Average female wage	\$15.70	\$15.85
Average hh income ^a	\$1,114	\$1,041
Share of matched pairs who marry	88.5%	77.1%

^aWeekly income of working-age married households, excluding benefits

4.3.1 Investigating the Dynamic, Joint labor Supply of Couples

The most powerful mechanism working against an added worker effect in the full model is the dynamic incentive for wives to maintain their husbands' human capital stock over the life-cycle, which leads to higher household wealth in the long run. When a husband receives a negative shock, wives may find it optimal to increase the amount of time devoted to the spouse (as a response to disability, we think of this as "caring"), even though a labor-limiting disability shock acts like a negative wage shock in the short run, and even though doing so naturally reduces the time available for her to increase her own labor supply. Perhaps unsurprisingly, we observe this behavior in both the unrestricted model and in the data.

Some evidence in support of the time transfers mechanism is reported in table 9, which contains the estimated change in h , the intrahousehold time transfer from wife to husband,

Table 9: Time Transfers and Disability Shocks

Table 10: Wife’s Time Transfer Response to Couple’s Wage Ratio and Husband’s ds

Estimates from PSID/ATUS				
	All hhs	hhs under 30	hhs 31-45	hhs 46-60
\bar{w}	1.90*	2.69*	2.53*	1.15*
$ds > 1$	2.91*	-0.48	3.99*	3.14*
Simulation results				
	All hhs	hhs under 30	hhs 31-45	hhs 46-60
\bar{w}	1.54*	5.82*	2.22*	1.20*
$ds > 1$	1.11*	0.38*	1.90*	1.09*

associated to changes in husband’s disability status and the couple’s wage ratio, for both model and data. The top panel reports results from a PSID sample of husbands and wives, using the imputed measures of *nll* (non labor, non leisure) time from the ATUS, as described in section 3.2 and Appendix 2; the bottom panel shows the results for the same regression using simulated data on couples in which the husband is currently working.³¹ The magnitudes of the effects are smaller for the simulated data, but the broad patterns are similar: on average, total time transfers from wife to husband are higher when the husband is currently disabled, controlling for the couple’s relative wages. Moreover, the effects are always strongest for households in the prime of the life cycle, between ages 30 and 45, because of the high marginal value of human capital at those ages. In contrast, the response of h to the couple’s wage ratio \bar{w} declines monotonically over the life cycle in both data and model. The difference between the timing of the strongest responses to ds and \bar{w} has to do with the magnitude of the average disability shock, which is increasing into middle age. Realizations of relatively high values of δ_n and δ_l become more common among middle age men receiving any disability shock and necessitate a stronger response from the wife to prevent their labor supply from falling. This effect counters the declining importance of the husband’s human capital to household permanent income as the couple ages.

4.4 Insurance within Marriage

The results described above suggest many interesting questions about the insurance role of marriage: how do spouses fare under the marriage contract? Who gains the most from marriage and who is excluded? The labor supply and marital sorting results outlined above suggest that men who receive bad shocks early in life may lose the insurance offered by marriage – that is, they are excluded on the extensive margin. However, conditional on being *within* a marriage, the implied insurance value is different between spouses and over the

³¹The last restriction is imposed because we only observe the couple’s wage ratio if both members worked in the previous year. Very severely disabled husbands who are not currently working are therefore excluded from both estimations.

life-cycle.

Figure 18 shows the life-cycle profile of the sharing rule parameter λ for couples³², disaggregated by whether the husband is currently in the low risk (blue line) or high risk (red line) health state. The first figure shows the λ s averaged over all marriages in the model; the right hand side figure considers only couples who, once married, never divorce. The similarity between the pictures suggests that divorce patterns are not driving most of the life-cycle variation in λ directly. Unsurprisingly, the λ for high-risk husbands (excluding the chronic disabled) lies above the line for the lower health risk husbands, especially for younger households for whom the husband's high health risk is a bad sign for future earnings: recall that men move from the low to the high-risk state monotonically and at an increasing rate over the life cycle. Early transitions to the high-risk, poor-health state signal to spouses or potential spouses that the husband will not be a productive mate.

Figure 19 provides a window into the main mechanism driving the change in λ : renegotiation. The left-hand panel shows the share of successful female-triggered (the blue line) and male-triggered (the red line) renegotiations for low-risk households; the right-hand panel shows the same thing for high-risk households. The figures suggest that the path of the mean λ by age plotted in figure 18 actually masks much of the movement in λ among married couples. For instance, at age 30, about 9% of high-risk households and 7% of low-risk households renegotiate their marriage contract successfully, compared to less than 2% of households that divorce. The U-shape of renegotiations over the life cycle suggests (consistent with the divorce results) that marriages become more stable the longer the couple is together. In part this is due to the consumption costs of children. Couples who are close to splitting do so before children arrive in order to avoid ending up as single parents with the associated high consumption costs \tilde{n} . However, it is also due to the fact that, as shocks are realized and uncertainty resolved over the early years of a marriage, additional information becomes less and less important and increasingly unlikely to break up the match. This patterns hold until changes in the relative value of marriage and worsening health of husbands once again destabilize the match close to retirement. However, the spate of late-in-life renegotiations generate little divorce. From the figures, it is also clear that low-risk marriages (i.e. marriages in which the husband is in the good or low-risk health state) are more stable overall: less renegotiation is triggered within these marriages at every age relative to high-risk marriages.

Another interesting question is *who* triggers renegotiation. For younger households, wives in low-risk marriages trigger renegotiations slightly more than their husbands, producing the very gradual rise in mean λ observed in figure 18. The reason is that, in low-risk marriages, the initial λ is relatively lower. If the wife receives a series of positive shocks relative to the husband, she is more likely to threaten a divorce than the other way round - that is, if the husband, who is already relatively powerful within the marriage, receives a series of good shocks relative to the wife.

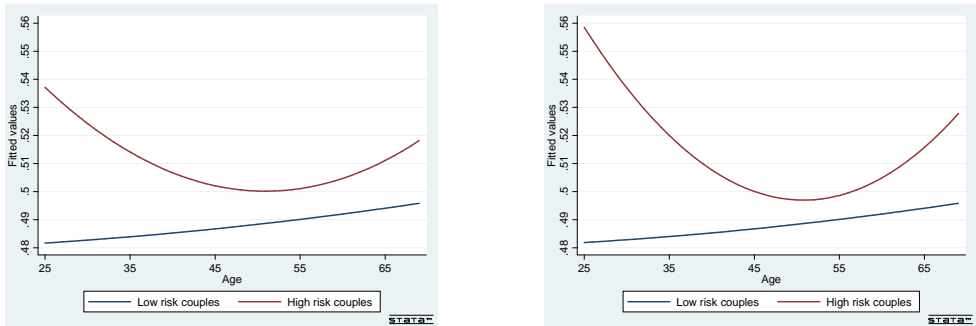
The opposite occurs for high-risk marriages. Within these marriages, a greater share of renegotiation (about 0.5% of marriages per year) is triggered by high-risk husbands than by their wives. High-risk men derive large utility gains from marriage early in life due to the

³²This parameter indicates how large a weight wife's utility has within a household.

availability of spousal transfers that aid in the accumulation of human capital and provide insurance versus expected early shock realizations. As these husbands age, their gains from marriage decrease as “buffer stocks” of human capital and assets are accumulated. When a husband’s gains from the marriage at the initial high λ become non-positive, a renegotiation is triggered and a new lower λ is agreed upon.

Finally, as couples age toward retirement, renegotiations are increasingly triggered by wives. The effect varies relatively little across the husband’s health type since the distinction between health types fades close to retirement; by age 60, most husbands have already moved or are expected to soon move into the high health-risk state. Husbands’ relative gains from marriage increase as they approach retirement: this is due to the approaching loss of labor income, the fact that assets would be split evenly in divorce and the increasing health risk due to age.³³

Figure 18: Life cycle paths of λ



(a) λ : All couples

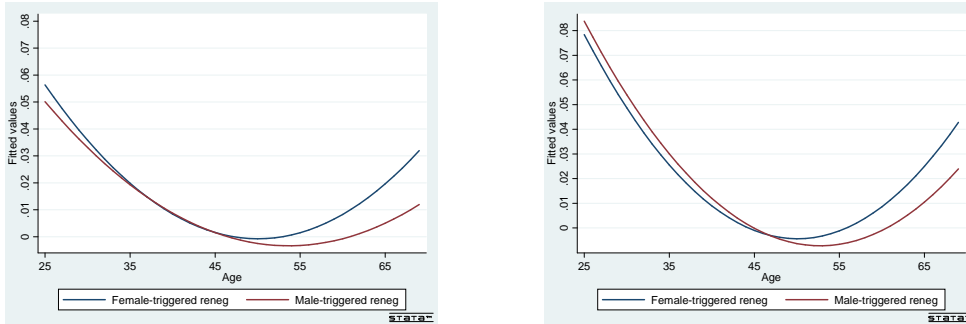
(b) λ : Couples who do not divorce

The gains from marriage to men are large: in the model, the mean wage of a married man is \$20.50 per hour, compared to \$15.25 for an unmarried man, which is slightly smaller than the difference of about \$6 observed in the data.³⁴ Figure 20 shows some additional cross-sectional results. We report only results for men, for whom the model performs best and for whom we have specified the full spectrum of life-cycle risk. The top panels of figure 20 plot the distribution of log wages for men by marital status: the left-hand panel shows the fitted

³³Wives are not subject to leisure shocks in the model, only to wage shocks. Since all uncertainty in the wife’s life, save mortality risk, is resolved by age 65, and she is assured of her Old Age Security benefit as well as half of accumulated household assets in the event of divorce, her gains from marriage fall approaching retirement and renegotiation drives the average λ back up. It would be interesting to what extent old age dynamics change if we relax the assumption that women do not receive leisure-limiting disability shocks in old age.

³⁴For women, being married slightly reduces the average wage, but the effects are much smaller than the gains for men. In the data there is almost no difference between married and single women’s wages.

Figure 19: Husband- and wife-triggered renegotiations



(a) Renegotiations: low risk couples

(b) Renegotiations: high risk couples

variances generated by the model and the right-hand panel shows the corresponding fitted variances from the full sample of men in the SLID, in both cases for the bottom 98% of the wage distribution. The model provides a good fit to the data, generating total cross-sectional variances in log wages of about 0.245 for single men and 0.247 for married men, which is nearly identical to the measure calculated from the sample of men in the SLID. In both the data and the model, the variance of wages is slightly higher for single men over most of the life cycle, before the pattern reverses around age 45 to 50.

The bottom right panel of figure 20 shows the corresponding cross-sectional variances of the log of consumption for men generated by the model.³⁵ Note that, while the life cycle variance of log wage varies only slightly by marital status, the life cycle variance of log consumption is quite different for marrieds and singles: in general, the variance of log consumption is larger for single men than for married men, but it remains relatively flat through middle age, one most marriage market sorting has been realized. The differences between marrieds and singles and the patterns over the life cycle can be ascribed to three main effects: (1) sorting into and out of marriage; (2) the increasing reliance of poorer single men on either disability or early retirement benefits in middle age; and (3) the absence of intra-household consumption sharing by a spouse. The first effect will tend to cause a contraction in the variance if sorting out of marriage generates a relatively homogeneous group of less endowed males. The second effect is responsible for the flattening of the consumption profile for singles in middle age.

The third effect is our main concern, however. To this end, figure 20(d) in the bottom right hand panel of figure 20 provides a somewhat cleaner indication of the consumption insurance value of marriage for men. It plots the absolute value of deviations in the ratio of men's year-to-year marginal utilities of consumption ($\frac{u_{ct}}{u_{ct+1}}$) from the theoretical full-insurance value. In

³⁵Unfortunately, no corresponding consumption data is available in the SLID.

a world with full insurance (including fair annuities to insure against mortality risk), the intertemporal ratio of marginal utilities would be a constant equal to the product of the male discount rate β_m and the interest rate. In the model, deviations from this value (Dev-RMUC in shorthand) represent deviations from full insurance. Figure 20(d) suggests that deviations from full insurance are much larger for single men than for married men, by more than 100% around age 40, using the Dev-RMUC metric. For both marrieds and singles, the Dev-RMUC falls over middle age as uncertainty is resolved and assets are accumulated. The sharper fall for singles is again due to their inferior labor market outcomes, as observed in section 4.2: in fact, single men in the model are more than three times more likely not to be working than married men between the ages of 40 and 55.³⁶ Early retirement with a fixed benefit naturally resolves much of the life-cycle uncertainty for single men and thereby brings them closer to the full-insurance benchmark.

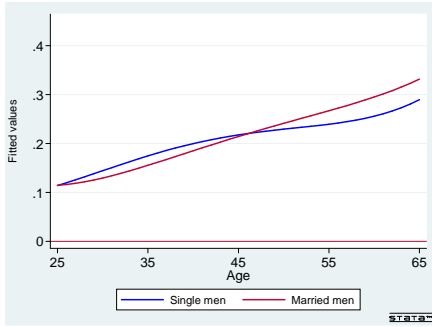
5 Conclusions

In this paper we study the dynamic responses of households to a specific class of idiosyncratic shocks. By using detailed panel information about the timing, persistence and responses to disability (health) shocks we document the patterns of dynamic optimal responses by gender and marital status. We then show that most predictions of a ‘basic’ household model, with persistent idiosyncratic risk, are at odds with the observed data patterns: we find that this model fails in several dimensions, including timing, magnitude and persistence of the optimal household responses to health shocks. We argue that this puzzle can be solved if we remove some of the restrictions implicit in the ‘basic’ model and allow for a richer set of intra-household allocations and marriage participation decisions. We show that the dynamic responses observed in the data are consistent with a model in which people sort endogenously into marriage and, conditional on being married, people can provide insurance through transfers of time, as well as consumption. We provide micro-data evidence in support of both these mechanisms: in particular we show that increased transfers of time tend to occur when disability shocks hit, and that endogenous marriage sorting goes a long way in explaining divorce rates and disability report rates by marital status. More importantly, the unrestricted model generates dynamic labor supply responses which are extremely close to the ones observed for both singles and couples.

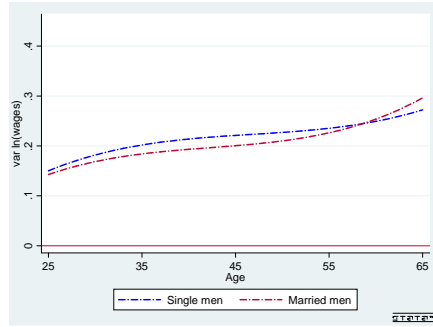
We also investigate how the ‘marriage’ contract is chosen to accommodate the possibility of idiosyncratic shocks and to provide insurance which is unavailable to single agents. We find that main-earners (husbands) tend to transfer consumption to second-earners (wives) who are happy to trade additional consumption insurance for time transfers in periods of need. The latter transfers serve as an important mechanism within the household, as they allow the main earner to smooth labor supply and achieve higher levels of human capital

³⁶The participation rates generated by the model are 82% for single men, 96% for married men and 84% for divorced men between 25 and 54. These compare quite well to observed participation rates by marital status in Canada of approximately 84%, 94% and 86% for single, married and divorced men respectively.

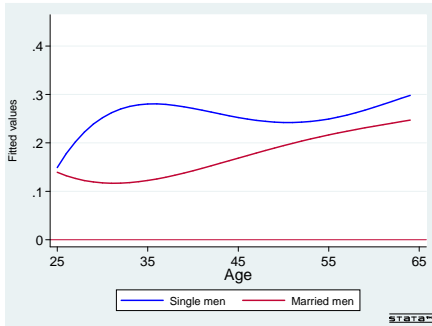
Figure 20: Cross sectional variances of men's ln wages and consumption



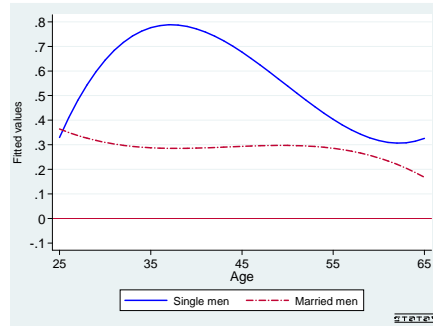
(a) Ln(wages): model



(b) Ln(wages): data



(c) Cross sectional variance of ln(consumption)



(d) Dev-RMUC

and earnings. However, in a limited commitment environment, we also show that incentives to renegotiate (and possible terminate) the contract can arise over the life-cycle because of the history of idiosyncratic shocks: in particular, we show that the relative value of marriage changes for men and women as they age. Men who are at high risk of receiving idiosyncratic shocks value marriage early in life, when they are poor in assets and human capital, while all men value marriage at the late stages of their working life as they approach retirement and periods of high health risk. Accordingly, sharing and allocations within households change, in some cases non-monotonically, over the working life. These changes, typically unobserved in micro data, play an important role in explaining individual and household responses to shocks, and labor supply decisions, over the life cycle.

Appendix 1: Estimation of longitudinal effects of disability using the SLID

Disability questions in the SLID

In the SLID, an individual is classified as disabled if he or she reports a limitation along any of the following dimensions: (1) easily completing one or more routine physical activities such as climbing stairs; (2) accomplishing required or desired activities ‘at work’ or ‘at a job, business or school’; (3) accomplishing required or desired activities ‘at home’; or (4) completing required or desired ‘other activities’ such as those associated with transportation or leisure. The questions about disability limitations ‘at work’ are asked of respondents or about subjects under 70 who worked in the reference year. The question about disability limitations ‘at a job or business or at school’ is asked of respondents under 70 who did not work in the previous year. In the longitudinal file, the responses to these questions are combined into a single variable reported for the entire sample population under age 70.

For the latter three types of limitation, individuals can respond that their were limited in the respective type of activity ‘sometimes’ or ‘often’. Additional questions are asked of workers who report any type of disability: whether the condition made it difficult to change jobs or find a better job; and whether individuals wanted to work more or fewer hours due to their condition. Individuals who report a limitation but did not work in the previous year are additionally asked if their condition ‘completely prevents’ them from ‘working at a job or business or looking for work’. Finally, individuals who report a current disability along at least one of the dimensions detailed above are asked about the duration of their condition: how many years they have had it and whether it was present at birth.

Sample adjustment and creation of “disabled” and “chronic disabled” samples

Two adjustments to the 1999-2004 SLID sample are required before we can be confident that equation (20) is giving us consistent estimates of the longitudinal effects of disability onset on individual and/or family responses:

- first, since questions about the history of disability are asked only of those who report a current limitation, we expect that the chronicity of an individual’s disability to be correlated with the reported duration of the condition.³⁷ We define chronicity as the proportion of periods in which a current disability (latent or limiting) is reported. Duration is the number of periods (years) the currently-disabled person claims to have had the condition previous to joining the SLID panel.³⁸

³⁷This problem does not arise in the U.S. studies cited above because the assumed year of onset is directly observed.

³⁸The first report of long-term disabilities do not necessarily occur in the first year of the panel. We also assume the actual duration of the condition is the first one reported when the condition becomes current.

To deal with this problem, we randomly select individuals from different chronicity-duration “bins” in order to equate chronicity across duration. We define two samples, called our ‘disabled’ and ‘chronic’ disabled samples. For the ‘disabled’ sample, we randomly drop some individuals with especially high chronicity and long reported durations, and some individuals with especially low chronicity and short reported durations. The mean chronicity for this sample is .45 (close to three of six), which is the average chronicity of the aggregate currently-disabled SLID population. Our “chronic” sample roughly corresponds to the chronic sample defined by Meyer and Mok (2006) and Charles (2003). For this sample, we drop all individuals in all duration bins whose chronicities are less than .25. For this group, the resulting mean chronicity is .62 or nearly four of six panels.³⁹

- the second sample-selection issue deals with establishing control groups which we can use to estimate the $\widehat{\delta}_k^h$ parameters in (20). If disability shocks are correlated with other demographic or unobserved effects in the population, the estimated $\widehat{\delta}_k^h$ values will be inconsistent estimates of the time-from-onset effects because they will also contain information about being the sort of individual who reports a disability in the first place. The U.S. studies cited above avoid this problem by using as their control groups the future-disabled who are at least four or five years from onset and omitting everybody else, the ‘never-disabled’, from the regressions. However, our short panel makes this an unfeasible exclusion. Instead, we use logit-based propensity score matching, with k nearest neighbors, to establish control groups that combine all future-disabled individuals two or more years from onset with a comparable subset of the never-disabled. The value of k , as well as the set of control variables used in the regressions, is set optimally through trial and error for each gender and marital status category. Matching is performed on selected demographic (e.g. education, urbanization, region), life-cycle (e.g. age, size of household, whether paying or receiving private support) and subjective (e.g. health status, level of stress) variables, as well as occupation dummies. Table 11 reports the results from this selection procedure for the ‘average’ disabled sample. Results for the chronic group are broadly similar and are available in Gallipoli and Turner (2008). The columns report coefficients from dummy variable regressions of each dependent variable in the table on the indicator function associated to being a future-disabled. In the first and third columns, the omitted (benchmark) category is the unselected never-disabled (UND). In the second and fourth columns the omitted category is the selected never-disabled (SND) from the propensity matching. Asterisks indicate differences at the 5% significance level, based on boot-strapped standard errors from fifty resamplings.

The results in Table 11 suggest that propensity matching is largely effective: for all four gender-marital status groups there are almost no significant differences between the ‘selected-in never-disabled’ and the ‘future-disabled’ groups. This stands in sharp contrast to the

³⁹One drawback is that this procedure necessarily omits individuals reporting their first ever limitation after the final two years of the panel: for these individuals we lack sufficient post-onset information to determine chronicity.

‘unselected never-disabled’ who appear to differ significantly along labor supply, income and demographic dimensions from the ‘future-disabled’. The possible exception to this general point are single men, for whom the propensity matching decreases the similarity of the groups along some dimensions, such as own earnings, and improves it along others, such as transfers and household income. The results are especially strong and encouraging for the economic variables we report (hours worked, participation rates, individual earnings and transfers received and household-level income), which are not included as regressors in the propensity score logit regressions.

Table 11: Propensity matching results: Average disabled group

Variable	Singles		Married	
	UND	SND	UND	SND
Men				
Demographic				
Age	6.94*	2.31*	1.39*	0.84*
HH size	-0.34*	0.01	-0.079	-0.11*
Education	0.026	-0.038	-0.068	-0.029
Economic				
Hours	-2.29	-45.34	27.27	-19.63
Participation	0.019	0.028	.061*	0.024
HH income	-9285.39*	3115.38	-7503.93*	-4099.89
Ind transfers	924.42*	136.56	-174.49	-24.39
Ind earnings	554.12	1574.85	-5093.6*	-1799.93s
Women				
Demographic				
Age	12.99*	0.005	0.103	-0.15
HH size	-0.621*	0.007	0.041	0.034
Education	0.718	-0.16*	0.061	0.04
Economic				
Hours	195.54*	40.66	9.37	5.65
Participation	0.045*	-0.013	0.056*	-0.01
HH income	-11622.8*	-4.79	-9758.56*	-2066.7
Ind transfers	1111.03*	345.32	127.05	-126.14
Ind earnings	4540.24*	-2747.71	-3110.75*	-974.61

Appendix 2: Methodological issues and supplemental estimates for the numerical model

Specification for the human capital process for men and women

To estimate the dynamic human capital process given in table 4 in section 3.1.3, we adopt a four-stage estimation process that controls for possible endogeneity. In the first stage, we use Wooldridge (1995)’s panel data selection estimator to calculate inverse Mills ratios for the

likelihood of observing a wage for all individual-year observations in the sample. This method controls for individual fixed effects through the inclusion of time-averages of the covariates used to predict the wage. Our exclusion restrictions are: current household income from investments, household size and number of children, a dummy for having experienced a death in the family in the previous year, and measures of the annual provincial unemployment rate and the deviation in this rate from its 10-year average. In the second stage, we predict hours in 2001- 2003 using hours in 1999, and the human capital stocks in 2001-2003 by the wage observed in 1999 (adjusted for selection and fixed effects using the inverse Mills ratios and the time-averages of the covariates.) In the third stage, we estimate the main parameters of our human capital equation using our instrumented values of lagged hours and lagged wages as regressors. Our sample here is limited to the last three observations for men and women in the 1999 panel who appear and report all information needed for the regression in all six years of the panel. In the fourth stage, we substitute the actual lags of hours and wages for the instruments and use the consistent marginal effects from the third stage to compute a new set of residuals. In the fourth stage, we apply a non-linear least squares estimator to the residuals to estimate the parameters of the individual shock process. Our estimates at this stage strongly suggest that the error process is i.i.d. We also re-estimated the entire system without instrumenting for the lags of hours and wages and also find no evidence that the shock to human capital is autocorrelated.

Estimation of spousal time transfers: PSID and ATUS

To arrive at the results reported in table 5 in section 3.2, we start by imputing measures of *nll* time to our PSID sample. From the ATUS sample, we calculate housework as a share of total *nll* activity for individuals at different ages, genders, urban/rural classification, education, and number of children. We then use the reported annual hours of housework in the PSID and the demographics-coefficients to impute total *nll* time (i.e. \tilde{h}) to the entire PSID sample.

Next, we use the assumption that $\tilde{h}_i = \bar{h}_i$ for singles to calculate the determinants of \bar{h} for the PSID sample of unmarried men and women. Regressors in this stage again include age terms, education, gender, number of children and urban/rural classification, as well as number of pre-school age children. This allows us to assign a measure of \bar{h} to married individuals in the PSID sample, based on the assumptions that neither \bar{h} nor the share of housework in total *nll* activity varies directly with marital status (though both vary strongly by number of children, which is highly correlated with marital status).

Finally, we use our imputed/estimated measures of \bar{h} and \tilde{h} to calculate Δh_i for married individuals, using the identity $\Delta h_i = \tilde{h} - \bar{h}$. We derive estimates of a , b , α and β by applying a nonlinear systems of equations estimator to (35), using our constructed Δh_i and the couple's wage ratio \bar{w} as the dependent and independent variables. The sample of PSID couples used at this last stage meets the following conditions: (1) both members worked at least 200 hours in the previous year, (2) neither member reported a disability, (3) the husband is at least 45, and (4) \bar{w} falls between 0.25 and 4. The first two restrictions ensure that we are dealing with an interior solution and that the wage rate is an accurate measure of the consumption-leisure

trade off. The restriction to older couples is intended to minimize the possible bias introduced by the presence of human capital considerations in the FOCs, which are likely to be largest for young workers.

Single-shock “basic” model parameter estimates from the SLID

Table 12: Wage Process for the One-Shock Idiosyncratic Risk Model: Parameter Estimates

$(\ln(w_{it}))$	Male	Female
age_{it}	.172 (.021)	.245 (.028)
age_{it}^2	-.002 (.0005)	-.006 (.0006)
age_{it}^3	0.00002 (.00001)	.00005 (.000005)
$I_{\delta_n > 1}$	-.145 (.018)	-.131 (.023)
σ_{η}^{2E}	.021 (.004)	.018 (.005)
ρ	.944 (.007)	.957 (.006)
σ_{ρ}^{2E}	.019 (.003)	.015 (.002)
$cons$	1.44	1.21
$AdjR - sq$.384	.321
n	30879	31056

References

- ABOWD, J. M., AND D. CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57(2), 411–45.
- ATTANASIO, O., AND S. DAVIS (1996): “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy*, 104, 1227–62.

- ATTANASIO, O., AND N. PAVONI (2007): “Risk Sharing in Private Information Models with Asset Accumulation: Explaining the Excess Smoothness of Consumption,” NBER Working Papers 12994, National Bureau of Economic Research, Inc.
- ATTANASIO, O., AND J.-V. RÍOS-RULL (2000): “Consumption smoothing in island economies: Can public insurance reduce welfare ?,” *European Economic Review*, 44, 1225–58.
- AU, D. W. H., T. F. CROSSLEY, AND M. SCHELLHORN (2005): “The Effect of Health Changes and Long-Term Health on the Work Activity of Older Canadians,” *Health Economics*, 14(10), 999–1018.
- AUTOR, D., AND M. DUGGAN (2003): “The Rise in the Disability Rolls and the Decline in Unemployment,” *Quarterly Journal of Economics*, 118(1), 157–205.
- BLUNDELL, R., AND I. PRESTON (1998): “Consumption Inequality And Income Uncertainty,” *The Quarterly Journal of Economics*, 113(2), 603–640.
- BOLDUC, D., B. FORTIN, F. LABRECQUE, AND P. LANOIE (2002): “Workers Compensation, Moral Hazard and the Composition of Workplace Injuries,” *Journal of Human Resources*, 37(3), 623–652.
- BROWNING, M., P.-A. CHIAPPORI, AND V. LECHENE (2006): “Collective and Unitary Models: A Clarification,” *Review of Economics of the Household*, 4(1), 5–14.
- BURKHAUSER, R., J. BUTLER, AND G. GUMUS (2004): “Dynamic Programming Model Estimates of Social Security Disability Insurance Application Timing,” *Journal of Applied Economics*, 19, 671–685.
- BURKHAUSER, R., M. C. DALY, A. H. HOUTENVILLE, AND N. NARGIS (2002): “Self-Reported Work Limitation Data: What They Can and Cannot Tell Us,” FRBSF Working Paper 2002-22.
- CAMPOLIETI, M. (2002): “Disability and the Labor Force Participation of Older Men in Canada,” *Labour Economics*, 9(3), 405–432.
- CAMPOLIETI, M., AND H. KRASHINSKY (2006): “Disabled Workers and Wage Losses: Some Evidence from Workers with Occupational Injuries,” *Industrial and Labour Relations Review*, 60(1), 120–138.
- CAMPOLIETI, M., AND J. LAVIS (2000): “Disability Expenditures in Canada 1970-1996: Trends, Reform Efforts and a Path for the Future,” *Canadian Public Policy*, 26(2), 241–264.
- CHARLES, K. K. (1999): “Sickness in the Family: Health Shocks and Spousal Labor Supply,” Unpublished manuscript, University of Michigan.

- (2003): “The Longitudinal Structure of Earnings Losses Among Work-Limited Disabled Workers,” *Journal of Human Resources*, 38(3), 618–646.
- COYLE, C. C. (2004): “Health Shocks and Couples’ Labor Supply Decisions,” NBER Working Paper.
- DOMEIJ, D., AND M. FLODEN (2005): “Labour-supply elasticity and borrowing constraints: Why estimates are biased,” *Review of Economic Studies*, 9, 242–262.
- GALLIPOLI, G., AND L. TURNER (2008): “Disability in Canada: A Longitudinal Household Analysis,” UBC Working paper.
- GRONAU, R. (1973): “The Intrafamily Allocation of Time: The Value of the Housewives’ Time,” *The American Economic Review*, 63(4), 634 – 51.
- (1977): “Leisure, Home Production and Work. The Theory of the Allocation of Time Revisited,” *Journal of Political Economy*, 85(6), 634 – 51.
- GRUBER, J. (2000): “Disability Insurance Benefits and Labor Supply,” *Journal of Political Economy*, 108(6), 1162–1183.
- HEATHCOTE, J., G. VIOLANTE, AND K. STORESLETTEN (2008): “The Macroeconomic Implications of Rising Wage Inequality in the United States,” *Working Paper*.
- IMAI, S., AND M. P. KEANE (2004): “Intertemporal Labor Supply and Human Capital Accumulation,” *International Economic Review*, 45(2), 601–641.
- KNOWLES, J. (2007): “Why Are Married Men Working So Much? The Macroeconomics of Bargaining Between Spouses,” IZA Discussion Papers 2909, Institute for the Study of Labor (IZA).
- KOCHERLAKOTA, N. R. (1996): “Implications of Efficient Risk Sharing without Commitment,” *Review of Economic Studies*, 63(4), 595–609.
- KREIDER, B. (1999): “Social Security Disability Insurance: Applications, Awards and Lifetime Income Flows,” *Journal of Labour Economics*, 17(4), 784–827.
- LIGON, E., J. THOMAS, AND T. WORRALL (2002): “Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies,” *Review of Economic Studies*, 69(1), 209–44.
- LOW, H., C. MEGHIR, AND L. PISTAFERRI (2008): “Wage Risk and Employment Risk over the Life Cycle,” IZA Discussion Papers 3700, Institute for the Study of Labor (IZA).
- MACCURDY, T. E. (1985): “Interpreting Empirical Models of Labor Supply,” in *Longitudinal Analysis of Labour Market Data*, ed. by J. Heckman, and B. Singer, pp. 111–155. Cambridge University Press.

- MANSER, M., AND M. BROWN (1980): "Marriage and Household Decision-Making: A Bargaining Analysis," *International Economic Review*, 21(1), 31–45.
- MAZZOCCO, M. (2004): "Saving, Risk Sharing and Preferences for Risk," *American Economic Review*, 94(4), 1169–1182.
- MAZZOCCO, M., AND S. YAMAGUCHI (2007): "Labor Supply, Wealth Dynamics, and Marriage Decisions," University of Wisconsin Economics Department Working paper.
- MEGHIR, C., AND L. PISTAFERRI (2006): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.
- MEYER, B., AND W. MOK (2006): "Disability, Earnings, Income and Consumption," Working paper.
- MILLER, R., AND S. SANDERS (1997): "Human capital development and welfare participation," Carnegie-Rochester Conference Series on Public Policy 46.
- PENCAVEL, J. (2002): "A Cohort Analysis of the Association between Work Hours and Wages among Men," *The Journal of Human Resources*, 37(2), 251–274.
- PISTAFERRI, L., R. BLUNDELL, AND I. PRESTON (2004): "Consumption Inequality, Income Uncertainty and Insurance," 2004 Meeting Papers 215, Society for Economic Dynamics.
- RIOS-RULL, J.-V. (1993): "Working in the Market, Working at Home, and the Acquisition of Skills: A General-Equilibrium Approach," *American Economic Review*, 83(4), 893–907.
- SHAW, K. (1989): "Life Cycle Labor Supply with Human Capital Accumulation," *International Economic Review*, 30(2), 431–456.
- SPECTOR, A. (2006): *Disability and Low Income in the Non-Senior Adult Population*. Research Paper Series: Applied Research Branch of Strategic Policy, HRSDC.
- STEPHENS, M. (2001): "The Long-Run Consumption Effects of Earnings Shocks," *Review of Economics and Statistics*, 83(1), 28–36.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2004): "Consumption and risk sharing over the life cycle," *Journal of Monetary Economics*, 51(3), 609–633.
- TURNER, L. (2008): "Social Security, endogenous retirement and intrahousehold cooperation," UBC Working paper.
- WOOLDRIDGE, J. (1995): "Selection Corrections for Panel Data Models Under Conditional Mean Independence Assumptions," *Journal of Econometrics*, 68, 115–132.