

Household Responses to Individual Shocks: Disability and Labor Supply

Giovanni Gallipoli
University of British Columbia

Laura Turner
University of Toronto *

December 31, 2011

Abstract

How do people respond to idiosyncratic shocks? Using longitudinal data from the Canadian Survey of Labour and Income Dynamics we use information on health status to develop and estimate a life cycle framework which rationalizes the observed labor supply of single men and couples following disability shocks. Two puzzling findings associated with disability onset motivate our work: (1) the almost complete absence of observable ‘added worker’ effects within households and, (2) the fact that mean cross sectional declines in participation and hours worked following disability are larger and more persistent for single than for married men. We argue that these facts are consistent with optimal life cycle behavior when we account for the interaction of two familiar and important mechanisms: first, a dynamic human capital accumulation motive linking wages to labor supply; second, the ability of individuals to select into and out of marriage. Our findings suggest that the inclusion of both mechanisms in a standard life cycle model of the household can rationalize much of the evidence on post-onset outcomes among Canadian men. We also show that introducing the ability of couples to optimally allocate home production duties among themselves further increases the model’s predictive power, and in fact allows us to fully account for the well-known lack of observed (market) added worker effects. This is especially important if, as our data suggests, disability limits individuals’ capacity in the home as well as in the labor market.

*We gratefully acknowledge financial support for this project from the Canadian Labour Market and Skills Researcher Network (CLSRN). The statistical analysis was produced from Statistics Canada microdata. The opinions expressed are those of the authors and do not represent those of Statistics Canada.

1 Introduction

Intrafamily insurance is arguably the most common form of protection most individuals have against welfare losses generated by adverse idiosyncratic shocks. The implications of intrahousehold insurance have spurred much recent research in economics. As suggested by ?, the declining relevance of intrahousehold specialization over the past half century, due in large part to the decline in fertility and improvements in household technology, have made the insurance role of marriage perhaps the primary economic rationale for the formation and maintenance of multi-member households.

This paper examines the implications of intrahousehold insurance for couples facing a particular type of idiosyncratic shock: disability. Disability, in our framework, takes the form of partially anticipated and persistent shocks which change relative and absolute productivity of household members by increasing the time requirement for a given amount of market (or non-market) work. An immediate prediction of the intertemporal household model with multiple members, as surveyed for instance in ?, is that the family unit will substitute labor supply between a newly disabled and a still healthy spouse, thereby mitigating the direct disutility of disability for the affected member and also reducing potential consumption losses for all household members. As a result, conditional on stability of the household unit, we should observe that married men, while disabled, exhibit relatively larger declines in labor supply than otherwise similar, but “uninsured”, single men. Similarly, the wives of disabled married men should exhibit higher labor supply, on average, relative to similar wives of healthy or not-yet-disabled men. Using a longitudinal sample of Canadian men with relatively accurate information on health and disability, we begin by showing that neither of these predictions hold true empirically. Furthermore, we document that their failure to hold cannot be ascribed entirely to average differences in health between married and single men and, also, cannot easily be explained by differential preferences for leisure. The relatively muted cross-sectional implications of disability for labor reallocation within marriage therefore pose a challenge for the standard intertemporal household model and raise important questions about the exact channels through which marital insurance operates.

We argue, however, that these seemingly incongruous patterns are in fact entirely consis-

tent with life cycle theory of multi-member households, and with the implications of marital insurance, once we introduce two very well-known and plausible extensions to the basic intertemporal model. The first is an endogenous process for selection into and out of marriage. Such ‘dynamic collective’ models of the household, in which the continuation of the family depends on neither partner finding it optimal to terminate the current union and re-enter the marriage market, have been shown to account for a variety of interesting labor market and saving phenomena: see, for example ? on labor supply and savings decisions surrounding divorce, or ? for the case of differential saving rates between married and singles. It is intuitively appealing that the correlations of health and other outcomes with marital status are in fact determined through rational optimization by individuals in marriage markets. Of course, the possibility that poor health or disability makes an individual less attractive as a mate has major implications for the quality of intrahousehold insurance and merits serious attention.¹

The second change to the basic model is an endogenous process of life-cycle human capital accumulation through learning-by-doing. Beginning with Becker (1965) and ?, the idea that individuals accumulate market human capital as formal investments, or as a by-product of experience, on the job has become broadly accepted by economists. Several recent papers use models of endogenous on-the-job wage formation to study various labor market phenomena,² Such a model is obviously an appealing candidate for studying the evolution of labor supply and wages following disabling health shocks. In fact, we show that human capital dynamics may be the primary factor in explaining both observed timing of labor supply responses to disability onset and differential labor supplies by disabled men across types of household.

The interaction of endogenous marital formation and endogenous wage determination is sufficient to explain nearly all of the puzzling-at-first-blush evidence on economic outcomes

¹Limited evidence on health-based sorting through divorce is mixed. ? find no evidence that divorce hazards increase following disability shocks to husbands. Singleton (2010), using the SIPP, finds that divorce hazards spike following onset and, interestingly, rise in the years preceding onset, though his estimates are not always statistically significant. Weiss and Wallace (1997) find that earnings shocks do precipitate divorce.

²(? and Belley (2010) parameterize and estimate learning by doing models using U.S. data. ? uses a learning-by-doing model to study the evolution of the male-female wage gap, while ? study its implications for the estimation and measurement of Frisch elasticities. ? suggest that differences in estimated returns to on-the-job experience can account for some of the observed difference in employment and hours worked between the U.S. and Europe.

following disability among Canadian males. In a learning-by-doing model of wages, the need to protect previously accumulated human capital stocks from depreciating gives disabled men a strong incentive to maintain pre-onset levels of labor supply soon after onset, particularly if they do not yet know the extent or permanence of their disability. This effect is strengthened within an endogenous marriage framework. Specifically, since role-reversal of the main and secondary earners within the household may jeopardize the stability of the union, married couples have a strong incentive ex-ante to prevent human capital losses that would arise through weakened labor market attachment. As well, because we limit our analysis to men who do not change marital status during the time we observe them, selection on disability and health status plays a significant role in explaining the results. We document that, consistent with what we see in the data, most of this selection hinges on expectations about disability, with a smaller amount occurring through an increase in the probability of divorce, or reduction in the probability of marrying, following disability shocks. This is consistent both with the idea that marriage constitutes a costly investment (in terms of opportunity cost of searching) and with the existing U.S. evidence on disability and divorce.

Finally, while marital selection and endogenous wages are capable of rationalizing observed responses among men, they cannot fully account for observed spousal responses to disability onset. In particular, our model predicts an eventual increase in spousal participation rates following a husband's disability onset, while the data suggest a decline in spousal participation for wives of long-time disabled men. In the final section of the paper, and as a final 'piece' of the puzzle, we extend the model in one more dimension: by introducing and estimating a process for intrahousehold task sharing (a form of home production) by which couples allocate home production, and other unavoidable non-labor non-leisure ("*nll*") tasks between them, so as to maximize the expected value of their combined disposable time. This final extension allows us to fully account for the absence of an observed spousal labor response to disability onset, or 'added worker effect', in the data. Moreover this allows us to match the distribution of reported disability by type of limitation: specifically, whether the condition is primarily work-limiting or limiting in home production. These results follow from the fact that many second earners find it optimal to absorb part of the time cost of

husband’s disability onset by participating more in home production rather than market work, a sort of hidden or ‘informal’ sector added worker effect. We conclude by reporting microdata evidence from the PSID which is consistent with this mechanism.

The layout of the paper is as follows. Section 2 reports empirical evidence on disability and its observed labor market consequences among Canadian households. Section 3 introduces a basic life cycle model of marriage, augmented with a specific process for health and disability shocks. Health status, broadly defined, affects men’s likelihood of becoming and remaining disabled, while disability itself is a stochastic process of persistent, ‘time-stealing’ shocks. We show that a household model combining endogenous selection into marriage and learning-by-doing human capital accumulation generates analytic predictions which are potentially consistent with observed patterns, while a simpler ‘workhorse’ model with exogenous marriage and wages generates counterfactual predictions. Section 4 introduces the numerical counterpart of the model and describes its estimation and parametrization. Section 5 presents simulation results which broadly confirm the analytical predictions. We document the quantitative effects of selection and learning-by-doing and show that the model can match empirical observations on total hours worked, participation, wages and marital transitions. Finally, in section 6 we consider a final extension to the model, in which couples engage in intrahousehold task-sharing, and examine the implications of this additional mechanism. Section 7 concludes and outlines avenues for future work.

2 Disability and labor supply: empirical evidence

The main data source for our study is the Canadian Survey of Labour and Income Dynamics (SLID), a longitudinal survey of Canadian Households maintained by Statistics Canada. The SLID follows a sample of about 8,500 households per wave, for a period of six years, with the majority of income data taken directly from tax records. A new wave begins every three years, so two waves are always active. We use the 1999-2004, 2002-2007 and the first five years of the 2005-2010 waves. Compared to other income panel studies, the SLID contains detailed information about the *economic* consequences and severity of disability, though,

like most income studies, it provides little information on actual *physical* features of the limitation. Our measures of disability are constructed from responses to a series of questions in the SLID disability module. We classify disability into three broad types. Disability is ‘latent’ if it limits physical activity but does not directly limit capacity at work, school, home or other activities such as transportation or leisure. Disability is ‘work-limiting’ if it limits the respondent at work or in other work- or human capital-based activities such as school or job-search. Finally, disability is ‘home-limiting’ if the individual reports being limited in home-based or other non-work activities such as transportation or leisure.³ Table 1 reports some summary results on the incidence, persistence, and effects of disability among Canadian men from 2002 to 2009.⁴ The first five columns report the incidence of disability by type for men between the ages of 25 and 59—‘prime-age’—for all men and for men disaggregated by marital status.⁵ Columns two through six report the average share of all reported disabilities by type and thus sum to one horizontally. The sixth column reports a measure of disability persistence over the panel: specifically the mean share of periods in which a man makes an affirmative disability report, conditional on reporting a disability at least once and being observed at least five years. We note three important facts from the table. First, more than half of all reported disabilities are both work- and home-limiting, while relatively few are exclusively work-limiting.⁶ Second, the incidence of disability is high: around 20% of prime-age respondents reported some dimension of disability, with the proportion increasing

³In the SLID, the question about limitations “in the home” and limitations in “other activities such as transportation and leisure” are asked separately. Here and in what follows, we combine them into a single category.

⁴There is evidence of an uptick in disability reporting between 1999 and 2002 that may have to do with the administering of the disability questionnaire, which changed permanently beginning in 1999. There is also a small uptick in disability reporting in 2008 and 2009 that may reflect a response to the recession, an example of ‘justification bias’. Studies based on Canadian data, by ? (for subjective health measures) and ? (using disability measures from the National Population Health Survey) find that the effects of self-reported health status on labor market outcomes are likely to suffer, on net, from attenuation or measurement error bias, which leads to underestimates of health effects, rather than justification bias. Similar evidence for the U.S. is discussed in ?. We consider a model of justification bias explicitly in section 5. For our data analysis, we simply assume the effects of justification bias are small and average out over time.

⁵Incidence statistics are weighted using cross-sectional weights provided by Statistics Canada.

⁶We would argue that this result intuitive since many home production activities are quite physical in nature while many work-based activities are not. An office worker with asthma may be unable to mow his lawn or chop firewood at home, but is typically not prevented from working on his computer; though the same case of asthma would limit a construction worker on the job.

in age (not shown). Third, prime-age single men report significantly higher incidence and persistence of disability than married men. They also more likely to report their disabilities as being both work- and home-limiting and less likely to characterize their disabilities as having no effects on economic life (i.e. latent).

Table 1: Incidence and persistence of disability among prime age Canadian men 2002-2009

(%)	Any disability	Latent disability	h -limiting disability	n -limiting disability	h and n -limiting	Frequency of reports in 6 year interval
All	18.5	19.1	16.2	6.9	57.3	.471
Married	17.8	21.0	18.3	6.2	54.5	.528
Single	21.7	14.7	12.9	8.7	63.7	.448

2.1 Labor supply effects of disability onset

To measure the effects of disability onset on predicted labor supply over time from onset, we adopt a methodology based closely on that employed by ?, who in turn build on work by ? and ?. We focus on results from two basic estimating equations:

$$y_{it} = \pi_{\alpha}\alpha_i + \pi_t + \pi_X X_{it} + \mathcal{F}(age_i) + \sum_k \hat{\delta}_k A_{kit} + e_{it} \quad (1)$$

$$y_{it} = \pi_{\alpha}\alpha_i + \pi_t + \pi_X X_{it} + \mathcal{F}(age_i, \eta_i) + \pi_X X_{it} + \sum_k \tilde{\delta}_k A_{kit} + \sum_k \bar{\delta}_k A_{kit} \eta_{it} + e_{it} \quad (2)$$

In section 2.1 below, we report results where y_{it} measures own or spousal average weekly hours worked during the reference year or an indicator of own or spousal participation. X contains time-varying demographic and life cycle information including household size, number of children, a dummy for living in a city of at least 50,000, provincial minimum wage, a measure of self-assessed health, an indicator for having ever worked full time, and regional dummies.⁷ α_i contains fixed individual variables including education, dummies indexing SLID panel, and indicators for whether the individual is a visible minority, an aboriginal or a first-generation immigrant, as well as time averages of all covariates in X for each

⁷In robustness regressions we ran specification 1 for married men with and without children and found no systematic differences in estimated responses.

individual over the years we observe him. Time effects, π_t , are captured with year dummies. The marital status index η is equal to 1 for married and 0 for non-married. Importantly, we limit the sample to individuals for whom η is unchanged over the years he is observed. $\mathcal{F}(age)$ is a cubic function, with each term also interacted with η in (2). Index k , ranging from -1 to +10, denotes the number of years elapsed from initial disability onset (observations prior to one year from onset are the control group), and A is an indicator variable indexing k .

We report OLS estimates of (1) and (2) with errors clustered at the individual level. It is important to emphasize that we focus here on reduced-form (cross sectional) patterns. The six-year panel dimension, and the relatively small sample of single men, makes within-effects estimation of long-run responses difficult.⁸ A second, more potentially problematic, issue with the data is that the short panel dimension forces us to rely on individuals' reports of the duration of current disabling condition—the number of years they had the condition before reporting it—to identify individuals at long durations from onset. While duration reporting is typically internally consistent among men at all years from onset, we are still vulnerable to selection bias since men are only asked about the duration of their conditions as a follow-up question conditional on reporting a *currently* active disability. This will lead us to miss those men at longer durations from onset whose disabilities are temporary, or those who fully recover and cease to report them after a few years. To correct against this potential bias, we re-weight the data to create a post-onset sample that approximates as closely as possible the post-onset population of Canadian men based on information in the National Health and Population Survey (NPHS), which relies on nearly identical disability questions but for which we observe the same sample of men biannually for twelve years. These sample adjustments, and various tests of their validity, are described in detail in appendix A.

Labor supply following disability. Figures 1 and 2 report mean differences in labor supply following disability onset ($\widehat{\delta}_k$), estimated from equation (1), for married (including

⁸This is especially true since, as shown in ?, the major effects of disability are concentrated among the subgroup of men whose disabilities become chronic and severely work-limiting. We are not able to identify a man as belonging to this group if we observe him only at, or around, time of onset. In unreported robustness checks, however, we demonstrate that our model replicates fixed-effect estimates of post-onset changes in labor supply as well as the cross-sectional estimates reported here.

common law) and unmarried (never-married, separated/divorced and widowed) men, and for spouses of married men. The left panel of each figure reports post-onset differences in total annual hours worked, divided by 50 to give an average weekly level. The middle and right panels of each figure decompose the total hours difference into differences in (annual) participation and hours conditional on participation pre- and post-onset. In each figure, large dots denote estimates which are significant at the 5% level, whereas small dots indicate significance at the 10% level. The grey dashed lines in Figure 1 plot *differences*, $(\bar{\delta}_k)$ estimated from equation (2), between post-onset differences in labor supplies of married and single men, i.e. a marital status difference-in-difference.

Figure 1: Male labor supply by year from onset onset by marital status




Figure 2: Spousal labor supply by year from husband's onset




Figure 1 shows that cross-sectional declines in labor supply are persistent and increasing by reported year from onset, but that relative declines are much larger for single men than for married men, with the difference-in-differences becoming significant at 10% around five (three) years post-onset for mean weekly hours (participation), and becoming larger and more

significant further from onset. For both married and single men, the post-onset difference in labor supply is concentrated on the extensive (participation) margin. Differences in hours conditional on working exist, and larger for single than for married men, but the effects are muted and only intermittently significant, and dissipate over time from onset for married men.

Figure 2 demonstrates an absence of any meaningful added worker effect for wives of disabled men, either in terms of participation or hours conditional on working. In fact, there is a significant negative added worker effect on the participation margin at very long durations from onset, though this effect disappears when the recession years 2008 and 2009 are dropped from the sample. The finding of negligible added worker effects is consistent with U.S. findings by ? using the HRS and ? using the PSID.⁹

2.2 Other post-onset patterns

We plot two additional results that are of potential interest in formulating a dynamic model of disability and labor supply. The first is the evolution for observed ln hourly wage following onset, which we show separately for single and married workers in figure 3. For neither type of worker is there an observable decline in wages following onset; for single men, the observed wages of the post-onset sample are actually higher. This is likely, in good part, a result of selection: since time from disability onset reduces participation rates, the wages of many of the post-onset men, including the most seriously disabled, are not observed. Nevertheless, the complete absence of an observable wage effect poses a challenge for model of learning-by-doing human capital that we propose in the next section.¹⁰

Second, in figure 4, we plot the cross-sectional evolution of the likelihood of divorcing on a sample of married or newly divorced males by year from onset (left panel) and the

⁹However ? reports PSID evidence that long-run added worker effects after job displacement are significant, eventually replacing 25% of husband's lost annual income, a fact we will argue is consistent with our model's predictions when home production is endogenous.

¹⁰In a previous draft of this paper, we also presented selection-adjusted wage results using measures of household composition and age as instruments. These selection-adjusted results offer some evidence that, in fact, average wages of single males, but not married males, do fall substantially post-onset, on the order of ten to twenty percent over ten years. However, the difficulty in finding strong instruments, especially for single men, leads us to omit these results. They are available from the authors.

cross-sectional evolution of the likelihood of marrying on a sample of single or newly married males (right panel) by year from onset. The results suggest that the probability of divorcing rises significantly at onset and remains elevated for around four years. This result is broadly consistent, though modestly stronger, than recent estimates of the effect of work-limiting disability on divorce, reported by ?, who finds that divorce hazards rise in the two years following disability onset, with the effect strongest for young and well-educated men. By contrast, we find no effect of disability onset on the likelihood of marrying in cross-section, though it is unlikely we would identify a modest effect given the relatively small sample size of single men and the fact that marriage after age 30 is relatively less frequent.

Figure 3: Male ln wages by year from onset

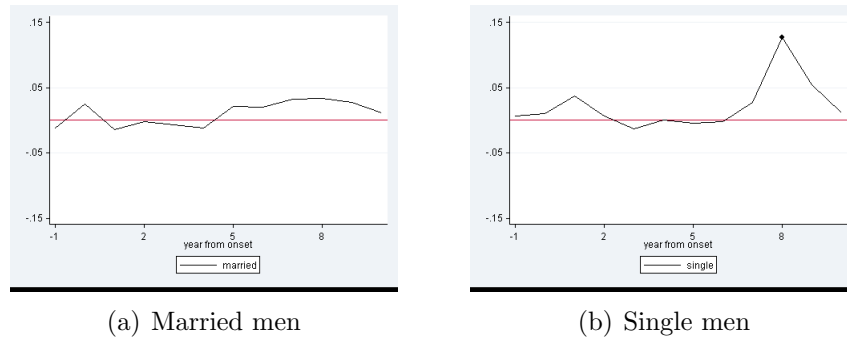
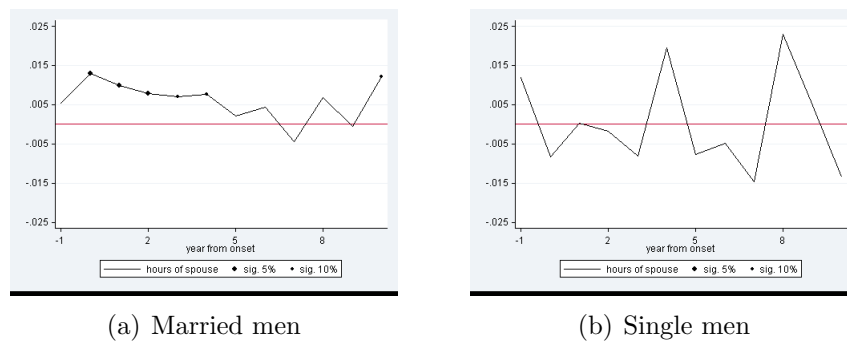


Figure 4: Probability of marrying and divorcing by year from onset



3 A household model and two puzzles

In this section we develop a simple life cycle model with both selection into marriage and learning-by-doing human capital accumulation, and demonstrate that its predictions are consistent with the empirical evidence reported above. In section 3.1 we set up the life cycle problem for men as they move through the married and single states. In section 3.2 we examine predictions from this model when men are confronted with an explicit process of time-stealing disability as well as productivity risk. We show that the analytical predictions of this model are consistent with the evidence on the labor supply of men across different marital states, but become far less consistent with the data once either marital selection or wage endogeneity is shut down.

3.1 A simple life cycle model of marriage

Individuals begin adulthood as single and, over the life cycle, live by turns as singles and as members of married households.¹¹ The matching and separation process is fully endogenous; however, the main focus of this paper is on understanding the responses to shocks exhibited by agents of a *given* marital status, rather than on the broader issue of how marital status evolves in response to shocks, which we explore in related work. At each point in the life cycle, a man is described by his age j and a state vector x indexing non-human wealth, productivity and health. At the start of each period, a single man meets a marriageable partner—that is, a person whom he both wishes to marry and who wishes to marry him—with probability q . We assume perfect assortative mating by age and education, but otherwise the partner is drawn from a known distribution of wealth and productivity, $X_f^S(j)$.¹² Household-level wealth a_M at time of marriage is equal to the sum of both members’ personal assets a and a_f where f subscript denotes a female variable. Denoting the x -conditional expectation operator as \mathbb{E} , an age- j single male has value function

¹¹Our focus on males is not based on the assumption that men and women are inherently different in their economic responses to disability. However, we are mainly interested in household responses following shocks to main-earners. In most prime-age Canadian married households (75% of which are dual-earner) the husband is still the primary earner, a fact replicated in the model.

¹²Education subscripting is omitted throughout for notational clarity.

$$V^S(j, x) = \max_{\{c, l, a'\}} \left(u(c, l) + \beta \mathbb{E}[(1 - q)V^S(j+1, x') + q \int_{X_f^M} V^M(j+1, x'_M(x_f, x')) dX_f^M | j, x] \right) \quad (3)$$

where β is a composite individual discount rate, x (x_f) is the vector of a man's (matched female partner's) characteristics; x_M indexes the characteristics of a married household (the union of the individuals' productivity and health states, plus a_M). The 'primes' denote next-period values. $X_f^M = X_f^M(j+1, x')$ is the subset of single women within $X_f^S(j+1)$ such that, for any $x_f \in X_f^M$, a match results in marriage given x' . V^M is the value of marriage conditional on a couple's x_M . For any possible match, a marriage occurs if and only if:

1. $V_f^M(j, x_M(x_f, x)) > V_f^S(j, x_f)$
2. $V^M(j, x_M(x_f, x)) > V^S(j, x)$

Therefore, $q = q(j, x, a') = q^* \frac{\sigma_{j+1}^S X_f^M(j+1, x')}{X_f^S(j+1)}$, where q^* is a state-invariant probability that a man meets *any* woman, and σ_{j+1}^S is the relative share of the age $j+1$ female population that is single.¹³ The value function $V_f(j, x_f)$ is the female analogue of (3), with $q_f(j, x_f, a'_f)$ similarly depending on a woman's current state vector, intertemporal choices, expected evolution of x_f conditional on choices, and the distribution of single men. In choosing whether to marry a current match, men and women consider their future marriage prospects and the (stationary) distribution of matches in the population.

Next, we turn to married households. Members of a married household maximize a joint objective function U ,

$$U(j, x_M, \lambda, \theta^f, \theta) = \max_{\{c_f, l_f, c, l, a'\}} \left((1 - \lambda)V^M + \lambda V_f^M \right) \quad (4)$$

for a given 'marriage contract' λ , specifying the relative weight of the wife's individual value function in the joint maximization problem. The husband's (wife's) non-pecuniary value of marriage is given by θ (θ^f) which enters utility additively in every period he (she) is married.

A married man's individual value function is

¹³At age 19, $\sigma_j^S = 1$ since all individuals enter adulthood as singles. Thus, the probability of meeting a suitable partner q changes over the life cycle with the share of the population that is single at each age.

$$V^M(j, x_M) = u(\mathbf{c}, \mathbf{l}, \theta) + \beta \mathbb{E} \left[\varpi \max \{ V^M(j+1, x'_M), V^S(j+1, x'(x'_M)) \} \right. \\ \left. + (1 - \varpi) V^S(j+1, x'(x'_M)) | j, x_M \right] \quad (5)$$

where $\varpi = \varpi(j, x_M, a'_M)$ and $(1 - \varpi)$ is the probability that the marriage terminates without husband's consent. This could be due to wife's death, to an exogenous separation shock, or to the wife no longer finding it optimal to be married given x'_M , that is, $V_f^S(j+1, x'_f(x'_M)) > V_f^M(j+1, x'_M)$. The husband may also independently opt to separate from his wife if next-period value of re-entering the marriage market is higher than the continuation value of the current marriage. His post-divorce state vector $x'(x_M)$ includes half the total assets of the married household. Note that in equation (5) there is no maximization: the individual basket of consumption and leisure (\mathbf{c} and \mathbf{l}) and household savings \mathbf{a}'_M are the policy functions that solve the joint problem in (4).

To keep the joint allocation problem tractable we set $\lambda = .5$ so that marriages allocate utility in an egalitarian way. This fairly common assumption substantially simplifies the numerical implementation of the couples' problem. Nevertheless, in a more sophisticated model of marriage, λ would be determined in an initial period through bargaining and then would evolve over time to prevent 'inefficient' divorces from occurring (see ?, ?).¹⁴

¹⁴That is, the λ in the initial period of marriage (λ^*) could be the solution to the following cooperative (Nash) bargaining problem:

$$\lambda^* = \arg \max_{\lambda} S(j, x_f, x | \lambda, \theta^f, \theta) \\ \text{s.t.} \\ S(\cdot) = [V^M(j, x_M(x_f, x) | \lambda, \theta) - V^S(j, x)] [V_f^M(j, x_M(x_f, x) | \lambda, \theta^f) - V_f^S(j, x_f)] \\ V^M(j, x_M(x_f, x) | \lambda, \theta) - V^S(j, x_M(x_f, x)) \geq 0$$

where $S(\cdot)$ is the product of the partners' individual surplus from marrying, conditional on at least one partner's surplus (here, the male's) being positive. The evolution of λ has interesting implications for the insurance value of marriage among spouses over the life cycle. However, since we estimate the model to generate the empirically observed share of divorces, the first-order effect of introducing this more sophisticated marriage model in this paper would be on the estimated values of θ and θ_f . However, in other work we return to directly examine the effects of evolving λ on labor supply following health shocks.

3.2 Disability, optimization and comparative statics

Next, we focus on the solution of the intertemporal problems described above in a world in which individuals are subject both to productivity and time-stealing disability shocks. We first explicitly define the personal state spaces x , x_f and x_M . Vector x is given by $\{a, w, i, \delta_h, \delta_n\}$, x_f by $\{a_f, w_f\}$, and x_M by $\{a_M, w_f, w, \delta_h, \delta_n\}$.¹⁵ As before, a denotes household-level assets and is a function of households' saving choices. w is the current wage commanded in the labor market, assumed to be first-order independent of $\{rsk_i, \delta_h, \delta_n\}$, which index health and disability status. Health index rsk does not limit a man directly, but captures the risk that he becomes and/or remains disabled. The δ 's are the disability factors. They operate as multiplicative factors which 'steal' time from individuals by increasing the total amount of time required to complete a given activity: 'labor limiting' (δ_n) shocks increase the amount of time required to complete a unit of market work, while 'home-limiting' (δ_h) shocks increase the time required to complete a unit of non-market, home-based *nll* activity.¹⁶ At each point in the life cycle, a married couple solves (4) subject to constraint set

$$\begin{aligned}
 \xi_1 : \frac{(T - \delta_h \bar{h} - l)}{\delta_n} w + (T - \bar{h}_f - l_f) w_f + (1 + r) a_M + b(\cdot) - c_f - c - a'_M &= 0 \\
 \xi_{2a} : l \leq T - \delta_h \bar{h} & \qquad \qquad \qquad \xi_{2b} : l_f \leq T - \bar{h}_f \\
 \xi_{3a} : \varpi = \varpi(a'_M; \delta_h, \delta_n, w_f, w) & \qquad \qquad \qquad \xi_{3b} : \varpi^f = \varpi^f(a'_M; \delta_h, \delta_n, w_f, w) \\
 \xi_{4a} : \mathbb{E} w' = H(l, h | w, \delta_h, \delta_n) & \qquad \qquad \qquad \xi_{4b} : \mathbb{E} w'_f = H^f(l_f, h_f | w_f)
 \end{aligned} \tag{6}$$

where ξ_1 is the shadow value of the household budget constraint, ξ_2 are the shadow values of the couple's time constraints, ξ_3 are the shadow values of the probability of avoiding a forced separation in the next period, and ξ_4 are the shadow values of the learning-by-doing human

¹⁵For numerical tractability, and for reasons discussed in footnote 11, only men are assumed to be subject to a random health process. Therefore x_f contains only asset holdings and wage.

¹⁶The notion of time-stealing disability is well known in the health literature: see ?. An alternative approach from the economics literature is to model a stock of health (or 'energy') and let people make endogenous investments affecting its status, as in ? and ?. Given non-observability of health stocks, this approach would make the model substantially more complicated to estimate. At the same time it would have very similar implications to the model augmented with intrahousehold task sharing, discussed in section 6.

capital functions for each partner. The function $b(\cdot)$ captures all benefit entitlements as a function of the household's state space and labor choices. T is a fixed weekly time endowment and \bar{h} is the amount of time that must be devoted to *nll* activities, such as errands or work at home. For now we assume that \bar{h} is a fixed requirement; in section 6 we will relax this assumption for married couples. Hours of market work are denoted n .¹⁷ The problem is exactly identical for a single male household, with a replacing a_M , $w_f = \lambda = \xi_{2b} = \xi_{3b} = 0$, and the multiplier ξ_{3a} referring to $q = q(a'; \delta_h, \delta_n, w)$ and giving the shadow value of the likelihood of successfully matching as a function of expected next-period state variables.

The first order conditions for each partners' leisure are given by

$$(1 - \lambda)u_l = \xi_1 \frac{w}{\delta_n} + [\xi_{3a}\varpi_H + \xi_{4a}] \frac{H_n(\cdot)}{\delta_n} \quad (7)$$

$$\lambda u_l^f = \xi_1 w_f + [\xi_{3a}\varpi_{H^f} + \xi_{4b}] H_n^f(\cdot) \quad (8)$$

where $H_n(\cdot) \equiv \frac{\partial H(\cdot)}{\partial n}$ gives the marginal change in the expected next-period wage from a small increase in hours worked this period. Let Ω_{δ_n} give the expression for $\frac{\partial n}{\partial \delta_n}$; Ω_{δ_h} give the expression for $\frac{\partial n}{\partial \delta_h}$; and Ω_{δ}^f give the expression for $\frac{\partial n_f}{\partial \delta_i}$ for $i \in \{n, h\}$. Then the responses can be written as follows:

$$\begin{aligned} \Omega_{\delta_n} &= A + B_{\delta_n} + C_{\delta_n} + D_{\delta_n} + E_{\delta_n} \\ \Omega_{\delta_h} &= B_{\delta_h} + C_{\delta_h} + D_{\delta_h} + E_{\delta_n} \\ \Omega_{\delta}^f &= B_{\delta}^f + D_{\delta}^f + E_{\delta}^f \end{aligned} \quad (9)$$

where

1. $A < 0$ is an intertemporal substitution effect
2. $B > 0$ is the wealth effect
3. $C < 0$ is the 'time-loss' effect
4. $D \gtrless 0$ is the human capital effect

¹⁷We omit the constraint on borrowing, $a \geq \underline{a}$. In our numerical implementation we set \underline{a} relatively low to allow for borrowing up to roughly 1.5 times average household income. The borrowing interest rate is estimated through model simulation and lies above the saving interest rate.

5. $E \gtrless 0$ is the marital stability effect

Detailed analytical expressions given our assumed separable preference functions, and a technical exposition of the various effects, are provided in appendix B.¹⁸ We begin by focussing on the first three components, which take the same form regardless of whether or not wages or the probability of marital dissolution are endogenous. A and B give, respectively, the usual intertemporal substitution and wealth effects of a negative shock. There is no substitution effect associated to home-limiting disability in Ω_{δ_h} : a δ_h shock reduces household resources and restricts the choice over time allocations, but does not affect relative prices of labor and leisure. C is a ‘time-loss’ effect reflecting the reduction in disposable time experienced by disabled individuals, the major way in which δ_n shocks differ from standard productivity shocks. The net effect of either type of disability on male labor supply is technically ambiguous since the time-loss and substitution effects work in opposite directions from the wealth effect. In general, however, we would expect the time-loss effect to dominate for large shocks, and for $|A_{\delta_n} + B_{\delta_n} + C_{\delta_n}| > |B_{\delta_h} + C_{\delta_h}|$ for shocks of similar magnitude and following similar processes.

There are three major implications of the expressions in (20) when D and E are ignored. The first is that disability onset has the strongest negative effect on labor supply immediately after a shock is realized, since substitution and time-loss effects are temporary, lasting only while $\delta > 1$; on the other hand, the potential change in ξ_1 (wealth effect) has positive effects on n and n_f that outlast the shock itself. Second, conditional on the husband’s shock being at least partially unexpected and sufficiently large, we should observe an increase in wife’s labor supply—an ‘added worker effect’—since $C_{\delta}^f > 0$. Third, and finally, labor supply should fall more, on average, for married men than for single men following a shock of similar magnitude and persistence. This prediction arises from the insurance provided by

¹⁸To derive and interpret the analytical expressions, we maintain two assumptions: (1) $\frac{\partial H(n;w)}{\partial \delta}|_n \equiv H_{\delta}(\cdot)|_n = 0$, i.e. there is no direct partial effect of disability on the evolution of wages so long as effective labor supply is held constant. This is consistent with our estimates of equation (11) described in section 4 and reported in appendix C.1. (2) $\frac{\partial \varpi^f}{\partial w_f} \equiv \varpi_{H^f}^f = \frac{\partial \varpi^f}{\partial w'} \equiv \varpi_H^f = 0$. This assumption implies that women are the main drivers of divorce, and that the likelihood of a man choosing to divorce his wife is independent of his own and of his wife’s wage. This assumption does not hold explicitly in the model, but is consistent with the fact that $\theta > \theta^f$ in SMM estimation of all our models, as reported in table 3.

marriage: since husbands' potential earnings account for a smaller share of total household resources, we should expect the wealth effect B to be smaller for marrieds. This implication is further strengthened when considering that married households typically have higher wealth-to-income ratios (about 70% higher at the peak of the life cycle in Canadian data) than single households. All of these predictions are counterfactual to the patterns reported in the previous section, implying a role for the two major mechanisms of our model: marital selection and endogenous wages.

We next consider the role of endogenous human capital, effect D . The overall effect of a shock on the necessity and ability to accumulate human capital turns out to be ambiguous, as discussed in technical detail in appendix B. However, it is easy to see how an endogenous wage process with depreciation improves the prediction of the model taken across the full population of men. The desire to protect the existing human capital stocks from excessive depreciation gives all men at or soon after onset a strong incentive to forego very large drops in labor supply, even in the face of large immediate utility costs. This incentive will decline at later years from onset for the subset of men whose disability worsens as their human capital stocks decline. Therefore, we can predict that the *average* post-onset profiles of labor supply and participation should look more like the profiles shown in figure 1.

Intuitively, the role of endogenous human capital can also go some ways toward accounting for difference-in-differences in labor supplies between married and single men. Married men typically have higher wages (larger human capital stocks) than single men, which will increase their relative intertemporal incentive to suffer low utility from leisure today in order to protect their stock. However, the ability to accumulate human capital also makes it easier for wives to transition into main earner status by increasing their own labor supply—and wages—over time, relieving chronically disabled husbands of the need to work. Indeed, the results we report in section 5 suggest both these forces—selection and substitution within stable marriages—are indeed at work in the model.

Finally, we consider the role of marital selection on post-onset labor supply: effect E . Its predictions for the model are two-fold. First, if couples value their marriages but wives are likely to divorce husbands with low human capital, then this further strengthens the couple's

incentive to maintain a husband’s labor supply the face of a (potentially) temporary disability shock.¹⁹ Second, over time, a selection effect can emerge by which severely disabled, and consequently wage-poor, men are more likely to remain single or change marital status, which will also improve the model’s ability to match the data.

An obvious caveat is that the predictions made directly above refer to married and single men’s reactions to similar δ shocks, following similar *risk* processes. In our data, married and single men face clearly different disability risk processes, as documented in section 2. However, in section 5, we show that these differences *alone* are insufficient to fully explain the data when marital status evolves exogenously and health depends on marital status. Furthermore, in section 6 we show that the observed differential health outcomes arise endogenously through the marriage and separation process.²⁰

4 The numerical counterpart of the model

The interaction of human capital accumulation and selection into marriage provide a channel by which the data patterns described in section 2 can be reconciled with standard economic theory. However, since the analytical predictions are suggestive, and in most cases ambiguous, we require a numerical implementation to assess the model’s ability to rationalize the cross sectional evidence. Next, we turn to describing our numerical implementation in detail.

The simulated economy is populated by one-member (single) and two-member (married) households who solve the period-by-period optimization problem described in section 3. The individual state space for males is $x = \{a, w, risk, ds\}$ where *ds* is a discrete set of combinations of $\{\delta_h, \delta_n\}$. Given the large number of parameters needed for a numerical implementation, we proceed in two steps. First, a number of parameters are estimated directly from microdata. These estimates are then employed to numerically simulate the model and estimate the remaining parameters through simulated method of moments (SMM). We use

¹⁹A similar effect operates on single men, but many singles in their 40s and 50s already face relatively low likelihood of remarriage, particularly if selection effects are operational.

²⁰Another possible concern regards the exogeneity of health shocks. However, the existing evidence of direct effects of income on health is weak (for a discussion see ?). Using HRS data, ? shows instead that, among otherwise similar people, those suffering health events drop out of the labor force more frequently and work fewer hours, resulting in long-term earnings’ losses.

a downhill simplex method to minimize the sum of squared percentage deviations between simulated and empirical targets.²¹

Demographics, marital evolution and preferences. A model period is one year. Individuals enter the model as singles at age 19, work to the age of 65 and live to a maximum age of 99. Education, discretized to be either high or low, is exogenously assigned. The matching function for singles is exactly as described in section 3.1, with two additional restrictions: (1) an individual is only matched with partners whose non-human wealth is at least one quarter and not more than four times their own;²² (2) once married, couples face an exogenous divorce probability of \lceil per period, which allows some share of divorces to occur for “non-economic” (or at least unobservable) reasons.

Effective discount rates are age-varying, the product of an invariant gender-specific geometric discounting factor $\{\mathcal{B}_f, \mathcal{B}\}$ and age/gender-specific survival probabilities taken from Canadian vital statistics for individuals over 60. Households save at real interest rate $r = 3.8\%$. Households age 65 and younger can also borrow at an interest rate of r_B which is estimated in the model.

Individual period utility functions at age j are given by

$$\begin{aligned} u_j^f &= (1 - \gamma_j^f) \frac{(c_f \tilde{n}_\eta)^{1-\omega}}{1-\omega} + \gamma_j^f \frac{l^{1-\psi}}{1-\psi} + I_\eta \theta^f \\ u_j &= (1 - \gamma_j) \frac{(c \tilde{n}_\eta)^{1-\omega}}{1-\omega} + \gamma_j \frac{l^{1-\psi}}{1-\psi} + I_\eta \theta \end{aligned} \tag{10}$$

for women and men respectively, where ω and ψ govern the intertemporal elasticity of consumption and labor, respectively; \tilde{n} captures changes in marginal utility of consumption due to economies of scale in marriage; and I_η is an indicator function for being married. We set $\omega = \psi = 1.5$, which results in men’s Frisch elasticities broadly consistent with empirical

²¹As the model is exactly identified, we iterate over the parameter space until the mean percentage difference between simulated and empirical targets drops to approximately .05%. Using starting parameter values that generate an average 5% difference, this typically requires around 200 iterations. Each iteration requires approximately ten minutes to run on an Xion 360 processor.

²²For those in debt or with very low assets, matches occur only with partners whose wealth levels differ by less than approximately the mean household income in the model. This matching restriction substantially reduces the computational burden and has little effect on the distribution of successful matches.

evidence.²³ For couples \tilde{n} is set to .85, following the OECD equivalence scale. The remaining parameters are estimated through SMM as described in section 4.1. The relative preferences for leisure, $\{\gamma^f, \gamma\}$, follow separate cubics in age for men and women. Age-varying γ captures changing consumption needs due to deterministic fertility,²⁴ and allows us to generate reasonable life cycle profiles of hours worked in the presence of endogenous human capital.

(*nll*) time requirements. We calibrate individuals’ ‘*nll*’ time costs or requirements (\bar{h} and \bar{h}_f), using data from the public use 2005 Canadian General Social Survey (GSS), which provides us with which provides time diary data for a representative cross-section of Canadians using a methodology similar to the American Time Use Survey. Our set of *nll* activities is very close to those activities, other than market work, excluded from ‘Leisure Measure 2’ in ?²⁵ and includes all traditional home production activities, yard and vehicle maintenance, procurement of services, and child care other than play. In the models, we allow *nll* requirements to evolve separately for high- and low-educated men and women by age and the presence of children. Table 2 reports means and standard deviations of weekly *nll* hours of prime-age individuals by gender and marital status.²⁶ Both women and men spend substantially more time on *nll* tasks when married than when single.

Disability process. The process of disability is summarized by two state variables: *ds*,

²³In the workhorse model the analytical Frisch elasticity is given by $\frac{1}{\psi} \frac{l}{\delta_n n}$. For a 40-year-old healthy married male working 45 hours per week, this corresponds to a Frisch elasticity of labor of around .83. Various authors, e.g. ?, ? and ?, have shown that conventional estimation methods may underestimate the Frisch elasticity by 50% or more. Our findings confirm this problem: using ?’s method of estimating ψ from the intratemporal optimality condition for married men results in an *estimated* Frisch elasticity of only .27, and uncompensated labor supply elasticity of .16 in a model in which human capital is exogenous. In our full model, the analytical Frisch elasticity is more complicated, depending on ψ , δ_n and other variables. However, the same estimation procedure leads to an estimate of .14 (and an uncompensated elasticity of .10), both within the range suggested in the micro empirical literature on intertemporal labor supply. The estimated Frisch elasticity for married women is much larger at .60, which is intuitively appealing, although there is less of a consensus on its ‘true’ value in the literature.

²⁴All women in the model, whether married or single, have children, which also impose time costs.

²⁵The main difference is that we include personal care, pet care and gardening/grounds maintenance in our definition of *nll*, while ? include them in their Leisure Measure 2 (LM2). We do not put education in *nll*, but we do adjust the total time endowments by gender, age, and education to account for time lost to study early in the life cycle. Finally, Aguiar and Hurst include sleep in LM2, while we exogenize a ‘core’ amount of sleep (42 hours) and consider the remainder to be leisure. Thus, individuals in the model divide 126 hours weekly between *n*, *l*, *h* and education.

²⁶The GSS reports data from diary days which can be summed within cells to give the mean weekly *nll*. However, the standard deviations reported at the weekly level are of limited value. We therefore report them only for illustrative purposes and omit them as moments in the SMM estimation of the model.

Table 2: Time devoted to *nll* task by gender and marital status in 2005: General Social Survey

Men		Women	
Married	Single	Married	Single
26.1 (19.5)	21.1 (17.6)	39.5 (20.6)	32.1 (21.8)

which indexes current disability status, and *rsk*, which captures a man’s underlying risk of becoming, or remaining, disabled. We include *rsk* to capture the idea that, while disability shocks arrive as ‘news’, men differ observably in their *susceptibility* to *ds* shocks at each given age.²⁷ The *rsk* matrix has three states. *rsk* 1 and 2 denote high and low disability risk; men move randomly from *rsk* state 1 into 2 as they age.²⁸ A third *rsk* state captures ‘chronic’ disability, in which a limiting *ds* state becomes permanent.

The *ds* vector comprises six states in ascending order of severity: healthy (*ds* 1); latent disability (*ds* 2); *h*-limiting only (*ds* 3); *n*-limiting only (*ds* 4); *h*- and *n*-limiting, both milder (*ds* 5), and severe (*ds* 6). For non-chronic disabled men, transitions across *ds* states follow an age- and *rsk*-specific Markov transition matrix estimated directly from the SLID. The probability that any limiting disability becomes chronic (*rsk* 3) is 3.7% per period, chosen to match the frequency of disability across six-year intervals in the 2002-2007 SLID panel (row one of column 7 from table 1).

Wage dynamics and human capital accumulation.

To parameterize our learning-by-doing human capital process, we adopt the following equation for evolution of wages:

²⁷? argues that paths of consumption growth among disabled households are consistent with the idea that disability shocks may be partially anticipated, particularly by older workers.

²⁸To estimate the process for *rsk*, we run a probit regression on SLID data in which the dependent variable is an indicator for having a disability during the course of the panel and the regressors include a variety of standard demographic controls, including age terms and self-assessed health. We then split predicted probabilities at the median so that half of the male SLID population is ‘high’ *rsk* and half ‘low’ *rsk*, and these groups are used to estimate the high- and low-risk *ds* matrices described above. The age-dependent probabilities of transitioning into *rsk* 2 from *rsk* 1 are chosen to replicate the shares of men in *rsk* 1 and 2 at ages 20-25, 40-45 and 60-65. By age 66, more than 99% of men have transitioned into *rsk* 2 and about 5% are chronically disabled.

$$\begin{aligned}
H_{it+1} &= \mathcal{H}(n_{i,t}|H_{it}, ed_i, age_{t+1}, X_{it+1}, \nu_{it+1}) \\
&= \kappa_{it} + (\alpha_1 + \alpha_2 n_{it} + \alpha_3 n_{it}^2) H_t + \nu_{it+1} \\
w_{it} &= R_t H_{it}
\end{aligned} \tag{11}$$

where R is the wage rate per unit of human capital, n is average weekly hours worked in the previous year, X is a vector of other observable covariates (including δ_n) and H is start-of-period human capital stock. The individual- and age-specific intercept $\kappa = \{\kappa_0 + \kappa_1 \times ed + \kappa_2 \times age + \kappa_3 \times age^2\}$ approximates the minimum human capital level a person can have, given age and education. Again, the κ 's are estimated in reduced form as part of our SMM procedure. The rate of depreciation of H is given by $(1 - \alpha_1)$ where $\alpha_1 = \alpha_{11} + \alpha_{12} \times ed$ varies across individuals by education level. $\alpha_k = \{\alpha_{k1} \times age + \alpha_{k2} \times age^2\}$, for $k = 2, 3$, governs the rate at which H is replenished through market work. The i.i.d. shock ν is heteroskedastic in age and current human capital stock. Following ?, we let $R_t = R = 1$ for all years in our sample. Estimates of (11) for men and women are reported in appendix C.1, table 10, along with details of the estimation procedure, including controls for selection into the labor market and the likely endogeneity of n . It is worth noting that we find no evidence of direct effects of disability on the human capital production function, which allows us to use current disability status to control directly for selection into work.

Policy environment. Government policy in the model approximates the existing patchwork of Canadian federal and provincial disability, retirement and anti-poverty programs. In every period, working-age households receive a basic transfer (\underline{bb}) capturing unemployment benefits, child benefits, transfers from family members outside the household and other residual public and private transfers. These are targeted for single male, single female and married households as part of the SMM model estimation described in the next section, and their values are reported in table 3. Low-income households with assets below an annual cut-off level of \$20000 (\$30000 for marrieds) receive a minimum-income benefit equal to basic welfare transfers available under Ontario Social Assistance.²⁹ Non-workers over 60 receive

²⁹Information on welfare benefits and eligibility comes from the Canadian Council of Welfare. The minimum asset levels include an adjustment for housing and vehicle wealth. The level of the cutoff has essentially no effect on the results.

a benefit worth 25% of expected weekly earnings based on the wage at age 60, up to the average weekly manufacturing wage of \$22/hour. At 65 all individuals receive a flat-rate transfer equal to the federal Old Age Security (OAS) benefit, and low-income households receive an additional benefit representing the Guaranteed Income Supplement (GIS).

Disability entitlements of working-age men depend on chronicity and type of disability. Men with non-chronic work-limiting disabilities receive a benefit approximately equal to one eighth of 75% of their current expected earnings if they reduce their work time by at least five hours per week below their healthy weekly average. This benefit approximates a temporary workers' compensation (WC) benefit, with average duration of about two months in Canada and available to roughly 80% of workers. Alternatively, any work-limited man can drop out of the labor force and receive a benefit proportional to his wage (up to \$22 per hour) in the last period worked, an approximation to disability insurance under the Canadian Pension Program. We estimate replacement rates (rr) for men in, respectively, ds 4 or 5 and ds 6, to match the average transfers received by prime-age single men in the corresponding ds category in the SLID; the estimated values are reported in table 3. Finally, many Canadian workers receiving temporary WC benefits (including workers in Ontario at establishments with more than twenty workers) are legally entitled to be rehired by their previous employer at a comparable wage once the disability has passed. ? show that workers who return to their previous jobs after suffering a workplace injury suffer negligible wage losses compared to workers who change employers. We capture this fact by allowing disabled workers who would otherwise be full time (at least 35 hours per week) and who reduce their labor supply by less than 20% in the current year to receive a human capital return equivalent to their full time hours.³⁰

³⁰To fund retirement and disability benefits, individuals face a payroll tax of 9.9%, which is the CPP payroll tax rate for covered workers. There is also a progressive income tax with eight brackets approximating the rates and brackets (federal + provincial) for a taxpayer living in Ontario, including deductions for children and a consumption tax of 8%.

4.1 Estimation: Simulated Method of Moments.

In addition to the policy parameters discussed in the previous section, there are five sets of parameters to estimate in the full model. They are: (1) relative preference for leisure $\{\gamma_j^f, \gamma_j\}$ which evolve according to a cubic in age; (2) disability time-costs $\{\delta_h, \delta_n\}$ in each limiting *ds* state; (3) gender-specific time-discount rates $\{\mathcal{B}^f, \mathcal{B}\}$, and borrowing interest rate r_B ; (4) education- and age-specific intercepts for male and female human capital processes $\{\kappa^f, \kappa\}$; and (5) female and male non-pecuniary value of marriage $\{\theta^f, \theta\}$, exogenous matching rate and random divorce hazard $\{q^*, \mathcal{D}\}$.

Identification. The parameters are estimated jointly through SMM. The γ 's and κ 's are primarily identified by the predicted evolution of hours and ln wages over the working life, and by variation in average ln wages by education.³¹ The δ 's in each limiting *ds* state are identified by variation in labor supplies of prime-age men in each state, under the restrictions that $\delta_n(5) = \delta_n(4)$ and $\delta_n(6) = \delta_h(6)$, which is consistent with disability severity reports in the SLID. Time-discount rates \mathcal{B}^f and \mathcal{B} are primarily identified by wealth to after-tax income ratios in households with a female and male member respectively. Borrowing interest rate r_B is pinned down by the debt-to-income ratio for all households with debt. The parameters governing marital formation and divorce are identified by the evolution of marriage observed in Canadian data. Specifically, the parameters adjust so that (a) the mean age of first marriage is 26.5; (b) 71.5% of working-age households are married at any point in time; (c) 1.7% of this stock of marriages terminate each period; and (d) prime-age married men earn on average .18 log points per hour more than their single counterparts. Table 3 summarizes this information and report parameter values for the model. SMM estimates for all other models presented in section 5 and 6 are available from the authors.

³¹Because we are using a data set covering eleven years, we are not able to directly account for cohort effects. Implicitly, we invoke balanced growth path assumptions: real wages and benefits grow over time at a constant rate for every cohort while all other variables, and ratios of male to female or married to single values, do not change across cohorts. Cross-cohort wage growth is accounted for in both the estimation of \mathcal{H} and the targeted paths of lifecycle wages targeted in the model. For men, balanced growth assumptions seem reasonable as the estimated life cycle profile of hours worked and the profile of hours worked by *ds* status used as data targets is virtually unaffected by the inclusion of controls for year and ten-year birth cohort. For women, the assumption is clearly more problematic, which is another reason for focussing the analysis on men. However, changing cohort effects for females do not appear to have any major implications for our results.

Table 3: Estimated parameters

Parameter	Value: Model	Target ^a
γ^f	$.186(1.0 + .100age - 5.57E^{-3}age^2 + 8.48E^{-5}age^3)$	mean wkly hours worked for married women 20-59 ^b
γ	$.364(1.0 + .46E^{-2}age - 6.28E^{-4}age^2 + 6.28E^{-6}age^3)$	mean wkly hours worked for healthy single men 20-59
$\delta_h(3)$	1.53	n for prime-age single men in $ds3^c$
$\delta_n(4)$	1.21	n for prime-age single men in $ds4$
$\delta_h(5)$	1.23	n for prime-age single men in $ds5$
$\delta_h(6)$	1.61	n for prime-age single men in $ds6$
\mathcal{B}^f	0.992	median wealth/income ratio for hhs with a female member
\mathcal{B}	0.958	median wealth/income ratio for hhs with a male member ^d
r_B	$5.39E^{-2}$	median debt/income ratio for all households
κ^f	$-14.7 + .41ed + 1.04age - 2.03E^{-2}age^2 + 1.19E^{-4}age^3$	mean wages by age and education for female workers
κ	$-12.7 + .21ed + 1.22age - 3.18E^{-2}age^2 + 2.35E^{-4}age^3$	mean wages by age and education for male workers ^e
θ^f	$4.11E^{-4}$	share of divorces per 1000 married households
θ	$1.96E^{-2}$	average wage differential between prime-age married and single men
q^*	.745	mean age of first marriage for males
\lceil	.121	share of households that are married ^f
$rr\ ds\ 4,5$.16	average benefit received by prime-age single men in $ds\ 4$ & 5
$rr\ ds\ 6$.26	average benefit received by prime-age single men in $ds\ 6$
b^{fh}	31.3	universal benefit for single women under 65
\bar{b}^m	70.5	universal benefit for single men under 65
\bar{b}^M	84.1	universal benefit for couples under 65

(a) All labor and income targets come from the 2002-2007 cross sectional files of the SLID. Wealth targets are taken from public use version of the 2005 Survey of Financial Security. Time-use targets are taken from the 2004 General Social Survey on Canadian time use. ‘Prime age’ refers to individual and households age 25-59. (b) Hours worked for non-disabled single males age 20-29/ 30-39/ 40-49/ 50-59 are 28.8/ 38.5/ 37.8/ 31.7. For married women, the corresponding hour targets are 24.8/ 26.2/ 28.0/ 21.9. (c) The targets for each ds state are 30.9/ 24.4/ 20.9/ 9.7. (d) ‘Wealth’ includes all financial assets plus the termination value of pension entitlements. Income refers to after-tax income. The median wealth to income ratio for all prime-age hhs with a male (female) member is 3.04 (3.06), and the median debt-to-income ratio for all hh’s is .37. (e) The specific targets are: (1) the raw educational difference in prime-age wages by education category for male (\$6.05) and female (\$5.83) workers; wages for young 25-34 male (\$19.6) and female (\$16.6) workers; wages for young 35-44 male (\$23.1) and female (\$18.3) workers; and wages for older 45-55 male (\$24.7) and female (\$18.7) workers. (f) 1% of marriages (10 in 1000) among households under 66 terminate in steady state, while prime age married men earn \$4.30 per hour more on average than prime age single men. 71.5% of households between 25 and 60 are married. (g) Wage replacement under CPP are received by non-workers in the respective ds state and are targeted to match, in combination with universal and welfare benefits, total public and private benefits of \$97 (\$150) received by prime-age single men in $ds\ 4$ and 5 ($ds\ 6$) (h) Universal benefits are targeted to match, in combination with welfare entitlements, total per-capita public and private benefits received of \$35, \$130 and \$44.1 for prime-age non-disabled single male, single female, and non-disabled married households respectively. They are received regardless of other benefit entitlements and are replaced by GIC and OAS benefits at age 65.

5 Results

We now turn to assessing models' performance in replicating the patterns discussed in section 2, and the implications of human capital and marital selection for the ability of households to handle disability shocks over the life cycle.

Labor supply over the life cycle. Figures 5 and 6 show how the model replicates the mean life cycle hours and wage profiles of single men, married men and married women. Differences in labor supply and wages *across* marital status are not targeted in the model, but emerge endogenously.³² The fit of profiles suggest that the following plots of post-onset labor supply are not driven by spurious life cycle effects.

Figure 5: Life cycle profiles of hours worked





Figure 6: Life cycle profiles of mean wages



Labor supply following disability shocks. Next, Figure 7 plots the central result of this section, showing estimates of the $\hat{\delta}$'s from equation (1) and $\bar{\delta}$'s from (2), using average


³²Profiles of participation give similar fits. The variance of \ln wages is .25 for men and .24 for women.

weekly hours (left panel), participation (middle panel) and hours conditional on participation (right panel) as dependent variables,³³ against the corresponding estimates from the SLID.³⁴ The top panel shows the $\hat{\delta}$'s for the pooled sample of men, the middle panel plots $\bar{\delta}$ from the same sample, and the bottom panel plots $\hat{\delta}$'s estimated on the sample of wives. In all graphs solid lines show simulation results and dotted lines show corresponding results from the SLID. Small nodes indicate where the estimates differ from zero at the 10% level, while large nodes indicate significance at the 5% level. Broadly speaking, the numerical results confirm the analytical predictions from section 3 for the post-onset labor supply of men. The life cycle model with endogenous marriage and learning-by-doing human capital explains both the persistence of the (mean) decline in labor supply across years from onset and the larger difference in post-onset hours of single, as compared to married, men, with the difference concentrated on the participation margin. The major way in which the model does not fit the empirical evidence is for the post-onset labor supplies of spouses. The model predicts a small but noticeable and intermittently significant added worker effect at long durations from onset, on both the participation and intensive hours margins. We return to address this problem of fit that we return to in section 6. In the next subsection, we explore the model's implications in more detail and examine the roles of learning-by-doing human capital and marital selection in generating our results.


³³The simulated weekly hours variable is actually the average of current period hours and lagged hours. This is because, in the SLID population, disability shocks arrive continuously, while in the model they can arrive only at the beginning of a period, leading to a potential mismatch between model and data in the period of time affected by a 'reported' shock.

³⁴To construct model samples we simulate 40,000 male and female individuals and follow them, and their partners while married, for (at most) 81 periods. From this simulated sample, we keep only males between 25 and 59, and randomly sample five consecutive observations from this interval to generate a sample that approximates the SLID sample. The control group consists of men observed between 2 and 5 years before onset. The treatment, or 'post-onset', group consists of men who report at least one current disability during the five-year window in which we observe them (that is, who would be identified as disabled in our SLID sample). We weight the simulated data to replicate the weighting of the SLID sample (as described in section 2) over a ten-year post-onset period. Finally we run OLS regressions controlling for education, *rsk*, and a cubic in age.


Figure 7: Average weekly hours and annual participation rates following onset




(a) All men: Avg wkly hrs




(b) All men: Participation




(c) All men: Hours of workers




(d) Married x onset: Avg wkly hrs




(e) Married x onset: Participation




(f) Married x onset: Hours of workers



(g) Wives: Avg wkly hrs



(h) Wives: Participation




(i) Wives: Hours of workers

5.1 Learning-by-doing human capital and wages

Figure 8 shows the cross-sectional profile of post-onset hours worked by year from onset for all men and wives of disabled men, and the difference in post-onset differences in hours worked by marital status (the same figures in the first column of figure 7) when we replace our estimated human capital accumulation process with a standard exogenous wage process in which wages vary deterministically with age and education and are subject to a standard

process for persistent idiosyncratic shocks.³⁵ The results, unsurprisingly, show that human capital is necessary to explain the post-onset labor supply patterns in the data. Without a process of human capital deterioration and replenishment, the model generates patterns that are consistent with the analytical predictions from section 3: men drop their labor supply precipitously at onset and married men more so than single men.

Figure 8: Differences in mean wkly hours when wages are exogenous




The role of human capital is important because the changing stock of human capital creates a large incentive to maintain pre-onset labor supply levels when a disability shock first occurs, but reduces the work incentive over years from onset for the chronically/severely disabled as their stocks of human capital gradually depreciate. The effect is stronger for married men who have, on average, higher pre-onset wages. As well, since (some) marriages are destabilized by losses of earning potential for the main earner, married main earners (typically husbands) have an extra incentive to prevent human capital losses by maintaining labor supply. This incentive effect is sufficient to account for the patterns of male labor supply observed in the data by household type.

The previous argument is, at first blush, inconsistent with the fact that that we do not see any wage declines among workers in the data. However, Figure 9 plots the evolution of log wages after disability onset in our model against corresponding patterns in the data, disaggregated by marital status, and shows that both model and data show no average effect of post-onset duration on wages among the sample of post-onset men who stay in the labor force. This is consistent with the fact that the labor supply responses to disability onset in

³⁵We set autocorrelation coefficient of .95, which is close to similar estimates derived from the SLID for men and women and adjust the variance of idiosyncratic wage shocks to replicate the same variance of ln wages as in the human capital model.

the model are concentrated on the extensive (participation) margin. Men who experience large wage declines due to severe disability are out of the labor force leaving a pool of higher-wage post-onset men who retain their labor market attachment.

Figure 9: Wage declines in model vs. data



5.2 Selection into marriage and health outcomes


Table 4 show the distribution of disability by type and chronicity generated by the benchmark model. The upper panel of the table shows the distribution across different ds states among currently disabled males by marital status. The lower panel reports the measure of the chronicity of disability from table 1: specifically, the mean share of periods within a five-year interval in which a randomly drawn prime-age individual reports a limiting disability, conditional on reporting *at least one* disability during the interval. The model closely replicates both the distribution of chronic disability by marital status and the distribution of reported disability types by marital status: married men more likely to report exclusive h -limitations and single men more likely to report dual limitations. This is unsurprising: h -limiting shocks are generally more ‘manageable’ when the husband is the main earner since they do not affect his productivity, while the severest disabilities reduce men’s capacity in the labor market substantially and are more likely to lead to divorce or to a man remaining single.

An obvious question is therefore what the model implies about marriage and divorce rates following disability onset. Our estimates, and those by other authors, suggest at most

Table 4: Type of limiting *ds* by marital status for prime-age men

	SLID		Model	
	Married	Single	Married	Single
<i>ds</i> 3 (%)	22.2	13.2	20.1	15.4
<i>ds</i> 4 (%)	7.6	9.0	9.2	7.8
<i>ds</i> 5 (%)	45.4	47.5	45.9	46.3
<i>ds</i> 6 (%)	24.8	30.3	24.8	30.6
share of <i>ds</i> reports over 6 years	.44	.52	.43	.53

Figure 10: Changes in marriage and divorce probabilities by year from onset



a modest effect of disability onset on the likelihood of marital dissolution. Figure 10 shows the model’s predictions for the likelihood of divorcing and marrying by year from onset against those estimated from the data. The model generates a small, temporary uptick in the likelihood of divorce in the first two years post-onset which is consistent with our estimates from the SLID. There is no evidence of a change in the likelihood of marrying following onset, though again this is probably due to low variation in the dependent variable over the sample as marriage becomes an increasingly uncommon event after age 25. Either way, the model does not overpredict the rate at which marriages dissolve following disability onset.

Two other findings about the regressions reported in figure 10 are worth noting. First, the (unreported) coefficient on *rsk*, which is positive and significant and indicates that high-disability risk men are about half a percentage point (about 30%) more likely to divorce conditional on age than low-disability-risk men. This suggests that much health-based marital sorting on may actually occur before the onset of a limitation as partners perceive the risk that a potential mate becomes sick or incapacitated later on. Second, the results in

table 4 suggest that what extra post-onset divorces occur are likely concentrated among the most severe disabilities. This is consistent with results reported by ?, who finds that divorce hazards are largest for relatively young men who report the onset of severe work-limiting disabilities.

Finally, while the role of selection is clearly important, we have not formally addressed the possibility that selection into marriage is based not on health outcomes but rather on unobservables that are correlated with disability reporting in the data. Since our years-from-onset regressions use a cross-sectional estimator rather than within-estimator, we cannot rule out that post-onset differences-in-differences between married and single men are due to unobserved characteristics, such as preferences for leisure, that differ substantially between single and married men. Men who value leisure highly are less likely to marry and may also be more likely to report having limiting disabilities, a form of justification bias.

To assess this possibility, we recalibrate a version of the model in which men do not vary by rsk , but instead by their relative preference for leisure. Specifically, we assume that 25% of low-educated men value leisure more than other men, by an additive term \uparrow . Their preference function at age j takes the form

$$u_j = (1 - \gamma_j) \frac{(c\tilde{n}_\eta)^{1-\omega}}{1 - \omega} + (\gamma_j + \uparrow) \frac{\uparrow_{\infty-\psi}}{\infty - \psi} + \mathcal{I}_\eta \theta \quad (12)$$


Men with “normal” preference for leisure face the disability risk process of low- rsk men in the benchmark model. Men with high preference for leisure face the disability process of high- rsk men. Since men with strong preference for leisure also experience more disability, the estimated time costs by ds state are much lower than in the benchmark model. Table 5 shows the estimated costs from the two models and the value of \uparrow that just replicates the distribution of chronicity across marital status as defined in table 1. Figure 11 shows the estimated cross-sectional results for average weekly hours worked following onset, which are equivalent to the three figures in the first column of figure 7. The model performs quite poorly: the estimated $\hat{\delta}$ s for the pooled sample of men are very small in absolute value and mostly insignificant, consistent with the fact that disability imposes quite small costs except in ds 6. The married-single difference in difference is also muted since most pre-onset single

Table 5: Disability cost in the benchmark and heterogenous-preference models

	Benchmark	Heterogenous preference for leisure
$\delta_h(3)$	1.53	1.19
$\delta_n(4)$	1.21	0.82
$\delta_h(5)$	1.23	1.22
$\delta_h(6)$	1.61	1.40
\updownarrow	n/a	-.22

men are already working few hours. The shape of the post-onset profiles are difficult to reconcile with a model in which disability costs are not real and large.

Figure 11: Differences in mean wkly hours when preferences are heterogenous



5.3 Comparing the models

To conclude this section, we formalize the results shown above by pooling together real and simulated data and estimate the following reduced-form system of equations for men and women:

$$\begin{aligned}
 y_{it} = & \zeta_0(X_{it}, \mathcal{F}(age_{it}, \eta_i)) + S_i + \eta_i + \sum_k [\zeta_{1k} A_{kit} + \zeta_{2k} A_{kit} \times S_i \\
 & + \zeta_{3k} A_{kit} \times \eta_i + \zeta_{4k} A_{kit} \times S_i \times \eta_i] + \epsilon_{it} \\
 y_{it}^f = & \zeta_0^f(X_{it}, \mathcal{F}(age_{it})) + S_i + \eta_i + \sum_k [\zeta_{2k}^f A_{kit} \times \eta_i + \zeta_{4k}^f A_{kit} \times S_i \times \eta_i] + \epsilon_{it}^f \quad (13)
 \end{aligned}$$

where y_{it} is average weekly hours worked; S is an indicator variable taking the value one if the observation is from the SLID and zero otherwise; η is an indicator variable taking

Table 6: Tests of differences between model and data

Benchmark	Results forthcoming
Model with exogenous wages	Results forthcoming
Model with heterogenous preferences	Results forthcoming

the value one if the individual is married and zero otherwise; \mathcal{F} is a cubic in age, alone and interacted with η for males; and X contains all the demographic variables from equations (1) and (2) interacted with S , plus health status and a constant. As in equation (1), A_k denotes an indicator for year from onset with $k \in \{-1, 10\}$. Coefficient vector ζ_1 roughly captures the change in y relative to the pre-onset control group at each post-onset year ($\Delta^k y$) for single men in the simulation; ζ_2 captures the difference in single male $\Delta^k y$ between the real and simulated samples; ζ_3 captures the difference in $\Delta^k y$ by marital status in the simulated data; ζ_4 captures the difference-in-difference in $\Delta^k y$ between single and married men, taken across the real and simulated samples. For women, ζ_1^f and ζ_3^f are constrained to be zero since there are no spouses in single male households by definition; ζ_2^f captures changes in female y relative to the pre-onset control group by year from onset; and ζ_4^f captures the difference in wives' $\Delta^k y$ between the real and simulated samples.³⁶ Table 6 reports Ξ^2 test statistics and p-values from tests of the joint null hypothesis that $\{\zeta_{2k}, \zeta_{4k}, \zeta_{4k}^f\}_{k=0}^{10} = 0$ for all four models examined so far.

The results in this section imply that our model of learning-by-doing human capital combined with marital selection on health is able to reconcile the cross-sectional evidence on post-onset labor supply by men by marital status, and provide a plausible framework for assessing the role and limitations of intrahousehold insurance against different types of shock. However, the model performs relatively poorly on one key dimension: the labor supplies of married women following a husband's disability onset. In the next section, we propose a final extension to the model that allows us to account for the absence of an added worker effect: endogenous task sharing in home production. We will show that this extension eliminates an observable added worker effect among wives of disabled men even in an economy in which disability shocks impose large costs on households, but—crucially—in which those shocks

³⁶We apply a seemingly unrelated regression estimator adjusted to allow for covariation in the error terms across equations.

incapacitate individuals in both the labor market and in home production.

6 Extension: Intrahousehold task sharing

To reconcile the lack of added worker effect and absence of a wage-onset correlation among married men with the predictions of our model, we consider one final, very intuitive, extension: making time spent in home production endogenous for couples who combine their individual time inputs to complete a given bundle of *nll* requirements. Analytically the problem is very similar to the one discussed in ? in which the couple solves a two-stage intratemporal problem in which time in which the partners' home production time is chosen first so as to maximize the (permanent income) value of their disposable time, after which labor supply and saving decisions are made.³⁷ The relevance of this extension to us is that, in the presence of disability—and particularly of home-limiting disability—this flexibility allows a healthy spouse to alleviate the time loss experienced by a disabled partner, raising both his utility and his potential for market work.

6.1 Modeling endogenous task sharing

Recall that, in our model, individuals must complete a set of ‘non-labor, non-leisure’ (*nll*) activities before labor and leisure decisions can be made. The time requirement for such activities was defined in section 3.1 as \bar{h} for men and \bar{h}_f for women. We now define the time a man (or woman) *actually devotes* to *nll* tasks as h (or h_f). With both task sharing and endogenous wages, the couple's full problem is a maximization of (4) with respect to $\{c_f, c, l_f, l, h_f, h, a'_M\}$, subject to (6) and

$$\xi_5 : G(\bar{h}_f, \bar{h}) = \phi_f(h_f) + \phi(h) \tag{14}$$

where $G(\cdot)$ defines the couple's aggregate *nll* obligations as a function of the partners' autarkic *nll* requirements. The FOCs for leisure are unchanged from those given in section 3

³⁷Technically, because of human capital production function, the choices of h and l must be made simultaneously in our model, but the set-up is otherwise identical.

and appendix B. The additional FOC for the intra-household *nll* allocation choice is

$$\frac{\xi_1(\frac{\delta_h w}{\delta_n}) - \xi_{4a} \frac{\partial H}{\partial h}}{\xi_1 w_f - \xi_{4b} \frac{\partial H^f}{\partial h_f}} = \frac{\phi'(h)}{\phi'_f(h_f)} \quad (15)$$

which simply equates the ratio of the partners' marginal human capital costs of h to the ratio of their marginal *nll* productivities. Assuming incomplete specialization, so that both spouses engage in some market and some house work, the FOC implies that couples choose housework time inputs h to maximize the (permanent) value of their pooled disposable time, with labor and savings decisions made according to the remaining FOCs from (6).

To identify and estimate a intra-household task sharing technology for the simulations, we introduce the following functional form assumptions:

$$\begin{aligned} \phi(h) &= \underline{h} + a\tilde{h}^q \\ \phi_f(h_f) &= \underline{h}_f + a_f\tilde{h}_f^{q_f} \\ G(\bar{h}_f, \bar{h}) &= \bar{h}_f + \bar{h} \end{aligned} \quad (16)$$

where $h = (\underline{h} + \tilde{h})$, and \tilde{h} (\tilde{h}^f) are the husband's (wife's) time input into the 'allocatable' component of household *nll* obligations. The parameters $\{\underline{h}_f \leq h^f, \underline{h} \leq h\}$ denote the amounts of the household *nll* burden which cannot be shared and must be completed by each member using the same linear technology available to singles. Both \underline{h} and \underline{h}_f are estimated, as a proportion of total household per-period *nll*, as part of our SMM procedure. The last line of (16) states that a married household's total *nll* obligations are simply the combined autarkic *nll* obligations of both partners.³⁸

One can estimate task-sharing functions $\{\phi_f, \phi\}$ using: (i) the predicted combined sharable *nll* time of observationally similar singles as a dependent variable and, (ii), observed \tilde{h} time inputs of married couples as right-hand side independent variables. We are not aware of any Canadian data providing information about *nll* activities for more than one household member. Therefore we estimate the following empirical counterpart using data on 'housework'

³⁸Note that this is a special case of the set of constant-proportion parameterizations $G(\bar{h}_f, \bar{h}) = b(\bar{h}_f + \bar{h})$, with $b = 1$. If we did not restrict b , we could estimate a and a_f up to a multiplicative constant.

from the 1999-2005 biannual waves of the Panel Study of Income Dynamics:

$$\bar{\mathbf{h}}_f + \bar{\mathbf{h}} = a_f \mathbf{h}_f^{q_f} + a \mathbf{h}^q \quad (17)$$

where $\bar{\mathbf{h}}$ and \mathbf{h} are the empirical counterparts of the (allocatable) housework component of \bar{h} and h respectively. Non-linear estimates of (17) are reported in table 7, along with robust standard errors. Our estimates suggest that men’s and women’s time are relatively good substitutes in household production, with women becoming relatively more efficient than men at high levels of housework, around fifteen hours per week. We provide additional estimation details are in appendix C.2.³⁹

a_f	1.494 (0.285)	q_f	0.776 (0.026)
a	1.569 (0.517)	q	0.741 (0.051)
$n = 10,172$		$r^2 = .969$	

6.2 Results

Figure 12 plots post-onset labor supplies (again focussing on weekly hours of work averaged over a year) and wages, in the model augmented with endogenous intrahousehold task-sharing. To save space, we combine the lines from figures in the top left panel of figure 7 into a single image, showing the average cross-sectional post-onset differences in labor supply for all men, lying below the x-axis, and the difference-in-difference between marrieds and singles, lying above the x-axis. The top right panel shows the change in average weekly hours of

³⁹The decision to leave total *nll* activity fixed, rather than allowing households to choose it, is made for tractability. Our approach allows us to add only one extra choice variable to the simulation for marrieds rather than two (and zero rather than one for singles), and to avoid estimating all four parameters as well as preferences for home produced goods. Experiments with other household production functions suggest that substitute inputs work better for our purposes than allowing home time inputs to be complements; however the precise parametrization of the home technology is not important. In a major extension provided in the working paper version of this paper, we also experimented with allowing households access to a ‘*nll* service market’, which captures much of the same idea as allowing total *nll* time inputs to be endogenous. We find that adding this ‘market for time’ actually improves our final results marginally but in a statistically significant way. We hope to return to these issues in future work.

wives. The bottom panel shows the post-onset difference plots for ln wage similar to those in figure 9.

Comparing the figures in the top panel to those in 7 it is clear that adding a process for task sharing further improves the fit of the model. Specifically, it completely eliminates the estimated (cross-sectional) added worker effect at long durations from onset.⁴⁰ This improvement, and the tighter fit of the inter-marital difference-in-difference, is confirmed by our tests on the pooled sample of SLID and simulated data, reducing the χ^2 statistic from the test of cross-sample difference in differences in labor supply to [forthcoming and the corresponding p-value to [forthcoming].

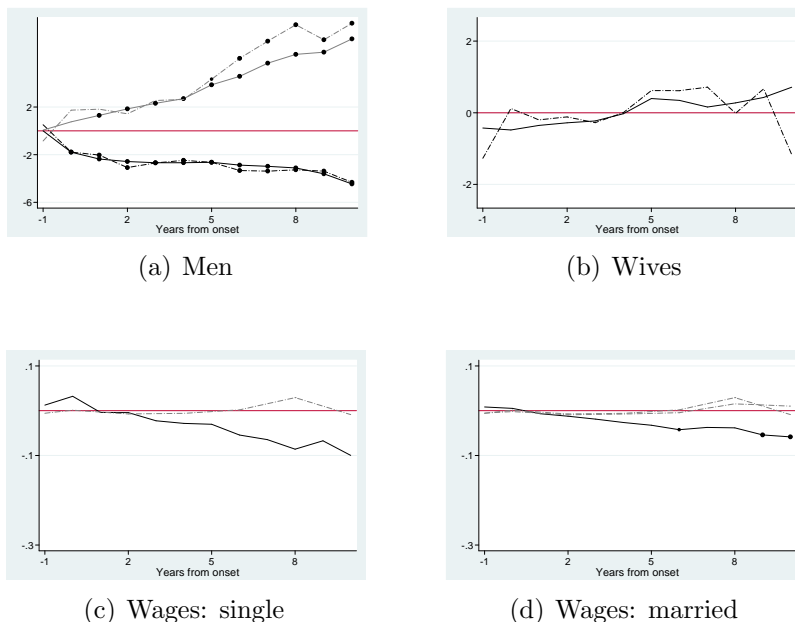
The role of endogenous home production turns in large part on the fact that most disability affects people not only in the formal labor market, but also in the home, with the estimated severity of home limitations larger, as seen in table 3. From equation (15), it is clear that home-limitations induce (to the extent possible) a substitution from male to female home production time inputs. There is also a second effect operating through the dynamic human capital motive and applicable to both δ_h and δ_n shocks: if the husband is initially the higher-wage earner, the couple may be inclined to temporarily allocate more of the *nll* burden to the wife in order to support the husband's labor supply, which in turn preserves his wages and, potentially, the stability of the union. In this respect, the ability to reallocate home production duties provides a mechanism by which couples are able to implement their best outcome without a large immediate welfare loss to the husband: preserving main-earner labor supply and human capital stocks through (usually) persistent but non-permanent spells of disability.

6.3 Disability and home production in the PSID

The previous section suggests that worker effects may in fact exist, but are concentrated in the household sector, rather than the market sector. We next provide some data that is

⁴⁰The added-worker effect based on participation and hours conditional on working are also both insignificant and close to zero.

Figure 12: Differences in hours and wages following onset, with endogenous task-sharing



broadly consistent with this idea.⁴¹ Table 8 reports results from some simple regressions run on the same sample from the 1999-2005 PSID used to estimate our *nll* sharing technology.⁴² We regress wives' average weekly reported hours of housework on an indicator for dual-earner status, the couple's wage ratio (interacted with dual-earner status; for single-earner or non-working households '*wr*' is not observed), family income, a dummy for husbands reporting work-limiting disability, husband's average weekly market hours, and the interaction of disability status and market hours ($ds \times n$). Column 1 reports results from this regression for all couples under 66. The next three columns document this relationship by age category: 51-65; 36-50 and <36. The first six rows report estimated coefficients (and standard errors); the seventh row presents average reported hours of housework for wives in the age range; the last row shows results of an F -test for joint significance of ds and $ds \times n$.

The results support the general implication of the model: for men who work very little,

⁴¹Note that the model makes no direct predictions about the time that a disabled individual would spend on household tasks, even conditional on a given level of market work. A disabled husband substitutes away from *nll* activities because they are costly in terms of time and effort (substitution effect). However *nll* activities in which he does engage consume more of his time (time-loss effect).

⁴²The sample is the same, except that we restrict spouses to be no more than five years apart in age.

Table 8: PSID: Wives' housework and husband's disability status

Wife's hours of housework h_f	All	51-65	36-50	20-35
n	0.01 (0.55)	0.10 (0.02)	0.11 (0.02)	0.08 (0.04)
ds	-3.33 (1.24)	-1.82 (1.76)	-5.44 (1.62)	-5.54 (3.56)
$ds \times n$	0.11 (0.03)	0.09 (0.05)	0.14 (0.04)	0.16 (0.08)
average h_f	18.8	19.2	18.6	18.7
F_2	1%	10%	1%	20%

a husband's disability reduces the wife's time devoted to nll , but this effect is reduced, and eventually reversed, for disabled husbands contributing sufficient hours of market work. To put these results into perspective, the wife of a currently disabled husband working 45 hours a week, performs on average 1.6 hours—or about 9%—more housework per week than the average wife. The joint effect of disability, and disability interacted with hours, on predicted housework is jointly significant at conventional levels for all groups except the youngest, for which we have only 150 observations on disabled husbands; the joint test of n , ds and ds_n (not reported) is always significant at 1%.

7 Conclusion

[Forthcoming]

References

- Aguiar, M., Hurst, E., 2007. Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics* 122 (3), 969–1006.
- Altonji, J., 1986. Intertemporal substitution in labor supply: Evidence from micro data. *Journal of Political Economy* 94 (2), 176–215.
- Au, D. W. H., Crossley, T. F., Schellhorn, M., 2005. The effect of health changes and long-term health on the work activity of older Canadians. *Health Economics* 14 (10), 999–1018.
- Batavia, A. I., Beaulaurier, R. L., 2001. The financial vulnerability of people with disabilities: Assessing poverty risk. *Journal of Sociology and Social Welfare* 38 (1), 139–162.

- Becker, G., 1985. Human capital, effort, and the sexual division of labor. *Journal of labor economics* 3 (1), 33–58.
- Becker, G., 1991. *A Treatise on the Family*. Harvard Univ Pr.
- Ben-Porath, Y., 1967. The production of human capital and the life cycle of earnings. *The Journal of Political Economy* 75 (4), 352–365.
- Blundell, R., MaCurdy, T., 1999. Labor supply: A review of alternative approaches. *Handbook of labor economics* 3, 1559–1695.
- Campolieti, M., 2002. Disability and the labor force participation of older men in canada. *Labour Economics* 9 (3), 405–432.
- Campolieti, M., Krashinsky, H., 2006. Disabled workers and wage losses: Some evidence from workers with occupational injuries. *Industrial and Labour Relations Review* 60 (1), 120–138.
- Chang, Y., Kim, S., 2006. From Individual to Aggregate Labor Supply: A Quantitative Analysis Based On a Heterogenous Agent Macroeconomy. *International Economic Review* 47 (1), 1–27.
- Charles, K. K., 1999. *Sickness in the family: Health shocks and spousal labor supply*, unpublished manuscript, University of Michigan.
- Charles, K. K., 2003. The longitudinal structure of earnings losses among work-limited disabled workers. *Journal of Human Resources* 38 (3), 618–646.
- Charles, K. K., Stephens, M., 2004. Job displacement, disability and divorce. *Journal of Labor Economics* 22 (2), 489–522.
- Coyle, C. C., 2004. *Health shocks and couples' labor supply decisions*, nBER Working Paper.
- Cutler, D., Deaton, A., Lleras-Muney, A., Summer 2006. The determinants of mortality. *Journal of Economic Perspectives* 20 (3), 97–120.
- Domeij, D., Floden, M., 2005. Labour-supply elasticity and borrowing constraints: Why estimates are biased. *Review of Economic Studies* 9, 242–262.
- Guner, N., Knowles, J., 2003. *Marital instability and the distribution of wealth*. manuscript, University of Pennsylvania.
- Imai, S., Keane, M. P., 2004. Intertemporal labor supply and human capital accumulation. *International Economic Review* 45 (2), 601–641.
- Knowles, J., Jul. 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).

- Low, H., Pistaferri, L., 2009. Disability Risk, Disability Insurance and Life Cycle Behavior. Tech. rep., Stanford mimeo.
- Mazzocco, M., Yamaguchi, S., 2007. Labor supply, wealth dynamics, and marriage decisions, university of Wisconsin Economics Department Working paper.
- Meyer, B., Mok, W., 2008. Disability, earnings, income and consumption, harris School Working Paper 06.10.
- Michelacci, C., Pijoan-Mas, J., May 2007. The effects of labor market conditions on working time: the us-eu experience. Discussion Papers 6314, CEPR.
- Olivetti, C., 2006. Changes in womens hours of market work: The role of returns to experience. *Review of Economic Dynamics* 9, 557–587.
- Rogerson, Rupert, Wright, 2000. Homework in labor economics: Household production and intertemporal substitution. *Journal of Monetary Economics* 46 (3), 557–579.
- Shaw, K., 1989. Life cycle labor supply with human capital accumulation. *International Economic Review* 30 (2), 431–456.
- Singleton, P., 2009. Insult to injury: Disability, earnings and divorce. Working Papers, Center for Retirement Research at Boston College.
- Smith, J. P., 2005. Consequences and predictors of new health events. In: *Analyses in the Economics of Aging*. NBER Chapters. National Bureau of Economic Research, pp. 213–240.
- Stephens, M., 2001. The long-run consumption effects of earnings shocks. *Review of Economics and Statistics* 83 (1), 28–36.
- Stephens, M., 2002. Worker displacement and the added worker effect. *Journal of Labor Economics* 20 (3), 504–537.
- Stevenson, B., Wolfers, J., 2007. Marriage and divorce: Changes and their driving forces. *The Journal of Economic Perspectives* 21 (2), 27–52.

A Estimation of longitudinal effects of disability using the SLID

In this appendix, we describe in more detail the process for deriving the estimates reported in section 2. In section A.1 we explain the definition of disability implied in the SLID. In section A.2 we describe the data adjustments required to estimate equations 1 and 2.

A.1 Disability in the SLID

In each year of 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 about individuals 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.

Individuals who report having one of the latter three types of limitation further report whether they were limited in the respective type of activity ‘sometimes’ or ‘often’. Non-workers under 70 who report a work limitation are additionally asked if their condition ‘completely prevents’ them from ‘working at a job or business or looking for work’. We do not use severity information explicitly in our estimates of post-onset behavior, but we do use it to inform the different disability states in the simulation.⁴³ 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. In our analysis, men who report having their condition from birth or before age 25 are discarded from the sample. We use this retrospective data to identify individuals’ year from onset and discard individuals who report current limitations but do not provide duration information.

A.2 Comparing pre- and post-onset behaviour

A crucial advantage of the SLID is its relatively large cross-sectional dimension, which gives us the sample sizes needed to estimate pre- and post-onset differences in labor supply, participation and wages disaggregated by marital status. Unfortunately, the large cross-section

⁴³Additional 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. We do not make use of this data.

is compensated by a relatively short panel dimension, especially compared to U.S. income studies like the PSID. The six-year panel, combined with the use of retrospective data on the duration of disabilities, leads us to make two sample adjustments prior to estimating (1) and (2). Because information about the duration of a condition is reported only by individuals who report a *current* limitation (i.e. whose disability is currently active) we will identify post-onset individuals only if the disability affects them in physical or economic activities in one of the $N \leq 6$ periods we observe them). This will lead us to misidentify some portion of the post-onset population, those with the least chronic disabilities, as being in the never-disabled population.

To adjust for this potential selection problem and create a disability sample that is representative of the post-onset Canadian population, we use information from a second panel data set of Canadian households, the National Health and Population Survey (NHPS) to identify the average chronicity, or report frequency, at different years from onset among Canadian men. The NHPS has followed the same panel of individuals for seven waves, with data currently available biannually from 1994 to 2008. Crucially, it asks the same questions about limiting conditions as the SLID. Taking individuals who report a new limiting condition in the first three waves of the NHPS, we can therefore observe the average re-occurrence rate at each year $k \in \{-1, 10\}$ following onset among Canadian men. Specifically, we estimate the average number of reports (out of three) in the current and subsequent two interviews for individuals observed at year k from onset in the NHPS. We then re-weight our SLID sample to reflect the distribution of chronicity at each post-onset year.⁴⁴

Our second sample adjustment is to create a weighted control group of individuals two or more year pre-onset. Our sample consists of all individuals in the SLID observed (up to five years) pre-onset, and we re-weight each pre-onset observation by the estimated probability that he would be observed in the (weighted) post-onset population based on pre-determined characteristics such as educational attainment, parents' education, region of residence and age, as well as interactions of these variables. These weighted observations are then used in the final estimations of (1) and (2) reported in section 2. As a test of the weights, we regressed several micro-variables we might plausibly expect to be exogenous from disability history on dummies for when the observation is observed: before onset, at or one year after onset, two to five years after onset or six or more years post-onset, along with age and panel effects. The test variables are indicators for having high school or more; having a bachelors or trade degree beyond high school; having held a blue collar job during in at least one year in the sample; living in a Maritime province; and living in a city of 50,000 people or more. The results are reported in the table 9 below. To save space, we report results based on the pooled sample of men; however, the results are essentially unchanged when we run the tests separately for married and single men. The only significant difference observed by period from onset is the likelihood of living in a Maritime province with increases by roughly 10% at eight or more years post-onset. Dropping Maritime residents at long durations from onset from the regressions reduces the precision of the estimates but otherwise had little

⁴⁴Some individuals in the SLID still have zero "chronicity" according to the NHPS measure since they may report only latent disability. There is no corresponding summary question about latent disability in the NHPS.

effect on the results. Taken together, these regression results offer some confidence that re-weighting approach reduces or eliminates sample selection problems that would lead us to attribute differences in labor supply at long durations of onset to other factors that are simply correlated with the likelihood of being observed as disabled at long durations from onset.

Table 9: Estimated differences in covariates by year from onset

	High school or more	Bach. degree or more	Held blue- collar job	Lives in Maritime prov	Lives in city ≥ 50,000
Year post onset (ypo)					
0-2 ypo	-.022 (.012)	-.015 (.010)	.004 (.012)	.007 (.007)	.004 (.012)
3-7 ypo	-.028 (.017)	-.022 (.014)	.010 (.017)	.012 (.014)	-.016 (.017)
8-10 ypo	-.020 (.020)	-.030 (.016)	-.014 (.020)	.013 (.017)	-.007 (.020)
F-test (3, ~ 20000)	1.09 (.35)	1.25 (.29)	.89 (.44)	2.13 (.09)	1.46 (.22)

Finally, we also ran, but omit for space, corresponding regressions using unweighted SLID data. The post-onset differences in hours and participation are (unsurprisingly) larger for both single and married men since they reflect a more chronically disabled population at longer durations from onset. However, the general post-onset labor supply, participation, and wage differences are basically identical: the married-single difference in difference is large and significant and no added worker effect is evident.

B Optimal responses to disability shocks with a dynamic human capital motive

This section provides a more detailed technical analysis of the analytic predictions of labor supply response of men to to work-limiting (Ω_{δ_n}) and home-limiting disability shocks (Ω_{δ_h}) and of female spouses to their partner's disability shock (Ω_{δ}) in the model with both endogenous (learning-by-doing) human capital accumulation and marital selection. The analysis is based on a model in which individuals' utility functions are defined over, and fully separable in, consumption and leisure as used in the numerical analysis.

To derive and interpret the analytical expressions, we maintain two assumptions, repeated here from footnote 18: (1) $\frac{\partial H(n;w)}{\partial \delta}|n \equiv H_{\delta}(\cdot)|n = 0$, i.e. there is no direct partial effect of disability on the evolution of wages so long as effective labor supply is held constant. This is consistent with our estimates of equation (11) described in section 4 and reported below in appendix C.1. (2) $\frac{\partial \varpi^f}{\partial w^f} \equiv \varpi_{H^f}^f = \frac{\partial \varpi^f}{\partial w'} \equiv \varpi_H^f = 0$. This assumption implies that women

are the main drivers of divorce, and that the likelihood of a man choosing to divorce his wife is independent of his own and of his wife's wage. As previously discussed, this assumption does not hold explicitly or entirely in the model, but is consistent with the fact that $\theta > \theta^f$ in SMM estimation of all our models, as reported in table 3, and with the idea that, to the extent it is related to economic factors, marital stability should in general be more responsive to the earning power of the main than the secondary earner.

We begin by restating the first order conditions governing spouses' leisure choices in a married household. Again, the problem, and solution, for single men has exactly the same format, conditional on $w_f = \lambda = \xi_{2b} = \lambda = \xi_{3b} = \xi_{4b} = 0$ and $\xi_3 : q = q(a', w'; \delta_h, \delta_n)$.

$$(1 - \lambda)u_l = \xi_1 \frac{w}{\delta_n} + [\xi_{3a}\varpi_H + \xi_{4a}] \frac{H_n(\cdot)}{\delta_n} \quad (18)$$

$$\lambda u_l^f = \xi_1 w_f + [\xi_{3a}\varpi_{H^f} + \xi_{4b}] H_n^f(\cdot) \quad (19)$$

Differentiation of (18) with respect to δ_n gives an expression for $\frac{\partial l}{\partial \delta_n} \equiv \mu_{\delta_n}$ that can be decomposed into four components:

- $A = \left(-\xi_1 \frac{w}{\delta_n^2} - \xi_{3a}\varpi_H \frac{H_n(\cdot)}{\delta_n^2} - \xi_{4a} \frac{H_n(\cdot)}{\delta_n^2} \right) / \Delta > 0$ [substitution effect]
- $B_{\delta_n} = \left(\frac{\partial \xi_1}{\partial \delta_n} \frac{w}{\delta_n} \right) / \Delta < 0$ [wealth effect]
- $D_{\delta_n} = \left((-\xi_{3a}\varpi_H - \xi_{4a}) \frac{H_{nnn}}{\delta_n^2} + \frac{\partial \xi_{4a}}{\partial \delta_n} \frac{H_n}{\delta_n} \right) / \Delta \gtrless 0$ [human capital effect]
- $E_{\delta_n} = \left(\frac{H_n}{\delta_n} \xi_{3a}\varpi_{H\delta_n} + \frac{\partial \xi_{3a}}{\partial \delta_n} \varpi_H \frac{H_n}{\delta_n} \right) / \Delta \gtrless 0$ [marital stability effect]

where $\Delta = (1 - \lambda)u_{ll} + (\xi_{3a}\varpi_H + \xi_{4a}) \frac{H_{nn}}{\delta_n^2} < 0$, $\varpi_H \equiv \frac{\partial \varpi}{\partial w'}$, and $\varpi_{H\delta} \equiv \frac{\partial^2 \varpi}{\partial w' \partial \delta}$.⁴⁵

Differentiation of (19) with respect to δ_h gives an expression which can be decomposed in only three parts, as there is no substitution effect associated to δ_h shocks:

- $B_{\delta_h} = \left(\frac{\partial \xi_1}{\partial \delta_h} \frac{w}{\delta_n} \right) / \Delta < 0$ [wealth effect]
- $D_{\delta_h} = \left((-\xi_{3a}\varpi_H - \xi_{4a}) \frac{H_{nnh}}{\delta_n^2} + \frac{\partial \xi_{4a}}{\partial \delta_h} \frac{H_n}{\delta_n} \right) / \Delta \gtrless 0$ [human capital effect]
- $E_{\delta_h} = \left(\frac{H_n}{\delta_n} \xi_{3a}\varpi_{H\delta_h} + \frac{\partial \xi_{3a}}{\partial \delta_h} \varpi_H \frac{H_n}{\delta_n} \right) / \Delta \gtrless 0$ [marital stability effect]

The uncompensated response of leisure hours for a husband facing δ_n (20) and δ_h (21) shocks are:

$$\mu_{\delta_n} = A + B_{\delta_n} + C_{\delta_n} + D_{\delta_n} \quad (20)$$

$$\mu_{\delta_h} = B_{\delta_h} + C_{\delta_h} + D_{\delta_h} \quad (21)$$

Corresponding optimal responses of wives to either type of shock are simpler:

⁴⁵Throughout, we also suppress the expectation operator on H' and w' . We can think of individuals choosing some \hat{H}' which is deterministic function of n and w . Then, $H' = \hat{H}' + \nu'$.

$$\mu_{\delta}^f = B_{\delta_i} + C_{\delta_i} + D_{\delta_i} \quad (22)$$

for $i \in \{n, h\}$, where

- $B_{\delta} = \left(\frac{\partial \xi_1}{\partial \delta} w_f\right) / \Delta^f$ is the wealth effect;
- $C_{\delta} = \left(\frac{\partial \xi_{4b}}{\partial \delta} H_n^f\right) / \Delta^f$ is the human capital effect; and
- $D_{\delta} = \left(H_n^f \xi_{3a} \varpi_{H^f \delta} + \frac{\partial \xi_{3a}}{\partial \delta} \varpi_{H^f \delta} H_n^f\right) / \Delta^f$ is the marital stability effect.

and $\Delta^f = \lambda u_{ll}^f + (\xi_{3a} \varpi_{H^f \delta} + \xi_{4b}) H_{nn}^f(\cdot) \gtrless 0$ and $\varpi_{H^f \delta} \equiv \frac{\partial^2 \varpi}{\partial w_f \partial \delta}$. Finally, we transform semi-elasticities of leisure μ into semi-elasticities of labor Ω (as reported in section 3.2) imposing the time-use constraint ξ_{2a} for husbands and ξ_{2b} for wives.

$$\Omega_{\delta_n} = -\frac{\mu_{\delta_n}}{\delta_n} - \frac{n}{\delta_n} \gtrless 0 \quad \Omega_{\delta_h} = -\frac{\mu_{\delta_h}}{\delta_n} - \frac{\bar{h}}{\delta_n} \gtrless 0 \quad (23)$$

$$\Omega_{\delta_n}^f = -\mu_{\delta_n}^f \gtrless 0 \quad \Omega_{\delta_h}^f = -\Omega_{\delta_h}^f \gtrless 0 \quad (24)$$

where $C_{\delta_n} = -\frac{n}{\delta_n}$ and $C_{\delta_h} = -\frac{\bar{h}}{\delta_n}$ are the time-loss effects associated to disability onset, reducing the disposable time available to the disabled individual for labor and leisure.

We next explain the roles of human capital (effect D) and marital stability (effect E) for male labor supply and the implications for disability for spousal labor Ω_{δ} in more detail.

1. *The human capital effect on post-onset male labor supply* As defined above, the total effect of a δ shock on labor supply, working through human capital motives (D_{δ}), can not be unambiguously signed. This is due to the fact that the term $\frac{\partial \xi_{4a}}{\partial \delta}$ may be positive or negative depending (mainly) on expectations about δ' .⁴⁶ If a disability shock is expected to be very persistent or ‘chronic’, the individual may expect to be driven to a corner solution for l in all future periods, for example through early retirement. In this case, $\frac{\partial \xi_{4a}}{\partial \delta}$ will be negative since future wages are irrelevant given the large disutility of working. $\frac{\partial \xi_{4a}}{\partial \delta}$ is also likely to decrease in absolute value with age and increase arithmetically in w since low human capital decreases the attractiveness of remaining in the labor force while disabled. As a result, human capital concerns may become smaller as a disability spell progresses or worsens, and as an individual ages, leading to gradual reductions in labor supply following onset.
2. *The marital stability effect on post-onset male labor supply* The marital stability effect E_{δ_h} is also of unambiguous sign because it is not in general possible to sign $\varpi_{H\delta}$ (or $q_{H\delta}$ for singles). In the extreme case, if the disability is persistent and wives always divorce disabled husbands, then $\varpi_{H\delta}$ will be negative, inducing larger drops in labor supply (usually followed by changes in marital status). If on the other hand, the presence of disability induces divorce only when husbands don’t work (because, for instance, the δ shocks also make husbands less useful in the home), then $\varpi_{H\delta} > 0$, leading to

⁴⁶The first terms in D_{δ} simply capture the fact that at an initial allocation of time across l and h , the effective time devoted to n will be reduced, raising the marginal human capital return to n if $H(\cdot)$ is concave.

smaller reductions in labor supply following shocks. Since in this paper we study only relatively stable marriages, this effect helps explain the large role that human capital plays in (potentially) solving the ‘puzzle’ of married men’s responses to disability.

3. *Spousal responses to disability onset* Finally, the net effect of disability on the labor supply spouses is very ambiguous, to the point where none of the terms in Ω_δ^f are directly signable. However, for stable marriages (or in a model with exogenous marriage, with effect E shut down) the human capital-augmented optimal responses of spousal labor are unambiguously positive, and in fact larger than in a model with exogenous wages. This is because the term $\frac{\partial \xi_{ab}}{\partial \delta}$ in D_{delta} is unambiguously positive: the wife’s human capital is more important to the household when her husband is limited by a disability.⁴⁷ When marriage is endogenous, however, the terms in Ω_δ^f relating to marital stability will generally work in the opposite direction, particularly if, as is intuitively likely, (1) the presence of disability raises the λ -weighted welfare cost of an endogenous separation ($\frac{\partial \xi_{3a}}{\partial \delta} > 0$); and (2) increases in the wife’s wage make a separation more likely ($\varpi_{H^f} < 0$). This effect is further reinforced if $\varpi_{H^f \delta} < 0$, meaning increases in spousal wages are more likely to induce a divorce when the husband is disabled. We note that this is especially likely to be the case when home production is endogenous, since there is limited potential for δ_h -limited men to re-specialize in the home.

The previous discussion suggests that both endogenous wages and marital selection are needed to resolve the “puzzles” laid out in section 3. However, they are only expository. We leave a more detailed study of marital stability and idiosyncratic risk to future work.

C Methodological issues and supplemental estimates for the numerical model

In this section we discuss issues related to the measurement and estimation of the home production and human capital technologies in our extended model.

C.1 Human capital technology

We estimate a process for human capital accumulation from the 1999-2009 longitudinal files of the SLID, specifically the six-year panels beginning in 1999 and 2002 and the (currently) five year panel beginning in 2005. To do so we use a three-stage estimation process that exploits the panel structure of the data to control for selection into the labor market (ability to observe the wage) and for the likely endogeneity of lagged hours worked.

The equations at each stage (suppressing individual and time subscripts) are:

⁴⁷Note also that under Canadian benefits policy, a man’s CPP and workers’ compensation benefit entitlements do not depend on the spouse’s labor supply.

$$\begin{aligned} Prob(\text{see wage} | X, \bar{X}, Z, \bar{Z}, P, Q, \bar{Q}, N_{3L}) \\ = \Phi(\beta_1 X, \beta_2 \bar{X}, \beta_3 Z, \beta_4 \bar{Z}, \beta_5 P, \beta_6 Q, \beta_7 \bar{Q}, \beta_8 N_{3L}) \end{aligned} \quad (25)$$

$$\hat{N}_{1L} = g(\pi_1 X, \pi_2 \bar{X}, \pi_3 Z, \pi_4 \bar{Z}, \pi_7 P, \pi_8 \bar{Q}, \pi_9 N_{3L}) \quad (26)$$

$$H = \mathcal{H}(\alpha_1 H_L, \Psi \hat{N}_{1L}, \Lambda_1 X, \Lambda_2 \bar{X}, \Lambda_4 \bar{Z}, \Lambda_5 \bar{P}, \Lambda_6 \bar{Q}) \quad (27)$$

where X , Z and Q are vectors of exogenous time-varying coefficients subject to exclusions described below; P is a vector of time-invariant individual fixed effects (five-year cohort, SLID panel indicators, indicators of aboriginal status, visible minority status, and first-generation immigrant status); \bar{X} , \bar{Z} and \bar{Q} are vectors of lags of X , Z and Q averaged in two-year intervals and spanning the previous four years of each observation; N_{iL} is a vector of variables containing hours worked n at lag $i \in \{1, 4\}$; and H is human capital stock, which is equivalent to the observed wage at lag iL . In particular:

$$\pi_7 N_{3L} = \pi_{7,1}(age_{3L} \times n_{3L}) + \pi_{7,2}(age_{3L} \times n_{3L}^2) + \pi_{7,3}(age_{3L}^2 \times n_{3L}) + \pi_{7,4}(age_{3L}^2 \times n_{3L}^2) \quad (28)$$

$$\kappa = \Lambda_1 X + \Lambda_2 \bar{X} + \Lambda_4 \bar{Z} + \Lambda_5 P + \Lambda_7 \bar{Q} \quad (29)$$

$$\alpha_1 H_{1L} = (\alpha_{11} + \alpha_{12} \times ed) H_{1L} \quad (30)$$

$$\Psi \hat{N}_{1L} = (\alpha_{21}(age \times n_{1L}) + \alpha_{22}(age^2 \times n_{1L}) + \alpha_{31}(age \times n_{1L}^2) + \alpha_{22}(age^2 \times n_{1L}^2)) H_{1L} \quad (31)$$

where the last three equations give the human capital accumulation function from section 4.

In the first stage, we run a selection equation and calculate inverse Mills ratios for the likelihood of observing a wage for all individual-year observations in the sample. The selection variable set Q includes current variables that affect the likelihood of being the labor force but are not expected to have any independent effect on the wage, conditional on our specified functions of lagged human capital (lagged deflated wages) and lagged hours worked. These indicators include the number of children in the family, current marital status, an indicator for a recent death and/or recent birth in the family, and current limiting disability status.

In the second stage, we use a heteroskedasticity-robust GMM estimator to estimate current wages as a function of lagged wages, H_{1L} , lagged hours worked, N_{1L} , X and lag vectors \bar{X} , \bar{Z} , and \bar{Q} . To control for the likely endogeneity of lagged hours, we instrument for all terms involving lagged hours with the three-year lag of the same term H_{3L} , as well as various demographic effects contained in vector Z , which are likely to affect the path of hours worked but have no additional independent effect on the observed wage. These controls include start-of-period health, household size and marital status for both men and women. The averaged lags of these covariates are, in general, proxy for individual fixed effects and are therefore not valid exclusions, as confirmed by a Hansen's J -test. Covariates that affect the current wage directly as well as through their effect on lagged hours are contained in X , which includes urban and region effects, provincial minimum wage, one-digit occupational

dummies, average male wage in each province and year and a cubic term in age.

Our sample for the second-stage regressions therefore consists of the last three observations for each individual in the 1999 and 2002 panels, and the last two observations in the 2005 panel, restricted to those individuals are also observed in each of the previous three years. The main assumptions underlying our instrumental variables approach is that lagged hours are likely to be simultaneously determined with the next-period wage due to the timing of reported hours, lagged hours and wage, or to anticipation of shocks near in the future (such as a factory closing). However, hours should be uncorrelated with iid error realizations far in the future. The additional demographic controls in Z have no effect on hours outside determining the intensity of labor market participation.⁴⁸ A Hansen J-test of the residuals from this third-stage regression fails to reject the validity of the instrument set. A further test, which includes the exclusions from the selection equation in the residual regression finds no evidence that any of these variables are individually or jointly correlated with the residuals from the human capital regression.

Finally, in the last stage of the estimation we use our residuals to estimate the heterogeneity structure of ν , which we allow to be heteroskedastic in age and H . Results from a linear regression of σ_ν^2 on a cubic in age and lagged values of human capital H are reported in the last two columns of table 10 below.

Table 10: Endogenous wage parameter estimates

Parameters			Heterogeneity structure for σ_ν^2	
	Male	Female	Male	Female
α_{11}	0.557 (0.117)	0.849 (0.0737)	H 0.527 (1.68)	-0.013 (3.46)
α_{12}	0.0450 (0.0213)	0.0320 (0.0340)	H^2 0.00108 (0.0546)	-0.0288 (0.1290)
α_{21}	$4.18E^{-4}$ ($2.08E^{-4}$)	$-5.57E^{-4}$ ($.14E^{-4}$)	H^3 $2.99E^{-4}$ ($4.90E^{-4}$)	0.0014 (0.0013)
α_{22}	$-3.24E^{-6}$ ($2.45E^{-6}$)	$1.26E^{-6}$ ($3.84E^{-6}$)	age 1.110 (3.71)	-2.460 (4.25)
α_{31}	$-3.24E^{-6}$ ($2.23E^{-6}$)	$1.26E^{-6}$ ($3.79E^{-6}$)	age^2 -0.0577 (0.0865)	0.0463 (0.104)
α_{32}	$1.35E^{-8}$ ($3.08E^{-8}$)	$-2.23E^{-7}$ ($7.29E^{-8}$)	age^3 $6.36E^{-4}$ ($6.51E^{-4}$)	$-2.72E^{-4}$ ($8.06E^{-4}$)
	.	.	$intrcpt$ 8.28 (45.2)	52.5 (55.6)
<i>Adj.R - sq</i>	0.849	0.789		

⁴⁸For instance, discrimination based on visible minority status could affect an individual's likelihood of being hired and fired, but, conditional on his hours worked and previous wage realizations, has no additional effects on his predicted current wage.

Parameter vectors α_1 , α_2 and α_3 are each jointly significant for both men and women at the 1% level.

C.2 Home production technology

We obtain the estimates reported in table 7 using a two stage approach. In the first stage, we estimate married men and women’s inputs into housework—the left-hand side variables in equation (7)—as a function of their own, and their partner’s, characteristics. These include age of both spouses and its square, years of education of both spouses and its square, number of children in the household, presence of children under six, self-reported health of both spouses, whether or not the individual receives help from relatives, whether or not the individual is Catholic, whether the individual or the spouse are union members, own and spousal race (white or non-white), individual and spousal wage and its square, and year dummies corresponding to the given wave. Instrumenting is necessary to deal with measurement error in reported housework which can arise from the fact that only one member of the couple reports housework hours for *both* himself and his spouse, as well as from potentially varying interpretations of the term ‘housework’ as defined in the questionnaire.

In the second stage, the (non-linear) estimating equation is a version of (17) except that the hats indicate that we are using fitted values from the first stage,

$$\dot{\hat{h}}_f + \dot{\hat{h}} = a_f \hat{h}_f^{q_f} + a \hat{h}^q \quad (32)$$

The dependent variables $\dot{\hat{h}}_f + \dot{\hat{h}}$ are the predicted amount of time that the partners would have to devote to housework in autarky. These values are estimated for single men and women in our PSID sample using the same set of regressors used to instrument for married hours of housework (omitting child indicators for single men, under the assumption that children live with their mothers in the event of divorce or widowhood).

The estimation sample includes all households in the 1999-2005 PSID panels for which at least one member does positive hours of housework, and total reported housework is less than 150 hours.

C.3 Cross-sectional estimates of post-onset behavior: data and models

[Final tables forthcoming]