THE UNIVERSITY OF CHICAGO

SPOUSAL LABOR SUPPLY AND THE WELFARE IMPLICATIONS OF DISABILITY INSURANCE REFORM

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE SOCIAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

KENNETH C. GRIFFIN DEPARTMENT OF ECONOMICS

BY
MAXWELL DYLAN KELLOGG

CHICAGO, ILLINOIS
AUGUST 2021

Table of Contents

List of Tables

List of Figures

Acknowledgments

This research was supported by funding from the National Institute on Aging (T32AG000243). The HRS (Health and Retirement Study) is sponsored by the National Institute on Aging (grant number NIA U01AG009740) and is conducted by the University of Michigan. I am deeply grateful to Stéphane Bonhomme, Manasi Deshpande, and Magne Mogstad for their guidance and support. I am grateful as well to Bradley Setzler for valuable advice on my structural model, and Jack Light for his assistance with the PSID data. I thank Derek Neal, Alessandra Voena, and participants of the Public/Labor Lunch for valuable feedback and suggestions. I would also like to thank many of my colleagues and classmates—too many to name individually—for the friendship and intellectual camaraderie that has made my years in Chicago so rewarding. Last but certainly not least, I thank my wife, Grace, for her support, and my parents, Ward and Debora, for the sacrifices that have made my education possible.

Abstract

This paper uses a life cycle model to study interactions between household self-insurance and the U.S. Disability Insurance (DI) system. The model is motivated and guided by evidence from panel data on disability onset in U.S. households, showing that married workers benefit from both higher self-insurance capacity and higher utilization of DI compared to unmarried workers—who are left, by contrast, more exposed to the costs of disability. These responses are consistent with adverse selection, whereby the long application process and strict work limitations of the DI system screen out worse self-insured workers. Accounting for household self-insurance and the implicit costs of utilizing the DI system, the model delivers novel insights into the welfare implications of DI reform. Welfare gains from DI reforms are large, especially ones that lower the costs of acquiring DI benefits and consequently provide income support to households that value it highly. Accounting for the substantial insurance value that expansionary reforms provide is important for drawing these welfare conclusions. On the other hand, accounting for the self-insurance provided by spousal labor supply and pooled family savings is also important, as it reduces welfare gains from DI reforms by as much as 25 percent.

Chapter 1

Introduction

The Disability Insurance system (DI) is a large component of the U.S. social safety net. In 2018, it provided over $185 billion in cash transfers to 12.5 million disabled workers, representing approximately 6.2 percent of the working-age U.S. population (Social Security Administration, 2019). Like many other OECD countries, the DI system in the U.S. experienced substantial growth over the past 50 years, which has spurred attempts at retrenchment and debate over whether and how to reform it. While DI reform is a topic of active policy interest, there is limited research to date on the social value of DI reform and the underlying trade-off between the insurance value and fiscal costs of the DI system. Low and Pistaferri (2015) concludes that DI has high social value to individuals compared to its costs, using a model of single earners. At the same time, other recent work suggests that accounting for spousal labor supply and pooled family savings is important for understanding the insurance value of DI: they buffer the effects of DI allowance on income and consumption in Norway, at least among individuals who apply for DI and then appeal an initial rejection (Autor, Kostøl, Mogstad, and Setzler, 2019). To the best of my knowledge though, no research to date has incorporated spousal labor supply and family savings in examining the utilization—and the social value—of the DI system.[1]

Household self-insurance mechanisms can theoretically crowd out the insurance value of cash transfers to disabled workers, but there are reasons to suspect that they are uniquely important for understanding the DI system of the United States. The U.S. DI system is characterized by a lengthy allowance process and strict work limitations, both of which are intended to limit the system's fiscal size and screen out workers who are capable of gainful employment. However, these costs might also perversely screen out prospective beneficiaries with poor household self-insurance capacity, who value benefit receipt highly but cannot smooth consumption over the application process. Moreover, work limitations on applicants and beneficiaries may drive stronger added worker effects in the United States than other countries that have more generous general social safety nets. Consequently, accounting for spousal labor supply and pooled family savings may be especially important to understanding the welfare value of the U.S. DI system.

The purpose of this paper is to assess the welfare implications of DI reform using a model of disability and the working life cycle that accounts for selection into DI, work disincentives of DI, and household self-insurance provided by spousal labor supply and pooled family savings. The first set of analyses examine the status quo costs of disability for married and unmarried workers. These estimates deliver three key results that are together consistent with adverse selection into DI based on self-insurance capacity. First, married workers are roughly 50 percent more likely to receive DI benefits after a severe work-limiting disability than are unmarried workers. Unmarried workers are instead more likely to continue working full-time or to move into part-time employment. Second, the presence of a spouse is, for married workers, a continued source of work income after disability onset and a source of higher precautionary savings compared to unmarried workers. Third, married workers are comparatively well-insured against declines in consumption due to disability onset. Unmar-

[1] Autor et al. study the welfare consequences of the appeals process for DI applicants in Norway and do not model pre-application behavior or non-applicant behavior.

ried workers, by contrast, experience a decline in consumption even though they more often continue working. These facts suggest that accounting for household self-insurance is critical to understanding how workers utilize the DI system and the welfare implications of DI reform.

Motivated and guided by the descriptive evidence of the first part of the paper, the second part of the paper develops, estimates, and validates a dynamic model of disability, work, DI, and the household life cycle. Drawing from Low and Pistaferri (2015), the model includes self-insurance through savings (with borrowing constraints), the DI system, and several sources of exogenous risk faced by ex-ante heterogeneous workers: permanent wage shocks unrelated to health, risk of job loss, and disability risk (distinguishing moderate from severe disability). The DI system partially insures workers against risk, with a year-long application process during which the household must fund consumption with accumulated savings, spousal work income, and general safety net transfers. Differences in the timing of disability risk, the presence of a partner, and earnings potential within the household generate heterogeneity in self-insurance capacity across households by which (adverse) selection into DI may arise. DI examiners do not observe the health or self-insurance capacity of applicants or beneficiaries. They instead make allowance and reassessment decisions based on noisy signals of individual work capacity that depend on the individual's true health and current employment. The model therefore allows for moral hazard in the sense that workers may forgo gainful employment in hopes of acquiring or maintaining DI benefits. Risks faced by partners are correlated, which limits the capacity for household self-insurance. Household preferences in the model are non-separable over work, health, and consumption.

The model is empirically fit and validated using panel data on joint consumption, savings, work, health, earnings, and household composition from the PSID. Fitting the model in steps, I first estimate parameters of exogenous processes governing health, potential wages, and marital status in reduced-form, accounting for joint selection into work and latent heterogeneity across households. The model then translates data related to consumption, work, health, and DI utilization into two key sets of parameters that are estimated by simulated method of moments: latent allowance rules for the DI system and health-dependent household preferences over leisure and consumption. To validate the fitted model, I show that it reproduces life cycle patterns in savings and earnings, qualitative patterns in joint work decisions, and the household responses to disability onset documented in the first part of the paper. Benchmarking the behavior of households in the model against the literature, it generates plausible responses to variation in wages and DI benefit generosity.

The third and final part of the paper uses the estimated model to provide several valuable insights into the welfare effects of alternative DI reforms.[2] One key insight is that expansionary DI reforms tend to be welfare-improving, consistent with previous results by Low and Pistaferri (2015). Reforms that reduce the implicit application costs of DI, though, are especially valuable. These reforms reduce adverse selection into DI on the basis of self-insurance capacity, at the expense of increasing work disincentives (including moral hazard costs). Increasing the leniency of the DI application process or implementing a national temporary DI program are two such reforms, which generate twice the social surplus of alternative reforms. The national temporary DI program, modeled after temporary DI programs that have been recently implemented at the state level, reduces the costs of acquiring DI benefits by providing applicants with immediate income support.[3] Compared to the alternative reforms, both it and increasing the leniency of the DI system reduce negative selection into DI, disproportionately increasing social welfare and DI take-up among poorly self-insured households. However, increasing the benefits provided by food stamps (SNAP) has higher per-dollar welfare gains than any of the reforms to the DI system considered in this paper,

[2] I measure the welfare effects of reform by the consumption equivalent that keeps the lifetime expected utility, behind a veil of ignorance, constant as the reform is implemented.

[3] Washington, Massachusetts, Connecticut, Oregon, and Washington DC have all recently implemented temporary DI programs. Temporary DI programs have existed in California, Rhode Island, New Jersey, and New York since the 1940s.

with its benefits being most concentrated among the worst-off households.

Another insight this paper provides is that accounting for insurance value is important for identifying optimal policy recommendations. A naïve welfare measure that abstracts away from the insurance provided by the DI system would simply recommend whichever expansionary reform has the smallest fiscal externality (Hendren and Sprung-Keyser, 2020). Such a rule would conclude (erroneously, according to the results of this paper) that expansionary DI reforms are not worth their fiscal costs to implement . The results of this paper imply that similar concerns may arise when studying other social welfare programs. For instance, results from the model suggest that insurance value plays an important role in driving the value of an expansion to food stamp (SNAP) benefits.

A final insight from the model is that accounting for household self-insurance, on top of individual-level self-insurance mechanisms, is important for understanding the value of DI reform. Neglecting to incorporate spousal labor supply can cause one to overestimate the welfare gains of DI reform by as much as 25 percent. Relatedly, the model-based results imply that the social value of expansionary DI reform comes, only to a limited extent, from insuring income against disability risk. The value of these reforms comes also from broader income reallocations and from allowing households to forgo costly work when suffering from disability and when earnings are low.

Taken together, the evidence built in this paper further develops two segments of the DI literature. The first segment focuses on the fiscal costs of DI, especially its moral hazard costs (see Autor and Duggan, 2006 and Liebman, 2015 for reviews). An important aspect of this segment, dating back to Bound (1989), includes characterizing the outcomes of DI applicants and the causal effects of DI receipt on work and income (Autor et al., 2019; Autor, Maestas, Mullen, and Strand, 2015; Deshpande, Gross, and Su, 2021; French and Song, 2014; Kostol and Mogstad, 2015; Maestas, Mullen, and Strand, 2013). This paper shows that self-insurance capacity is one important margin on which those DI applicants and beneficiaries may be self-selected. Consequently, relaxing the screening devices designed to contain the work disincentive costs documented in the literature may be welfare-improving because those devices exacerbate negative selection. More indirectly, the evidence for selection into DI on self-insurance capacity implies that exogenous changes over time in female earnings and labor supply may dynamically affect the DI applications of male workers—a potential contributor to the historical rise in DI rolls not accounted for in the literature (Autor and Duggan, 2003,0; Duggan and Imberman, 2009; Liebman, 2015).

For the second segment, this paper corroborates and builds upon recent developments in the empirical literature on insurance against disability and the role that DI plays for workers. Meyer and Mok (2019) documents negative effects of disability on consumption, work, income, and savings. Reduced-form evidence presented here reveals that marital status is an important source of heterogeneity underlying those results. Next, whereas the literature has naturally focused on how DI insures against disability risk, Deshpande and Lockwood (2020) show using reduced-form methods that DI insures household consumption substantially against non-health risks as well. The model-based results of this paper reinforce the conclusions of Deshpande and Lockwood (2020) and suggest that non-health risks play an important role in driving the social value of DI reform.

In addition to its contributions to the DI literature, a final contribution of this paper is to link work on the DI system to the broader empirical literature on screening and the targeting efficiency of social programs. Theoretically, the screening devices used by the DI system should improve the targeting efficiency of DI to the extent that they impose relatively higher costs on more work-capable and less-in-need workers (Nichols and Zeckhauser, 1982; Nichols, Smolensky, and Tideman, 1971). The recent empirical literature on waiting periods and work restrictions largely concludes that this holds in the context of other social programs. Similar application costs have been found to improve targeting efficiency for SNAP (Finkelstein and Notowidigdo, 2019) and for the Indonesian welfare system (Alatas, Banerjee, Hanna, Olken, Purnamasari, and Wai-Poi, 2016). Of special note, de Jong, Lindeboom, and van der Klaauw (2011) find wait times improve targeting of the Dutch DI system, though the authors point out that the result may not generalize to a country like the U.S. where DI applicants receive little income support (see Parsons, 1991 for a closely related argument). To my knowledge, Deshpande and Li (2019) is the first study to show empirical evidence of "perverse" screening, finding that local SSA office closures reduce the targeting efficiency of DI in terms the allowance rate of applicants and benefit receipt among low-earning, less educated individuals. This paper's analysis of a temporary DI program complements the results of Deshpande and Li (2019) and validates the concern expressed in de Jong et al. (2011), suggesting that waiting times screen out worse-insured individuals from the U.S. DI system.

The remainder of the paper proceeds as follows. Chapter 2 provides background information on the U.S. DI system and summarizes the institutional reasons that its utilization and value may depend on household self-insurance capacity—and spousal labor supply in particular. Chapter 3 introduces the subsample of the PSID which is used in my event study analysis and for fitting my structural model. Chapter 4 lays out my event study design and estimates the effects of disability onset on married and unmarried households. Chapter 5 describes the setup of my life cycle model of the household. Chapter 6 describes the mapping of model parameters to data and describes the process of empirically fitting the model. Chapter 7 fits the model parameters to data and shows that it reproduces the empirical behaviors of households well. Chapter 8 documents the key insights from the model regarding the welfare value of DI reforms and their relationships with household self-insurance.

Chapter 2

Institutional Background

The public Disability Insurance system (DI) in the United States consists of two programs, both administered by the Social Security Administration. Benefits from both of these programs consist of cash transfers and health insurance coverage.[1] Social Security Disability Insurance (SSDI) is the larger of these two programs, and it is described in detail by Liebman (2015). Eligibility and benefit amounts for SSDI are both linked to a worker's earnings history, mimicking Social Security retirement benefits. The second disability program, Supplemental Security Income (SSI), is explicitly means-tested, requires no work history, and provides standardized benefits to enrollees. The asset test and income limits of SSI are strict.[2] Earned income (including spousal income) is reduced from SSI benefits at a rate of 50 percent, after the first $65 each month.[3] SSI benefits in 2018 provided $750 monthly, though states can supplement this amount.

Both DI programs have criteria intended to limit take-up to individuals with both poor

[1]SSI provides health insurance coverage to beneficiaries through Medicaid, and SSDI provides health insurance coverage through Medicare. Medicare coverage for SSDI beneficiaries only begins two years after the establishment of entitlement to SSDI benefits.

[2]Couples can own no more than $3,000 in assets, excluding one's home and the household's first car. Since 2014, the ABLE Act allows beneficiaries to protect up to $100,000 in savings for a broad class of qualifying disability-related expenses.

[3]After exempting the first $20 of unearned income, any other unearned income is reduced dollar for dollar from SSI benefits.

health and poor employment prospects, consistent with their statutory purpose. Benefits are only provided to DI applicants who suffer from a work-limiting disability, according to the assessment of DI examiners and administrative law judges. An applicant's employment prospects are also explicitly considered as part of the determination process, based on their age, education, and occupational skill set (Chen and van der Klaauw, 2008). Furthermore, applying for and collecting DI benefits imposes strict limitations on employment—a feature of the system intended to discourage take-up from individuals with relatively strong work capacity.[4] Two other features of DI make these work restrictions more onerous. First, the process of acquiring DI benefits can be time-consuming. Including appeals and reapplications, eventual DI beneficiaries wait on average 15.3 months for an allowance decision (Autor et al., 2015). Second, most DI beneficiaries are only eligible to begin collecting benefits five months after the onset of their qualifying disability even if their decision arrives quickly.

Household self-insurance arising from spousal labor supply and pooled family savings may therefore affect the private insurance value and public costs of the DI system in theoretically ambiguous ways. On the private insurance value side, DI take-up may crowd out self-insurance from spousal labor supply and pooled family savings, reducing the insurance value households derive from DI. However, spousal work can fund consumption for DI applicants and beneficiaries, potentially reducing the welfare cost of the long application process and work restrictions of DI. By reducing the utility cost of acquiring and maintaining DI benefits, spousal labor capacity can consequently drive differences in DI utilization that are unrelated to the health or work capacity of the potential beneficiary. On the public cost side, the work disincentive costs of DI depend on the labor supply decisions of spouses. Whereas applicants and beneficiaries face strict work limitations, no formal restrictions are imposed by DI on the work of other household members (except that SSI benefits are reduced by spousal earnings). The DI system could still crowd out spousal labor supply, but it could alternatively drive

[4] Applicants and beneficiaries are generally required to earn no more work income than "substantial gainful activity." This amount was $1,180 per month in 2018.

added worker effects.

The relationship DI has with the federal income tax system can also affect, for married workers, both the consumption cost of forgoing work to apply for DI and the spousal work (dis)incentives of DI benefit receipt. Married partners in the United States are jointly taxed on their income and SSDI payments are only partly subject to income taxes, so a worker who substitutes from employment into DI receipt can reduce the average tax rate faced by their working spouse even if cash benefits from DI are high.[5] SSDI contrasts with general welfare programs like the Supplemental Nutrition Assistance Program (SNAP, or formerly Food Stamps) and Temporary Assistance for Needy Families (TANF, or formerly AFDC), which will claw back benefits based on the earned income of partners.[6] Compared to a household benefiting from those general safety net programs, a household benefiting from SSDI will face a substantially lower average tax rate on spousal labor income.

[5]Between 15 percent and 75 percent of SSDI income is exempt from federal income taxation depending on the amount of DI benefits and other taxable income received.

[6]Both of these programs condition eligibility on work requirements and income restrictions. SNAP benefits are reduced by 30% of all household income (including spousal earnings and DI receipt) net of deductions. Rules for TANF vary across states but are typically consistent with SNAP; benefits are reduced by other sources of income.

Chapter 3

Data

The main analytic sample throughout this paper is drawn from the Panel Study of Income Dynamics (PSID), though certain event study results instead rely on data from the Health and Retirement Study (HRS). The PSID follows a nationally representative longitudinal sample of households, collecting particularly detailed information on the "reference person" and their partner (if one exists). These are the two people most financially responsible for the household, and they are the focus of this paper. The HRS instead longitudinally follows a sample of older US residents (age 50 and above) and their households, collecting detailed information on health and non-work income. For the purpose of defining the analytic sample, "reference persons" in the HRS will refer to in-sample individuals.

I follow Low and Pistaferri (2015) in imposing two restrictions on the analytic sample drawn from both surveys. I first restrict to households in which the reference person has no more than a high school education. In the PSID, these reference persons are more than 2.5 times more likely to receive DI than their counterparts with more education. They are plausibly the observable subgroup in which the welfare value of DI is most concentrated, as well as the most relevant subgroup for considering the effects of marginal reforms to DI. Secondly, I further restrict to households with a male reference person who does not engage in self-employment. The restriction on sex excludes only female reference persons with no partner. While women compose almost half of DI rolls, women also tend to have less consistent labor force attachment over the life cycle which presents unique challenges for modeling interactions with the DI system.

I use waves of the PSID from 1986 through 2017, where the PSID collects information annually from 1986 through 1997 and biannually from 1997 onward. Relative to comparable longitudinal and publicly available datasets, the PSID is uniquely suited to the exercises of this paper. It contains measures of key household outcomes related to disability and DI receipt: work, income, self-reported health, consumption, program take-up, and assets. The PSID does not collect information about applications for DI benefits—only receipt of benefits. To examine DI applications and more detailed health information, I turn to the HRS data. I use all available waves of the RAND HRS file, which runs biannually from 1992 to 2018.

I distinguish households with a partnered reference person from other households. In the PSID, partnered households are those in which the reference person is married or in which the partner has been cohabiting for at least one year. For the older individuals of the HRS, these are any married or cohabiting sample persons. As most of these households involve a married couple, I refer to them throughout the remaining paper as "married" households or workers, contrasting them with the remaining "unmarried" households or workers. Table 3.1 reports average values for key measures of interest in my analytic sample of PSID households, by marital status of the reference person. Compared to unmarried reference persons, married reference persons are older, more white, wealthier, more likely to have a high school degree than have no degree, and less likely to suffer a work-limiting disability. Married households tend to contain more adults and more children than unmarried households. I observe an average of 920 married reference persons and 305 unmarried reference persons per survey wave. Similar statistics are reported for the analytic sample of the HRS in Table A.2.

Work-limiting disability is measured in the PSID by self-reported subjective assessments.

A disability for the PSID sample is classified as "severe" if an individual reports that it prevents work altogether or limits "a lot" the amount of work they can do. If a disability does not meet those criteria, it is "moderate." These definitions follow Low and Pistaferri (2015), and are provided fully in Appendix B, alongside definitions for wealth and consumption measures. Appendix Table A.1 shows how these subjective assessments capture differences in functional limitations and objective diagnostic criteria that are collected across intermittent waves of the PSID. Individuals who report a severe work-limiting disability are more likely to report all these measures of poor health than are those who report a moderate work-limiting disability, who in turn report them more often than those who report no limitation whatsoever. In the HRS, similar diagnostics are measured more consistently over the course of the survey and can be used directly to define disability events.

Table 3.1: *PSID sample summary statistics*

	Single	Married	Married-Single
Age	37.64***	40.37***	2.730***
	(0.341)	(0.183)	(0.377)
HS Grad	0.673***	0.686***	0.0132
	(0.0182)	(0.0103)	(0.0195)
White	0.435***	0.599***	0.163***
	(0.0195)	(0.0111)	(0.0207)
Number of Adults	1.231***	2.299***	1.068***
	(0.0137)	(0.00967)	(0.0165)
Number of Kids	0.253***	1.290***	1.037***
	(0.0173)	(0.0208)	(0.0258)
Mod. Dis.	0.0807***	0.0723***	-0.00838
	(0.00634)	(0.00337)	(0.00700)
Sev. Dis.	0.108***	0.0611***	-0.0466***
	(0.00967)	(0.00366)	(0.0101)
Has Partner	0.406***	1.000***	0.594***
	(0.0160)	(0.000106)	(0.0160)
Wealth	22.39***	60.18***	37.79***
	(2.271)	(2.581)	(3.328)
Years Worked	18.17***	21.90***	3.729***
	(0.340)	(0.186)	(0.377)
Years Worked Full-Time	15.84***	20.37***	4.534***
	(0.356)	(0.193)	(0.396)
Eligible for DI	0.981***	0.995***	0.0134***
	(0.00415)	(0.00144)	(0.00432)
# Households/Wave	224.23	956.91	
# Waves/Household	4.55	7.41	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level. Wealth is reported in thousands of 1996 US dollars. Partners include married spouses and non-marital partners.

Chapter 4

Event studies of disability onset

This section estimates the status quo costs of disability and their relationship with DI utilization and household self-insurance. The focal event is the first time that a disability is reported within a household, to avoid confounding the effects of current disabilities with the residual effects of past ones. No restrictions are imposed on how health evolves within households in the years following the first reported disability, though this section does restrict to workers who are fully employed in one of the two years prior to that disability.[1] In separating the consequences of disability for married versus unmarried households, I condition on the household's marital status two years prior to the first disability onset and allow it to evolve over time as an outcome.

Figure 4.1 plots the share of PSID households that are collecting DI benefits in the years surrounding first disability onset, where that disability is a severe work limitation for the household's reference person. In the four years prior to the onset of that condition, no work limitations are reported in each household (by construction) and virtually no households collect SSDI or SSI benefits.[2] For unmarried workers, DI receipt quickly plateaus within

[1] Workers who do not work in the lead up to a self-reported disability often report collecting SSI benefits before the onset of their condition and not after, which calls into question how these individuals may interpret the meaning of a work limitation.

[2] The PSID does not consistently distinguish SSDI receipt from SSI receipt across survey waves.

a year after disability onset, so that roughly 25 percent of unmarried households collect benefits between two and four years later. Married workers, by contrast, steadily increase their take-up of DI over time, so that 45 percent of married workers report receiving DI benefits four years later.

The underlying application and appeal behaviors of workers are important for understanding these differences in DI take-up. While information about DI applications and appeals is not collected in the PSID, it is collected in the HRS. As results to follow will show, unmarried workers are in fact more likely to apply for DI and appeal and initial allowance decision. However, they are less likely to ultimately receive DI benefits than married workers. Moreover, the unmarried workers who do collect DI benefits experience a shorter waiting period for those benefits than married workers. Together, these facts are consistent with unmarried workers being more often screened out by a long DI application and appeals process compared to married workers.

One can interpret the post-event averages plotted for married and unmarried households in Figure 4.1 as average treatment effects on the treated (ATTs), to the extent that households virtually never collect DI benefits without suffering from some type of work-limiting condition. Counterfactual values are less obvious, though, for other key outcomes such as employment, earnings, or consumption. To provide a more complete picture of the costs of disability onset, effects for these measures are estimated using an event study design.

4.1 Event Study Design

Let D denote the year in which a household faces disability onset, and let $Y_t(\ell)$ denote the potential outcome Y of a household in calendar year t if it were to face disability onset ℓ years in the past. Let $Y_t(-\infty)$ denote the outcome for a household which has not yet faced a disability by year t (i.e., $Y_t(-\infty) = Y_t(\ell)$ for all $\ell < 0$). The parameters of interest can be defined as follows:

$$\delta_\ell \equiv E[Y_t(\ell) - Y_t(-\infty)|D = t - \ell], \quad \ell \in \{-6, \ldots, 6\} \tag{4.1}$$

so that δ_ℓ denotes the average effect of onset ℓ years after disability onset.

To estimate each δ_ℓ, the event study design requires comparing trends over time in outcomes among work-limited households to trends over time among comparable control households. The estimators take the form:

$$\hat{\delta}_\ell \equiv \overline{\Delta Y_{-2,\ell}}^T - \overline{\Delta Y_{-2,\ell}}^C \tag{4.2}$$

where $\overline{\Delta Y_{-2,\ell}}^T$ and $\overline{\Delta Y_{-2,\ell}}^C$ denote the average change in outcome Y between 2 years prior to disability onset (the reference year) and year ℓ from disability onset, for disabled households and their matched non-disabled peers respectively.

The control group consists of contemporaneous peer households that have not yet faced (or never face) their own work-limiting condition. These peer households match exactly to one (or more) work-limited household on observable characteristics: birth cohort and marital status in the work-limited households' reference year, race of the reference person, and access to a state-level temporary disability program. For the control households, ℓ denotes time relative to disability onset for the households to which they were matched.[3]

Households that suffer a work-limiting disability at a young age will tend to have more not-yet-disabled peer households to serve as control units than households that suffer a work-limiting disability at a later age. For that reason, I reweigh the control group so that the distribution of its observable characteristics are identical to that of the households suffering a disability. Applying that weighting scheme to Equation 4.2 results in an estimator equivalent

[3]A non-disabled household may match to multiple households with the same observable characteristics but facing disability in different years. In that case, the non-disabled household will appear as a control unit multiple times, but ℓ for each duplicate will be indexed to a different age of disability onset.

to averaging together (for a given ℓ) many conventional difference-in-difference estimators, one for each set of observably identical work-limited households and their controls.

Another problem with the PSID and HRS samples is that the household panels are imbalanced. In constructing each $\hat{\delta}_\ell$, I only restrict to households observed both ℓ years from disability onset and 2 years prior to disability onset. This is a partial balance restriction, and imposing only partial balance means that the composition of households underlying $\hat{\delta}_\ell$ can vary over ℓ. PSID household panels are heavily imbalanced due to a combination of attrition, the design of the PSID, and the ending of the sample period. Imposing full balance for even a modest range of years relative to disability onset is restrictive of the (small) sample of disabled households. I observe severe disability onset, for instance, for a total of 59 unmarried households and 85 married households.

As a result of the partial balance restriction, each $\hat{\delta}_\ell$ is equivalent to the coefficient on an indicator for disability in a two-way fixed effect regression estimated on a dataset consisting of the year-ℓ sample (which includes only observations in year ℓ from disability onset and in the comparison year) with weights applied based on observable characteristics as described above. Figure 4.2 presents an example event study, reporting the average effects of a severe disability on reference person's equivalized work income from the PSID.[4] Both unmarried and married reference persons face a large and immediate drop in their own work income in the year following disability (\$9,500 and \$8,200 respectively). While there are some qualitative but noisy recovery of work income for unmarried reference persons, the decline is persistent and stable for married reference persons.

The event study design relies on a key common trends assumption: outcomes within disabled households would tend to change over time in a manner comparable to contemporaneous non-disabled peer households in the absence of disability. This will hold if, for

[4] Equivalized measures are divided by the OECD-modified scale, which approximates how income can map into individual-level consumption. It assigns a value of 1 for the reference person, 0.5 for each additional adult, and 0.3 for each child.

Figure 4.1: *Event study estimates of the effect of severe disability onset on reference person's DI receipt*

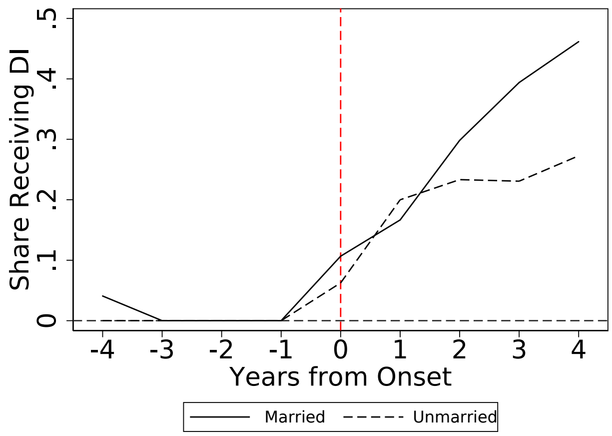

Notes: Dashed lines indicate 90% confidence intervals, with standard errors clustered at the household level. Point estimates and standard errors are zero in $\ell = -3$ and $\ell = -1$ because no work-limited or control household is ever observed on DI in those years.

lower-educated workers, work-limiting disabilities and their timings are randomly assigned conditional on other observable characteristics. The design also relies on an assumption of no anticipation of disability. This is implicit in writing $Y_t(-\infty) = Y_t(\ell)$ for $\ell < 0$. If it does not hold, then anticipation of future work-limiting disabilities will contaminate treatment effects by changing the behavior of households in the control group.

As the basis for testing the common trends assumption, I estimate $\hat{\delta}_\ell$ for $\ell < 0$ and report these coefficients, as in Figure 4.2. These effect estimates compare pre-disability trends in outcomes for disabled workers to the contemporaneous trends in outcomes of their later-disabled (or never-disabled) peers. Absent anticipation of disability onset, $\hat{\delta}_\ell$ should be close to zero for all $\ell < 0$.

Table 4.1: Effects of a severe disability onset for the reference person

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Panel A: Disability Insurance Receipt					
SSI and SSDI Receipt	0	0.00320	0.178***	0.270***	0.0924
	(.)	(0.00228)	(0.0540)	(0.0466)	(0.0714)
Panel B: Employment					
Ref. Full-Time Work	0.845***	0.861***	-0.363***	-0.373***	-0.0103
	(0.0494)	(0.0362)	(0.0983)	(0.0530)	(0.112)
Ref. Part-Time Work	0.105**	0.0784***	0.0834	0.0711*	-0.0122
	(0.0411)	(0.0265)	(0.0619)	(0.0428)	(0.0754)
Partner Full-Time Work	0.0155*	0.525***	-0.0834***	-0.0376	0.0457
	(0.00807)	(0.0577)	(0.0317)	(0.0501)	(0.0594)
Partner Part-Time Work	0.0124**	0.194***	-0.0134	-0.00157	0.0119
	(0.00589)	(0.0429)	(0.00938)	(0.0449)	(0.0452)

Notes: Point estimates in columns (3) and (4) represent the average effect of disability onset over the first seven years, on households which suffer a disability onset. They are equivalent to averaging the coefficients of Equation 4.2, proportional to the distribution of household-year observations across the post-disability years ℓ. Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Full-time employment corresponds to working on average at least 30 hours per week over the year. Gross income includes informal transfers such as assistance from family, which are accounted for in the first four income measures. Standard errors are clustered at the household level. Post-tax income is imputed using NBER TAXSIM Version 27

Table 4.1: *(Continued) Effects of a severe disability onset for the reference person*

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Panel C: Household Income					
Ref. Earned Inc.	20.20***	13.86***	-5.891***	-4.924***	0.967
	(1.827)	(0.848)	(2.084)	(0.870)	(2.260)
Social Security Income	0.285	0.441**	1.162***	1.445***	0.283
	(0.184)	(0.218)	(0.383)	(0.332)	(0.507)
Other Transfers	1.343*	1.178***	1.481***	0.783**	-0.698
	(0.749)	(0.315)	(0.517)	(0.343)	(0.621)
Work Inc. (Others)	1.064***	7.612***	0.534	-0.299	-0.832
	(0.378)	(0.948)	(0.977)	(0.701)	(1.205)
Gross Income	22.89***	23.09***	-2.646	-3.006***	-0.361
	(1.708)	(1.348)	(1.979)	(0.974)	(2.210)
Post-Tax Income	17.74***	17.22***	-1.077	-1.396**	-0.319
	(1.173)	(0.894)	(1.495)	(0.625)	(1.624)
Panel D: Consumption and Savings					
Household Consumption	13.07***	11.98***	-1.150	-0.501	0.649
	(1.819)	(0.961)	(2.192)	(1.159)	(2.480)
Household Wealth	31.05**	41.17***	-7.401	0.841	8.241
	(14.63)	(10.98)	(4.574)	(10.25)	(11.23)

Notes: Point estimates in columns (3) and (4) represent the average effect of disability onset over the first seven years, on households which suffer a disability onset. They are equivalent to averaging the coefficients of Equation 4.2, proportional to the distribution of household-year observations across the post-disability years ℓ. Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Full-time employment corresponds to working on average at least 30 hours per week over the year. Gross income includes informal transfers such as assistance from family, which are accounted for in the first four income measures. Standard errors are clustered at the household level. Post-tax income is imputed using NBER TAXSIM Version 27

Figure 4.2: Event study estimates of the effect of severe disability onset on head's work income

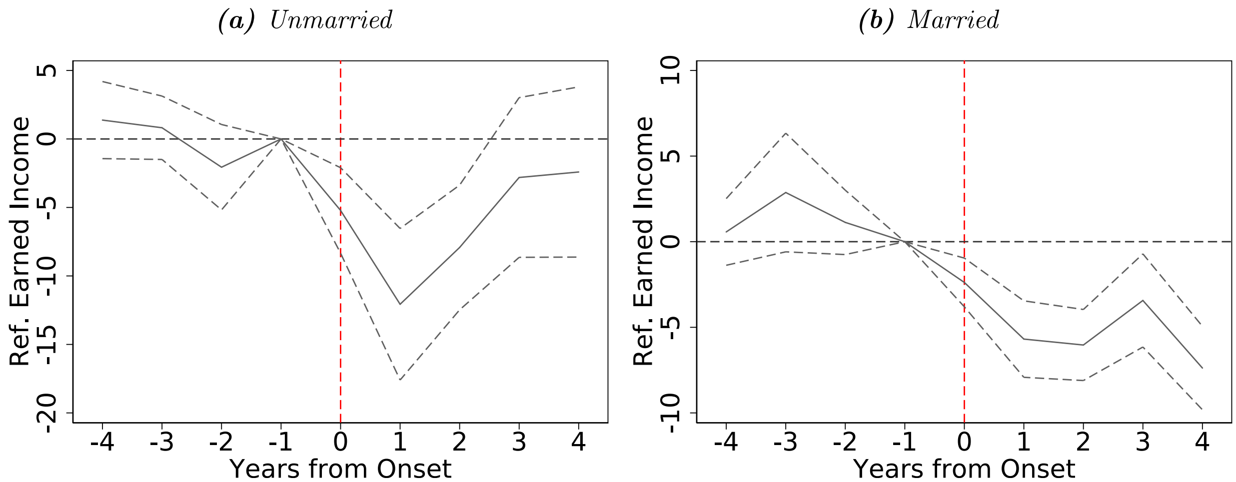

Notes: Dashed lines indicate 90% confidence intervals, with standard errors clustered at the household level. Income is reported in thousands of 1996 dollars, and adjusted using the OECD-modified equivalence scale to account for how household size affects the mapping of income into personal consumption.

4.2 Effects on DI receipt and work

Panels A and B of Table 4.1 summarize the effects of severe disability onset on household work behaviors and DI receipt in the PSID. Columns 1 and 2 report the pre-disability averages by marital status among the to-be-disabled households. These contextualize the effect sizes reported in the table, and these averages are taken over the six years prior to disability onset. Columns 3 and 4 report the average effect of disability onset by marital status. These coefficients are constructed by averaging the annual effects of disability onset (the coefficients $\hat{\delta}_\ell$) over the first seven years after disability onset.[5] Column 5 reports the average effect differences between married and unmarried workers. While differences in Column 5 are not statistically significant for any outcomes reported in Table 4.1, these event studies nevertheless reveal consistent qualitative differences between married and unmarried workers that are informative for motivating, developing, and assessing a model of disability

[5]More households are observed in years closer to disability onset, so these receive greater weight. The coefficient reported represents the average treatment effect on the treated, averaging over person-years.

and the working life cycle.

Panel A summarizes the effects on DI receipt, analogous to the level outcomes plotted in Figure 4.1. It show that married reference persons are over 50 percent more likely than unmarried reference persons to respond to severe disability onset by taking up DI benefits (27 percent and 17.8 percent respectively).

Panel B shows that married reference persons are slightly more likely to leave full-time work (37.3 percentage points) than are unmarried reference persons (36.3 percentage points), in line with the higher take-up of DI among married reference persons. Workers shift from full-time work into part-time work to a limited extent; disability onset causes the share of married reference persons working part-time to increase by 7.1 percentage points, versus (an imprecisely estimated) 8.3 percentage points for unmarried reference persons.

The final two rows of Panel B 4.1 summarize the effects of disability onset on future spousal employment. Note that these measures do not condition on the presence of a marital or non-marital partner in the post-period, so changes in partner employment can reflect the entry or exit of a partner from the household. Married workers see no detectable effects on either full-time or part-time employment of partners. By contrast, unmarried workers lose (future) spousal employment. Unmarried workers are less likely to have a full-time working partner (8.3 percentage points) due to disability. This decline is related to the decrease in the likelihood of marriage for unmarried workers, shown in Appendix C.3.

4.3 Effects on measures of income

Panel C of Table 4.1 summarizes the effects of severe disability onset for the reference person on measures of household income over the most recent tax year. Consistent with the declines in full-time employment, married reference persons see a $4,924 decline in annual equivalized work income. Owing to the fact that they tend to live in smaller households, unmarried workers see a larger decline in equivalized work income ($5,891).

The impact of lost earned income of the reference person is attenuated in two ways. First, reference persons receive income from DI programs. Married household income from these programs increases by $1,445 in annual equivalized dollars, compared to $1,162 for unmarried households and reflecting the higher take-up of DI among married households.

Second, households receive slightly more income through broader safety net programs such as unemployment insurance, worker's compensation, SNAP/Food Stamps, and housing support. These effects are reported as "Other Transfers," and are a more important channel than DI receipt for unmarried workers. Married households receive on average an additional $783 per year through these channels after disability and unmarried households receive an additional $1,481 per year. Unemployment insurance and worker's compensation drive most of the differences marital status; unmarried households collect $884 in annual equivalized income through those programs, whereas married households do not increase unemployment or worker's compensation receipts after disability onset.

Adjustments in the work of other household members due to disability generally do not provide additional (or less) income. The work income of other household members decreases by $299 for married households and increases by $534 for unmarried households, though these effects are also not statistically distinguishable from zero. This is not consistent with added worker effects as a self-insurance mechanism against disability. However, more than half of spouses in married households were working full-time before disability onset, and another 19 percent were working part-time. Consequently, with a pre-disability average of $7,612 in equivalized dollars, work income from the spouse and other household members is an important continued source of income for married households.

After accounting for transfers and other earned income in the household, married and unmarried reference persons lose $3,006 and $2,636 respectively in equivalized gross annual income due to disability onset. Owing to reductions in household tax obligations, married households face a $1,396 decline in equivalized post-tax income, versus $1,077 for unmarried households.

4.4 Effects on consumption and savings

Panel D of Table 4.1 summarizes effects of a severe disability onset on household consumption and savings. Whereas unmarried households experience a decline in consumption commensurate with the decline they face in their income, married households are better able to insulate their consumption against the decline in income due to disability onset. Average consumption declines due to severe disability by $501 equivalized dollars for married households, versus $1,150 for unmarried households. While noisily estimated, the declines in spending for unmarried households are not economically small, amounting to 9 percent of their average pre-disability consumption. Appendix C.1 decomposes spending into component categories, showing that the decline for unmarried households comes from decreases in transportation, rent, food child care, and home insurance expenditures.

With married workers insulating their consumption against income loss, they would necessarily spend down their savings if Table 4.1 accounted for all sources of income and spending. However, married households instead see a slight increase in average wealth due to disability, and it is unmarried workers (for whom declines in income pass directly through into consumption) that instead experience an economically large but imprecisely estimated decline in wealth. Appendix C.2 examines how the changes in wealth, when annualized, account for the changes in income and consumption discussed previously. It finds that household wealth declines more sharply with disability than would be implied by the declines they face in spending and income. One potential explanation is that households face disability-related costs that are not captured by PSID consumption categories, such as spending on in-home care or non-prescription treatments. Those costs are consistent with the model in Section 5.

Figure 4.3: *Effects of disability onset for the reference person in the HRS*

(a) *Effect of a work-limiting disability*

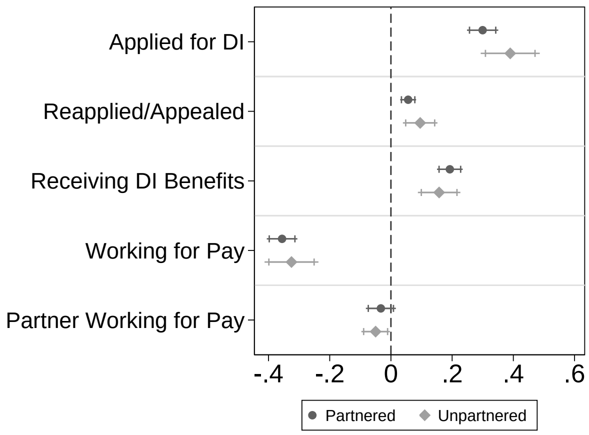

(b) *Effect of back problems*

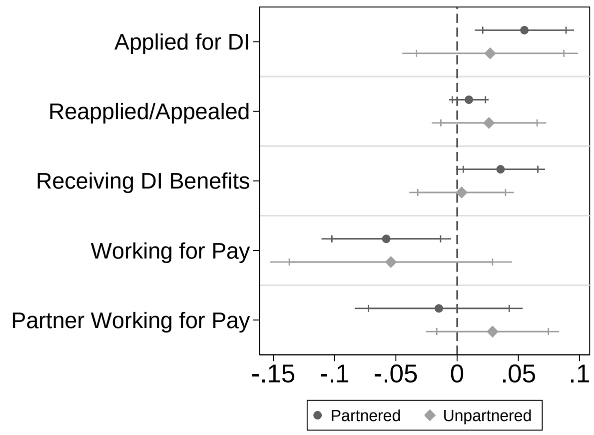

Notes: coefficients represent the average effect of disability onset over the first seven years, on households which suffer a disability onset. It is equivalent to averaging the coefficients of Equation 4.2, proportional to the distribution of household-year observations across the post-disability years ℓ. Lines indicate confidence 95 percent confidence intervals. Caps on the lines indicate 90 percent confidence intervals. Standard errors are clustered at the household level.

4.5 DI applications, appeals, and allowance after disability onset

The event studies of the PSID suggest that married households benefit from comparatively high DI take-up and stable spousal labor supply after severe disability onset, so that they are able to smooth consumption after suffering a severe work-limiting disability. Unmarried households, by contrast, rely on other transfer programs and experience a decline in consumption due to disability, despite the fact that they experience a smaller decline in equivalized post-tax income.

A limitation of the PSID, though, is that it lacks information about DI appeals or applications. To evaluate whether the higher take-up of DI by married households is driven by applications, appeals, or allowance decisions, Figure 4.3a summarizes effects of a work-limiting disability on these measures for the older participants of the HRS. As the HRS does not record measures of subjective severity, the events in Figure 4.3a represent work-limiting disabilities of any severity. It is first worth noting that even after restricting to adults over the age of 50, it remains true that unmarried workers are less likely to take up DI benefits in response to disability than are married workers. However, the differences in DI take-up appear to be driven by allowance probabilities; unmarried workers are more likely to both submit an application for DI after disability onset and to report reapplying or appealing after a rejection. One potential reason that unmarried workers are less successful in the DI application process is that they are quicker to drop their applications and appeals. Married workers who ultimately receive DI benefits report waiting an average of 316 days between application and allowance, whereas unmarried workers who ultimately receive DI benefits wait only 265 days on average.

Figure 4.3b appears to be consistent with that explanation. It reports the effects of suffering specifically from back problems on the same set of outcomes reported in Figure 4.3a. Back problems—or musculoskeletal conditions more broadly—are difficult conditions for examiners to verify and have driven an increasing share of DI claims, raising concern among some economists that they represent strategic behavior on the part of applicants. Figure 4.3b shows that both married and unmarried workers are about 5 percentage points less likely to work after back problems arise. However, only married workers correspondingly apply for (and eventual receive) DI benefits after suffering from back problems. Virtually no unmarried workers are allowed onto DI in the years following the onset of a back problem.

4.6 Comparisons with the literature

Where comparisons can be made, the event study evidence is broadly consistent with recent estimates of the effects of work-limiting disability from the literature. Averaging over married and unmarried households, the results indicate that disability causes a reduction in employment and income that is only partially offset by savings and DI receipt, therefore passing through into consumption. This characterization is consistent with the analysis of the PSID from Meyer and Mok (2019), though methodological differences, sample differences, and differences in the key outcomes of interest prevent a one-to-one comparison of results.[6]

Effects of a moderate disability on work and DI receipt are smaller compared to a severe disability, which is also consistent with Meyer and Mok (2019). The event studies of moderate disability are reported in Appendix C.4. Effects on work and DI receipt are qualitatively indistinguishable by marital status. However, unmarried households experience a qualitatively large decline in income and large declines in consumption and wealth that married households tend not to experience.

4.7 Effects on household structure and attrition

Appendix C.3 presents effect estimates for measures of attrition from the PSID and changes in household structure. It shows that unmarried households are slightly more likely to disappear from the PSID sample after disability onset compared to their peers. With regard to household composition, it shows that disability reduces the chances that a unmarried worker will get married.

4.8 Robustness checks and pre-trend tests

Married and unmarried households differ in several ways that do not relate to household self-insurance per se, yet could explain some of the differences in the effects of disability onset by marital status. One potential explanation for the differential DI take-up by marital status

[6] Meyer and Mok (2019) use a conventional fixed effects design to estimate the year-by-year effects of disability on outcomes, relative to the year of disability onset. This approach may fail to return a coherent treatment effect estimator if treatment effects are heterogeneous (de Chaisemartin and D'Haultfoeuille, 2019). My event study design returns the average treatment effect on the treated by design, and is similar in spirit to the estimation procedure proposed by de Chaisemartin and D'Haultfoeuille (2019). My sample also differs from the sample of Meyer and Mok (2019). I include the 2017 wave of the PSID, but restrict to households in which the male has no more than a high school education.

is that unmarried workers, because they are younger and work slightly less often prior to disability onset, are less likely to have a qualifying work history for SSDI. To determine if this explanation is consistent with the data, I combine historical recall about past employment with the contemporaneous employment records of the PSID to approximate eligibility for SSDI. I find virtually all reference persons in my sample should have enough work to qualify for SSDI.

Married workers also tend to be older than their unmarried counterparts when they face disability onset. Age is known to factor into the DI application process, and older workers have had more time to accumulate wealth from their own earnings to self-insure against disability. Appendix C.5 shows that differences by marital status persist when fixing age between married and unmarried households.

While this paper focuses on lower-educated workers among whom the welfare implications of DI reform are plausibly concentrated, it may also be informative to contrast the effects of disability onset on married versus unmarried workers with higher educational attainment. To the extent that higher-educated workers are better-insured against risks such as disability, one might expect that these workers respond to disability by taking up DI at a similar rate regardless of marital status. Indeed, Appendix Table C.5 shows that married and unmarried workers with more than a high school education utilize DI at similar rates after a severe work-limiting disability. However, it remains true that married workers do not experience a decline in consumption, whereas unmarried workers experience a large decline.

Appendix C.6 presents tests of the common trends assumption, finding little evidence that the assumption is violated. If pre-trends are taken seriously where evidence of them arises, accounting for them would increase the differences in how disability affects married versus unmarried households.

Chapter 5

Life cycle model of the household

The event study results suggest striking differences in how married versus unmarried workers take up DI in response to disability onset, as well how they smooth consumption around a disability event. However, these results are not sufficient to draw welfare conclusions regarding the design of the DI system or the social safety net more broadly. The welfare implications of DI reform depend critically on the cost of work, the marginal utility of consumption, and how these two features of preferences depend on disability and marital status. If, for instance, unmarried households experience a large decline in the marginal utility of consumption that married households do not experience (perhaps partners are still able to spend money in ways that benefit the household), then both may be optimally-insured against disability risk despite the fact that disability causes a decline in average consumption for unmarried households. In that case, it is unlikely that expansionary DI reform will be worth its fiscal costs to households.

To pin down these features of household preferences, I develop a model of work, DI applications, health, household formation, wage risks, and employment risks over the life cycle. The model accounts for the work restrictions and long application process of DI, self-insurance against risks through spousal work, the possibility of assortative matching on potential wages and risks, and the manner in which marriage may affect how households use consumption spending versus informal care to offset costs of work and disability. Workers face heterogeneous exposure to three sources of exogenous risk: permanent wage shocks, disability, and potential job loss. The DI system provides partial insurance against risks, but has a year-long application process during which household self-insurance mechanisms and general social safety net transfers must support household consumption. Differences in marital status, savings, and potential earnings within the household are not accounted for by the DI system and may consequently drive (adverse) selection into DI. While it may play a role in driving adverse selection into DI, the allowance process helps to control the fiscal size of the DI system, including the moral hazard costs that arise when individuals choose to forgo otherwise gainful employment in hopes of acquiring or maintaining DI benefit receipt. Moral hazard motives are derived from the allowance process for DI applicants and continuing disability reviews for DI beneficiaries, both of which examiners base on noisy signals of work capacity that depend on the individual's current employment and true health.

5.1 The household problem

In each year t, a household of marital status M chooses work, whether to apply for DI (for the male), and consumption to maximize the present discounted flow of utility:

$$\max_{C_t, L_t^1, L_t^2, App_t^1} E_t \sum_{r=t}^{72} \beta^{r-t} U(C_r, L_r^1, L_r^2; M, \boldsymbol{H}_r^1, \boldsymbol{H}_r^2) \qquad \text{(Household objective)}$$

where β is a discount parameter, L^s is an indicator for full-time employment of partner s, M is the marital status of the household, and $\boldsymbol{H}^s$ is a vector of indicators for health status for partner s ($s \in \{1, 2\}$).Health status takes three potential values for the reference person (no disability, moderate disability, severe disability). Spousal health takes two potential values (no disability or some disability). Spouses do not exist in unmarried households ($M = 0$), so spousal work and health status do not enter into unmarried household preferences.

Households take expectations over several sources of exogenous risk in the model. Household members are exposed to joint health risks that affect potential earnings and preferences over work and consumption. They also face risks of job loss and correlated permanent shocks to earnings over the life cycle. Reference persons who choose to apply for DI benefits face an allowance probability depending on age, disability status, and marital status. DI beneficiaries are randomly selected for reassessment and potential removal from DI. Lastly, households face marriage and divorce risks.

The model allows households to insure against these risks in three ways. First, households can accrue savings to hedge against the realization of future risks. Second, individuals can adjust labor supply in response to risks, such as the onset of a partner's work-limiting disability. Third, households receive partial insurance through DI and the general social safety net.

A household's state space consists of a fixed "type," savings, age in the life cycle, marital status, DI beneficiary status, and individual-specific states for the head and their spouse: health status, wage, and employment status. The household's type characterizes the profile of marriage and divorce risks, health risks, and potential wages a household will face over the life cycle. It is known to the household but unobserved in the data, and it reflects the already realized differences across workers as they enter adulthood. It captures the idea that there may be a group of individuals who have low productivity, a low likelihood of marrying a partner, and a high propensity to suffer disability. Through household type, workers may anticipate their future risks and behave accordingly.

Household choices are made subject to several constraints. If exogenously fired, an individual loses the option of working for the year. Reference persons may only apply for DI if they meet eligibility criteria, and the application process takes a full year. While applying for or receiving DI benefits, the reference person faces work restrictions. These work restrictions drive moral hazard in the DI system, in the sense that workers may choose to limit their (otherwise gainful) employment in order to apply for or maintain DI benefits. Households can save but cannot borrow against future income, a restriction which is common in the literature and which prevents households from borrowing against disability or retirement income. During the working life, the savings equation takes the form

$$S_{t+1} = R[S_t + \tau(W_t^1, W_t^2, L_t^1, L_t^2, DI_t) - C_t^* - F_{t,H_t^1}^{1,M} L_t^1 - F_{t,H_t^2}^{2} L_t^2] \qquad \text{(Savings)}$$
$$F_{t,H_t^s}^{s,M} \equiv F^{old,s,M} \mathbb{1}\{t \geq 45\} + F^{s,M\prime} \boldsymbol{H}_t^s$$

where S_t are savings at the beginning of year t, R is a fixed interest rate, W_t^s are the hourly wages of partner s, C_t^* is spending on consumption, DI_t an indicator for the reference person receiving DI benefits, and $F_{t,H_t^s}^{s,M}$ is the monetary fixed cost of work that depends on health, current household age, and marital status. The fixed cost of work consists of two additively separable terms: a cost associated with old age, and a cost associated with current health. Both costs depend, for the reference person, on marital status (spouses do not exist in unmarried households) and they reflect how work requires accommodations not captured by PSID consumption categories (e.g., spending on over-the-counter treatments). Spending on consumption is adjusted for family size using the OECD-modified equivalence scale (the distinction between C_t^* and C_t), reflecting that household consumption is an imperfectly public good. Gross income which includes earnings, DI receipt, and other transfers is mapped into net income by the function τ which approximates the joint tax and transfer system for households.

The timing of the model is as follows:

i) DI allowance are made for applicants based on health information and employment from year $t-1$. Those allowed onto DI become beneficiaries, receiving cash payments in year t.

ii) Individuals who worked in year $t-1$ may lose their job, leaving them unemployed for year t. All other individuals may choose to seek a job in year t, based on their information from year $t-1$. This job seeking is costless and provides the option of working in year t, but it is observed by DI examiners.

iii) New and continuing DI beneficiaries face potential reassessment and removal from DI depending on their health status from year $t-1$ and their job seeking behavior, where removal leaves them without benefits in year t.

iv) Exogenous health, wage, employment shocks, and changes in marital status are next realized for year t.

v) The household then makes employment decisions, DI application choices, and consumption decisions with full knowledge of its current state. The household receives net income from work, along with transfers from DI or other programs.

The household life cycle begins at age 23, and ends with exogenous retirement at age 62 followed by 10 years of post-retirement life.[1] The life cycle terminates with no bequest motive. The timing of the model allows moral hazard to arise through job-seeking for DI beneficiaries and through actual employment for DI applicants. Work-capable DI beneficiaries may choose to forgo job seeking in order to influence the reassessment process. However, in doing so, they run the risk of being removed from DI anyway and left unemployed for the year. DI applicants face a similar trade-off when choosing whether to work while applying; they may choose to forgo work income during the application year in order to increase the chances that the examiner will allow them DI benefits. Adverse selection may stem from the link between work and DI as well, to the extent that forgoing work is necessary for DI receipt and costlier for workers with worse household self-insurance capacity.

[1] At age 62, households become eligible for early Social Security retirement benefits.

5.2 Preferences

Household preferences are non-separable in consumption and the health and leisure of the two partners. The parametric specification for preferences is an extension of Low and Pistaferri (2015) and is similar to Autor et al. (2019):

$$U(C_t, L_t^1, L_t^2; \boldsymbol{H}_t^1, \boldsymbol{H}_t^2, M_t, t) = \frac{(C_t \exp(\mu M_t + \theta_1' \boldsymbol{H}_t^1 + \theta_2' \boldsymbol{H}_t^2 + \eta_1(M_t, \boldsymbol{H}_t^1) L_t^1 + \eta_2(\boldsymbol{H}_t^1) L_t^2))^{1-\gamma}}{1-\gamma} \tag{5.1}$$

$$\eta_1(M_t, \boldsymbol{H}_t^1) \equiv \eta_1^{dis'} \boldsymbol{H}_t^1 + \eta_1^M \qquad \eta_2(\boldsymbol{H}_t^2) \equiv \eta_2^{dis'} \boldsymbol{H}_t^1 + \eta_2$$

The fixed cost of work, η_s, is a function of health and (for the reference person) depends also on marital status. It consists of two additively separable terms: a base disutility of work term η_s^M, which varies by marital status (for the reference person), and an additional disability-related disutility of work term η_s^{dis}, which is common across marital status.

The marital intercept μ allows for the possibility that households prefer more (or less) equivalized consumption when married. Additionally, these preferences have three key features supposing that, as is implied by the data, work and any level of disability are both "bads" ($\eta_s < \mathbf{0}$ and $\theta_s < \mathbf{0}$, for $s \in \{1, 2\}$ and $M \in \{0, 1\}$). First, work and disability will both increase the marginal utility of consumption in the household; households will prefer to shift consumption into periods of work and/or disability. For that reason, households will prefer to self-insure against disability risk even if disability has no consequences for potential income.

Second, poor health increases the disutility of work for the individual as well as their partner, even if $\eta_s^{dis} = \mathbf{0}$. Individuals will prefer to avoid working when the household suffers poor health, which is consistent with both disability making work more difficult and with home care arising as a potential response to a partner's disability.

Third and lastly, households will prefer to avoid having partners work jointly at the same time. This is consistent with home production (which is not distinguishable from leisure in this model) reducing the cost of work for one's spouse. On the other hand, this is not consistent with leisure complementarities in utility. Instead, empirical correlations in work within the household will be driven by life cycle dynamics and by correlated health and wage risks.

5.3 Taxes, transfers, and Disability Insurance

The tax and transfer system τ maps earned income and DI beneficiary status into net income, closely following the U.S. tax code and rules for major transfer programs. The tax and transfer system is described in full detail in Appendix D. It includes federal income taxes (joint for married partners), child tax credits, FICA taxes, the EITC, unemployment insurance, SNAP (Food Stamps), SSDI, and SSI.

The DI system in the model is simplified in a few ways. For computational reasons, eligibility for DI requires only that a reference person has worked full-time for at least one year before applying. The only cost of applying for DI is, potentially, forgone work income during the application process and during benefit receipt if allowed onto DI. Owing to a lack of information about DI applications and appeals, I do not fully model the appeals process. Reference persons instead make a single year-long decision to apply for DI. Allowance probabilities depend on marital status, current disability, and age in a manner which is consistent with the Vocational Grid used by examiners. In addition, they depend on employment during the application process (for applicants) or current job-seeking (for beneficiaries). One can interpret the resulting allowance probabilities as an equilibrium between DI examiners, who observe noisy signals of work capacity that are a function of health and work-related behaviors, and applicants, who have fully committed to the application process.

An applicant who is allowed on DI begins collecting benefits in the year after applying; a rejected applicant may return to work. The model does include reapplications; rejected applicants may chose to re-initiate the application process after returning to work. The model also includes Continuing Disability Reviews, whereby DI beneficiaries are reassessed by examiners. Like the allowance decision to follow, the probability of reassessment depends on health status and current job seeking. As I do not allow for an intensive margin to the employment decision, I do not model the substantial gainful activity threshold over which the DI system formally limits work. Instead, I suppose the DI system views all work (regardless of the earnings produced by it) equally.

5.4 Wage equation

Individuals face a potential wage that evolves over time and depends on their health status. Potential wages of partner s at age t are characterized by the equation:

$$\ln W_t^s = \beta_1^{s\prime} \boldsymbol{A}_t^s + \beta_2^{s\prime} \boldsymbol{H}_t^s + \beta_3^s M_{it} + f_g^s + \omega_t^s, \quad \omega_t^s = \omega_{t-1}^s + \xi_t^s, \quad \omega_{22}^s = 0 \tag{5.2}$$

where $\boldsymbol{A}_t^s$ is a square polynomial in age, $\boldsymbol{H}_t^s$ denotes the vectorized health status of person s, and f is an intercept that is common among all households of the same household type g.[2]

The permanent wage innovations ω_t^s follow a random walk and are potentially correlated within households, reflecting variation in potential earnings which is unrelated to health. Correlations in observed earnings within households will additionally be driven by life cycle dynamics and by correlations in health shocks. With heterogeneity in initial wages being driven by the type-specific intercept f, the random walk begins at $\omega_{22}^s = 0$.

[2]In estimation, the wage equation also contains calendar year fixed effects.

5.5 Employment

Individuals may choose each year whether or not to seek a job, and if they seek a job, whether to actually work full-time or ultimately not work that year. The choice to seek a job is costless and trivial except for DI beneficiaries, for whom it affects reassessment and examiner allowance probabilities. At the end of each year, those who choose to work are subject to a risk of exogenous job destruction with probability δ_s^M which depends on marital status M (for the reference person) and differs across partners. If that occurs, the worker is subject to an unemployment spell for the following year, and the worker receives Unemployment Insurance benefits. This job destruction process provides one mechanism by which individuals may choose to apply for DI as a consequence of employment shocks rather than health shocks.

5.6 Household formation

Marital status can vary over the life cycle as workers exogenously marry and divorce partners. The probability of finding a partner when single, or separating from a partner when married, depends on the age and household type. Persons of the same household type match with a common pool of potential spouses, allowing for flexible assortative matching on characteristics which are realized before early adulthood. The health and wages of a potential spouse are carried with a worker throughout the life cycle, and they evolve whether the worker is married or unmarried. This means that divorce does not affect the household's expectations about spousal health or spousal potential wages if the household later re-enters marriage.[3]

The number of adolescent and adult children within the household evolves exogenously and deterministically over time, depending on household type. Household size is allowed to vary as a function of marital status, reflecting how dependents are split among households in

[3]Savings follow the male partner across marriage and divorce, as I do not model behaviors of unmarried spouses separately.

separation and in non-marital relationships. As in Blundell, Pistaferri, and Saporta-Eksten (2016b), counterfactual simulations based on the model will hold these processes fixed and abstract away from any implications that DI reform might have for household formation. Variation in the number of children over the life cycle, by marital status, and household type serves two purposes. First, it drives variation in the cost of consumption, since children must be financially supported. All else equal, households will prefer to spend more when they have more children. Second, children enter into the tax and transfer system. EITC payments, tax deductions, tax credits, and personal tax exemptions all depend on the number of children in the household. Similarly, SNAP/Food Stamp eligibility and payments both depend on the number of children in the household.

5.7 Health process

Individual health can take one of three states for reference persons (healthy, moderately work-limited, severely work-limited) or two states for partners (healthy, work-limited). Health within a household evolves jointly over time, according to an exogenous first order Markov process. Transition probabilities in this process depend flexibly on age and household type. This allows health shocks to be flexibly correlated within households, in a manner that may change over time. As suggested by the results in Table C.1, health is not allowed to affect marriage or separation probabilities.

The conditional exogeneity of health risks requires that work does not affect health. This restriction is reasonable, as effects of non-employment on health found in the literature have been economically small (Bound and Waidmann, 2007; Stern, 1989). Furthermore, shocks to wealth through inheritance or stock market variation, and variation in income across birth cohorts all have little effect on health (Adda, Banks, and von Gaudecker, 2009; Kim and Ruhm, 2012; Schwandt, 2018). Income more generally appears to have at most modest effect on health fixing education (Smith, 2004).

Chapter 6

Taking the model to data

Using the household data from the PSID, the model is estimated in three steps. In the first step, some model parameters are fixed without estimation, at values determined by policy. I fix the probabilities of reassessment for DI beneficiaries to match SSA guidelines. These guidelines are based on likelihood of recovery ("expected to improve", "not permanent", or "permanent"), but I follow Low and Pistaferri (2015) in mapping them to current health status in the model if an individual chooses not to seek a job (giving probabilities of 1, 1/3, and 1/7 in order of decreasing health). As very few individuals ever concurrently work (in any capacity) and receive DI benefits in the data, I assume that the DI system is strict with individuals who engage in work. If an individual seeks a job, they will certainly be reassessed regardless of their health status. Moreover, applicants who worked while applying for DI and beneficiaries who seek a job will both be rejected by DI examiners with certainty. Consequently, work and DI receipt are mutually exclusive choices for an individual in the fitted model. I also fix some preference parameters at values commonly used in the literature (e.g., Autor et al. 2019; Low and Pistaferri 2015). Specifically, I fix the interest rate R at 1.016, the discount factor β at 0.9756, and the coefficient of relative risk aversion γ at 1.5.

The remaining two steps fit parameters of the model to data. The second step involves analytically identifying and estimating health and marriage transition probabilities, household types, and wage equation parameters with empirical moments which do not depend on the structure of the model. Finally, the third step estimates remaining parameters by simulated method of moments (SMM).

6.1 Ex-ante heterogeneity across households

Households in the model have some ex-ante information about their self-insurance capacity, their exposure to health risks, their household formation prospects, and their potential wages. To estimate this heterogeneity, households are allowed to draw one of three potential types at the beginning of the life cycle. These types are estimated by k-means clustering on reference persons' average log wages and the share of years spent working full-time.[1] With this approach, estimating the model amounts to the two-step procedure proposed by Bonhomme, Lamadon, and Manresa (2019), except that the second step here will be moment-based. In classifying household types, I restrict to households observed in at least three survey waves. Reference persons with no earnings in any survey wave are assigned to the lowest type (as all these households tend to have low marriage rates and high disability risk).

The patterns across types in health risks, marital risks, and potential wages arises non-parametrically from the data. The estimates presented in the sub-sections to follow show that households with the least fortunate childhood draw ("low-type" households) face high risks of disability, are unlikely to get married, and have low expected potential wages. High-type households are likely to be married, have low disability risks, and have high expected potential wages. Mid-type households fall between high-type households on each of these characteristics.

Because household-level averages are taken over potentially short panels, the average characteristics used to classify households may be noisily measured. Consequently, k-means may partition household types with excess error, as it constructs these partitions by min-

[1] Characteristics are residualized with respect to age of the reference person and calendar year, to control for differences arising due to when the household was sampled by the PSID.

imizing within-group variances rather than using information from the model to maximize likelihoods. Fortunately, excess classification error does not appear to be a great concern; Table A.3 shows that k-means clustering assigns 83 percent of households to the type most likely to generate its observed average characteristics, according to the fitted model. Most of the classification error comes from reassigning some low-type households to instead be mid-type households.

6.2 Health risks, household formation risks, and household demographics

Taking household type as given, the size of the household (the number of adolescent and adult dependent children), marital transition probabilities, and health transition probabilities are independent of any household choices and can be estimated by their natural sample analogues.[2] To illustrate heterogeneity in the health risks faced by households, Figure 6.1 depicts the estimated probability that healthy reference persons suffer some kind of negative health shock over the life cycle given their household type. Reference persons in low-type households are most likely to experience a negative health shock, facing an annual disability probability that rises from less than 2.5 percent at age 23 to 15 percent at age 62. Reference persons in mid-type and high-type households, by contrast, face an annual disability probability that remains well below 5 percent for the whole life cycle. Figure 6.2 depicts the probability that a reference person is married at each age, reflecting flows into and out of marriage across the life cycle. Low-type households are roughly half as likely as others to start the life cycle with a married partner. All households types experience net divorce in the first 3 years of adulthood, before marriages begin outpacing divorces for the remainder of the life cycle. Marriage among low-type households accelerates in the 40's, so that low-types

[2]The PSID survey moves from annual to biannual waves in 1996. In estimating transition probabilities for marriage and health, I follow Low and Pistaferri (2015) in abstracting away from the change in panel structure.

Figure 6.1: *Negative health transitions among healthy households, by type*

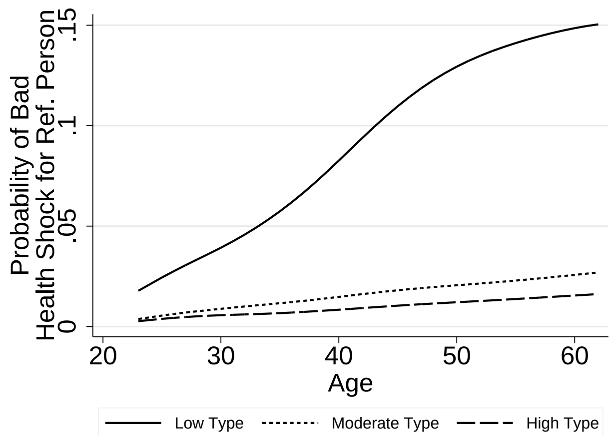

catch up with other households at the end of the working life cycle.

Consistent with the lack of evidence in Table C.1 of substantial correlation between health and marital status, health and marital risks are allowed to evolve independently conditional on household type and age.

6.3 Wage process

I follow Low and Pistaferri (2015) in augmenting Equation 5.2 with an idiosyncratic error term which can be interpreted as measurement error:

$$\ln W_{it}^{s} = \beta_1^{s\prime} \boldsymbol{A}_{it}^{s} + \beta_2^{s\prime} \boldsymbol{H}_{it}^{s} + f_{g(i)}^{s} + \omega_{it}^{s} + \epsilon_{it}^{s} \tag{6.1}$$

$$\omega_{it}^{s} = \omega_{it-1}^{s} + \xi_{it}^{s} \tag{6.2}$$

$$\begin{pmatrix} \xi_{it}^{1} \\ \xi_{it}^{2} \end{pmatrix} \sim N\left(0, \begin{bmatrix} {\sigma_\xi^1}^2 & \theta\sigma_\xi^1\sigma_\xi^2 \\ \theta\sigma_\xi^1\sigma_\xi^2 & {\sigma_\xi^2}^2 \end{bmatrix}\right), \quad \epsilon_{it}^{s} \sim N(0, {\sigma_\epsilon^s}^2) \tag{6.3}$$

Estimation of the wage equation proceeds in two steps. In the first step, I estimate the

Figure 6.2: *Marital transitions by productive type*

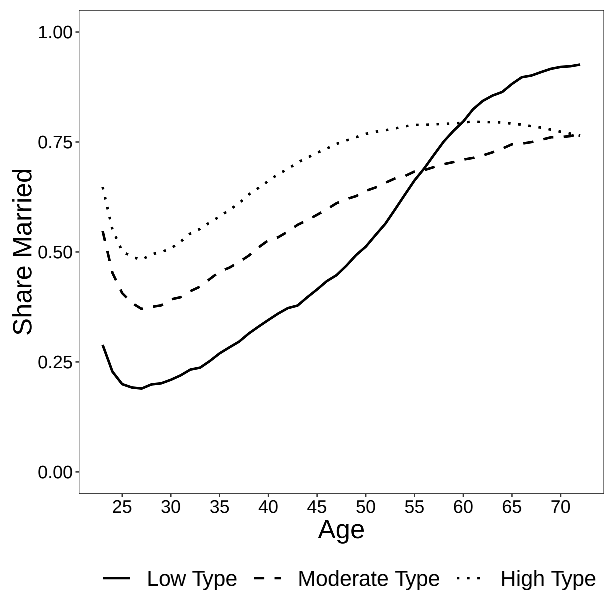

coefficients β_1^s and β_2^s using first differences, with the estimating equation:

$$\Delta \ln W_{it}^s = \beta_1^{s\prime} \Delta \boldsymbol{A}_{it}^s + \beta_2^{s\prime} \Delta \boldsymbol{H}_{it}^s + \Delta \omega_{it}^s + \Delta \epsilon_t^s \tag{6.4}$$

These coefficients are estimated by OLS, using a Heckman selection correction (Heckman, 1979). To implement the correction, I follow Low and Pistaferri (2015) by instrumenting for selection into work with a simulated measure of the generosity of potential welfare benefits, in the event that an individual has low work earnings. This measure is in the spirit of simulated IV, and uses variation over states and time in welfare program policies (see Appendix E). With household types identified separately from the parameters of the wage equation, I allow households of different types to differentially select into work in the first stage.

The type-specific fixed effects of the wage equation are estimated in the second step. The estimation of these fixed effects relies on the residual log wage, removing the effect on wage of observable characteristics:

$$\widetilde{\ln W_{it}^s} = \ln W_{it}^s - \hat{\beta}_1^{s\prime} \boldsymbol{A}_{it}^s - \hat{\beta}_2{}' \boldsymbol{H}_{it}^s \tag{6.5}$$

Estimates for the group fixed effects $\hat{f}_g^s$ are taken for each g using an average of $\widetilde{\ln W_{it}^s}$ over all i such that $g(i) = g$, and applying a selection correction to account for the truncation of the empirical distribution of wage innovations ω_{it}^s due to selection into work.[3]

Appendix F provides more information regarding the estimation and selection correction procedure, including the modeling of the work decision. Appendix F also derives an estimator of the variance-covariance matrix of the wage innovations. This estimator accounts for selection into work on the part of both partners, extending from the manner in which the Heckman correction addresses selection on the part of a single individual. The current estimator relies on the assumption that measurement error terms are exogenous, and it imposes some restrictions on the joint selection problem of partners: it imposes that shocks to preferences for work are uncorrelated across partners, and that one's wage shocks are unrelated to the latent preferences for work of one's partner.

Table 6.1 reports estimates for regression coefficients. Disability has at most a small and imprecise effect on the potential wages of reference persons or their partners; a severe disability for the head reduces his work income by about 9 percentage points. The coefficients on the age polynomials indicate that the profile of wages over age for reference persons is steeper than for their partners. Figure 6.3 shows this by presenting permanent wages over the life cycle, for reference persons and their (potential) partners of each productive type. The plot and table also show that differences in initial wages across productive types are larger for reference persons compared to their partners. Both reference person and spousal potential earnings are higher for high-type households than low-type or mid-type households,

[3] The selection correction in the estimator $\hat{f}_g$ is consistent with the correction applied in Equation 6.4 and the distributional assumptions imposed on ω_{it}^s. It is provided in Appendix F.2. As mentioned in Appendix F.3, I assume workers do not select on measurement error.

Table 6.1: Wage equation parameters

	Head	Partner
Sev. Dis.	-0.0485	
	(0.0529)	
Mod. Dis.	-0.00205	0.00120
	(0.0169)	(0.0343)
Age	0.0537***	0.0462***
	(0.00572)	(0.00996)
$Age^2/100$	-0.0487***	-0.0510***
	(0.00731)	(0.0121)
Married	0.00477	
	(0.0129)	
f_1	0.867	1.282
f_2	0.921	1.229
f_3	1.552	1.512
N_1	634	
N_2	1,184	
N_3	1,387	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level. Wealth is reported in thousands of 1996 US dollars. Partners include married spouses and non-marital partners.

consistent with assortative matching.

Table 6.2 reports estimates for the variance-covariance matrix characterizing the wage innovation process for reference persons and their partners. Year-to-year wage innovations are large, with a standard deviation of 14.5 log points for reference persons and a standard deviation of 25.9 log points for partners. Wage innovations are moderately and positively correlated within households, with a correlation coefficient of 0.259.

Table 6.2: *Variance-covariance parameters of the wage innovation process*

σ_ξ^1	0.0209***
	(0.00163)
σ_ξ^2	0.0672***
	(0.00631)
ρ_ξ	0.259**
	(0.0882)

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Figure 6.3: *Permanent earnings from full-time work over the life cycle, by household type*

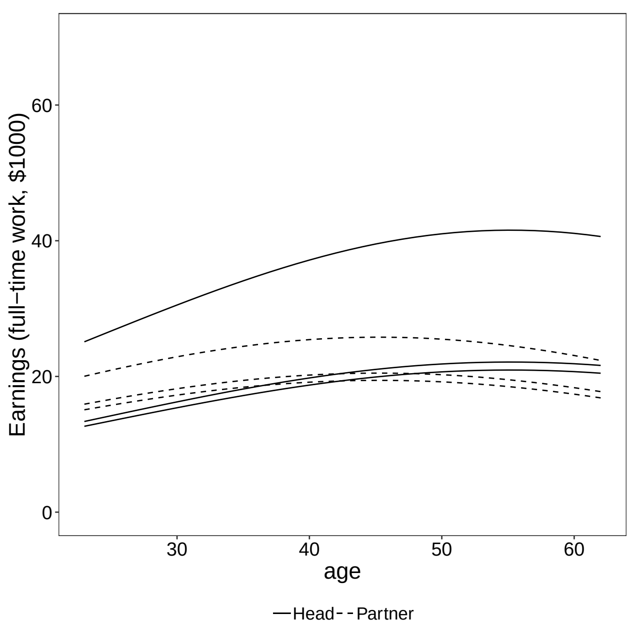

Notes: The three lines depict the expected path of earnings (absent a work-limiting disability) for low, medium, and high household types in increasing order.

__Table 6.3:__ Targeted employment and DI moments and their simulated fit

	Data	Simulation	Diff	Std. Error
	Panel A: Employment Shares, Married Workers			
Healthy, Young	0.93	0.921	0.009*	0.004
Healthy, Old	0.853	0.855	-0.002	0.008
Dis., Young	0.736	0.77	-0.034	0.024
Dis., Old	0.502	0.488	0.013	0.031
Sev. Dis., Young	0.275	0.179	0.096*	0.032
Sev. Dis., Old	0.098	0.131	-0.033*	0.014
	Panel B: Employment Shares, Single Workers			
Healthy, Young	0.813	0.812	0.001	0.01
Healthy, Old	0.744	0.691	0.053*	0.024
Dis., Young	0.469	0.487	-0.018	0.049
Dis., Old	0.301	0.296	0.005	0.06
Sev. Dis., Young	0.199	0.178	0.022	0.033
Sev. Dis., Old	0.1	0.091	0.009	0.029
	Panel C: Employment Shares, Spouses			
Healthy, Young	0.58	0.626	-0.046*	0.01
Healthy, Old	0.558	0.514	0.044*	0.017
Dis., Young	0.324	0.264	0.061*	0.022
Dis., Old	0.193	0.268	-0.075*	0.021

Reference persons who report receiving only SSI (and not SSDI) benefits did not, by definition, work regularly before entering the DI system and are dropped from these calculations. Confident Stars denote statistical rejection of equality of simulated and empirical moments: * $p < 0.05$. Standard deviations are constructed by adjusting empirical standard error estimates (clustered at the household level) by a factor $(1 + 1/S)$, where $S = 10$ is the number of times the PSID dataset is re-simulated (McFadden, 1989)

6.4 Preferences, frictions, and Disability Insurance parameters

The remaining model parameters (household preferences, work frictions, and DI allowance probabilities when not working) are estimated by simulated method of moments. The moments related to work and DI take-up which I target are reported in the first column of Table 6.3. The moments related to consumption which I target are reported in the first column of Table 6.4. While all the parameters are estimated jointly and their mappings into data are indirect, each parameter can be linked to specific moments which intuitively provide

Table 6.3: *(Continued) Targeted employment and DI moments and their simulated fit*

	Data	Simulation	Diff	Std. Error
	Panel D: Flows into DI			
Healthy, Young	0.001	0.002	-0.001*	0
Dis., Young	0.025	0.027	-0.002	0.007
Sev. Dis., Young	0.156	0.098	0.058*	0.023
Healthy, Old	0.008	0.009	0	0.001
Dis., Old	0.083	0.078	0.005	0.014
Sev. Dis., Old	0.259	0.252	0.007	0.027
	Panel E: Stocks on DI, Married Workers			
Healthy, Young	0.002	0.004	-0.002*	0.001
Dis., Young	0.049	0.057	-0.008	0.014
Sev. Dis., Young	0.288	0.185	0.103*	0.044
Healthy, Old	0.015	0.023	-0.008*	0.002
Dis., Old	0.204	0.229	-0.025	0.025
Sev. Dis., Old	0.598	0.555	0.043	0.034
	Panel F: Stocks on DI, Single Workers			
Healthy, Young	0.008	0.005	0.003	0.002
Dis., Young	0.187	0.086	0.101*	0.041
Sev. Dis., Young	0.411	0.24	0.171*	0.056
Healthy, Old	0.058	0.037	0.021	0.012
Dis., Old	0.384	0.321	0.063	0.057
Sev. Dis., Old	0.603	0.584	0.019	0.062

Reference persons who report receiving only SSI (and not SSDI) benefits did not, by definition, work regularly before entering the DI system and are dropped from these calculations. Confident Stars denote statistical rejection of equality of simulated and empirical moments: * $p < 0.05$. Standard deviations are constructed by adjusting empirical standard error estimates (clustered at the household level) by a factor $(1 + 1/S)$, where $S = 10$ is the number of times the PSID dataset is re-simulated (McFadden, 1989)

Table 6.4: *Targeted log consumption regression coefficients and their simulated fit*

	Data	Simulation	Diff	St. Dev.
Disab (mod), single	-0.192	-0.194	0.002	0.097
Disab (sev), single	-0.144	-0.25	0.106	0.128
Head Work (FT), single	0.568	0.648	-0.081	0.063
DI (healthy), single	0.089	-0.076	0.164	0.089
DI X Disab (mod), single	0.537	0.295	0.242	0.143
DI X Disab (sev), single	0.346	0.225	0.122	0.147
Married	0.289	-0.238	0.526*	0.251
Disab (mod), married	-0.16	-0.111	-0.049	0.074
Disab (sev), married	-0.168	-0.062	-0.107	0.075
Head Work (FT), married	0.214	0.24	-0.026	0.043
Work x Disab (mod)	0.235	0.296	-0.061	0.083
Work x Disab (sev)	0.133	0.208	-0.075	0.105
Spouse Disab	-0.09	-0.039	-0.05	0.033
Spouse Work (FT)	0.161	0.517	-0.356*	0.02
Spouse Work (FT) X Disab	0.144	0.326	-0.182*	0.045
DI (Healthy), married	-0.234	-0.281	0.048	0.08
DI X Disab (mod), married	-0.032	0.055	-0.088	0.088
DI X Disab (sev), married	0.039	-0.222	0.261*	0.091

Reference persons who report receiving only SSI (and not SSDI) benefits did not, by definition, work regularly before entering the DI system and are dropped from these calculations. Confident Stars denote statistical rejection of equality of simulated and empirical moments: * $p < 0.05$. Standard deviations are constructed by adjusting empirical standard error estimates (clustered at the household level) by a factor $(1 + 1/S)$, where $S = 10$ is the number of times the PSID dataset is re-simulated (McFadden, 1989)

identifying information. It is straightforward to see how these moments partially identify their corresponding parameters, fixing the remaining parameters of the model.

The fixed costs of work associated with health (the vector $F^{s,M}$) drive down employment among both younger and older disabled workers, by reducing the payoff from work when one suffers a disability. The fixed cost of work associated with old age ($F^{old,s,M}$) mediates the difference in employment rates between younger workers and older workers fixing health. The exogenous job loss rate δ^s drives down employment among young and healthy workers specifically, who face no fixed costs of work and typically prefer to work if given the option. These parameters are all intuitively pinned down by employment shares for reference persons and their partners, conditional on health status and age bin.

The DI allowance probabilities for individuals who do not engage in work are naturally pinned down by the flow of reference persons into DI receipt, conditional on health status, employment, and age group. They can alternatively be pinned down by the share of reference persons on DI conditional on age and health status. For estimation, I fit the allowance probabilities to target all of these moments.

The parameters for disutility of work and poor health (the functions η_s and the parameters θ_s) are linked to empirical moments from a regression of log household consumption on work and household characteristics conditional on marital status:[4]

$$\begin{aligned}\ln C_{it} =& \alpha_0^{M\prime} + \alpha_1^{M\prime}\boldsymbol{H}_{it}^1 + \alpha_2^{\prime}\boldsymbol{H}_{it}^2 + \alpha_3^{M\prime}\boldsymbol{H}_{it}^1 DI_{it} + \alpha_4^{M\prime}L_{it}^1 \\ &+ \alpha_5 L_{it}^2 + \alpha_6^{M\prime}t + \alpha_7^{M\prime}t^2 + \epsilon_{it}^{cons}\end{aligned} \tag{6.6}$$

The motivation for using this equation comes from the first order condition for optimal consumption each period, taking work decisions and health as given. Differentiating the preferences of Equation 5.1 with respect to consumption makes clear the log-additive rela-

[4]When I estimate the consumption regression in sample, I account for part-time work in the same way that full-time work enters the regression equation. I also estimate the consumption regression with calendar year effects.

tionship underlying the consumption regression.

$$\frac{\partial U(C_t, L_t^1, L_t^2; \boldsymbol{H}_t^1, \boldsymbol{H}_t^2)}{\partial C_t} = C_t^{-\gamma}(\exp(\mu M_t + \theta_1' \boldsymbol{H}_t^1 + \theta_2' \boldsymbol{H}_t^2 + \eta_1(M_t, \boldsymbol{H}_t^1) L_t^1 + \eta_2(\boldsymbol{H}_t^2) L_t^2))^{1-\gamma} \tag{6.7}$$

With full insurance, households would equalize marginal utility of consumption (Equation 6.7) across potential work and health states within each time period, and set them proportionally across time according to the relative size of the interest rate and discount factor, R/β. This implies that the work preference parameters defining the functions η_s (taking γ as known) should map to the appropriate coefficients on work (or work interacted with health) for person s in the log consumption regression (Equation 6.6). The same is true for the disability parameters θ_s and the coefficient on health status of person s. However, this identification result relies on full insurance against work and health risks for households.

In reality households are not perfectly insured against disability (or wage and work risks), so log consumption is contaminated with income effects with respect to both health and work. These can threaten the identification of the disutility parameters if the income effects break the injective mapping of model parameters to observed empirical moments. While it is difficult to formally dismiss this threat to identification, it can be instructive to focus one parameter at a time. As shown for select parameters in Figure A.1, preference parameters tend to have a monotonic relationship with their linked empirical moment, fixing all other model parameters.

Chapter 7

Parameter estimates and fit from simulated method of moments

The parameter values solved by SMM are intuitively sensible. Insurance against health risk is valuable because work-limiting disability reduces take-home earnings from work, increases the disutility of work, and increases the marginal utility of consumption. The extent to which a disability has each of these effects depends on its severity. The main effect of a moderate work-limiting disability is to increase the marginal utility of consumption. A moderate work-limiting disability does not substantially affect the disutility of work, but it does modestly reduce take-home income from work. A severe work-limiting disability, on the other hand, increases the disutility of work greatly, and in doing so, increases the marginal utility of consumption to a greater extent if the individual chooses to work.

Appendix G reports the full set of fitted parameter values and discusses their interpretation. Where I can make direct comparisons to the fitted values of Low and Pistaferri (2015) and Autor et al. (2019), the fitted values of my model look similar. A notable exception is that I find relatively low fixed costs of work associated with a severe disability compared to the value reported by Low and Pistaferri (2015). I instead load the work-related costs of a severe disability into the interaction between work and disability in household preferences.

I find that allowing that interaction is important for reproducing the consumption patterns observed in the data.

7.1 Fit of model to targeted moments

The complete set of empirical moments targeted by SMM are reported in Tables 6.3 and 6.4, and the model is able to fit them quite well.[1] Table 6.3 presents the moments related to work and DI receipt, while Table 6.4 presents the consumption moments.

Panels A and B of Table 6.3 show that the model replicates employment patterns as a function of age and health for both married and unmarried workers. However, simulated married workers with a severe work-limiting tend to work too little when younger and too much when older. The model fits the patterns in work behaviors for spouses (reported in Panel C) in a similar way; simulated work-limited spouses work too little when younger and too much when older. These are the only employment-related moments that economically deviate from their empirical values in Table 6.3, though it is reassuring to see that, as shown below, simulated spousal labor supply responds to disability of the reference person in a manner which is consistent with the data.

Empirical moments related to DI are reported in Panels D, E, and F of 6.3. The model matches how flows into DI increase with age and disability severity in the data. It also recreates the patterns in DI benefit receipt by marital status, health, and age well. DI stocks increases with age and work-limiting disability, and tends to be higher among unmarried workers than married workers.[2] However, simulated DI stocks among young people (espe-

[1] To form the objective function minimized by SMM, I use a diagonal weighting matrix which is equivalent to weighting each deviation between an observed and simulated moment by the inverse of the standard deviation of the observed moment. This is the weighting matrix used by Autor et al. (2019) and Blundell, Costa-Dias, Meghir, and Shaw (2016a), and it is motivated by evidence from Altonji and Segal (1996) that the asymptotically efficient weighting matrix has poor small-sample properties. To numerically solve the objective by SMM, I use a particle swarm algorithm.

[2] Unmarried person receiving DI more often than married persons is an apparent contradiction of the event study results, but it can be explained by the fact that unmarried persons tend to belong to lower household types which are more persistently exposed to health risks.

cially unmarried young people) fit the data less well than simulated DI flows. The mismatch between DI stocks and flows can, perhaps, be attributed to the highly simplified reassessment and continuation process for DI beneficiaries in this model. Alternatively, DI stocks among young unmarried people may be driven by individuals for whom disability arose in childhood, who are not accounted for in the model and who sample restrictions have failed to screen out of the data.

Table 6.4 shows how the model recreates responses of log consumption to work, marital status, health, and DI receipt which are observed in the data. Consumption increases with work, particularly if the worker suffers a work-limiting disability. Consumption decreases with disability for all households in the absence of DI receipt, though DI receipt offsets that decline for unmarried households. The consumption of married households does not increase with DI receipt, suggesting consumption insurance is not the main motive for their take-up of DI. The model is able to recreate these patterns, though it erroneously finds a decline in consumption with DI benefit receipt among severely work-limited married households.

7.2 Untargeted fit and implications of the model

The fitted model is able to reproduce an array of moments that are not targeted when fitting the parameters of the model. Figure 7.1 compares the simulated and empirical patterns in household wealth and consumption. Median simulated wealth compares well to the data until age 40. From age 40 to 45, the median simulated household tends to save too much relative to the data. The two begin to converge again at age 50 where the difference peaks, with the median household having $56,000 in savings in the model versus $40,000 in savings empirically. As we approach retirement, the data and model converge once again. Rather than reflecting a problem with the model per se, this pattern could reflect issues with wealth measurement in the PSID, which is known to underestimate wealth compared to the Survey of Consumer Finances (Pfeffer, Schoeni, Kennickell, and Andreski, 2016). Simulated average

Figure 7.1: *Simulated and empirical life cycle household outcomes*

(a) *Median Household Wealth* ***(b)*** *Log of Average Earnings*

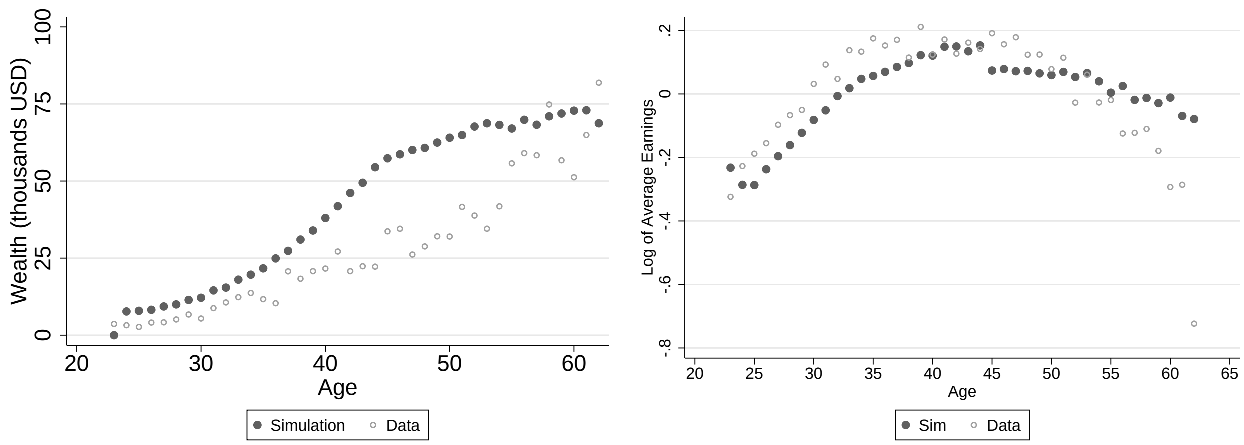

Note: Panel (a) plots wealth in 1996 US dollars. Panel (b) plots demeaned log average earnings.

earnings, on the other hand, replicates the hump-shaped life cycle curve observed in the data. Simulated earnings are low relative to the data in the 20's and 30's when individuals sometimes engage in part-time work (not allowed for in the model), and they are high relative to the data when individuals sometimes retire early starting in the late 50's.

The model also recreates the effects of disability onset documented in Section 4. Figure 7.2 compares the effects of disability onset for the reference person (of any severity) from empirical event studies to effects from simulated event studies. The model is able to recreate effects by marital status on DI receipt, log consumption, spousal work, and reference person work. The empirical estimates lie within 95 percent confidence intervals implied by the model for all outcomes.

When effects for moderate and severe disabilities are separated, the simulated effects of disability continue matching well to the empirical effects discussed in Section 4. Those results are reported in Figure A.2. The model recreates differences in effects for married versus unmarried households, of both a moderate disability and a severe disability. No coefficients in these event studies statistically differ from their empirical values.

Figure 7.2: *Empirical and simulated event studies of reference person disability*

(a) *Any severity, Single* ***(b)*** *Any severity, Married*

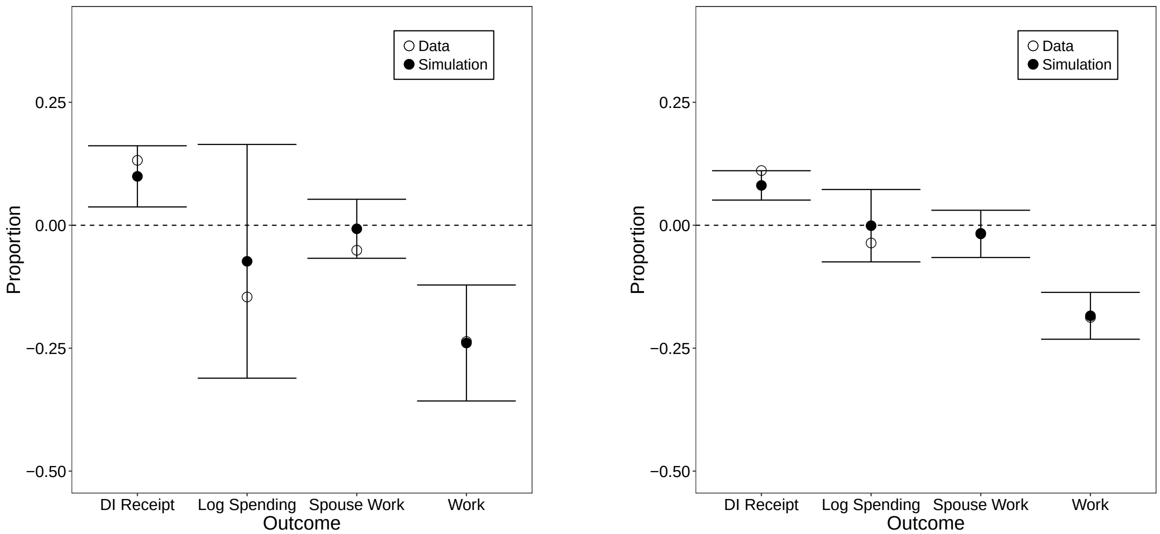

Error bars denote 95% confidence intervals for simulated parameter values. Confidence intervals are constructed by adjusting empirical standard error estimates by a factor $(1{+}1/S)$, where $S = 10$ is the number of times the PSID dataset is re-simulated (McFadden, 1989)

Table 7.1: *Empirical and simulated joint labor supply behaviors in married households*

Full-Time	Data	Simulation	Diff	SE
Neither Head nor Spouse Works	0.12	0.176	-0.056*	0.005
Only Spouse Works	0.056	0.076	-0.02*	0.003
Only Head Works	0.439	0.52	-0.081*	0.007
Both Head and Spouse Work	0.384	0.228	0.156*	0.008

Reference persons who report receiving only SSI (and not SSDI) benefits did not, by definition, work regularly before entering the DI system and are dropped from these calculations. Confident Stars denote statistical rejection of equality of simulated and empirical moments: * $p < 0.05$. Standard deviations are constructed by adjusting empirical standard error estimates (clustered at the household level) by a factor $(1 + 1/S)$, where $S = 10$ is the number of times the PSID dataset is re-simulated (McFadden, 1989)

While the model is targeted to fit only marginal moments related to the employment of the reference person and spouse, Table 7.1 shows that it is able to capture the qualitative joint work behaviors within the household well. The most common work arrangement among married couples is for the reference person to work and the partner to not work. However, reference persons and spouses often jointly work at the same time. Compared to the data, the model implies reference persons work on their own 8.1 percentage points too often, and work jointly with their spouse 15.6 percentage points too infrequently. Households more often do not work at all than do spouses work alone, in both the fitted model and the data.

In order to understand the magnitude of the fitted utility parameters, I compare elasticities implied by the model to values commonly found in the literature in Table 7.2. The model suggests non-employment and DI applications are rather inelastic to DI benefit amounts, when compared to what the literature would suggest. This is not surprising since the elasticity estimates in the literature are based on models which do not account for spousal labor supply. I find an elasticity of non-employment to DI benefits of 0.067, which lies on the low end of the range of values in the literature (Bound and Burkhauser, 1999; Haveman and Wolfe, 2000) and which compares closely to the value of 0.056 simulated by Low and Pistaferri (2015). I find an elasticity of DI applications to DI benefits of 0.261, which is also low (Low and Pistaferri cite a central value of 0.6) but which still fits well within the range of values cited by Bound and Burkhauser (1999). It lies below the value of 0.62 simulated by Low and Pistaferri (2015).

The model produces sensible elasticities of labor supply to earnings, on the other hand. Marshallian elasticities of labor supply for men and women (reference persons and their partners) in my model are 0.081 and virtually 0 respectively. The former is comparable to the central value of 0.06 reported by Keane (2011), and the latter is smaller than the values of 0.358, 0.42, and 0.475 found more recently by models from Autor et al. (2019), Blundell et al. (2016b), and Blundell et al. (2016a) respectively but within the range of the broader

Table 7.2: *Simulated and empirical labor supply elasticities*

	Sim	Lit (central)	Lit (range)
Nonwork to benefits	0.067	0.7	(0.06,0.93)
Apply to benefits	0.261	0.6	(0.2,1.3)
Head Work, Marsh	0.081	0.06	(-0.47,0.51)
Spouse Work, Marsh	0	0.4	(-0.2,1.36)
Head Work, Frisch	0.972	0.85	(0.03,2.75)
Spouse Work, Frisch	0.507	0.7	(0.03,2.35)

Notes: Elasticities are calculated in this model by finite-difference evaluated at a 10% increase in DI benefits, or 1% of a standard deviation permanent shock to log wages. Elasticities of non-employment to DI benefits are drawn from Table 16 of Bound and Burkhauser (1999) and Table 10 of Haveman and Wolfe (2000). The central value is borrowed from Low and Pistaferri (2015), which finds a value of 0.056 in its model. Elasticities of DI applications to DI benefits are drawn from Table 13 of Bound and Burkhauser (1999). The central value is borrowed from Low and Pistaferri (2015), which finds a value of 0.62 in its model. Labor supply elasticities for the reference person, and their central values, are drawn from Table 6 of Keane (2011). Frisch labor supply elasticities are also drawn from Blundell et al. (2016b) and Blundell et al. (2016a). The model of the former finds Frisch elasticities of 0.58 and 0.59. Spousal labor supply elasticities are drawn from Table 7 of Keane (2011), from Blundell et al. (2016b), and from Blundell et al. (2016a). The model of Blundell et al. (2016a) finds a female Marshallian elasticity of 0.42, and the model of Blundell et al. (2016b) finds values of 0.475. The model of Blundell et al. (2016b) finds a female Frisch elasticity of 0.88, and the model of Blundell et al. (2016a) finds a value of 0.63.

literature. The Frisch elasticity of labor supply for men is 0.972, lying well within the range of estimates reviewed by Keane (0.03-2.75) and compares well to more recent work. Blundell et al. (2016b) find a Frisch elasticity of labor supply of 0.58 for men and 0.88 for women. The Frisch elasticity for women is 0.507, comparable to the value of 0.63 found by Blundell et al. (2016a).

Chapter 8

Reforms to Disability Insurance

With the model fit to data, I now use it to assess the welfare value of alternative reforms to the DI system. The welfare value of a reform is measured by the household's willingness to pay (WTP) to implement it, paid by a proportional reduction to all household consumption. This measure of WTP allows the welfare gains of reform to be isolated from and compared against the reform's fiscal costs. The model implies that this payment mechanism causes no inefficiencies by distorting the non-consumption behaviors of households in the model. In this sense, it is equivalent to directly extracting utility from the household. It also has a natural cash value: the expected present value of the stream of forgone consumption flows.[1]

I consider the welfare implications of four alternative reforms to the DI system. To provide a benchmark, these reforms are compared against the welfare implications of a reform to the general social safety net: increasing the cash transfers provided by SNAP (formerly Food Stamps).

The first proposed reform is the national roll-out of a temporary DI program, which provides transfers to individuals during the period where they would typically be refraining from

[1] This contrasts with, for example, a lump sum measure of willingness to pay. Borrowing constraints cause the welfare cost (or gain) of lump sum transfers to highly vary across the life cycle, above and beyond variation related to the household's discount factor. A household at the beginning of life, for instance, has no savings and can mechanically pay no more than the household's income that year to implement any reform.

work and awaiting the outcome of their long-term DI applications. Temporary DI programs have existed in California, Rhode Island, New Jersey, and New York since the 1940s and have been recently implemented in 4 additional states as well as Washington DC (National Partnership for Women and Families, 2019). These programs provide benefits for 13 to 52 weeks, generally replacing roughly two-thirds of the worker's recent wages.[2] I approximate the key features common across these programs with the national temporary DI program in my model by assuming any worker with a moderate or severe work-limiting disability who is applying for long-term DI can meet the more lenient health certification standards for temporary DI benefits. Therefore, any DI applicant who suffers from a moderate or severe work-limiting disability will be allowed temporary DI benefits. An applicant who is disallowed from long-term DI cannot continue collecting temporary DI benefits, and must return to work before potentially re-initiating the application process.

A second reform considered in this section is to completely relax work restrictions on DI applicants. This reform is in the spirit of work incentive programs that have been implemented in the past (e.g., the Ticket to Work program, implemented in 1999). It can alternatively be viewed as approximating, for applicants but not beneficiaries, the effects of the gradual increase over time in the Substantial Gainful Activity threshold.[3] Like a national temporary DI program, this reform will reduce the cost of acquiring DI benefits by providing a source of income for applicants.

The final two reforms involve adjusting the leniency of the DI application process (the probability with which applicants are allowed onto DI), and adjusting the generosity of DI benefits. These are historically implemented reforms that have been assessed by Low and Pistaferri (2015). In revisiting these reforms, I account for their interactions with household self-insurance. The former reform approximates a number of explicit and implicit changes

[2]New York's program is one notable exception, only providing benefits up to $170 weekly.

[3]SGA, the maximum income an applicant or beneficiary is allowed to earn while remaining eligible for DI, is indexed to national mean wages. With mean wages growing faster than median wages, this threshold is increasing at a faster rate than earnings for most workers.

to DI, taking place from the 1980's through the 2010's.[4] It can also represent adjusting allowance guidelines for DI examiners, holding fixed the behavior of applicants at the pre-reform equilibrium (conditional on applying).[5] The latter approximates the smooth increase over time in potential DI benefits relative to potential earnings for most workers, owing to rising wage inequality and the indexation of the DI benefit formula to mean wage growth.

For the two reforms that involve adjusting cash payments (adjusting the generosity of DI benefits and the benchmark reform to SNAP), my main results study the effect of increasing all potential payments by 10 percent. For adjusting the leniency of the application process for DI, my main results study the effect of a proportional 10 percent increase in all allowance probabilities.

8.1 The economic costs of disability insurance reform

Table 8.1 summarizes the effect of each reform on measures of DI take-up and economic output. In making the DI system more attractive, each reform disincentivizes work to a certain extent, as shown by the reductions in employment in Column 1. Relative to welfare gains provided by the reform, the disemployment costs of temporary DI are low compared to other reforms, at 1.31 years of lost employment per $100,000 in welfare gains. The reform that relaxes work restrictions on applicants is one exception, providing an additional 75 years of employment per $100,000 in welfare gains. This is because relaxing work restrictions on applicants slightly increases overall employment, and because the welfare gains from that reform are small.

[4]The retrenchment of DI from 1979 to 1984, the subsequent liberalization of DI from 1984-1998, and the retraining of judges in the 2010's can all be viewed as adjusting the probability with which applicants are allowed DI benefits Autor and Duggan (2003); Schwartz (2018). The closing of SSA offices from 2009 to 2014 can also be seen as (perhaps unintentionally) reducing allowance probabilities Deshpande and Li (2019).

[5]Proportionally increasing the leniency of DI allowance is equivalent to reducing the time examiners spend on each case in the following model of the application process: DI applicants send exponentially distributed signals of good health, truncated based on their health status. If an examiner sees a signal of positive health while reviewing an applicant's case, the application is rejected. Otherwise, the applicant is allowed DI benefits.

The declines in employment from expanding DI reflect moral hazard, but they also reflect income effects of benefit receipt among poorly self-insured workers and changes in the precautionary behaviors of households. To isolate the moral hazard costs of DI reforms, the final column of Table 8.1 shows how reforms change the percent of applicants and beneficiaries for whom the work restrictions of DI bind. While the temporary DI program has low disemployment costs relative to its welfare gains, it comes at a higher moral hazard cost than other reforms. It increases the proportion of DI beneficiaries who would work (if DI allowed it) by 0.64 percentage points. Other reforms tend to increase moral hazard a small amount. The reform that allows DI applicants to work is once again an exception, as lifting that restriction reduces moral hazard by 12.17 percentage points

The distortionary cost of each reform's welfare gains in terms of economic output and years spent applying for or receiving DI are summarized in the intermediary columns of Table 8.1. The output costs of reform tend to be lower for the national temporary DI program than for other reforms. Column 2 shows that a temporary DI program decreases total output by $0.53 for each dollar of ex-ante welfare value it provides to households. Columns 3 show that it causes households to spend 3.93 additional years applying for DI for each $100,000 in ex-ante welfare gains, which is a larger application response than created by increasing the generosity of DI benefits (1.5 years) or increasing the generosity of SNAP benefits (0.03 years). Increasing the leniency of DI reduces the number of years individuals spend (re-)applying by 1.31 years per $100,000 in ex-ante welfare gains.

8.2 The social benefits and costs of disability insurance reform

To consider the welfare effects of the alternative reforms, Column 1 of Table 8.2 presents the ex-ante WTP, where households take expectation (fully behind a veil of ignorance) over the

Figure 8.1: *Effect of self-insurance mechanisms on willingness to pay for DI reforms*

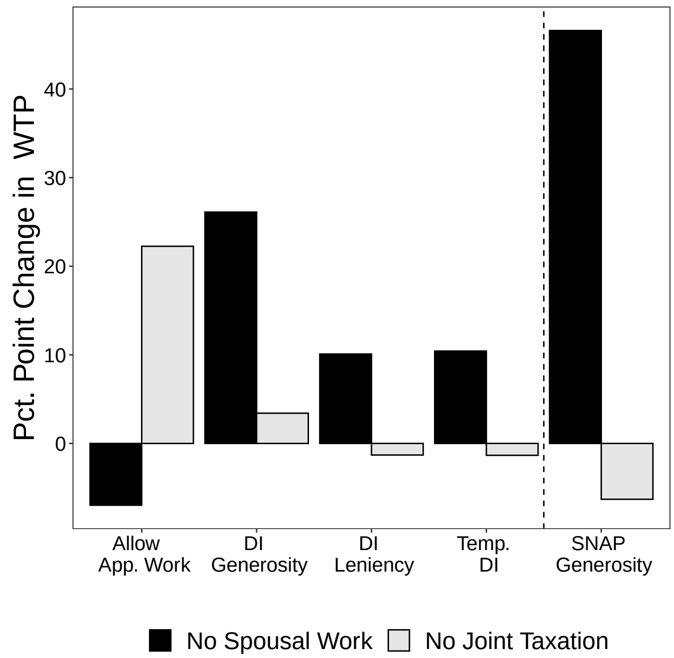

Table 8.1: *The economic costs of reforms to disability insurance*

	(1)	(2)	(3)	(4)		(5)
	Costs per $100k Welfare Gain					
				DI Rolls		
Reform	Employment	Output ($100k)	DI Applications	Single	Married	Moral Hazard
Temp. DI	-1.31	-0.53	3.93	0.26	-0.04	0.64
Allow App. Work	75.65	23.47	344,874.58	-89.75	0.00	-12.17
DI Leniency	-4.97	-0.97	-1.31	11.94	10.80	-0.12
DI Generosity	-2.19	-0.75	1.50	1.43	3.64	0.21
SNAP Generosity	-3.29	-0.59	0.03	-0.03	0.01	0.05

Notes: Marginal value of public funds (MVPF) is calculated by dividing the cash value of household WTP for a reform by its fiscal costs. The effect of DI generosity is measured by increasing potential cash payments to DI beneficiaries by 10 percent. The effect of Welfare generosity is measured by increasing potential cash payments to SNAP beneficiaries by 10 percent. The effect of DI leniency is measured by increasing all allowance probabilities for DI applicants proportionally by 10 percent. Columns (3)-(5) report the expected change in years spent applying, collecting DI, or working per $100,000 in ex-ante welfare gains provided to households.

Table 8.2: *Social benefits and costs of reforms to disability insurance*

	(1)	(2)	(3)			(4)			(5)
			WTP by Type			Singlehood Premium			
Reform	WTP	Costs	Low	Mid	High	Low	Mid	High	**MVPF**
Temp. DI	3.19	1.81	6.81	2.67	1.24	16.7%	13.9%	8.6%	**1.77**
Allow App. Work	0.00	0.00	0.00	0.00	0.00	135.3%	-35.8%	122.4%	**-0.07**
DI Leniency	0.19	0.13	0.58	0.10	0.03	-0.9%	0.0%	-5.7%	**1.43**
DI Generosity	0.31	0.25	0.95	0.14	0.06	-3.3%	-2.8%	-6.1%	**1.24**
SNAP Generosity	0.90	0.39	1.54	1.16	0.15	23.9%	23.2%	-20.0%	**2.32**

Notes: Marginal value of public funds (MVPF) is calculated by dividing the cash value of household WTP for a reform by its fiscal costs. The effect of DI generosity is measured by increasing potential cash payments to DI beneficiaries by 10 percent. The effect of Welfare generosity is measured by increasing potential cash payments to SNAP beneficiaries by 10 percent. The effect of DI leniency is measured by increasing all allowance probabilities for DI applicants proportionally by 10 percent. Columns (3)-(5) report the expected change in years spent applying, collecting DI, or working per $100,000 in ex-ante welfare gains provided to households.

potential type they could be assigned, initial marital status, and the various life cycle risks.[6] The national temporary DI program, for instance, generates the largest WTP, equal to 3.19 percent of annual household consumption. However, differences in WTP across reforms partly reflect differences in their fiscal scope. Column 2 presents the fiscal costs of each reform, showing that the national temporary DI program is also the most expensive reform to implement (with fiscal costs equal to 1.81 percent of household consumption).

Comparing the WTP and fiscal costs of each reform suggests that, though each reform has economic costs, they are all nevertheless welfare-improving. This result is consistent with Low and Pistaferri (2015), and it is made explicit by examining the marginal value of public funds (MVPF), reported in Column 5 of Table 8.2. The MVPF is calculated by dividing

[6]This ex-ante measure of WTP is monotonically linked to the proportional increase in expected utility caused by the reform. Letting $\text{EV}_{\text{baseline}}$ and $\text{EV}_{\text{reform}}$ denote expected utility under the baseline model and under the reform respectively, WTP is defined as:

$$\text{WTP} = \left(\frac{\text{EV}_{\text{reform}}}{\text{EV}_{\text{baseline}}}\right)^{\frac{1}{\gamma-1}} - 1$$

household WTP for a reform (Column 1) by its fiscal costs (Column 2), and it represents the utility that households gain, in dollar terms, for each dollar spent on a reform. All reforms have a MVPF larger than 1, meaning that each would be welfare-improving if funded with a non-distortionary tax. However, there is substantial variation across reforms in welfare gains relative to fiscal costs. The most valuable DI reform is the national temporary DI program with a MVPF of 1.77. The DI reform with the next highest MVPF is increasing the leniency of DI allowance (1.43), followed by increasing the generosity of payments to DI beneficiaries (1.24). Households would prefer spending marginal resources on the benchmark reform to SNAP benefits over any of the DI reforms, though, as it has a MVPF of 2.32.

These welfare gains come disproportionately from households that are worse-insured against disability risk in the baseline model. Column Group 3 of Table 8.2 shows how WTP differs conditional on household type. Low-type households (with low potential earnings, low probability of marriage, and high health risks) have a WTP for temporary DI that is over twice as high as the ex-ante WTP, whereas the WTP of mid-type and high-type households are both below ex-ante WTP. Similar patterns hold for other reforms. However, Column Group 4 of Table 8.2 reports differences in WTP by marital status within type, showing that temporary DI is the only reform that is also progressive with respect to marital status for all household types. The benchmark reform to SNAP benefits is not progressive with respect to marital status among high-type households, but it is even more progressive than DI reforms among low-type and mid-type households.

It is also the case that the DI reforms with the highest MVPF are precisely the ones which reduce adverse selection into DI the most. This is clear from returning to Table 8.1 and examining how reforms change DI receipt among unmarried versus married workers in Column 4. Implementing a national temporary DI program and increasing the leniency of the DI system are the only two reforms which cause more unmarried workers than married workers to be added to the DI rolls. The benchmark reform to SNAP, though, causes some unmarried workers to substitute from DI receipt to relying on the general social safety net.

8.3 Insurance and redistribution motives for reform

Confirming Low and Pistaferri (2015), these model-based results indicate that households are willing to cover the costs of several expansionary reforms to the DI system. However, they contrast with results coming from a nascent but influential literature in public finance that has used envelope theorem logic to compare the ex-post willingness to pay of beneficiaries to the fiscal costs of policy changes (Finkelstein and Hendren, 2020; Hendren, 2016; Hendren and Sprung-Keyser, 2020). That literature has found that private benefits from expansionary DI reforms are consistently smaller their fiscal costs Hendren and Sprung-Keyser (2020).

Echoing the point made by Hendren (2021) on reforms to health insurance markets, the key reason that model-based welfare gains from DI reform are so high relative to their fiscal costs is that they provide valuable insurance to households. However, while the statutory purpose of DI is to protect households against disability risk, the insurance value of DI reform may not come from disability alone. It may also insure against non-health risks that both directly affect households and that limit their capacity to self-insure against disability. To quantify the value of the insurance provided by DI reform, and the extent to which it comes from health versus non-health risks, Appendix H decomposes each reform into five intermediary policies that reallocate income across households, following Shaw (2014). The first two intermediary policies capture how a reform redistributes income (between households and across the adult life cycle) and the remaining three intermediaries capture the insurance it provides to adults (against disability, singlehood, and productivity risk in adulthood). The components of the reforms that relate to insurance against adulthood risks explain 50 percent of WTP for temporary DI, 86 percent of WTP for increasing leniency of DI allowance, 25 percent of WTP for increasing the generosity of DI benefits, and 91 percent of WTP for the benchmark reform to SNAP benefits. Rather than more effectively insuring household income against negative shocks, the insurance value of reforms comes mostly from relieving households from engaging in costly work when suffering from disability and when earnings are low.

The role of both health and non-health risks in driving the value of DI reform can be seen another way, by showing how WTP changes as particular risk mechanisms are shut down one at a time. This is shown in Figure A.3, which illustrates that WTP for reforms is substantially lower if disability risks or divorce risks are reduced. One important distinction between this exercise and the Shaw decomposition, though, is that the current exercise does not purely capture the value of insurance against these risks, since shutting them down also affects households' expected income.

8.4 Quantifying the importance of spousal labor supply and household self-insurance

Together, the results so far indicate that marital status interacts in important ways with DI reform: initial marital status drives substantial heterogeneity in ex-ante welfare gains from reforms, and those welfare gains come in no small part from the insurance that they provide against marital risks over the life cycle. However, these results do not necessarily imply that accounting for household self-insurance arising from spousal labor supply is important for fully assessing the costs and benefits of DI reform. For instance, married and unmarried households in the model also face different costs of work, different prices for (equivalized) consumption, and different allowance probabilities when they apply for DI.

To explicitly gauge the importance of accounting for spousal labor supply and household self-insurance, Figure 8.1 shows how WTP for each reform changes as we remove particular household self-insurance mechanisms from the model. Failing to account for spousal labor supply would cause the welfare gains from increasing the generosity of the DI benefits to increase by 26 percent, whereas WTP for increasing the leniency of the DI allowance process or for a temporary DI program would increase by 10 percent. These DI reforms are less sensitive to spousal labor supply than the benchmark reform to SNAP benefits, the welfare gains of which would increase by 60 percent. The welfare gains of allowing DI applicants to work, by contrast, would decrease by 7 percent if the model failed to account for spousal labor supply.

Chapter 9

Conclusion

Given the size of public DI programs in the United States, the value of DI reform is an important policy question. This paper provides reduced-form evidence that workers adversely select into DI, in the sense that workers with stronger self-insurance capacity (married workers) are more likely to take up DI benefits in response to a work-limiting event. Despite leaving work at greater rates due to disability onset, married workers consequently see smaller declines in consumption and wealth compared to unmarried workers.

To explore the welfare implications these findings have for DI reform, I estimate a life cycle model which accounts for household self-insurance and selection into DI. The model-based results suggest expansions to DI tend to be welfare-improving, especially reforms that reduce the cost of acquiring DI benefits, such as a national temporary DI program. These reforms provide substantial insurance value to individuals, which can play an important role in driving high willingness to pay. While failing to account for the insurance value of DI reforms may cause one to erroneously conclude that households are unwilling to fund those reforms, failing to account for household self-insurance may cause one to overestimate the their insurance value.

When considering the interpretation of my results, I emphasize three caveats. First, my model abstracts away from Medicare, Medicaid, private health insurance, and their relationships with work and DI. Private health insurance in the United States is closely tied to employment and marriage, and health insurance provision through DI is another incentive for take-up of the program. Health insurance is also another way in which the application process of DI may be disproportionately costly for unmarried workers; Unmarried workers can neither rely on their own employer-provided health insurance during the application process, nor on health insurance coverage provided by a spouse. In that sense, the model should underestimate the value of DI benefits, but also underestimate the opportunity cost of applying for unmarried workers.

Second, I do not fully model the application and appeals process of DI. Individuals can, in reality, choose to initiate an application and then choose whether to continue the process based on the signals they receive about their allowance probability as review and appeals proceed. The static, year-long application decision in my model likely overestimates application costs for some workers, though it roughly matches the average wait time for benefits.

Lastly, this study focuses on the value of DI for male workers with lower educational attainment. While this is a subsample of the U.S. labor force that is especially prone to DI receipt, and perhaps the group in which the welfare effects of DI may be most concentrated, the restriction does limit the scope of my welfare analysis. The interactions of female workers with DI is an understudied topic in the literature that warrants further attention.

Chapter 10

References

Adda, J., J. Banks, and H.-M. von Gaudecker (2009): "The Impact of Income Shocks on Health: Evidence from Cohort Data," *Journal of the European Economic Association*, 7, 1361–1399. 39

Alatas, V., A. Banerjee, R. Hanna, B. A. Olken, R. Purnamasari, and M. Wai-Poi (2016): "Self-Targeting: Evidence from a Field Experiment in Indonesia," *Journal of Political Economy*, 124, 371–427. 6

Altonji, J. G. and L. M. Segal (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14, 353–366, publisher: [American Statistical Association, Taylor & Francis, Ltd.]. 54

Andreski, P., G. Li, M. Z. Samancioglu, and R. Schoeni (2014): "Estimates of Annual Consumption Expenditures and Its Major Components in the PSID in Comparison to the CE," *American Economic Review*, 104, 132–135. 85

Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019): "Disability Benefits, Consumption Insurance, and Household Labor Supply," *American Economic Review*, 109, 2613–2654. 1, 2, 5, 35, 40, 53, 54, 58, 113

Autor, D. H. and M. G. Duggan (2003): "The Rise in the Disability Rolls and the Decline in Unemployment," *The Quarterly Journal of Economics*, 118, 157–206. 6, 63

——— (2006): "The Growth in the Social Security Disability Rolls: A Fiscal Crisis Unfolding," *Journal of Economic Perspectives*, 20, 71–96. 5, 6

Autor, D. H., N. Maestas, K. J. Mullen, and A. Strand (2015): "Does Delay Cause Decay? The Effect of Administrative Decision Time on the Labor Force Participation and Earnings of Disability Applicants," Working Paper 20840, National Bureau of Economic Research. 5, 9

Benítez-Silva, H., M. Buchinsky, H. M. Chan, S. Cheidvasser, and J. Rust

(2004): "How Large Is the Bias in Self-Reported Disability?" *Journal of Applied Econometrics*, 19, 649–670. 87

Blundell, R., M. Costa-Dias, C. Meghir, and J. Shaw (2016a): "Female Labor Supply, Human Capital, and Welfare Reform," *Econometrica*, 84, 1705–1753. 54, 58, 59, 60

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016b): "Consumption Inequality and Family Labor Supply," *American Economic Review*, 106, 387–435. 39, 58, 59, 60

Bonhomme, S., T. Lamadon, and E. Manresa (2019): "Discretizing Unobserved Heterogeneity," Becker Friedman Institute for Economics Working Paper No. 2019-16, University of Chicago. 41

Bound, J. (1989): "The Health and Earnings of Rejected Disability Insurance Applicants," *American Economic Review*, 79, 482–503. 5

Bound, J. and R. V. Burkhauser (1999): "Economic analysis of transfer programs targeted on people with disabilities," in *Handbook of Labor Economics*, Elsevier, vol. 3, 3417–3528. 58, 59, 86

Bound, J. and T. Waidmann (2007): "Estimating the Health Effects of Retirement," Tech. Rep. 2007-168, University of Michigan. 39

Chen, S. and W. van der Klaauw (2008): "The work disincentive effects of the disability insurance program in the 1990s," *Journal of Econometrics*, 142, 757–784. 9, 99, 114

de Chaisemartin, C. and X. D'Haultfoeuille (2019): "Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects," Working Paper 25904, National Bureau of Economic Research. 28

de Jong, P., M. Lindeboom, and B. van der Klaauw (2011): "Screening Disability Insurance Applications," *Journal of the European Economic Association*, 9, 106–129. 7

Deshpande, M., T. Gross, and Y. Su (2021): "Disability and Distress: The Effect of Disability Programs on Financial Outcomes," *American Economic Journal: Applied Economics*, 13, 151–178. 5

Deshpande, M. and Y. Li (2019): "Who Is Screened Out? Application Costs and the Targeting of Disability Programs," *American Economic Journal: Economic Policy*, 11, 213–248. 7, 63

Deshpande, M. and L. Lockwood (2020): "Beyond Health: Non-Health Risk and the

Value of Disability Insurance," *Working Paper*. 6, 86

Duggan, M. and S. A. Imberman (2009): "Why Are the Disability Rolls Skyrocketing? The Contribution of Population Characteristics, Economic Conditions, and Program Generosity," in *NBER Chapters*, National Bureau of Economic Research, Inc, 337–379. 6

Finkelstein, A. and N. Hendren (2020): "Welfare Analysis Meets Causal Inference," *Journal of Economic Perspectives*, 34, 146–167. 68

Finkelstein, A. and M. J. Notowidigdo (2019): "Take-Up and Targeting: Experimental Evidence from SNAP*," *The Quarterly Journal of Economics*, 134, 1505–1556. 6

French, E. and J. Song (2014): "The Effect of Disability Insurance Receipt on Labor Supply," *American Economic Journal: Economic Policy*, 6, 291–337. 5, 100

Garner, T., G. Janini, L. Paszkiewicz, and M. Vendemia (2006): "The CE and the PCE: a comparison," *Monthly Labor Review*, 129, 20. 85

Haveman, R. and B. Wolfe (2000): "Chapter 18 The economics of disability and disability policy," in *Handbook of Health Economics*, Elsevier, vol. 1, 995–1051. 58, 59

Heckman, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47, 153–161, publisher: [Wiley, Econometric Society]. 44

Hendren, N. (2016): "The Policy Elasticity," *Tax Policy and the Economy*, 30, 51–89, publisher: The University of Chicago Press. 68

——— (2021): "Measuring Ex Ante Welfare in Insurance Markets," *The Review of Economic Studies*, 88, 1193–1223. 68

Hendren, N. and B. Sprung-Keyser (2020): "A Unified Welfare Analysis of Government Policies," *Quarterly Journal of Economics*, 135, 1209–1318. 5, 68

Hoynes, H. W. and D. W. Schanzenbach (2009): "Consumption Responses to In-Kind Transfers: Evidence from the Introduction of the Food Stamp Program," *American Economic Journal: Applied Economics*, 1, 109–139. 102

Keane, M. P. (2011): "Labor Supply and Taxes: A Survey," *Journal of Economic Literature*, 49, 961–1075. 58, 59, 60

Kim, B. and C. J. Ruhm (2012): "Inheritances, health and death," *Health Economics*, 21, 127–144. 39

Kostol, A. R. and M. Mogstad (2015): "Earnings, Disposable Income, and Consumption of Allowed and Rejected Disability Insurance Applicants," Tech. Rep. 8893, Institute for the Study of Labor (IZA). 5

Kreider, B. (1999): "Latent Work Disability and Reporting Bias," *The Journal of Human Resources*, 34, 734–769. 87

Liebman, J. B. (2015): "Understanding the Increase in Disability Insurance Benefit Receipt in the United States," *Journal of Economic Perspectives*, 29, 123–150. 5, 6, 8

Low, H. and L. Pistaferri (2015): "Disability Insurance and the Dynamics of the Incentive Insurance Trade-Off," *American Economic Review*, 105, 2986–3029. 1, 3, 4, 11, 13, 35, 40, 42, 43, 44, 53, 58, 59, 62, 66, 68, 85, 86, 99, 100, 101, 103, 104, 105, 110, 113

Maestas, N., K. J. Mullen, and A. Strand (2013): "Does Disability Insurance Receipt Discourage Work? Using Examiner Assignment to Estimate Causal Effects of SSDI Receipt," *American Economic Review*, 103, 1797–1829. 5

McFadden, D. (1989): "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica*, 57, 995–1026. 48, 49, 50, 57

Meyer, B. D. and W. K. C. Mok (2019): "Disability, earnings, income and consumption," *Journal of Public Economics*, 171, 51–69. 6, 28, 86

National Partnership for Women and Families (2019): "State Paid Family and Medical Leave Insurance Laws," Tech. rep., National Partnership for Women and Families. 62

Nichols, A. L. and R. J. Zeckhauser (1982): "Targeting Transfers through Restrictions on Recipients," *The American Economic Review*, 72, 372–377, publisher: American Economic Association. 6

Nichols, D., E. Smolensky, and T. N. Tideman (1971): "Discrimination by Waiting Time in Merit Goods," *American Economic Review*, 61, 312–23, publisher: American Economic Association. 6

Parsons, D. O. (1991): "Self-Screening in Targeted Public Transfer Programs," *Journal of Political Economy*, 99, 859–876, publisher: University of Chicago Press. 7

Pfeffer, F. T., R. F. Schoeni, A. Kennickell, and P. Andreski (2016): "Measuring wealth and wealth inequality: Comparing two U.S. surveys," *Journal of Economic and Social Measurement*, 41, 103–120. 55, 85

Schwandt, H. (2018): "Wealth Shocks and Health Outcomes: Evidence from Stock Market

Fluctuations," *American Economic Journal: Applied Economics*, 10, 349–377. 39

Schwartz, N. D. (2018): "Disability Applications Plunge as the Economy Strengthens (Published 2018)," *The New York Times*. 63

Shaw, J. (2014): "The redistribution and insurance value of welfare reform," IFS Working Papers W14/21, Institute for Fiscal Studies (IFS), London. 68, 117

Smith, J. P. (2004): "Unraveling the SES: Health Connection," *Population and Development Review*, 30, 108–132. 39

Social Security Administration (2019): "Annual Statistical Report on the Social Security Disability Insurance Program, 2018," Tech. rep., Social Security Administration. 1, 86

Stern, S. (1989): "Measuring the Effect of Disability on Labor Force Participation," *The Journal of Human Resources*, 24, 361–395. 39

A Appendix Figures

Table A.1: *Health Characteristics by Self-Reported Subjective Disability Status*

	Healthy	Mod. Dis.	Sev. Dis.
Panel A: Functional Limitations			
Walking or Climbing Stairs (1986, 2017)	0.0514***	0.550***	0.775***
	(0.00508)	(0.0379)	(0.0358)
Bending or Lifting (1986, 2017)	0.0742***	0.626***	0.812***
	(0.00601)	(0.0372)	(0.0335)
Driving (1986)	0.00471**	0.127***	0.378***
	(0.00210)	(0.0332)	(0.0568)
Seeing (1986)	0.0235***	0.0784***	0.162***
	(0.00465)	(0.0268)	(0.0431)
Preparing Own Meals (2003-2017)	0.00258***	0.0424***	0.170***
	(0.000622)	(0.0107)	(0.0213)
Personal Shopping (2003-2017)	0.00421***	0.0627***	0.241***
	(0.000929)	(0.0116)	(0.0245)
Managing Money (2003-2017)	0.0134***	0.101***	0.189***
	(0.00166)	(0.0156)	(0.0223)
Using the Telephone (2003-2017)	0.00177***	0.0203***	0.0867***
	(0.000488)	(0.00596)	(0.0179)
Doing Heavy Housework (2003-2017)	0.0148***	0.319***	0.672***
	(0.00163)	(0.0244)	(0.0247)
Doing Light Housework (2003-2017)	0.00543***	0.0867***	0.262***
	(0.000898)	(0.0170)	(0.0233)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level.

Table A.1: *(Continued) Health Characteristics by Self-Reported Subjective Disability Status*

	Healthy	Mod. Dis.	Sev. Dis.
Panel B: Medical Diagnoses			
Arthritis (1999-2017)	0.0778***	0.298***	0.448***
	(0.00505)	(0.0229)	(0.0276)
Asthma (1999-2017)	0.0634***	0.140***	0.119***
	(0.00525)	(0.0184)	(0.0198)
Hypertension (1999-2017)	0.233***	0.467***	0.564***
	(0.00915)	(0.0266)	(0.0305)
Cancer (1999-2017)	0.0142***	0.0342***	0.0503***
	(0.00239)	(0.00788)	(0.00976)
Diabetes (1999-2017)	0.0804***	0.187***	0.247***
	(0.00619)	(0.0220)	(0.0258)
Psychiatric Problems (1999-2017)	0.0275***	0.167***	0.278***
	(0.00279)	(0.0198)	(0.0263)
Heart Attack (1999-2017)	0.0220***	0.118***	0.158***
	(0.00326)	(0.0181)	(0.0218)
Heart Disease (1999-2017)	0.0213***	0.118***	0.177***
	(0.00264)	(0.0160)	(0.0221)
Lung Disease (1999-2017)	0.0200***	0.0769***	0.152***
	(0.00222)	(0.0125)	(0.0213)
Stroke (1999-2017)	0.00659***	0.0627***	0.109***
	(0.00127)	(0.0144)	(0.0192)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level.

Table A.2: *HRS sample summary statistics*

	Single	Partnered	Partnered-Single
Age	58.40***	57.79***	-0.609***
	(0.0484)	(0.0350)	(0.0583)
Has HS Degree	0.717***	0.723***	0.00523
	(0.00938)	(0.00572)	(0.0104)
White	0.585***	0.724***	0.139***
	(0.0102)	(0.00562)	(0.0111)
# Adults in Household	1.090***	2.007***	0.917***
	(0.0100)	(0.00233)	(0.00993)
# Living Children	2.278***	3.415***	1.138***
	(0.0450)	(0.0251)	(0.0495)
Has Back Problem	0.212***	0.256***	0.0439***
	(0.00518)	(0.00335)	(0.00609)
Has Hypertension	0.399***	0.440***	0.0415***
	(0.00863)	(0.00531)	(0.00990)
Has Diabetes	0.150***	0.174***	0.0240***
	(0.00638)	(0.00419)	(0.00743)
Has Cancer	0.0466***	0.0523***	0.00575
	(0.00350)	(0.00234)	(0.00408)
Has Lung Disease	0.0802***	0.0694***	-0.0109**
	(0.00459)	(0.00268)	(0.00514)
Has Heart Problem	0.134***	0.157***	0.0230***
	(0.00599)	(0.00398)	(0.00699)
Had a Stroke	0.0558***	0.0435***	-0.0123***
	(0.00405)	(0.00215)	(0.00444)
Has Psych. Problem	0.165***	0.102***	-0.0627***
	(0.00681)	(0.00335)	(0.00735)
Has Arthritis	0.317***	0.363***	0.0462***
	(0.00821)	(0.00511)	(0.00940)
Has Work Limitation	0.314***	0.246***	-0.0682***
	(0.00767)	(0.00428)	(0.00866)
# Households/Wave	666.8571	2015.357	
# Waves/Household	3.046997	3.604369	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level. Wealth is reported in thousands of 1996 US dollars. Partners include married spouses and non-marital partners.

Table A.3: *Validating k-means type assignments with the fitted model*

		Type (K-means)		
		Low	Moderate	High
	Low	334	42	34
Type (Model)	Moderate	219	1053	82
	High	81	89	1271

Columns divide sample households according to the type assigned to them by k-means clustering. Rows divide sample households according to the type that is most likely to generate the average characteristics used in k-means classification (log wages and employment rate when healthy), based on the fitted model. The two type classifications agree for 83 percent of households.

Figure A.1: *Local identification of select parameters by log consumption regression*

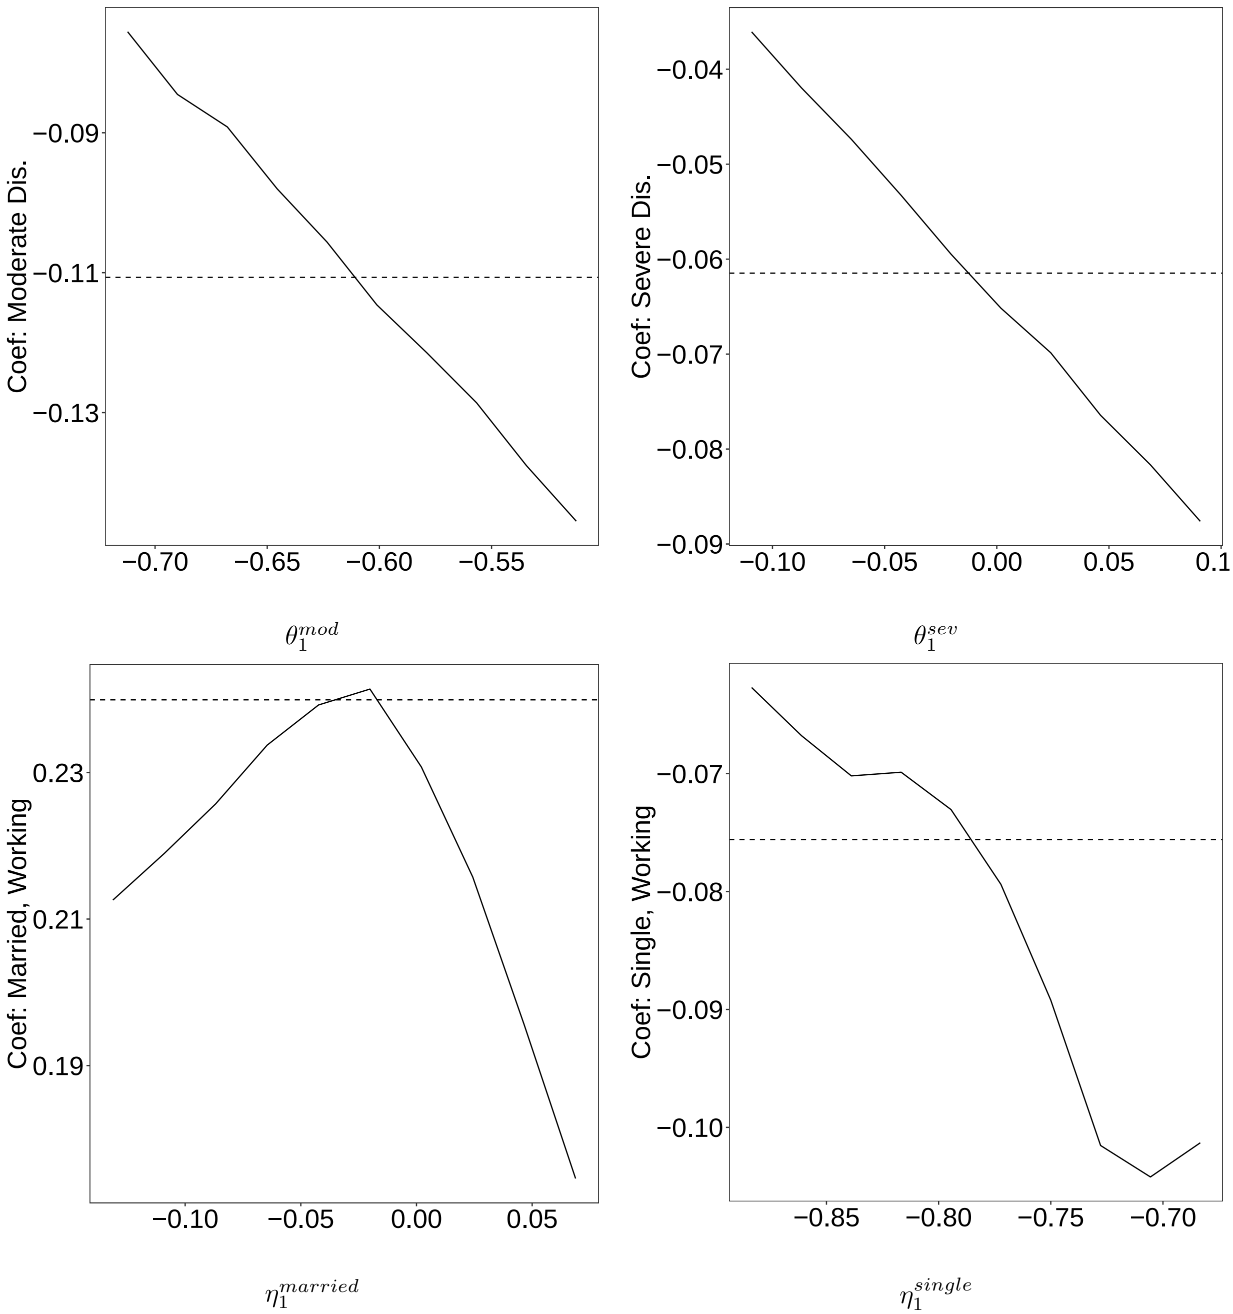

Figure A.2: *Empirical and simulated event studies of reference person disability*

(a) *Severe Disability, Single*

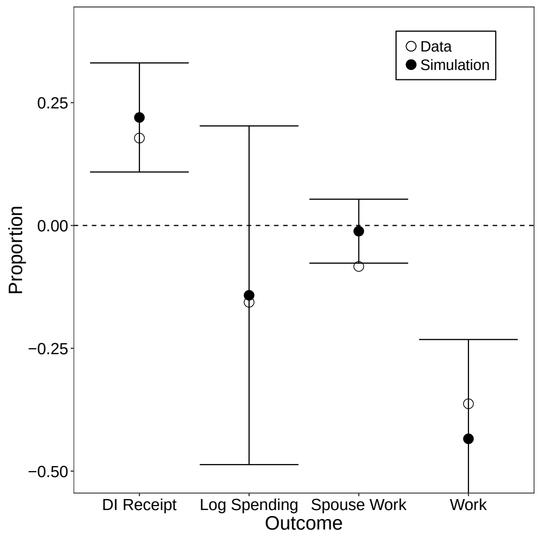

(b) *Severe Disability, Married*

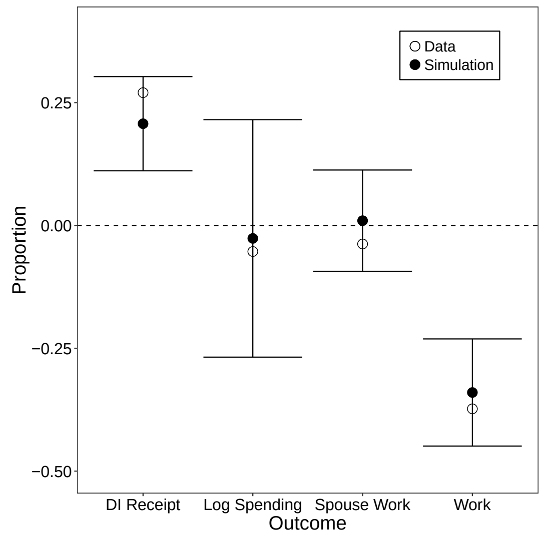

(c) *Moderate Disability, Single*

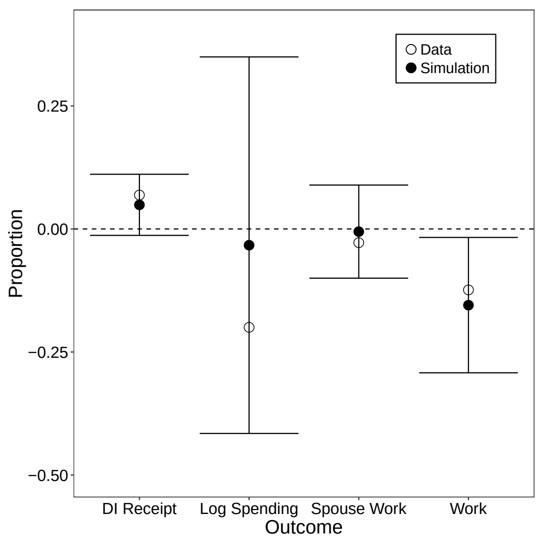

(d) *Moderate Disability, Married*

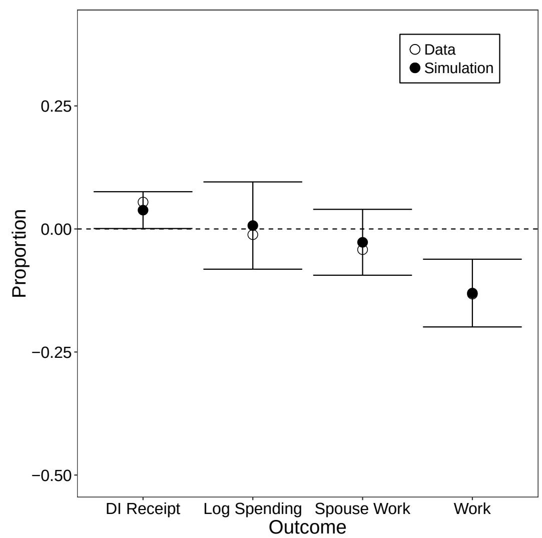

Error bars denote 95% confidence intervals for simulated parameter values.

Figure A.3: *Effect of risks on willingness to pay for DI reforms*

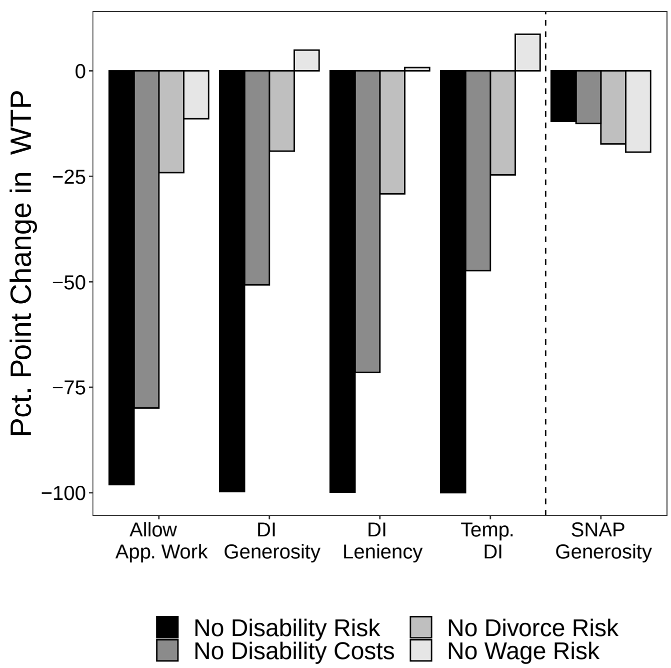

Notes: shutting down disability risk and shutting down wage risk cause the WTP for relaxing work restrictions to increase by 150 percent and 99 percent respectively. This is because the gains from working while applying for DI depend on potential wages. One of the costs of disability is that it decreases net potential earnings. Shutting down disability risk removes all WTP for DI reforms (except relaxing work restrictions on applicants) because allowance probabilities for healthy workers are sufficiently low that they rarely choose to apply.

B Definitions of Key PSID Measures

B.1 Consumption

Limited consumption information is available for all waves of the PSID. I follow Low and Pistaferri (2015) in using more comprehensive measures of household consumption which are only first collected in the 1999 survey wave. This includes information on food, health care, housing and rent, utilities, transportation, education, and child care spending over the year.[1] Of these categories, only food spending and rent are available in years prior to 1999. While additional categories are introduced in later years, I restrict to the previous measures for consistency over time.[2] Average total consumption based on these measures compares closely to the more detailed consumption information in the Consumer Expenditure Survey, and underestimates aggregated Personal Consumption Expenditures data of the National Income and Product Accounts by about 35% (Andreski, Li, Samancioglu, and Schoeni, 2014; Garner, Janini, Paszkiewicz, and Vendemia, 2006).[3]

B.2 Wealth

Wealth information is also collected at the time of survey from 1999 onward in nine categories: business assets, transaction accounts (including savings), home equity, equity in vehicles, stocks, other real estate equity, retirement accounts, other assets, and other debt. Compared to the more detailed information collected by the Survey of Consumer Finances, reported wealth tends to be slightly lower in the PSID (Pfeffer et al., 2016). Differences between the SCF and PSID are concentrated in the right tail of wealth, and among households with business assets. Given that the households of interest for this paper tend to have lower wealth and are not self-employed, PSID measures should serve reasonably well.

[1] Implied rent for homeowners is imputed as 6% of the home's reported value.

[2] The new measures introduced in 2005 include home repairs, household furnishing, clothing, vacations, and recreation.

[3] The CES underestimates the NIPA PCE by 25% in more recent preliminary tables published by the BLS.

B.3 Disability

The PSID contains self-reported measures of work-limiting disability for reference persons and their partners. Heads and their partners are first asked (Q1): *Do you have any physical or nervous condition that limits the type of work or the amount of work you can do?* If responding with "Yes", the individual is asked (Q2): *Does this condition keep you from doing some types of work*, to which respondents may reply "Yes", "No," or "Can do nothing." To those who reply to the previous question with "Yes" or "No", the final question is posed (Q3): *For work you can do, how much does it limit the amount of work you can do?* Individuals may respond with " A lot," "Somewhat," "Just a little," or "Not at all."

I code these responses into three exhaustive and mutually exclusive disability categories following Low and Pistaferri (2015), Meyer and Mok (2019), and Deshpande and Lockwood (2020). I classify those who report "No" to Q1 or "Not at all" to Q3 as having no work limitation ($H_{it} = 0$). I classify those who respond "Yes" to Q1 and either "Somewhat" or "Just a little" to Q3 as moderately limited ($H_{it} = 1$). I classify those who report "Yes" to Q1 and either "Can do nothing" in Q2 or "A lot" to Q3 as severely limited ($H_{it} = 2$).

These self-reported measures are commonly used in the literature. On the one hand, many work-limiting disabilities stem from conditions that may be more reliably captured with subjective measures like these, as opposed to objective measures. For example, almost 63% of all disabled workers on DI rolls in 2018 were awarded benefits for a mental or musculoskeletal disorder (e.g., back pain) (Social Security Administration, 2019). On the other hand, subjective measures of health suffer from two potential weaknesses.

First, subjective evaluations of health may not capture variation in true health. To the extent that true health is measured by clinical diagnostics, there are several reassuring results on this concern. A review by Bound and Burkhauser (1999) finds that self-reported measure of health are highly correlated with clinical measures of disability. Low and Pistaferri (2015) and Meyer and Mok (2019) also provide evidence that the PSID measures are correlated with an array of objective health measures among working-age men.

Second, subjective health may be endogenously determined to rationalize work decisions and DI receipt ex-post. Benítez-Silva, Buchinsky, Chan, Cheidvasser, and Rust (2004) alleviate that concern, showing that subjective measures of health are unbiased predictors of disability determinations made by the SSA among those who apply for DI. Kreider (1999) suggests on the other hand that disability is over-reported among the unemployed, compared to the employed who suffer from similar clinical conditions.

As the results of Section 4 and Appendix C show, individuals often continue working full-time even after reporting a severe work-limiting disability. That is not consistent with self-reports being primarily determined to ex-post rationalize non-employment and DI receipt. If the ex-post rationalization of work choices and DI receipt drives some self-reported disability, its consequences for my model are unclear. To the extent that self-reported measures actually captures non-health shocks in addition to or instead of true health shocks, this should simply alter the interpretation of my event studies and model results.

C Additional Event Study Results

C.1 Effects of disability on consumption categories

Figure C.1: *Decomposing consumption effects of disability onset into component categories*

(a) *Single* ***(b)*** *Married*

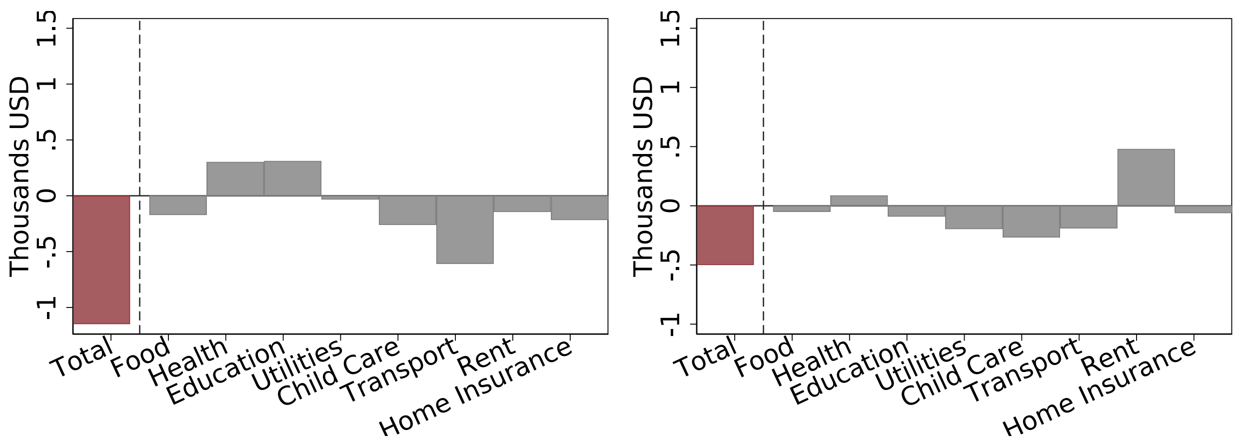

Notes: Dashed lines indicate 90% confidence intervals, with standard errors clustered at the household level.

Figure C.1 explores the spending categories which drive declines in consumption, in dollars than log terms. Unmarried households see category-specific drops in spending on food, utilities, child care, transportation, rent, and home insurance expenditures, while spending on health and education increases. Married households, on the other hand, see declines in all categories of spending except for health expenditures and rent. Thy see a large increase in rent spending.

C.2 Comparing effects of disability on income, consumption, and wealth

Figure C.2: *Comparing income and consumption effects of disability onset to wealth effects*

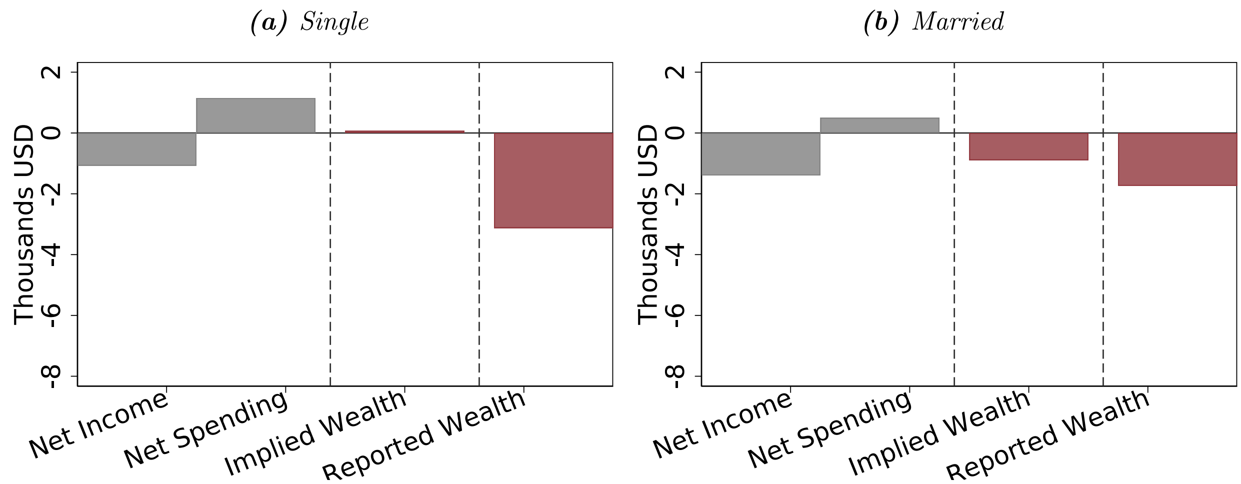

Notes: Dashed lines indicate 90% confidence intervals, with standard errors clustered at the household level. Reported wealth in this figure represents the average annual change in wealth, implied by the year-by-year event study results for level wealth.

A simple accounting equation links the effects of disability on total consumption, wealth, and net income.[4] Figure C.2 depicts the components of this accounting equation. It suggests that the decline in wealth implied by changes in consumption and income underestimates the decline that households actually report. This could reflect costs faced by households of disabled individuals that are not captured by the measured consumption categories (e.g., in-home care or non-prescription treatments). Such costs are consistent with the structural model I build in Section 5.

C.3 Effects on household composition and survey nonresponse

We can make sense of the drop in partner employment among unmarried workers observed in Panel A of Table 4.1 by examining the effects of disability onset on future marital status. This is done in Table C.1, mirroring the format of Table 4.1. Unmarried workers face a 15.2

[4]Beyond concerns that these outcomes may not be accurately measured, another reason the accounting equation may not hold exactly is that total spending and wealth measures are only available in 1997 onward, so the event studies use a different (smaller) sample than those for income.

percentage point drop ($p < 0.01$) in the probability of being married after disability onset relative to their peers, which is consistent with the drop in partner employment discussed previously. Married workers face no detectable change in separation risk. Instead, they tend to have 0.24 fewer children as a consequence of disability onset ($p < 0.01$).

Differential attrition from the sample is one concern with using the PSID data for this event study design. While I cannot detect if this happens contemporaneously with disability onset, I check if reference persons differentially exit the sample after disability onset in Table C.1. Attrition for married households following disability onset is comparable to that of their non-disabled peers. Attrition for unmarried households is slightly higher, having a 9 percentage point higher chance of dropping from the sample compared to their non-disabled peers ($p < 0.05$). The differential attrition of unmarried households is not driven by death or moving into a household with a different reference person. Rather, it is driven by nonresponse to the PSID.

C.4 Effect of a moderate work-limiting disability

Table C.2 summarizes the key event study results for a moderate work-limiting disability, mirroring the structure of Table 4.1. It shows that responses of households to a moderate disability are less sharp in terms of DI take-up, work, and income when compared to the responses to a severe disability. However, unmarried households face a substantial decline in consumption due to moderate disability onset of -25.7 log points, which is comparable to the decline they face due to severe disability onset. Married households see no such decline (they instead spend down their savings), in the same way they saw no decline in consumption due to a severe disability onset. Table C.3 reports effects of a moderate work-limiting disability on household composition and survey non-response. A moderate work-limiting disability has no detectable effects on marital status, the number of children or adults in the household, or attrition.

C.5 Fixing the timing of disability onset by marital status

As shown by Figure C.3, unmarried households tend to face disability onset at systematically younger ages than married households (on average at age 38.6 and 44.7 respectively). This raises the question: how much of the differences in effects of disability by marital status can be explained by birth cohort of the reference person and the age of disability onset? Table C.4 reports the main event study results for the onset of a severe disability, reweighing the data so that unmarried households match married households on the distribution of these observable characteristics. Holding fixed the distribution of birth cohort and age of disability onset does not explain away the differential responses to disability onset across marital status.

C.6 Testing parallel trends

As a test of the parallel trends assumption, I report the year-to-year average differential trend implied by the pre-period event study coefficients in Tables C.6. Column 1 reports this value for unmarried households, and Column 2 reports it for married households. Compared to a joint null test of all the pre-period event study coefficients, this test should have higher power for detecting monotonic differential trends between work-limited households and their peers. I overall find little evidence of violations of parallel trends. The pre-trend test is weakly rejected ($p < 0.1$) for the part-time work of unmarried reference persons. It is not rejected, at any conventional significance level, for any key measures of income, wealth, or consumption of unmarried households.

The pre-trend test is rejected at conventional levels ($p < 0.01$) for the post-tax income of married households, suggesting that married households are already losing household income compared to their peers prior to disability onset. However, this can be naturally explained by mismeasurement in the timing of disability onset. A work-limiting disability first reported in year t could have in truth began at any point after their $t-1$ interview, including sometime during year $t-1$. Consistent with this explanation, Figure 4.2 shows

Figure C.3: *Distribution of reference person severe disability onset, by marital status at time of onset*

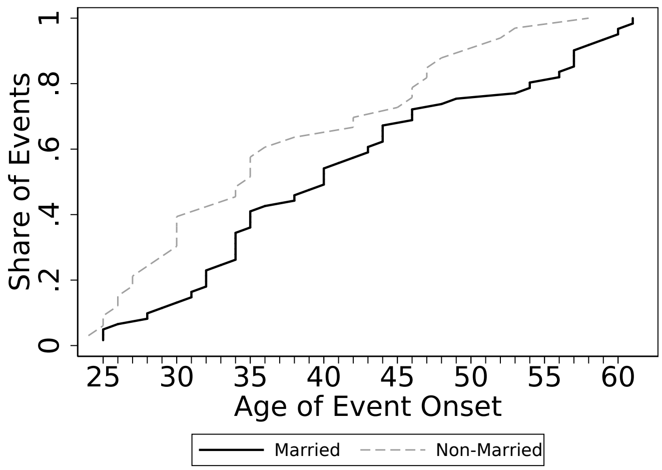

The total subsample of households facing a severe disability onset includes 55 unmarried households and 66 married households.

that the decline in earned income for married workers is driven by a precise dip in the final year prior to disability onset. If we were to take the pre-trend seriously and account for it though, it would suggest that married households lose even less income due to disability than unmarried households.

Table C.1: *Effect of a severe disability onset on household composition and survey attrition*

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Marital Status	0.0372***	0.978***	-0.0590	-0.0287	0.0303
	(0.0121)	(0.0116)	(0.0513)	(0.0265)	(0.0576)
Number of Kids	0.214***	1.286***	0.159	-0.0957	-0.255*
	(0.0734)	(0.154)	(0.102)	(0.0943)	(0.139)
Number of Adults	1.313***	2.166***	-0.238*	0.109*	0.347**
	(0.0906)	(0.0511)	(0.124)	(0.0639)	(0.139)
Leaves household	0	0	0.00730	-0.000177	-0.00748
	(.)	(.)	(0.00722)	(0.00839)	(0.0111)
Dies	0	0	0.0143	0.00575	-0.00850
	(.)	(.)	(0.0102)	(0.00628)	(0.0120)
Drops from Sample	0	0	0.0924***	-0.0137	-0.106***
	(.)	(.)	(0.0313)	(0.0156)	(0.0349)
Sample Data Ends	0	0	-0.0218	0.0136	0.0354
	(.)	(.)	(0.0264)	(0.0122)	(0.0291)
Data is Missing	0	0	0.0922**	0.00548	-0.0867*
	(.)	(.)	(0.0414)	(0.0214)	(0.0466)

Dollar values (other than wages) are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. .

Table C.2: Effects of a moderate disability for the reference person

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Panel A: Employment					
Ref. Full-Time Work	0.745***	0.862***	-0.124*	-0.132***	-0.00851
	(0.0611)	(0.0229)	(0.0669)	(0.0335)	(0.0747)
Ref. Part-Time Work	0.161***	0.108***	-0.0151	0.0201	0.0352
	(0.0519)	(0.0208)	(0.0498)	(0.0236)	(0.0550)
Partner Full-Time Work	0.0342**	0.572***	-0.0279	-0.0420	-0.0141
	(0.0139)	(0.0337)	(0.0460)	(0.0326)	(0.0564)
Partner Part-Time Work	0.00311	0.192***	0.0285	0.0489	0.0205
	(0.00311)	(0.0242)	(0.0246)	(0.0314)	(0.0398)
Panel B: Disability Insurance Receipt					
DI Receipt	0.0528	0.00887	0.0689**	0.0501***	-0.0189
	(0.0365)	(0.00627)	(0.0282)	(0.0180)	(0.0335)
Panel C: Household Income					
Ref. Earned Inc.	15.71***	15.56***	-2.696*	-1.820***	0.877
	(1.558)	(0.740)	(1.387)	(0.628)	(1.522)
Social Security Income	0.166	0.0845***	1.028***	0.491***	-0.536
	(0.127)	(0.0288)	(0.310)	(0.125)	(0.334)
Other Transfers	1.095***	0.881***	-0.0856	0.666***	0.752
	(0.265)	(0.142)	(0.433)	(0.225)	(0.488)
Work Inc. (Others)	1.698**	8.933***	-1.651	-0.319	1.332
	(0.710)	(0.616)	(1.198)	(0.452)	(1.280)
Gross Income	18.63***	25.48***	-3.386**	-1.019	2.367
	(1.665)	(1.066)	(1.665)	(0.695)	(1.803)
Post-Tax Income	14.89***	18.72***	-2.421*	-0.192	2.229
	(1.280)	(0.730)	(1.405)	(0.507)	(1.493)
Panel D: Consumption and Savings					
Consumption	11.63***	11.26***	-4.774	0.206	4.981
	(2.422)	(0.539)	(3.687)	(0.532)	(3.725)
Wealth	23.94	28.83***	-19.51**	-8.002	11.50
	(26.42)	(6.775)	(9.109)	(6.885)	(11.45)

Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Gross income includes informal transfers such as assistance from family, which are not included in the first four income measures.

Table C.3: *Effect of a moderate disability onset on household composition and survey attrition*

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Marital Status	0.0497***	0.990***	0.0394	-0.0225	-0.0619
	(0.0170)	(0.00310)	(0.0655)	(0.0176)	(0.0679)
Number of Kids	0.357***	1.318***	0.119	-0.0365	-0.156
	(0.107)	(0.0893)	(0.128)	(0.0773)	(0.149)
Number of Adults	1.165***	2.346***	0.110	-0.0812*	-0.191*
	(0.0536)	(0.0542)	(0.0982)	(0.0480)	(0.109)
Leaves household	0	0	0.0100	0.0105**	0.000539
	(.)	(.)	(0.00879)	(0.00533)	(0.0103)
Dies	0	0	-0.000686	0.00292	0.00361
	(.)	(.)	(0.000688)	(0.00263)	(0.00268)
Drops from Sample	0	0	0.0439	-0.00492	-0.0488*
	(.)	(.)	(0.0279)	(0.00916)	(0.0293)
Sample Data Ends	0	0	-0.0318	0.000377	0.0322
	(.)	(.)	(0.0213)	(0.00721)	(0.0225)
Data is Missing	0	0	0.0214	0.00892	-0.0125
	(.)	(.)	(0.0347)	(0.0127)	(0.0369)

Dollar values (other than wages) are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. .

Table C.4: *Effects of a severe disability onset for the reference person, fixing the distribution of onset age across marital status*

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Panel A: Employment					
Ref. Full-Time Work	0.833***	0.933***	-0.366**	-0.719***	-0.353*
	(0.110)	(0.0327)	(0.179)	(0.0974)	(0.204)
Ref. Part-Time Work	0.117	0.0667**	0.161	0.267***	0.106
	(0.0777)	(0.0327)	(0.107)	(0.0687)	(0.127)
Partner Full-Time Work	0.0333	0.733***	-0.191**	-0.306***	-0.114
	(0.0351)	(0.110)	(0.0808)	(0.115)	(0.140)
Partner Part-Time Work	0	0.100	0.00221	0.169	0.167
	(.)	(0.0841)	(0.0355)	(0.110)	(0.115)
Panel B: Disability Insurance Receipt					
DI Receipt	0	0	0.179*	0.214**	0.0357
	(.)	(.)	(0.0995)	(0.103)	(0.143)
Panel C: Household Income					
Ref. Earned Inc.	20.28***	14.61***	-3.620	-8.084***	-4.464
	(5.569)	(3.024)	(4.140)	(1.649)	(4.457)
Social Security Income	0.390	0.192	1.044	1.417**	0.373
	(0.411)	(0.192)	(0.728)	(0.692)	(1.004)
Other Transfers	0.823**	0.230	0.989	1.306**	0.317
	(0.355)	(0.162)	(0.725)	(0.530)	(0.898)
Work Inc. (Others)	0.315	6.100***	2.811	-2.343***	-5.154
	(0.332)	(1.936)	(3.472)	(0.884)	(3.582)
Gross Income	21.79***	21.13***	1.063	-7.658***	-8.721**
	(5.310)	(3.150)	(3.798)	(1.947)	(4.270)
Post-Tax Income	16.38***	16.21***	2.842	-4.945***	-7.787**
	(3.401)	(2.410)	(3.148)	(1.170)	(3.360)
Panel D: Consumption and Savings					
Spending	11.34**	9.822***	3.397	-0.791	-4.188
	(5.444)	(3.120)	(2.752)	(3.924)	(4.793)
Wealth (no Business)	-6.889	12.72*	-0.398	4.295	4.694
	(7.812)	(6.664)	(4.007)	(5.120)	(6.502)

This table reports event study coefficients after reweighing unmarried households to match the distribution of observable characteristics for married households. Note that coefficients for married households also differ from Table 4.1; these event studies restrict to the common support of age of onset and birth cohort across marital status. Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table C.5: Effects of a severe disability on higher-educated reference persons

	Pre-Disability Means		Effects of Disability Onset		
	(1)	(2)	(3)	(4)	(5)
	Single	Married	Single	Married	(4) - (3)
Panel A: Employment					
Ref. Full-Time Work	0.969***	0.975***	-0.447***	-0.387***	0.0595
	(0.0147)	(0.00643)	(0.0643)	(0.0427)	(0.0774)
Ref. Part-Time Work	0.0242**	0.0226***	0.0869*	0.139***	0.0516
	(0.0116)	(0.00573)	(0.0456)	(0.0259)	(0.0526)
Partner Full-Time Work	0.0277**	0.543***	-0.0187	0.0143	0.0330
	(0.0112)	(0.0498)	(0.0540)	(0.0539)	(0.0764)
Partner Part-Time Work	0.0346*	0.229***	-0.00330	-0.0516	-0.0483
	(0.0184)	(0.0423)	(0.0202)	(0.0515)	(0.0553)
Panel B: Disability Insurance Receipt					
DI Receipt	0	0	0.173***	0.164***	-0.00868
	(.)	(.)	(0.0429)	(0.0297)	(0.0522)
Panel C: Household Income					
Ref. Earned Inc.	29.58***	23.91***	-7.739**	-2.786	4.954
	(3.199)	(2.057)	(3.408)	(3.582)	(4.945)
Social Security Income	0.284	0.211*	1.280**	1.274***	-0.00657
	(0.234)	(0.111)	(0.518)	(0.289)	(0.593)
Other Transfers	1.394**	0.670***	2.161**	1.748***	-0.413
	(0.543)	(0.141)	(0.915)	(0.542)	(1.064)
Work Inc. (Others)	3.131***	11.22***	-2.669*	0.0196	2.689
	(1.160)	(1.255)	(1.613)	(1.115)	(2.004)
Gross Income	34.34***	36.01***	-6.945*	0.266	7.211
	(4.000)	(2.509)	(4.095)	(3.686)	(5.536)
Post-Tax Income	27.29***	25.97***	-8.388*	1.506	9.894*
	(3.824)	(2.375)	(4.388)	(3.857)	(5.858)
Panel D: Consumption and Savings					
Spending	12.96***	10.30***	-3.967**	1.013	4.980**
	(1.509)	(0.661)	(1.844)	(0.630)	(1.949)
Wealth (no Business)	44.07**	82.21***	-24.28	-21.56	2.724
	(18.95)	(21.33)	(44.03)	(20.05)	(48.40)

Sample persons in this event study exercise have more education than a high school degree. Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Gross income includes informal transfers such as assistance from family, which are not included in the first four income measures.

Table C.6: *Testing Pre-trends for Event Studies of a Severe Disability Onset*

	Pre-Trend Test	
	(1)	(2)
	Single	Married
Panel A: Disability Insurance Receipt		
SSI and SSDI Receipt	2.53e-34	0.0137
	(4.87e-19)	(0.00917)
Panel B: Employment		
Ref. Full-Time Work	-0.00504	0.0357
	(0.0563)	(0.0575)
Ref. Part-Time Work	0.0345	0.0369
	(0.0605)	(0.0539)
Partner Full-Time Work	-0.0328	0.0215
	(0.0306)	(0.0531)
Partner Part-Time Work	-0.00339	-0.00783
	(0.0234)	(0.0584)
Panel C: Household Income		
Ref. Earned Inc.	1.404	1.451
	(1.362)	(1.118)
Social Security Income	-0.132	-0.352
	(0.0895)	(0.242)
Other Transfers	1.303**	-0.307
	(0.528)	(0.452)
Work Inc. (Others)	-0.0339	1.038
	(0.550)	(0.706)
Gross Income	2.549*	1.840
	(1.352)	(1.234)
Post-Tax Income	2.275**	1.392
	(1.031)	(0.858)
Panel D: Consumption and Savings		
Household Consumption	-1.631	0.294
	(2.239)	(0.918)
Household Wealth	-18.30	9.610
	(18.79)	(8.834)

Dollar values are reported in thousands, and are adjusted for the size of the household using the OECD-modified equivalence scale. Wealth excludes any assets associated with an owned business. Confidence stars indicate: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The pre-trend test reports the average year-to-year differential trend of treated households relative to control households, implied by the pre-period event study coefficients.

D Modeling the Tax and Transfer System

The tax and transfer system in my life cycle model of the household incorporates federal income taxes, child tax credits, FICA taxes, the EITC, unemployment insurance, Food Stamps, SSDI, and SSI. To the extent it is feasible, each of these features follows its real program rules as of 1996.

D.1 Disability Insurance Program

Benefit payments through DI follow the true program rules for SSDI and SSI described in Section 2. Eligibility requires working full-time for at least one year, and individuals must choose to apply. Applicants are admitted to DI with a probability depending on current health and age, consistent with the Vocational Grid used by examiners (see Chen and van der Klaauw 2008). While applying for or receiving benefits, individuals are subject to work constraints. Beneficiaries may leave DI in two ways. The first is voluntary, as individuals choose to return to work. The second is by being involuntarily removed from DI by reassessment. Reassessments occur with a probability based on one's health status. The penalty for being removed from DI is that one cannot work and receives no DI benefits in the period of removal.

The model of DI abstracts from reality in a handful of ways. First, DI applicants and beneficiaries in my model are restricted from performing any work. In reality, workers may earn up to Substantial Gainful Activity (SGA), which was $700 per month in 2019. In practice, few DI beneficiaries or applicants engage in work. I rationalize this constraint as the result of strategic behavior (e.g., working being viewed by examiners as a positive signal of health).

Second, I follow Low and Pistaferri (2015) in making some concessions which greatly simplify the state space of the model. This includes simplifying DI eligibility; a reference person is eligible to apply for DI if he has worked for at least a year since his previous application. It also includes simplifying DI payments, which are made based on the worker's permanent wage rather than their true earnings history. Ex-post, applicant reference persons in the model tend to satisfy the true, more extensive work history requirements for DI regardless. As in Low and Pistaferri (2015), SSI enters as an interaction between general welfare programs and DI receipt.

Third, I allow only the (male) reference persons in this model to apply for DI, not their (typically female) partners. This is mainly to avoid concerns arising from how I approximate eligibility with recent work history (since labor force attachment among these partners is weaker). Though overall women take up DI almost as much as men, these partners in the PSID are less than half as likely to take up DI as reference persons, and joint take-up of DI among both reference persons and their partners is exceedingly rare. While females represent a rising share of overall DI rolls, the take-up decision of women tends to be more challenging to credibly model and less salient within partnered households of the PSID.

Fourth, I approximate the uncertain and potentially lengthy application process. Eligible workers must choose to apply, which takes a full year (during which the applicant may not work). At the end of the application year, an allowance decision is realized. Beneficiaries begin receiving payments in the following year, whereas disallowed applicants may return to work (before reapplying in the future, if they choose to do so). This is intended to approximate the five month waiting period, alongside any necessary appeals to reach a final decision.[5] In this model, choosing to apply for DI represents applying and committing to the potentially lengthy appeals process.

Once allowed onto DI, beneficiaries are subject to random continuing disability reviews conditional on their current health status. If selected, a beneficiary is re-examined and may be removed from DI.

Lastly, I abstract away from the health insurance aspect of work and DI receipt. In

[5] About 30% of applicants are allowed onto DI within a year of applying. Another 30% are allowed within 1 to 3 years of applying, with a right tail of applicants who remain in the process even longer (French and Song, 2014).

reality, health insurance is closely tied to both employment and marriage. DI receipt is an alternative mechanism by which individuals may acquire public health insurance through Medicare or Medicaid. Health insurance coverage may also be an important consideration for work-limited individuals, especially prior to the passage of the Affordable Care Act and the protections it provides nationally for pre-existing conditions. Limited data on health insurance coverage in the PSID precludes an adequate accounting of access to health insurance coverage for workers. Examining the role of health insurance in the welfare value of DI is consequently beyond the scope of this study, as was the case for Low and Pistaferri (2015).

D.2 Joint Tax and Transfer System

Household income is subject to key features of the tax and transfer system: the EITC, federal income taxes, food stamps, and unemployment insurance. Parameters of the tax and transfer system are set to their 1996 values.

Transfers through the EITC are determined by household earned income, the maximum credit amount, phase-in rate, and phase-out rate which are all a function of the number of children in the household (consistent with real program rules). The earned income level at which EITC phase-out begins is a function of both marital status and the number of children in the household.[6]

Federal taxes imposed on pre-tax income also follow the true progressive tax schedules from 1996.[7] Households take non-refundable standard deductions ($6,700 for married or $4,000 for a single filers), personal exemptions ($2,550 per person), and child tax credits ($400 per child).[8] Social security payments are subject to federal income taxes, with deductions

[6]Maximum credit amounts were $347, $2,312, or$3,816 for 0, 1, or more than 1 child. Phase-in rates are 7.65%, 34%, or 40% for the same. Phase-out rates were 7.65%, 15.98%, or 21.06%. Phase-out begins at $5,700 for households with no children or $12,500 for households with children.

[7]For married households, the 1996 tax schedule consisted of marginal tax rates of 15%, 28%, 31%, 36%, and 39.6% beginning at $0, $40,100, $96,900, $147,700, and $263,750 respectively.

[8]I assume that unmarried reference persons file with single status rather than head of household status, and do not claim children for the purposes of the EITC. The reason for this is that most unmarried reference persons empirically have no children and would not qualify. Among the 28% of household-years in which unmarried reference persons are observed with children, 80% of households contain at least one other adult which renders unclear who may claim the child for head of household status.

depending on the household's earned income.

The above description of the tax system assumes that married couples file jointly. The model restricts to less educated and relatively lower-earning households which would tend to benefit from joint filing (filing separately would disqualify households for the EITC). Empirically, the large majority of married households file jointly across the income distribution. Through the jointness of the tax system and the EITC, a worker's net tax obligation is weakly decreasing in the labor force participation of their partner.

Food stamps enter into the model as the main general social safety net program. Eligibility and cash benefits are modeled according to real program rules. Eligibility requires that gross household income falls below 130% of the Federal Poverty Level (FPL), and income net of deductions falls below 100% FPL. Deductions include a standard \$1,608 deduction, a 20% deduction for earned income, and a \$2,100 deduction per dependent child. Food stamp transfers to the household are increasing in the size of the household and reduced by 30% of household income net of deductions, including DI or retirement benefits or employment earnings of a spouse.

Food stamps depart in the model depart from the true program in three ways. First, take-up in the model is automatic among eligible households and households automatically satisfy any work or work training requirements. Second, means tests in the model are performed only on income and not on wealth.[9] Third, food stamps in the model provides a cash transfer, rather than in-kind transfers.[10]

Unemployment insurance (UI) is the final component of the general tax and transfer system. It provides a cash transfer in the event of exogenous job destruction.[11] Unemployment insurance rules in practice vary from state to state. Its role in this model is to dampen the

[9]In reality, food stamps includes restrictions on liquid wealth, but not on durables, housing, or pension wealth. The model only includes a single liquid asset.

[10]As with liquid vs. illiquid assets, this model does not differentiate between food and other consumption goods. Furthermore, see Hoynes and Schanzenbach (2009) for evidence that households treat food stamp transfers as equivalent to cash.

[11]voluntarily leaving work does not qualify one for UI, I assume this is observable in the model.

consequences of job destruction. For this reason, I model unemployment insurance as a single cash payment equal to 25% of the previous year's earnings (up to a limit). This roughly approximates the average 47% replacement rate reported by the Department of Labor's Benefit Accuracy Measurement survey, with benefits typically collectible for a maximum of two quarters. With job destruction causing a fixed year-long involuntary unemployment spell, unemployment insurance has no moral hazard consequences in this model.

E Simulating Potential Welfare Benefits

As an exogenous shifter of the propensity to work, I construct a measure consistent with the measure used by Low and Pistaferri (2015). This is a measure of potential welfare benefits in the event that a reference person works an average of 20 hours per week at the federal minimum wage, with no other earned income in the household. This measure is constructed using safety net programs which experience reforms over time, and some of which vary across states as well. The programs included in this measure are: the EITC (including state supplements), SNAP/food stamp payments, AFDC/TANF benefits administered by states, and unemployment insurance benefits administered by states.[12]

$$Z_{it} = Z_{it}^{EITC} + Z_{it}^{AFDC} Z_{it}^{TANF} + Z_{it}^{FS/SNAP} + Z_{it}^{UI}$$

The EITC payments are a kinked function of gross earned income w_t, which depends on t through the federal minimum wage:

$$Z_{it}^{EITC} = \tau_{it}^1 w_t \mathbb{1}\{0 < w_t \leq k_t^1\} + \tau_{it}^1 k_t^1 \mathbb{1}\{k_t^1 < w_t \leq k_t^2\} + (\tau_{it}^1 k_t^1 - \tau_{it}^2 (w_t - k_t^2)) \mathbb{1}\{k_t^2 < w_t < k_t^3\}$$

[12]The federal minimum wage enters the computation of these benefits, as each of them depens on earned income.

The parameters k_t^1, k_t^2, and k_t^3 are kinks in the EITC formula determined by federal rules and which vary from year to year. The parameters τ_{it}^1 and τ_{it}^2 are the phase-in and phase-out rates for EITC payments respectively. They vary across households i through both the number of dependent children in the household and state of residence, as many states supplement the federal EITC payment. Both federal phase-in and phase-out rates, as well as state supplements, vary over time. I follow Low and Pistaferri (2015) in disallowing EITC payments from households belonging to demographic groups (based on family size, education, state of residence, age, and calendar year) with average assets that exceed EITC limits.[13]

Prior to 1996, AFDC payments are determined as a share $r_{s(i)}^{NS}$ of the state-determined need standard (NS_{it}, varying over states, time, and family size), reduced by countable income (y_{it}^{AFDC}) according to a state-specific rate $r_{s(i)}^{y}$. Payments are also subject to an upper limit M_{it}^{AFDC} which depends on state and family size, and varies over time.

$$Z_{it}^{AFDC} = \min\{\max\{0, r_t^{NS} NS_{it} - r_t^y y_{it}^{AFDC}\}, M_{it}^{AFDC}\}$$

Countable income y_{it}^{AFDC} represents earned household income net of deductions. Those deductions include a \$90 flat deduction, in addition to a \$30 deduction for the first year of benefit receipt and a deduction equal to 1/3 of earned income for the first 4 months on a job. For the purpose of simulated benefits, a worker is always assumed to be on the first year of their job. Consistent with real program rules, benefits are 0 for households with no children, gross income above 185 percent of the need standard, or countable income above 100 percent of the need standard. AFDC benefits are only non-zero prior to welfare reform in 1996.

TANF replaces AFDC after 1996 welfare reform. TANF payments are equal to some

[13]Historical federal EITC parameters are taken from Congressional Research Service reports. Historical state EITC parameters are taken from NBER TAXSIM documentation.

maximum M_{it}^{TANF} which depends on state of residence, family size, and year. This amount is reduced by earned income net of a state-specific deduction formula (y_{it}^{TANF}). Households earning above an income limit (varying by state and family size, and over time) are ineligible for TANF payments.

$$Z_{it}^{TANF} = \max\{0, M_{it}^{TANF} - y_{it}^{TANF}\}$$

Food stamp benefits are equal to a maximum amount varying by year and family size ($M_{it}^{FS/SNAP}$), which is reduced by $0.30 for each $1.00 of family income net of deductions ($y_{it}^{FS/SNAP}$). Family income includes earned income, as well as other transfer payments. Deductions include a standard deduction, and deductions for: earned income (20%), dependent care, medical costs, child support payments, and excess shelter expenses. The excess shelter deduction is equal to any rent payments over 50% of family income net of other deductions and subject to a limit which varies over time. The standard deduction and dependent care deduction vary over time. Following Low and Pistaferri (2015), I set dependent care and medical deductions to zero because information on these expenses is unavailable prior to 1999. Child support and rent payments are set at average values within demographic cells (based on family size, education, state of residence, age, and calendar year).

$$Z_{it}^{FS/SNAP} = \max\{10, M_{it}^{FS/SNAP} - 0.3y_{it}^{FS/SNAP}\}$$

For unemployment insurance, I follow Low and Pistaferri (2015) in assuming full coverage and receipt of payments for 26 weeks. UI pays a minimum amount $m_{s(i)t}^{UI}$, up to a maximum amount $M_{s(i)t}^{UI}$ which varies across state of residence $s(i)$ and time. Payments are based on a replacement rate $r_{s(i)t}^{UI}$ for reference earnings $y_{s(i)t}^{UI}$. Additionally, some states provide an allowance for dependents, denoted b_{it}, which may vary over time. The definition of reference earnings $y_{s(i)t}$ vary across states and over time.[14] The replacement rate $r_{s(i)t}^{UI}$ likewise varies

[14]More than half of states now base payments on a worker's highest quarter of earnings, converted to average weekly earnings. The period over which earnings are taken is most often the first four of the last

over states, and within states, over time.

$$Z_{it}^{UI} = \max\{\min\{r_{s(i)t}^{UI} y_{s(i)t}^{UI}, M_{s(i)t}^{UI}\}, m_{s(i)t}^{UI}\} + b_{it}$$

F Selection Correction in the Wage Equation

The expression for the selection correction is motivated by distributional assumption on unobserved errors in the wage equation, based on the standard Roy model. Work decisions and the log wage process are respectively defined as follows:

$$\begin{aligned}
L_{it}^{s} =& \mathbb{1}\{\alpha_0^{s\prime} \boldsymbol{Z}_{it}^{s} + \alpha_1^{s\prime} \boldsymbol{A}_{it}^{s} + \alpha + 2^{s\prime} \boldsymbol{H}_{it}^{s} + k_{g(i)}^{s} + \nu_{it}^{s} > 0\} \\
\ln W_{it}^{s} =& \beta_1^{s\prime} \boldsymbol{A}_{it}^{s} + \beta_2^{s\prime} \boldsymbol{H}_{it}^{s} + f_{g(i)}^{s} + \omega_{it}^{s} + \epsilon_{it}^{s} \\
\omega_{it}^{s} =& \omega_{it-1}^{s} + \xi_{it}^{s}, \quad \xi_{it}^{s} = \sum_{\ell} \Gamma_{\ell} \nu_{it-\ell}^{s} + \eta_{it}^{s} \\
\nu_{it}^{s} \sim& N(0,1) \\
\begin{pmatrix} \epsilon_{it}^{1} \\ \epsilon_{it}^{2} \end{pmatrix} \sim& N\left(\mathbf{0}, \begin{pmatrix} (\sigma_{\epsilon}^{1})^2 & \theta_{\epsilon} \\ \theta_{\epsilon} & (\sigma_{\epsilon}^{2})^2 \end{pmatrix}\right), \quad \begin{pmatrix} \xi_{it}^{1} \\ \xi_{it}^{2} \end{pmatrix} \sim N\left(\mathbf{0}, \begin{pmatrix} (\sigma_{\xi}^{1})^2 & \theta_{\xi}\sigma_{\xi}^{1}\sigma_{\xi}^{2} \\ \theta_{\xi}\sigma_{\xi}^{1}\sigma_{\xi}^{2} & (\sigma_{\xi}^{2})^2 \end{pmatrix}\right)
\end{aligned}$$

The unobservable terms $(\epsilon_{it}^{1}, \epsilon_{it}^{2})$ are interpreted as measurement error. They are normally distributed and strictly exogenous with mean zero, variance $(\sigma_{\epsilon}^{s})^2$, and inter-partner correlation θ_{ϵ}. The terms ω_{it}^{s} are also unobservable, and follow a random walk.

The innovations in the random walk are denoted ξ_{it}^{s}, and are drawn from a multivariate normal distribution, with variances $(\sigma_{\xi}^{1})^2$ and $(\sigma_{\xi}^{2})^2$, and inter-partner correlation θ_{ξ}. The stated structure allows for flexible correlations in some important ways. The innovations of partner s may be correlated with past, contemporaneous, and future shocks to preferences

five completed quarters prior to the filing of a claim. State parameters are taken from reports of the US Department of Labor.

of work for partner s (ν_{it}^s) according to the correlation parameters Γ_ℓ. This allows for the classic selection problem, whereby shocks to potential earnings may affect one's decision to work. However, the innovations of partner s may not be correlated with the shocks to work preferences of the other partner.

The structure also imposes restrictions on the joint selection problem of spouses. It rules out the possibility that shocks to preferences for work, ν_{it}^s, are correlated across partners. It also implies that innovations to one's wage do not affect selection into work for one's partner, hence why ξ_{it}^s can be expressed as a function of only $\boldsymbol{\nu}_i^s$ and not both $\boldsymbol{\nu}_i^1$ and $\boldsymbol{\nu}_i^2$. The term η_{it}^s represents variation in wages which is uncorrelated with one's own work preferences by construction. The specification of ξ_{it}^s and ν_{it}^s together imply that η_{it}^1 and η_{it}^2 are jointly normally distributed with a correlation that depends on θ_ξ.

F.1 Estimation

The model for work is estimated for both partners using a probit regression. The main coefficients are presented in Appendix Table F.1. I estimate parameters of the wage process by taking first differences, which eliminates the group effects $f_{g(i)}^s$ and accumulated values of the random walk error term. To accommodate the movement from annual to biannual surveying in 1997, I take differences over p periods $p \in \{1, 2\}$:

$$\Delta^p \ln W_{it}^s = \beta_1^{s\prime} \Delta^p \boldsymbol{A}_{it}^s + \beta_2^{s\prime} \Delta^p \boldsymbol{H}_{it}^s + \sum_{j=0}^{p-1} \xi_{it-j}^s + \Delta^p \epsilon_{it}^s$$

An individual chooses to work only when $\nu^s_{it} > -\alpha^{s\prime}_0 \boldsymbol{Z}^s_{it} - \alpha^{s\prime}_1 \boldsymbol{A}^s_{it} - \alpha^{s\prime}_2 \boldsymbol{H}^s_{it} - k^s_{g(i)} \equiv -h^s_{it}$. Selection in the wage equation enters through the ξ^s_{it-j} terms.

$$
\begin{aligned}
E[\Delta^p \ln W^s_{it} | L_{it} = 1, L_{it-p} = 1] =& \beta^{s\prime}_1 \Delta^p \boldsymbol{A}^s_{it} + \beta^{s\prime}_2 \Delta^p \boldsymbol{H}^s_{it} \\
&+ \sum_{j=0}^{p-1} E[\xi^s_{it-j} | \nu^s_{it} > -h^s_{it}, \nu^s_{it-p} > -h^s_{it-p}] \\
=& \beta^{s\prime}_1 \Delta^p \boldsymbol{A}^s_{it} + \beta^{s\prime}_2 \Delta^p \boldsymbol{H}^s_{it} \\
&+ \sum_{j=0}^{p-1} E[\sigma^s_\xi \rho^s_j \nu^s_{it} + \sigma^s_\xi \rho^s_{j-p} \nu^s_{it-p} | \nu^s_{it} > -h^s_{it}, \nu^s_{it-p} > -h^s_{it-p}] \\
=& \beta^{s\prime}_1 \Delta^p \boldsymbol{A}^s_{it} + \beta^{s\prime}_2 \Delta^p \boldsymbol{H}^s_{it} \\
&+ \sum_{j=0}^{p-1} \sigma^s_\xi \rho^s_j \lambda^s_{it} + \sigma^s_\xi \rho^s_{j-p} \lambda^s_{it-p} \qquad (1)
\end{aligned}
$$

where $\lambda^s_{it} = \frac{\phi(h^s_{it})}{\Phi(h^s_{it})}$ is the inverse Mills' ratio and ρ^s_ℓ is the correlation between the shock to preferences for work and the ℓth lag (or lead if $k < 0$) of the permanent wage shock.[15]

The wage growth parameters are estimated consistently by including estimates of λ^s_{it} and λ^s_{it-p} (derived from the model for the first stage decision to work) as covariates on the right hand side of the log wage growth regression:

$$
\Delta^p \ln W^s_{it} = \beta^{s\prime}_1 \Delta^p \boldsymbol{A}^s_{it} + \beta^{s\prime}_2 \Delta^p \boldsymbol{H}^s_{it} + \sum_{j=0}^{p-1} \alpha^j_1 \hat{\lambda}^s_{it} + \alpha^j_2 \hat{\lambda}^s_{it-p} + \text{residual}_{it} \qquad (2)
$$

For the (male) reference persons who have high labor force participation, there is little room for selection to introduce bias into the wage regressions. For the (female) partners though, labor force participation is lower and the potential for selection is greater.

[15] With an instrument used to estimate the first stage decision to work, identification of the work and wage system does not rely on the distributional assumption of normality. The selection expression can be viewed as a non-parametric approximation of a true selection term in the case that the distributional assumptions are violated.

F.2 Estimation of Productive Fixed Effects

Given consistent estimators of the wage growth parameters as described in Section F, consider the expected residual log wage for partner s of household i:

$$\begin{aligned} E[\widetilde{\ln W_{it}^s}|i, L_{it}=1] = \; & E[\ln W_{it}^s - \hat{\beta}_1^{s\prime} \boldsymbol{A}_{it}^s - \hat{\beta}_2{}' \boldsymbol{L}_{it}^s | i, L_{it}=1] = f_{g(i)}^s + E[\omega_{it}^s + \epsilon_{it}^s | i, L_{it}=1] \\ = \; & f_{g(i)}^s + \sum_{\ell=0}^{t} E[\xi_{it-\ell}^s | i, L_{it}=1] \\ = \; & f_{g(i)}^s + \underbrace{\lambda_{it}^s \sum_{\ell=0}^{t} \sigma_\xi^s \rho_\ell^s}_{\text{bias}} \end{aligned}$$

The second equality uses that productive type effects are constant conditional on i. The second line uses the exogeneity of ϵ_{it} and the recursive definition of ω_{it}. The final equality provides an expression for the bias term arising from selection in $\widetilde{\ln W_{it}^s}$.

I account for selection bias in the household average of $\widetilde{\ln W_{it}^s}$ by subtracting $\hat{\alpha}_1^0 \hat{\lambda}_{it}^s$ from it, which is an estimate for $\sigma_\xi^s \rho_0^s \lambda_{it}^s$. This is only part of the full bias term, but it is the largest of the components if shocks to work preferences are more correlated with contemporaneous shocks to wages than shocks to wages in past years.

F.3 Estimation of Wage Variance

As described above, the two partners draw their innovations for the wage equation from a joint normal distribution, allowing for a correlation, θ_ξ, in contemporaneous productivity shocks in the household. This distribution has the following variance-covariance matrix.

$$\Sigma \equiv \begin{bmatrix} (\sigma_\xi^1)^2 & \theta_\xi \sigma_\xi^1 \sigma_\xi^2 \\ \theta_\xi \sigma_\xi^1 \sigma_\xi^2 & (\sigma_\xi^2)^2 \end{bmatrix}$$

Estimation of the three parameters defining Σ, and estimating the terms of the variance-covariance matrix of the measurement error process, is based on the p-period difference of residual log wages ($p \in \{1, 2\}$). That observed variable is defined in Equation 6.5, and its mean is an estimator for $\Delta^p(\omega_{it}^s + \epsilon_{it}^s)$.

Given the assumptions on ϵ_{it}^s, ξ_{it}^s, and ν_{it}^s stated at the beginning of this section, the variance-covariance parameters are estimated alongside several nuisance parameters by generalized method of moments. The system used in GMM includes the second order moment, first order moment and autocovariance $\Delta^p \widetilde{\ln W_{it}^s}$. These respectively take the form:

$$
\begin{aligned}
E[(\Delta^p(\omega_{it}^s + \epsilon_{it}^s))^2 | \mathbf{L}_{i,(t-p,t)}^s = \mathbf{1}] &= \sigma_\xi^{s2}\left(p - h_{it}^s \lambda_{it}^s \sum_{j=0}^{p-1} (\rho_j^s)^2\right) \\
&- \sigma_\xi^{s2}\left(h_{it-p}^s \lambda_{it-p}^s \sum_{j=1}^{p} (\rho_{-j}^s)^2\right) \\
&+ 2\sigma_\xi^{s2}\left(\sum_{j=0}^{p-1} \rho_j^s \rho_{j-p}^s \lambda_{it}^s \lambda_{it-p}^s\right) \\
&+ \sigma_\xi^{s2}\left(\sum_{j=0}^{p-1} \sum_{\substack{\ell=0 \\ \ell \neq j}}^{p-1} (\rho_j^s \lambda_{it}^s + \rho_{j-p}^s \lambda_{it-p}^s)(\rho_\ell^s \lambda_{it}^s + \rho_{\ell-p}^s \lambda_{it-p}^s)\right) \\
&+ 2\sigma_\epsilon^{s2} \\
E[\Delta^p(\omega_{it}^s + \epsilon_{it}^s) | \mathbf{L}_{i,(t-p,t)}^s = \mathbf{1}] &= \sum_{j=0}^{p-1} \sigma_\xi^s \rho_j^s \lambda_{it}^s + \sigma_\xi^s \rho_{j-p}^s \lambda_{it-p}^s \\
E[\Delta^p(\omega_{it}^s + \epsilon_{it}^s)\Delta^r(\omega_{it-p}^s + \epsilon_{it-p}^s)\mathbf{L}_{i,(t-p,t)}^s = \mathbf{1}] &= -\sigma_\epsilon^{s2}
\end{aligned}
$$

The vector $\mathbf{L}_{i,(t-p,t)}^s$ contains indicators for work of partner s of household i in periods $t-p$ and t.[16] These three equations along with instruments derived from variation in simulated welfare benefits are sufficient to pin down the variance of the wage innovations and the variance of the noise terms, as in Low and Pistaferri (2015).

In order to pin down the intra-household correlation in wage innovations and in measurement errors (the parameters θ_ξ and θ_ϵ), the above system of equations is augmented with a

[16]Note that for some random variable X, $E[X|\mathbf{L}_{(t-p,t)}^s = \mathbf{1}] \equiv E[X|\nu_{it}^s > -h_{it}^s, \nu_{it-p}^s > -h_{it-p}^s]$

final set of equations: the intra-household correlation in $\Delta^p \widetilde{\ln W_{it}^s}$:

$$E[\Delta^p(\omega_{it}^1 + \epsilon_{it}^1)\Delta^p(\omega_{it}^2 + \epsilon_{it}^2)|\mathbf{L}_{(t-p,t)}^{(1,2)} = \mathbf{1}] = E\left[\left(\sum_{j=0}^{p-1} \xi_{it-j}^1\right)\left(\sum_{j=0}^{p-1} \xi_{it-j}^2\right)|\mathbf{L}_{(t-p,t)}^{(1,2)}\right]$$
$$+ E[\epsilon_{it}^1\epsilon_{it}^2 + \epsilon_{it-p}^1\epsilon_{it-p}^2]$$
$$= \left(\sum_{j=0}^{p-1} \rho_j^1\lambda_{it}^1 + \rho_{j-p}^1\lambda_{it-p}^1\right)\left(\sum_{j=0}^{p-1} \rho_j^2\lambda_{it}^2 + \rho_{j-p}^2\lambda_{it-p}^2\right)$$
$$+ p\sigma_\xi^1\sigma_\xi^2\theta_\xi + 2\sigma_\epsilon^1\sigma_\epsilon^2\theta_\epsilon$$

where $\mathbf{L}_{(t,t-p)}^{(1,2)}$ is a vector containing indicators for work of both partners in household i in both periods $t-p$ and t.

Table F.1: *Work probit model parameters*

	Head	Partner
Low Type	$-1.334{*}{*}{*}$	-0.0921
	(0.0666)	(0.0889)
Medium Type	0.0660	0.0318
	(0.0591)	(0.0606)
High Type	0	0
	(.)	(.)
Sev. Dis.	$-1.521{*}{*}{*}$	
	(0.440)	
Mod. Dis.	$-1.021{*}$	$-2.261{*}{*}{*}$
	(0.410)	(0.385)
Age	$0.119{*}{*}{*}$	$0.0890{*}{*}{*}$
	(0.0149)	(0.0230)
$Age^2/100$	$-0.185{*}{*}{*}$	$-0.129{*}{*}{*}$
	(0.0174)	(0.0292)
Married	$0.316{*}{*}{*}$	
	(0.0561)	
N	24491	14064

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors (in parentheses) are clustered at the household level. Wealth is reported in thousands of 1996 US dollars. Partners include married spouses and non-marital partners.

G Parameter values estimated by simulated method of moments

Table G.1 reports the fitted values for parameters selected by simulated method of moments. Where I can compare these parameter values directly with values from Low and Pistaferri (2015) (hereafter in this section called LP) and Autor et al. (2019), they look similar.

Panel A of Table G.1 reports the estimated disutilities associated with work and health. They imply that disability and work both negatively affect household utility. Disutility of work when healthy is very low for married reference persons ($\eta_1^{mar} = -0.031$) and high for unmarried reference persons ($\eta_1^{sing} = -0.783$), suggesting marriage reduces the utility cost of work. A large disutility of spousal labor supply ($\eta_2^{mar} = -0.449$) partially offsets this effect if the spouse is working. A moderate disability has only a minor effect on the disutility of work for the reference person ($\eta_1^{mod} = -0.057$). I can therefore compare the direct effect of a moderate disability to the one reported by LP, which fixes $\eta_1^{mod} = 0$. The direct effect I find ($\theta_1^{mod} = -0.612$) effectively reduces utility by 54 percent in terms of consumption. This is comparable to the 36 percent effect found by LP. A severe work-limiting disability for the reference person has a less harmful direct effect on utility when compared to a moderate work-limiting disability ($\theta_1^{sev} = -0.009$). It instead greatly increases the disutility of work ($\eta_1^{sev} = -0.694$). Combining that with the baseline cost of work, this implies a total cost of work when suffering a severe disability of -0.703 and -1.48 for married and unmarried workers respectively, which are comparable in magnitude to the values of -0.829 and -1.076 reported by Autor et al. (2019).

Panel B of Table G.1 reports estimated values for the exogenous job loss risk and the fixed costs of work. I find a value of 9.3 percent for married reference persons, nearly 0 percent for unmarried reference persons, and 1.3 percent for spouses. The married and unmarried job loss probabilities for reference persons imply a 5.8 percent change of job loss for reference persons, unconditional on marital status. The probabilities of job loss found for workers in LP is 6.2 percent.

Fixed costs of work are reported as a share of annual earnings for the median reference person at age 23. Fixed costs associated with disability and older age both positive. The fixed costs associated with a moderate disability of the reference person (0.448 for married households and 0.312 for unmarried households) are of comparable magnitude to the fixed cost of 0.547 reported by LP. The fixed costs associated with a severe disability are small (0.233 for married and 0.062 for unmarried reference persons), whereas LP reports a large cost of 0.952. This can be explained by the fact that LP do not allow for interaction terms between work and health in household preferences. The fitted value of that interaction term (η_1^{sev}) is very large in my model, suggesting that severe disability onset causes a large change in consumption which my model would not explain if I were to load costs of work onto F_1^3 (the monetary fixed cost of work with a severe disability which appears in the savings equation) rather than η_1^{sev}.

Panel C of Table G.1 reports the fitted allowance probabilities for DI applicants by age and health. Both married and unmarried workers face allowance probabilities which mostly increase with disability severity and age, consistent with the Vocational Grid used by examiners (see Chen and van der Klaauw 2008).

The allowance probabilities for married workers in my model tend to be lower than the annualized allowance probabilities implied by LP. This is reasonable, as spousal labor supply reduces the cost of applying for DI in my model relative to a single-earner model such as LP, meaning that lower allowance probabilities are necessary in order to rationalize empirical DI take-up rates. A hypothetical young applicant in LP can apply twice in a single year and would be admitted with a probability of 0.31 with a moderate disability and 0.55 with a severe disability. This compares to respectively probabilities of 0.07 and 0.36 for married workers in my model. Those values, if the applicant is over the age of 45, are 0.33 and 0.86

from LP compared to 0.02 and 0.79 from my model.

Table G.1: *Parameter values estimated by simulated method of moments*

Panel A: Preferences		
	Head	Spouse
μ	-0.021	
θ^{mod}	-0.612	-0.016
θ^{sev}	-0.009	
η^{mar}	-0.031	-0.449
η^{sing}	-0.783	
η^{mod}	-0.057	-0.306
η^{sev}	-0.694	

Panel B: Frictions and Fixed Costs			
	Head (Married)	Head (Single)	Spouse
δ	0.093	0.000	0.013
F^2	0.448	0.312	0.012
F^3	0.233	0.062	
F^{old}	0.368	0.092	0.000

Panel C: Allowance Probabilities		
	Single	Married
π_1^{young}	0.002	0.050
π_2^{young}	0.025	0.070
π_3^{young}	0.359	0.359
π_1^{old}	0.001	0.000
π_2^{old}	0.286	0.002
π_3^{old}	0.686	0.794

Notes: the preference parameters θ_1 and θ_2 denote disutilities of health associated with moderate and severe disability. The parameters η^{mar} and η^{sing} are disutilities of work, depending on marital status. The parameters η^{mod} and η^{sev} are additional disutilities of work which depend on disability status and are common across married and unmarried workers. δ denotes the probability of exogenous job destruction. The fixed costs of work for moderate and severe disability (F_2 and F_3) and for old age (F^{old}) are reported as a share of median income at age 23. The parameters π_h^a denote the probability that an applicant of age group a and health status h is allowed onto DI. The *old* age group consists of workers age 45 and older.

Table H.1: *Redistribution and insurance motives for willingness to pay for DI reforms*

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Redistribution				Insurance	
	Give-Away	Between	Within	Health	Singlehood	Productivity	Non-Pecuniary
Temp. DI	0.37	0.16	-0.03	0.09	-0.05	-109.65	110.10
Allow App. Work	30.51	-16.18	4.36	8.32	-7.34	-6.16	-12.51
DI Leniency	0.14	0.13	-0.13	0.01	-0.25	-0.11	1.22
DI Generosity	0.65	0.39	-0.29	0.03	-0.29	-0.12	0.62
SNAP Generosity	0.06	0.07	-0.04	0.00	-0.09	-1.79	2.80

H Methods and results from decomposing insurance and redistribution motives

This section analyzes how WTP for reforms to DI is driven by insurance versus redistribution motives, and how WTP is specifically driven by motives to insure against health risk. To do this, I decompose each reform into five components following Shaw (2014). The first component (called the "give-away" by Shaw) measures how the reform tends to increase the average income of households. The second component (redistribution between households) measures how the reform tends to reallocate income across low-type, mid-type, and high-type households. The third component (redistribution within households) measures how the reform tends to reallocate income across the life cycle, within households of common type. The fourth component (insurance against health risks) measures how the reform shifts income from the healthy to the work-limited, fixing age and type. The fifth component (insurance against singlehood risk) measures how the reform shifts income from married to unmarried workers, fixing health and the previous characteristics. The sixth component (insurance against productivity risk) measures how the reform shifts income from individuals with high potential wages to individuals with low potential wages, fixing marital status and the previous characteristics. The final component (non-pecuniary effects) completes the reform, including how changes in transfers brought about by the reform depend on the household's choices.

For a given reform, these components are defined formally by the following decomposition:

$$\tau^{\text{giveaway}} = \tau + \underbrace{\Delta^{\text{reform}} E[\tau]}_{\text{Give-away}} \tag{3}$$

$$\tau^{\text{Redist. Between}} = \underbrace{\Delta^{\text{reform}} (E[\tau|g] - E[\tau])}_{\text{Redistribution Between Households}} \tag{4}$$

$$\tau^{\text{Redist. Within}} = \tau^{\text{Redist. Between}} + \underbrace{\Delta^{\text{reform}} (E[\tau|g,t] - E[\tau|g])}_{\text{Redistribution Within Households}} \tag{5}$$

$$\tau^{\text{Health Insur.}} = \tau^{\text{Redist. Within}} + \underbrace{\Delta^{\text{reform}} (E[\tau|\boldsymbol{H}_t, g, t] - E[\tau|g,t])}_{\text{Insurance Against Health Risks}} \tag{6}$$

$$\tau^{\text{Singlehood Insur.}} = \tau^{\text{Health Insur.}} + \underbrace{\Delta^{\text{reform}} (E[\tau|S_t, \boldsymbol{H}_t, g, t] - E[\tau|\boldsymbol{H}_t, g, t])}_{\text{Insurance Against Singlehood Risk}} \tag{7}$$

$$\tau^{\text{Productivity Insur.}} = \tau^{\text{Singlehood Insur.}} + \underbrace{\Delta^{\text{reform}} (E[\tau|\boldsymbol{W}_t, S_t, \boldsymbol{H}_t, g, t] - E[\tau|S_t, \boldsymbol{H}_t, g, t])}_{\text{Insurance Against Productivity Risk}} \tag{8}$$

$$\tau^{\text{reform}} = \tau^{\text{Productivity Insur.}} + \underbrace{\Delta^{reform} (\tau - E\left[\tau|\boldsymbol{W}_t, S_t, \boldsymbol{H}_t, g, t\right])}_{\text{Non-pecuniary effects}} \tag{9}$$

where τ is the function mapping employment earnings and DI receipt into net transfer receipts from the government (with arguments suppressed for brevity), S_t is an indicator for being unmarried, $\boldsymbol{W}_t$ contains the potential wages of the reference person and partner, and $\boldsymbol{H}_t$ is the health status of the reference person and partner at age t. Table H.1 shows the share of household WTP for each reform that is explained by the five components. Each row decomposes a different reform, with the entries in each row summing to 1.

The importance of the "give-away" component varies from reform to reform, though it represents less than half of the overall value of most reforms. For those reforms, the monetary value of the transfers alone do not explain the welfare gains they provide. Instead, most of the welfare gains are determined by who receives those transfers and when they are provided. The give-away can explain 37 percent of WTP for a temporary DI program and 65 percent of WTP for increasing the generosity of DI benefits. On the other hand, it explains only 14 percent of WTP for increasing the leniency of the DI application process, and 6 percent of WTP for increasing the generosity of SNAP benefits. The give-away explains far more than the full welfare gains from allowing applicants to work, owing to the fact that the increased income this reform provides to households comes at the cost of requiring DI applicants to work and earn it.

Redistribution between households reflects preferences for reallocating income from higher-type households to lower-type households, which are characteristics that are realized before adulthood begins. It represents 16 percent of WTP for a temporary DI program, 13 percent of WTP for increasing leniency of DI allowance, and 39 percent of WTP for increasing the generosity of DI benefits. It represents a smaller share of WTP for the benchmark reform to SNAP benefits (7 percent). For the reform that allows DI applicants to work, redistribution between households more than cuts the welfare gains from the give-away in half, because DI applicants who receive the most income from this reform come from higher-earning (higher-type) households.

Owing to the fact that disability tends to happen later in life, DI reforms disproportionately reallocate income from younger ages to older ages. Redistribution within households consequently reduces household willingness to pay for most reforms, as households face borrowing constraints, and would generally prefer receiving income earlier in life. Redistribution within households reduces WTP for temporary DI and the bencharm reform to SNAP by small amounts, equal to 3 percent and 4 percent of their total values respectively. WTP for increasing the leniency of DI allowance and increasing the generosity of DI benefits are reduced by larger amounts, equal to 13 pecent and 29 percent of their total values respectively. The reform that allows applicants to work has a positive welfare gain associated with redistribution within households, owing to the fact that younger DI applicants (having lower savings) more often choose to work while applying for DI. Unlike other reforms, allowing applicants to work therefore provides more income to younger individuals than older ones.

Cumulatively, the three non-insurance components of reforms explain 50 percent of WTP for temporary DI, 14 percent of WTP for increasing leniency of DI allowance, and 75 percent of WTP for increasing the generosity of DI benefits. This means that the value of most reforms cannot be explained entirely by the manner in which they redistribute income based on the information set available to households at the beginning of the life cycle. By conditioning transfers on the realizations of adulthood risk, reforms also provide households with valuable insurance.

Insurance against health risks, singlehood risks, and productivity do not explain a large part of the welfare gains from reform, and if anything, reduce the welfare gains that reforms provide. To understand why, consider the households whose behavior is affected by the reform. When each reform is implemented, it causes these households, typically having faced negative shocks to health or wages, to forgo work income in order to instead collect DI benefits. These households, for whom income is highly valuable, consequently lose income from the pre-reform world to the post-reform world, though they typically lose less than if they had forgone work in the pre-reform world. The first three insurance-based components of each reform account for this decline in income from leaving work, but they do not allow households the corresponding benefit of leaving work that motivates the post-reform income decline.

That explanation also makes clear why the final component, consisting of non-pecuniary effects, is highly valuable for most reforms. This is the step in the decomposition where households are provided the benefit of leaving work in bad states of the world that comes with the relative decline in income households face in bad states versus other states, owing to DI reform. In this sense, the value of most DI reforms comes from insuring households against enduring work when suffering disability and when earnings are low, rather than insuring household income. The one exceptional reform is allowing applicants to work, which has the opposite effects—households in bad states of the world benefit from the higher income that comes from working while applying for DI in the first three insurance components, but the final component requires that they actually pay the cost of working in order to earn that income.