Household Labor Supply Responses to Severe Health Shocks and the Gains from Social Insurance*

Itzik Fadlon and Torben Heien Nielsen†

Abstract

This paper studies how households respond to severe health shocks and the insurance role of spousal labor supply. In the empirical part of the paper, we provide new evidence on individuals' labor supply responses to spousal mortality and health shocks. Analyzing administrative data on over 500,000 Danish households in which a spouse dies, we find that survivors immediately increase their labor supply and that this effect is entirely driven by those who experience significant income losses due to the shock. Notably, widows—who experience large income losses when their husbands die—increase their labor force participation by more than 11%, while widowers—who are significantly more financially stable—decrease their labor supply. In contrast, studying over 70,000 households in which a spouse experiences a heart attack or a stroke but survives, we find no economically significant spousal labor supply responses to non-fatal health shocks, consistent with the adequate insurance coverage for their associated income losses in our setting. In the theoretical part of the paper, we show that spousal labor supply responses have direct welfare implications for social insurance against mortality and health shocks. In light of this theoretical result, our empirical findings imply large welfare gains from transfers to survivors and identify efficient ways for targeting government transfers.

*We thank David Autor, Martin Browning, Raj Chetty, Jeffrey Clemens, Julie Cullen, David Cutler, Gordon Dahl, Ariel Dora Stern, Raluca Dragusanu, Liran Einav, Eric French, Elhanan Helpman, Nathan Hendren, Simon Jäger, Xavier Jaraverl, Lawrence Katz, Claus Thustrup Kreiner, David Laibson, Jessica Laird, Søren Leth-Petersen, Magne Mogstad, Joana Naritomi, Arash Nekoei, Jann Speiss, Yoram Weiss, Gal Wettstein, Danny Yagan, and Eric Zwick for helpful discussions and comments. We gratefully acknowledge financial support from the Alfred P. Sloan Foundation through the NBER Household Finance working group (under the direction of Brigitte Madrian and Stephen Zeldes) and the NBER Pre-doctoral Fellowship Program on the Economics of an Aging Workforce (under PI David Card), the Boston College Center for Retirement Research, the NBER Roybal Center for Behavior Change in Health and Savings, and EdLabs at Harvard University. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA or any agency of the Federal Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. The research reported in this publication was also supported by the National Institute on Aging of the National Institutes of Health under Award Number P30AG034532. The content is the sole responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health.

†Fadlon: University of California, San Diego and NBER, address: Department of Economics, University of California, San Diego, 9500 Gilman Drive #0508, La Jolla, CA 92093-0508 (email: fadlon@ucsd.edu); Nielsen: University of Copenhagen, address: Department of Economics, University of Copenhagen, Øster Farimagsgade 5 building 26, DK-1353, Copenhagen K (email: thn@econ.ku.dk).

1 Introduction

Severe illnesses and the subsequent deaths of primary earners are among the most devastating shocks that households face and are a major source of economic risk. The social insurance programs that protect households against the financial vulnerabilities imposed by these shocks—namely, survivors benefits and disability insurance—have become among the largest safety-net programs in most OECD countries in recent decades (OECD 2014).[1]

In this paper we study how households respond to fatal and non-fatal severe health shocks and use that to draw implications for the design of social insurance. In the empirical component of the paper, we provide new and clear evidence on individuals' labor supply responses to spousal mortality and health shocks and analyze the mechanisms that underlie these responses. Overall, we find significant increases in spousal labor supply when income losses are large and households lack adequate formal insurance, consistent with a self-insurance role of household members' labor market behavior. Then, in a simple theoretical framework, we show that beyond their relevance for understanding households' behavior over the life-cycle, these responses have direct welfare implications. We show that spousal labor supply responses to adverse events can be mapped into the household's valuation of additional government transfers. Hence, household labor supply behavior can be used for assessing the benefits on the margin from the sizable social insurance programs that protect against income losses due to mortality and health shocks.

To recover the causal effects of these shocks we develop a new quasi-experimental research design that constructs counterfactuals to affected households by using households that experience the same shock but a few years in the future. We combine event studies for these two experimental groups and identify the immediate and longer-run treatment effects using a "dynamic" differences-in-differences estimator. The identification strategy that we offer relies on the assumption that the timing of the shock within a short period of time is as good as random. Therefore, it is applicable to the analysis of a wide range of other common economic shocks. We conduct the analysis by exploiting long panels of administrative data on the entire Danish population (from the years 1980-2011). These data provide comprehensive information on health-care utilization, income, wealth, and labor market behavior with spousal linkages, and are therefore ideal for our purposes.

Our empirical analysis focuses on the extreme shock of the death of a spouse, which can lead to significant and permanent income losses. Analyzing over 500,000 Danish households of married and cohabiting couples in which a member has died, we find an immediate and persistent increase in survivors' labor supply following their spouse's death. This response amounts to an average increase of 7.6% in survivors' labor force participation and 6.8% in annual labor income by the fourth year after the shock.

The average effects that we find are entirely driven by households that experience substantial

[1] For example, in 2014 the US government paid \$93 billion to more than 6 million surviving spouses and \$132 billion to 9 million disabled workers through the Old-Age, Survivors, and Disability Insurance program, compared to \$46 billion paid in unemployment benefits (SSA 2015; White House 2015).

income shocks due to the loss of a spouse, and therefore have greater need for self-insurance through labor supply. In particular, we show that the mean increase in labor supply is attributable to survivors whose deceased spouses had earned a large share of the household's income, who have less disposable income at the time of the shock, and who are less formally insured by government transfers. We also find that high-earning survivors decrease their labor supply, consistent with the conjecture that their high income is no longer necessary to support two people. Notably, widowers—who tend to be financially stable when losing their wives—decrease their labor supply, while widows—who tend to experience considerably larger income losses when losing their husbands—significantly increase their labor supply. By the fourth year after their husbands die, widows increase their participation by 11.3%, which translates into a 10.1% increase in their annual earnings.

We additionally analyze alternative hypotheses other than self-insurance for the mechanisms that may underlie the average increase in survivors' labor supply. In particular, we find that the evidence is inconsistent with the conjecture that this response is driven by a lower cost of supplying labor (or a higher willingness to work) following the death of a spouse, e.g., due to loneliness and the desirability of social integration or because the survivor no longer has to care for an ill spouse.

In contrast to mortality shocks, we show that for shocks that are well-insured in our setting (through social and private insurance) and require no additional informal insurance, there are no economically significant spousal labor supply responses. Studying over 70,000 households in which one member experiences a non-fatal heart attack or stroke, we find that the earnings of the sick individuals drop by 19% after the shock, while these households' post-transfer income declines by only 3.3%. Consistent with this lack of a significant income drop, there are no notable changes in the spouses' participation with an economically small decline in their labor earnings (of about 1%). The combination of our quasi-experimental research design and rich administrative data allows us to precisely estimate this small response, which has proven difficult in previous studies (e.g., Coile 2004; Meyer and Mok 2013).

Following our empirical analysis, we highlight in a simple theoretical model the importance of studying the labor supply responses of spouses by showing they can inform the design of social insurance. Specifically, we show that these ex-post responses to shocks are an integral component of the assessment of social gains from additional government benefits and, thus, have direct welfare implications. The intuition of this theoretical result is that the extent to which households self-insure income losses using spousal labor supply is directly related to the degree to which they lack formal insurance. Hence, the relative changes in spousal labor supply in response to death and health shocks can reveal the scope for additional welfare-improving government benefits through the social insurance schemes that protect against them.

We conclude with a discussion on the qualitative welfare implications of our findings in light of the simple theoretical framework. First, our results show that the relative increase in survivors' labor supply is substantially larger for older widows. Driven by the differential attachment to the labor force over the life-cycle, which effectively attenuates the financial vulnerability of younger survivors, this finding suggests that survivors benefits should be age-dependent. Second, survivors' labor supply is strongly increasing in the share of the household's income that the deceased had earned, which suggests that survivors benefits should be a function of the deceased spouse's work history. Evidently, both of these features characterize the current large survivors benefits scheme within the Social Security system in the US.

This paper relates to several strands of the literature. First, a significant and active body of research studies the effects of adverse health on *own* labor market outcomes. This includes smaller-sample studies using survey data in the US (e.g., Charles 2003; Chung 2013; Meyer and Mok 2013; Dobkin et al. 2016) as well as recent large-scale studies in countries in which administrative data that link health and own labor market outcomes are available (e.g., Lundborg et al. 2011; Halla and Zweimüller 2013; Pohl et al. 2013; Gupta et al. 2015). Due to the unavailability of large-scale data that combine both administrative health-care utilization and labor market information *and* include household linkages, there is very little direct evidence on the effects of health shocks on *household* labor supply (e.g., Coile 2004 and Meyer and Mok 2013 who use survey data in the US). Our paper provides novel empirical results for the impact of fatal and non-fatal health shocks on spousal labor supply using large-scale data with objective health measures, and to the best of our knowledge, is among the first papers to study labor supply responses to spousal death.

Second, numerous past empirical studies have analyzed spousal labor supply responses to wage and unemployment shocks in order to uncover the extent to which it is used as insurance (what is known as the "added worker effect"). However, while spousal labor supply is commonly modeled as an important self-insurance mechanism against adverse shocks to the household (e.g., Ashenfelter 1980; Heckman and Macurdy 1980; Lundberg 1985), this prior empirical work has been largely unable to find evidence of significant labor supply responses to *temporary* spousal unemployment (e.g., Heckman and Macurdy 1980, 1982; Lundberg 1985; Maloney 1987, 1991; Gruber and Cullen 1996; Spletzer 1997). The leading explanation for this lack of evidence has been that, in the context of temporary unemployment, income losses are small relative to the household's lifetime income and are already sufficiently insured through formal social insurance (Heckman and Macurdy 1980; Cullen and Gruber 2000). Consistent with this explanation, Stephens (2002) and Blundell et al. (2015) find that wives' labor supply is an important consumption insurance device against permanent shocks to husbands' wages, as opposed to (wealth-constant) transitory shocks. To uncover the self-insurance role of spousal labor supply using a different strategy, Cullen and Gruber (2000) study whether it is crowded out by unemployment insurance benefits and find a large crowd-out effect; and in recent work Autor et al. (2015) find similar crowd-out effects in the context of disability insurance. We take an alternative empirical approach in a new setting, different from unemployment shocks, and directly study the effects of severe health shocks with different degrees of income loss: fatal shocks, which impose large and permanent income losses that are only partially insured, and non-fatal shocks, which are well-insured in our setting.

Third, prior related work on estimating welfare gains from social insurance has focused on studying its "consumption smoothing" effects, mostly in the context of disability insurance with little direct work on the gains from survivors benefits.[2] Our conceptual framework offers an alternative approach that uses data from the labor market instead of consumption data, and studies smoothing of spousal labor supply. In practice, each approach has its comparative advantages which we discuss in the paper. In our main context of mortality shocks, studying consumption would under-estimate the benefits from additional survivors benefits: the large increase in widows' labor supply mitigates the consumption drop they would otherwise experience, but comes at a cost of reduced leisure which is accounted for in the analysis of labor supply. Additionally, to address the change in the household's composition following a spouse's death, consumption analysis would require accounting for economies of scale in the household, and would therefore rely on estimates for adult equivalence scales (Blundell and Lewbel 1991). This would be essential for translating changes in household expenditure into actual changes in survivors' individual consumption, in contrast to the analysis of labor supply behavior that is directly assignable to individual household members.

Our theoretical analysis also relates to and builds on recent work on labor market methods for welfare analysis in the context of unemployment, which rely on the labor market behavior of the unemployed (Shimer and Werning 2007; Chetty 2008; Landais 2015). In the shocks that we consider, in which adverse health hurts individuals' ability to work, these methods cannot be applied. The directly affected individual may be unresponsive to economic incentives (or even deceased) and hence cannot reveal the household's preferences through labor market behavior. Exploiting the interplay between the labor supply decisions of household members, our approach uses only the responses of the indirectly affected spouse. As such, our analysis offers a labor market method in a household setting that is also applicable to other economic shocks in which the individual who is directly impacted may be unresponsive to economic incentives or at a corner solution (of either working full-time or not working at all). For example, relevant to the debate on the privatization of Social Security, the value of protecting against pension-wealth losses in the 401(k) account of a working individual can be assessed by the labor supply response of his or her spouse.

The paper is organized as follows. We begin with Section 2 that sets the conceptual framework for the empirical analysis by theoretically illustrating the self-insurance role of spousal labor supply using a model of household labor force participation. Prior to our empirical analysis, Section 3 describes the private and social institutional environment in Denmark and the data sources that we use. In Section 4 we develop our empirical research design for recovering the causal effects of adverse shocks. Our core empirical analysis is presented in Section 5. In this section we estimate individuals' labor supply responses to spousal mortality and health shocks and analyze the heterogeneity in these responses to study their underlying mechanisms. In Section 6 we use our simple

[2] This work includes reduced-form studies in the context of health shocks and the death of a spouse (e.g., Myers et al. 1987; Hurd and Wise 1989; Auerbach and Kotlikoff 1991; Cochrane 1991; Stephens 2001; Bernheim et al. 2003; Meyer and Mok 2013; Chung 2013; Ball and Low 2014; Dobkin et al. 2016) and studies that rely on structural economic modeling in the context of disability insurance and Social Security (e.g., İmrohoroğlu et al. 1995, 2003; Huang et al. 1997; Kotlikoff et al. 1999; Bound et al. 2004; Benitez-Silva et al. 2006; Nishiyama and Smetters 2007; Chandra and Samwick 2009; Bound et al. 2010; Low and Pistaferri 2015; Autor et al. 2015).

theoretical framework to highlight the normative relevance of our estimates, and we briefly discuss the qualitative welfare implications of our findings in Section 7. Section 8 concludes.

2 Conceptual Framework: A Simple Model of Household Labor Force Participation

We begin with analyzing a static unitary model of household extensive labor supply decisions, which demonstrates how both death and health shocks can be analyzed within the same framework. The purpose of this section is to motivate our empirical analysis by formalizing how spousal labor supply can be used as insurance against income shocks to the household. Intuitively, when individuals experience severe health shocks that cause them to decrease their labor supply and earn less income—or when they die—their spouses can compensate for the imposed income loss by increasing their own labor supply. Moreover, the relative increase in spousal labor force participation in response to shocks increases with the income loss, which can reveal the extent to which the household needs to self-insure.

The model that we study here is the simplest possible model that demonstrates the insurance role of spousal labor supply in the context of this paper. Our qualitative arguments, however, extend to much more general settings. We discuss important generalizations to the highly-stylized model, alternative assumptions about the household's behavior and preferences, and their impact on our theoretical results when we return to the normative analysis later in the paper. These extensions include adding dynamics, analyzing intensive vs. extensive margin responses, using the collective rather than the unitary approach to the household's behavior, as well as generalizing the household's production technologies and preferences and allowing for additional margins of response.

Setup. We study labor force participation decisions of a two-person household, which consists of individuals 1 and 2. Our focus is on the extensive margin rather than work intensity to allow for labor market frictions, such as hour requirements set by employers, which can limit employees' ability to optimize on the intensive margin.[3]

We consider a world with two states of nature: a "good" state, state g, in which member 1 is in good health and works; and a "bad" state, state b, in which member 1 experiences a shock and drops out of the labor force. Households spend a share of μ^g of their adult life in state g and a share of μ^b in state b (with $\mu^g + \mu^b = 1$). In what follows, the subscript $i \in \{1, 2\}$ refers to the household member and the superscript $s \in \{g, b\}$ refers to the state of nature.

[3] Evidence for such frictions in the Danish context is found in Chetty et al. (2011). Note that the choice of the appropriate model for welfare analysis should depend on the context. In our specific context, we saw empirically that survivors' responses were concentrated on the participation margin. However, other applications, such as studying a sub-population with full employment before a shock occurs, would call for the intensive-margin model because work intensity is expected to be the operative margin. The simple participation model of this section is most closely related to Kleven et al. (2009) and Immervoll et al. (2011), who study optimal taxation of couples with extensive-margin labor supply responses.

Household Preferences. Denote by c_i^s and l_i^s the consumption and labor supply of member i in state s, respectively. Let $U(c^s; l_1^s, l_2^s)$ represent the household's utility as a function of aggregate consumption, $c^s = c_1^s + c_2^s$, at its optimal allocation across spouses, and the household members' labor force participation, l_1^s and l_2^s, in state s. Since we study extensive margin behavior, $l_i^s = 1$ if i works and $l_i^s = 0$ otherwise. For simplicity, we assume that $U(c^s; l_1^s, l_2^s) = u(c^s) - v_1 \times l_1^s - v_2 \times l_2^s$, where $u(c^s)$ is the household's utility from consumption, and v_i represents each member i's disutility from labor (including the utility value of direct work costs and the opportunity costs of lost home production). To let the model incorporate both the case in which the bad state is when member 1 is sick and the case in which member 1 is deceased, we set $u(c^s) = Q(c^s)$ when both spouses are alive and $u(c^s) = q(c^s)$ when member 1 does not survive. We do not carry a state superscript for $u(c^s)$ for notational convenience, although we allow $Q(c^s)$ and $q(c^s)$ to be arbitrarily different. [4] The couple's disutilities from labor (v_1, v_2) are drawn from a continuous distribution defined over $[0, \infty) \times [0, \infty)$. We denote the marginal probability density function of v_2 by $f(v_2)$ and its cumulative distribution function by $F(v_2)$.

Household Budget Constraint. Each choice of individual 2's employment determines the household's overall income in state s, $y^s(l_2^s)$, such that $y^s(l_2^s) = A^s + \bar{z}_1^s \times l_1^s + \bar{z}_2^s \times l_2^s + B^s(l_2^s)$, where A^s is the household's state-contingent wealth and non-labor income (including life insurance payouts, transfers from any other source of individually-purchased or employer-provided private insurance, transfers from relatives, and medical expenses), and $\bar{z}_i^s$ is i's net-of-tax labor income in state s. $B^s(l_2^s)$ represents transfers from the government in state s, which we allow to depend on 2's participation, so that transfers can be state-dependent as well as earnings-tested at the household level. It is also straightforward to include economies of scale in the household's resource constraint (using an arbitrarily general technology, G, that transforms income into consumption as in Browning et al. 2013, so that $c^s = G(y^s(l_2^s)))$, as well as differential tax rules for joint filing.

Household Behavior. The household consumes its entire disposable income in each state of nature and in each of member 2's employment statuses, so that $c^s(l_2^s) = y^s(l_2^s)$. There are no savings decisions involved in the baseline static model, which we introduce in the dynamic extension to the model. Hence, in the current setting, the household's choices reduce to the labor force participation decision of the unaffected spouse (member 2) in each state s: l_2^s. The unaffected spouse works in state s if and only if

$$v_2 < \bar{v}_2^s \equiv u(y^s(1)) - u(y^s(0)). \tag{1}$$

That is, the spouse works if the household's valuation of the additional consumption coming from his or her labor income compensates for his or her utility loss from working.

Spousal Labor Supply as Insurance. At this point it is easy to see the self-insurance role of spousal labor supply responses to shocks, our outcome of interest. Define y_{-2}^s as the household's resources excluding those directly attributed to 2's labor supply decision—i.e., $y_{-2}^s \equiv A^s + \bar{z}_1^s \times l_1^s$—

[4] We assume that these functions are well-behaved—i.e., that $Q'(c^s) > 0$, $Q''(c^s) < 0$, $q'(c^s) > 0$, and $q''(c^s) < 0$—which implies that $u'(c^s) > 0$ and $u''(c^s) < 0$ in each scenario.

such that the income loss from the shock is $L \equiv y_{-2}^{g} - y_{-2}^{b}$. In addition, for each state s, denote the spouse's probability of participation (or the participation rate of spouses in the population) by $e_2^s \equiv F(\bar{v}_2^s)$. In each state, the spouse's probability of participation decreases in his or her unearned income:

$$\frac{\partial e_2^s}{\partial y_{-2}^s} = f(\bar{v}_2^s)[u'(y^s(1)) - u'(y^s(0))] < 0. \tag{2}$$

This implies that we should expect a relative increase in spousal labor force participation, $\frac{e_2^b}{e_2^g} > 0$, whenever there is some income loss that is not fully insured, $L > 0$. That is, in our simplified model, income shocks lead to self-insurance through the spouse's labor force participation. Furthermore, the spouse's labor supply response to the shock $(\frac{e_2^b}{e_2^g})$ increases in the income loss L:

$$\frac{\partial \left(\frac{e_2^b}{e_2^g}\right)}{\partial L} = \frac{f(\bar{v}_2^b)}{F(\bar{v}_2^g)}[u'(y^b(0)) - u'(y^b(1))] > 0, \tag{3}$$

so that the extent of spousal self-insurance increases in the degree of income loss imposed by the shock. The comparative statics in (2) and (3) are no more than simple income effects at the household level, which are a direct implication of the concavity of $u(c^s)$. With these comparative statics at hand and after theoretically illustrating the insurance role of spousal labor supply, we now turn to the empirical analysis of the impact of spousal mortality and health shocks.

3 Data and Institutional Background

To study labor supply responses to spousal mortality and severe health shocks we leverage rich administrative data from Denmark. The Danish setting is unique in providing full-population register-based data on both health and labor market outcomes, as well as household linkages, and is therefore ideal for our purposes. In this section, we describe the Danish insurance environment, both social and private, as it relates to sick individuals and surviving spouses, and list our data sources.

Institutional Background. It is useful to distinguish between two types of insurance: health insurance (coverage of medical care) and income insurance (insurance against income losses in different health states). Health insurance in Denmark is a universal scheme in which almost all costs are covered by the government, with a few exceptions such as dental care, chiropractic treatments, and prescription drugs that entail a limited degree of out-of-pocket expenses. Therefore, the Danish setting allows us to concentrate on (social and private) income insurance for losses that go beyond immediate medical expenses, as we describe below.

In Denmark, income insurance against severe health shocks and the death of a spouse consists of four main components that are typical of systems in developed countries: temporary sick-pay benefits, permanent Social Disability Insurance, privately purchased insurance policies, and other indirect social insurance programs (such as early retirement and old-age pensions).

During the first four weeks after a health shock occurs, workplaces are obliged to provide the sick employee with sick-pay benefits, which fully replace wages as long as the employee is ill within this period. Some common agreements and work contracts insure wage earnings against sicknesses of longer duration. For example, some blue-collar common agreements in the private sector provide wages during periods of sickness for up to one year. If the sick worker's contract does not provide such a scheme, then the local government must provide flat-rate sick-pay benefits from the fifth up to the fifty-second week after the worker has stopped working. In 2000, for example, a sick worker received a fixed daily rate that added up to DKK 11,400 ($1,425) per month (the same as the prevailing unemployment benefit rate).

If the worker remains sick and is unable to work, he or she can apply at the municipality level for Social Disability Insurance (Social DI) benefits that will provide income permanently. For example, in 2000, subject to income-testing against overall household income, a successful application amounted to DKK 110,400 ($13,800) per year for married or cohabiting individuals and DKK 144,500 ($18,000) for single individuals.

Importantly for our analysis, the Danish Social DI program has a broad social insurance scope. It can be awarded for vaguely defined social reasons, to individuals who are deemed to be in need of government assistance. In 1984 the notion of social reasons came to replace a complex mix of programs, such as survivors benefits for women and special old-age pensions for single women (where the motive behind this rule change was that the pre-1984 rules discriminated between genders). Therefore, Social DI is the effective social insurance mechanism for surviving spouses who are unable to maintain their standard of living after losing their partners. Indeed, we find sharp increases in the take-up rate of Social DI by survivors immediately after their spouses die. Hence, we refer to Social DI in the context of spousal mortality shocks as social survivors benefits.

While Social DI and its surviving benefits component are state-wide schemes, they are locally administered. Regional councils (in a total of 15 regions) decide whether to approve or reject an individual's application, and municipal caseworkers (in a total of 270 municipalities) administer the application and handle all aspects of each case—including any contact with the applicant, preparation of the application, and collection of financial and health status records. The local administration of the program, combined with the vague notion of awarding benefits for social reasons, has led to differential application processing behavior across municipalities. In turn, it has resulted in substantial variation in rejection rates—ranging from 7% to 30%—and thus in the mean receipts of the program's benefits across the different municipalities over time (Bengtsson 2002). We exploit this municipality-year variation in the awarding of the survivors benefits component of the program later in the paper.

Another source of income to a household that experiences health shocks or in which a member dies is payments from an employer-based insurance policy, an element that is standard in labor-market pension plans. Since 1993, most sectors covered by common agreements (75% of the labor force) have mandatory pension savings, part of which consists of life insurance and insurance against specific health shocks. These pay out a lump-sum to the sick worker, as long as he or she is making contributions to the pension plan, or to the surviving spouse in case the plan member dies. The rates of these payouts are set by the individual pension funds. In addition, some individuals can purchase insurance policies of a similar structure in the private non-group market.[5]

Lastly, there are old-age social insurance programs that can indirectly protect eligible survivors or households that experience other shocks, who can decide to take them up at different ages according to their financial needs. At age 60 and until they reach their old-age pension retirement age, individuals who have (voluntarily) been members of an unemployment fund for a sufficiently long period (10 years before 1992 and gradually increasing to 20 years thereafter) are eligible for the Voluntary Early Retirement Pension (VERP). Approximately 80% of the population is eligible for VERP, which provides a flat-rate annual income of roughly DKK 130,000 ($16,250). At the "full-retirement" age of 67 (or 65 for those born after July 1, 1939) all residents become eligible for the Old-Age Pension (OAP), which provides income-tested annuities of up to DKK 99,000 ($12,375) per year for singles and DKK 75,000 ($9,375) for coupled individuals (at 2000 rates). Note that the benefits to single survivors who qualify for social survivors benefits are reduced at the full-retirement age (from $18,000 to $12,375) when the program transitions into the Old-Age Pension.[6]

Data Sources. We have merged data from several administrative registers to obtain annual information on Danish households of married and cohabiting couples from 1980 to 2011. We use the following registers: (1) the national patient register, which covers all hospitalization records (from both private and public hospitals), and from which we extract information on all the individuals that experienced a heart attack or a stroke; (2) the cause of death register, from which we identify death dates; (3) income registers, which include *all* sources of household income—e.g., labor income, capital income, annuity payouts, and government benefits from any program—as well as annual measures of gross wealth and liabilities that include bank-account balances and lump-sum transfers from insurance companies;[7] and (4) the Integrated Database for Labor Market Research, which includes

[5] Note that the private market for life insurance in Denmark is not negligible. A recent study by the Danish government suggests that approximately two thirds of the survivors surveyed received or expected to receive some transfer from a life insurance policy (see http://www.raadtilpenge.dk/~/media/PPP/Mister%20aegtefaelle/Grafikrapport.ashx). However, there is a potentially important role for government intervention in the context of mortality shocks for the following reasons. First, purchasing life-insurance products in Denmark requires answering a health status and behaviors questionnaire (and even undergoing medical exams) and, therefore, applications by unhealthy or older households are likely to be rejected. These rejections by the insurance companies can be explained by private information that is held by these rejected households (Hendren 2013). Second, it is common that even when the life-insurance product is purchased by younger and healthy households (both in group and non-group markets) the coverage sharply declines with age, leaving older and unhealthy Danish households with no coverage through the private market against spousal death. Moreover, in our sample, which spans 1980-2011, the life-insurance coverage rate is likely smaller compared to recent years because life-insurance holdings experienced a significant increase in the last decade due to the expansion of the labor market pension schemes that sometimes include a life-insurance component.

[6] An additional small government-mandated pension scheme (for all wage earners in Denmark) that supplements the OAP is the ATP program. This program pays out a life annuity to individuals who reached full-retirement age, based on the number of years they contributed to the scheme. In 2003, for example, the average annual payout from the scheme amounted to DKK 4,900 ($612). Unlike the OAP, there is a life insurance element tied to this scheme, albeit negligible relative to the other labor-market based (as well as privately-purchased) life insurance policies. Until 2002 a surviving spouse was eligible for 30% of the capitalized value of the deceased spouse's remaining ATP benefits. Since 2002 survivors are instead eligible for a lump sum of DKK 40,000 ($5,000), taxed at 40%, if the deceased spouse is younger than 67 at death (which progressively reduces with the deceased's age at death and entirely lapses if the spouse dies after age 70).

[7] In our main analysis sample of spousal mortality shocks, the net assets of the median household amount to only DKK

demographic variables for the entire population as well as measures for full-time and part-time employment for individuals younger than 60. These full-time and part-time employment measures are constructed using records of employees' payments to the government-mandated ATP pension scheme. The mandatory level of payments into this program is a one-to-one function of employment status, where full-time employment is defined as working at least 30 hours per week all 12 months of the calendar year ("full-time full-year"), and part-time employment is defined as working at some point during the year but either fewer than 30 hours per week or fewer than 12 months within the calendar year.

All monetary values are reported in nominal Danish Kroner (DKK) deflated to 2000 prices using the consumer price index. In that year the exchange rate was approximately DKK 8 per US $1. We postpone describing the summary statistics of the analysis sample to the next section since they directly relate to the discussion on the advantages of our research design.

4 Research Design

The goal of our empirical analysis is to identify the causal (ex-post) effects of spousal mortality and health shocks on individuals' labor supply, $\frac{e_w^b}{e_w^g}$. In this section we describe the empirical strategy that we develop for overcoming the selection challenges inherent in the identification of these effects. We also report summary statistics of the analysis sample to display the comparability of our treatment and control groups.

4.1 Quasi-Experiment

The ideal experiment would randomly assign shocks to households and track labor supply responses over time. Therefore, we need to compare the responses to shocks of affected households to the counterfactual behavior of ex-ante similar unaffected households. This requires comparing households with same expectations over the distribution of future paths, but with different realizations, to isolate the "unanticipated" component of the shock. The access to over three decades of administrative panel data on the universe of Danish households allows us to develop a quasi-experimental research design that mimics this ideal experiment, by exploiting the potential randomness of the timing of a severe health shock or death within a short period of time.

To do so, we look only at households that have experienced the shocks that we consider at some point in our sample period, and identify the treatment effect from the timing at which the shock was realized. We construct counterfactuals to affected households using households that experience the same shock a few years in the future, and recover the treatment effect by performing event studies for these two experimental groups. Note that a simple event study, which analyzes the evolution of outcomes of a treated group around the time of a shock, is not appropriate for our application. Pure event studies identify short-run responses, while we are interested in identifying longer-run effects because of potential delays in adjustment (due to, e.g., labor market frictions). This necessitates

13,236 ($1,655) while the median annual household-level income is DKK 239,922 ($29,990). Therefore, our analysis of labor supply responses focuses on income losses, and we use the wealth data in robustness checks.

a control group, as we construct in our design, that can account for complex life-cycle trends in the counterfactual behavior absent a shock (in, e.g., spousal labor force participation as depicted in Appendix Figure 1).

Before formally describing our research design, we illustrate with a concrete example its basic intuition of the similarity of households that experience shocks close in time.

Illustrative Example. Let us focus on a treatment group of individuals born between 1930 and 1950 who experienced a severe health shock, in particular, a heart attack or a stroke, in 1995. Consider studying the effect of the shock on some economic outcome of these individuals, e.g., their labor force participation. Panel A of Figure 1 plots the outcome for these households against the outcomes for households that have not experienced this shock in our sample period, and reveals very different life-cycle patterns across the two groups prior to 1995. This suggests that traditional matching estimators, which use these unaffected households as a control group, are inappropriate for our application as their validity will rely on the set of available controls and on the unconfoundedness assumption.[8] Panel B of Figure 1 plots the outcome for the treatment group of households that experienced a shock in 1995 as well as for households that experienced the same shock in 2010 (15 years later), in 2005 (10 years later), in 2000 (5 years later), and in 1996 (1 year later). Studying the behavior of households that experienced the shock in different years reveals increasingly comparable patterns to those of the treatment group's behavior—in terms of trends before 1995—the closer the year in which the individual experienced the shock was to 1995. These patterns confirm our intuition and suggest using households that experienced a shock in 1995+$\triangle$ as a control group for households that experienced a shock in 1995. Panel D of Figure 1 displays a potential control group when we choose $\triangle = 5$.

Our method generalizes this example by aggregating different calendar years. Simply put, our design conducts event studies for two experimental groups: a treatment group composed of households that experience a shock in year τ, and a matched control group composed of households from the same cohorts that experience the same shock in year $\tau + \triangle$. We identify the treatment effect purely from the trend in the difference in outcomes in each year across the two groups.

The trade-off in the choice of $\triangle$, which captures the main weakness of our design, can be seen in Panel C of Figure 1. On the one hand, we would want to choose a smaller $\triangle$ such that the control group is more closely comparable to the treatment group, e.g., year 1996 which corresponds to $\triangle = 1$. On the other hand, we would want to choose a larger $\triangle$ in order to be able to identify longer-run effects of the shock, up to period $\triangle - 1$. For example, using those who experienced a shock in 2005 ($\triangle = 10$) will allow us to estimate the effect of the shock for up to 9 years. However, this entails a potentially larger bias since the trend in the behavior of this group is not as tightly parallel to that of the treatment group. Our choice of $\triangle$ is five years, such that we can identify effects up to four years after the shock. We assessed the robustness of our analysis to this choice

[8] This assumption requires that conditional on observed covariates there are no unobserved factors that are associated both with the treatment assignment and with potential outcomes (Imbens and Wooldridge 2009).

and found that local perturbations to $\triangle$ provide very similar results.

Formal Description of the Design and Estimator. Fix a group of cohorts, denoted by Ω, and consider estimating the treatment effect of a shock experienced at some point in the time interval $[\tau_1, \tau_2]$ by individuals who belong to group Ω. We refer to these households as the treatment group and divide them into sub-groups indexed by the year in which they experienced the shock, $\tau \in [\tau_1, \tau_2]$. We normalize the time of observation such that the time period, t, is measured with respect to the year of the shock—that is, $t = year - \tau$, where $year$ is the calendar year of the observation. As a control group, we match to each treated group τ the households among cohorts Ω that experienced the same shock but at $\tau + \triangle$ for a given choice of $\triangle$. For these households we assign a "placebo" shock at $t = 0$ by normalizing time in the same way as we do for the treatment group ($t = year - \tau$).[9]

Denote the mean outcome of the treatment group at time t by y_t^T and the mean outcome of the control group at time t by y_t^C and choose a baseline period (or periods) prior to the shock (e.g., period $t = -2$), which we denote by p (for "prior"). For any period $n > 0$, the treatment effect can be simply recovered by the differences-in-differences estimator

$$\beta_n \equiv \left(y_n^T - y_n^C\right) - \left(y_p^T - y_p^C\right). \tag{4}$$

The treatment effect in period n is measured by the difference in outcomes between the treatment group and control group at time n, purged of the difference in their outcomes at the baseline period, p. Note that the choice of $\triangle$ puts an upper bound on n such that $n < \triangle$.

The identifying assumption is that, absent the shock, the outcomes of the treatment and control groups would run parallel. In particular, in accordance with the differences-in-differences research design, there is no requirement regarding the levels of outcomes. The plausibility of this assumption relies on the notion that within the short window of time of length $\triangle$ the particular year at which the shock occurs may be as good as random. To test the validity of our assumption, we accompany our empirical analysis with the treatment and control groups' behavior in the five years prior to the shock year 0 in order to assess their co-movement in the pre-shock period. By showing that there are virtually no differential changes in the trends of the treatment and control groups before period 0, we alleviate concerns that the two groups may still differ by, e.g., their expectations over the particular year of the shock within our chosen five-year window of $\triangle$.[10]

Other papers, which use empirical strategies different from ours, rely on similar identifying assumptions. These include, for example, studies in the context of the long-run effects of job

[9] By construction, their actual shock occurs at $t = \triangle$.

[10] Conceptually, as long as there is no perfect foresight we can use our strategy with the appropriate choice of $\triangle$. This choice is context dependent and requires empirical investigation, where any potential difference across the experimental groups would be included in the bias consideration in the choice of $\triangle$. Comparability is then an empirical question that can be investigated in several ways, such as: (1) analyzing sub-samples of shocks that are more likely to come as a surprise; (2) studying the robustness of the results to a rich set of controls; and the strategies we mentioned above: (3) testing for parallel trends in the pre-period; and (4) investigating the sensitivity of the results to the chosen control group by changing $\triangle$. We conduct this set of tests in our application and verify the robustness of our results in support of our underlying identifying assumption. The analysis of tests (1)-(3) appears in the paper, and the analysis of (4) is available from the authors on request.

displacement (Ruhm 1991) and the effect of arrests on employment and earnings (Grogger 1995).[11] Our quasi-experimental design can be applied to these shocks and to any other shock whose exact timing is likely random, which can be validated in any particular setting by studying the pre-trends of the experimental groups.

Estimating Equation. We use the simple "dynamic" differences-in-differences estimator of equation (4) to study the evolution of household responses in the graphical event studies that we conduct below. To quantify the mean treatment effects, we estimate the regression counterpart of this estimator, averaged over the years after the shock. In the regression results we also explicitly report statistical significance and account for controls as a robustness check. Our baseline estimating equation is of the form:

$$y_{i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_3 treat_i \times post_{i,t} + \beta_4 X_{i,t} + \alpha_i + \varepsilon_{i,t}. \quad (5)$$

In this regression, $y_{i,t}$ denotes an outcome for household i at time t; $treat_i$ denotes an indicator for whether a household belongs to the treatment group; $post_{i,t}$ denotes an indicator for whether the observation belongs to post-shock periods; $X_{i,t}$ denotes a vector of potential controls; and α_i is a household fixed effect. The parameter β_3 represents the average causal effect of spousal shocks on household outcomes.

4.2 Analysis Sample and Summary Statistics

Appendix Table 1 displays key summary statistics for the analysis sample. The sample of our main analysis includes households in which one spouse died between ages 45 and 80 and is comprised of 310,720 households in the treatment group and 409,190 households in the control group.

The table reveals the advantage of our research design—the comparability of the year of observation and the age of unaffected spouses across experimental groups. The average survivor in the treatment group loses his or her spouse in 1993 at age 62.86 and the average unaffected spouse in the control group experiences the placebo shock in year 1993 at age 62.27. The sub-sample of survivors under age 60, the age at which there is a large drop in labor force participation (due to eligibility for early retirement benefits as shown in Appendix Figure 1), displays even closer similarities. By construction, the research design nets out calendar year effects. However, due to the randomness of the exact timing of the shock (and without directly matching on age), it also effectively nets out life-cycle effects by comparing groups of almost identical ages.

The sample for our secondary analysis of severe health shocks includes households in which one spouse experienced a heart attack or a stroke (for the first time) and survived for at least three years. These shocks are among the leading causes of death in the developed world and their timing within a short period of time is likely unpredictable. Since the average age of spouses precisely at the time of these health shocks is just over 60 (60.67), we focus on households with both spouses

[11] More recent examples include the study by Hilger (2014), who exploits variation in the timing of fathers' layoffs to study the effect of parental income on college outcomes, and the work by Persson and Rossin-Slater (2016) who exploit variation in the exact timing of deaths in the family to study the effect of mothers' stress during pregnancy on children's well-being.

under 60 to ensure that the results we document are driven only by the health shocks and not by eligibility for early retirement benefits.[12] The sample consists of 37,432 households in the treatment group and 54,926 households in the control group. The unaffected spouse is on average 45.7 years old in the treatment group at the time of the shock and 45.3 years old in the control group, where the mean calendar year of the shock is around 1992 for both groups.[13]

5 Spousal Labor Supply Responses

In this section, we present our primary analysis of the impact of mortality and health shocks on spousal labor supply. We begin with the focus of our study, the extreme shock of spousal death. We present the average labor supply responses of survivors and then analyze the heterogeneity of these responses to uncover the mechanisms through which they may operate. In particular, guided by the comparative statics in (2) and (3), we study how survivors' behavior varies by the degree of income loss imposed by the death of their spouse and by the extent of coverage through survivors benefits to study the self-insurance role of spousal labor supply. We also analyze other potential mechanisms using a simple test that aims at assessing the extent to which survivors' willingness to work may change in response to the shock. Then, we briefly study households' labor supply responses to non-fatal severe health shocks—specifically, heart attacks and strokes—for which we do not expect spousal labor supply responses as self-insurance since their resulting labor income losses are formally well-insured in our setting.

5.1 Labor Supply Responses to the Death of a Spouse

5.1.1 Mean Responses

Panels A and B of Figure 2 plot the average labor supply responses of individuals whose spouse died between ages 45 and 80. Panel A reveals an immediate increase in labor force participation (defined as having any positive level of annual earnings) following the death of a spouse. By the fourth year after the shock, the surviving spouses' participation increases by 7.6% – an increase of 1.6 percentage points (pp) on a base of 20.6 pp. Panel B of Figure 2 shows that this response translates into a 6.8% increase in annual earnings (including zeros for those who do not work). Appendix Figure 2 repeats this analysis for survivors whose spouses experienced a heart attack or a stroke for the first time and died within the same year. This allows us to focus on deaths that are more likely to be "unexpected" and come as a surprise, and for which we have a sufficient number of observations. As seen in the figure, the pre-trends, levels, pattern of response, and response magnitudes are all very similar to those in the sample of deaths from any cause.[14]

With significant disparities in baseline participation rates and labor income, men and women

[12] The results do not change, however, when we look at the unconstrained sample.

[13] We also report the means of main labor supply outcomes in Appendix Table 1 for completeness. Note that participation and earnings are slightly higher for the control group, which poses no threat to the validity of the design since comparability requires similar trends and not similar levels.

[14] Similar results were found for the small sample of "accidental" deaths. The analysis is available from the authors on request.

may face substantially different financial distress when they lose their spouse and, therefore, may respond differently to this shock. Indeed, Panels C and D of Figure 2 reveal clear differences in the responses of widowers and widows. While on average widowers do not change their labor force participation when their wife dies, widows immediately and significantly increase their labor force participation when they lose their husband. Four years after the shock, widows' labor force participation increases by 2.2 pp from a baseline participation rate of 19.5 pp, which amounts to a large increase of 11.3% in their labor force participation.

This differential response suggests that female survivors have greater need to self-insure through labor supply and that they experience greater income losses when they lose their spouse as compared to their male counterparts. To test this conjecture, we plot the evolution of overall household income (from any source) around the death of a spouse, including earnings, capital income, annuity payouts, and benefits from social programs. We begin by plotting the household's income in the absence of behavioral responses from the unaffected spouse in order to capture the income loss directly attributable to the loss of an earning spouse. To do so, we plot in Panel A of Appendix Figure 3 the household's overall income, holding the unaffected spouse's earnings and social benefits at their pre-shock level.[15]

Before discussing this figure, it is useful to mention benchmarks for the changes that we observe in household income in order to interpret their magnitude. Following a spousal death, the household's composition changes so that insuring the private consumption of surviving spouses does not require the entire 100% of the household pre-shock income. At the same time, potential economies of scale within the household can make 50% of the household's income before the shock insufficient for the survivors to maintain the same level of utility (see, e.g., Nelson 1988; Browning et al. 2013). The share of the household's income that would keep individuals' consumption utility at its pre-shock level is usually assumed to lie between 50% and 100% and is commonly referred to as the adult "equivalence scale". Some commonly used scales are the modified OECD equivalence scale of 67% and the square-root scale of 71%.[16] One would expect surviving spouses to broadly compensate for income losses with respect to this general benchmark, such that declines in household income on the order of 29%-33% would not require self-insurance through labor supply. Note that we merely mention this benchmark to gauge magnitudes. Our analysis of income losses compares relative losses across gender sub-groups, and does not rely on assumptions regarding the exactly appropriate compensation level.[17]

Panel A of Appendix Figure 3 shows that widowers, who do not change their labor supply on

[15]Specifically, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their level in $t = -1$.

[16]The implicit equivalence scale in the Danish Social DI is approximately 0.65 and is 0.66 in the Old-Age Pension (see Section 3). The relevant equivalence scales that we mention here are for adults because the median age of the youngest child of our treated individuals born after 1930 (for whom we have data on children) is 30, with only 10% having a youngest child under 18.

[17]Indeed, as we mentioned in the introduction, the difficulty of knowing the correct benchmark for full compensation is one of the issues that make the analysis of consumption fluctuations (relying on changes in income flows) challenging for welfare evaluations in the context of spousal mortality.

average, experience an overall decline of 32% in household income. However, as expected by their different labor supply responses, widows experience a significant relative loss of 8 pp compared to widowers, so that the potential decrease in household income is 25% larger for female survivors. To study the **actual** change in household income, Panel B of Appendix Figure 3 takes into account the surviving spouses' labor supply responses and any change in the benefits they may receive from social or private insurance. The figure shows that widowers experience an actual decline of 31% and that widows manage to decrease their additional potential loss – through the increase in labor supply and higher take-up of social insurance – to incur an actual decline of 35%. Overall, widows' labor supply responses account for 22% of the 5 pp shrinkage in their potential income loss (from 40% to 35%).Note that surviving spouses do not fully compensate for a loss in household income (to 100%) since as singles they do not need the full pre-shock level of income.

Younger Households. The baseline labor supply behavior of individuals below and above age 60 substantially differ. At age 60 there is a sharp drop in participation when 80% of the labor force becomes eligible for early retirement benefits. This implies that surviving spouses under age 60 have a much stronger attachment to the labor force and significantly higher labor earnings and are, therefore, implicitly more financially resilient after the loss of an earning spouse. Consistent with this view, Panels A and B of Figure 3 reveal that widows under 60 exhibit a much smaller relative increase in labor supply compared to the universe of widows. Their "shock elasticities" amount to a mere increase of 3.3% in participation and 3.2% in annual earnings. Interestingly, widowers under 60 respond with a decrease in their labor supply, which amounts to a 4.1% decline in their annual earnings. The majority of these widowers (74%) earned more than their deceased wives and, compared to widows, they have significantly larger baseline participation rates (0.78 vs. 0.715) and average labor income (DKK 227,560 vs. DKK 138,232). Therefore, their behavior is consistent with the notion that they no longer support two people in the household and that, as singles, they may not require the entire amount of their high labor income to meet their consumption needs. Put together, the behavior of younger households suggests that their higher participation rates and annual earnings effectively insure them better against losing an earning spouse. We further test this hypothesis directly in Section 5.1.2, where we analyze heterogeneity in survivors' responses by the level of their own pre-shock earnings.

For younger survivors, the registers additionally consist of administrative measures for full-time and part-time employment as defined in the data section. This allows us to further investigate the dynamics of spousal labor supply behavior and to distinguish between extensive and intensive margin responses. As we show in Figure 4, in periods 0 and 1 there are temporary transitions to part-time work, consistent with spending time with the dying spouse and mourning his or her loss. These transitions stabilize thereafter such that the active decision margin becomes full-time work vs. non-participation.

For completeness, we report in Appendix Table 2 estimates for the regression counterparts of the main figures that we presented so far by using the specification of equation (5). To account for the full-time/part-time dynamic responses that we documented above, $post_{i,t}$ assumes the value 1 for periods 2 to 4 (when the response stabilizes on the full-time margin). We present in the table the average treatment effects and their statistical significance, and verify the robustness of our results to the inclusion of year, age, and household fixed effects.

5.1.2 Heterogeneity in Responses by Income Losses and Degree of Insurance

We continue with further investigation of the heterogeneity in the survivors' labor supply responses across different subgroups to provide evidence that is consistent with the insurance mechanism hypothesis. Importantly, using different strategies we show that the responses are proportional to the loss of income that survivors experience when their spouse dies, and depend on the survivors' degree of financial stability and level of insurance.

Within-Gender Analysis of Heterogeneity by Income Loss. We begin by studying the effect of the death of a spouse on labor force participation by the degree of income loss for each gender separately. To this end, for each household we calculate the *potential* income loss due to the shock in the following way.

First, similarly to Panel A of Appendix Figure 3, we calculate for each household the overall income (from any source) holding the unaffected spouse's earnings and social benefits at their pre-shock level (in $t = -1$). Second, we calculate the ratio of this "potential income" measure in $t = 1$ to the household's income in $t = -1$. Third, we normalize this ratio for the treated households by the mean ratio of the control households in order to purge life-cycle and time effects. This leaves us with a measure of the potential income replacement rate for each treated household, which we denote by rr_i, that captures the change in household income directly attributed to (and only to) the loss of a spouse. This measure is smaller whenever the deceased spouse's relative contribution to the household's income was larger. Notably, it accounts for the deceased spouse's income from any source: labor earnings, private or social retirement income, and benefits from government programs. Therefore, as required for our analysis, it also captures the income loss imposed by the death of older non-working individuals who receive income from sources other than the labor market.

To study the heterogeneity in labor supply responses by the income replacement rate (rr_i), we augment the baseline differences-in-differences model of equation (5) and estimate the following specification:

$$l_{i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_{3i} treat_i \times post_{i,t} + \beta_4 X_{i,t} + \alpha_i + \varepsilon_{i,t}, \qquad (6)$$

where

$$\beta_{3i} = \beta_{30} + \beta_{31} rr_i + \beta_{32} Z_{i,t}.$$

In this regression $l_{i,t}$ denotes an indicator for the labor force participation of the unaffected spouse in household i at time t. We adjust the basic differences-in-differences design by allowing the treatment effect, β_{3i}, to vary across households and model it as a function of the household's potential replacement rate rr_i. Our parameter of interest is β_{31}, which captures the extent to which the surviving spouse's labor supply response correlates with the income loss he or she experiences. Since β_{31} can capture other dimensions of heterogeneity beyond the income replacement rate, we let the treatment effect vary with additional household-level characteristics, $Z_{i,t}$, such that β_{31} further isolates the treatment effect's partial correlation with the loss of household income.[18]

Table 1 reports the results of estimating (6) separately for each gender, with and without $Z_{i,t}$, for the entire sample of surviving spouses and for only the sub-sample of survivors under age 60. The results consistently show throughout the specifications the strong correlation between labor supply responses and income losses: survivors in households with lower potential income replacement rates (lower rr_i), who experience larger income losses, are much more likely to increase their labor force participation in response to the shock. Specifically, it implies larger increases in spousal labor supply among households in which the deceased had earned a larger share of the household's income. Since controlling for the additional interactions with $Z_{i,t}$ does not change the results much, the evidence suggests that the heterogeneous responses are indeed driven by differential income replacement rates. In addition, the estimation results reveal very similar sensitivity to income losses across genders; so that re-weighting the female and male sub-samples using the regression in (6) to match on pre-shock own and spousal income would lead to similar average responses across genders. This strengthens the conjecture that unobserved gender differences (e.g., in preferences) do not drive the differential average labor supply responses across female and male survivors, but rather their divergent income losses.

Responses by Own Earnings. The heterogeneity in responses with respect to the loss in household income that we have analyzed so far has focused on income changes *relative* to pre-shock flows. An additional strategy for studying this sort of heterogeneity focuses on the *levels* of the surviving spouses' disposable income available at the time of the shock. To do this, we turn to analyze how labor supply responses of surviving spouses may vary with their own level of earnings when their spouses die, since higher-earning survivors have more disposable income and are therefore effectively better insured.

We constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than that of their experimental-group-specific 20th percentile. Then, for each household we calculate the pre-shock labor income share of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force. Specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households in which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household.[19]

We divide the remaining sample into five equal-sized groups according to the survivors' pre-shock level of earnings, and plot in Panel A of Figure 5 the average labor income response (as well

[18]The variables we include in $Z_{i,t}$ are age dummies for the surviving spouse, dummies for the age of the deceased at the year of death, year dummies, indicators for the number of children in the household as well as the surviving spouse's months of education (and its square). The results are also robust to the inclusion of a quadratic in the household's net wealth. Note that $X_{i,t}$ always includes the variables in $Z_{i,t}$ as well as their interaction with $treat_i$ and $post_{i,t}$.

[19]These restrictions also imply that the results below for this sub-sample are mainly driven by intensive-margin responses.

as its 95-percent confidence interval) against the pre-shock mean earnings for each group.[20] The figure reveals a strong gradient of labor supply responses with respect to the survivors' own level of earnings when the shock occurs. Survivors at the bottom of the income distribution increase their annual earnings by 7.79% in order to meet their consumption needs, while those at the top decrease their earnings by 2.93%. As for widowers younger than 60, the behavior of the high-earning survivors is consistent with the notion that their high income is no longer necessary to support two people and that they may find lower levels of income sufficient for their consumption needs as singles.

Since the household's pre-shock labor income is composed of two earners, we need to account for the pre-shock earnings of the dying spouse. Hence, we divide the sample into two groups – households in which the dying spouse's pre-shock labor income fell within the bottom three quintiles of its group-specific distribution, to which we refer as "low earners", and households in which the dying spouse's pre-shock labor income fell within the top two quintiles, to which we refer as "high earners". Panels B and C of Figure 5 reveal that the gradient prevails in both sub-samples, such that surviving spouses with lower earnings are much more likely to increase their labor supply when their spouse dies, regardless of whether their spouse was a high or low earner. Panel A of Appendix Table 3 shows that this relationship is robust to the inclusion of dummy variables for age and year (as well as to the inclusion of a quadratic in the household's net wealth) by separately estimating the corresponding differences-in-differences equation for each surviving spouses' quintile. Note that merely analyzing the average earnings response in this sample would have masked the substantial heterogeneity we documented. Panel B of Appendix Table 3 shows that the average labor income increase for this sub-sample is DKK 585 (0.39%) and is not statistically different from zero.

Spatial Variation in Social Insurance over Time. Lastly, we take advantage of spatial variation in the administration of social survivors benefits to study survivors' labor supply responses by the generosity of social insurance. This allows us to test the hypothesis that the self-insurance mechanism underlies spousal labor supply responses using variation in the household's income that is plausibly exogenous. It also allows us to analyze whether better social insurance crowds out labor supply increases in response to shocks. Consistent with our heterogeneity analysis so far, we find that the increase in survivors' participation due to the shock declines in the formal insurance they receive from the government which mitigates their income loss.

For this analysis, we constrain the sample to survivors under 67 (the age at which the program transitions into the Old-Age Pension) and to the period prior to 1994 due to a data break in the reporting method of survivors benefits received through Social DI. In addition, since for this sample the increase in the take-up of the program following the shock is attributable to females, we focus the analysis on widows. Inclusion of widowers does not change the qualitative results. Panel A of Appendix Figure 4 clearly displays the insurance role of Social DI for widows, whose take-up of the program increases by more than 50% in the year that their husbands die.

[20] Note that pre-shock earnings of a survivor is calculated as average earnings in periods -5 to -2. This measure smooths out transitory wage shocks and excludes periods just before the spouse's death so that potential responses in anticipation of it do not affect this measure.

Recall that while Social Survivors Benefits is a state-wide program, it is locally administered so that regional councils decide whether to approve or reject an application and municipal caseworkers (in a total of 270 municipalities) administer the application and handle all aspects of each case. Since this structure and the vague definitions for eligibility criteria have led to substantial variation in rejection rates across municipalities, it has created significant variation in the mean receipts of the program's benefits across the different municipalities over time (Bengtsson 2002).

We use these year-by-municipality average receipts as an instrument for actual receipts. In particular, we calculate for each municipality the average survivors benefits received by non-working surviving spouses through Social DI in each year. Then, we assign to each widow of household i in the treatment group the respective mean in municipality m at time t excluding her own benefits (the "leave-one-out" mean), denoted by $\overline{SB}_{-i,t,m}$.[21] We estimate the following augmented differences-in-differences regression:

$$l_{i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_{3i} treat_i \times post_{i,t} + \beta_4 X_{i,t} + \varepsilon_{i,t}, \quad (7)$$

where

$$\beta_{3i} = \beta_{30} + \beta_{31} SB_{i,t}.$$

In this regression, $l_{i,t}$ denotes the participation of the unaffected female spouse of household i at time t, and $X_{i,t}$ includes municipality m's unemployment rate and average earnings (and their interaction with $treat_i$, $post_{i,t}$, and $treat_i \times post_{i,t}$), as well as age, year, and municipality fixed effects. $SB_{i,t}$ are actual social survivors benefits receipts, measured in annual DKK 1,000 ($125) terms, for which we instrument using $\overline{SB}_{-i,t,m}$ (where the F-statistic on the excluded instrument in the first stage is 24.3). The identifying assumption is that, given our set of controls, the average of social survivors benefits transferred to other widows in a municipality in a given year affects a widow's participation only through its influence on her own survivors benefits receipts. Note that the source of variation we use is within municipalities over time since we include municipality and calendar year fixed effects as controls.

The two-stage least squares results are presented in Table 2. The estimate for our parameter of interest, β_{31}, is -.0057. With an average of DKK 23,262 ($2,908) in actual survivors benefits receipts by widows in the analysis sample (including zeros for those not on the program) and a participation rate of 0.5054, this estimate translates to a participation elasticity with respect to social benefits of -0.26 for widows under 67.[22] This implies that formal social insurance crowds out labor supply responses which otherwise provide an informal self-insurance mechanism against loss of income following the death of a spouse.

In summary, the results reveal a clear pattern: there are significant average increases in labor supply in response to losing a spouse, which are entirely driven by households that experience large income losses due to this shock. At the same time, survivors whose pre-shock earnings were high and represented a large share of the household's income decrease their labor supply, consistent with the change in the household's composition so that they no longer financially support their spouse.

[21] The variation in this instrument is displayed in Panel B of Appendix Figure 4.

[22] For a sense of scale, estimates for the net-of-tax participation elasticity are on the order of 0.25 (see, e.g., Chetty 2012).

Put together, the results provide strong evidence of the self-insurance role of spousal labor supply in the extreme case of the death of a spouse, which translates into large and permanent income losses for most households in our setting due to incomplete insurance of spousal mortality.

5.1.3 Alternative Mechanisms: Survivors' Labor Disutility State Dependence

Besides income losses, there are other important ways in which households can be directly affected by mortality shocks that can drive our results. For example, potential changes in the unaffected spouse's labor disutility (or willingness to work) can directly lead to spousal labor supply responses even when households are well-insured. We are specifically interested in testing the hypothesis that the increase in survivors' labor supply can be attributable to lower costs of supplying labor following the death of a spouse, due to loneliness and the desirability of social integration or because the survivor no longer has to care for an ill spouse. In this section, we briefly discuss a simple and intuitive strategy to test this conjecture. In Fadlon and Nielsen (2015) we develop a complementary formal test that leads to similar conclusions.

Consider widows, for whom we find an increase in participation in response to spousal death, who did not work before the shock in a model where time is divided between labor and leisure (or any other use of "time at home"). Widows in households in which the deceased spouse did not work before his death experience smaller income losses (taking into account the deceased's income from any source including government transfers), but consumed more joint leisure and hence may be more likely to experience loneliness. Similarly, if the deceased spouses in these households did not work because they were potentially ill in the years preceding their death, their widows are also more likely to have taken care of them, thereby having more time available for market work when their husbands die. Overall, survivors in these households are more likely to experience a decrease in the utility cost of labor supply. In contrast, widows in households in which the deceased spouse worked before his death consumed less joint leisure (or care-giving time), but experience larger income losses. The social integration (or "loneliness") hypothesis and the hypothesis of decreased care-giving time are consistent with spouses in the first group of households increasing their labor supply more than spouses in the latter group do, while the self-insurance hypothesis is consistent with the opposite pattern.

In Table 3 we first verify the differential level of income losses across the two groups of widows (column (3)). Studying their labor supply, we find that the increase in the labor force participation of survivors in households in which the deceased worked is much larger (by 4.61 pp), with a negligible effect for survivors in households in which the deceased did not work (0.78 pp). Moreover, among households in which the deceased did not work and received low levels of income, there was no increase in the widows' labor supply. Hence, the evidence is inconsistent with the conjecture that the mean increase in the surviving spouses' labor supply is driven by lower cost of labor following the shock, so that the analysis consistently supports the view that this increase is driven by self-insurance of large income losses.

5.2 Labor Supply Responses to Spousal Health Shocks

As a complementary analysis, we study in this section households' labor supply responses to non-fatal severe health shocks. Recall that our analysis sample for this shock consists of households in which a spouse experienced a heart attack or a stroke (for the first time) and survived for at least four years (until $t = 3$), and in which both spouses were under age 60.

Panel A of Figure 6 shows that within three years of the shock, the affected spouses' participation sharply falls, which translates into a large loss of annual earnings. Appendix Table 4 quantifies these effects by estimating a differences-in-differences regression, in which we allow for differential treatment effects in the "short run" (periods 1 and 2) and the "medium run" (period 3), to account for the gradual responses documented in Panel A of Figure 6.[23] Columns 2 and 4 of Appendix Table 4 reveal that by the third year after the shock the labor force participation rate of the sick spouses drops by 12 pp – about 17% – and that annual earnings drop by DKK 36,015 ($4,500) – a significant drop of 19%.

However, while there is a significant drop in the sick spouses' earnings, Columns 5 and 6 of Appendix Table 4 show that the actual loss of income that these households experience is much smaller and amounts to only 3.3% of overall household income. That is, taking into account the entire household income, including any transfers from social or private sources (particularly Disability Insurance), reveals that these shocks are very well-insured in our Danish setting. Consistent with this lack of a notable income drop, there are no economically significant labor supply responses among unaffected spouses (as shown in Panel B of Figure 6 and Columns 7 to 10 of Appendix Table 4) as there is no significant need to self-insure.[24] In line with our previous findings, this behavior is consistent with self-insurance being the driving mechanism of spousal labor supply responses to shocks, here in a context where this form of informal insurance is not exercised since households are formally well-insured.

Note that the rich data and our research design allow for a precise estimation of these economically insignificant spousal responses to shocks.[25] In particular, our results imply a small but positive degree of complementarity in spouses' labor supply in response to health shocks, with an estimate of 0.065 for the unaffected spouse's earnings elasticity with respect to the affected spouse's earnings. Since the household's income is not perfectly insured, this response likely implies some degree of health-state dependence of the household's utility. Intuitively, the fact that given a small loss of income due to the shock the unaffected spouses' decrease in labor supply involves an additional (very

[23] We estimate the following specification

$$y_{i,t} = \beta_0 + \beta_1 treat_i + \beta_{2a} post^a_{i,t} + \beta_{3a} treat_i \times post^a_{i,t} + \beta_{2b} post^b_{i,t} + \beta_{3b} treat_i \times post^b_{i,t} + \alpha_i + \varepsilon_{i,t}, \quad (8)$$

where $y_{i,t}$ denotes an outcome for household i at time t, $post^a_{i,t} = 1$ in periods 1 and 2 and zero otherwise, and $post^b_{i,t} = 1$ in period 3 and zero otherwise. Therefore, β_{3a} captures the "short run" effect, and β_{3b} captures the "medium run" effect.

[24] As for spousal mortality shocks, we find a strong correlation between spousal labor supply responses and income losses (expressed in household income replacement rates) in the context of health shocks using a specification similar to equation (6). The results are reported in Appendix Table 5.

[25] Note that Meyer and Mok (2013) similarly find that the typical disabled individual in the US loses about 21% in earnings but only 6.75% in post-transfer household income by the fourth year after the shock.

small) loss (through their lower earnings) is consistent with two main state dependence channels. First, it is consistent with households in the bad state valuing income less than do households in the good state – i.e., a consumption utility state dependence. Second, it is consistent with an increase in the unaffected spouses' utility loss from time spent away from home either because they would like to take care of their sick spouse or due to preferences for joint leisure – i.e., a labor disutility state dependence. With no additional assumptions, one can only reach conclusions about the ratio of these two types of potential state dependence (see Fadlon and Nielsen 2015 for a formal analysis).

6 Theory: Welfare Implications of Spousal Labor Supply Responses

Having studied households' labor supply responses to mortality and severe health shocks, we now move on to their normative interpretation. Using our stylized model from Section 2, the purpose of this section is to illustrate that the degree to which spouses respond to shocks is an integral component in assessing the welfare gains from social insurance. We show that the relative change in spousal labor force participation in response to shocks can reveal the extent to which households lack formal insurance and need to self-insure, because the magnitude of this response increases with the degree of income loss. Hence, spousal labor supply responses can inform us of the scope for additional welfare-improving government transfers through more generous social insurance. We begin with the simple model and then discuss important extensions and generalizations.

6.1 Gains from Social Insurance

In this section we analyze the planner's problem of designing social insurance programs that protect households against the financial vulnerabilities imposed by death and health shocks. Our goal is to map how individuals respond to spousal shocks into predictions about the welfare gains from providing more generous social insurance through these programs.[26]

Policy Tools. The planner observes the state of nature as well as the employment status of each spouse. Since some spouses work and earn more than others do, the optimal policy is dependent on whether the spouse is employed. We denote the tax on spouse i's labor income in state g by T_i^g and the monetary benefits given to non-working spouses in state g by m^g. In state b, households receive transfers of the amount M^b, with additional benefits of the amount m^b to lower-income households in which the indirectly-affected member 2 does not work. This tax-and-benefit structure allows for the analysis of flexible policy designs and mimics features of existing social insurance programs in most developed countries (such as income-testing that characterizes the Supplemental Security Income program within the Old-Age, Survivors and Disability Insurance in the US and the Social Disability Insurance in Denmark).[27] We denote taxes by $T \equiv (T_1^g, T_2^g)$ and benefits by

[26] The main rationale for government intervention in our setting is private information. See Hendren (2013) for relevant evidence in the context of life and disability insurance markets.

[27] The exact way in which we model transfers simplifies the analysis and is not necessary for our results. Any system that conditions transfers on the state of nature and employment, as well as on age in the dynamic model, can be analyzed within our framework.

$B \equiv (m^g, M^b, m^b)$.

Planner's Problem. Let $W^s(v_2)$ denote the household's value function in state s such that

$$W^s(v_2) \equiv \begin{cases} u(y^s(1)) - v_1 \times l_1^s - v_2 & \text{if } v_2 < \bar{v}_2^s \\ u(y^s(0)) - v_1 \times l_1^s & \text{if } v_2 \geq \bar{v}_2^s. \end{cases}$$

Therefore, the household's expected utility is $J(B,T) \equiv \mu^g \int_0^\infty W^g(v_2) f(v_2) dv_2 + \mu^b \int_0^\infty W^b(v_2) f(v_2) dv_2$. The social planner's objective is to choose the tax-and-benefit system that maximizes the household's expected utility subject to the requirement that expected benefits paid, $\mu^g(1-e_2^g)m^g + \mu^b\left(M^b + (1-e_2^b)m^b\right)$, equal expected taxes collected, $\mu^g\left(T_1^g + e_2^g T_2^g\right)$. Hence, the planner chooses the benefit levels B and taxes T that solve

$$\max_{B,T} J(B,T) \quad \text{s.t. } \mu^g(1-e_2^g)m^g + \mu^b(M^b + (1-e_2^b)m^b) = \mu^g(T_1^g + e_2^g T_2^g). \tag{9}$$

Welfare Benefits from Social Insurance. What is the welfare gain from providing more generous benefits if the bad state occurs? To answer this question, consider transferring resources from the good state g to the bad state b, by a small decrease in the benefit m^g to finance a corresponding balanced-budget increase in m^b.[28]

The social gain from this perturbation consists of the household's valuation of additional insurance. To construct a measure for this valuation, consider first the household's utility loss from a \$1 reduction in m^g to finance the additional insurance. This loss is captured by $\frac{\partial J(T,B)}{\partial m^g} = \mu^g \frac{\partial}{\partial m^g}\left(\int_0^\infty W^g(v_2)f(v_2)dv_2\right) = \mu^g(1-e_2^g)u'(c^g(0))$, since every dollar taken produced a value of $u'(c^g(0))$ and was transferred to to household with probability $\mu^g(1-e_2^g)$. Partially differentiating the government's budget, this reduction in m^g allows a balanced-budget increase in m^b of the amount $\left|\frac{\partial m^b}{\partial m^g}\right| = \frac{\mu^g(1-e_2^g)}{\mu^b(1-e_2^b)}$. The household's valuation per \$1 increase in m^b is given by $\frac{\partial J(T,B)}{\partial m^b} = \frac{\partial}{\partial m^b}\left(\int_0^\infty W^b(v_2)f(v_2)dv_2\right) = \mu^b(1-e_2^b)u'(c^b(0))$, as it produces a value of $u'(c^b(0))$ with probability $\mu^b(1-e_2^b)$. The overall utility gain from the increase in benefits when the shock occurs is, therefore, $\frac{\partial J(T,B)}{\partial m^b} \times \left|\frac{\partial m^b}{\partial m^g}\right|$.

Put together, the welfare benefits from a (balanced-budget) increase in m^b financed by a \$1 decrease in m^g is $\frac{\partial J(T,B)}{\partial m^b} \times \left|\frac{\partial m^b}{\partial m^g}\right| - \frac{\partial J(T,B)}{\partial m^g}$. To gain cardinal interpretation for this expression, we follow the recent social insurance literature and normalize it by the baseline welfare gain from a \$1 transfer through m^g (Chetty and Finkelstein 2013). That is, the normalized welfare benefit from our policy change is[29]

$$MB \equiv \frac{\frac{\partial J(T,B)}{\partial m^b} \times \left|\frac{\partial m^b}{\partial m^g}\right| - \frac{\partial J(T,B)}{\partial m^g}}{\frac{\partial J(T,B)}{\partial m^g}} = \frac{u'(c^b(0)) - u'(c^g(0))}{u'(c^g(0))} = \frac{u'(c^b(0))}{u'(c^g(0))} - 1. \tag{10}$$

[28] In the simple model, this perturbation concerns the distribution of transfers to low-income households across different health states. Other perturbations to the system will follow the steps of the analysis conducted below. We focus on this particular aspect of the policy since it captures the essence of insuring households against shocks in a mathematically simple way.

[29] Our analysis resembles the derivation of individuals' willingness to pay for additional unemployment insurance in Hendren (2015). Deriving these benefits can be also achieved by characterizing the first-order conditions of the planner's problem as in Chetty (2006a), Chetty and Finkelstein (2013), and Fadlon and Nielsen (2015) (see also our analyses of the dynamic model and the intensive margin model in the online appendix).

The marginal benefit from a balanced-budget increase in m^b is captured by the insurance value of transferring resources from the good to the bad state, which is measured by the gap in marginal utilities of consumption across the two states. This "rate of return" on shifting funds, which is zero in the first-best allocation in which marginal utilities are smoothed across states of nature, measures market inefficiency and quantifies the potential benefit from government intervention. This expression in exactly the benefit side of Baily's (1978) and Chetty's (2006a) formula for the optimal level of social insurance applied to our setting.[30]

Consumption-Based Method for Identifying Welfare Benefits. The main reduced-form approach for directly assessing the welfare gains from social insurance in (10) is the method that was developed by Baily (1978) and Chetty (2006a) and was first implemented by Gruber (1997) in the context of unemployment insurance. To provide a simple description of it here, consider the case in which the household's consumption utility exhibits constant relative risk aversion, so that $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$ and $\gamma \neq 1$. In this case, we can express (10) by $MB = \left(\frac{u''(c^b(0))}{u''(c^g(0))}\right) \times \left(\frac{c^b(0)}{c^g(0)}\right) - 1 \cong \gamma \times \left(\frac{c^g(0)-c^b(0)}{c^g(0)}\right)$ (where the latter approximation follows from a Taylor expansion). This formula evaluates fluctuations in the consumption of goods across states, $\frac{c^b(0)}{c^g(0)}$, with the relative rate of change in the utility from marginal dollars, captured by the curvature of the utility function (measured either with $u''(c^s(0))$ or γ). Put differently, the benefits from insurance can be evaluated by the analysis of "quantity" fluctuations in consumption, which are then "priced" in utility terms.

Labor Supply Representation of Welfare Benefits. We show next that simple implications of the household's labor supply decisions allow us to rewrite the marginal benefit in (10) in terms of the unaffected spouse's labor supply. By doing so, we show that the gains from additional insurance can be alternatively measured by evaluating changes in the consumption of the spouse's leisure instead of changes in the household's consumption of goods. The following proposition summarizes this welfare result. We provide a simple proof and then discuss the intuition behind the welfare formula.

Proposition 1. *The marginal benefit from raising m^b through a balanced-budget decrease in m^g is*

$$MB = \Phi \times \left(\frac{e_2^b}{e_2^g}\right) - 1, \tag{11}$$

where $\Phi \equiv \phi^b/\phi^g$, $\phi^s \equiv \frac{|\varepsilon(e_2^s, m^s)|}{m^s \times f(\bar{v}_2^s)}$, and $\varepsilon(e_2^s, m^s)$ is the unaffected spouse's participation elasticity with respect to the policy tool m^s.

Proof. Recall that the unaffected spouse works when the value of additional consumption from his or her labor income, $\bar{v}_2^s \equiv u(y^s(1)) - u(y^s(0))$, outweighs his or her disutility from labor, v_2.

[30] When contemplating a policy change, the planner must weigh this benefit against the associated cost. The cost of transferring \$1 across states is due to within-state behavioral responses, which lead to fiscal externalities that households impose on the government budget when changing their participation decisions. For example, in our case the government's revenue could decrease since more generous social insurance will lead to fewer spouses choosing to work in the bad state (see Fadlon and Nielsen 2015 for details). Identifying marginal costs is conceptually straightforward and much of the social insurance literature has focused on their estimation in different contexts. Therefore, we abstract from their analysis in this paper.

This decision rule reveals the household's consumption value of an additional dollar, $u'(y^s(0))$, through the change in the critical labor-disutility threshold below which the spouse works ($\bar{v}_2^s$) in response to an increase in benefits, since $\left|\frac{\partial \bar{v}_2^s}{\partial m^s}\right| = u'(y^s(0))$.[31] Hence, we can rewrite the marginal benefit from social insurance using the change in the marginal entrant's disutility of labor, and to represent MB with labor supply responses of the unaffected spouse using the equalities $e_2^s = F(\bar{v}_2^s)$ and their semi-elasticities, $\varepsilon(e_2^s, m^s)/m^s = \frac{f(\bar{v}_2^s)}{F(\bar{v}_2^s)}\frac{\partial \bar{v}_2^s}{\partial m^s}$. Conceptually, the proof shows that extensive margin responses can be used for welfare analysis of **marginal** policy changes for the following reason: extensive margin responses are governed by the labor disutility of the marginal entrant; and the labor disutility of the marginal entrant can be mapped into consumption preferences through the household's optimization. Within a state, marginal entrants are always indifferent regarding whether to work or not; it is the cross-state analysis that entails information on the benefits from additional insurance.[32]

Intuition. This formula shows that the marginal net benefit from social insurance can be expressed using different moments of the unaffected spouse's labor supply. The leading term, $\Phi \times \left(\frac{e_2^b}{e_2^g}\right)$, which captures the gross marginal benefit, has two parts: $\frac{e_2^b}{e_2^g}$ that consists of a cross-state response; and Φ that consists of within-state responses. This leading term essentially expresses the relative gain from additional consumption of goods in the bad state by using its mirror image of the relative gain from additional consumption of leisure in the bad state, due to less need to compensate for the income loss associated with the shock through spousal labor supply. Put differently, the formula assesses the benefits of incrementally smoothing labor supply across states as a result of additional formal social insurance that reduces costly self-insurance.

Similar to the consumption representation, our formula evaluates benefits by multiplying the change in the "quantity" of spousal labor supply in response to shocks, $\frac{e_2^b}{e_2^g}$, by the change in the relative "price" of labor force participation across the two states, Φ. In the comparative statics of our model in equation (3) we saw that this "quantity" term increases with income losses and captures the self-insurance role of spousal labor supply. Intuitively, the welfare benefits from additional transfers are higher whenever this quantity is larger, since stronger need to self-insure implies lack of adequate formal insurance and a greater scope for welfare-improving social insurance.

To see how Φ captures the relative price, or utility cost, of spousal labor supply, let us consider its components. Within a state s, the extent to which spousal labor force participation responds to

[31] This is isomorphic to differentiating the first-order condition for the spouse's search effort in a search model of labor force participation (see Online Appendix A).

[32] In our simple model with one-dimensional heterogeneity over labor disutility, the marginal entrant identifies consumption preferences for infra-marginal households (by extrapolation). With multi-dimensional heterogeneity, so that households also differ by consumption preferences, identification requires having marginal households for each "type" of consumption preferences. The average marginal utility of consumption is then assessed using the average participation responses taken over all households. This is equivalent to having all types at an interior solution in search or intensive margin models (such as in Chetty 2008 or in Online Appendices A and C). See Chetty (2008) and Hendren (2013) for such aggregations of private willingness to pay for social insurance. In addition, for other applications in which the identification of household preferences and welfare evaluations rely on **discrete** choices similar to labor force participation responses see Chetty (2006b) and Chetty (2009).

a policy variation in the benefits m^s is captured by $\varepsilon(e_2^s, m^s)/m^s$, the percent change in within-state participation. This measure is proportional to $\frac{\partial \bar{v}_2^s}{\partial m^s}$, the change in the labor disutility of the marginal entrant, which identifies the cost of additional labor supply on the margin in the extensive margin model. Hence, $\varepsilon(e_2^s, m^s)/m^s$ reveals information on the social marginal cost of survivors' labor supply in each state. However, this semi-elasticity is also proportional to the share of marginal households, $f(\bar{v}_2^s)$. Therefore, we must normalize it by $f(\bar{v}_2^s)$ to achieve a within-state "price" change attributable to preferences only, which yields the within-state measure $\phi^s \equiv \frac{|\varepsilon(e_2^s, m^s)|}{m^s \times f(\bar{v}_2^s)}$. The relative change in prices across states of nature, required for evaluating cross-state transfers, is then $\Phi \equiv \phi^b/\phi^g$.[33]

Note that the cross-state term $\frac{e_2^b}{e_2^g}$ is not policy specific and is present whenever the planner considers additional insurance by transferring resources from the good to the bad state. However, the price term, Φ, adjusts according to the policy change and involves within-state semi-elasticities with respect to the specific policy tools that the planner considers changing.

Comparison of Methods. The two methods, which are equivalent in the simple model, have practical comparative advantages in different settings. The main advantage of analyzing welfare benefits by assessing the utility cost of consumption fluctuations is that it aims at directly estimating the gap in marginal utilities of consumption. As such, it does not require that households' optimization leads them to an "interior" solution in other choice variables which, in our case, means that the threshold value of the marginal entrant's labor disutility is within the support of the labor disutility distribution (see Finkelstein et al. 2015 for a related discussion). The main limitations of this approach are twofold: accurate longitudinal data on overall consumption expenditure across health states is largely uncommon (Chetty and Finkelstein 2013; Bee et al. 2013; Pistaferri 2015); and welfare calculations are sensitive to the value of risk aversion, for which there is a wide range of context-dependent estimates (Chetty and Szeidl 2007; Chetty and Finkelstein 2013; Low and Pistaferri 2015).[34] The important advantage of evaluating benefits by studying labor supply responses is the wide availability of large-scale accurate data from the labor market, and the use of directly-estimated changes in participation rates and labor supply elasticities. However, this approach relies

[33] Note the comparison to the intensive-margin version of our model in Online Appendix C. In the model here, the marginal entrant, $\bar{v}_2^s$, reveals the cost of labor supply on the margin, and the change in the marginal entrant's disutility, $\frac{\partial \bar{v}_2^s}{\partial m^s}$, evaluates the relative cost of labor supply across states of nature. In an intensive-margin model in which $v(l_2^s)$ is the spouse's labor disutility function and l_2^s is the choice of work intensity, it is the cost of the marginal hour of work, $v'(l_2^s)$, that reveals the cost of labor supply on the margin, and the change in the cost of the marginal hour, $v''(l_2^s)$ (or $\varphi \equiv \frac{v''(l_2^g)}{v'(l_2^g)} l_2^g$), is used to evaluate the cross-state labor supply responses. This is equivalent to the use of $u''(c^s)$ (or γ) in the consumption-based approach. In the appendix, we show that for general utility functions $MB \cong \gamma \times \left(\frac{c^g - c^b}{c^g}\right) \cong \varphi \times \left(\frac{l_2^b - l_2^g}{l_2^g}\right)$. As in the extensive-margin formula, the intensive-margin formula consists of the "quantity" of spousal labor supply responses to shocks, $\left(\frac{l_2^b - l_2^g}{l_2^g}\right)$, evaluated in utility terms using φ, which can be estimated using within-state elasticities (similar to the estimation of risk aversion in Chetty 2006b).

[34] There have been recent important advances in the analysis of consumption which deal with some of these data limitations, both by using survey data that cover a wide set of non-durable and services expenditure (Blundell et al. 2015; Low and Pistaferri 2015), and by creating consumption measures from income and wealth registers (Browning and Leth-Petersen 2003; De Giorgi et al. 2012; Koijen et al. 2014; Autor et al. 2015; Kolsrud et al. 2015).

on the presence of marginal households in the participation model. That is, identification would not be achieved for sub-populations in which all unaffected spouses never work (e.g., due to significant labor market frictions) or work full-time prior to the shock. Additionally, for some policy changes, such as the simple one that we analyze here, full identification of benefits would require additional calibrations of the labor disutility distribution.[35]

In summary, we have shown using a simple model that the ex-post causal effects of mortality and health shocks on spousal labor supply, $\frac{e_2^b}{e_2^g}$, have direct welfare implications. Even in the dynamic extension to our model, where ex-ante responses are explicitly accounted for, ex-post responses to a shock (rather than in expectation of it) assess the gain from additional *conditional* benefits. The intuition behind this conceptual result is as follows: when forward-looking households make adjustments in anticipation of shocks (according to their expectations), their responses after the shock is experienced recover its residual uninsured risk. This leftover risk assesses the insurance gap that the government can efficiently fill. Empirically, when ex-ante responses are possible, one has to carefully choose research designs that cleanly recover the causal impact of shocks, which are required for identifying households' willingness to pay for benefits (see discussions on these issues in Chetty 2008 and Hendren 2015). Identifying these exact responses was the goal of our empirical analysis.

6.2 Extensions and Generalizations

The stylized model that we analyzed is the simplest possible model that demonstrates our positive and normative arguments. However, as we mentioned earlier, the qualitative arguments that we made so far extend to much more general settings. We discuss below some main extensions to the simple model.

Specifically, in the context of social insurance over the life-cycle, it is important to consider households' self-insurance through ex-ante mechanisms such as precautionary savings. In Online Appendix A, we analyze a generalized, fully-dynamic life-cycle model. This model allows for endogenous savings, as well as private and informal insurance arrangements, and can incorporate a general class of arbitrary choice variables (such as time investment in home production). Our formula extends to this model with the adjustment that post-shock responses in the static case are replaced by *mean* responses when a shock occurs in the dynamic case. Since this is what we recover

[35] Specifically, in our case, since Φ involves the ratio $\frac{f(\bar{v}_2^g)}{f(\bar{v}_2^b)}$, full identification of welfare benefits would require calibrating this ratio. This can be done in different ways. In our application, the empirical analysis of spousal participation across states of nature is consistent with $\bar{v}_w^g$ and $\bar{v}_w^b$ being within a small region of the support $[0, \infty)$. Hence, one can plausibly invoke a locally linear approximation of F in the threshold region $(\bar{v}_w^g, \bar{v}_w^b)$, so that $\frac{f(\bar{v}_2^g)}{f(\bar{v}_2^b)} \cong 1$. Note that this approximation is isomorphic to a second-order approximation of the search effort function in a search model of participation that we analyze in Online Appendix A. However, if one wishes to avoid approximations, one can accompany the analysis with assumptions regarding the family of distributions to which F belongs, and then calibrate its parameters with the participation rates observed in the data. For example, in our analysis of mortality shocks, assuming that F is log-normally distributed and normalizing its mean and standard deviation to 0 and 1, respectively, implies that $\frac{f(\bar{v}_2^g)}{f(\bar{v}_2^b)} = 0.99$.

in our empirical analysis, our findings readily apply to the dynamic setting. The generality of our normative results stems from the fact that they are derived using optimality conditions implied by the household's labor supply choices and using the envelope theorem. Since these conditions hold in more complex models, the economic forces that underlie the assessment of the gains from social insurance using spousal labor supply remain similar in more general settings.[36]

We consider additional noteworthy extensions to the model. In Online Appendix B, we introduce state dependence in preferences and allow for flexible consumption-leisure complementarities. In Online Appendix C, we analyze an intensive-margin version of our model. In Online Appendix C we also show how our results apply to models of the household other than the unitary framework by using the collective approach to household behavior (Chiappori 1988, 1992; Apps and Rees 1988). Further extensions, details, and discussions are available in the NBER working paper version of this paper (Fadlon and Nielsen 2015).

7 Welfare Implications of Empirical Results

We have established that the importance of studying spousal labor supply goes beyond understanding households' behavior in response to shocks over the life-cycle. We have shown that spousal responses to mortality and health shocks have direct welfare implications, which was the purpose of our normative theoretical analysis of the previous section. In particular, we have illustrated within our simple model that the gains from providing additional social insurance are higher when spousal relative labor supply responses to shocks are larger. In this section, we rely on this comparative statics to discuss the potential *qualitative* welfare implications of our empirical findings on the target of the analysis: $\frac{e_2^b}{e_2^g}$. *Quantitatively* assessing the welfare gains from more generous benefits is not our goal. This would require identifying the components in Φ which is not the aim of this paper. [37]

The heterogeneity analysis of survivors' labor supply revealed two notable patterns. First, studying the behavior of widows below and above 60, the age at which there is a sharp decline in labor force participation, we found a significantly smaller relative increase in labor supply for younger survivors. Appendix Figure 5 explicitly compares the relative change in participation for widows below and above 60. It shows that younger widows increase their labor force participation by 3.3%, while older widows increase their participation by 45.4%. To gauge magnitudes, consider the case when $\Phi = 1$, so that the formula in equation (11) reduces to $MB = \frac{e_2^b}{e_2^g} - 1$. In this case, an additional \$1 transferred to survivors could create welfare gains equivalent to 3.3 cents and 45.4 cents for widows younger and older than 60, respectively. This result suggests that the differential attachment to the labor force over the life-cycle, which reduces the financial vulnerability of younger

[36] This is a feature of the sufficient statistic approach to welfare analysis (see Chetty 2006a).

[37] We refer the interested reader to Fadlon and Nielsen (2015) for suggestive calibrations. Note also that the nature of the following discussion implicitly makes the unsubstantiated assumption that Φ is constant across sub-groups. This is for illustrative purposes only.

survivors, may justify age-dependent social insurance for spousal mortality shocks.

Second, we found that increases in the surviving spouse's labor supply are strongly correlated with the share of the household's income the deceased had earned. This suggests that it may be welfare improving to let survivors benefits increase in the deceased spouse's pre-shock share of annual household earnings. A similar pattern of heterogeneity in responses was found for non-fatal spousal health shocks, so that disability benefits may also be more efficiently distributed if dependent on the disabled spouse's work history.[38]

It is worth noting that all of these features characterize the American system. Survivors in the US are eligible for benefits through their deceased spouse's Social Security entitlement only after age 60, so that survivors benefits in the US discontinuously increase in age. These benefits are also a function of the deceased's work history and are therefore increasing in the labor income the deceased had earned. The case is similar for disability benefits, that are based on the disabled worker's pre-shock average earnings.

Overall, our findings revealed significant heterogeneity in responses across different *pre-shock* dimensions of household characteristics, which suggests that enriching the policy tools to condition transfers on these observable characteristics may be welfare improving. Of course, any policy consideration that conditions transfers on pre-shock characteristics should take into account potential ex-ante behavioral responses with respect to these margins.

8 Conclusion

This paper provides new evidence on households' labor supply responses to fatal and non-fatal severe health shocks and uses these responses to draw implications for the design of social insurance. Studying the critical event of the death of a spouse, we find large increases in the surviving spouses' labor supply driven by households for whom this event imposes significant income losses. Analyzing households in which an individual has experienced a severe health shock but survived, for whom income losses are well-insured, we find no significant spousal labor supply responses. Together, the results provide strong support for self-insurance against large and persistent income losses as the mechanism that underlies spousal labor supply responses to shocks. They also point to a potential explanation for the elusiveness of the insurance role of spousal labor supply in previous literature. In support of the hypotheses raised by Heckman and MaCurdy (1980) and Cullen and Gruber (2000), we find that spousal labor supply plays a significant self-insurance role when the income loss incurred by the shock is large relative to the household's lifetime income – as in the death of a spouse – and is irrelevant when the loss is sufficiently insured through formal social insurance – as in spousal health shocks.

We show the importance of studying spousal labor supply behavior in response to shocks for

[38]There may be countervailing arguments when distributional considerations are involved. We abstract from these potentially important considerations since our focus, and the main goal of this section, is to draw qualitative conclusions about how to efficiently target benefits for the purpose of insuring households against the income loss imposed directly by mortality and health shocks.

its welfare implications. The degree to which households self-insure through spousal labor supply reveals their lack of formal insurance and, hence, the scope for more generous social insurance. Based on this result, we discuss the implications of our findings for potentially improving efficiency in the distribution of government benefits. The significant heterogeneity in responses that we find across different pre-determined dimensions of household characteristics imply that richer policy tools that condition transfers on these observable characteristics may be welfare improving.

More broadly, our quasi-experimental design for identifying the effect of shocks as well as our method for welfare analysis can be applied to other important economic questions. Our research design, which relies on comparing households that are affected only a few years apart, can be applied to estimating the effect of a shock in any setting in which its particular timing is likely to be random. When large-scale labor market data are readily available, our welfare analysis, which relies on spousal labor supply, can be used to assess the gains from social insurance in settings in which the directly affected individual may be at a corner solution. For example, in the debate on the privatization of Social Security, the value of protecting against pension-wealth losses in the 401(k) account of a working individual can be recovered by the labor supply response of his or her spouse. Spousal labor supply can also be used to evaluate the welfare losses caused by the discontinuation of an employee's compensation, such as health insurance, as well as the value of unemployment insurance for the long-term unemployed (whose long durations of unemployment significantly harm their employment prospects).[39]

[39] See Kroft, Lange, and Notowidigdo (2013) on the adverse effect of longer unemployment spells.

References

Aguila, E., O. Attanasio, and C. Meghir (2011). Changes in consumption at retirement: Evidence from panel data. *Review of Economics and Statistics 93*(3), 1094–1099.

Apps, P. F. and R. Rees (1988). Taxation and the household. *Journal of Public Economics 35*(3), 355–369.

Ashenfelter, O. (1980). Unemployment as disequilibrium in a model of aggregate labor supply. *Econometrica 48*(3), 547–564.

Auerbach, A. J. and L. J. Kotlikoff (1991). The adequacy of life insurance purchases. *Journal of Financial Intermediation 1*(3), 215–241.

Autor, D., A. Kostøl, and M. Mogstad (2015). Disability benefits, consumption insurance, and household labor supply. Unpublished manuscript.

Baily, M. N. (1978). Some aspects of optimal unemployment insurance. *Journal of Public Economics 10*(3), 379–402.

Ball, S. and H. Low (2014). Do self-insurance and disability insurance prevent consumption loss on disability? *Economica 81*(323), 468–490.

Bee, A., B. D. Meyer, and J. X. Sullivan (2013). The validity of consumption data: Are the consumer expenditure interview and diary surveys informative? In *Improving the Measurement of Consumer Expenditures*. University of Chicago Press.

Bengtsson, S. (2002). Bestemmer forvaltningen om du faar foertidspension? -kommunens forvaltningspraksis og tilkendelse af foertidspension. Technical Report 02:15, Socialforskningsinstituttet.

Benitez-Silva, H., M. Buchinsky, and J. Rust (2006). Induced entry effects of a $1 for $2 offset in ssdi benefits. Unpublished manuscript.

Bernheim, B. D., K. G. Carman, J. Gokhale, and L. J. Kotlikoff (2003). Are life insurance holdings related to financial vulnerabilities? *Economic Inquiry 41*(4), 531–554.

Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2007). Collective labour supply: Heterogeneity and non-participation. *Review of Economic Studies 74*(2), 417–445.

Blundell, R. and A. Lewbel (1991). The information content of equivalence scales. *Journal of Econometrics 50*, 49–68.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2015). Consumption inequality and family labor supply. *American Economic Review (forthcoming)*.

Bound, J., J. B. Cullen, A. Nichols, and L. Schmidt (2004). The welfare implications of increasing disability insurance benefit generosity. *Journal of Public Economics 88*(12), 2487–2514.

Bound, J., T. Stinebrickner, and T. Waidmann (2010). Health, economic resources and the work decisions of older men. *Journal of Econometrics 156*(1), 106–129.

Browning, M., P.-A. Chiappori, and A. Lewbel (2013). Estimating consumption economies of scale, adult equivalence scales, and household bargaining power. *Review of Economic Studies 80*(4), 1267–1303.

Browning, M. and S. Leth-Petersen (2003). Imputing consumption from income and wealth information. *The Economic Journal 113*(488), F282–F301.

Chandra, A. and A. A. Samwick (2009). Disability risk and the value of disability insurance. In *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*, pp. 295–336. University of Chicago Press.

Charles, K. K. (2003). The longitudinal structure of earnings losses among work-limited disabled workers. *Journal of Human Resources 38*(3), 618–646.

Chetty, R. (2006a). A general formula for the optimal level of social insurance. *Journal of Public Economics 90*(10-11), 1879–1901.

Chetty, R. (2006b). A new method of estimating risk aversion. *American Economic Review 96*(5), 1821–1834.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy 116*(2), 173–234.

Chetty, R. (2009). Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods. *Annual Review of Economics 1*, 451–488.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica 80*(3), 969–1018.

Chetty, R. and A. Finkelstein (2013). Chapter 3 - social insurance: Connecting theory to data. Volume 5 of *Handbook of Public Economics*, pp. 111–193. Elsevier.

Chetty, R., J. N. Friedman, T. Olsen, and L. Pistaferri (2011). Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records. *Quarterly Journal of Economics 126*(2), 749.

Chetty, R. and A. Szeidl (2007). Consumption commitments and risk preferences. *Quarterly Journal of Economics*, 831–877.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica 56*(1), 63–90.

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy 100*(3), 437–467.

Chung, Y. (2013). Chronic health conditions and economic outcomes.

Cochrane, J. H. (1991). A simple test of consumption insurance. *Journal of political economy*, 957–976.

Coile, C. C. (2004). Health shocks and couples' labor supply decisions. Working Paper 10810, National Bureau of Economic Research.

Cullen, J. B. and J. Gruber (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics 18*(3), 546–572.

Cutler, D. M. and L. F. Katz (1992). Rising inequality? changes in the distribution of income and consumption in the 1980's. *American Economic Review 82*(2), 546–551.

De Giorgi, G., A. Frederiksen, and L. Pistaferri (2012). Consumption network effects. Technical report.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2016, May). The economic consequences of hospital admissions. Working Paper 22288, National Bureau of Economic Research.

Fadlon, I. and T. H. Nielsen (2015). Household responses to severe health shocks and the design of social insurance. Working Paper 21352, National Bureau of Economic Research.

Finkelstein, A., N. Hendren, and E. F. Luttmer (2015). The value of medicaid: Interpreting results from the oregon health insurance experiment. Working Paper 21308, National Bureau of Economic Research.

Grogger, J. (1995). The effect of arrests on the employment and earnings of young men. *Quarterly Journal of Economics 110*(1), 51–71.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review 87*(1), 192–205.

Gruber, J. and J. B. Cullen (1996). Spousal labor supply as insurance: Does unemployment insurance crowd out the added worker effect? Working Paper 5608, National Bureau of Economic Research.

Gupta, N. D., K. J. Kleinjans, and M. Larsen (2015). The effect of a severe health shock on work behavior: Evidence from different health care regimes. *Social Science & Medicine 136*, 44–51.

Halla, M. and M. Zweimüller (2013). The effect of health on earnings: Quasi-experimental evidence from commuting accidents. *Labour Economics 24*, 23–38.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 153–161.

Heckman, J. J. and T. E. Macurdy (1980). A life cycle model of female labour supply. *Review of Economic Studies 47*(1), 47–74.

Heckman, J. J. and T. E. Macurdy (1982). Corrigendum on a life cycle model of female labour supply. *Review of Economic Studies 49*(4), 659–660.

Hendren, N. (2013a). The policy elasticity. Working Paper 19177, National Bureau of Economic Research.

Hendren, N. (2013b). Private information and insurance rejections. *Econometrica 81*(5), 1713–1762.

Hendren, N. (2015). Knowledge of future job loss and implications for unemployment insurance. Unpublished manuscript.

Hilger, N. G. (2014). Parental credit constraints and college attendance: Evidence from timing of fathers' layoffs. Unpublished manuscript.

Huang, H., S. İmrohoroğlu, and T. Sargent (1997). Two computations to fund social security. *Macroeconomic Dynamics 1*(1), 7–44.

Hurd, M. D. and D. A. Wise (1989). The wealth and poverty of widows: Assets before and after the husband's death. In *The Economics of Aging*, pp. 177–200. University of Chicago Press.

Imbens, G. W. and J. M. Wooldridge (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature 47*(1), 5–86.

Immervoll, H., H. J. Kleven, C. T. Kreiner, and N. Verdelin (2011). Optimal tax and transfer programs for couples with extensive labor supply responses. *Journal of Public Economics 95*(11), 1485–1500.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (1995). A life cycle analysis of social security. *Economic theory 6*(1), 83–114.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (2003). Time-inconsistent preferences and social security. *Quarterly Journal of Economics 118*(2), 745–784.

Keane, M. P., P. E. Todd, and K. I. Wolpin (2011). Chapter 4 - the structural estimation of behavioral models: Discrete choice dynamic programming methods and applications. Volume 4, Part A of *Handbook of Labor Economics*, pp. 331–461. Elsevier.

Kleven, H. J., C. T. Kreiner, and E. Saez (2009). The optimal income taxation of couples. *Econometrica 77*(2), 537–560.

Koijen, R., S. Van Nieuwerburgh, and R. Vestman (2014). Judging the quality of survey data by comparison with" truth" as measured by administrative records: Evidence from sweden. In *Improving the Measurement of Consumer Expenditures*. University of Chicago Press.

Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2015). The optimal timing of unemployment benefits: Theory and evidence from sweden. Technical report, IZA Discussion Papers.

Kotlikoff, L. J., K. Smetters, and J. Walliser (1999). Privatizing social security in the united states – comparing the options. *Review of Economic Dynamics 2*(3), 532–574.

Kroft, K., F. Lange, and M. J. Notowidigdo (2013). Duration dependence and labor market conditions: Evidence from a field experiment. *Quarterly Journal of Economics 128*(3), 1123–1167.

Landais, C. (2015). Assessing the welfare effects of unemployment benefits using the regression kink design. *American Economic Journal: Economic Policy 7*(4), 243–278.

Low, H. and L. Pistaferri (2015). Disability insurance and the dynamics of the incentive-insurance tradeoff. *American Economic Review (forthcoming)*.

Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics 3*(1), 11–37.

Lundborg, P., M. Nilsson, and J. Vikström (2011). Socioeconomic heterogeneity in the effect of health shocks on earnings: Evidence from population-wide data on swedish workers. Unpublished manuscript.

Maloney, T. (1987). Employment constraints and the labor supply of married women: A reexamination of the added worker effect. *Journal of Human Resources 22*(1), 51–61.

Maloney, T. (1991). Unobserved variables and the elusive added worker effect. *Economica 58*(230), 173–187.

Meyer, B. D. and W. K. Mok (2013). Disability, earnings, income and consumption. Working Paper 18869, National Bureau of Economic Research.

Myers, D. A., R. V. Burkhauser, and K. C. Holden (1987). The transition from wife to widow: The importance of survivor benefits to widows. *Journal of Risk and Insurance*, 752–759.

Nelson, J. A. (1988). Household economies of scale in consumption: Theory and evidence. *Econometrica 56*(6), 1301–1314.

Nishiyama, S. and K. Smetters (2007). Does social security privatization produce efficiency gains? *Quarterly Journal of Economics 122*(4), 1677–1719.

OECD (2011). Divided we stand: Why inequality keeps rising. Technical report, OECD Publishing.

OECD (2014). Society at a glance 2014. Technical report, OECD Publishing.

Persson, P. and M. Rossin-Slater (2016). Family ruptures, stress, and the mental health of the next generation. *American Economic Review (forthcoming)*.

Pistaferri, L. (2015). Household consumption: Research questions, measurement issues, and data collection strategies. *Journal of Economic and Social Measurement 40*(1-4), 123–149.

Pohl, V., C. Neilson, and F. Parro (2013). The effect of health shocks on employment: Evidence from accidents in chile.

Ruhm, C. J. (1991). Are workers permanently scarred by job displacements? *American Economic Review 81*(1), 319–324.

Shimer, R. and I. Werning (2007). Reservation wages and unemployment insurance. *Quarterly Journal of Economics 122*(3), 1145–1185.

Social Security Administration (2015). Annual statistical supplement, 2014. Technical report, SSA.

Spletzer, J. R. (1997). Reexamining the added worker effect. *Economic Inquiry 35*(2), 417–427.

Stephens, M. J. (2001). The long-run consumption effects of earnings shocks. *Review of Economics and Statistics 83*(1), 28–36.

Stephens, M. J. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics 20*(3), 504–537.

White House (2015). Fiscal year 2016: Historical tables – budget of the us government.

Figure 1:
Illustration of the Empirical Research Design

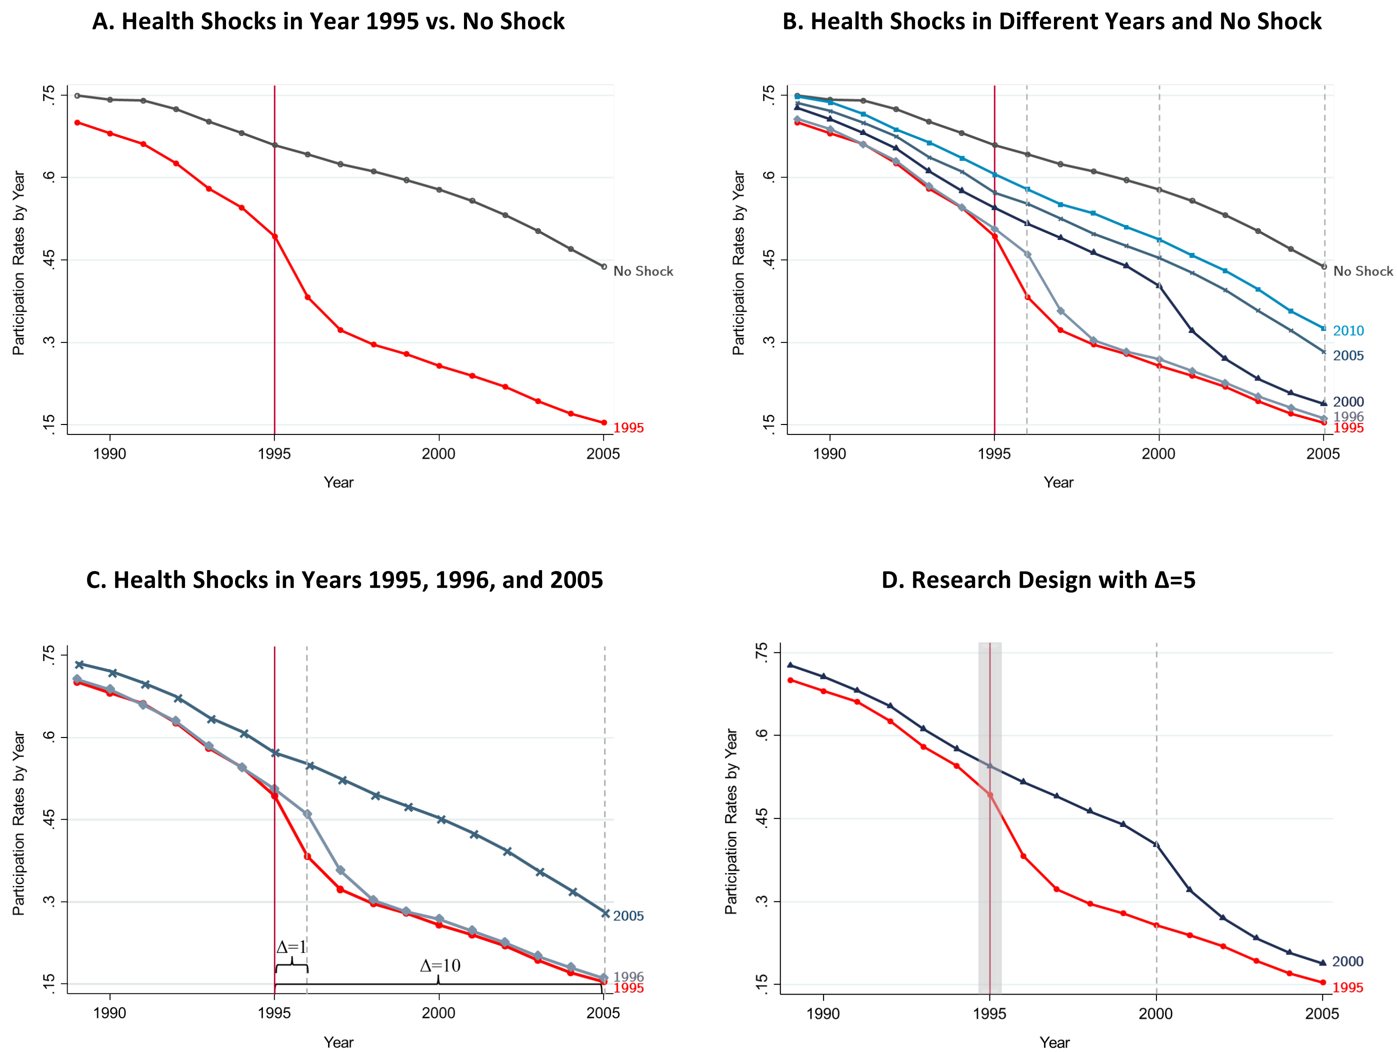

Notes: These figures compare the labor force participation of a treatment group of individuals who were born between 1930 and 1950 and experienced a heart attack or a stroke in 1995 to that of potential control groups. Panel A compares the treatment group to those who belong to the same cohorts but did not experience a shock in our data window, years between 1980 and 2011, and shows that the pre-1995 patterns of these groups are far from parallel. Panel B adds the behavior of households that experienced the same shock but in different years, and shows that the groups are becoming increasingly comparable to the treatment group – in terms of parallel trends before 1995 – the closer the year in which the individual experienced the shock was to the year the treatment group experienced the shock (1995). The figures suggest using households that experienced a shock in year 1995+Δ as a control group for households that experienced a shock in 1995. The trade-off in the choice of Δ is presented in Panel C. On the one hand, we would want to choose a smaller Δ such that the control group would be more closely comparable to the treatment group, e.g., year 1996 which corresponds to Δ=1. On the other hand, we would want to choose a larger Δ in order to be able to identify longer-run effects of the shock, up to period Δ-1. Using those that experienced a shock in 2005, which corresponds to Δ=10, will allow us to estimate up to the 9-year effect of the shock. However, this entails a potentially greater bias since the trend in the behavior of this group prior to 1995 is not as tightly parallel to that of the treatment group. Panel D displays the potential control group for this example when we choose Δ=5. Our research design generalizes this example by aggregating different calendar years.

Figure 2:
Survivors' Labor Supply Responses to the Death of Their Spouse

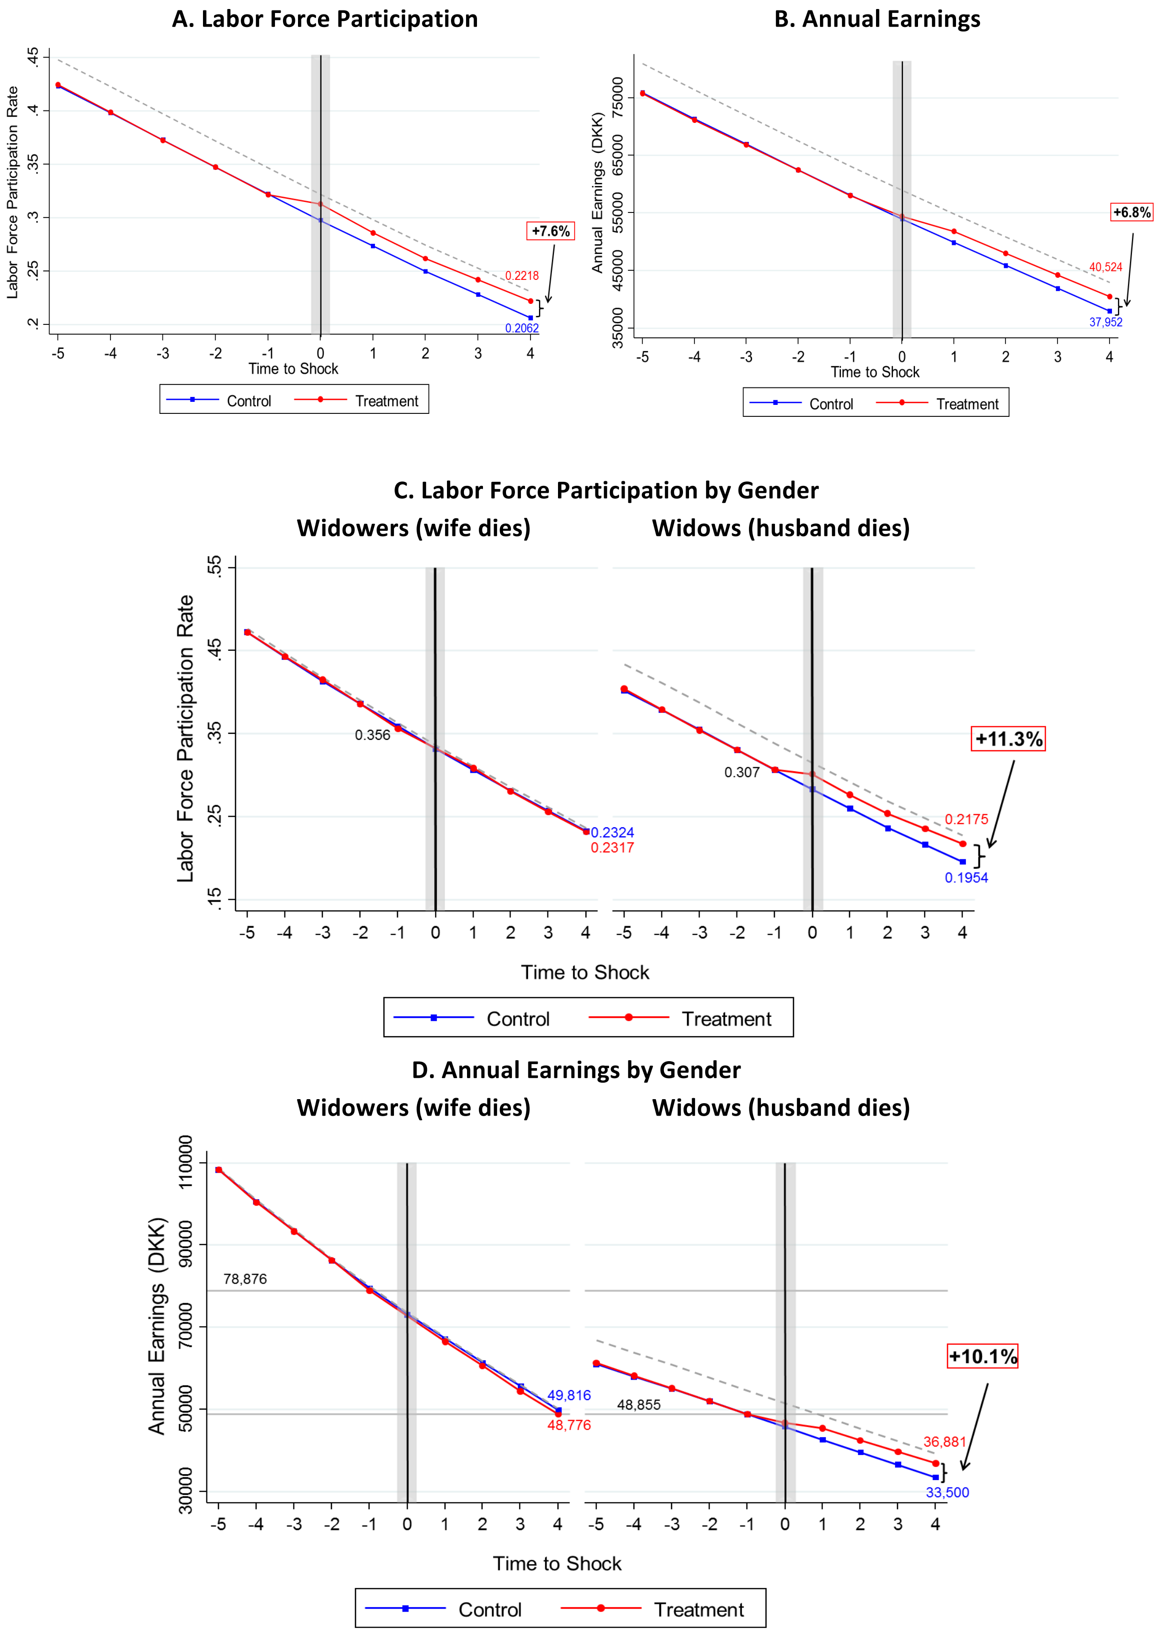

Notes: These figures plot the labor supply responses of survivors to the death of their spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011. Panels A and B depict the behavior of labor force participation and annual earnings, respectively, for the entire sample. Panels C and D break down these average responses by the gender of the surviving spouse. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

Figure 3:
Labor Supply Responses of Survivors under Age 60 to the Death of Their Spouse by Gender

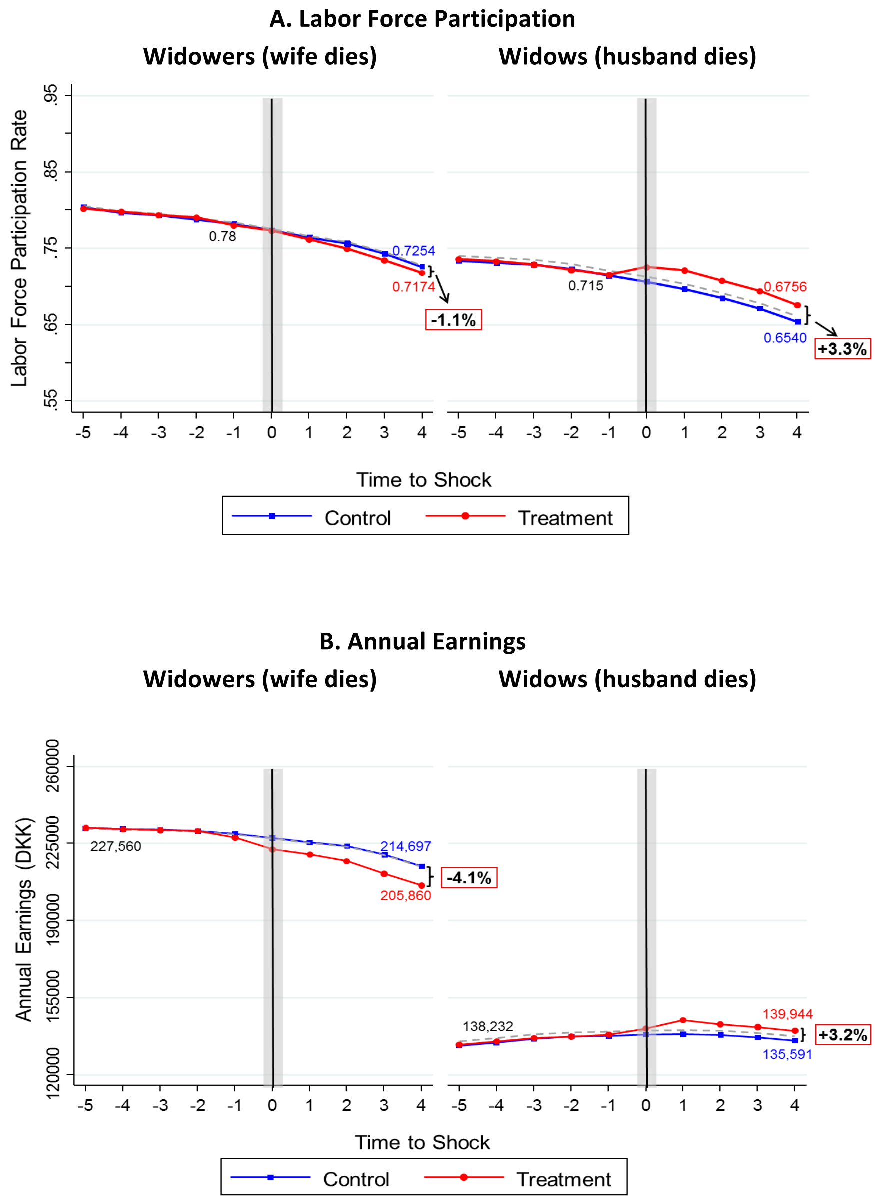

Notes: These figures plot the labor supply responses of survivors under age 60 to the death of their spouse by the gender of the surviving spouse. The sample includes individuals under 60 whose spouses died between ages 45 and 80 from 1980 to 2011. Panel A depicts the behavior of labor force participation, and Panel B depicts the behavior of annual earnings. The lines in each graph are constructed as described in the notes of Figure 2.

Figure 4:
Labor Supply Responses of Survivors under Age 60 to the Death of Their Spouse:
Full-Time/Part-Time Employment

A. Labor Force Participation

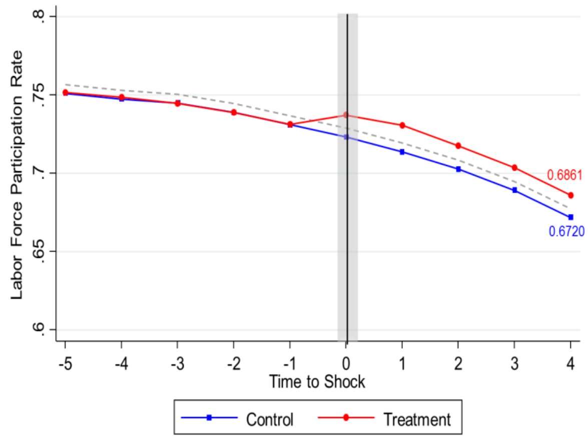

B. Full-Time Employment

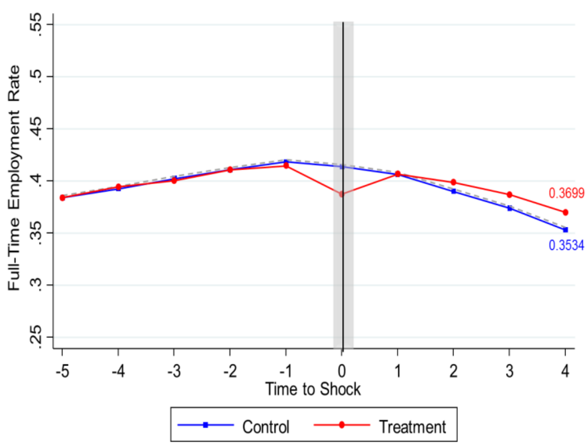

C. Part-Time Employment

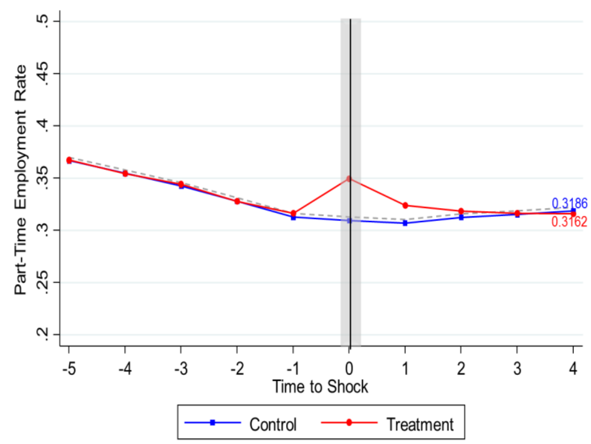

Notes: These figures plot labor supply responses of survivors under age 60 to the death of their spouse. The sample includes individuals under 60 whose spouses died between ages 45 and 80 from 1980 to 2011. Panel A depicts labor force participation; Panels B and C depict the fraction of surviving spouses who are employed full time and part time, respectively. The pictures are constructed from ATP data available for workers under 60. Full-time employment is defined as working at least 30 hours per week all 12 months of the calendar year ("full-time full-year"); part-time employment is defined as working at some point during the year, but either fewer than 30 hours per week or fewer than 12 months within the calendar year. The lines in each graph are constructed as described in the notes of Figure 2.

Figure 5:
Survivors' Annual Earnings Responses to the Death of Their Spouse
by the Level of their Own Pre-Shock Earnings

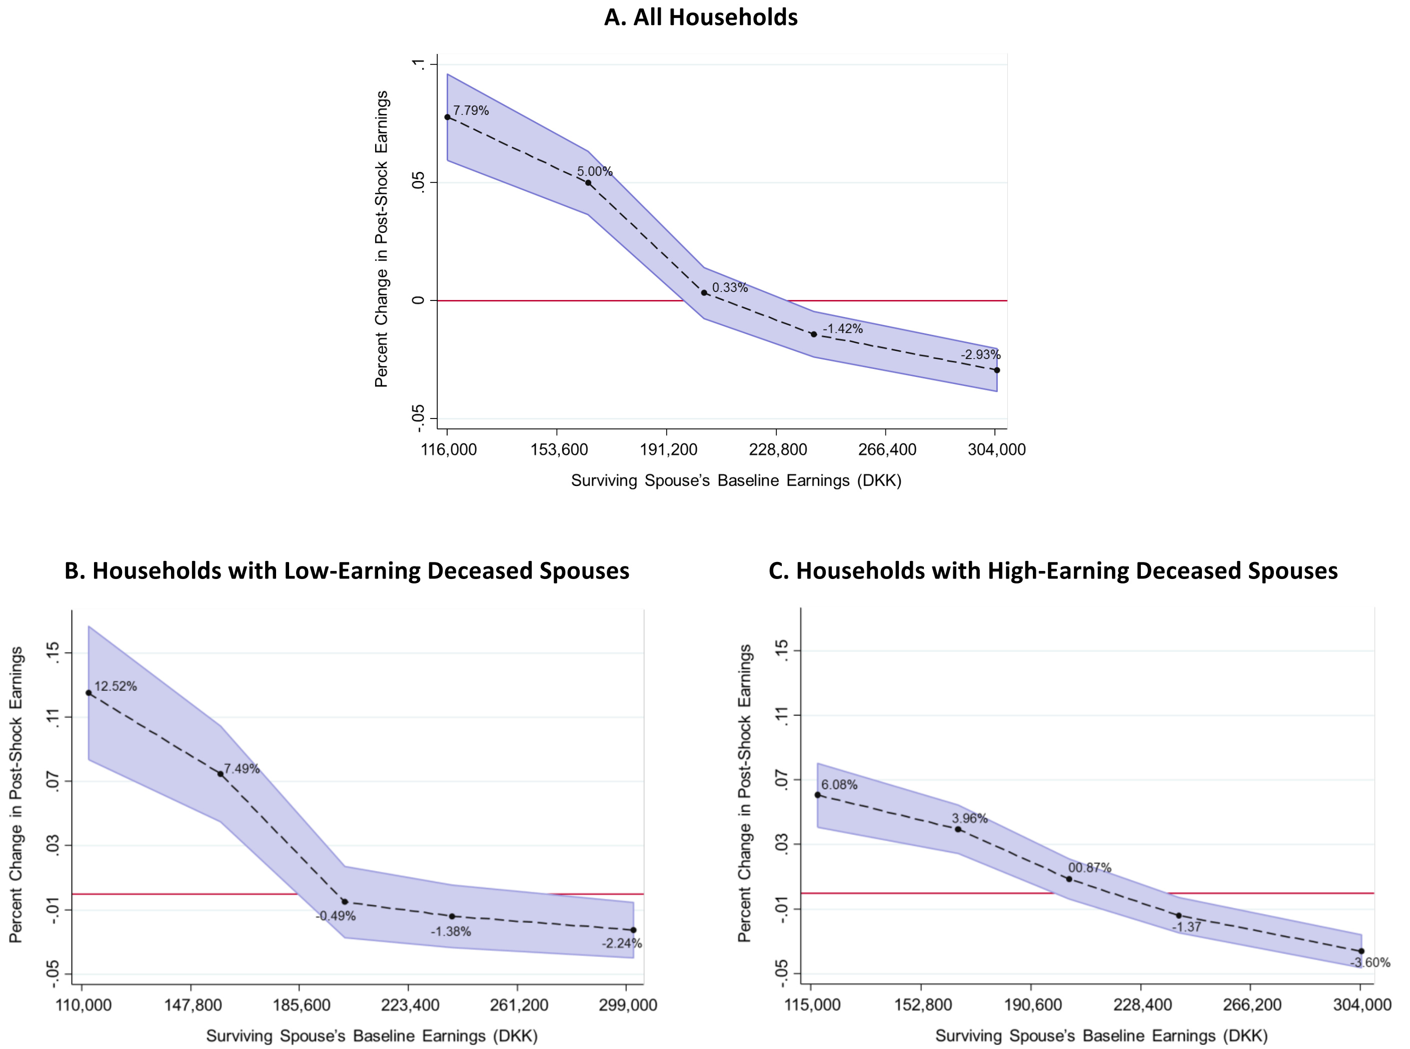

Notes: These figures include individuals whose spouses died between ages 45 and 80 from 1980 to 2011, where we constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than their experimental-group-specific 20th percentile. Then, we calculate for each household the pre-shock labor income share of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force; specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households for which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household. In addition, to guarantee that our results are not driven by outliers, we exclude households with dying spouses whose mean pre-shock earnings did not fall within their group-specific 5th and 95th percentiles as well as households with unaffected spouses whose mean pre-shock earnings were higher than those of their group-specific 95th percentile. We divide the remaining sample into five equal-sized groups by their pre-shock level of earnings and plot the average labor income response as well as its 95-percent confidence interval (where standard errors are calculated using the Delta method) against the pre-shock mean earnings for each group. Panel A includes all households; Panel B includes households in which the dying spouse's pre-shock labor income fell within the bottom three quintiles of its group-specific distribution, to which we refer as "low-earners"; Panel C includes households in which the dying spouse's pre-shock labor income fell within the top two quintiles, to which we refer as "high-earners". The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4.

Figure 6:
Household Labor Supply Responses to Non-Fatal Severe Health Shocks

A. Affected Spouse

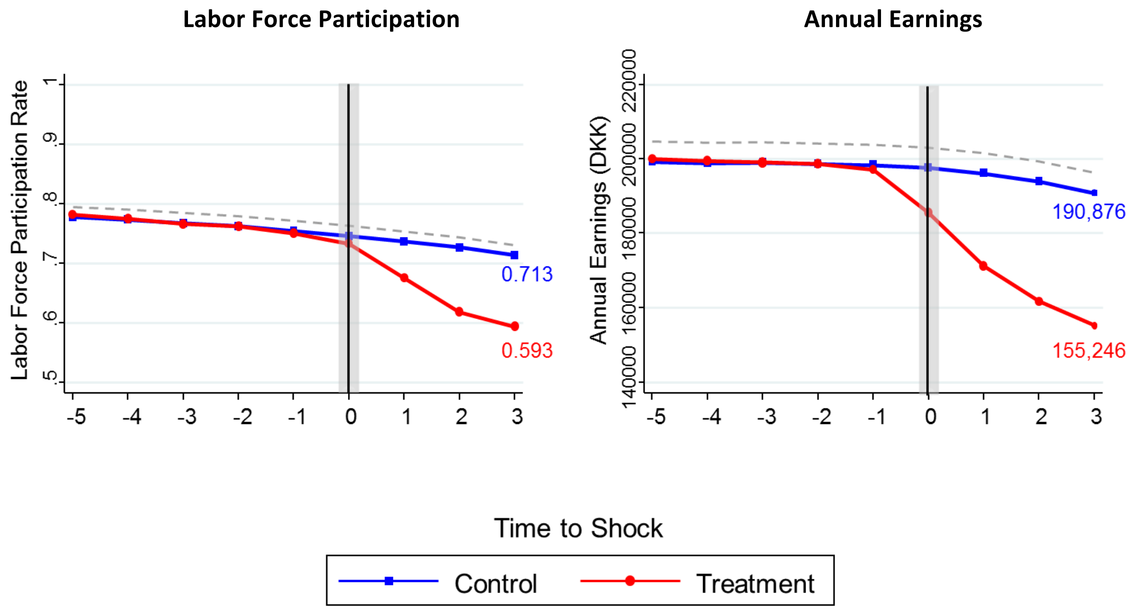

B. Unaffected Spouse

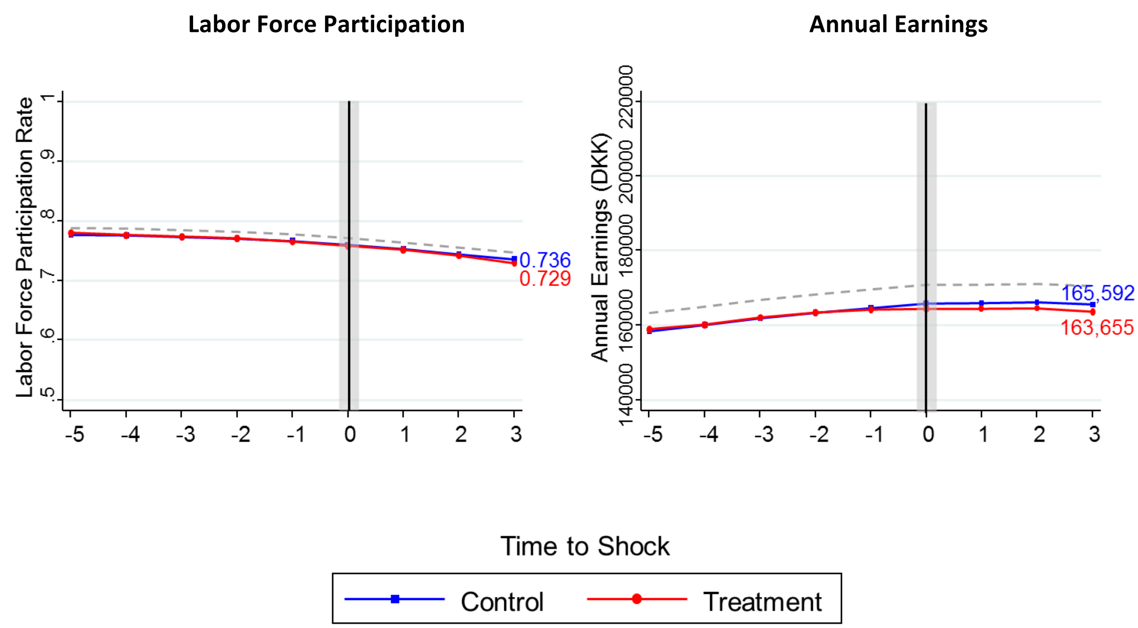

Notes: These figures plot the labor supply responses of households in which an individual experienced a heart attack or a stroke between 1980 and 2011 and survived for at least three years. The sample includes households in which both spouses were under age 60. Panel A depicts the labor force participation and annual earnings of the individual that experienced the shock. Panel B depicts the labor force participation and annual earnings of the unaffected spouse. The lines in each graph are constructed as described in the notes of Figure 2.

Table 1:
Survivors' Labor Force Participation Responses to the Death of Their Spouse by the Degree of Income Loss

A. Surviving Spouses of All Ages			
1. Baseline Regression	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post	0.1265***	0.1220***	0.1170***
	(0.0023)	(0.0042)	(0.0027)
Treat × Post ×	-0.1889***	-0.1894***	-0.1744***
Replacement Rate	(0.0035)	(0.0061)	(0.0044)
Number of Obs.	4,288,621	1,387,615	2,901,006
Number of Households	714,892	231,318	483,574
2. Regression with Interactions	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post ×	-0.1989***	-0.2021***	-0.1927***
Replacement Rate	(.0045)	(.0081)	(.0056)
Number of Obs.	2,741,690	821,742	1,919,948
Number of Households	459,622	137,724	321,898
Regression 1 for Sub-Sample of Regression 2			
Treat × Post ×	-0.1922***	-0.1918***	-0.1832***
Replacement Rate	(.0043)	(.0077)	(.0054)
B. Surviving Spouses under 60			
1. Baseline Regression	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post	0.0883***	0.0652***	0.0954***
	(0.0054)	(0.0125)	(0.0063)
Treat × Post ×	-0.1270***	-0.1081***	-0.1338***
Replacement Rate	(0.0083)	(0.0168)	(0.0101)
Number of Obs.	803,158	201,487	601,671
Number of Households	134,199	33,720	100,479
2. Regression with Interactions	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post ×	-0.1481***	-0.1375***	-0.1499***
Replacement Rate	(0.0091)	(0.0186)	(.0110)
Number of Obs.	704,370	173,620	530,750
Number of Households	118,812	29,288	89,524
Regression 1 for Sub-Sample of Regression 2			
Treat × Post ×	-0.1377***	-0.1236***	-0.1430***
Replacement Rate	(.0088)	(.0184)	(.0107)

Notes: This table reports the interaction of the treatment effect of the death of a spouse with the household's post-shock income replacement rate (equation (3)). The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011.The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Panel A reports estimates for the sample of all survivors by gender; Panel B reports estimates for the sample of survivors under age 60 by gender. In each panel, we report estimates of two specifications. Specification 1 in each panel estimates a baseline differences-in-differences specification which interacts the treatment effect with the replacement rate variable. This replacement rate is calculated as follows. First, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their pre-shock levels (in period -1). Then, we calculate the ratio of this adjusted household income in period 1 (post-shock) to that in period -1 (pre-shock), and normalize it by the average ratio for the control group in order to account for calendar year trends as well as for life-cycle effects. Specification 2 in each panel extends specification 1 to include interactions of the treatment effect with additional household characteristics: age dummies for the surviving spouse, dummies for the age of the deceased at the year of death, year dummies, indicators for the number of children in the household as well as the surviving spouse's months of education (and its square). The results are also robust to the inclusion of a quadratic in the household's net wealth. All the variables that are interacted with "Treat × Post" are interacted with "Treat" and "Post" and enter the regressions separately as well. Since there are households with missing values for some of the controls (that are therefore included in the estimation of specification 1 but not 2), we show the robustness of our estimate of interest ("Treat × Post × Replacement Rate") to the inclusion of this set of controls by reporting estimates for specification 1 for the sub-sample of households that are included in the estimation of specification 2. All specifications include year, age, and household fixed effects. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 2:
Widows' Labor Force Participation Responses to the Death of Their Spouse by Social Survivors Benefits

Dependent Variable:	Widows' Participation
Treat × Post × Survivors Benefits	-.0057*** (.0020)
Average Treatment Effect	1.8 pp
Counterfactual Participation	48.7 pp
Number of Obs.	364,100
Number of Clusters	268

Notes: This table reports the interaction of the treatment effect of the death of a spouse with the actual survivors benefits widows received through the Social Disability Insurance (Social DI) program (equation (4)). The regression is estimated by two-stage least squares, where the instrument for actual benefits is constructed as follows. In each year we calculate for each municipality the average benefits received by non-working surviving spouses through Social DI. Then, we assign to each widow in the treatment group her respective municipality-year leave-one-out mean. The sample includes widows under age 67 (the age at which the program transitions into the Old-Age Pension) in years prior to 1994 (when there is a data break in the reporting method of survivors benefits received through Social DI). The controls included in the estimation are municipality unemployment rate and average earnings (and their interaction with "Treat", "Post", and "Treat × Post") as well as age, year, and municipality fixed effects. The identifying assumption is that, given our set of controls, the average social survivors benefits transferred to widows in a municipality in a given year affects a widow's participation only through its influence on her own survivors benefits receipts. Note that the source of variation we use is within municipalities over time since we include municipality and calendar year fixed effects as controls. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the municipality level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 3:
Labor Force Participation Responses of Widows Who Did Not Work before the Shock

	Mean Spousal Labor Force Participation	Spousal Participation by the Deceased's Employment History	Overall Household Income by the Deceased's Employment History	Deceased Did Not Work	
				Deceased's Income Less than 10th Percentile	Deceased's Income Less than 5th Percentile
	(1)	(2)	(3)	(4)	(5)
Treat × Post	0.0132*** (0.0005)	0.0078*** (0.0005)	-72,326*** (841)	0.0018 (0.0012)	0.0021 (0.0018)
Treat × Post × Deceased Worked		0.0461*** (0.0027)	-59,208*** (6,438)		
Number of Obs.	1,320,908	1,320,908	1,320,908	114,851	57,381
Number of Households	220,270	220,270	220,270	19,160	9,577
Number of Treated Households with Non-Working Deceased	90,686	90,686	90,686		
Number of Treated Households with Working Deceased	11,257	11,257	11,257		

Notes: This table reports the differences-in-differences estimates of the labor force participation responses of widows who did not work during the five-year period preceding their spouse's death. The sample includes households in which the husband died between ages 45 and 80 from 1980 to 2011 and in which he either worked throughout the entire five-year period preceding his death (periods -5 to -1) or did not work altogether during this period. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Column 1 reports the simple differences-in-differences estimate in a regression in which the outcome variable is spousal labor force participation. Column 2 adds an interaction of the treatment effect with an indicator for whether the husband worked before his death. Column 3 runs the same specification as in Column 2 but where the outcome variable is the household's overall income. Columns 4 and 5 report the spousal labor force participation effect for sub-samples of the households in which the husband did not work before his death. Column 4 reports the treatment effect for households in which the non-working deceased's overall income before the shock (periods -3 to -1), including any transfer from government programs, was lower than the 10^{th} percentile of this sample's income distribution; Column 5 reports the treatment effect for households in which the non-working deceased's overall income before the shock was lower than the 5^{th} percentile. All specifications include year, age, and household fixed effects. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Figure 1:
Life-Cycle Labor Force Participation of the Unaffected Spouses in the Death Event Sample

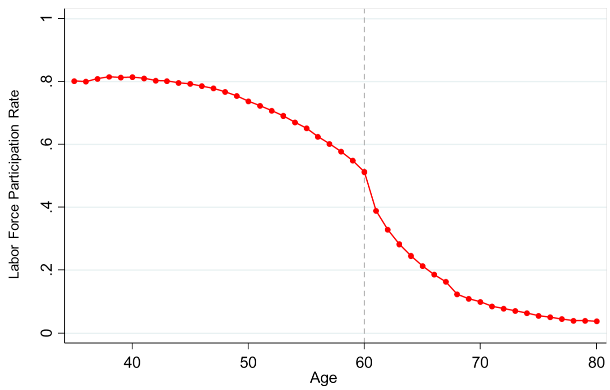

Notes: This figure displays the life-cycle labor force participation of the unaffected spouses that are included in the death event sample (i.e., individuals whose spouses died between ages 45 and 80 from 1985 to 2011). The observations include the pre-shock periods (specifically, periods -5 to -2). The sharp drop at age 60 corresponds to eligibility for the Voluntary Early Retirement Pension (VERP). The figure shows the complex life-cycle trends in labor supply and illustrates why an extrapolation based on behavior in previous years is a poor predictor of future behavior.

Appendix Figure 2:
Survivors' Labor Supply Responses to the Death of Their Spouse: Fatal Heart Attacks and Strokes

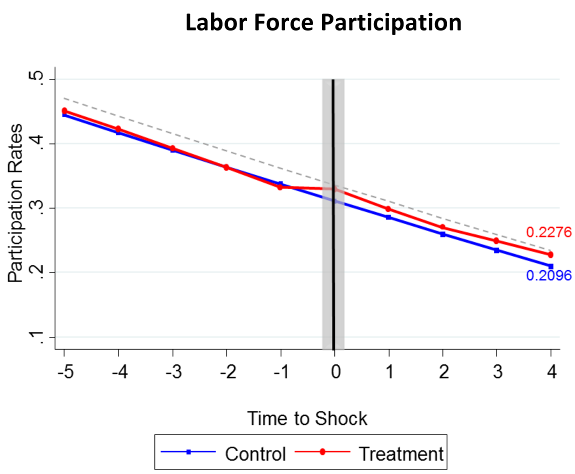

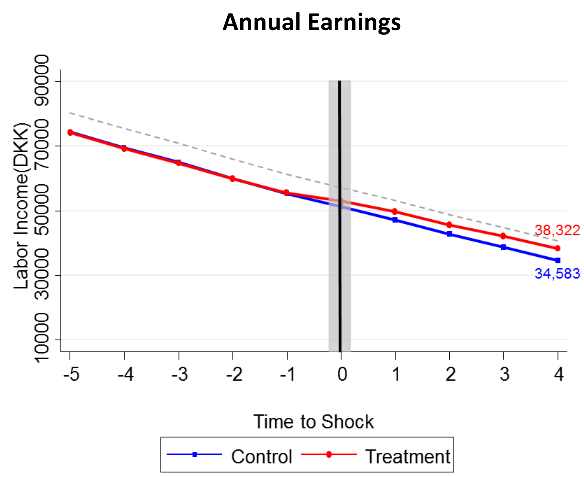

These figures plot the labor supply responses of survivors to the death of their spouse. The sample includes households in which an individual experienced a heart attack or a stroke between 1980 and 2011 and died in the same year. The panel on the left depicts the behavior of labor force participation, and the panel on the right depicts the behavior of annual earnings. The lines in each graph are constructed as described in the notes of Figure 2.

Appendix Figure 3:
Household Income around the Death of a Spouse

A. Potential Household Income

Widowers (wife dies) Widows (husband dies)

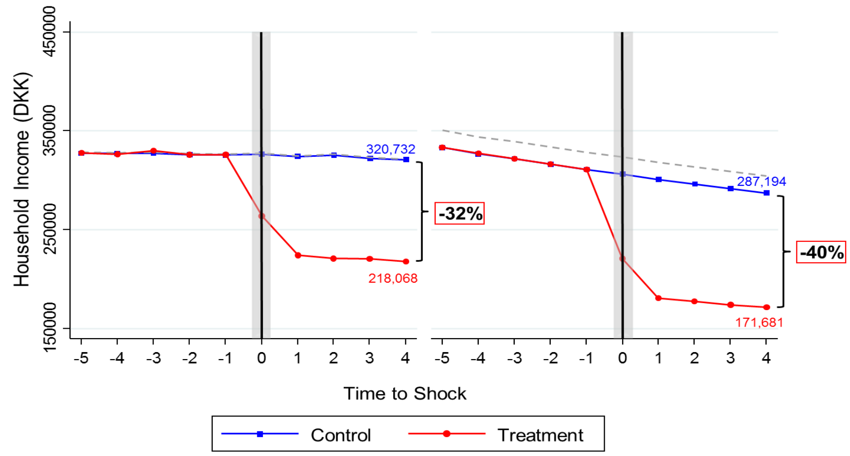

B. Actual Household Income

Widowers (wife dies) Widows (husband dies)

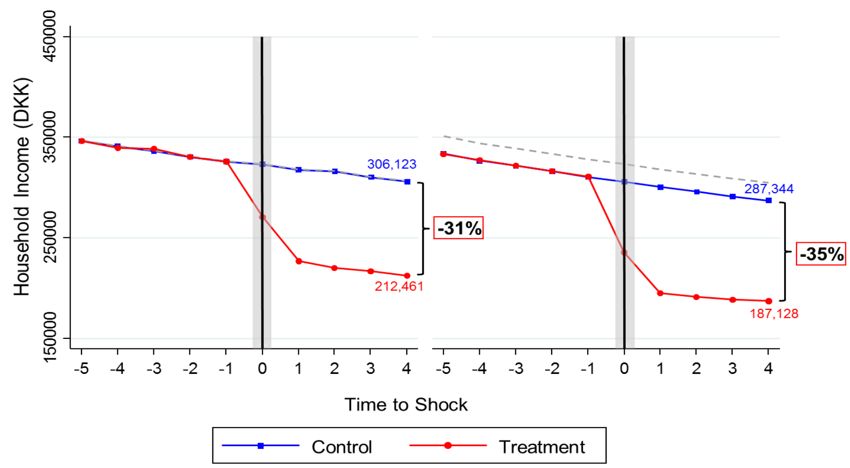

C. Potential Household Income of Survivors under Age 60

Widowers (wife dies) Widows (husband dies)

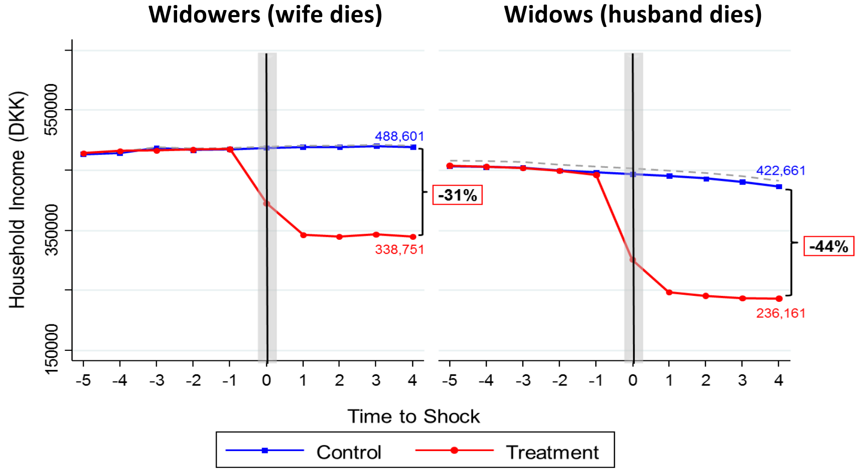

Notes: These figures plot different measures of household-level income around the death of a spouse by the gender of the surviving spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011. Panel A plots an adjusted measure of household income. Specifically, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their pre-shock levels (in period -1). Hence, this measure captures the income loss that is directly attributed to the loss of a spouse. Panel B plots the actual household income that is observed in the data, which takes into account the surviving spouse's behavioral responses. Panel C plots the adjusted measure from Panel A for survivors under age 60. The lines in each graph are constructed as described in the notes of Figure 2.

Appendix Figure 4:
Social Survivors Benefits for Widows under Age 67

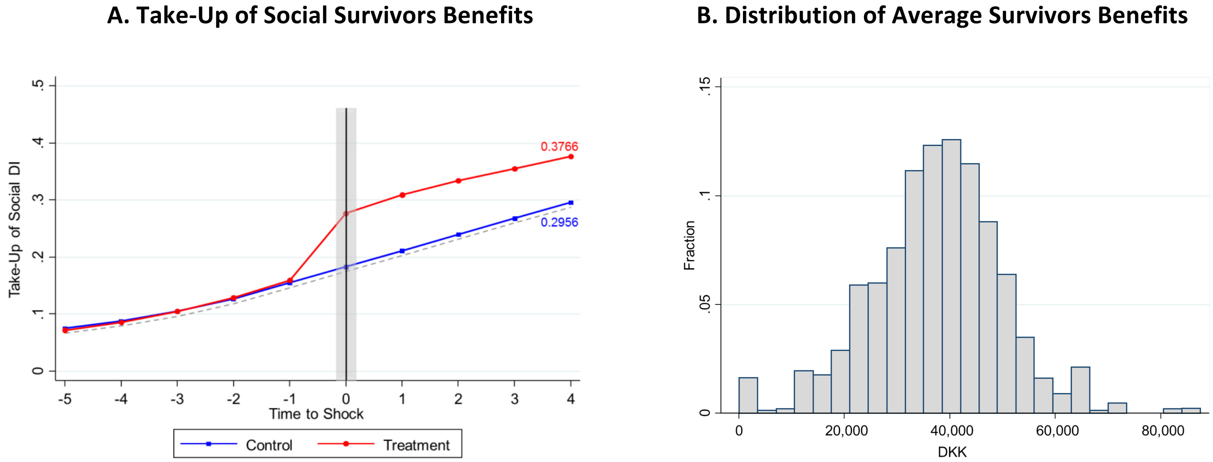

Notes: These figures include widows younger than 67 (the age at which the Social Disability Insurance transitions into the Old-Age Pension) in years prior to 1994 (when there is a data break in the reporting method of benefits received through Social Disability Insurance). Panel A plots these widows' take-up of survivors benefits through the Social Disability Insurance program around the death of their spouse. Panel B displays the distribution of the instrument that we use in estimation of equation (4). These are the year-by-municipality average benefits received by non-working survivors through Social DI.

Appendix Figure 5:
Widows' Labor Force Participation Responses to the Death of Their Spouse by Age

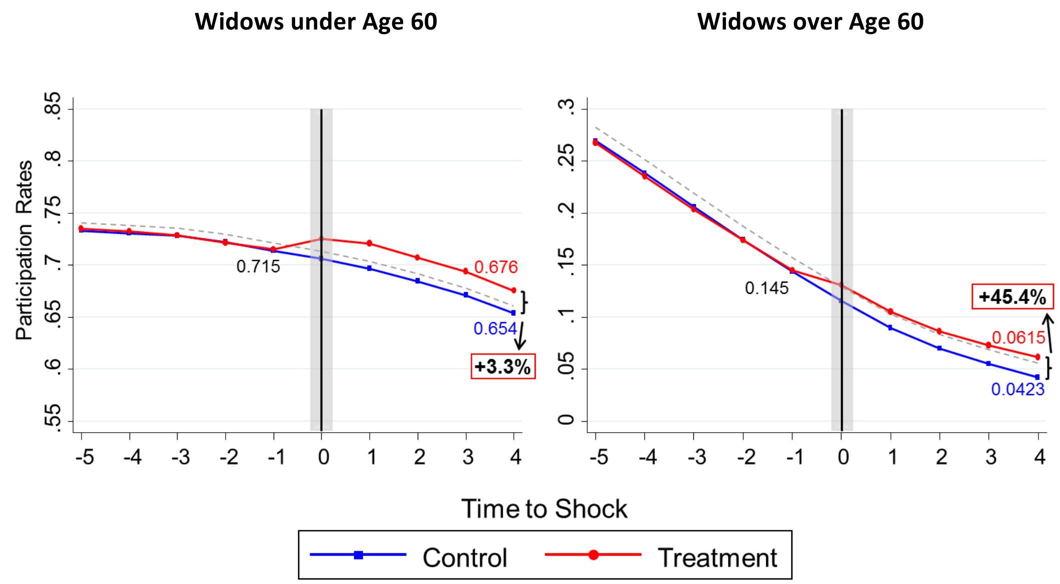

Notes: These figures plot the labor force participation responses of widows to the death of their spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011. The panel on the left depicts the behavior of widows younger than 60, the age at which there is a sharp decline in labor force participation due to eligibility for early retirement benefits, and the panel on the right depicts the behavior of widows older than 60. The lines in each graph are constructed as described in the notes of Figure 2.

Appendix Table 1: Summary Statistics of Analysis Sample

		Death Event Sample				Health Shock Sample	
		All Ages (1)		Under 60 (2)		Under 60 (3)	
		Treatment	Control	Treatment	Control	Treatment	Control
	Characteristics						
	Year of Observation	1993.13	1993.09	1992.74	1992.75	1991.83	1991.95
Unaffected Spouse	Age	62.86	62.27	47.60	47.48	45.69	45.30
	Education (months)	118.66	119.94	129.19	129.38	130.94	132.48
	Percent female	0.6937	0.6632	0.7485	0.7485	0.7551	0.7367
Affected Spouse	Age	64.84	64.01	52.51	52.14	47.80	47.27
	Education (months)	123.57	124.05	131.80	132.22	134.90	136.31
	Outcomes						
Unaffected Spouse	Participation	0.3474	0.3719	0.7389	0.7445	0.7709	0.7820
	Earnings (DKK)	62,455	67,452	160,799	162,094	163,336	168,311
Affected Spouse	Participation	0.2723	0.3211	0.6033	0.6560	0.7621	0.7790
	Earnings (DKK)	51,579	61,791	143,118	158,447	198,723	204,191
Number of Households		310,720	409,190	55,103	80,578	37,432	54,926

Notes: This table presents means of key variables in our analysis sample. All monetary values are reported in nominal Danish Kroner (DKK) deflated to 2000 prices using the consumer price index. In this year the exchange rate was approximately DKK 8 per US $1. For each event, the treatment group comprises households that experienced a shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Columns 1 and 2 report statistics for the death event sample of households in which a spouse died of any cause between ages 45 and 80 from 1980 to 2011. Column 1 reports statistics for the entire sample, and Column 2 reports statistics for the sub-sample of surviving spouses under age 60. Column 3 reports statistics for the health event sample. It includes households in which one spouse experienced a heart attack or a stroke between 1980 and 2011 and survived for at least three years, and in which both spouses were under age 60. The values reported in the table are based on data from two periods before the shock occurred (period t = -2).

Appendix Table 2: Survivors' Labor Supply Responses to the Death of Their Spouse

A. Surviving Spouses of All Ages

	Widowers				Widows			
Dependent variable:	Participation (1)	Participation (2)	Earnings (3)	Earnings (4)	Participation (5)	Participation (6)	Earnings (7)	Earnings (8)
Treat × Post	-.0016 (.0017)	-.0017 (.0016)	-939* (485)	-906** (448)	.0188*** (.0011)	.0164*** (.0010)	2,957*** (201)	2,707*** (188)
Household FE	X	X	X	X	X	X	X	X
Year and Age FE		X		X		X		X
Number of Obs.	1,397,030	1,397,030	1,397,030	1,397,030	2,919,946	2,919,946	2,919,946	2,919,946
Number of Households	232,973	232,973	232,973	232,973	486,890	486,890	486,890	486,890

B. Surviving Spouses under 60

	Widowers				Widows			
Dependent variable:	Participation (1)	Participation (2)	Earnings (3)	Earnings (4)	Participation (5)	Participation (6)	Earnings (7)	Earnings (8)
Treat × Post	-.0075** (.0036)	-.0071** (.0036)	-7,902*** (1444)	-7,730*** (1439)	.0207*** (.0023)	.0219*** (.0023)	4,093*** (522)	4,423*** (516)
Household FE	X	X	X	X	X	X	X	X
Year and Age FE		X		X		X		X
Number of Obs.	203,569	203,569	204,438	204,438	607,437	607,437	608,742	608,742
Number of Households	34,104	34,104	34,118	34,118	101,529	101,529	101,562	101,562

Notes: This table reports the differences-in-differences estimates of the surviving spouses' labor supply responses (equation (2)). The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Panel A reports the responses of all survivors by gender, where widowers are those who lost their wives and widows are those who lost their husbands. Panel B reports the responses of survivors under 60 by gender. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 3:
Survivors' Annual Earnings Responses to the Death of Their Spouse

A. Mean Responses by Quintiles of Own Pre-Shock Earnings

		All Survivors (1)		Low-Earning Deceased (2)		High-Earning Deceased (3)	
Quintile 1	Treat × Post	6,062***	8,847***	7,237***	9,034***	5,105***	8,565***
		(1,211)	(978)	(2,194)	(1,784)	(1,481)	(1,199)
	Mean Earnings	75,092		58,025		84,202	
	Percent Change	**8.07%**	**11.78%**	**12.47%**	**15.57%**	**6.06%**	**10.17%**
Quintile 2	Treat × Post	5,946***	7,283***	7,012***	7,120***	4,919***	6,860***
		(1,348)	(1,070)	(2,530)	(2,014)	(1,641)	(1,313)
	Mean Earnings	115,830		92,992		123,835	
	Percent Change	**5.13%**	**6.26%**	**7.54%**	**7.66%**	**3.97%**	**5.54%**
Quintile 3	Treat × Post	1,154	3,744***	-667	2,341	1,370	3,919***
		(1,369)	(1,049)	(2,505)	(1,893)	(1,674)	(1,305)
	Mean Earnings	148,700		128,151		156,070	
	Percent Change	**0.78%**	**2.52%**	**-0.52%**	**1.83%**	**0.88%**	**2.51%**
Quintile 4	Treat × Post	-2,203	-934	-2,224	-986	-2,644	-1,484
		(1,495)	(1,157)	(2,746)	(2,095)	(1,818)	(1,416)
	Mean Earnings	185,311		162,883		192568	
	Percent Change	**-1.19%**	**-0.50%**	**-1.37%**	**-0.60%**	**-1.37%**	**-0.77%**
Quintile 5	Treat × Post	-7,494***	-5,846***	-4,872	-3,703	-8,877***	-7,466***
		(1,765)	(1,399)	(3,211)	(2,498)	(2,170)	(1,718)
	Mean Earnings	239,994		217,992		246,641	
	Percent Change	**-3.12%**	**-2.45%**	**-2.23%**	**-1.7%**	**-3.60%**	**-3.03%**
Household FE		X	X	X	X	X	X
Age and Year FE			X		X		X

B. Mean Responses by Gender

	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post	585	-6,623***	3,405***
	(667)	(1,342)	(729)
Counterfactual Earnings	150,994	163,010	145,969
Household FE	X	X	X
Number of Obs.	686,521	220,125	466,392
Number of Households	114,462	36,705	77,756

Notes: This table reports the differences-in-differences estimates of the surviving spouses' annual earnings by the level of their own earnings when their spouses died. The sample includes individuals whose spouses died between ages 45 and 80 from 1980 to 2011, where we constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than their experimental-group-specific 20th percentile. Then, we calculate for each household the pre-shock labor income share of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force; specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households for which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household. In addition, to guarantee that our results are not driven by outliers, we exclude households with dying spouses whose mean pre-shock earnings did not fall within their group-specific 5th and 95th percentiles or households with unaffected spouses whose mean pre-shock earnings were higher than their group-specific 95th percentile. We divide the remaining sample into five equal-sized groups by their pre-shock level of earnings. Panel A separately estimates a differences-in-differences specification for each surviving spouses' quintile. Column 1 includes all surviving spouses; Column 2 includes households in which the dying spouses' pre-shock labor income fell within the bottom three quintiles of its group-specific distribution, to which we refer as "low-earners"; Column 3 includes households in which the dying spouses' pre-shock labor income fell within the top two quintiles, to which we refer as "high-earners". The gradient of survivors' labor supply responses with respect to their own level of pre-shock earnings is also robust to the inclusion of a quadratic in the household's net wealth. Panel B reports the average treatment effect for this sample. The second row reports the counterfactual outcome based on the differences-in-differences estimation. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 4:
Household Responses to Non-Fatal Severe Health Shocks

	Affected Spouse				Household Income		Unaffected Spouse			
Dependent variable:	Participation		Earnings				Participation		Earnings	
	Short Run	Medium Run	Short Run	Medium Run	Short Run	Medium Run	Short Run	Medium Run	Short Run	Medium Run
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treat × Post	-.0861***	-.1212***	-29,012***	-36,015***	-12,114***	-18,665***	-.0018	-.0071***	-1,712***	-2,041***
	(.0023)	(.0027)	(741)	(879)	(2168)	(2380)	(.0020)	(.0024)	(538)	(628)
Household FE	X	X	X	X	X	X	X	X	X	X
Counterfactual Post-Shock Mean of Dependent Var.	.7328	.7147	195,433	191,225	503,460	503,318	.7489	.7366	166,216	165,756
Percent Change	-12%	-17%	-15%	-19%	-2.4%	-3.7%	0	-1%	-1.03%	-1.23%
Percent Change Excluding the Unaffected Spouse's Responses					-2.1%	-3.3%				
Number of Obs.	644,699		646,272		645,817		644,359		645,817	
Number of Households	92,349		92,358		92,356		92,324		92,356	

Notes: This table reports the differences-in-differences estimates of household labor supply responses to severe health shocks in which the affected spouse survived and the effect of these shocks on overall household income (equation (5) in footnote 20). The sample includes households in which one spouse experienced a heart attack or a stroke and survived for at least three years, and in which both spouses were under age 60. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). We allow for differential treatment effects for the "short run" – periods 1 and 2 – and the "medium run" – period 3, to account for the gradual responses documented in Figure VII. The pre-shock periods include periods -5 to -2. Household income (Columns 5 and 6) includes income from any source – including earnings, capital income, annuity payouts, and benefits from any social program. The third row reports the counterfactual outcome based on the differences-in-differences estimation. Robust standard errors clustered at the household level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 5:
Spousal Participation Responses to Non-Fatal Severe Health Shocks by the Degree of Income Loss

A. Baseline Regression	Participation (1)	Participation (2)
Treat × Post	0.3423***	0.3327***
	(0.0147)	(0.0147)
Treat × Post × Replacement Rate	-0.3529***	-0.3692
	(0.0149)	(0.0152)
Treat × Post × Affected Spouse Works		0.0420***
		(0.0048)
B. Regression with Interactions	Participation (1)	Participation (2)
Treat × Post × Replacement Rate	-0.3603***	-0.3720***
	(0.0151)	(0.0153)
Treat × Post × Affected Spouse Works		0.0373***
		(0.0050)
Number of Obs.	236,897	236,897
Number of Households	47,459	47,459

Notes: This table reports the interaction of the treatment effect of non-fatal spousal health shocks with the household's post-shock income replacement rate (similar to equation (3)). The sample includes households in which one spouse experienced a heart attack or a stroke and survived for at least three years, and in which both spouses were under age 60. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Column (1) in Panel A estimates a baseline differences-in-differences specification which interacts the treatment effect with the replacement rate variable. This replacement rate is calculated as follows. First, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their pre-shock levels (in period -1). Then, we calculate the ratio of this adjusted household income in period 1 (post-shock) to that in period -1 (pre-shock), and normalize it by the average ratio for the control group in order to account for calendar year trends as well as for life-cycle effects. Column (1) in Panel B extends this specification to include interactions of the treatment effect with additional household characteristics: age dummies for both spouses, year dummies, indicators for the number of children in the household, the healthy spouse's months of education (and its square), and a quadratic in the household's net wealth. All the variables that are interacted with "Treat × Post" are interacted with "Treat" and "Post" and enter the regressions separately as well. Column (2) in Panels A and B replicates column (1) in Panels A and B, respectively, but adds interactions with an indicator for the sick spouse's labor force participation. All specifications include year, age, and household fixed effects. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 and 3. Robust standard errors clustered at the household level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Online Appendices

In this online appendix, we provide extensions to the simple static model that we analyze in the paper. In Online Appendix A, we analyze a fully-dynamic life-cycle model that allows for endogenous savings, as well as private and informal insurance arrangements, and can incorporate a general class of arbitrary choice variables (such as time investment in home production). In Online Appendix B, we introduce state dependence in preferences and allow for flexible consumption-leisure complementarities. In Online Appendix C, we analyze an intensive-margin version of our model. We use Online Appendix C to also show that the results extend to models of the household other than the unitary framework by using the collective approach to household behavior (Chiappori 1988, 1992; Apps and Rees 1988).

Online Appendix A: A Dynamic Model of Household Labor Force Participation

In the context of social insurance over the life-cycle, it is important to consider households' self-insurance through ex-ante mechanisms such as precautionary savings. In this appendix, we analyze life-cycle participation decisions using a dynamic search model that allows for endogenous savings.[1] The general result of this analysis is that our welfare formula extends to the dynamic case with the adjustment that post-shock responses in the static case are replaced by mean responses when a shock occurs.[2] The intuition behind this theoretical result is that responses of forward-looking households to shocks internalize the full expected path of future consumption and leisure. Therefore, responses in periods right after a shock occurs reveal the household's life-time welfare implications of additional transfers.

The proposition for representing the welfare gains from social insurance using spousal labor supply relies on optimality conditions, which are implied by the household's labor supply choices and are derived using the envelope theorem. Since these conditions must hold whenever households make optimal choices, we can still represent the gains from insurance by using the unaffected spouse's labor supply when arbitrary decision variables are added to the decision making process. Therefore, our welfare results hold in the generalized setting that accounts for other decision variables that may be included in the household's optimization problem. Specifically, we can incorporate the household's time use decisions, e.g., of how to allocate non-work time between home production and leisure. The robustness of our formula to the inclusion of additional margins of response is a general feature of the sufficient statistic approach to welfare analysis (see Chetty 2006a).

Setup. We consider a discrete-time setting in which households live for T periods $\{0, ..., T-1\}$ (where T is allowed to go to infinity) and set both the interest rate and the agents' time discount rate to zero for simplicity. Households consist of two individuals, members 1 and 2. We assume that at time 0 households are in the "good health" state ,state g, in which member 1 is in good health and works. In each period, the household transitions with probability ρ_t to the "bad health" state ,state b, in which member 1 experiences a health shock and drops out of the labor force. Conditional on being sick, 1 may die in period t with probability λ_t in which case the household transitions to the state where j is a widower or a widow, state d. In what follows, the subscript $i \in \{1, 2\}$ refers to the spouse and the superscript $s \in \{g, b, d\}$ refers to the state of nature.

At the beginning of the planning period, member 2 does not work and searches for a job. When 2 enters period t in state s without a job he or she chooses search intensity, e^s_{2t}, which we normalize to equal the probability of finding a job in the same period. If 2 finds a job, the job begins at time t and is assumed to last until the end of the planning period once found.[3]

Household Preferences. Let $u_t(c^s_t)$ represent the household's flow consumption utility at time t as a function of aggregate consumption at time t in state s, c^s_t. To let the model incorporate both the case in which the bad state is when member 1 is sick and the case in which member 1 is deceased, we set

[1] This model also allows for delays in the adjustment of labor supply by incorporating search frictions.

[2] It is important to note that when transfers are conditional on experiencing a shock, the moment of interest that comes out of the dynamic model is the labor supply response when the shock actually occurs rather than in expectation of it. The reason is that it captures the residual risk that was chosen not to be insured through the existing institutions given the probability of experiencing the shock in each period.

[3] This simplifies the algebra of the analysis. In the appendix of Fadlon and Nielsen (2015) we allow for job separations such that employment is absorbing within a health state but not across health states.

$u_t(c_t^s) = Q_t(c_t^s)$ when both spouses are alive and $u_t(c_t^s) = q_t(c_t^s)$ when member 1 does not survive. We assume that these functions are well-behaved – i.e., that $Q_t'(c_t^s) > 0$, $Q_t''(c_t^s) < 0$, $q_t'(c_t^s) > 0$, and $q_t''(c_t^s) < 0$ – which implies that $u_t'(c_t^s) > 0$ and $u_t''(c_t^s) < 0$ in each scenario. We denote member 2's cost of search effort at time t when unemployed by $\kappa(e_{2t}^s)$, which we assume to be strictly increasing and convex.

Policy Tools. The planner observes the state of nature as well as the employment status of each spouse. Since some spouses work and earn more than others do, the optimal policy is dependent on whether the spouse is employed. We denote the tax on spouse i's labor income in state g by T_i^g and the benefits given to non-working spouses in state g by m^g. In state $\sigma \in \{b, d\}$, households receive transfers of the amount M^σ, and households in which member 2 does not work receive additional benefits of the amount m^σ. We denote taxes by $T \equiv (T_1^g, T_2^g)$ and benefits by $B \equiv (m^g, M^b, m^b, M^d, m^d)$, and let $B^s(l_{2t}^s)$ represent the actual transfers received by a household in state s as a function of 2's participation, l_{2t}^s.

Household's Problem. At the beginning of each period t in state s, the household chooses the consumption flow, c_t^s, as well as member 2's search effort if 2 is unemployed, e_{2t}^s. In each period and state, 2's employment status, l_{2t}^s, determines the household's income flow, $y_t^s(l_{2t}^s)$, such that $y_t^s(l_{2t}^s) = \bar{z}_{1t}^s \times l_{1t}^s + \bar{z}_{2t}^s \times l_{2t}^s + B^s(l_{2t}^s) + I_t^s$, where z_{it} is i's labor income; $\bar{z}_{it}^s = z_{it} - T_i^s$ is i's labor income net of taxes in state s (with $T_i^\sigma = 0$, $\sigma \in \{b, d\}$); l_{1t}^s is an indicator for member 1's labor force participation; and I_t^s is the household's time-state contingent non-labor income or expenses (including life insurance payouts, transfers from any other source of individually-purchased or employer-provided private insurance, transfers from relatives, and medical expenses). This implies that each period's consumption as well as the next period's wealth – where we denote assets in period t by A_t – are functions of 2's participation, which we denote by $c_t^s(l_{2t}^s)$ and $A_{t+1}(l_{2t}^s)$, respectively. Therefore, the value function for households in state s who enter period t when 2 is without a job and with household assets A_t is

$$V_t^{s,0}(B,T,A_t) \equiv \max \left\{ \begin{array}{c} e_{2t}^s \left(u(c_t^s(1)) + W_{t+1}^{s,1}(B,T,A_{t+1}(1)) \right) \\ +(1-e_{2t}^s) \left(u(c_t^s(0)) + W_{t+1}^{s,0}(B,T,A_{t+1}(0)) \right) - \kappa(e_{2t}^s) \end{array} \right\},$$

where the budget constraints satisfy

$$c_t^s(l_{2t}^s) + A_{t+1}(l_{2t}^s) = A_t + y_t^s(l_{2t}^s),$$

and $W_{t+1}^{s,l_{2t}^s}(B,T,A_{t+1})$ are the continuation value functions which depend on whether the job search was successful or not in time t. The continuation functions are defined by

$$W_{t+1}^{g,l_{2t}^g}(B,T,A_{t+1}) \equiv (1-\rho_{t+1}) V_{t+1}^{g,l_{2t}^g}(B,T,A_{t+1}) + \rho_{t+1} V_{t+1}^{b,l_{2t}^g}(B,T,A_{t+1}),$$

$$W_{t+1}^{b,l_{2t}^b}(B,T,A_{t+1}) \equiv (1-\lambda_{t+1}) V_{t+1}^{b,l_{2t}^b}(B,T,A_{t+1}) + \lambda_{t+1} V_{t+1}^{d,l_{2t}^b}(B,T,A_{t+1}),$$

$$W_{t+1}^{d,l_{2t}^d}(B,T,A_{t+1}) \equiv V_{t+1}^{d,l_{2t}^d}(B,T,A_{t+1}),$$

where $V_t^{s,1}(B,T,A_t)$ is the value of entering period t when 2 is employed in state s which is defined by

$$V_t^{s,1}(B,T,A_t) \equiv \max \left\{ u(c_t^s(1)) + W_{t+1}^{s,1}(B,T,A_{t+1}(1)) \right\}.$$

The optimal search effort is chosen according to the first-order condition

$$\left(u(c_t^s(1)) + W_{t+1}^{s,1}(B,T,A_{t+1}(1)) \right) - \left(u(c_t^s(0)) + W_{t+1}^{s,0}(B,T,A_{t+1}(0)) \right) = \kappa'(e_{2t}^s), \tag{1}$$

where the effect of a \$1 increase in the benefit level m^s on search intensity in state s is

$$\frac{\partial e_{2t}^s}{\partial m^s} = -\frac{1}{\kappa''(e_{2t}^s)} \left(u'(c_t^s(0)) + \frac{\partial W_{t+1}^{s,0}}{\partial m^s} \right). \tag{2}$$

Planner's Problem. We define the household's expected utility at the beginning of the planning period by $J_0(B,T) \equiv (1-\rho_0) V_0^{g,0}(B,T,A_0) + \rho_0 V_0^{b,0}(B,T,A_0)$. The social planner's objective is to choose the tax-and-benefit system that maximizes the household's expected utility subject to a balanced-budget constraint.

For simplicity, we assume there is some expected revenue collected from each household and study the optimal redistribution of this revenue. We abstract from the specific way in which revenue is collected (or, similarly, assume a lump-sum tax that is determined outside of our problem) since our focus is on the benefits from social insurance and not its fiscal-externality costs. The perturbations we study involve increasing m^σ, $\sigma \in \{b, d\}$, by lowering m^g. Therefore, to further simplify the analysis we assume that $M^b = M^d = 0$, as well as that $m^d = 0$ when we perturb m^b, and that $m^b = 0$ when we perturb m^d.

Let D^s denote the expected share of the household's life-time in state s and let $\hat{e}_2^s$ denote the conditional probability of member 2 being employed if observed in state s. To construct the budget constraint, consider randomly choosing a household at a random point in its life-cycle. The probability of choosing a household in state s is D^s and, hence, the probability of choosing a household in state s in which 2 is unemployed is $D^s\left(1 - \hat{e}_2^s\right)$. If the government collects revenues of the amount r per household, a balanced budget requires that the expected transfer to a random household is equal to this amount. That is, $D^g\left(1 - \hat{e}_2^g\right) m^g + D^b\left(1 - \hat{e}_2^b\right) m^b + D^d\left(1 - \hat{e}_2^d\right) m^d = r$. Hence, the planner chooses the benefit levels B that solve

$$\max_B J_0(B, T) \quad \text{s.t. } D^g\left(1 - \hat{e}_2^g\right) m^g + D^b\left(1 - \hat{e}_2^b\right) m^b + D^d\left(1 - \hat{e}_2^d\right) m^d = r. \tag{3}$$

Optimal Social Insurance

We consider the optimal distribution of benefits to households with non-working spouses across health states σ and g ($\sigma \in \{b, d\}$). First, consider a \$1 increase in m^b financed by lowering m^g. The net welfare gain from this perturbation is

$$\frac{dJ_0(T, B)}{dm^b} = S_1^b + S_2^b \frac{dm^g}{dm^b}, \tag{4}$$

where $S_1^b \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial m^b} + (1 - \rho_0) \frac{\partial V_0^{g,0}}{\partial m^b}\right)$ and $S_2^b \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial m^g} + (1 - \rho_0) \frac{V_0^{g,0}}{\partial m^g}\right)$. The following proposition provides an approximated formula for a normalized version of this gain.

Proposition A1. *Under a locally quadratic approximation of the effort function around e_{20}^g, the marginal net benefit from raising m^b through a balanced-budget decrease in m^g is*

$$M_w(m^b) \cong MB(m^b) - MC(m^b),$$

with

1. $MB(m^b) \equiv \Phi^b \times \left(\frac{\hat{e}_{20}^b}{e_{20}^g}\right) - 1$, *where* $\Phi^b \equiv \frac{\left|\varepsilon(\hat{e}_{20}^b, m^b)\right| / m^b}{\left|\varepsilon(e_{20}^g, m^g)\right| / m^g}$, $\varepsilon(x, y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$, e_{20}^g *is 2's participation rate at the beginning of the planning period, and* $\hat{e}_{20}^b$ *is 2's mean participation rate in households that transition to state b.*

2. $MC(m^b) \equiv \beta_0^b + \beta_1^b \epsilon(1 - \hat{e}_2^g, m^b) + \beta_2^b \epsilon(1 - \hat{e}_2^b, b^b)$, *where the coefficients* β_0^b, β_1^b, *and* β_2^b *are functions of the transition probabilities, average participation rates and benefits, and* $\epsilon(x, y) \equiv \frac{dx}{dy}\frac{y}{x}$.[4]

Proof. The general logic of the proof is to characterize the derivatives of the value functions in their sequential problem representation – that is, as a sum of derivatives over time and over different states of nature. To do so, we work backwards from period $T - 1$ to period 0. Taylor approximations then lead to our results.

[4] Specifically, $\beta_0^b \equiv \frac{\sigma^b D^b\left(1 - \hat{e}_2^b\right) - D^g\left(1 - \hat{e}_2^g\right)}{D^g\left(1 - \hat{e}_2^g\right)}$, $\beta_1^b \equiv \sigma^b \frac{m^g}{m^b}$, and $\beta_2^b \equiv \sigma^b \frac{D^b\left(1 - \hat{e}_2^b\right)}{D^g\left(1 - \hat{e}_2^g\right)}$, where $\sigma^b \equiv (1 - p_0)\left(1 - e_{20}^g\right)/\rho(1 - \hat{e}_{20}^b)$ and $\rho \equiv \sum_{i=0}^{T-1}\left(\prod_{j=0}^{i-1}(1 - \rho_j)\rho_i\right)$. Note that the elasticities in $MC(b^b)$ consist of the total effect of increasing m^b, which takes into account the effect of lowering the level of the financing tool, m^g. Also note that with forward-looking households, transfers in states not yet encountered can have effects through ex-ante responses. For example, individuals in state g can lower labor supply and savings today in response to larger benefits in state b.

We begin by providing expressions for $\frac{\partial V_0^{b,0}}{\partial m^b}$ and $\frac{\partial V_0^{g,0}}{\partial m^b}$ in order to characterize S_1^b. First, we have that $\frac{\partial V_t^{b,0}}{\partial m^b} = (1-e_{2t}^b)\left(u'(c_t^b(0)) + \frac{\partial W_{t+1}^{b,0}}{\partial m^b}\right)$ and $\frac{\partial W_{t+1}^{b,0}}{\partial m^b} = (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial m^b}$, which imply that $\frac{\partial V_t^{b,0}}{\partial m^b} = (1-e_{2t}^b)\left(u'(c_t^b(0)) + (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial m^b}\right)$. Working backwards one can show that $\frac{\partial V_t^{b,0}}{\partial m^b} = (1-e_{2t}^b)\left[u'(c_t^b(0)) + \sum_{i=t+1}^{T-1}\left(\prod_{j=t+1}^{i}(1-e_{2j}^b)(1-\lambda_j)\right)(u'(c_i^b(0)))\right]$.

Next, since $\frac{\partial W_{t+1}^{g,1}}{\partial m^b} = 0$ we obtain $\frac{\partial V_t^{g,0}}{\partial m^b} = (1-e_{2t}^g)\frac{\partial W_{t+1}^{g,0}}{\partial m^b}$, where $\frac{\partial W_{t+1}^{g,0}}{\partial m^b} = (1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial m^b} + \rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial m^b}$. Therefore, we get that $\frac{\partial V_t^{g,0}}{\partial m^b} = (1-e_{2t}^g)(1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial m^b} + (1-e_{2t}^g)\rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial m^b}$, which implies by working backwards from period $T-1$ to period 0 that $\frac{\partial V_t^{g,0}}{\partial m^b} = (1-e_{2t}^g)\sum_{i=t+1}^{T-1}\left(\prod_{j=t+1}^{i-1}(1-e_{2j}^g)(1-\rho_j)\right)\rho_i\frac{\partial V_i^{b,0}}{\partial m^b}$.

Putting the terms together, it follows that

$$S_1^b = \left(\rho_0\frac{\partial V_0^{b,0}}{\partial m^b} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial m^b}\right) = \sum_{i=0}^{T-1}\left[\prod_{j=0}^{i-1}(1-e_{2j}^g)(1-\rho_j)\rho_i\right]\frac{\partial V_i^{b,0}}{\partial m^b}. \tag{5}$$

Using equation (2) and $\frac{\partial V_t^{b,0}}{\partial m^b} = (1-e_{2t}^b)\left(u'(c_t^b(0)) + \frac{\partial W_{t+1}^{b,0}}{\partial m^b}\right)$, we get that $\frac{\partial V_t^{b,0}}{\partial m^b} = -\kappa''(e_{2t}^b)\frac{\partial e_{2t}^b}{\partial m^b}(1-e_{2t}^b)$. Plugging this expression into (5) yields the following result

$$S_1^b = -\sum_{i=0}^{T-1}\left[\prod_{j=0}^{i-1}(1-e_{2j}^g)(1-\rho_j)\rho_i\right](1-e_{2i}^b)\kappa''(e_{2i}^b)\frac{\partial e_{2i}^b}{\partial m^b}. \tag{6}$$

To understand the meaning of this formula let us break it down into its components. First, note that it is a weighted sum of a function of the change in effort (or participation rate), $\frac{\partial e_{2i}^b}{\partial m^b}$. The weight, the term in brackets, is the probability of reaching period i with 2 unemployed and transitioning to state b exactly in that period. For households that transition to state b in period i when 2 is employed, the change in effort and participation rates is zero (because they stay employed and do not engage in search effort). Therefore, dividing the probability weights by the chance of transitioning to state b at some point throughout the planning horizon, $\rho \equiv \sum_{i=0}^{T-1}\left(\prod_{j=0}^{i-1}(1-\rho_j)\rho_i\right)$, and rewriting (6) in terms of elasticities (with $\varepsilon(x,y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$) yield $S_1^b = \rho E_b\left\{(1-\bar{e}_{20}^b)\kappa''(\bar{e}_{20}^b)\left|\varepsilon(\bar{e}_{20}^b, m^b)\right|\frac{\bar{e}_{20}^b}{m^b}\right\} \equiv \rho E_b(g(\bar{e}_{20}^b))$, where $\bar{e}_{20}^b$ denotes participation in the period the household transitions to state b and E_b is the expectation operator conditional on being in state b. By expanding $g(e)$ around 2's average participation in households in which 1 becomes sick – which we denote by $\hat{e}_{20}^b$ – such that $g(e) \cong g(\hat{e}_{20}^b) + g'(\hat{e}_{20}^b)(e-\hat{e}_{20}^b)$, we approximate $E_b(g(\bar{e}_{20}^b)) \cong E_b(g(\hat{e}_{20}^b)) = g(\hat{e}_{20}^b)$ and obtain the approximation

$$S_1^b \cong \rho(1-\hat{e}_{20}^b)\kappa''(\hat{e}_{20}^b)\left|\varepsilon(\hat{e}_{20}^b, m^b)\right|\frac{\hat{e}_{20}^b}{m^b}. \tag{7}$$

We now turn to provide expressions for $\frac{\partial V_0^{b,0}}{\partial m^g}$ and $\frac{V_0^{g,0}}{\partial m^g}$ in order to characterize S_2^b. Since households that transitioned to state b either stay in state b or transition to state d, we have that $\frac{\partial V_0^{b,0}}{\partial m^g} = 0$. In addition, $\frac{V_t^{g,0}}{\partial m^g} = (1-e_{2t}^g)\left(u'(c_t^g(0)) + \frac{\partial W_{t+1}^{g,0}}{\partial m^g}\right)$, which combined with equation (2) yields $\frac{V_t^{g,0}}{\partial m^g} = -(1-e_{2t}^g)\left(\kappa''(e_{2t}^g)\frac{\partial e_{2t}^g}{\partial m^g}\right)$. Put together, we get that

$$S_2^b = (1-\rho_0)(1-e_{20}^g)\kappa''(e_{20}^g)\left|\varepsilon(e_{20}^g, m^g)\right|\frac{e_{20}^g}{m^g}. \tag{8}$$

To complete the proof we need to calculate $\frac{dm^g}{dm^b}$. Total differentiation of the simplified budget constraint $D^g(1-\hat{e}_2^g)m^g + D^b(1-\hat{e}_2^b)m^b = r$ with respect to m^b gives us

$$\frac{dm^g}{dm^b} = -\frac{m^g}{m^b}\epsilon(1-\hat{e}_2^g, m^b) - \frac{D^b(1-\hat{e}_2^b)}{D^g(1-\hat{e}_2^g)}\epsilon(1-\hat{e}_2^b, m^b) - \frac{D^b(1-\hat{e}_2^b)}{D^g(1-\hat{e}_2^g)}, \tag{9}$$

where $\epsilon(x,y) \equiv \frac{dx}{dy}\frac{y}{x}$. Plugging (7), (8), and (9) into (4), and using a quadratic approximation of the effort function around e_{20}^g, we obtain the approximated formula for the normalized welfare gain $M_w(m^b) \equiv \frac{\frac{dJ_0(T,B)}{dm^b}/\rho(1-\hat{e}_{20}^b)}{\frac{\partial J_0(T,B)}{\partial m^g}/(1-\rho_0)(1-e_{20}^g)}$ that is stated in the proposition, which completes the proof. ■

Next, consider a \$1 increase in m^d financed by lowering m^g. We analyze this perturbation separately from the former since the sequential nature of the model requires a more careful investigation of transfers to different "bad" states (as shown in the following proof), although the approximated formulas turn out to be similar. The net welfare gain from this perturbation is

$$\frac{dJ_0(T,B)}{dm^d} = S_1^d + S_2^d \frac{dm^g}{dm^d}, \tag{10}$$

where $S_1^d \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial m^d} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial m^d}\right)$ and $S_2^d \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial m^g} + (1-\rho_0)\,\frac{V_0^{g,0}}{\partial m^g}\right)$. We present the approximated formula in the following proposition.

Proposition A2. *Under a locally quadratic approximation of the effort function around e_{20}^g, the marginal net benefit from raising m^d through a balanced-budget decrease in m^g is*

$$M_w(m^d) \cong MB(m^d) - MC(m^d),$$

with

1. $MB(m^d) \equiv \Phi^d \times \left(\frac{\hat{e}_{20}^d}{e_{20}^d}\right) - 1$, *where* $\Phi^d \equiv \frac{\left|\varepsilon(\hat{e}_{20}^d, m^d)\right|/m^d}{\left|\varepsilon(e_{20}^g, m^g)\right|/m^g}$, $\varepsilon(x,y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$, e_{20}^g *is 2's participation rate at the beginning of the planning period, and* $\hat{e}_{20}^d$ *is 2's mean participation rate in households that transition to state d.*
2. $MC(m^d) \equiv \beta_0^d + \beta_1^d \epsilon(1-\hat{e}_2^g, m^d) + \beta_2^d \epsilon(1-\hat{e}_2^d, m^d)$, *where the coefficients* β_0^d, β_1^d, *and* β_2^d *are functions of the transition probabilities, average participation rates and benefits, and* $\epsilon(x,y) \equiv \frac{dx}{dy}\frac{y}{x}$.[5]

Proof. We first find expressions for $\frac{\partial V_0^{b,0}}{\partial m^d}$ and $\frac{\partial V_0^{g,0}}{\partial m^d}$ in order to characterize S_1^d. With $\frac{\partial V_t^{b,0}}{\partial m^d} = (1 - e_{2t}^b)\left(\frac{\partial W_{t+1}^{b,0}}{\partial m^d}\right)$ and $\frac{\partial W_{t+1}^{b,0}}{\partial m^d} = (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial m^d} + \lambda_{t+1}\frac{\partial V_{t+1}^{d,0}}{\partial m^d}$ we have that $\frac{\partial V_t^{b,0}}{\partial m^d} = (1-e_{2t}^b)\left((1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial m^d} + \lambda_{t+1}\frac{\partial V_{t+1}^{d,0}}{\partial m^d}\right)$. Working backwards from period $T-1$ to period 0 one can show that $\frac{\partial V_t^{b,0}}{\partial m^d} = \sum_{i=t+1}^{T-1}\prod_{j=t}^{i-1}(1-e_{2j}^b)\prod_{j=t+1}^{i-1}(1-\lambda_j)\lambda_j\frac{\partial V_i^{d,0}}{\partial m^d}$.

In state g we have $\frac{\partial V_t^{g,0}}{\partial m^d} = (1-e_{2t}^g)\left(\frac{\partial W_{t+1}^{g,0}}{\partial m^d}\right)$ and $\frac{\partial W_{t+1}^{g,0}}{\partial m^d} = \rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial m^d} + (1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial m^d}$, which imply that $\frac{\partial V_t^{g,0}}{\partial m^d} = (1-e_{2t}^g)\left(\rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial m^d} + (1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial m^d}\right)$. Define the probability of transitioning to state d exactly at time i while 2 is unemployed by $\mu_i^{d,0}$ (which takes into account all the possible transition paths). Then, combining the results so far one can show by working backwards that $S_1^d = \left(\rho_0\frac{\partial V_0^{b,0}}{\partial m^d} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial m^d}\right) = \sum_{i=t}^{T-1}\mu_i^{d,0}E_{\mu_i^{d,0}}\left[\frac{\partial V_i^{d.0}}{\partial m^d}\right]$, where $E_{\mu_i^{d,0}}$ is the expectation operator conditional on arriving at period i with 2 unemployed and transitioning to state d then (taken over all possible paths).

Since $\frac{\partial V_t^{d.1}}{\partial m^d} = 0$ we have that $\frac{\partial V_t^{d.0}}{\partial m^d} = (1-e_{2t}^d)\left(u'(c_t^d(0)) + \frac{\partial V_{t+1}^{d.0}}{\partial m^d}\right)$. Combined with (2) it can be expressed as $\frac{\partial V_t^{d.0}}{\partial m^d} = -(1-e_{2t}^d)\kappa''(e_{2t}^d)\frac{de_{2t}^d}{dm^d}$. Putting the terms together we obtain

$$S_1^d = \sum_{i=t}^{T-1}\mu_i^{d,0}E_{\mu_i^{d,0}}\left[(1-e_{2i}^d)\kappa''(e_{2i}^d)\left|\varepsilon(e_{2i}^d, m^d)\right|\frac{e_{2i}^d}{m^d}\right]. \tag{11}$$

[5] Specifically, $\beta_0^d \equiv \frac{\sigma^d D^d\left(1-\hat{e}_2^d\right) - D^g\left(1-\hat{e}_2^g\right)}{D^g\left(1-\hat{e}_2^g\right)}$, $\beta_1^b \equiv \sigma^d\frac{m^g}{m^d}$, and $\beta_2^b \equiv \sigma^d\frac{D^d\left(1-\hat{e}_2^d\right)}{D^g\left(1-\hat{e}_2^g\right)}$, where $\sigma^d \equiv (1-p_0)\left(1-e_{20}^g\right)/\lambda(1-\hat{e}_{20}^d)$, $\lambda \equiv \sum_{i=0}^{T-1}\mu_i^d$, and μ_i^d is the probability of transitioning to state d in period i.

Define the probability of transitioning to state d in period i by μ_i^d and note that for those households that arrive at this period with 2 employed the change in participation is zero. Dividing the probabilities in (11) by the chance of transitioning to state d at some point, $\lambda \equiv \sum_{i=0}^{T-1} \mu_i^d$, we can rewrite S_1^d as $S_1^d = \lambda E_\lambda \left\{ (1-\bar{e}_{20}^d)\kappa''(\bar{e}_{20}^d) \left| \varepsilon(e_{2i}^d, m^d) \right| \frac{\bar{e}_{20}^d}{m^d} \right\} \equiv \lambda E_\lambda (g(\bar{e}_{20}^d))$, where $\bar{e}_{20}^d$ denotes participation in the period the household transitions to state d and E_λ is the expectation operator conditional on being in state d. Expanding $g(e)$ around 2's average participation upon the transition to state d – which we denote by $\hat{e}_{20}^d$ – we can approximate S_1^d by

$$S_1^d \cong \lambda(1-\hat{e}_{20}^d)\kappa''(\hat{e}_{20}^d) \left| \varepsilon(\hat{e}_{20}^d, m^d) \right| \frac{\hat{e}_{20}^d}{m^d}. \tag{12}$$

In addition, as in the proof of Proposition A1

$$S_2^d = \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial m^g} + (1-\rho_0) \frac{V_0^{g,0}}{\partial m^g} \right) = (1-\rho_0)(1-e_{20}^g)\kappa''(e_{20}^g) \left| \varepsilon(e_{20}^g, m^g) \right| \frac{e_{20}^g}{m^g}. \tag{13}$$

To complete the proof we differentiate the budget constraint with respect to m^d which yields

$$\frac{dm^g}{dm^d} = -\frac{m^g}{m^d}\epsilon(1-\hat{e}_2^g, m^d) - \frac{D^d\left(1-\hat{e}_2^d\right)}{D^g\left(1-\hat{e}_2^g\right)}\epsilon(1-\hat{e}_2^d, m^d) - \frac{D^d\left(1-\hat{e}_2^d\right)}{D^g\left(1-\hat{e}_2^g\right)}. \tag{14}$$

Plugging (12), (13), and (14) into (10), and using a quadratic approximation of the effort function around e_{20}^g, we obtain the approximated formula for the normalized welfare gain $M_w(m^d) \equiv \frac{\frac{dJ_0(T,B)}{dm^d}/\lambda(1-\hat{e}_{20}^d)}{\frac{\partial J_0(T,B)}{\partial m^g}/(1-\rho_0)(1-e_{20}^g)}$ that is stated in the proposition, which completes the proof. ■

Online Appendix B: Generalized Preferences

Besides income losses, there are other important ways in which households can be directly affected by the shocks that we analyze. In particular, preferences can change in several dimensions, which can lead to spousal labor supply responses even in the presence of full insurance. In this appendix, we generalize the preference structure in the model of Section 5 by considering different potential types of such state dependence in preferences. We also generalize preferences by allowing for flexible consumption-leisure complementarities.

Let $U^s(c^s; l_1^s, l_2^s)$ represent the household's utility as a function of aggregate consumption, c^s, and the household members' labor force participation, l_1^s and l_2^s, in state s, where $l_i^s = 1$ if i works and $l_i^s = 0$ otherwise. We assume that $U^s(c^s; l_1^s, l_2^s) = u^s(c^s; l_1^s, l_2^s) - v_1^s \times l_1^s - v_2^s \times l_2^s$, where $u^s(c^s)$ is the household's utility from consumption in state s, and v_i^s represents each member i's disutility from labor in state s.

This formulation generalizes preferences as follows. First, it allows for a completely flexible dependence of consumption utility on the state of nature. Second, it allows for flexible consumption-leisure complementarities by allowing the consumption utility to depend freely on participation. Third, we allow labor disutility, v_i^s, to change across states of nature. For the affected spouse 1, this captures the direct effect of health on the ability to work when state b is 1's sickness. For the unaffected spouse 2, this generalization captures the potential state dependence in the utility cost of supplying labor or the willingness to work. For example, when the bad state is 1's sickness, v_2^b might be greater than the baseline labor disutility v_2^g if 2 places greater value on time spent at home – e.g., to take care of his or her sick spouse. When the bad state is 1's death, working may become less desirable if the surviving spouse experiences depression and has difficulties working, or conversely, working may become more desirable if the surviving spouse feels lonely and wishes to seek social integration. For simplicity, we model this type of state dependence as $v_2^g = v_2$ and $v_2^b = \theta^b \times v_2^g$, such that θ^b captures the mean percent change in the utility cost of labor compared to the baseline state g.[6]

With these generalized preferences, potential changes in the unaffected spouse's labor disutility can directly lead to spousal labor supply responses. Even with complete insurance ($L = 0$), a decrease in spouse

[6] In the appendix for the dynamic search model in Fadlon and Nielsen (2015) we show that this is a simplification and that it is not necessary to define such a global parameter for our theoretical results, by illustrating how it can be locally and non-parametrically defined. In addition, we offer there an example for allowing heterogeneity in θ^b.

2's labor disutility in the transition from state g to state b (i.e., $\theta^b < 1$) will cause an increase in spousal labor force participation (such that $\frac{e_2^b}{e_2^g} - 1 > 0$).

Labor Supply Representation of Welfare Benefits with State-Dependent Preferences. The generalized preference structure can also affect the normative results. Since our welfare formula represents gains from additional social insurance using the labor supply behavior of the unaffected spouse, the generalization that impacts the normative analysis is confined to potential changes in the unaffected spouse's labor disutility. This means that the dependence of consumption utility on the state of nature and the consumption-leisure complementarities that we introduced have no effect on the welfare formula. It also implies that allowing the labor disutility of the affected sick spouse, v_1^s, to change completely across states of nature does not affect the analysis. It is indeed the underlying motive for studying the unaffected spouse's behavior in the first place since in the case of health shocks the affected spouse's preferences can change in many unidentifiable ways as a result of the shock. The potential state dependence of the unaffected spouse's labor disutility requires adjusting the welfare formula in the following way:

Proposition B1. *With the generalized preference structure, the marginal benefit from raising m^b through a balanced-budget decrease in m^g is*

$$MB = \theta^b \times \Phi \times \left(\frac{e_2^b}{e_2^g}\right) - 1, \tag{15}$$

where $\Phi \equiv \phi^b/\phi^g$, $\phi^s \equiv \frac{|\varepsilon(e_2^s, m^s)|}{m^s \times f(\bar{v}_2^s)}$, and $\varepsilon(e_2^s, m^s)$ is the unaffected spouse's participation elasticity with respect to the policy tool m^s.

Proof. The proof follows the proof of Proposition 1 with the adjustment of using the generalized preference structure.

Intuition. Compared to the formula in Proposition 1, the formula in the generalized case changes such that the "price" component, Φ, adjusts to $\theta^b \times \Phi$. This adjustment is driven by the relative cost of the unaffected spouse's labor supply across states of nature which becomes larger by a factor of θ^b. Intuitively, since the formula assesses benefits from social insurance by evaluating the change in the consumption of leisure, higher valuation of leisure, that is, $\theta^b > 1$, renders leisure more valuable in state b, which makes the transfer of resources from state g to state b more socially desirable. On the other hand, lower cost of labor supply following a shock ($\theta^b < 1$) can lead to an increase in labor force participation even if households are well insured. The welfare implications of labor supply responses in this case are different (so that additional transfers to the bad state may even become undesirable) since these responses would be driven by preferences for work and not by under-insurance.

Online Appendix C: An Intensive-Margin Model of Household Labor Supply

The choice of the appropriate model for welfare analysis should depend on the context. In our specific context, we saw empirically that survivors' responses were concentrated on the participation margin. However, other applications, such as studying a sub-population with full employment before a shock occurs, would call for an intensive margin model because work intensity is expected to be the operative margin. In this appendix, we present a baseline static model that is the intensive margin counterpart to the participation model in the text. The analysis of the dynamic version of this model follows the logic of the analysis in Online Appendix A and is available from the authors on request. In the model that we analyze here we chose to use the collective approach to household behavior (Chiappori 1988, 1992; Apps and Rees 1988) to illustrate that the welfare results extend to models of the household other than the unitary framework.

Setup. Households consist of two individuals, 1 and 2. We consider a world with two states of nature: a "good" state, state g, in which 1 is in good health, and a "bad" state, state b, in which 1 experiences a shock. Households spend a share of μ^g of their adult life in state g and a share of μ^b in state b ($\mu^g + \mu^b = 1$). In what follows, the subscript $i \in \{1, 2\}$ refers to the spouse and the superscript $s \in \{g, b\}$ refers to the state of nature.

Individual Preferences. Let $U_i(c_i^s, l_i^s)$ represent i's utility as a function of consumption, c_i^s, and labor supply, l_i^s, in state s. We assume that $\frac{\partial U_i}{\partial c_i^s} > 0$, $\frac{\partial^2 U_i}{\partial (c_i^s)^2} < 0$, $\frac{\partial U_i}{\partial l_i^s} < 0$, and $\frac{\partial^2 U_i}{\partial (l_i^s)^2} < 0$.

Household Preferences. We follow the collective approach to household behavior and assume that household decisions are Pareto efficient and can be characterized as solutions to the maximization of $\beta_1 U_1(c_1^s, l_1^s) + \beta_2 U_2(c_2^s, l_2^s)$, where β_1 and β_2 are the Pareto weights on 1 and 2, respectively. For simplicity, we assume equal Pareto weights ($\beta_1 = \beta_2 = 1$), which is without loss of generality as long as the spouses' relative bargaining power is stable across states of nature.[7]

Policy Tools. Households in state b receive transfers of the amount B, which are financed by a linear tax rate τ_i^s on i's labor income in state s. We denote taxes by $T \equiv (\tau_1^g, \tau_2^g, \tau_1^b, \tau_2^b)$ and actual transfers by B^s such that $B^g = 0$ and $B^b = B$.

Household's Problem. In each state s the household solves the following problem

$$V^s(B,T,A) \equiv \max_{c_i^s, l_i^s} U_1(c_1^s, l_1^s) + U_2(c_2^s, l_2^s)$$

$$\text{s.t. } c_1^s + c_2^s = A^s + w_1^s\left(1-\tau_1^s\right) l_1^s + w_2\left(1-\tau_2^s\right) l_2^s + B^s,$$

where A^s is the household's state-contingent wealth and non-labor income, w_1^s is 1's wage rate in state s, and w_2 is 2's wage rate. The household's first-order conditions imply that $\frac{\partial U_1}{\partial c_1^s} = \frac{\partial U_2}{\partial c_2^s} = -\frac{\partial U_2}{\partial l_2^s} \frac{1}{w_2\left(1-\tau_2^s\right)}$. Importantly, note that we allow 1 to be at a corner solution in state b – that is, $l_1^b = 0$ – and use only 2's labor supply first-order conditions.

Planner's Problem. The social planner's objective is to choose the tax-and-benefit system that maximizes the household's expected utility, $J(B,T) \equiv \mu^g V^g(B,T,A) + \mu^b V^b(B,T,A)$, subject to the requirement that expected benefits paid, $\mu^b B$, equal expected taxes collected, $\mu^g(\tau_1^g w_1^g l_1^g + \tau_2^g w_2 l_2^g) + \mu^b(\tau_1^b w_1^b l_1^b + \tau_2^b w_2 l_2^b)$. Hence, the planner chooses the benefit level B and taxes T that solve

$$\max_{B,T} J(B,T) \quad \text{s.t. } \mu^b B = \mu^g(\tau_1^g w_1^g l_1^g + \tau_2^g w_2 l_2^g) + \mu^b(\tau_1^b w_1^b l_1^b + \tau_2^b w_2 l_2^b). \tag{16}$$

Optimal Social Insurance

Consider a \$1 increase in B financed by an appropriate increase in taxes, e.g., through τ_1^g. To simplify notation we assume that $\tau_2^g = \tau_1^b = \tau_2^b = 0$, which allows us to obtain concise welfare formulas.[8] The welfare gain from this perturbation is $\frac{dJ(B,T)}{dB} = \mu^b \frac{\partial V^b}{\partial B} + \mu^g \frac{\partial V^g}{\partial \tau_1^g} \frac{d\tau_1^g}{dB}$, which we normalize by the welfare gain from raising 1's net-of-tax labor income in state g by \$1 (scaled by the targeted population) to gain a cardinal interpretation.[9] Exploiting the envelope theorem (in the differentiation of the household's value functions) and using the household's first-order conditions, we obtain $\frac{\partial V^g}{\partial \tau_1^g} = -w_1^g l_1^g \frac{\partial U_2}{\partial c_2^g}$ and $\frac{\partial V^b}{\partial B} = \frac{\partial U_2}{\partial c_2^b}$. Differentiating the budget constraint with respect to B we get $\frac{d\tau_1^g}{dB} = \frac{\mu^b}{\mu^g z_1^g}\left(1 + \frac{\varepsilon(z_1^g, 1-\tau_1^g)\frac{\tau_1^g}{1-\tau_1^g}}{1-\varepsilon(z_1^g, 1-\tau_1^g)\frac{\tau_1^g}{1-\tau_1^g}}\right)$, where $z_1^g \equiv w_1^g l_1^g$ is 1's taxable income and $\varepsilon(z_1^g, 1-\tau_1^g) \equiv \frac{\partial z_1^g}{\partial(1-\tau_1^g)} \frac{1-\tau_1^g}{z_1^g}$ is the commonly estimated net-of-tax taxable income elasticity. Put together, it follows that the normalized welfare gain from a marginal increase in B is $M_W(B) = MB(B) - MC(B)$, where $MB(B) \equiv \frac{\frac{\partial U_2}{\partial c_2^b} - \frac{\partial U_2}{\partial c_2^g}}{\frac{\partial U_2}{\partial c_2^g}}$ and $MC(B) \equiv \frac{\varepsilon(z_1^g, 1-\tau_1^g)\frac{\tau_1^g}{1-\tau_1^g}}{1-\varepsilon(z_1^g, 1-\tau_1^g)\frac{\tau_1^g}{1-\tau_1^g}}$.

[7] Browning et al. (2014) discuss the important distinction in the collective model between ex-post realizations of different states of nature, which should not affect the spouses' relative bargaining power under efficient risk sharing, and ex-ante distributions of income shocks, which may affect the Pareto weights. Similar to Chiappori (1992), baseline weights do not affect our welfare results.

[8] Relaxing this assumption would result in additional elasticities in $MC(B)$ below. In particular, when calculating the change in government revenues, we would need to take into account *any* possible margin that can respond to the change *and* is being taxed. For example, if we added taxes on 2, we would need to include his or her labor supply responses to changes in 1's tax rate.

[9] The formula for the normalized gain is $M_W(B) \equiv \frac{\frac{dJ(B,T)}{dB}/\mu^b}{\frac{\partial J(B,T)}{\partial z_1^g(1-\tau_1^g)}/\mu^g}$, where $z_1^g \equiv w_1^g l_1^g$.

Labor Supply Representation of Welfare Benefits. The representation of the gap in marginal utilities of consumption using the unaffected spouse's labor supply responses in the intensive-margin model is summarized in the following proposition:

Proposition C1. *Assuming consumption-leisure separability,*[10] *the marginal benefit from raising B is approximately*

$$MB(B) \cong \varphi \times \left(\frac{l_2^b - l_2^g}{l_2^g}\right), \tag{17}$$

where $\varphi \equiv \frac{\partial^2 U_2/\partial\left(l_2^g\right)^2}{\partial U_2/\partial l_2^g} l_2^g$.

Proof. Recall that the household's first-order conditions imply that $\frac{\partial U_1}{\partial c_1^s} = \frac{\partial U_2}{\partial c_2^s} = -\frac{\partial U_2}{\partial l_2^s}\frac{1}{w_2}$. This allows us to map i's marginal utility from consumption to the unaffected spouse's marginal disutility from labor, such that $MB(B) = \frac{\left|\frac{\partial U_2}{\partial l_2^b}\right| - \left|\frac{\partial U_2}{\partial l_2^g}\right|}{\left|\frac{\partial U_2}{\partial l_2^g}\right|}$. Following Gruber's (1997) analysis for estimating the consumption representation of the welfare formula (see also Chetty and Finkelstein 2013), we take a second-order approximation of 2's labor disutility function around l_2^g. The consumption-leisure separability assumption yields the result. ■

Note the comparison to the extensive-margin model, in which the marginal entrant, $\bar{v}_2^s$, reveals the cost of labor supply on the margin, and the change in the marginal entrant's disutility, $\frac{\partial \bar{v}_2^s}{\partial m^s}$, evaluates the relative cost of labor supply across states of nature. In the intensive margin model here, it is the cost of the marginal hour of work, $\frac{\partial U_2}{\partial l_2^s}$, that reveals the cost of labor supply on the margin, and the change in the cost of the marginal hour, $\frac{\partial^2 U_2}{\partial\left(l_2^s\right)^2}$ (or φ), is used to evaluate the cross-state labor supply responses. In the appendix of Fadlon and Nielsen (2015), we show that φ can be estimated using within-state labor supply elasticities (similar to the estimation of risk aversion in Chetty 2006b).

References

Aguila, E., O. Attanasio, and C. Meghir (2011). Changes in consumption at retirement: Evidence from panel data. *Review of Economics and Statistics 93*(3), 1094–1099.

Apps, P. F. and R. Rees (1988). Taxation and the household. *Journal of Public Economics 35*(3), 355–369.

Browning, M., P.-A. Chiappori, and Y. Weiss (2014). *Economics of the Family.* Cambridge University Press.

Chetty, R. (2006a). A general formula for the optimal level of social insurance. *Journal of Public Economics 90*(10-11), 1879–1901.

Chetty, R. (2006b). A new method of estimating risk aversion. *American Economic Review 96*(5), 1821–1834.

Chetty, R. and A. Finkelstein (2013). Chapter 3 - social insurance: Connecting theory to data. Volume 5 of *Handbook of Public Economics*, pp. 111–193. Elsevier.

Chiappori, P.-A. (1988). Rational household labor supply. *Econometrica 56*(1), 63–90.

Chiappori, P.-A. (1992). Collective labor supply and welfare. *Journal of Political Economy 100*(3), 437–467.

Fadlon, I. and T. H. Nielsen (2015). Household responses to severe health shocks and the design of social insurance. Working Paper 21352, National Bureau of Economic Research.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review 87*(1), 192–205.

[10] Recent research finds supportive evidence for consumption-leisure separability – e.g., Aguila et al. (2011) who find no change in consumption (defined as non-durable expenditure) around retirement. However, complementarities between consumption and leisure can be handled by estimating the cross-partial using the technique in Chetty (2006b).