NBER WORKING PAPER SERIES

FAMILY LABOR SUPPLY RESPONSES TO SEVERE HEALTH SHOCKS

Itzik Fadlon
Torben Heien Nielsen

Working Paper 21352
http://www.nber.org/papers/w21352

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge MA 02138
July 2015, Revised September 2017

An earlier version of this paper has been previously circulated as NBER Working Paper No. 21352 titled “Household Responses to Severe Health Shocks and the Design of Social Insurance.” We thank David Autor, Martin Browning, Raj Chetty, Jeffrey Clemens, Julie Cullen, David Cutler, Gordon Dahl, Ariel Dora Stern, Raluca Dragusanu, Liran Einav, Torben Fischer, Eric French, Elhanan Helpman, Nathan Hendren, Simon Jäger, Xavier Jaraverl, Lawrence Katz, Claus Thustrup Kreiner, David Laibson, Jessica Laird, Rafael Lalive, Jonathan Leganza, Søren Leth-Petersen, Magne Mogstad, Karthik Muralidharan, 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 or the National Bureau of Economic Research.

Family Labor Supply Responses to Severe Health Shocks
Itzik Fadlon and Torben Heien Nielsen
NBER Working Paper No. 21352
July 2015, Revised September 2017
JEL No. H0,I1,J1,J2

ABSTRACT

This paper provides new evidence on how household labor supply responds to fatal and severe non-fatal health shocks in the short- and medium-run. To identify the causal effects of these shock realizations, we leverage administrative data on families' health and labor market outcomes, and construct counterfactuals to affected households by using households that experience the same shock but a few years in the future. We find that fatal health shocks lead to an immediate increase in the surviving spouses' labor supply and that this effect is entirely driven by families who experience significant income losses. Accordingly, widows, who face large income losses when their husbands die, increase their labor force participation by more than 11%; while widowers, who are significantly more financially stable, slightly decrease their labor supply. Notably, however, the patterns of sensitivity to comparable income changes are similar across genders. In contrast to fatal shocks, we find that non-fatal health shocks—in particular, heart attacks or strokes—have no meaningful effects on spousal labor supply, consistent with the adequate insurance coverage for the associated foregone income. Overall, the results point to self-insurance as a driving mechanism for the family labor supply responses that we estimate. Combined with a stylized model, our findings suggest efficient ways to target government transfers through existing social insurance programs.

Itzik Fadlon
Department of Economics
University of California, San Diego
9500 Gilman Drive, #0508
La Jolla, CA 92093
and NBER
fadlon@ucsd.edu

Torben Heien Nielsen
University of Copenhagen
Department of Economics
Øster Farimagsgade 5, Building 25
DK-1353 Copenhagen K
torben.heien.nielsen@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 financial risk. Studying how households respond to severe adverse health events is therefore important for our understanding of self-insurance behavior over the life cycle, where a key potential self-insurance mechanism against income shocks is the labor supply of family members. Beyond its implications for household behavior, the degree to which households insure through labor supply is central for the design of social insurance programs. The programs that protect households against the potential income losses imposed by fatal and non-fatal health shocks—namely, survivors and disability insurance—have become among the largest safety-net programs in most OECD countries in recent decades (OECD 2014).[1]

Consequently, economists have been long interested in analyzing the effects of adverse health shocks on one hand, and in empirically uncovering the insurance role of spousal labor supply on the other hand. Yet, there is markedly limited direct evidence regarding the important link between these two significant strands of the literature. Specifically, we lack clear consensus about family members' labor supply responses to severe health shocks, and there is virtually no work on the impacts of fatal shocks in the modern literature.

Estimating these responses has been impeded by two main challenges. The first is the unavailability of large-scale *household-level* data on health and labor market outcomes, which are necessary for accurate estimation of family (rather than own) labor supply responses. The second obstacle is the difficulty of isolating causal effects of shocks in the presence of complex dynamics. Identification of impacts in our context requires constructing counterfactuals that account for life-cycle and time patterns in family labor supply which, among many other factors, are likely to depend on ex-ante expectations. Some papers have successfully done so in various contexts by using matched control groups from the pool of untreated units based on observables.[2] However, as we illustrate below, strategies that rely on unaffected households as controls are inadequate for our purposes. In particular, we show within our setting that affected and observably-similar unaffected households exhibit substantially different behavioral patterns over time, in violation of the requirement of parallel pre-trends across the two groups.

In this paper, we study how spousal labor supply responds to fatal and severe non-fatal health shocks by leveraging long panels of administrative data on families' health and labor market outcomes. The data—which encompass the entire Danish population from the years 1980-2011—provide register-based information on health-care utilization, income, wealth, and labor market behavior.

[1] For example, in 2014 the United States government paid 93 billion dollars to more than 4 million surviving spouses and 132 billion dollars to 9 million disabled workers through the Old-Age, Survivors, and Disability Insurance program. By comparison, 46 billion dollars were paid in unemployment benefits, and the outlays within the Supplemental Security Income (SSI) program, the Temporary Assistance for Needy Families (TANF), the Supplemental Nutrition Assistance Program (SNAP), and the Earned Income Tax Credit (EITC) scheme were 51, 20, 76, and 60 billion dollars, respectively (SSA 2015; White House 2015).

[2] See, for example, Goldschmidt and Schmieder (2017) who study the effects of job outsourcing on wages, and Jäger (2016) who studies the substitutability of workers by analyzing the impact of worker exits due to unexpected deaths.

This combination of large-scale objective health and labor market information and the ability to link across spouses provides us with a particularly well-suited setting for studying household labor supply responses in the context of developed economies. Starting from the universe of married and cohabiting couples, we study over 500,000 households in which one spouse experienced a fatal health shock (including a sub-sample of fatal heart attacks and strokes), and over 70,000 households in which one spouse experienced a non-fatal heart attack or stroke.

To identify the causal effects of experiencing these adverse shocks, we employ a 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 groups and estimate the short- and medium-run treatment effects using a straightforward dynamic (i.e., period-by-period) difference-in-differences estimator. Our identifying assumption is that, absent the realization of the shock, the outcomes of the treatment and control groups would run parallel. Among many other dimensions, this would ensure the similarity of the groups in terms of their expectations. Reassuringly, we demonstrate that the pre-trends run parallel for all of the outcomes we study. Of course, this approach has its own limitations; mainly, it places an upper bound on the analysis' time horizon since the control group becomes "treated" within a few years. The estimation strategy we use, that aims to estimate ex-post responses to realizations of shocks (rather than in anticipation of them), relies on the common notion that the timing of the shocks within a short period of time may be as good as random, which has been exploited for identification in a variety of settings. As such, our use of one-dimensional matching on timing to construct counterfactuals may also be useful for analyzing other economic events whose particular timing is likely unpredictable.[3]

With these data and design, we provide visually-clear estimates of individuals' labor supply responses to spousal mortality and health shocks, and we additionally exploit the richness of the data to analyze the potential mechanisms that may underlie these responses. Overall, we find significant and persistent increases in spousal labor supply when income losses are large and households lack adequate formal insurance. While a variety of potential forces may be at play in the different shocks that we consider, the findings are all consistent with family labor supply as a self-insurance mechanism.

We begin with the focus of our study, the extreme shock of the death of a spouse, which can lead to significant and permanent income losses. We find immediate increases in survivors' labor supply following their spouse's death, which persist through the duration of our time frame. By the fourth year after the shock, these responses amount to average increases of 7.6% in survivors' labor force participation and 6.8% in annual labor income.

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

[3] We know of several concurrent studies in different settings that have found it useful in practice. These studies follow an earlier version of this paper (which has been previously circulated as NBER Working Paper No. 21352) and include applications such as analyzing the effects of inheritance on wealth accumulation and inequality (Nekoei and Seim 2017; Martinello 2017), studying the effects of Social Security field office closings (Deshpande and Li 2017), and our own work on family spillovers in health behaviors (Fadlon and Nielsen 2017).

income shocks due to the death of a spouse, and may therefore have greater need for self-insurance through labor supply. In particular, we use different strategies and sources of heterogeneity to show that the mean increase in labor supply is fully 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. Notably, widowers, who tend to be primary earners and need to financially support one fewer person when losing their wives, if anything 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 to a 10.1% increase in their annual earnings. Importantly, though the exposure to risk is highly correlated with gender, female and male survivors exhibit similar sensitivity to comparable changes in household income. Further bolstering the plausibility of the self-insurance mechanism, we provide evidence that a fall in the cost of supplying labor following the death of a spouse does not appear to be an operative alternative explanation for the average responses we document.

In contrast to mortality shocks, non-fatal health shocks are well-insured in our setting through social and private insurance. Studying households in which one member has experienced a non-fatal heart attack or stroke, we first show that the earnings of the sick individuals drop by 19% after the shock, comparable to the findings of Meyer and Mok (2013) and Dobkin, Finkelstein, Kluender, and Notowidigdo (2017) in the US case. However, we then show that the average decline in these households' post-transfer income is only 3.3%. Consistent with this lack of a significant average income drop, which suggests there is no substantial need for self-insurance on average, there are no notable changes in spouses' labor supply. Yet while the average decline in household income is negligible, there is still cross-household variation in income replacement rates, and we find substantial heterogeneity in spousal labor supply responses as a function of this variation. Though fatal and non-fatal shocks differ in many aspects (such as in their impact on the composition of the household), our results suggest that self-insurance is a key motive for spousal labor supply responses to the financial aspects of both types of shocks.

As we described so far, our findings on family labor supply responses to adverse health shocks have direct implications for positive models of household self-insurance behavior over the life cycle. To additionally offer a normative interpretation of our results, we conclude with a brief discussion of their *qualitative* welfare implications. To do so, we adopt a stylized version of the model from Fadlon and Nielsen (2016), which illustrates that the extent to which households self-insure against *realized* income losses using spousal labor supply is directly related to the degree to which they lack formal insurance and would gain from more generous government benefits. Exploiting this logic, we argue that it may be welfare improving to let survivors benefits depend on age, since the relative increase in survivors' labor supply is substantially larger for older widows whose lower labor force attachment makes them more financially vulnerable. Also, since survivors' labor supply is strongly increasing in the share of the household's income that the deceased had earned, there is an argument for survivors benefits to depend on the deceased spouse's work history. This is consistent with Persson (2015) who finds higher valuation of survivors insurance by spouses with divergent levels of earned income, as manifested by their higher marriage rates when the marriage contract includes such a scheme. Additionally, as our findings indicate a similar pattern of heterogeneity in responses to non-fatal spousal health shocks, disability benefits may also be more efficiently distributed if made dependent on the disabled spouse's work history. Evidently, all of these features characterize the current large survivors and disability insurance schemes within the Social Security system in the US.

This paper relates to two main strands of the literature. First, significant research studies the effects of adverse health. The majority of this work focuses on the impacts on *own* labor market outcomes and includes studies that use survey data (such as Gertler and Gruber 2002, Charles 2003, Gallipoli and Turner 2011, Chung 2013, Meyer and Mok 2013, and Dobkin et al. 2017) as well as larger-scale studies in countries where administrative data that link health and own labor market outcomes are available (such as Lundborg et al. 2011, Halla and Zweimüller 2013, Pohl et al. 2013, and Gupta et al. 2015). However, despite the premise that households operate as tight economic units, there is much less (and mixed) direct evidence regarding the effects of non-fatal health shocks on *family* labor supply; and, to the best of our knowledge, there has been no work in the last few decades on the impacts of fatal shocks. The existing household-level studies on non-fatal shocks, which have generally utilized event-study type analyses and comparisons using unaffected households, have commonly used survey data;[4] and by comparison, we know of very little work that, similar to ours, uses rich administrative data and timing of objectively-identified health events.[5]

The combination of the data and research design that we use allows us to contribute to this literature by offering estimates for household labor supply responses to both fatal and non-fatal health shocks in the short- and medium-run. As the dataset includes comprehensive records of the different components of household income and wealth, it also lets us observe households' overall degree of formal insurance and income loss and analyze a variety of behavioral margins. This distinct aspect of the data provides us with the opportunity to additionally study the potential mechanisms that may underlie the labor supply responses to spousal shocks that we find.

Our paper also contributes to the numerous past empirical studies that have analyzed spousal labor supply responses to individuals' wage and unemployment shocks (what is known as the "added worker effect"). While spousal labor supply has been commonly modeled as an important self-insurance mechanism (e.g., Ashenfelter 1980; Heckman and Macurdy 1980; Lundberg 1985), this prior empirical work has been largely unable to find evidence of significant responses to *temporary*

[4] Recent related studies are Charles (1999), Jiménez-Martín et al. (1999), Johnson and Favreault (2001), Gertler and Gruber (2002), Coile (2004), Siegel (2006), Coe and Van Houtven (2010), Gallipoli and Turner (2011), Hollenbeak et al. (2011), Meyer and Mok (2013), Braakmann (2014), and Dobkin et al. (2017). Older studies include Parsons (1977), Berger (1983), Berger and Fleisher (1984), and Haurin (1989). To the best of our knowledge, only Berger (1983) and Haurin (1989) consider the case of spousal death.

[5] Indeed, we have been able to identify only two such papers, which rely on comparisons of affected households and observably-similar unaffected households over time. García-Gómez et al. (2013) study acute hospitalizations in the Netherlands, and in a concurrent paper Jeon and Pohl (2017) study cancer diagnoses in Canada. In related work using Swedish data, Nahum (2007) studies responses to spousal sickness absence from work as the identified shock, where the reference group is unaffected individuals with no sickness absence or those with short absence periods.

spousal unemployment.[6] A 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 life-time income and are already sufficiently insured through formal social insurance (Heckman and Macurdy 1980; Cullen and Gruber 2000). Consistent with this explanation, Stephens (2002), Autor et al. (2015), and Blundell et al. (2016) find that spousal labor supply is an important consumption insurance device against significant and permanent income losses (due to job displacements, wage shocks, or disability insurance denials). Our findings of no spousal responses to well-insured non-fatal shocks are consistent with previous studies, and we provide new evidence on the important role of spousal labor supply in offsetting permanent household income losses in the context of fatal shocks.

The remainder of the paper is organized as follows. We proceed with Section 2 that sets the conceptual framework for the empirical analysis by theoretically illustrating with a stylized model the self-insurance role of spousal labor supply. Then, prior to our empirical analysis, Section 3 outlines the institutional environment in Denmark and the data sources. In Section 4 we describe the empirical research design that we use 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 fatal and non-fatal spousal health shocks and analyze the heterogeneity in these responses to study their underlying mechanisms. In Section 6, we use the framework from Section 2 to briefly discuss the qualitative welfare implications of our findings. Section 7 concludes.

2 A Stylized Framework of Family Labor Supply

To set the conceptual framework for our empirical analysis and to later describe its potential normative implications, we start by analyzing a static model of household labor supply decisions. The purpose of this section is to formalize 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.[7] We show that the relative increase in spousal labor supply in response to shocks grows with the income loss, so this increase can reveal the extent to which the household needs to self-insure. While our empirical analysis includes both extensive and intensive margin labor supply responses, the stylized model of this section focuses on the intensive margin. As the goal of the model is merely to provide a simple conceptual economic framework, this modeling choice is motivated by simplicity and transparency.[8]

[6] See, for example, Heckman and Macurdy (1980, 1982), Lundberg (1985), Maloney (1987, 1991), Gruber and Cullen (1996), and Spletzer (1997).

[7] Of course, there are important non-financial linkages across spouses which we abstract from here since we want to highlight the financial channels that link household members. We return to mechanisms other than income loss that can drive spousal labor supply in response to shocks later in our empirical analysis.

[8] Specifically, the model below—including the setup and preference specification—provides the simplest possible framework to demonstrate the insurance role of spousal labor supply. It also provides a simple and intuitive formula that translates household labor supply responses to implications for the design of social insurance. In Fadlon and Nielsen (2016) we additionally study the participation decision counterpart of this section's model, and extensively analyze and discuss important generalizations to

Setup. We study labor supply decisions of a two-person household, which consists of individuals 1 and 2. We consider a world with two states of nature: a "good" state, state g, in which member 1 works; and a "bad" state, state b, in which member 1 experiences a shock—e.g., a severe health shock—and drops out of the labor force. We employ this extreme assumption regarding member 1's labor supply for simplification, but any shock that leads to some degree of exogenous decline in this member's labor earnings or income from other sources can be readily analyzed within the same framework. 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.

Household Budget Constraint. Denote by c_i^s and l_i^s the individual consumption and labor supply of member i in state s, respectively. Let A^s denote the household's state-contingent wealth and non-labor income—including transfers from any source of individually-purchased or employer-provided private insurance, out-of-pocket expenses (such as medical bills), and potentially transfers from pension schemes and retirement accounts and from relatives (in the simple case when those are not endogenous to spousal labor supply). We denote by $\bar{z}_i^s(l_i^s)$ i's net-of-tax labor income in state s, so that with a wage rate of w_i and a linear labor-income tax rate of τ_i we have $\bar{z}_i^s(l_i^s) = z_i^s \times (1 - \tau_i)$, where $z_i^s \equiv w_i l_i^s$ are gross earnings. Finally, let B^s represent benefits from the government in state s. These are not means-tested in the simple framework of this section, but the model can accommodate means-testing with some additional complexity. With this notation, the household's overall income in state s, y^s, satisfies $y^s = A^s + \bar{z}_1^s(l_1^s) + \bar{z}_2^s(l_2^s) + B^s$.

Preferences. Let $U^s(c_1^s, c_2^s; l_1^s, l_2^s)$ represent the household's utility as a function of consumption and labor supply of each member in each state. For simplicity, we assume that $U^s(c_1^s, c_2^s; l_1^s, l_2^s) = u_1^s(c_1^s) - v_1^s(l_1^s) + u_2^s(c_2^s) - v_2^s(l_2^s)$, where in each state $u_i^s(c_i^s)$ is member i's utility from consumption and $v_i^s(l_i^s)$ represents member i's disutility from labor (including the utility loss from direct work costs and the opportunity costs of lost home production). We employ the normalization $u_1^s(0) = v_1^s(0) = 0$. This lets the model incorporate the case in which the bad state is a *fatal* health shock (in which $c_1^b = l_1^b = 0$), so that the household's preferences reduce to the utility from member 2's allocation: $u_2^s(c_2^s) - v_2^s(l_2^s)$. Additionally, we assume that the consumption utility and the labor disutility functions are well-behaved—i.e., that $u_i^{s\prime}(c_i^s) > 0$, $u_i^{s\prime\prime}(c_i^s) < 0$, $v_i^{s\prime}(l_i^s) > 0$, and $v_i^{s\prime\prime}(l_i^s) > 0$.

Household Behavior. In our illustrative static model of this section, the household consumes its entire disposable income in each state of nature. Hence, the household's choices reduce to the labor supply and consumption allocation decisions. Formally, in each state s the household solves the problem:

$$\max\ U^s(c_1^s, c_2^s; l_1^s, l_2^s) \text{ s.t. } c_1^s + c_2^s = y^s.$$

the highly-stylized framework. These include a dynamic life-cycle model, general choice variables (that among other decisions encompass savings, life-insurance purchases, and informal insurance arrangements), alternative assumptions about the household's preference structure (with an explicit analysis of different types of state dependence and preference complementarities), different approaches to modeling the household's behavior (i.e., collective or unitary), means-testing in government transfers, as well as the presence of household public goods and economies of scale in the household's consumption technology.

When both spouses are alive, optimal consumption allocation across the two must satisfy $u_1^s{}'(c_1^s) = u_2^s{}'(c_2^s)$. Additionally, in each state of nature, the first-order condition with respect to the labor supply choice of the indirectly-affected member 2 satisfies $u_2^s{}'(c_2^s) = \frac{v_2^s{}'(l_2^s)}{w_2(1-\tau_2)}$. A similar condition holds for the labor supply of member 1, but only in state g.

Spousal Labor Supply as Self-Insurance. At this point it is straightforward to demonstrate the self-insurance role of spousal labor supply responses to shocks. Define y_{-2}^s as the household's resources excluding those directly attributed to 2's labor supply decision—i.e., $y_{-2}^s \equiv A^s + \bar{z}_1^s(l_1^s) + B^s$—so that the (exogenous) income loss from the shock is $L \equiv y_{-2}^g - y_{-2}^b$ (i.e., the gap in the spouse's unearned income across the two states). The household optimization conditions imply that the spouse's labor supply response to the shock, $\frac{l_2^b}{l_2^g} - 1$, is greater whenever the imposed income loss L is larger. That is,

$$\frac{\partial\left(\frac{l_2^b}{l_2^g} - 1\right)}{\partial L} = -\frac{u_i^b{}''(c_i^b)\frac{\partial c_i^b}{\partial A^b}}{l_2^g v_2^b{}''(l_2^b)/w_2(1-\tau_2)} > 0, \tag{1}$$

when consumption in the bad state is a normal good ($\frac{\partial c_i^b}{\partial A^b} > 0$). Intuitively, when individuals experience fatal or non-fatal health events that lead to shocks to the household's income (from foregone earnings or declines in other sources of income), their spouses can compensate for the associated income loss by increasing their own labor supply. Since the relative increase in spousal labor supply in response to shocks grows with the income loss, it can reveal the extent to which the household lacks formal insurance and needs to self-insure. With this simple comparative statics result at hand, which theoretically illustrates the insurance role of spousal labor supply in response to household shocks, we now turn to the empirical analysis of the impact of fatal and non-fatal spousal health shocks on individuals' labor supply.

3 Institutional Background and Data

To study labor supply responses to severe spousal health shocks we leverage rich administrative full-population data from Denmark. Compared to other countries, the Danish setting is unique in providing large-scale register-based data on both health and labor market outcomes, combined with spousal linkages. As such, it is a well-suited setting for the purpose of our study. In this section, we describe the Danish insurance environment (both social and private) as it relates to sick individuals and surviving spouses, and we list our data sources.

3.1 Institutional Setting

It is useful to distinguish between two types of insurance: health insurance (coverage of medical care) and income insurance (coverage of 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.

In Denmark, income insurance against severe health shocks and the death of a spouse consists of three main components that are generally typical of systems in developed countries: permanent Social Disability Insurance, privately-purchased insurance policies, and other indirect social insurance programs (such as early retirement and old-age pensions). In the rest of this section, we provide an overall description of the features and benefit levels for each of these components (which we also summarize in Appendix Table 7 for convenience). Later, within our empirical analysis, we exploit the data to show how the income streams provided to households from different sources (and their combination) behave in practice in our analysis sample.

Individuals who experience a shock that leads to a substantial reduction in their ability to work (of at least 50% as determined by the assigned evaluator) can apply at the municipality level for Social Disability Insurance (Social DI) benefits. An approved application will provide benefits permanently, which in 2000, for example, amounted to DKK 72,100 ($9,000) per year for married or cohabiting individuals and DKK 98,700 ($12,300) for single individuals (with potential supplements that depend on factors such as disability severity).

The basic eligibility criterion for this program is a prolonged need for support that is presumed to last until the transition into the old-age pension. Importantly for our analysis, since 1984 the Danish Social DI has a broad social insurance scope. That is, it can be awarded to individuals who prove that they are unable to engage in substantial gainful activity either for medical or for non-medical (vaguely defined) social reasons.[9] As an example for such non-medical reasons, Social DI benefits could be awarded to survivors who are out of the labor force and, upon their spouse's death, are deemed unfit for employment or training programs, e.g., due to their age (for additional details see Haanes-Olsen 1987 and Bingley et al. 2011). Note that within the context of social reasons, individuals will be automatically considered ineligible if their annual earnings in the years just before their application exceeded a certain threshold (which in 2000 was DKK 148,000 for married applicants and DKK 98,700 for singles). As there is no explicit survivors insurance program in Denmark, Social DI effectively acts as the relevant social insurance program that may support surviving spouses who, at the end of the cumbersome application process (which can take anywhere between 1 and 14 months), are determined unable to maintain their standard of living on their own. Indeed, we find (and document later in the paper) significant increases in the receipt rate of Social DI by survivors in the year their spouses die (which are predominantly driven by female survivors who did not work prior to the shock). Correspondingly, we henceforth refer to Social DI in the context of spousal mortality shocks as social survivors benefits.

While Social DI and its survivors 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. These include any contact with the applicant,

[9] 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).

the preparation of the application, and the 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-by-year variation in the awarding of the survivors benefits component of the program later in the paper. Specifically, we test the hypothesis that the generosity of this alternative insurance mechanism against fatal spousal shocks substitutes for self-insurance through labor supply.

Another source of income to households that experience health shocks or in which a member dies is payments from employer-based insurance policies, which have recently become standard in labor-market pension plans. Since 1993, most sectors covered by collective agreements (75% of the labor force) have introduced mandatory pension plans, which may include components of life insurance or insurance against specific health events. These schemes pay out a lump-sum to sick workers, as long as they make contributions to the pension plan, or to the surviving spouse in case the plan member dies, at rates that are set by the individual pension funds. In addition to employer-sponsored private insurance, and subject to health screenings, individuals may purchase similar insurance policies in the private non-group market. In our sample period, which spans 1980-2011 and in which the average year of spousal death is around 1995, the life-insurance coverage rate is generally low, since life-insurance holdings experienced significant increases only more recently with the gradual expansion of labor market pension schemes. As payouts from these private (group or non-group) insurance policies are observable in our data as part of households' liquid wealth, we later gauge the effective coverage rate within our sample by providing an analysis of actual changes in household liquid wealth balances around spousal death (both in terms of wealth levels and in terms of the fraction of households for whom shocks lead to wealth increases).

Note that despite the private market for life insurance in Denmark, there is still an important rationale for government intervention in the context of mortality shocks since unhealthy and older Danish households are largely uncovered through the private market. First, with health screenings required for purchasing life-insurance products in the Danish non-group market (which include answering health status and behaviors questionnaires and even undergoing medical exams), applications by unhealthy or older households are occasionally rejected.[10] Second, it is common in both group and non-group markets that even when life-insurance products are purchased by younger and healthy households, the coverage sharply declines with age (see a concrete example of such a group-market policy in Appendix Table 7). The combination of these features of the private insurance market and the lack of a universal-coverage social survivors insurance scheme leaves older Danish surviving spouses more vulnerable to the financial shocks imposed by spousal death. In our

[10] These rejections by insurance companies can be explained by private information that is held by rejected households, which provides a leading rationale for government intervention in our setting (Hendren 2013). Hendren (2013) also provides relevant evidence in the context of life insurance markets in the US.

setting, while survivors in older households experience smaller losses through foregone labor income (as their deceased spouses' earnings were lower), they are still exposed to substantial financial losses through the deceased spouses' non-labor income (including employer-based pension payments and benefits from government programs). This point will become relevant when we discuss the welfare implications of our findings in Section 6.

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 stages (within the range of eligible 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 (of 10 years before 1992 which has gradually increased 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 that amounted to roughly DKK 135,000 ($16,875) in 2000. At the full-retirement age of 67 (or 65 for those born after July 1st, 1939) all residents become eligible for the Old-Age Pension (OAP), which provides annuities that in 2000 amounted to DKK 72,100 ($9,000) for married individuals and DKK 98,700 ($12,300) for single individuals (similar to the benefit levels paid to Social DI beneficiaries). Note that DI and OAP are different components of the same social insurance program of Social Pensions, similar to Social Security in the US, and that Social DI recipients automatically transition into the Old-Age Pension program at their full-retirement age. Benefits through both Social DI and OAP are income-tested, with income thresholds and benefit reduction rates as depicted in Appendix Figure 9.[11]

3.2 Data Sources

We have merged several Danish registers that include individual-level records with household linkages that allow us to match spouses and cohabiting partners from 1980 to 2011. Doing so, our analysis uses long panels of detailed administrative datasets for the universe of Danish households with a wide range of objective measures of families' health and economic outcomes.

Health Data. To identify fatal and severe non-fatal health shocks we use two complementary datasets. Our first dataset is the *Cause of Death Registry*, which includes death dates and causes. Our second dataset is the *National Patient Registry*, which covers all hospitalization records (from both private and public hospitals) with exact timing and detailed diagnoses (using the International Statistical Classification of Diseases and Related Health Problems [ICD] system). The health shocks that we focus on are heart attacks and strokes, which are commonly-studied pervasive health events that are both sudden and severe (Chandra and Staiger 2007; Doyle 2011).

Economic Data. The economic data that we use cover years 1980-2011 and include comprehensive information on all sources of family income: earnings, government transfers from any program

[11] An additional small government-mandated pension scheme (for all wage earners in Denmark) that supplements the OAP and includes a limited amount of a one-time transfer to survivors is the ATP program. As this program represents a very small fraction of government transfers to older households and surviving spouses, we postpone its description to Appendix Table 7.

(including old-age pensions, disability insurance, welfare benefits, housing assistance, and unemployment benefits), payouts from retirement savings accounts, annuity payouts from insurance companies, and capital income. Importantly, we also observe third-party reported liquid wealth balances (measured annually on December 31st) from 1984-2011. Among other measures, these include bank-account balances and lump-sum transfers from insurance companies. In our main analysis sample of spousal mortality shocks, the baseline net asset *stock* of the median household amounts to only DKK 11,179 or $1,397 (with a *single* increase of $2,152 following the event) while the baseline median household-level income *flow* is DKK 253,843 or $31,730 annually (with an approximate decline of $10,739 *each year* after the event). Also, as we show later in the paper, any increase in net wealth around spousal death (through insurance payouts, etc) is attributable to only a very small share of households, as merely 5% of the affected households in our sample experience some growth in net wealth caused by the shock (as compared to the counterfactual). Therefore, our analysis of labor supply responses and their heterogeneity focuses on income losses. However, we use the wealth data in our robustness checks to account for life insurance payments and other potential changes in household wealth.

Our final dataset is the *Integrated Database for Labor Market Research*, which includes demographic variables for the entire population as well as administrative 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 describe our analysis sample and its summary statistics at the end of the next section after we present the research design and explain how we construct the treatment and control groups.

4 Research Design

The goal of our empirical analysis is to identify the dynamic causal effects of fatal and severe non-fatal spousal health shocks on individuals' labor supply. In this section we describe the empirical strategy that we use to overcome the selection challenges inherent in the identification of these effects. We then describe our analysis sample of treatment and control groups and report their summary statistics.

4.1 Quasi-Experiment

The ideal experiment for identifying the short- and medium-run effects of spousal shocks would randomly assign shocks to households and track labor supply responses over time. Therefore, we need to compare the ex-post responses to shocks of affected households to a counterfactual behavior of (hypothetical) ex-ante similar unaffected households. This requires comparing households with the same expectations over the distribution of future paths, but with different realizations, to isolate the unanticipated component of the shock. The access to three decades of administrative panel data on the universe of Danish households allows us to employ a quasi-experimental research design that mimics this ideal experiment, by exploiting the potential randomness of the timing of a severe (fatal or non-fatal) health shock 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. Specifically, we construct counterfactuals to affected households using households from the same cohorts that experience the same shock but a few years in the future. Then, we recover the treatment effect by performing traditional event studies for these two experimental groups and combining them into a straightforward dynamic difference-in-differences estimator. Before formally describing the 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 specific treatment group of married and cohabiting 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.[12]

In Panel A of Figure 1 we plot the outcome for this treatment group over time, and compare it to the time trend of the outcome for married and cohabiting individuals from the same cohorts who have not experienced this shock in our sample period. Inspection of this figure reveals considerably different behavioral patterns and visible non-parallel trends prior to 1995 across the two groups. The groups' divergent pre-trends persist even after we control flexibly for key variables, specifically, age, gender, and education (see Appendix Table 8). This motivates the consideration of alternative households, other than those who do not experience shocks, as potential control groups for the construction of the treatment group's counterfactual behavior in the absence of the shock.

We therefore proceed by looking only at affected households. Specifically, 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). Notably, 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 the intuition of comparability of households that experience shocks closer in time, 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 $\triangle = 5$.

[12] Illustrative examples that use other key household outcomes are available from the authors on request.

Our estimation strategy generalizes this example by aggregating different calendar years. Simply put, the 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 but in year $\tau + \triangle$. We identify the treatment effect purely from the change in the differences in outcomes (i.e., the difference-in-differences) across the two groups over time. By construction, the research design matches households (only) on the year the shock occurred, so it mechanically nets out calendar year effects. However, on top of that and *without* directly matching households on any other dimension, the design constructs in our setting experimental groups that are also very similar in the key dimension of age (as we show below). Doing so, it effectively nets out life-cycle effects, which are a main identification concern in the context of family labor supply.

The trade-off in the choice of $\triangle$, which captures the main limitation of the 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., those who experienced the shock in 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, since for each chosen $\triangle$ the estimation strategy provides estimates for 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 pre-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 and found that local perturbations to $\triangle$ provide very similar results.[13]

Formal Description of the Design and Estimator. Similar to common practice (for example, in the use of matching estimators; see, e.g., Imbens and Wooldridge 2009), our estimation procedure can be broken down into two steps. The first step constructs our treatment and control groups and, in the second step, estimation and inference are conducted using traditional methods. We describe the two steps successively.

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 individuals' households as the treatment group and divide them into sub-groups indexed by the year in which the shock was experienced, $\tau \in [\tau_1, \tau_2]$. 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 of individuals from the same cohort group Ω who experienced the same

[13] In some applications (e.g., with smaller samples or shorter panels), the researcher may wish to include in the control group households that experienced the same shocks less than $\triangle$ periods apart for improved efficiency. That is, for identifying the 3-year effect one can include both households that experience the shock 5 years later and 4 years later, for identifying the 2-year effect one can include households that experienced the shock 5, 4, and 3 years later, and so on. However, in our application, which has a large number of households, this is not needed (nor does it alter the results), and using a single $\triangle$ has the advantage of providing a transparent analysis of treatment and control groups that are both fully balanced for the entire analysis window.

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, i.e., $t = year - \tau$ (where, by construction, their actual shock occurs at $t = \triangle$).[14]

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 prior to the shock which we denote by p (for "prior"). For any period $n > 0$, the treatment effect γ_n can be simply recovered by the difference-in-differences estimator

$$\gamma_n \equiv \left(y_n^T - y_n^C\right) - \left(y_p^T - y_p^C\right). \tag{2}$$

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 $n < \triangle$ (since the control group becomes "treated" at $t = \triangle$).

The identifying assumption is that, absent the realization of the shock, the outcomes of the treatment and control groups would run parallel. 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. Similar timing-based variation has been exploited for identification in numerous previous papers within a variety of settings.[15] 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. We consistently show throughout the analysis that there are virtually no differential changes in the trends of the treatment and control groups before period 0. This validates the design and alleviates concerns that the groups may differ by, for example, their expectations over the *particular* year of the shock within our chosen five-year window of $\triangle$.[16]

It is worth noting that the research design does *not* preclude behavioral adjustments in expectation of a shock among treated households; nor do our results imply there are no such adjustments

[14] The same household can appear both in the treatment group and in the control group, but is never used as a control to itself. For example, if treated households that experienced a shock in 1990 (who are matched with households that experienced a shock in 1995 as controls) are included also in the control group, it is only since households that experience a shock in 1985 are included in the treatment group as well. We repeated our main analysis using treatment and control groups that do not overlap, either by including in the treatment group (and matching them with the corresponding control group) households that experience shocks in every other year, or by randomizing overlapping households to only one experimental group. The results remain similar (both qualitatively and quantitatively) and are available from the authors on request.

[15] Among many others, these include Ruhm (1991), Grogger (1995), Hilger (2016), and Persson and Rossin-Slater (2016) in the context of household shocks, and papers such as Guryan (2004) and Bailey and Goodman-Bacon (2015) in the context of program rollout.

[16] Conceptually, as long as there is no perfect foresight we can use the design with an 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 that 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.

in practice. Indeed, the differential pre-trends that we have seen across affected and unaffected households may be driven by exactly this type of anticipatory responses and these groups' diverging expectations. However, since our empirical target is *ex-post* responses to the realization of the shock, our aim has been to provide a control group with non-differential expectations. Thus, the parallel pre-trends across our constructed experimental groups do not mean that affected households do not exhibit anticipatory effects. Rather, the non-differential pre-trends signal that the research design achieves its goal: conducting comparisons across closely similar treatment and control households that hold comparable expectations.

Estimating Equation for Average Effects. Much of our analysis graphically tracks the simple dynamic difference-in-differences estimator of equation (2) to study the evolution of household responses. To quantify the mean treatment effects, we estimate the regression counterpart of this estimator, averaged over the years after the shock. These regressions are also useful for explicitly reporting statistical significance (where we cluster the standard errors at the household by experimental-group level) and for accounting for controls in robustness checks. Our baseline estimating equation is of the known difference-in-differences form:

$$y_{i,t} = \alpha_i + \beta post_{i,t} + \gamma treat_i \times post_{i,t} + \delta X_{i,t} + \varepsilon_{i,t}. \tag{3}$$

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 (time-variant) controls; and α_i is a household fixed effect (which absorbs any time-invariant characteristic including the "main effect" of $treat_i$). The parameter γ represents the average causal effect of spousal shocks on household outcomes.

4.2 Analysis Sample and Summary Statistics

Starting from the universe of married and cohabiting couples, our sample includes all households that experience shocks from year 1985 to 2011 and in which both spouses were between ages 45 and 80 in the year of the (actual or placebo) shock. Our main sample is comprised of all households in which one spouse experienced a fatal shock and includes 310,720 households in the treatment group and 409,190 households in the control group. Our secondary sample of non-fatal severe health shocks is comprised of all households in which one spouse experienced a heart attack or a stroke (for the first time) and survived for at least three years.[17] These health shocks are commonly studied as their timing within a short period of time is likely unpredictable (Chandra and Staiger 2007; Doyle 2011; WHO 2014). The average age of spouses precisely at the time of these cardiovascular health shocks is just over 60 (60.67), and recall that most individuals become eligible for early retirement benefits when they turn 60. Therefore, we focus in this second sample on households with both spouses under 60 to ensure that the results we document are driven only by the health shocks

[17] Our choice of the number of years of survival was motivated by balancing between the sample size (that shrinks when we condition on more years of survival) and the horizon of the analysis (that shrinks when we condition on less years of survival). Perturbations to this number do not alter the qualitative results. This analysis is available from the authors on request.

and not by eligibility for early retirement benefits.[18] The qualitative results and conclusions do not change, however, when we look at the unconstrained sample. Our sample of non-fatal shocks includes 37,437 households in the treatment group and 54,887 households in the control group. As heart attacks and strokes are among the leading causes of death in the developed world (WHO 2014), fatal cardiovascular events provide us with a large sub-sample of deaths that likely come as a surprise; which we additionally analyze in the estimation of average household responses to fatal shocks to validate that the estimation strategy isolates the ex-post responses to (unanticipated) shock realizations.

Panel A of Appendix Table 1 displays key summary statistics for the analysis samples and reveals an advantage of the research design. Specifically, the table shows the close comparability of the year of observation and of the age of spouses across the treatment and control groups. In the baseline period (chosen to be $t = -2$), the surviving spouse in the treatment group is observed on average in year 1993 at age 62.86 and the spouse in the control group is observed on average 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. Likewise, in the sample of non-fatal health shocks, at the baseline period the spouse is on average 45.7 years old in the treatment group and 45.3 years old in the control group, where the mean calendar year is around 1992 for both groups.[19]

5 Family Labor Supply Responses to Severe Health Shocks

In this section, we present our primary analysis of the impact of fatal and severe non-fatal health shocks on spousal labor supply. For extensive margin responses we analyze labor force participation, defined as having any positive level of annual earnings; for intensive margin responses we analyze annual earnings, and supplement the analysis with administrative measures for full-time and part-time employment. We begin with the focus of our study, the extreme shock of spousal death, which can lead to large and permanent income losses. We study the average labor supply responses of surviving spouses, 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 (1), we analyze how survivors' behavior varies by the degree of income loss their spouse's death imposes and by the extent of coverage through survivors benefits, to investigate self-insurance as a mechanism for family labor supply responses. We also analyze alternative potential mechanisms using a simple test that

[18] Eligibility for early retirement benefits leads to a sharp decline in labor force participation at age 60 (see an illustration in Appendix Figure 1). Since in the specific case of non-fatal shocks (within our setting and time frame) the average age at $t = 0$ is close to this age threshold, one might worry that in figures that depict *raw* means even very small age differences across the experimental groups may display small spurious responses to the shock. The purpose of our age restriction is to address such concerns.

[19] We additionally report in Appendix Table 1 the means of main labor supply outcomes at the baseline for completeness. Note that since comparability requires similar trends and not similar levels (like in any other difference-in-differences type research design), the slightly higher levels of participation and earnings for the control group do not pose a direct threat to the validity of the design.

aims to assess the extent to which survivors' willingness to work may change in response to the shock. Then, we study family labor supply responses to our second set of shocks, severe non-fatal health shocks; specifically, heart attacks and strokes. We show that in our setting the resulting income losses from the foregone earnings of the sick spouse are formally well-insured, primarily through disability benefits. Accordingly, in the context of non-fatal shocks, we do not expect spousal labor supply responses from self-insurance motives on average.

5.1 Spousal Labor Supply Responses to Fatal Health Shocks

5.1.1 Mean Responses

Panels A and B of Figure 2 plot the average labor supply responses of spouses in our overall sample of fatal health shocks. The structure of these and subsequent figures is as follows. 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, from which the treatment effect is identified, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome (in period $t = -2$).[20] This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

These panels first provide a visual verification of parallel trends across the treatment and control groups prior to period 0. Then, analyzing the effect of the shock, Panel A reveals an immediate increase in labor force participation following the death of a spouse. By the fourth year after the shock, the increase in the surviving spouses' participation amounts to 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 (where we include zeros for those who do not work at all). Appendix Figure 2 repeats this analysis for a sub-sample of these survivors, whose spouses experienced a heart attack or a stroke for the first time and died within the same calendar year. This allows us to focus on deaths whose particular timing is plausibly unexpected so that they are more likely to 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 overall sample of fatal shocks; further validating that the estimation strategy isolates the ex-post responses to shock realizations.[21]

With significant disparities in baseline participation rates and labor income, men and women may face substantially different financial distress when their spouse dies and, therefore, may respond differently to this shock. Indeed, Panels C and D of Figure 2 reveal clear differences in the responses of widowers (whose wife dies) and widows (whose husband dies). While on average widowers do

[20] We choose $t = -2$ as the baseline period in our figures to verify there are no differential trend breaks just before the shock (in $t = -1$) across the treatment and control groups. As will become visually clear in the figures, the results are very similar when we use any of the pre-shock periods (-5 to -1) as a baseline year.

[21] Similar results were found for the small sample of accident-related deaths (available on request).

not change their labor force participation when their wife dies (and, if anything, slightly decrease their annual earnings), widows immediately and significantly increase their labor supply following the death of 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 considerable increase of 11.3% in their labor force participation with a corresponding rise of 10.1% in their annual earnings.

This differential response suggests that female survivors may have greater need to self-insure through labor supply and that they may experience greater income losses when their spouse dies 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 on the part of survivors in order to capture the income loss directly attributable to their spouse's death. To do so, we plot in Panel A of Figure 3 the household's overall income, holding the surviving spouse's earnings and social benefits at their pre-shock level.[22]

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 fatal health shock, the household's composition changes so that insuring the consumption of surviving spouses as singles does not require the entire pre-shock level of household income. At the same time, potential economies of scale within the household can make half of the household's income before the shock insufficient for survivors to maintain their pre-shock utility levels after the shock (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 0.5 and 1 and is commonly referred to as the adult "equivalence scale". Some commonly used scales are the modified OECD equivalence scale of 0.67 and the square-root scale of 0.71. Hence, one would expect surviving spouses to broadly compensate for income declines with respect to this general benchmark, such that decreases in household income on the order of 29-33 pp would not require self-insurance through labor supply.[23]

Panel A of Figure 3 shows that widowers, who do not change their labor force participation on average, experience an overall decline of 32 pp in household income. However, as suggested by their different labor supply responses, widows experience a significant additional relative loss of 8 pp compared to widowers, so that the decrease in household income is 25% larger for female survivors. To study the actual (rather than the potential) change in household income, Panel B of Figure 3 takes into account the surviving spouses' labor supply responses and any change in the social benefits they may receive. The figure shows that widowers experience an actual decline of 31 pp 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 pp. Overall, widows'

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

[23] The relevant equivalence scales that we mention here as benchmarks for gauging magnitudes 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.

labor supply responses account for 22% of the 5 pp shrinkage (from 40 to 35 pp) in their potential income loss.[24] We return to additional male-female comparisons in Section 5.1.2, when we conduct a within-gender analysis of sensitivity to household income loss that flexibly accounts for a wide set of observables.

Younger Households. The baseline labor supply behavior of individuals below and above age 60 substantially differs. At age 60 there is a sharp drop in participation when 80% of the labor force becomes eligible for early retirement benefits (see Appendix Figure 1). This implies that surviving spouses under age 60 have a considerably stronger attachment to the labor force and hence notably higher disposable income from own labor earnings. In turn, it raises the possibility that they may be more financially resilient when their spouse dies. Consistent with this view and with the consequential notion that their need to self-insure through labor supply may be attenuated, Panels A and B of Figure 4 reveal that working-age widows (under age 60) exhibit a much smaller relative increase in labor supply compared to the universe of widows. Their "shock elasticities" amount to an increase of only 3.3% in participation and 3.2% in annual earnings. Interestingly, widowers under 60 even 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%) were the primary earners in their households and, compared to widows, they have significantly higher baseline participation rates (0.78 compared to 0.715) and average labor income (DKK 227,560 compared to DKK 138,232). Therefore, in contrast to widows, the behavior of working-age widowers is consistent with the idea 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 labor supply behavior of younger households suggests that higher participation rates and annual earnings may effectively insure against the potential income losses imposed by fatal spousal shocks. We further test this hypothesis more directly in Section 5.1.2, where we analyze heterogeneity in survivors' responses by the level of their own pre-shock earnings.

For younger surviving spouses, the data additionally consist of administrative measures for full-time and part-time employment (as described in the data section). This allows us to further investigate the dynamics and intensity of spousal labor supply behavior in response to fatal health shocks. As we show in Appendix Figure 3, 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 death. These short-term transitions stabilize thereafter so that the operative decision margin in the longer-run 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 (3). We present in this table the average treatment effects and their statistical significance, and verify the robustness

[24] As we described above, Panel B of Figure 3 depicts the household's *overall* income which is a composite of different sources. In Appendix Figure 6 we decompose this aggregate measure and depict how the different income sources evolve around the shock.

of our results to the inclusion of year, age, and household fixed effects.[25]

5.1.2 Heterogeneity in Responses by Income Losses and Degree of Social Insurance

We continue with further investigation of the heterogeneity in the survivors' labor supply responses across different subgroups which presents evidence that is consistent with the insurance mechanism hypothesis. 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 income insurance. We first provide a simple (more descriptive) analysis of heterogeneity across households (using cross-sectional variation with a rich set of controls). Then, we additionally exploit the municipality-by-year policy-induced variation in the awarding of social survivors benefits with the aim of isolating variation in households' income that is likely more exogenous.[26]

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.[27] 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 Figure 3, we calculate for each household the overall income (from any source) holding the surviving 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 for 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 death of a spouse. This measure is smaller whenever the deceased spouse's relative contribution to the household's income was larger. Importantly, 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 difference-in-differences model of equation (3) and estimate the following specification:

---

[25] These regressions present the medium-run effects (that are the focus of our analysis) so that $post_{i,t}$ assumes the value 1 for periods 2 to 4.

[26] Recall from Section 3.2 that our heterogeneity analysis focuses on income losses because the *one-time* change in households' net asset *stock* following the shock is dominated by the *annual* changes in households' income *flow* following the shock. We plot in Appendix Figure 7 the evolution of different moments and components of households' net wealth around the shock. The combination of Appendix Figure 7, Figure 3, and Panel F of Appendix Figure 6 provides these income vs. wealth comparisons for both medians and means. Appendix Figure 7 also highlights that any increase in net wealth caused by the shock pertains to only a very small fraction of treated households in our sample. Nonetheless, we account for baseline wealth levels and for changes in household net wealth in the regressions estimated in this section.

[27] We find similar patterns in the analysis of labor earnings, which includes both extensive and intensive margin responses (see mean results in Appendix Table 4; a more detailed analysis is available on request).

$$l_{i,t} = \alpha_i + \beta post_{i,t} + \gamma_i treat_i \times post_{i,t} + \delta X_{i,t} + \varepsilon_{i,t}, \quad (4)$$

where

$$\gamma_i = \gamma_0 + \gamma_1 rr_i + \gamma_2 Z_{i,t}.$$

In this regression, $l_{i,t}$ denotes an indicator for the labor force participation of the surviving spouse in household i at time t. We adjust the basic difference-in-differences design by allowing the treatment effect, γ_i, to vary across households and model it as a function of the household's potential replacement rate, rr_i. Our parameter of interest is γ_1, which captures the extent to which the surviving spouse's labor supply response (partially) correlates with the income loss he or she experiences. Since γ_1 can involve other dimensions of household heterogeneity beyond the income replacement rate (either preference related or insurance related), we let the treatment effect vary with additional household-level characteristics, $Z_{i,t}$, so that γ_1 would further isolate the treatment effect's partial correlation with the loss in household income. The variables that we include in $Z_{i,t}$ are age fixed effects for the surviving spouse, fixed effects for the age of the deceased at the year of death, year fixed effects, indicators for the number of children in the household and for the presence of adult and young children, as well as the surviving spouse's months of education (and its square).[28] The results are also robust to the inclusion of a quadratic in the household's net liquid wealth (which also accounts for liquidation of housing assets and changes in mortgage debt). Note that $X_{i,t}$ always includes the interaction of the variables in $Z_{i,t}$ with $post_{i,t}$, the variables in $Z_{i,t}$ that are time variant, and the interaction of these time-variant variables with $treat_i$.

Table 1 reports the results of estimating specification (4) 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 partial 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 may indeed be driven by differential income replacement rates. In addition, the estimation results reveal very similar sensitivity to comparable income losses across genders; so that re-weighting the female and male sub-samples using the regression in (4) 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) are unlikely to explain the observed differential average labor supply responses across female and male survivors, but rather their divergent income losses.

[28] Adult children are defined as being age 18 or older and young children are defined as being age 6 or younger (where the results are robust to perturbations to these specific age cutoffs). Among other dimensions, the children controls aim to allow for differential costs of supplying labor due to, e.g., the presence of young children, and differential informal insurance possibilities due to, e.g., potential support from adult children.

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 may therefore effectively be better insured.

We constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock 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 earned 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. Note that these restrictions also imply that labor supply increases within the results below for this sub-sample are mainly driven by intensive-margin responses.

We divide the remaining sample into five equal-sized groups according to the surviving spouses' pre-shock level of earnings, and plot in Panel A of Figure 5 the average labor income response (as well as its 95-percent confidence interval) against the pre-shock mean earnings for each group.[29] The figure reveals a strong negative gradient of labor supply responses with respect to the surviving spouses' own level of earnings when the shock occurs. In particular, survivors at the bottom of the labor income distribution increase their annual earnings by 7.79%, consistent with the view that it may be necessary for them to do so in order to meet their consumption needs; while those at the top of the labor income distribution decrease their earnings by 2.93%. Similar to what we found for widowers younger than 60, the behavior of these 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 these households' pre-shock labor income is composed of two earners, we need to also account for the pre-shock earnings of the deceased spouse. Hence, we divide the sample into two groups: households with low-earning deceased spouses whose pre-shock labor income fell within the bottom three quintiles of their group-specific distribution, and households with high-earning deceased spouses whose pre-shock labor income fell within the top two quintiles. Panels B and C of Figure 5 reveal that a negative 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 these relationships are robust to the inclusion of fixed

[29]To smooth out transitory wage or other labor income shocks, the pre-shock earnings of spouses are calculated as their average labor income in the sample years prior to period 0.

effects for age and year (as well as to the inclusion of a quadratic in the household's net wealth), by separately estimating the corresponding difference-in-differences specification of equation (3) for each surviving spouses' quintile. Note that merely analyzing the average earnings response in this sample would have masked the substantial heterogeneity that 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.

Before we proceed, we might wish to consider whether some of the response heterogeneity may be due to survivors' ability to respond. This could be the case when spouses in households with smaller income losses have high or close-to-full rates of participation or full-time employment before the shock. It is worth noting, however, that spouses in the sub-groups we have analyzed are generally sufficiently far from full participation or complete full-time employment (e.g., as seen in Figure 4 and Appendix Figure 3), so that they can meaningfully respond upward on the extensive or intensive margins. Also, we have found that high-earning and younger spouses, whose labor force attachment is stronger on average, actually exhibit labor supply decreases (as seen in Figures 4 and 5). In addition, we have broken down the heterogeneity analysis by household replacement rates (from Table 1) into sub-groups within which survivors have similar scope for labor supply increases. For younger households, for whom we have distinct measures of participation/part-time/full-time employment, we first constrained the sample to spouses who did not work before the shock and studied participation, for extensive margin responses; we then constrained the sample to spouses who worked part time before the shock and studied full-time work, for intensive margin responses. Likewise, for the entire sample, we ran similar specifications that separately study earnings responses by all spouses, spouses who did not work before the shock, and spouses who had positive earnings before the shock. In all these estimations we find negative correlations with household replacement rates as before (see Appendix Table 4). Finally, note that sources of variation such as the spatial variation in social insurance that we study next, which is uncorrelated with spouses' pre-shock participation, are not subject to this issue. Overall, while clearly one cannot rule out this alternative explanation, we think the evidence highlighted here reinforces self-insurance as a likely mechanism.

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 more plausibly exogenous. It also allows us to analyze whether better social insurance crowds out labor supply responses in our context of fatal spousal 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 the period prior to 1994 due to a data break in the reporting method of benefits received from Social DI, the program through which social survivors benefits are provided; and to survivors under 67, the age at which the program automatically transitions into the Old-Age Pension for the current sub-population. In addition, since the increase in the take-up of the program following the shock is attributable to females within this sample, we focus the analysis on widows (although the inclusion of widowers does not change the qualitative results). Panel A of Appendix Figure 4 displays the aggregate insurance role of Social DI for widows, whose take-up of the program increases by about 50% in the year their husbands die (primarily by survivors who did not work prior to the shock).

Recall that while the Social DI program and its survivors benefits component are state-wide schemes, they are locally administered. 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 in non-medical cases 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 consider 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, in each period t we assign to a widow in treated household i, who resided in municipality m prior to the shock, the respective mean of municipality m at time t excluding her own benefits (the "leave-one-out" mean), which we denote by $\overline{SB}_{-i,t,m}$. The variation in this instrument is displayed in Panel B of Appendix Figure 4. We estimate the following augmented difference-in-differences regression:

$$l_{i,t} = \alpha + \beta post_{i,t} + \gamma_i treat_i \times post_{i,t} + \delta X_{i,t} + \lambda treat_i + \varepsilon_{i,t}, \quad (5)$$

where

$$\gamma_i = \gamma_0 + \gamma_1 SB_{i,t}.$$

In this regression, $l_{i,t}$ denotes the labor force participation of the spouse in household i at time t, and $SB_{i,t}$ are actual social survivors benefits receipts, measured in annual DKK 1,000 ($125) units. $X_{i,t}$ includes age, year, and municipality fixed effects. To further control for key potential location-related (time-varying) confounders, we include in $X_{i,t}$ municipality m's unemployment rate and average earnings at time t (as well as their interaction with $treat_i$, $post_{i,t}$, and $treat_i \times post_{i,t}$). We instrument for $SB_{i,t}$ using $\overline{SB}_{-i,t,m}$ (where the F-statistic on the excluded instrument in the first stage is 24.25). 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 that 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 column 3 of Table 2 (where columns 1 and 2 present the reduced-form and the first-stage regressions, respectively). The estimate for our parameter of interest, γ_1, is -.0057. With an average of DKK 23,262 ($2,908) in actual survivors benefits received by widows in the analysis sample (including zeros for those not on the program) and with a baseline mean participation rate of 0.505, this estimate translates to a participation elasticity with respect to social benefits of -0.26 for widows under 67.[30] This suggests that formal social insurance provided to survivors crowds out labor supply increases on their part, which can otherwise provide an informal self-insurance mechanism against loss of income following the death of a spouse.

In summary, the findings of this section reveal the following pattern: there are significant average increases in labor supply in response to the death of 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. Put together, the results provide consistent evidence in support of self-insurance as a mechanism for spousal labor supply responses in the extreme case of fatal health shocks, which translate into large and permanent income losses for most households in our setting due to incomplete formal insurance.

5.1.3 Alternative Mechanism: Changes in Spousal Labor Disutility

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 surviving 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 intuitive strategy that provides a suggestive test for this conjecture.

Consider widows (for whom we find an increase in participation in response to spousal death), who did *not* work before the shock, within a simple framework in which time is divided between labor and leisure (or any other use of "time at home"). Among these widows, those in households where the deceased spouse did not work before his death, presumably experience smaller income losses (taking into account the deceased's income from any source including government transfers) but likely consumed more joint leisure; and hence may be more prone to experience social isolation or loneliness following the shock. Similarly, if the deceased spouses in these households did not work prior to the shock 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 this first sub-set of households (in which both spouses did not work prior to the shock) are presumably more likely to experience a decrease in the utility cost of labor supply after the shock. In contrast, widows in the second sub-set of households in which the deceased spouse worked before his death, likely consumed less joint leisure (or provided less care-giving time) but experience larger income losses. Intuitively, the hypotheses of social integration or of decreased care-giving time as the leading motives for the

[30] 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).

average increase in these survivors' participation, are consistent with spouses in the first group of households increasing their labor supply more than spouses in the second group. Conversely, the hypothesis of self-insurance as the leading motive for the estimated increase in these survivors' labor force participation is consistent with the opposite pattern. Note that for our current analysis sample of widows who did not work before the shock we include only those who were similarly detached from the labor market for the entire pre-shock period (i.e., did not work in all periods from -5 to -1), whether or not their husband worked before his death.[31] The overall average increase in these widows' labor force participation in response to their spouse's death is reported in column 1 of Table 3.

To proceed, we first verify in Table 3 the presumed differential level of income losses across the two groups of widows, by showing that households in which the deceased worked experience significantly larger drops in overall income (from any source; see column 3). Then, studying their labor supply in column 2, 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 very low levels of (non-labor) income, there was no increase in the widows' labor supply (see columns 4 and 5). Hence, in the context of this stylized suggestive test, the evidence is inconsistent with the conjecture that a lower cost of labor following the shock is what drives the estimated mean increase in surviving spouses' labor supply (where we cannot, of course, completely rule it out as a mechanism). Rather, like the analysis so far, the evidence supports the view that this increase is likely driven by self-insurance when the shock leads to large income losses.

5.2 Household Labor Supply Responses to Severe Non-Fatal Health Shocks

In our final set of results, we study in this section household labor supply responses to severe non-fatal health shocks. Recall that our analysis sample for these shocks consists of households with both spouses under age 60 in which a spouse experienced a heart attack or a stroke (for the first time) and survived for at least three years.

Panel A of Figure 6 shows that within three years following the shock, the sick individuals' participation sharply falls, which translates into a large loss of annual earnings. Table 4 quantifies these effects by estimating a difference-in-differences regression, in which we allow for differential treatment effects in the "short run" (periods 1 and 2) and in the "medium run" (period 3), to account for the gradual responses documented in Panel A of Figure 6.[32] Columns 2 and 4 of Table 4 reveal

---

[31] To further guarantee that we compare across two groups of spouses with similar labor force attachment, we repeated the analysis only for widows who exited the labor force exactly at the beginning of the pre-period ($t = -5$) within both households where the husband worked and those where the husband did not work, and reached similar conclusions.

[32] Specifically, we estimate the following specification

$$y_{i,t} = \alpha_i + \beta_a post^a_{i,t} + \gamma_a treat_i \times post^a_{i,t} + \beta_b post^b_{i,t} + \gamma_b treat_i \times post^b_{i,t} + \varepsilon_{i,t}, \tag{6}$$

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 that by the third year after the shock the labor force participation rate of the sick individuals drops by 12 pp—about 17%—and that annual earnings drop by DKK 36,070 ($4,500)—a significant drop of 19%. Meyer and Mok (2013) and Dobkin et al. (2017) find similar-magnitude effects of health shocks on own earnings in their US context.

However, while there is a significant drop in the sick individuals' earnings, columns 5 and 6 of 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 benefits that represent about 80% of the recovered loss), reveals that these shocks are well-insured in the Danish setting.

In line with this lack of a considerable income drop, which suggests there is no substantial need for self-insurance, there are no economically significant labor supply responses among spouses (see Panel B of Figure 6 and columns 7 to 10 of Table 4).[33] While fatal and non-fatal shocks differ in many aspects indeed (e.g., in their effects on the household's composition), in terms of the financial aspects of the shocks these results are consistent with our previous set of findings. That is, the behavior of spouses in our analysis of non-fatal health events is likewise in line with the notion of self-insurance as a driving mechanism for spousal labor supply responses to shocks; here in a context where there is no need to exercise this form of informal insurance since households are formally well-insured.

Consistent with this view, we provide two additional sets of results. First, we study how families differentially respond to non-fatal shocks with divergent degrees of severity (as defined by hospitalization days). We show that while greater severity leads to larger participation and earnings decreases by the sick individuals, the higher rate of social insurance (in part due to income testing) equalizes the overall income loss across households with different degrees of shock severity (in the sense that they exhibit no statistically different declines in post-transfer household income). If family labor supply responses are primarily governed by self-insurance motives (rather than, e.g., by preference complementarities in time spent away from work), one would expect no differential spousal labor supply reactions across households that experience shocks of different severity, which is what we find. See Appendix Figure 8 and Appendix Table 5 for this analysis. Second, while the mean decline in household income is negligible in the pooled sample, there is still cross-household variation in overall income replacement rates (which hold the spouse's earnings and social benefits receipts at their pre-shock level). Studying this variation, we find that also in the context of non-fatal health shocks there is a strong partial correlation between spousal labor supply responses and the imposed income losses. See Appendix Table 6 in which we estimate a specification similar to

period 3 and zero otherwise. Therefore, γ_a captures the "short run" effect, and γ_b captures the "medium run" effect.

[33] Specifically, for participation we find no effects in the short run and a small decline in the medium run (of less than 1%), and for earnings we find an overall decline of about 1%. The papers by Garćıa-Gómez et al. (2013) and Jeon and Pohl (2017) document qualitatively similar responses of labor supply declines, where the latter paper estimates responses of larger magnitudes in the context of cancer diagnoses in Canada.

equation (4). These additional sets of findings provide further evidence in support of the insurance mechanism hypothesis for spousal labor supply responses in the context of non-fatal shocks.

Before we continue to the next section, it is worth noting that the data's scale and the research design have allowed for a precise estimation of an economically insignificant average spousal response to non-fatal shocks. This result points to a small but positive degree of complementarity in spouses' labor supply in response to health shocks, with an estimate of 0.064 for the elasticity of the spouse's earnings with respect to the sick individual's earnings. Since the household's income is not perfectly insured, this response likely implies some degree of dependence of the household's utility on the health state. Intuitively, the fact that, given a small loss of income due to the shock, the spouses' decrease in labor supply involves an additional (very small) loss through their lower earnings is consistent with two main health 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 spouses' utility loss from time spent away from home either because they would like to take care of the sick individual or due to preferences for joint leisure—i.e., a labor disutility state dependence.

6 Qualitative Welfare Implications

Having analyzed family responses to severe health shocks from a purely positive standpoint, we conclude with a brief discussion of the potential normative implications of our findings. We do so in light of our theoretical results from Fadlon and Nielsen (2016). In that paper we demonstrate that, in frameworks of efficient household allocations, spousal labor supply responses to shocks have direct implications for the gains from more generous government benefits. Note, however, that a full quantitative welfare analysis is beyond the scope of this paper.[34]

Formula. To derive a concrete welfare formula for our application, we resort to the positive model from Section 2. We make the following additional structural assumption regarding the spouse's labor disutility function for the convenience of our illustrative exercise.

Assumption (spousal labor disutility). *Let $v_2^s(l_2^s)$ take the following constant-elasticity functional form:*

$$v_2^s(l_2^s) = f(X) \times \frac{(l_2^s)^{1+\varphi}}{1+\varphi},$$

[34] While the qualitative normative investigation of our results is only secondary to our analysis, it is worth mentioning its relation to prior work. Related studies that have assessed the welfare gains from social insurance programs (or have provided analysis relevant for such an assessment) have analyzed their consumption smoothing effects, mostly in the context of disability insurance with little direct work on the gains from survivors benefits. 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; McGarry and Schoeni 2005; Meyer and Mok 2013; Chung 2013; Ball and Low 2014; Dobkin et al. 2017) 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).

where $f(X)$ is a general function of the household characteristic vector X and $\varphi > 0$ governs the curvature of the labor disutility function (so that $\varphi = \frac{v_2^s{}''(l_2^s)}{v_2^s{}'(l_2^s)} l_2^s$).

This functional form assumption lets labor disutility vary by important variables, such as age and other personal household characteristics (that are unrelated to labor market opportunities), which may lead to variation in the (pre-shock) within-household division of labor across families. This allows us to make simple welfare comparisons across different sub-populations and households using the corresponding formula from Fadlon and Nielsen (2016) that is presented in the following Lemma (and is derived in Appendix A for completeness).

Lemma. *The marginal (gross) welfare gain from raising benefits in the bad state b to households (with characteristics X) can be represented by*

$$MB \cong \varphi \times \left(\frac{l_2^b}{l_2^g} - 1 \right). \tag{7}$$

This formula illustrates that the gains from additional social insurance can be measured by evaluating changes in spousal labor supply. The basic intuition behind this formula is as follows. More generous government benefits provide greater formal insurance, which decreases the need to compensate for shock-induced income losses using the self-insurance mechanism of spousal labor supply. The formula assesses the utility gain from this decreased labor supply (or, equivalently, from the additional consumption of spousal leisure) by evaluating the benefits from incrementally smoothing labor supply across states. These are captured by multiplying the change in the "quantity" of spousal labor supply in response to shocks, $\frac{l_2^b}{l_2^g} - 1$, by the rate at which the spouse's disutility from additional work changes, φ, which captures the utility "price" of labor-supply quantity fluctuations across the two states. All else equal, the welfare gains from additional social insurance are higher whenever spousal responses to shocks are larger—that is, whenever the household's baseline ability to smooth the spouse's consumption of leisure across states is lower. In the comparative statics of our model in equation (1) we saw that this quantity term increases with income losses and captures the self-insurance role of spousal labor supply. Therefore, intuitively, larger spousal labor supply responses—which correspond to a stronger need to self-insure—imply a greater scope for welfare-improving social insurance due to lack of adequate formal insurance. Similarly, the welfare gains from additional benefits are higher whenever φ is larger, as it implies that self-insurance through spousal labor supply is more costly.

The remainder of this section relies on the set of assumptions that we made so far. We use this simplified framework to draw the suggestive (and local) *qualitative* implications of our findings for the design of social insurance in the context of fatal and severe non-fatal health shocks. We wish to emphasize that the following discussion is limited to our set of assumptions, and also caution against drawing quantitative conclusions from it. The assumptions that we outline here let normative assessments using equation (7) focus only on the empirical moments that have been the target of this paper, i.e., spousal labor supply responses to shocks ($\frac{l_2^b}{l_2^g} - 1$) and the heterogeneity in these responses. Since the data do not contain separate information on wages and work hours, we measure in this section labor supply responses ($\frac{l_2^b}{l_2^g} - 1$) with earnings responses ($\frac{z_2^b}{z_2^g} - 1$) under the working assumption that $z_2^s \equiv w_2 l_2^s$ (so that spouses' potential wages are the same across states). Using the extensive-margin version of the model and the estimated participation responses would lead to similar qualitative conclusions.

Implications. 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. This finding was consistent with the notion that younger surviving spouses, whose participation rates and labor earnings are significantly higher, may be more financially resilient after the death of their spouse. Panel A of Appendix Figure 5 explicitly compares the relative change in earnings for widows below and above 60. It shows that younger widows increase their labor earnings by $\frac{l_2^b}{l_2^g} - 1 = \frac{139{,}944}{135{,}592} - 1 = 3.2\%$, while older widows increase theirs by $\frac{l_2^b}{l_2^g} - 1 = \frac{4{,}777}{2{,}466} - 1 = 94\%$. Under our set of assumptions, this implies that in the context of the current Danish system there are (locally) larger potential welfare gains from providing more generous transfers to older rather than younger widows. For cardinal interpretation, consider the case of quadratic labor disutility (i.e., $\varphi = 1$). In this case, an additional \$1 transferred to survivors could create welfare gains equivalent to 3.2 cents for widows younger than 60 and 94 cents for widows older than 60. More broadly, this result suggests that the differential attachment to the labor force over the life-cycle, which can effectively reduce the financial vulnerability of younger survivors, may justify age-dependent social insurance for spousal mortality shocks.[35]

Second, we found that increases in the surviving spouse's labor supply are strongly (partially) correlated with the share of the household's income the deceased had earned. As the marginal gains from social insurance in equation (7) grow in spouses' labor supply responses to shocks (all else equal), this result suggests a justification for letting survivors benefits depend on the deceased spouse's work history. 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.[36]

[35] At first glance, this might seem at odds with a common view that older households should be provided with less life-insurance coverage since their potential losses through foregone labor earnings are smaller. While this is indeed the case when only labor earnings are considered, there are several important additional aspects that we should take into account. First, a comprehensive view of households' financial resources requires the incorporation of non-labor income streams (such as private or social retirement income), which represent the majority of older individuals' resources and can involve significant losses upon spousal death. Second, the formal insurance environment should be considered. In many developed countries (including Denmark as described in Section 3.1), the coverage of unhealthy and older households through the private market can be substantially limited in practice due to screenings and rejections (see Hendren 2013 for evidence in the US context). In Denmark, this is exacerbated as it is common in both group and non-group markets that even when life-insurance products are purchased by younger and healthy households, the coverage sharply declines with age. Third, the differential ability and opportunities to self-insure through spousal labor supply and the divergent associated costs across age groups should be taken into consideration. A combination of these aspects seems to drive the qualitative normative result.

[36] There may be countervailing arguments when distributional considerations are involved. We abstract from these potentially important considerations since the focus of this section is to draw qualitative conclusions about how to efficiently target benefits

While caution is warranted (as the two settings differ in many aspects), it is worth noting that all of these features characterize the American Social Security system. Surviving spouses in the US are universally 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, which 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 also take into account potential ex-ante behavioral responses with respect to these margins (including household specialization and marriage market decisions).

7 Conclusion

This paper provides novel evidence on households' labor supply responses to fatal and severe non-fatal health shocks. Studying the critical event of the death of a spouse, we find large increases in the surviving spouses' labor supply that are 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. While households that experience fatal and non-fatal health shocks are affected in many (and different) ways, our analysis highlights how the self-insurance aspect of spousal labor supply can provide a unifying explanation for our set of findings. We additionally 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 suggest that richer policy tools that condition transfers on observable characteristics, such as age and household work history, may be welfare improving.

for the purpose of insuring households against the income loss imposed directly by mortality and health shocks.

References

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.

Bailey, M. J. and A. Goodman-Bacon (2015). The war on poverty's experiment in public medicine: Community health centers and the mortality of older americans. *American Economic Review 105*(3), 1067–1104.

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

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.

Berger, M. C. (1983). Labor supply and spouse's health: The effects of illness, disability, and mortality. *Social Science Quarterly 64*(3), 494.

Berger, M. C. and B. M. Fleisher (1984). Husband's health and wife's labor supply. *Journal of Health Economics 3*(1), 63–75.

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.

Bingley, P., N. D. Gupta, and P. J. Pedersen (2011). Disability programs, health and retirement in denmark since 1960. Working Paper 17138, National Bureau of Economic Research.

Blundell, R., L. Pistaferri, and I. Saporta-Eksten (2016). Consumption inequality and family labor supply. *American Economic Review 106*(2), 387–435.

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.

Braakmann, N. (2014). The consequences of own and spousal disability on labor market outcomes and subjective well-being: evidence from germany. *Review of Economics of the Household 12*(4), 717–736.

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.

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.

Chandra, A. and D. O. Staiger (2007). Productivity spillovers in health care: Evidence from the treatment of heart attacks. *Journal of Political Economy 115*(1), 103–140.

Charles, K. K. (1999). Sickness in the family: Health shocks and spousal labor supply.

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

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.

Coe, N. and C. Van Houtven (2010). Spousal health shocks and the timing of the retirement decision in the face of forward-looking financial incentives. Working Paper 2010-7, Boston College Center for Retirement Research.

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.

Deshpande, M. and Y. Li (2017). Who are we screening out? application costs and the targeting of disability programs.

Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2017). The economic consequences of hospital admissions. *American Economic Review (forthcoming)*.

Doyle Jr, J. J. (2011). Returns to local-area health care spending: evidence from health shocks to patients far from home. *American Economic Journal: Applied Economics 3*(3), 221–243.

Fadlon, I. and T. H. Nielsen (2016). Household labor supply and the gains from social insurance.

Fadlon, I. and T. H. Nielsen (2017). Family health behaviors.

Gallipoli, G. and L. Turner (2011). Household responses to individual shocks: Disability and labor supply.

García-Gómez, P., H. Van Kippersluis, O. O?Donnell, and E. Van Doorslaer (2013). Long-term and spillover effects of health shocks on employment and income. *Journal of Human Resources 48*(4), 873–909.

Gertler, P. and J. Gruber (2002). Insuring consumption against illness. *American Economic Review 92*(1), 51–70.

Goldschmidt, D. and J. F. Schmieder (2017). The rise of domestic outsourcing and the evolution of the german wage structure. *Quarterly Journal of Economics*.

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. 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.

Guryan, J. (2004). Desegregation and black dropout rates. *American Economic Review 94*(4), 919–943.

Haanes-Olsen, B. L. (1987). Notes and brief reports. *Social Security Bulletin 50*(11).

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

Haurin, D. (1989). Women's labor market reactions to family disruptions. *Review of Economics and Statistics 71*(1), 54–61.

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

Hilger, N. G. (2016, July). Parental job loss and children's long-term outcomes: Evidence from 7 million fathers' layoffs. *American Economic Journal: Applied Economics 8*(3), 247–83.

Hollenbeak, C. S., P. F. Short, and J. Moran (2011). The implications of cancer survivorship for spousal employment.

Journal of Cancer Survivorship 5(3), 226–234.

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.

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.

Jäger, S. (2016). How substitutable are workers? evidence from worker deaths.

Jeon, S.-H. and R. V. Pohl (2017). Health and work in the family: Evidence from spouses' cancer diagnoses. *Journal of Health Economics 52*, 1 – 18.

Jiménez-Martín, S., J. M. Labeaga, and M. Martínez-Granado (1999). Health status and retirement decision for older european couples.

Johnson, R. W. and M. Favreault (2001). Retiring together or working alone: The impact of spousal employment and disability on retirement decisions.

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.

Low, H. and L. Pistaferri (2015). Disability insurance and the dynamics of the incentive insurance trade-off. *American Economic Review 105*(10), 2986–3029.

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.

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.

Martinello, A. (2017). The effect of unexpected inheritances on wealth accumulation: Precautionary savings or liquidity constraints?

McGarry, K. and R. F. Schoeni (2005). Widow(er) poverty and out-of-pocket medical expenditures near the end of life. *Journals of Gerontology Series B: Psychological Sciences and Social Sciences 60*(3), S160–S168.

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.

Nahum, R.-A. (2007). Labour supply response to spousal sickness absence.

Nekoei, A. and D. Seim (2017). On the accumulation of wealth: The role of inheritance.

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

Parsons, D. O. (1977). Health, family structure, and labor supply. *American Economic Review 67*(4), 703–712.

Persson, P. (2015). Social insurance and the marriage market. Technical report, Working Paper, IFAU-Institute for Evaluation of Labour Market and Education Policy.

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

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.

Siegel, M. J. (2006). Measuring the effect of husband's health on wife's labor supply. *Health Economics 15*(6), 579–601.

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.

World Health Organization (2014). The top 10 causes of death.

Figure 1:
Illustration of the Quasi-Experimental Research Design

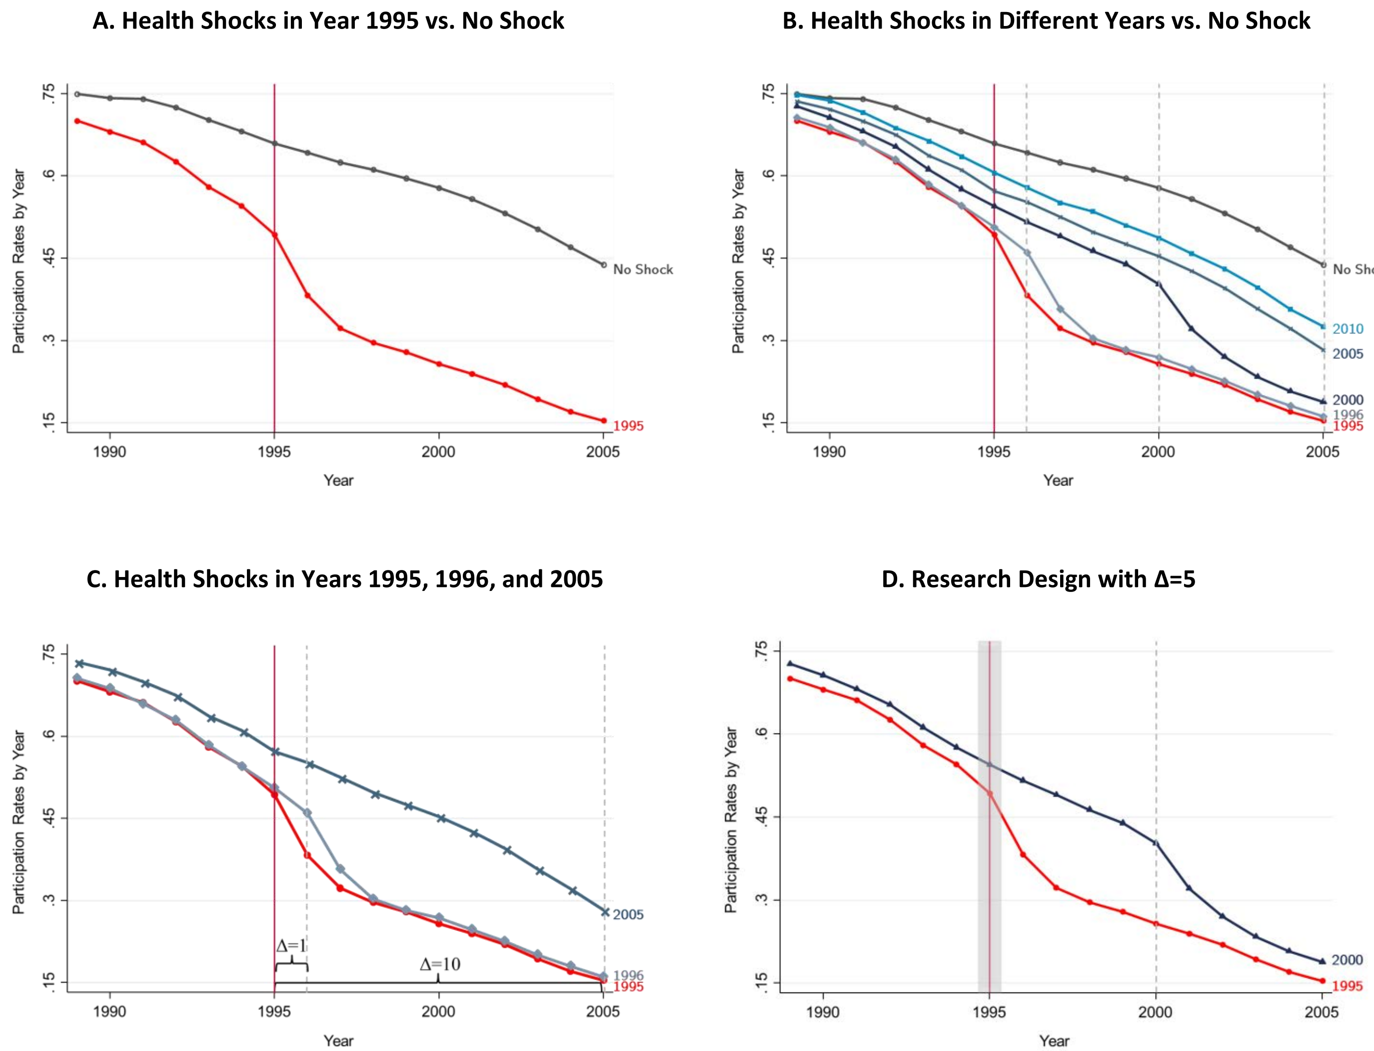

Notes: These figures compare the labor force participation of a treatment group of married and cohabiting individuals born between 1930 and 1950, who experienced a heart attack or a stroke in 1995, to that of potential control groups. Panel A compares the treatment group to married and cohabiting individuals from the same cohorts who have not experienced this shock within our sample period (1980-2011), and reveals considerably different behavioral patterns and visible non-parallel trends prior to 1995 across the two groups. 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 trends before 1995—the closer the year in which the individual experienced the shock was to 1995 (the year in which the treatment group experienced the shock). The figures reveal the similarity of households that experience shocks closer in time, and 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. For example, using those who experienced a shock in 2005, which corresponds to Δ=10, will allow us to estimate effects up to 9 years after 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:
Spouses' Labor Supply Responses to Fatal Health Shocks

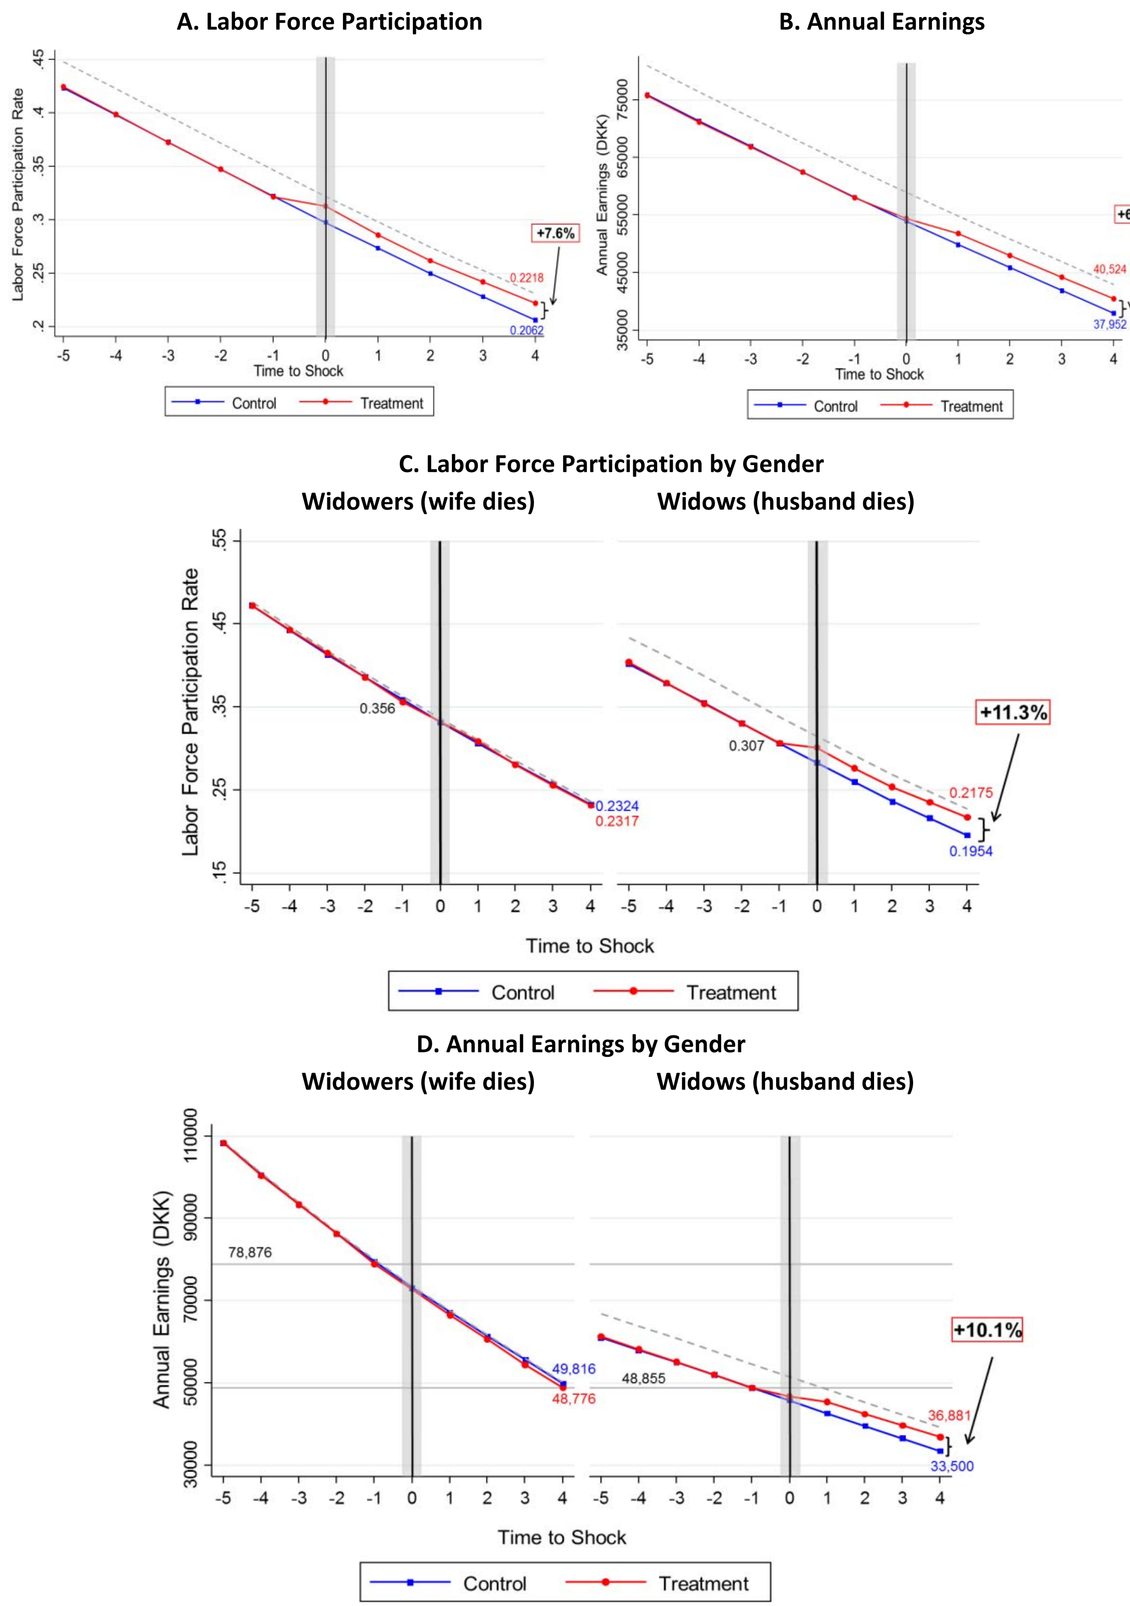

Notes: These figures plot spouses' labor supply responses to fatal health shocks. The sample includes households in which one spouse died between years 1985 and 2011 and in which both spouses were between ages 45 and 80 in the year of the (actual or placebo) shock. Panels A and B depict the behavior 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, from which the treatment effect is identified, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome (in period -2). 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:
Overall Household Income around Fatal Health Shocks

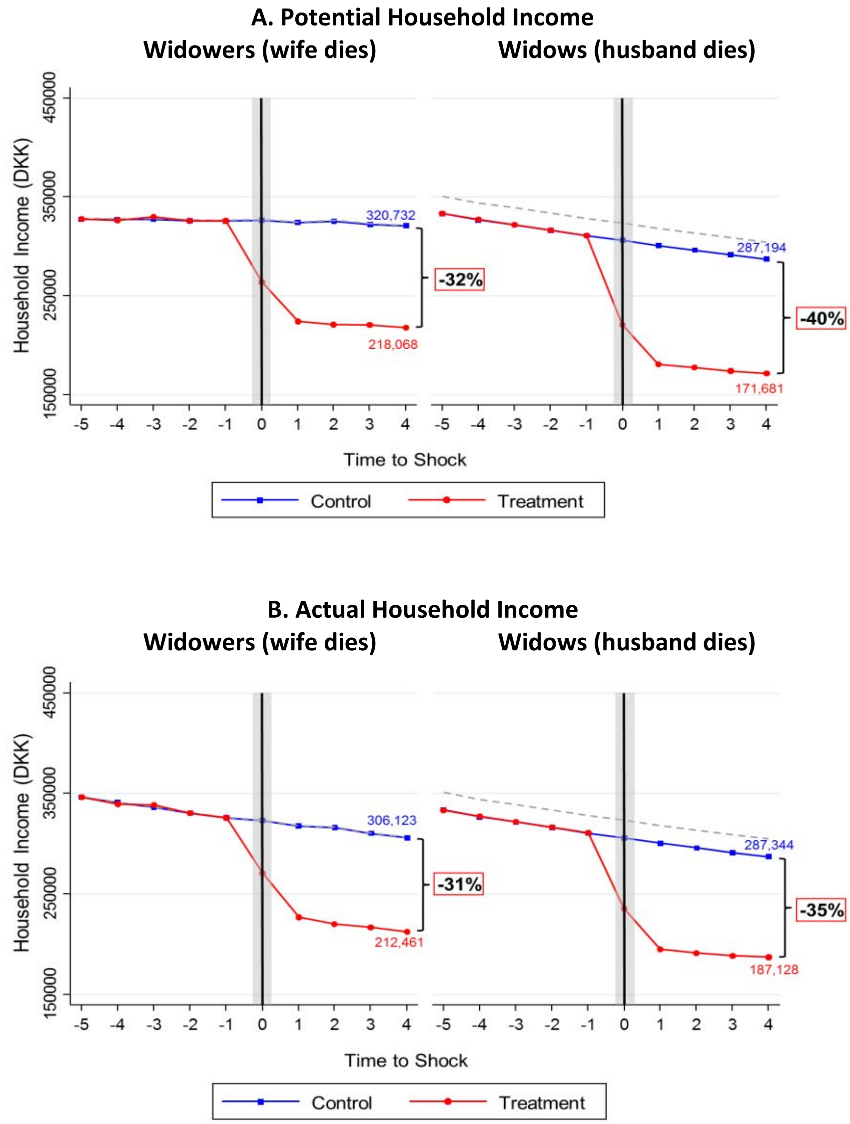

Notes: These figures plot different measures of household-level income around fatal health shocks by the gender of the surviving spouse. Panel A plots an adjusted measure of household income. Specifically, we fix the surviving spouse's labor income, Social Disability, and Social Security benefits at their pre-shock levels (in period -1). Hence, this measure captures the income loss that is directly attributed to the death 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. The figures are constructed as described in the notes of Figure 2.

Figure 4:
Labor Supply Responses of Spouses under Age 60 to Fatal Health Shocks by Gender

A. Labor Force Participation

Widowers (wife dies) **Widows (husband dies)**

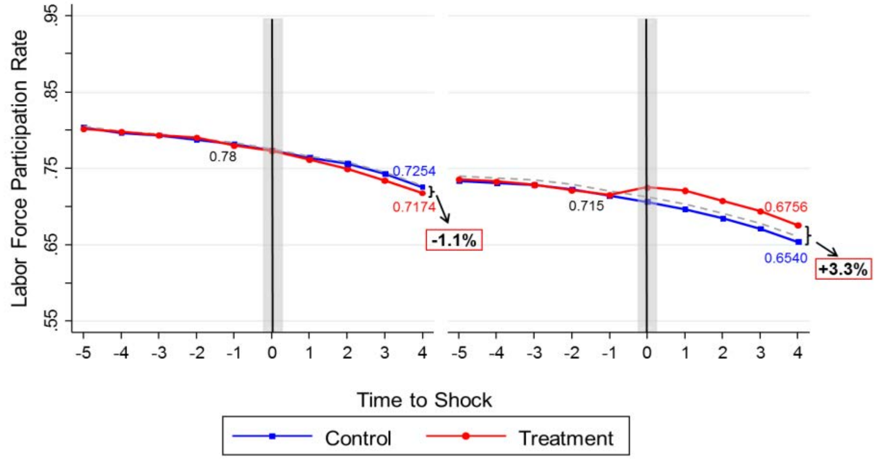

B. Annual Earnings

Widowers (wife dies) **Widows (husband dies)**

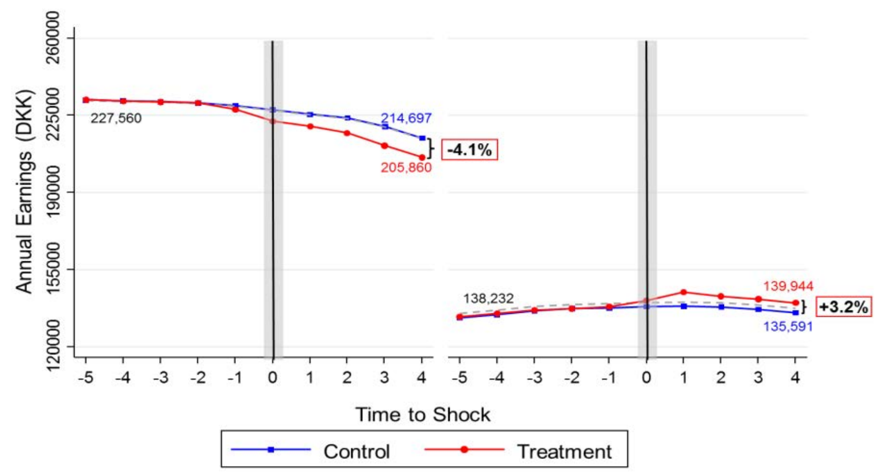

Notes: These figures plot the labor supply responses of spouses under age 60 to fatal health shocks by the gender of the surviving spouse. Panel A depicts the behavior of labor force participation, and Panel B depicts the behavior of annual earnings. The figures are constructed as described in the notes of Figure 2.

Figure 5:
Spouses' Annual Earnings Responses to Fatal Health Shocks
by the Level of their Own Pre-Shock Earnings

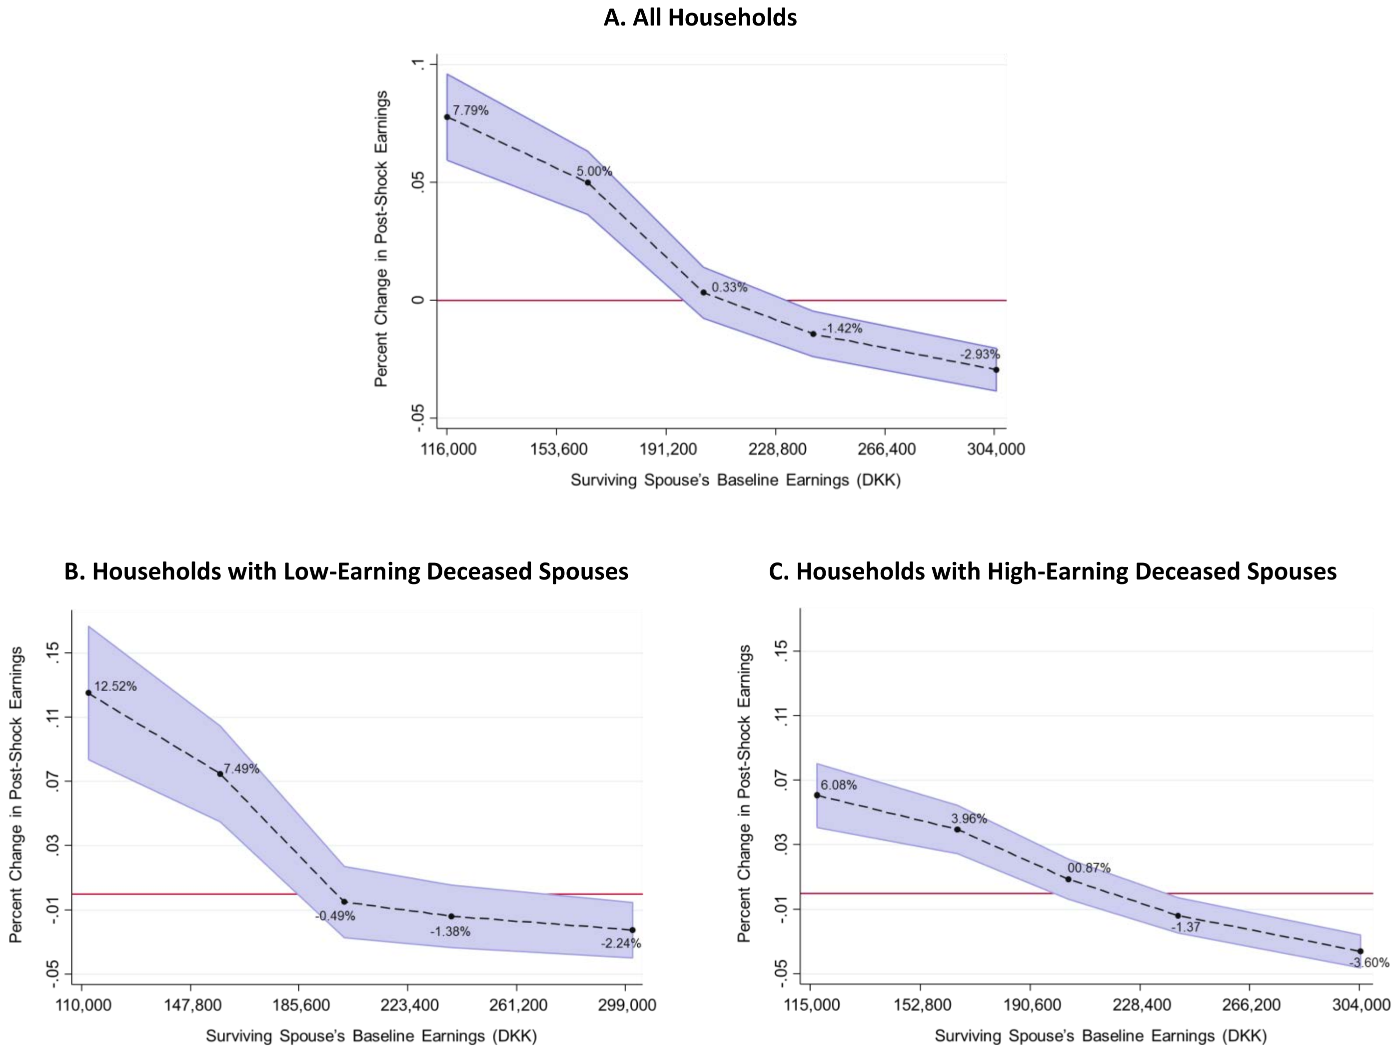

Notes: These figures plot spouses' annual earnings responses to fatal health shocks by the level of their own pre-shock earnings. The households included in the figures are a sub-sample of our sample of fatal health shocks which we construct in the following way. First, we exclude surviving spouses whose average labor income before the shock 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 in which there has been some loss of earned 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. We divide the remaining sample into five equal-sized groups by the surviving spouses' 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 quintile. Panel A includes all households; Panel B includes households with low-earning deceased spouses, whose pre-shock labor income fell within the bottom three quintiles of their group-specific distribution; Panel C includes households with high-earning deceased spouses, whose pre-shock labor income fell within the top two quintiles of their group-specific distribution.

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

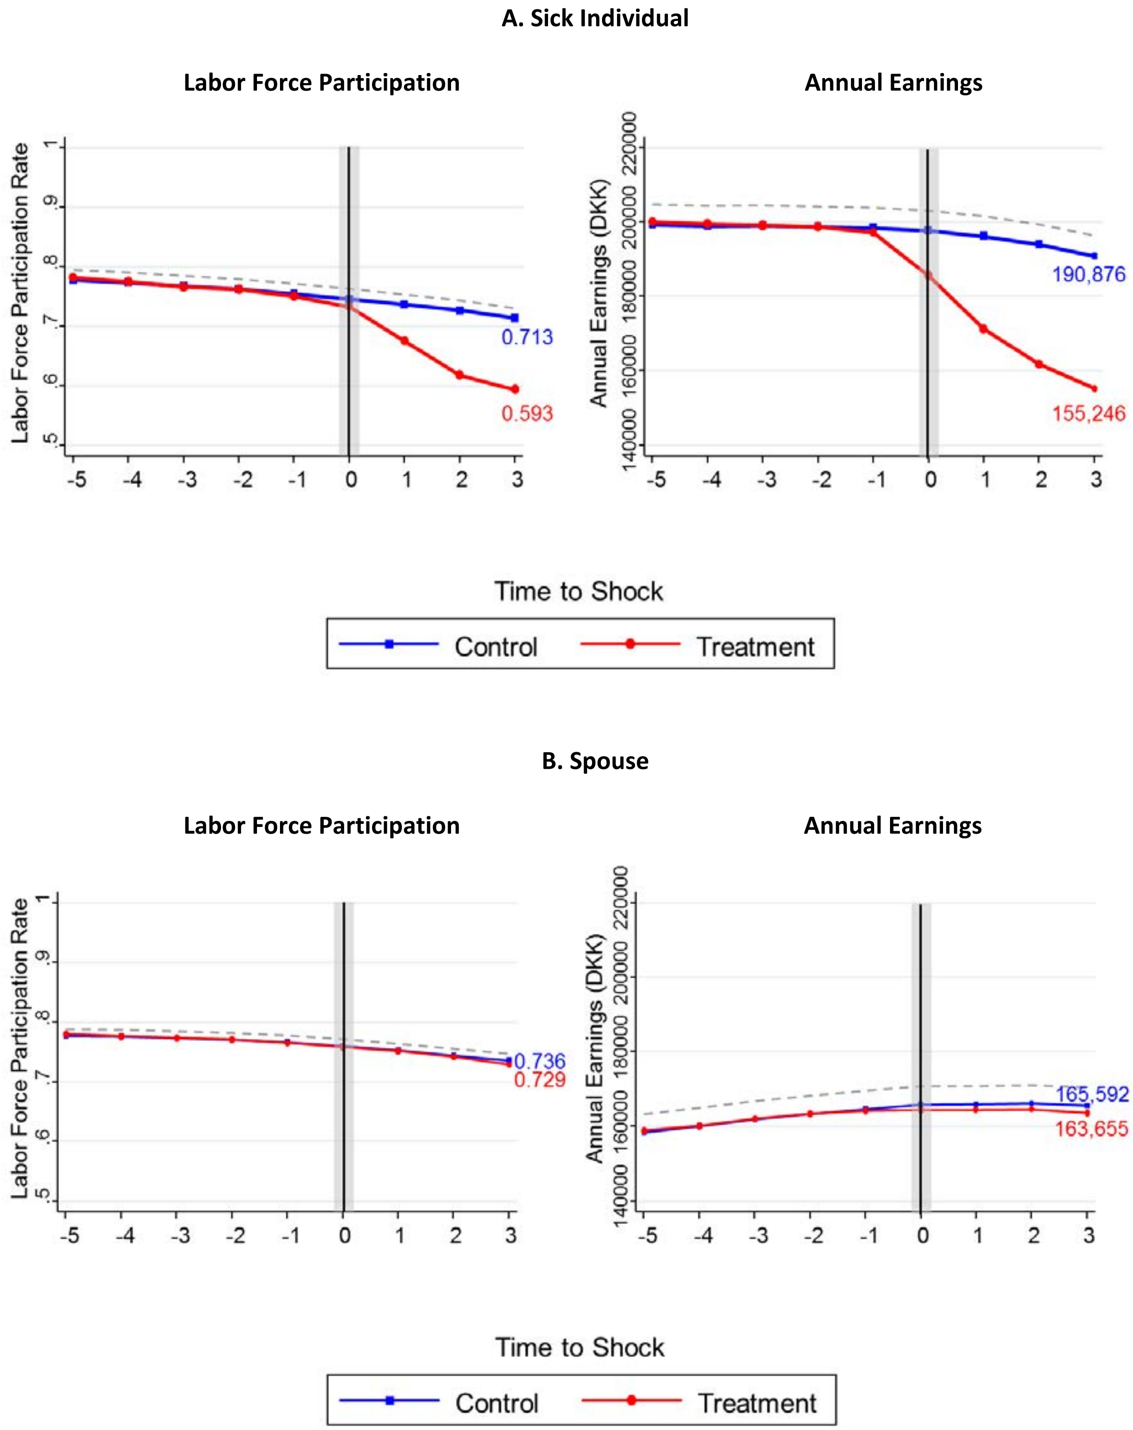

Notes: These figures plot household labor supply responses to severe non-fatal health shocks. The sample includes households in which one spouse experienced a heart attack or a stroke (for the first time) between 1985 and 2011 and survived for at least three years, with both spouses under age 60. Panel A depicts the labor force participation and annual earnings behavior of the sick individual; and Panel B depicts the labor force participation and annual earnings behavior of the spouse. The figures are constructed as described in the notes of Figure 2.

Table 1:
Spouses' Labor Force Participation Responses to Fatal Health Shocks 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*** (0.0023)	0.1220*** (0.0042)	0.1170*** (0.0027)
Treat × Post × Replacement Rate	-0.1889*** (0.0035)	-0.1894*** (0.0061)	-0.1744*** (0.0044)
Number of observations	4,288,621	1,387,615	2,901,006
Number of clusters	714,892	231,318	483,574
2. Regression with Interactions	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post × Replacement Rate	-0.1929*** (0.0046)	-0.2021*** (0.0081)	-0.1933*** (0.0056)
Number of observations	2,741,690	821,742	1,919,948
Number of clusters	459,622	137,724	321,898
Regression 1 for Sub-Sample of Regression 2			
Treat × Post × Replacement Rate	-0.1922*** (0.0043)	-0.1918*** (0.0077)	-0.1832*** (0.0054)
B. Surviving Spouses under 60			
1. Baseline Regression	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post	0.0883*** (0.0054)	0.0652*** (0.0125)	0.0954*** (0.0063)
Treat × Post × Replacement Rate	-0.1270*** (0.0083)	-0.1081*** (0.0168)	-0.1338*** (0.0101)
Number of observations	803,158	201,487	601,671
Number of clusters	134,199	33,720	100,479
2. Regression with Interactions	Both Genders (1)	Widowers (2)	Widows (3)
Treat × Post × Replacement Rate	-0.1471*** (.0094)	-0.1375*** (.0186)	-0.1515*** (.0110)
Number of observations	704,370	173,620	530,750
Number of clusters	118,812	29,288	89,524
Regression 1 for Sub-Sample of Regression 2			
Treat × Post × Replacement Rate	-0.1377*** (0.0088)	-0.1236*** (0.0184)	-0.1430*** (0.0107)

Notes: This table reports the interaction of the treatment effect of fatal spousal shocks with the household's post-shock income replacement rate using the specification of equation (4). This replacement rate is calculated as follows. First, we fix the surviving spouse's labor income, Social Disability, and Social Security 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 purge life-cycle and time effects. Panel A reports estimates for the sample of all surviving spouses by gender; Panel B reports estimates for the sample of surviving spouses 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. Specification 2 in each panel extends specification 1 to include interactions of the treatment effect with additional household characteristics: age fixed effects for the surviving spouse, fixed effects for the age of the deceased at the year of death, year fixed effects, indicators for the number of children in the household and for the presence of adult children (18 or older) and young children (6 or younger), 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 (which also accounts for liquidation of housing assets and changes in mortgage debt). 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, spouse's age, and household fixed effects, and additionally include the interactions with *Post* of covariates that are interacted with *Treat* × *Post*, the variables among these covariates that are time variant, and the interactions of the time-variant variables with *Treat*. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 2:
Widows' Labor Force Participation Responses to Fatal Spousal Shocks by the Generosity of Social Survivors Benefits

	Reduced Form (1)	First Stage (2)	Two-Stage Least Squares (3)
Treat × Post × Municipality-Specific Survivors Benefits Receipts	-.000913*** (.000332)		
Treat × Post × Survivors Benefits		.1468*** (.0298)	-.0057*** (.0020)
Mean Treatment Effect			1.8 pp
Counterfactual Mean Participation Rate			48.7 pp
Combined Mean Participation Rate			50.5 pp
Number of observations	364,100	364,100	364,100
Number of clusters	268	268	268

Notes: This table reports the interaction of the treatment effect of fatal health shocks with the generosity of survivors benefits that widows receive through the Social Disability Insurance (Social DI) program. We estimate the specification of equation (5), where the instrument we use for actual benefits received by widows 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 for the current sub-sample) 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 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 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 observations	1,320,908	1,320,908	1,320,908	114,851	57,381
Number of clusters	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 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 a husband died 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. Column 1 reports the simple differences-in-differences estimate in a regression similar to equation (3), 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 in a specification similar to equation (4). 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 mean spousal labor force participation effects (using the specification of equation (3)) for sub-samples of households in which the husband did not work and received very low levels of (non-labor) income before his death: column 4 reports the treatment effect for households in which the non-working deceased spouse's overall income before the shock, including any transfer from government programs, was lower than the 10th percentile of this sample's income distribution; column 5 reports the treatment effect for households in which the non-working deceased spouse's overall income before the shock was lower than the 5th percentile. All specifications include year, widow's age, and household fixed effects, and additionally include the interaction with *Post* of covariates that are interacted with *Treat* × *Post*. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

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

	Sick Individual				Household Income		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	-.0863***	-.1214***	-29,120***	-36,070***	-12,187***	-18,638***	-.0019	-.0071***	-1,754***	-2,012***
	(.0023)	(.0028)	(738)	(877)	(2161)	(2382)	(.0020)	(.0024)	(533)	(624)
Household FE	X	X	X	X	X	X	X	X	X	X
Counterfactual Post-Shock Mean of Dependent Var.	.7333	.7151	195,754	191,474	504,127	503,875	.7489	.7365	166,444	165,960
Percent Change	-12%	-17%	-15%	-19%	-2.4%	-3.7%	0	-1%	-1.05%	-1.21%
Percent Change Excluding Spousal Responses					-2.1%	-3.3%				
Number of observations	644,354		644,354		644,354		644,354		644,354	
Number of clusters	92,324		92,324		92,324		92,324		92,324	

Notes: This table reports estimates of household labor supply responses to severe non-fatal health shocks and the effect of these shocks on overall household income using the specification of equation (6) (in footnote 32). We allow for differential treatment effects for the "short run", periods 1 and 2, and for the "medium run", period 3, to account for the gradual responses documented in Figure 6. Household income (in 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 by experimental-group level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Figure 1:
Life-Cycle Labor Force Participation of Spouses

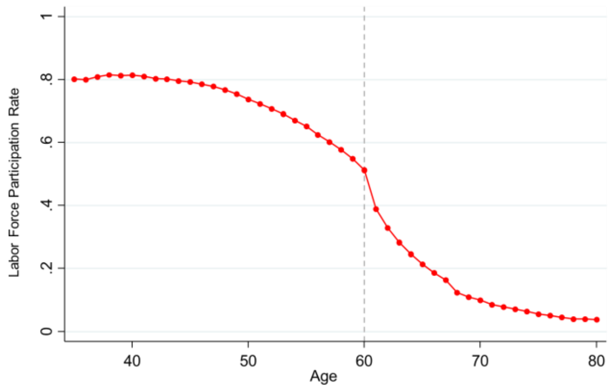

Notes: This figure displays the life-cycle trends in spousal labor supply. As an illustration, we depict the labor force participation of spouses included in the fatal health shock sample (of households in which one spouse experienced a fatal shock between years 1985 and 2011), where very similar patterns can be shown for spouses in the overall sample of non-fatal health shocks (of households in which one spouse experienced a heart attack or a stroke for the first time between 1985 and 2011 and survived for at least three years). The observations included in the figure are from the pre-shock periods. The sharp drop at age 60 corresponds to eligibility for the Voluntary Early Retirement Pension (VERP).

Appendix Figure 2:
Spouses' Labor Supply Responses to Fatal Heart Attacks and Strokes

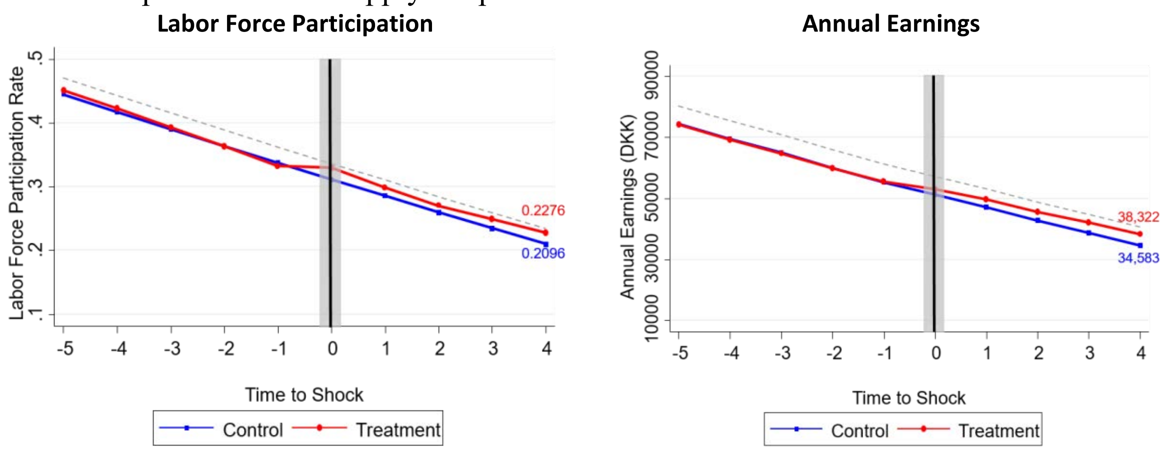

Notes: These figures plot spouses' labor supply responses to fatal heart attacks and strokes. Specifically, we constrain our sample of fatal health shocks to households in which an individual experienced a heart attack or a stroke (for the first time) between 1985 and 2011 and died within that year. 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 figures are constructed as described in the notes of Figure 2.

Appendix Figure 3:
Different Margins of Labor Supply Responses to Fatal Health Shocks of Surviving Spouses under Age 60

A. Labor Force Participation

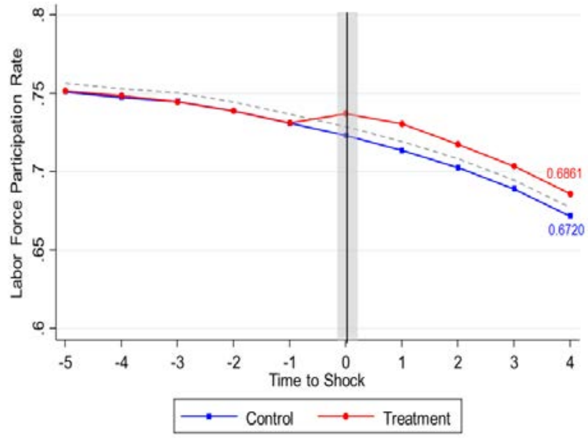

B. Full-Time Employment

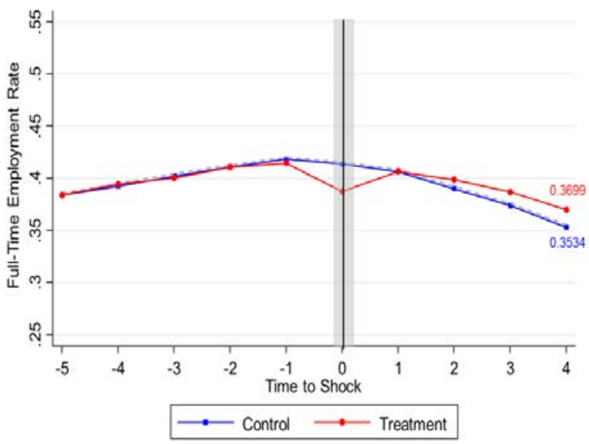

C. Part-Time Employment

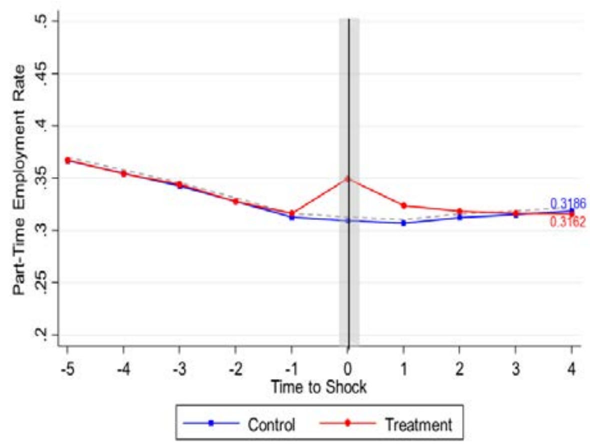

Notes: These figures plot different margins of labor supply responses to fatal health shocks of surviving spouses under age 60. Panel A depicts the behavior of labor force participation; Panels B and C depict the fraction of surviving spouses who are employed full time and part time, respectively. These administrative employment measures are available for workers under 60 based on records of their payments to the government-mandated ATP pension scheme. 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 figures are constructed as described in the notes of Figure 2.

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

A. Take-Up of Social Survivors Benefits **B. Distribution of Average Survivors Benefits**

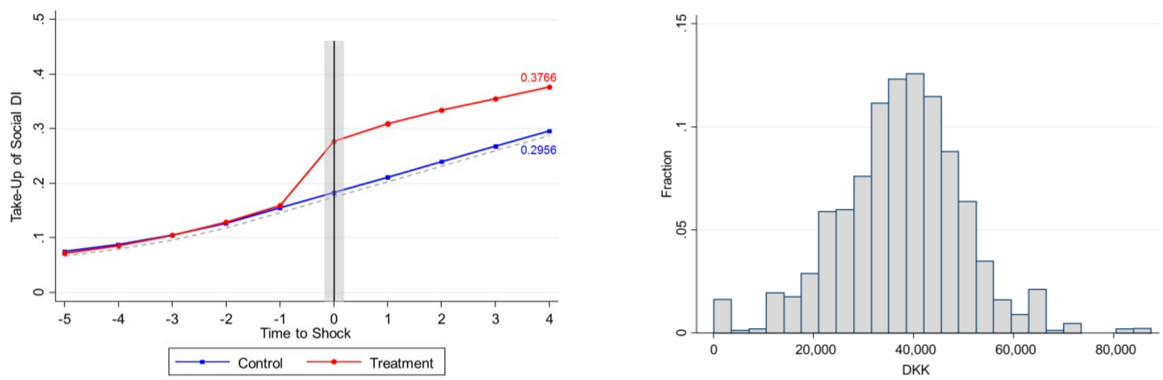

Notes: These figures include widows younger than 67 in years prior to 1994 (when there is a data break in the reporting method of benefits received through Social Disability Insurance). Panel A plots these widows' take-up of social survivors benefits through the Social Disability Insurance (Social DI) program around the death of their spouse. This figure is constructed as described in the notes of Figure 2. Panel B displays the distribution of the instrument that we use in the estimation of equation (5), i.e., the year-by-municipality ("leave-one-out") mean of survivors benefits received by non-working surviving spouses through Social DI.

Appendix Figure 5:
Widows' Labor Supply Responses to Fatal Health Shocks by Age

A. Annual Earnings

Widows under Age 60 **Widows over Age 60**

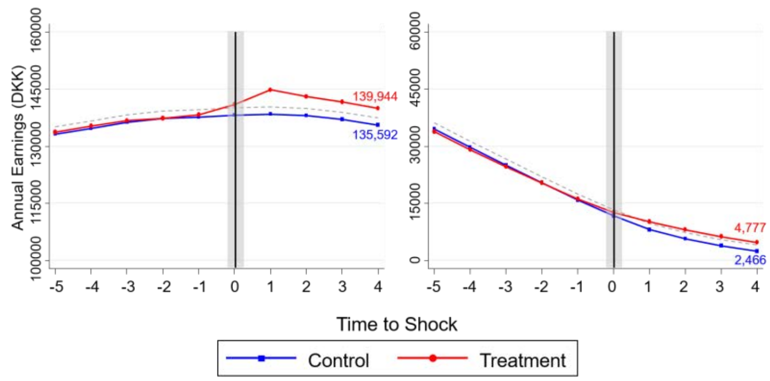

B. Labor Force Participation

Widows under Age 60 **Widows over Age 60**

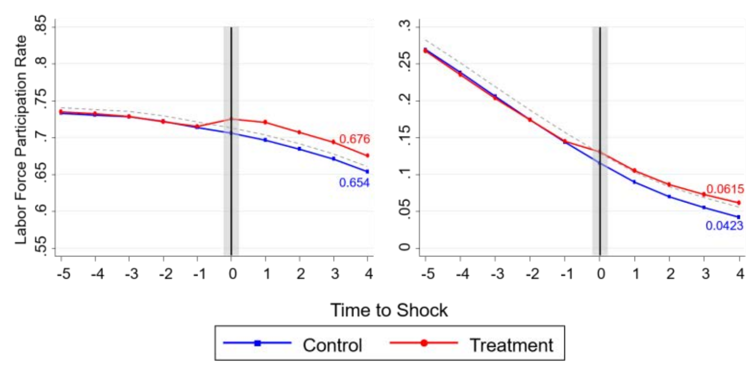

Notes: These figures plot widows' labor supply responses to fatal health shocks by their age. Panel A depicts the behavior of annual earnings, and Panel B depicts the behavior of labor force participation. In each panel, the figure on the left depicts the behavior of widows younger than 60, the age at which there is a sharp decline in labor force participation due to eligibility for early retirement benefits; and the figure on the right depicts the behavior of widows older than 60. The figures are constructed as described in the notes of Figure 2.

Appendix Figure 6:
Different Components and Moments of Household Income around Fatal Health Shocks

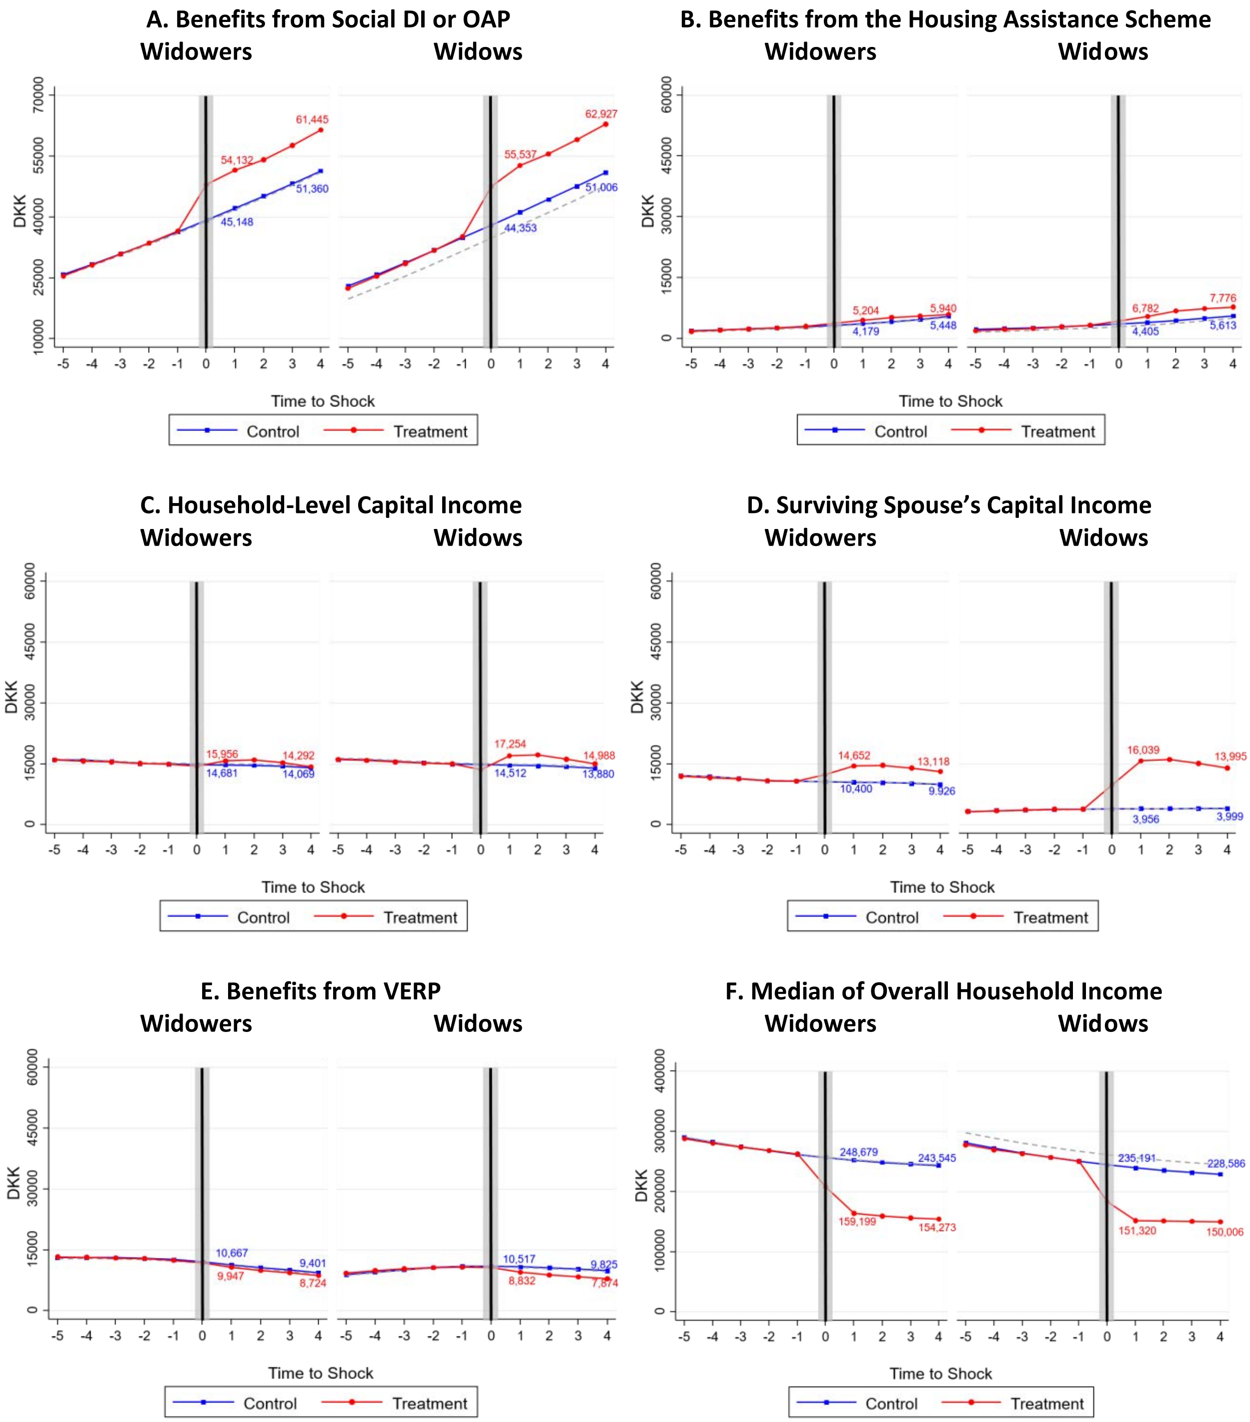

Notes: These figures plot the evolution of different components and moments of household income around fatal health shocks. The figures are constructed as described in the notes of Figure 2. Capital income is winsorized at its 99th percentile.

Appendix Figure 7:
Different Components and Moments of Household Wealth around Fatal Health Shocks

A. Median of Household Net Wealth
Widowers Widows

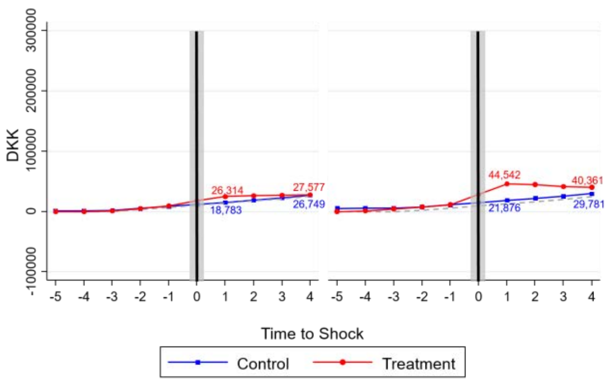

B. Mean of Household Net Wealth
Widowers Widows

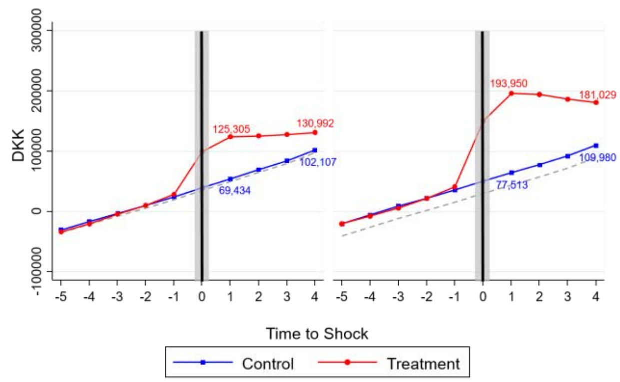

C. Median Change in Household Net Wealth
Widowers Widows

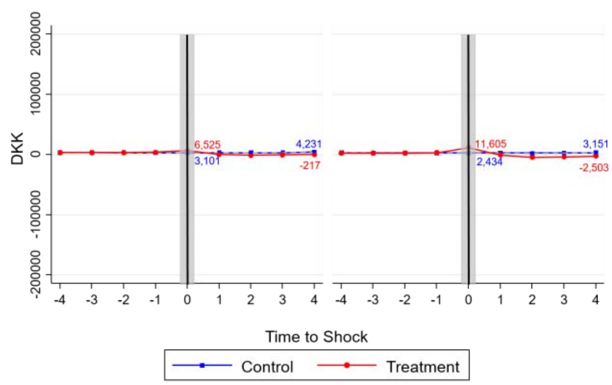

D. Mean Change in Household Net Wealth
Widowers Widows

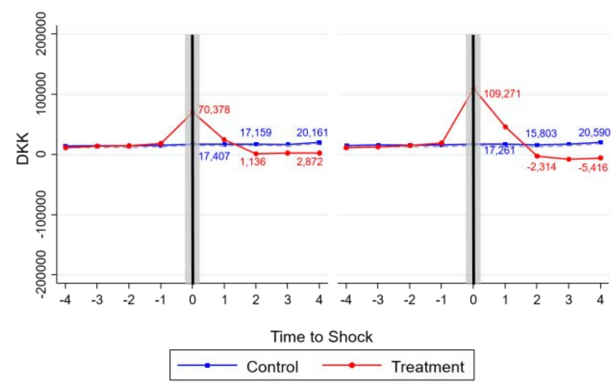

E. Rate of Year-to-Year Increase in Household Net Wealth
Widowers Widows

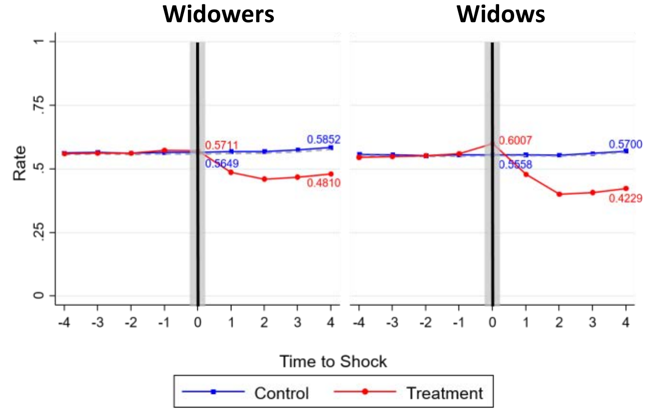

Appendix Figure 7—contd.:

Different Components and Moments of Household Wealth around Fatal Health Shocks

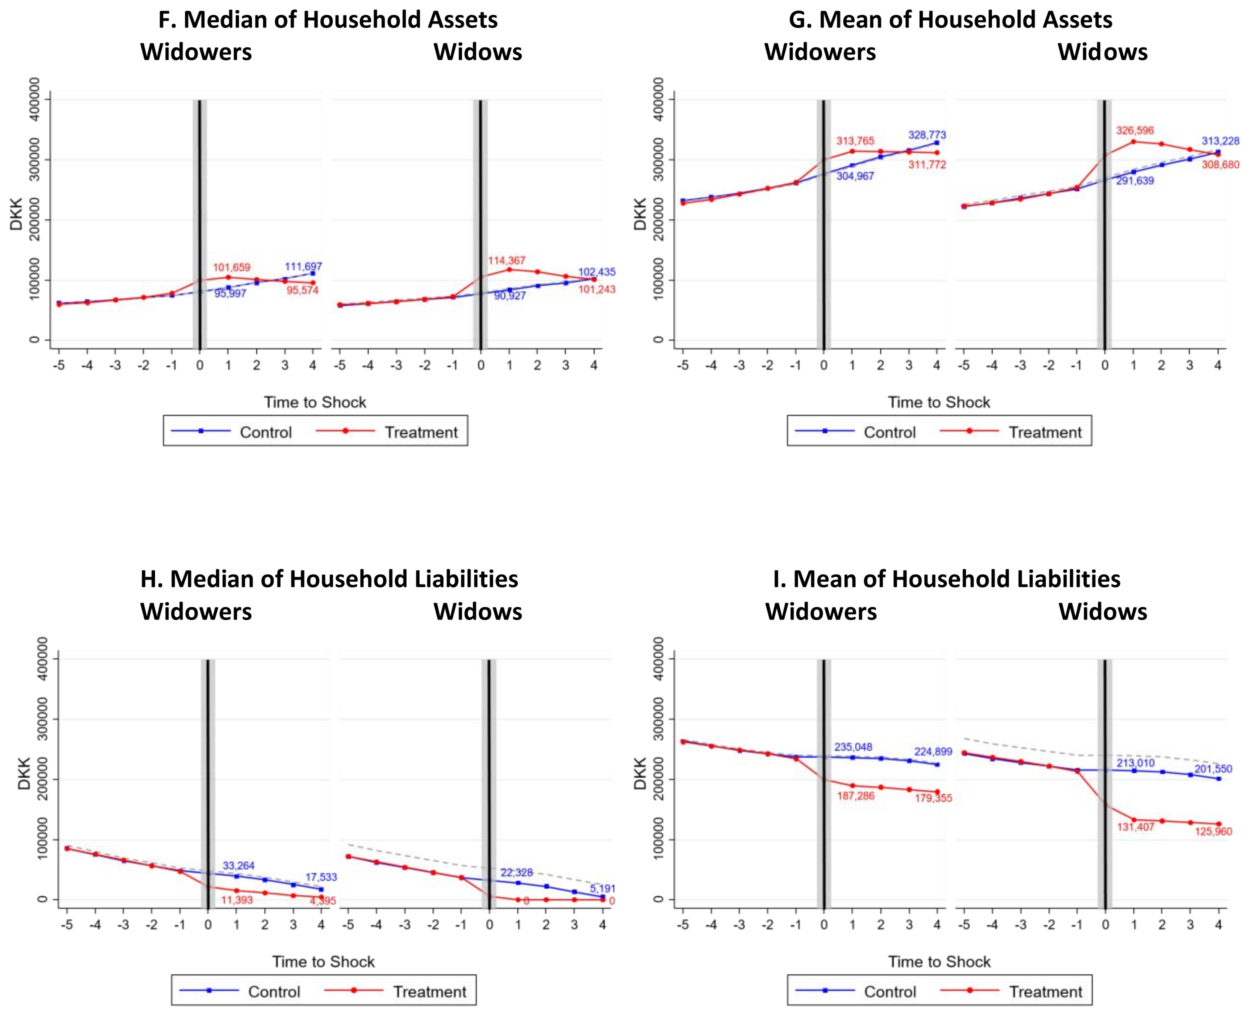

Notes: These figures plot the evolution of different components and moments of household wealth around fatal health shocks. The figures are constructed as described in the notes of Figure 2. All measures are winsorized at their 99th percentile.

Appendix Figure 8:
Household Labor Supply Responses to Non-Fatal Health Shocks of Differential Severity

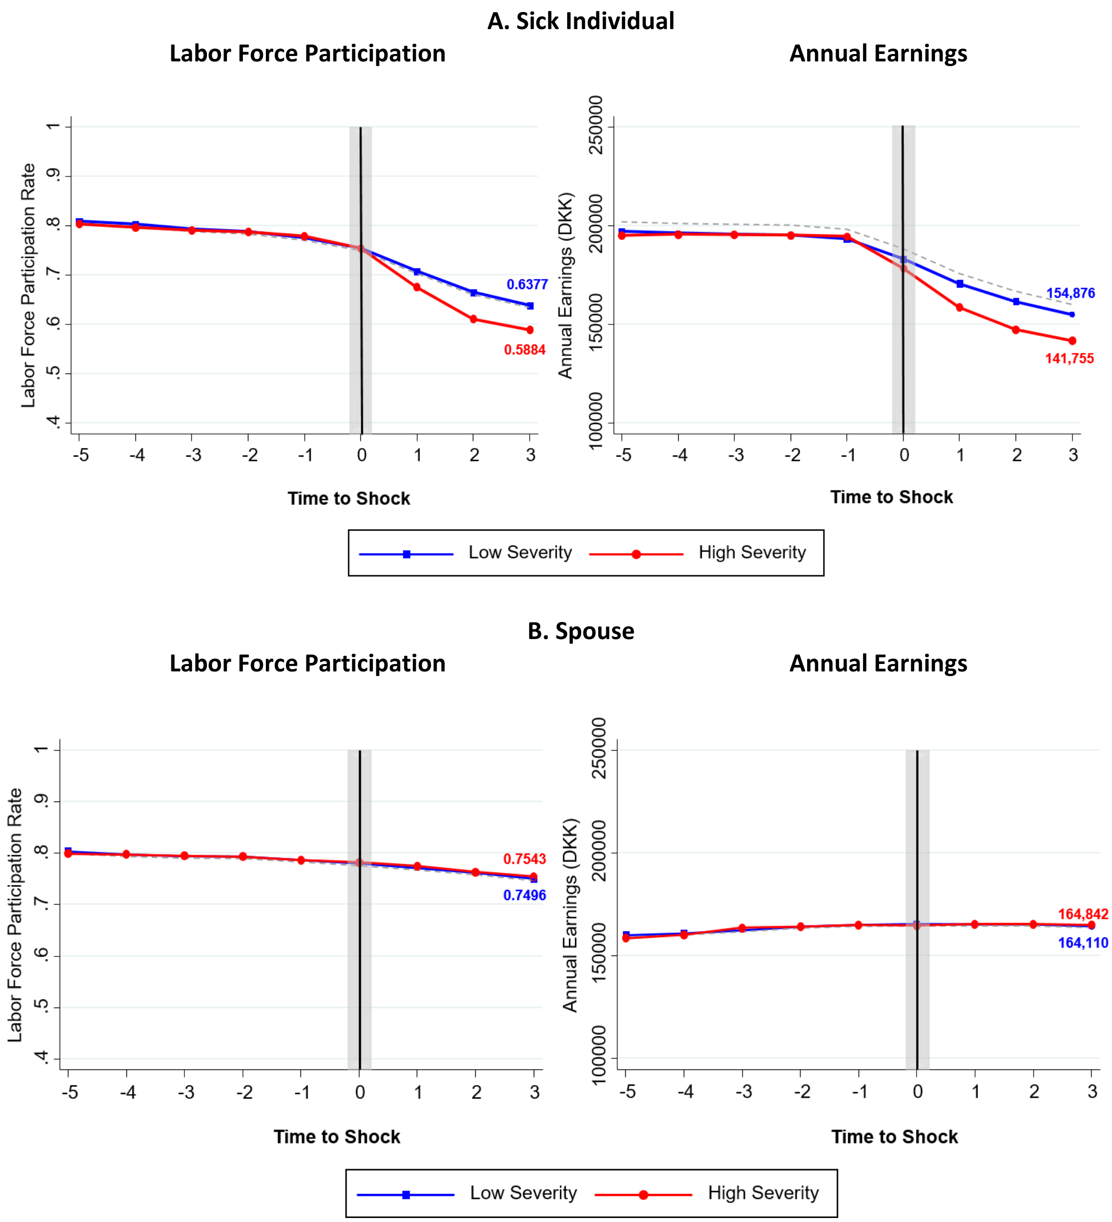

Notes: These figures plot household labor supply responses to non-fatal health shocks with different degrees of severity. The sample includes only households in the treatment group within our sample of non-fatal health shocks, which we divide by the shock's severity according to the 75th percentile of the distribution of hospitalization days associated with the shock (12 days in our sample). We assign households to the "low severity" group if the sick individual was hospitalized for less than 12 days following the shock, and assign households to the "high severity" group if the sick individual was hospitalized for 12 days or more following the shock. Panel A depicts the labor force participation and annual earnings behavior of the sick individual; Panel B depicts the labor force participation and annual earnings behavior of the spouse. The figures are constructed as described in the notes of Figure 2, but with the change that we treat the "low severity" households as the control group and the "high severity" households as the treatment group.

Appendix Figure 9:
Means-Testing of "Social Pensions" (Social Disability Insurance and Old-Age Pension)

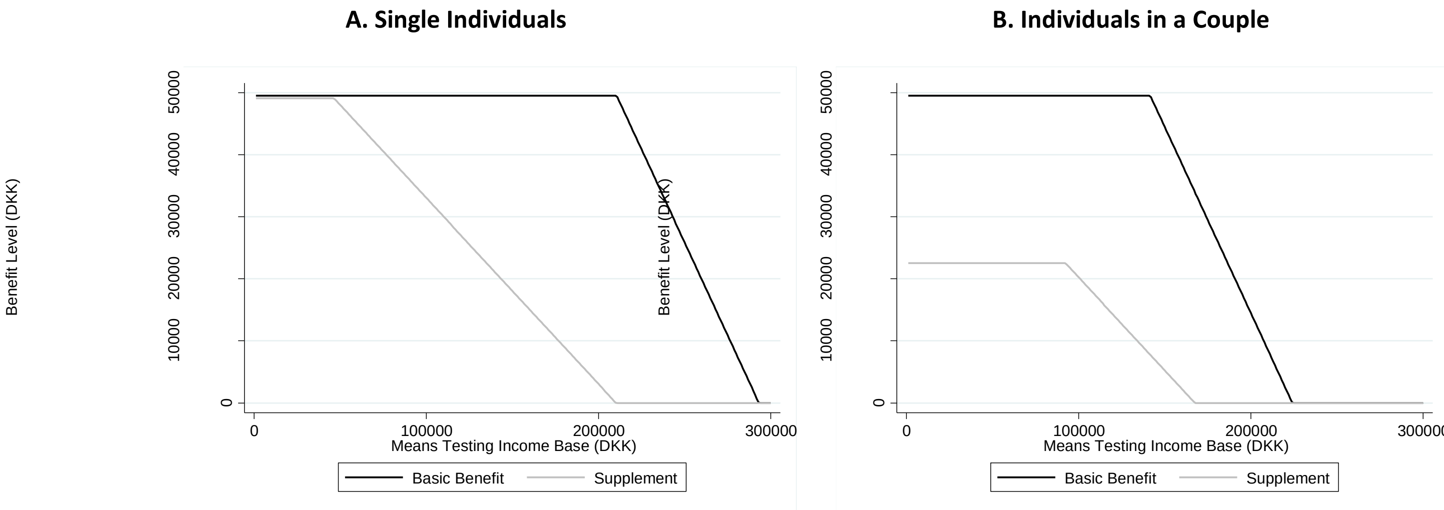

Notes: These figures plot the means-testing rules (at year 2000 rates and thresholds) in benefit levels transferred to households through the Social Pension (SP) scheme, which includes the Social Disability Insurance (DI) and the Old-Age Pension (OAP) programs. Panel A plots the rules for single recipients, and Panel B plots the rules for married or cohabiting recipients. In both DI and OAP, the income-tested transfers consist of a *basic benefit* of DKK 49,560 (solid black lines) and a *supplement* of DKK 49,140 for single individuals and DKK 22,536 for individuals in a couple (solid gray lines). The y-axis denotes (pre-income-tax) benefit levels; the x-axis denotes the corresponding means-testing (pre-tax) income bases. For the basic benefit, the income base for DI recipients is own overall non-SP income, and the income base for OAP recipients is own labor earnings. For the supplement, within both DI and OAP, the income base in Panel A is own overall non-SP income, and the income base in Panel B is the household's overall income (excluding own SP benefits).

Appendix Table 1:
Summary Statistics of Analysis Sample

		Fatal Health Shocks				*Non-Fatal Health Shocks*	
		All Ages (1)		Under 60 (2)		Under 60 (3)	
		Treatment	Control	Treatment	Control	Treatment	Control
Panel A: Household Characteristics							
Spouses	Year of Observation	1993.13	1993.09	1992.74	1992.75	1991.83	1991.95
	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
Individuals that Have Experienced the Shock	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
Panel B: Baseline Outcomes							
Spouses	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
Individuals that Have Experienced the Shock	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	37,437	54,887

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 that 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 as a control group households from the same cohorts that experienced the same shock but five years later (Δ=5). Column 1 reports statistics for the fatal health shock sample of households in which one spouse died between years 1985 and 2011 and in which both spouses were between ages 45 and 80 in the year of the (actual or placebo) shock. Column 2 reports statistics for the sub-sample of households with surviving spouses under age 60. Column 3 reports statistics for the non-fatal health shock sample. It includes households in which one spouse experienced a heart attack or a stroke (for the first time) between 1985 and 2011 and survived for at least three years, with both spouses under age 60. The values reported in the table are based on data from two years before the actual shock for the treatment and two years before the placebo shock for the control group (i.e., from period t = -2 for both groups).

Appendix Table 2:
Spouses' Labor Supply Responses to Fatal Health Shocks

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 observations	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 clusters	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 observations	203,569	203,569	204,438	204,438	607,437	607,437	608,742	608,742
Number of clusters	34,104	34,104	34,118	34,118	101,529	101,529	101,562	101,562

Notes: This table reports estimates of spouses' labor supply responses to fatal health shocks by the gender of the surviving spouse, using the differences-in-differences specification of equation (3). Panel A reports the responses of all surviving spouses in the sample of fatal health shocks, where widowers are those who lost their wives and widows are those who lost their husbands. Panel B reports the responses of surviving spouses under age 60. These regressions present the medium-run effects (that are the focus of our analysis) so that *Post* assumes the value 1 for periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 3:
Spouses' Annual Earnings Responses to Fatal Health Shocks

		All Surviving Spouses (1)		Low-Earning Deceased (2)		High-Earning Deceased (3)	
A. Mean Responses by Quintiles of Survivors' Own Pre-Shock Earnings							
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
Year and Age 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 observations	686,521	220,125	466,392
Number of clusters	114,462	36,705	77,756

Notes: This table reports estimates of spouses' annual earnings responses to fatal health shocks for each quintile of the surviving spouses' pre-shock level of own earnings, using the differences-in-differences specification of equation (3). The households included in these estimations are a sub-sample of our sample of fatal health shocks which we construct in the following way. First, we exclude surviving spouses whose average labor income before the shock 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 in which there has been some loss of earned 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. We divide the remaining sample into five equal-sized groups by the surviving spouses' pre-shock level of earnings. Panel A separately estimates equation (3) for each surviving spouses' earnings quintile. Column 1 includes all households; column 2 includes households with low-earning deceased spouses whose pre-shock labor income fell within the bottom three quintiles of their group-specific distribution; column 3 includes households with high-earning deceased spouses whose pre-shock labor income fell within the top two quintiles. The gradient of surviving spouses' 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 average treatment effects within this sample, where the second row indicates counterfactual outcomes based on equation (3). The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 4:
Spouses' Labor Supply Responses to Fatal Health Shocks by the Degree of Income Loss

A. Surviving Spouses under 60		
	Participation by spouses who Did Not Work	Full-Time by Spouses who Worked Part-Time
	(1)	(2)
1. Baseline Regression		
Treat × Post	0.1553***	0.0840***
	(0.0105)	(0.0094)
Treat × Post × Replacement Rate	-0.1577***	-0.0912***
	(0.0169)	(0.0145)
2. Regression with Interactions		
Treat × Post × Replacement Rate	-0.1455***	-0.0841***
	(0.0174)	(0.0150)
Number of observations	117,868	284,980
Number of clusters	19,966	48,124

B. Earnings Responses by Surviving Spouses of All Ages			
	Earnings		
	All Households	Spouses who Continuously Did Not Work	Spouses who Continuously Worked
	(1)	(2)	(3)
1. Baseline Regression			
Treat × Post	58,927***	4,955***	70,726***
	(703)	(432)	(1282)
Treat × Post × Replacement Rate	-93,161***	-6,638***	-101,556***
	(1194)	(700)	(1914)
2. Regression with Interactions			
Treat × Post × Replacement Rate	-84,296***	-7,321***	-91,682***
	(1198)	(739)	(1954)
Number of observations	2,741,690	1,182,622	1,240,640
Number of clusters	459,622	197,938	208,289

Notes: This table reports the interaction of the treatment effect of fatal spousal shocks with the household's post-shock income replacement rate using the specification of equation (4). This replacement rate is calculated as described in the notes of Table 1. Panel A reports estimates for the sample of surviving spouses under age 60: in column 1 we study participation responses by surviving spouses who did not work for the entire pre-period; and in column 2 we study full-time employment by surviving spouses who worked for the entire pre-period where at least one of these periods involved part-time work. Panel B reports estimates for the sample of all surviving spouses: in column 1 we study earnings responses by all surviving spouses; in column 2 we study earnings responses by surviving spouses who did not work for the entire pre-period; and in column 3 we study earnings responses by surviving spouses who worked for the entire pre-period. 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. Specification 2 in each panel extends specification 1 to include interactions of the treatment effect with additional household characteristics: age fixed effects for the surviving spouse, fixed effects for the age of the deceased at the year of death, year fixed effects, indicators for the number of children in the household and for the presence of adult children (18 or older) and young children (6 or younger), the surviving spouse's gender, and the surviving spouse's months of education (and its square). In both specifications 1 and 2 we include only households for whom there are no missing values for this entire set of controls. All specifications include year, spouse's age, and household fixed effects, and additionally include the interactions with *Post* of covariates that are interacted with *Treat* × *Post*, the variables among these covariates that are time variant, and the interactions of the time-variant variables with *Treat*. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Appendix Table 5:
Household Responses to Non-Fatal Health Shocks of Differential Severity

Dependent variable:	Sick Individual				Household Income		Spouse			
	Participation		Earnings				Participation		Earnings	
	Short Run (1)	Medium Run (2)	Short Run (3)	Medium Run (4)	Short Run (5)	Medium Run (6)	Short Run (7)	Medium Run (8)	Short Run (9)	Medium Run (10)
High Severity × Post	-.0398***	-.0462***	-12,406***	-12,469***	-2,573	-2,059	.0043	.0057	456	909
	(.0044)	(.0054)	(1,368)	(1,633)	(3,475)	(4,076)	(.0036)	(.0042)	(943)	(1,105)
Household FE	X	X	X	X	X	X	X	X	X	X
Number of observations	261,806		261,806		261,806		261,806		261,806	
Number of households	37,431		37,431		37,431		37,431		37,431	

Notes: This table reports estimates of changes in household labor supply and overall income in response to non-fatal health shocks with different degrees of severity. The sample includes only households in the treatment group within our sample of non-fatal health shocks, which we divide by the shock's severity according to the 75th percentile of the distribution of hospitalization days associated with the shock (12 days in our sample). We construct a binary variable of "high severity" by assigning the value 0 for households in which the sick individual was hospitalized for less than 12 days following the shock, and the value 1 for households in which the sick individual was hospitalized for 12 days or more following the shock. We then estimate a specification similar to equation (6) (in footnote 32), where we substitute the variable *Treat* with the variable *High Severity*. Household income (in columns 5 and 6) includes income from any source—including earnings, capital income, annuity payouts, and benefits from any social program. Robust standard errors clustered at the household level are reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 6:
Spouses' Labor Supply Responses to Severe Non-Fatal Health Shocks by the Degree of Income Loss

A. Baseline Regression	Participation
Treat × Post	0.3423***
	(0.0147)
Treat × Post × Replacement Rate	-0.3529***
	(0.0149)
B. Regression with Interactions	Participation
Treat × Post × Replacement Rate	-0.3612***
	(0.0152)
Number of observations	236,897
Number of clusters	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, using a specification similar to equation (4). This replacement rate is calculated as described in the notes of Table 1. Panel A estimates a baseline differences-in-differences specification which interacts the treatment effect with the replacement rate variable. Panel B extends the specification in Panel A to include interactions of the treatment effect with additional household characteristics: age fixed effects for both spouses, year fixed effects, indicators for the number of children in the household and for the presence of adult children (18 or older) and young children (6 or younger), the spouse's gender, the spouse's months of education (and its square), and a quadratic in the household's net wealth. Both specifications include age fixed effects for both spouses, year fixed effects, and household fixed effects, and additionally include the interactions with *Post* of covariates that are interacted with *Treat × Post*, the variables among these covariates that are time variant, and the interactions of the time-variant variables with *Treat*. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household by experimental-group level are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 7:
Summary of Income Insurance Schemes for Sick Individuals and Surviving Spouses

Program	Eligibility Criteria	Benefit Levels
Social Disability Insurance	The basic eligibility criterion is a prolonged need for support that is presumed to last until the transition into the old-age pension. Since 1984 the Danish Social DI has a broad social insurance scope: it can be awarded to individuals who prove that they are unable to engage in substantial gainful activity either for medical or for non-medical (vaguely defined) social reasons. In our setting, the program effectively applies to permanently sick individuals and to surviving spouses who are determined unable to maintain their standard of living on their own (primarily widows who had a weak attachment to the labor force prior to their husband's death).	Approved applications provide benefits permanently, which in 2000, for example, amounted to DKK 72,100 ($9,000) per year for married or cohabiting individuals and DKK 98,700 ($12,300) for single individuals. These benefits are income-tested as described in Appendix Figure 9.
Privately-purchased insurance policies (in group or non-group markets)	Since 1993, most sectors covered by collective agreements (75% of the labor force) have introduced mandatory pension plans, which may include components of life insurance or 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. In addition, subject to health screenings, individuals may purchase similar insurance policies in the private non-group market.	The rates of these payouts are set by the individual pension funds. For example, some large white-collar group-market policies guarantee DKK 1,076,000 ($162,050) if the insured employees die before age 45, DKK 853,000 ($128,460) if they die between ages 45 and 54, and DKK 538,000 ($81,025) if they die between ages 55 and 66, with no transfers if the insured die at or after they reach age 67.
Voluntary Early Retirement Pension (VERP)	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 (of 10 years before 1992 which has gradually increased to 20 years thereafter) are eligible for the Voluntary Early Retirement Pension (VERP). Approximately 80% of the population is eligible for VERP.	Flat-rate annual income that amounted to roughly DKK 135,000 ($16,875) in 2000.
Old-Age Pension (OAP)	At the full-retirement age of 67 (or 65 for those born after July 1st, 1939) all residents become eligible for the Old-Age Pension (OAP).	The program provides annuities that in 2000 amounted to DKK 72,100 ($9,000) for married individuals and DKK 98,700 ($12,300) for single individuals (similar to the benefit levels paid to Social DI beneficiaries). Note that DI and OAP are different components of the same social insurance program of Social Pensions, similar to Social Security in the US, and that Social DI recipients automatically transition into the Old-Age Pension program at their full-retirement age. Benefits are income-tested as described in Appendix Figure 9.
ATP	A small government-mandated pension scheme that applies to all wage earners in Denmark.	The 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). There is a small life insurance element tied to this scheme. 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).

Appendix Table 8:
Comparison of Pre-Trends in Labor Force Participation across Affected and Unaffected Households

Year and Experimental Group Interactions	(1)	(2)	(3)	(4)	(5)
1989	0.0345***	0.0239***	0.0221***	0.0297***	0.0333***
	(0.0055)	(0.0056)	(0.0056)	(0.0068)	(0.0069)
1990	0.0270***	0.0164***	0.0150***	0.0241***	0.0268***
	(0.0054)	(0.0054)	(0.0054)	(0.0066)	(0.0067)
1991	0.0246***	0.0173***	0.0160***	0.0273***	0.0291***
	(0.0052)	(0.0052)	(0.0052)	(0.0064)	(0.0064)
1992	0.0168***	0.0113**	0.0105**	0.0160***	0.0170***
	(0.0048)	(0.0049)	(0.0049)	(0.0060)	(0.0060)
1993	0.0028	0.0002	-0.0003	0.0069	0.0073
	(0.0038)	(0.0038)	(0.0038)	(0.0048)	(0.0048)
1994	0	0	0	0	0
	0	0	0	0	0
Age FE	X	X	X	X	X
Gender	X	X	X	X	X
Education			X	X	X
Gender interactions				X	
Full gender interactions					X
Number of households	1,229,514	647,681	647,681	647,681	647,681

Notes: This table compares the labor force participation patterns prior to 1995 of a treatment group of married and cohabiting individuals born between 1930 and 1950, who experienced a heart attack or a stroke in 1995, to that of a potential control group of married and cohabiting individuals from the same cohorts who have not experienced this shock in our sample period. We run a regression with calendar year fixed effects, an experimental group indicator, and the interaction of these two sets of variables, where the baseline year is 1994. The table reports the interaction terms in the periods prior to 1995 (the year at which the treatment group experiences the shock) to compare the pre-trends across the two groups. Column 1 includes age and gender fixed effects; column 2 replicates column 1 but for the sub-sample of individuals for whom we have non-missing values for education; column 3 adds controls for months of education and its square. Column 4 includes interactions of the baseline variables (year fixed effects, the experimental group indicator, and their interactions) with gender; and column 5 includes interactions of the full set of variables (including age and education) with gender. The table shows the divergent pre-trends across affected and unaffected households, which persist even after we account for key variables. Robust standard errors clustered at the household level are reported in parentheses.
*** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Appendix A: Proof of Lemma

Planner's Problem. Denote the vector of tax rates on labor income by $T \equiv (\tau_1, \tau_2)$, and the vector of state-contingent benefits by $B \equiv (B^g, B^b)$. Let W^s denote the household's value function in state s such that $W^s \equiv \max\ U^s(c_1^s, c_2^s; l_1^s, l_2^s)$ s.t. $c_1^s + c_2^s = y^s$. Therefore, the household's expected utility is $J(B,T) \equiv \mu^g W^g + \mu^b W^b$. 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, $E(B^s) \equiv \mu^g B^g + \mu^b B^b$, equal expected taxes collected, $E\left(\sum_{i=1}^2 w_i \tau_i l_i^s\right) \equiv \mu^g\left(w_1\tau_1 l_1^g + w_2\tau_2 l_2^g\right) + \mu^b\left(w_2\tau_2 l_2^b\right)$. Hence, the planner chooses the benefit levels B and taxes T that solve

$$\max_{B,T} J(B,T) \quad \text{s.t. } E(B^s) = E\left(\sum_{i=1}^2 w_i \tau_i l_i^s\right).$$

Welfare Gains from Social Insurance. What is the welfare gain from providing more generous benefits when the bad state occurs? To answer this question, consider transferring resources from the good state g to the bad state b through a small increase in, e.g., the tax rate τ_1 to finance a balanced-budget increase in benefits in the bad state B^b. 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 the marginal increase in τ_1 that finances the additional insurance. This loss is captured by $\left|\frac{\partial J(T,B)}{\partial \tau_1}\right| = \mu^g z_1^g u_i^{g\prime}(c_i^g)$, as the household's income in state g is reduced by z_1^g dollars, which are valued at $u_i^{g\prime}(c_i^g)$ per dollar. Partially differentiating the government's budget, this marginal increase in τ_1 allows a balanced-budget increase in B^b of the amount $\frac{\partial B^b}{\partial \tau_1} = \frac{\mu^g z_1^g}{\mu^b}$.[1] The household's valuation per \$1 increase in B^b is given by $\frac{\partial J(T,B)}{\partial B^b} = \mu^b u_i^{b\prime}(c_i^b)$, since it produces a value of $u_i^{b\prime}(c_i^b)$ and is transferred to the household with probability μ^b. The utility gain from the increase in benefits when the shock occurs is, therefore, $\frac{\partial J(T,B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1^g} = \mu^g z_1^g u_i^{b\prime}(c_i^b)$.

Put together, the welfare benefits from a (balanced-budget) increase in B^b financed by an increase in τ_1 are $\frac{\partial J(T,B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1} - \left|\frac{\partial J(T,B)}{\partial \tau_1}\right| = \mu^g z_1^g\left(u_i^{b\prime}(c_i^b) - u_i^{g\prime}(c_i^g)\right)$. To gain cardinal interpretation for this expression, we follow the recent social insurance literature and normalize it by the welfare gain from decreasing the labor income tax rate in the good state, τ_1 (Chetty and Finkelstein 2013). Overall, the normalized welfare benefit from our policy change is

$$MB \equiv \frac{\frac{\partial J(T,B)}{\partial B^b} \times \frac{\partial B^b}{\partial \tau_1} - \left|\frac{\partial J(T,B)}{\partial \tau_1}\right|}{\left|\frac{\partial J(T,B)}{\partial \tau_1}\right|} = \frac{u_i^{b\prime}(c_i^b) - u_i^{g\prime}(c_i^g)}{u_i^{g\prime}(c_i^g)}. \tag{1}$$

Labor Supply Representation of Welfare Gains. We show next that the household's labor supply decisions allow us to rewrite the marginal benefit in (1) in terms of the indirectly-affected spouse's labor supply.

Lemma. *The marginal (gross) welfare gain from raising benefits in the bad state b to households (with characteristics X) can be represented by*

$$MB \cong \varphi \times \left(\frac{l_2^b}{l_2^g} - 1\right).$$

Proof. The household's optimality conditions imply that each household member's marginal utility from consumption can be mapped to the spouse's marginal disutility from labor, since $u_1^{s\prime}(c_1^s) = u_2^{s\prime}(c_2^s) = \frac{v_2^{s\prime}(l_2^s)}{w_2(1-\tau_2)}$. This allows us to represent the marginal benefit from social insurance by $MB = \frac{v_2^{b\prime}(l_2^b) - v_2^{g\prime}(l_2^g)}{v_2^{g\prime}(l_2^g)}$. Intuitively, we use the household's optimality conditions to represent the degree to which households are able

[1] The partial differentiation of the government's budget allows us to focus on the gains from social insurance. To include the costs, we would analyze the total derivative $\frac{dB^b}{d\tau_1}$ which takes into account not only the required mechanical adjustments, but also the households' behavioral responses to the policy change that have an impact on the government's budget (the "fiscal externality").

to smooth the marginal utility from consumption, $\frac{u_i'(c_i^b)-u_i'(c_i^g)}{u_i'(c_i^g)}$, using the degree to which they are able to smooth the marginal disutility from the spouse's labor, $\frac{v_2'(l_2^b)-v_2'(l_2^g)}{v_2'(l_2^g)}$. A quadratic approximation to member 2's labor disutility function around l_2^g yields the result.[2] ■

References

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

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.

[2]When the third order terms of spousal labor disutility are not small, the labor supply representation of welfare gains requires an additional term, analogous to the additional term that involves the coefficient of relative prudence in the consumption smoothing representation of welfare gains (Chetty 2006).