

# Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records

Itzik Fadlon and Torben Heien Nielsen\*

## Abstract

We provide new evidence on households' labor supply responses to fatal and severe non-fatal health shocks in the short- and medium-run. To identify causal effects, we leverage administrative data on Danish families and construct counterfactuals using households that experience the same event a few years apart. Fatal events lead to considerable increases in surviving spouses' labor supply, which the evidence suggests is driven by families who experience significant income losses. Non-fatal shocks have no meaningful effects on spousal labor supply, consistent with their adequate insurance coverage. The results support self-insurance as a driving mechanism for the family labor supply responses.

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

---

\*Fadlon: University of California, San Diego and NBER, address: Department of Economics, University of California, San Diego, 9500 Gilman Drive #0508, La Jolla, CA 92093-0508 (email: fadlon@ucsd.edu); Nielsen: University of Copenhagen, Department of Economics and the Center for Economic Behavior and Inequality (CEBI), address: Øster Farimagsgade 5 building 26, DK-1353, Copenhagen K (email: thn@econ.ku.dk). The activities of CEBI are financed by the Danish National Research Foundation. 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 Jaravel, 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.

household behavior, the degree to which households insure through labor supply is central for the design of social insurance programs (Fadlon and Nielsen 2019b). 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).<sup>1</sup>

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, to the best of our knowledge, 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.<sup>2</sup> 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 Danish 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,

---

<sup>1</sup>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).

<sup>2</sup>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.

and labor market behavior. 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 event, and over 67,000 households in which one spouse experienced a severe non-fatal event, specifically, a heart attack or stroke.

To identify the causal effects of experiencing these adverse events, we employ a quasi-experimental research design that constructs counterfactuals to affected households by using households that experience the same event 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 event, 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.<sup>3</sup>

With these data and design, we provide visually-clear estimates of individuals' labor supply responses to spousal mortality and health shocks. We additionally exploit the richness of the data to offer suggestive analysis of 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 events that we consider, the findings are all consistent with family labor supply as a self-insurance mechanism.

---

<sup>3</sup>We know of several studies in different settings that have found it useful in practice and follow an earlier version of this paper (which has been previously circulated as NBER Working Paper No. 21352). These studies include applications such as analyzing the effects of inheritance on wealth accumulation and inequality (Martinello 2017; Nekoei and Seim 2019), studying the effects of Social Security field office closings (Deshpande and Li 2019), a follow-on to our study that analyzes surviving spouses' housing decisions among elderly homeowners in Denmark (Jensen 2017), and our own work on family spillovers in health behaviors (Fadlon and Nielsen 2019a).

We begin with the focus of our study, the extreme event of the death of a spouse, which can lead to significant and permanent income losses. We find 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 event, these responses amount to average increases of 7.6% in survivors' labor force participation and 6.8% in annual labor income.

The evidence suggests that the average effects that we find are driven by households that experience substantial income shocks due to the death of a spouse, and may therefore have greater need for self-insurance through labor supply. In particular, we show that the mean increase in labor supply is attributable to survivors whose deceased spouses had provided a large share of the household's income 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 larger income losses when losing their husbands, meaningfully 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 which suggests 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 fatal events, non-fatal health events 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 18% after the shock, comparable to the findings of Dobkin et al. (2018) and Meyer and Mok (2019) in the U.S. case. However, we then show that the average decline in these households' post-transfer income is only 3.4%. 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 evidence of 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 could be a key motive for spousal labor supply responses to the financial aspects of both types of events.

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 McClellan 1998,

Gertler and Gruber 2002, Charles 2003, Gallipoli and Turner 2011, Chung 2013, Dobkin et al. 2018, and Meyer and Mok 2019) 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 events. 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;<sup>4</sup> and by comparison, we know of very little work that, similar to ours, uses rich administrative data and timing of objectively-identified health events.<sup>5</sup>

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, it also lets us observe households' overall degree of formal insurance and income loss. This provides us with the opportunity to study the potential mechanisms that may underlie the labor supply responses to spousal shocks that we find.

Second, 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* spousal unemployment.<sup>6</sup> A leading explanation for this lack of evidence has been that, in the context of temporary unemployment, income

---

<sup>4</sup>Recent related studies are McClellan (1998), 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), Braakmann (2014), Dobkin et al. (2018), and Meyer and Mok (2019). 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.

<sup>5</sup>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.

<sup>6</sup>See, for example, Heckman and Macurdy (1980, 1982), Lundberg (1985), Maloney (1987, 1991), Gruber and Cullen (1996), and Spletzer (1997).

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), Blundell et al. (2016), and Autor et al. (2019) 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 events.

The remainder of the paper is organized as follows. Section I outlines the institutional environment in Denmark and the data sources. In Section II we describe the empirical research design that we use for recovering the causal effects of adverse health events. Our core empirical analysis is presented in Section III. In this section we estimate individuals’ labor supply responses to fatal and non-fatal spousal health shocks. We also analyze the heterogeneity in these responses to offer suggestive evidence of their underlying mechanisms. Section IV discusses our findings and concludes.

## **I Institutional Background and Data**

To study labor supply responses to spousal health events 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 environment as it relates to income coverage of losses among sick individuals and surviving spouses, and we list our data sources.

### **A 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 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 main features and benefit levels of these components, and we provide some more details of these programs in Online Appendix E.

Individuals who experience a health 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). We note that whereas Social DI benefits are income-tested against current income (as described in the online appendix), they are flat-rated with respect to the beneficiary's earnings history, unlike other countries such as the U.S. where DI benefits are a function of previous earnings.

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. 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. 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). 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 are determined unable to maintain their standard of living on their own. Indeed, we find (and document later) significant increases in the receipt rate of Social DI by survivors in the year their spouses die. The share of surviving spouses who end up on social DI (for any reason) in the year of the spousal death event is 25.5% among survivors younger than the full retirement age (when the program automatically transitions to the Old-Age Pension).

Another source of income to households that experience health shocks or in which a member dies is potential payments from employer-based insurance policies, which have recently become more prevalent through labor market pension packages that sometimes include them. Employer-based pension plans were generally common throughout our sample period in the public sector, which composes 30% of the Danish labor market and is covered by collective agreements. In addition, based on a reform that was announced in 1987 and was implemented in 1993, most of the segment of the private sector that is covered by collec-

tive agreements (75% of the labor force in the private sector) has introduced mandatory defined-contribution pension plans, some of which may include components of life insurance or insurance against specific health events. These latter schemes pay out to sick workers who experience a severe health event, or to a surviving spouse in case the plan member dies, at rates that are set by the individual pension funds. We note that such policies (when offered) are part of mandated plans, and are hence arguably less likely prone to adverse selection. In addition, these collective-agreement pension funds are non-profit organizations controlled by labor unions, their policies involve reduced administrative costs due to their size, and they do not require marketing expenditure. Taken together, these group-market policies are presumably closer to being actuarially fair. Subject to health screenings, individuals may also purchase insurance policies in the private non-group market. Overall, in combination and within our sample period, in which the average year of spousal death is around 1995, the life-insurance coverage rate is generally low compared to the larger market today, since life-insurance holdings have been increasing more recently with the gradual expansion of schemes through labor market pensions.<sup>7</sup>

It is worth noting that despite the private market for life insurance in Denmark, there is still rationale for government intervention in the context of spousal mortality, specifically 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. Such rejections by insurance companies, which represent a form of a market failure in life insurance, can be explained by private information held by rejected households as has been shown in other settings (Hendren 2013). 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. 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. The combination of these features of the private insurance market and the lack of a universal-coverage social survivors insurance scheme leaves Danish surviving spouses more vulnerable to the financial shocks imposed by spousal death as they age.<sup>8</sup>

---

<sup>7</sup>Payouts from these private (group or non-group) insurance policies are observable in our data as part of households' liquid wealth, and we later gauge the effective coverage rate within our sample by assessing actual changes in household liquid wealth balances around spousal death.

<sup>8</sup>In our setting, while survivors in older households experience smaller losses through foregone labor



Lastly, there exist 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 (amounting to approximately 80% of the population) are eligible for the Voluntary Early Retirement Pension (VERP). 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 structure and level of benefits provided through these two programs are detailed in Online Appendix E.

## B Data Sources and Variables

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 events we use two complementary datasets. Our first dataset is the *Death Registry*, which includes deceased individuals' date of death. Our second dataset is the *National Patient Registry*, which covers all hospitalization records 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 (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.<sup>9</sup>

---

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) that are not automatically transferred to them.

<sup>9</sup>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 events, 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, we have found that 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

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 the Early Retirement Age of 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.

Our main labor supply outcomes are labor force participation and annual earnings. Labor force participation is defined as having any positive level of annual earnings. For individuals younger than the Early Retirement Age we use the administrative measure of any employment (based on full-time/part-time measures). For earnings, we need to make an adjustment in the year of spousal death, since Statistics Denmark includes the deceased’s last month of earnings as part of the surviving spouse’s annual earnings that year. We subtract from the surviving spouse’s earnings in the year of the event an approximation to this amount using the pro-rated earnings of the deceased from that year based on their month of death. 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 U.S. \$1.

## II Research Design

The goal of our empirical analysis is to identify the dynamic causal effects of fatal and severe non-fatal spousal health events 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. Some additional illustrations and robustness analysis are summarized in Online Appendix B. We also describe in this section our resulting analysis samples of treatment and control groups for the different health events that we study.

### A Quasi-Experiment

The ideal experiment for identifying the short- and medium-run effects of spousal health events would randomly assign events to households and track labor supply responses over

---

caused by the event (as compared to the counterfactual). In our analysis of labor supply responses we therefore focus on imposed income losses.

time. Therefore, we need to compare the ex-post responses to events of affected households to a counterfactual behavior of (hypothetical) ex-ante similar unaffected households. This requires comparing households with similar expectations over the distribution of future paths, but with different realizations, to isolate the effect of the event. 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 aims to mimic this ideal experiment, by exploiting the potential randomness of the timing of a severe (fatal or non-fatal) health event within a short period of time.

To do so, we look only at households that have experienced the events that we consider at some point in our sample period, and identify the treatment effect from the timing at which the event was realized. Specifically, we construct counterfactuals to affected households using households from the same cohorts that experience the same event 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 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.

Figure 1 plots the outcome for this treatment group over time, and compares it to the time trend of the outcome for 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 online Appendix Table B.1). 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, Figure 1 additionally plots the outcome for households that experienced the same shock in 2010 (15 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  $1995+\Delta$  as a control group for households that experienced a shock in 1995.

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  $\tau$ , and a matched control group composed of households from the same cohorts that experience the same shock but in year  $\tau+\Delta$ . 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.

By design, the same household may appear both in the treatment group and in the control group, which we allow for increased statistical power. We note, however, that a household is never used as a control to itself. For example, if treated households that experienced an event in 1995 (who are matched with households that experienced an event in 2000 as controls) are included also in the control group, it is only since households that experience an event in 1990 are included in the treatment group as well. Nevertheless, in our estimations we report robust standard errors clustered at the household level, so that we avoid including the same household in multiple clusters. We also repeat our main analysis using treatment and control groups that do not overlap, by randomizing households to only one experimental group, with similar findings (see online Appendix Table B.2).

There is a key trade-off in the choice of  $\Delta$ . On the one hand, we would want to choose a smaller  $\Delta$  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  $\Delta = 1$ . On the other hand, we would want to choose a larger  $\Delta$  in order to be able to identify longer-run effects of the shock, since for each chosen  $\Delta$  the estimation strategy provides estimates for up to period  $\Delta - 1$ . In the current illustration, for example, using those who experienced a shock in 2005 ( $\Delta = 10$ ) will allow us to estimate the effect of the shock for up to 9 years. However, this could entail a potentially larger bias since the pre-trend in the behavior of such a group would not be as tightly parallel to that of the treatment group. Our choice of  $\Delta$  is five years, such that we can identify effects up to four years after the shock. We assess the robustness of our analysis to this choice and find that local perturbations to  $\Delta$  provide very similar

results (see online Appendix Figure B.1 and Appendix Table B.3).

*Formal Description of the Design and Estimator.* Fix a group of cohorts, denoted by  $\Omega$ , and consider estimating the treatment effect of an event experienced at some point in the time interval  $[\tau_1, \tau_2]$  by individuals who belong to group  $\Omega$ . We refer to these individuals’ households as the treatment group and divide them into sub-groups indexed by the year in which the event 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 event—that is,  $t = year - \tau$ , where  $year$  is the calendar year of the observation. As a control group, we match to each treated group  $\tau$  the households of individuals from the same cohort group  $\Omega$  who experienced the same event but at  $\tau + \Delta$ , for a given choice of  $\Delta$ . For these households we assign a “placebo” event 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 event occurs at  $t = \Delta$ ).

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 event which we denote by  $p$  (for “prior”). For any period  $n > 0$ , the treatment effect  $\gamma_n$  can be simply recovered by the difference-in-differences estimator

$$(1) \quad \gamma_n \equiv (y_n^T - y_n^C) - (y_p^T - y_p^C).$$

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  $\Delta$  puts an upper bound on  $n$  such that  $n < \Delta$  (since the control group becomes “treated” at  $t = \Delta$ ).

The identifying assumption is that, absent the realization of the event, 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  $\Delta$  the particular year at which the event occurs may be as good as random. Similar timing-based identifying assumptions have been exploited in numerous previous papers within a variety of settings.<sup>10</sup> To test the validity of our assumption in the current setting, we accompany our empirical analysis with plots of the treatment and control groups’ behavior in the five years prior to the event in order to assess their co-movement in the pre-event periods. We consistently show 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 event within

---

<sup>10</sup>Among many others, these include Ruhm (1991), Grogger (1995), Hilger (2016), and Persson and Rossin-Slater (2018) in the context of household shocks, and papers such as Guryan (2004) and Bailey and Goodman-Bacon (2015) in the context of program rollout.

our chosen five-year window of  $\Delta$ .

It is worth noting that the research design does not preclude behavioral adjustments in expectation of an event among treated households; nor do our results imply there are no such adjustments 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 event, 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 expectation paths.

*Estimating Equation for Average Effects.* To visualize the estimation strategy, we first provide figures that plot the raw data for the main analysis, together with estimates for the corresponding dynamic difference-in-differences regressions based on equation (1). For the remainder of the analysis, we then turn to regression estimations that quantify the mean treatment effects with an estimating equation of a standard difference-in-differences form:

$$(2) \quad y_{i,t} = \alpha_i + \beta post_{i,t} + \gamma treat_i \times post_{i,t} + \delta X_{i,t} + \varepsilon_{i,t}.$$

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-event periods;  $\alpha_i$  is a household fixed effect (which absorbs any time-invariant characteristic including the “main effect” of  $treat_i$ ); and  $X_{i,t}$  denotes a vector of potential (time-variant) controls for robustness checks. The parameter  $\gamma$  represents the average causal effect of spousal health events on household outcomes.<sup>11</sup>

## B Analysis Sample

Our sample is composed of all households in which one spouse experienced an event from year 1985 to 2011. Our main sample is comprised of all households in which one spouse died and was between ages 45 and 80 in the year of the (actual or placebo) event. It 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

---

<sup>11</sup>In these estimations we present the medium-run effects that are the focus of our analysis, so that we let  $post_{i,t}$  assume the value 1 for periods 2 to 4.

three years. These health shocks are commonly studied as their timing within a short period of time is likely unpredictable (Chandra and Staiger 2007; Doyle 2011). 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 and not by eligibility for early retirement benefits. Our sample of non-fatal shocks includes 35,143 households in the treatment group and 52,196 households in the control group. Online Appendix Table A.1 displays summary statistics of key demographics for the analysis samples and illustrates their comparability.

### III Family Labor Supply Responses to Severe Health Events

In this section, we present our primary analysis of the impact of fatal and severe non-fatal health events on spousal labor supply. We begin with the focus of our study, the extreme event of spousal death which can lead to large and permanent income losses, and we study the labor supply responses of surviving spouses. We provide complementary analysis of heterogeneity in these responses to shed light on potential mechanisms, focusing on self-insurance and the degree of income loss imposed by spousal death. Then, we study family labor supply responses to our secondary set of events, severe non-fatal health shocks. 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.

#### A Spousal Labor Supply Responses to Fatal Health Events

##### A.1 Overall Responses

Figure 2 plots the labor supply responses of surviving spouses in our sample of fatal health events. The structure of this figure is as follows. The x-axis denotes time with respect to the event, normalized to period 0. For the treatment group, period 0 is when the actual event occurs; for the control group period 0 is when a “placebo” event occurs (while their actual event occurs in period 5). The gray line plots the behavior of the control group (along with the corresponding 95-percent confidence intervals). 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-event level of the treatment group’s outcome (in period  $t = -2$ ). This normalized counterfactual is displayed by the blue line and squares. The red line and

circles plot the behavior of the treatment group (along with the corresponding 95-percent confidence intervals). Online Appendix Table B.2 additionally provides the estimations of the corresponding dynamic difference-in-differences regressions. Note that since individuals die at different points throughout the calendar year, period 0 is a “transitional” year, and period 1 is the first year in which all households in the treatment group have fully experienced the event. Therefore, one should consider period 1 as the initial period that captures a “full” impact of the event.

Figure 2 first provides a visual verification of parallel trends across the treatment and control groups prior to period 0 (as verified statistically in online Appendix Table B.2). Then, analyzing the effect of the event, panel A reveals an increase in survivors’ labor force participation which is already apparent in the year of spousal death. By the fourth year after the event, 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).

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 event. Indeed, Panels A-B of Figure 3 and panel A of Table 1 reveal clear differences in the responses of widowers (whose wife dies) and widows (whose husband dies). While on average widowers do not change their labor force participation when their wife dies (and, if anything, slightly decrease their annual earnings), widows meaningfully increase their labor supply following the death of their husband. Four years after the event, 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, panels C-D of Figure 3 and panel B of Table 1 analyze 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 with 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 first analyze the household’s overall income, holding the surviving spouse’s earnings and social benefits at their pre-event level.<sup>12</sup>

---

<sup>12</sup>Specifically, we fix the surviving spouse’s labor income, Social Disability, and Social Security benefits at their level in  $t = -1$ .



Before discussing the results, it is useful to mention benchmarks for the changes that we observe in household income in order to interpret their magnitude. Specifically, we interpret the magnitudes based on equivalence scaling which translates income at the household level into income at the individual-equivalent level, to account for the household’s compositional change following the fatal health event. 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.<sup>13</sup>

The results show 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 meaningful 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, we take into account the surviving spouses’ labor supply responses and any change in the social benefits they may receive. We find that widowers experience an actual decline of 30.5 pp; and that widows manage to decrease their additional potential loss that totaled to 40 pp—through the increase in labor supply and higher take-up of social insurance—to incur an actual decline of about 35 pp, while scale economies suggest a fall of 29-33 pp will not harm consumption.

## A.2 Potential Mechanisms and Mitigating Factors

We next offer heterogeneity analysis to provide suggestive evidence of the mechanisms through which survivors’ labor supply responses may operate. We first analyze how survivors’ behavior varies by the degree of income loss their spouse’s death imposes and by the extent of coverage through social insurance to investigate the potential role of self-insurance. We then use a simple test that aims to assess the extent to which survivors’ willingness to work may change in response to the event. Lastly, we investigate and discuss other alternative channels and characteristics that may induce or attenuate spousal labor supply responses to fatal health events. Overall, the results provide consistent evidence in support of the self-insurance mechanism hypothesis. Online Appendix C summarizes these results and provides additional estimation details.

*Responses by Degree of Income Loss.* We begin by studying the effect of the death of a spouse on labor supply by the degree of potential income loss. First, we analyze among

---

<sup>13</sup>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.

prime-age households prior to the early retirement age, how surviving spouses' labor supply responses may differ by whether the deceased was the primary earner. Online Appendix Figure C.1 and Appendix Table C.1 show that spouses who lose a primary earner increase their labor supply, consistent with a greater need to self-insure; whereas those who lose a secondary earner and were primary earners themselves even decrease their labor supply, consistent with the idea that their earnings are no longer necessary for consumption by two people.<sup>14</sup> Second, we provide analysis that estimates the heterogeneity in labor supply responses by the household's overall income replacement rate. This analysis incorporates income from any source (not only labor earnings), both in pre-event and the post-event years, and directly accounts for changes (i.e., losses) in income at the household level due to the event. The results in online Appendix Table C.2 consistently show the strong partial correlation between labor supply responses and income losses. Survivors in households with lower potential income replacement rates, who experience larger income losses, are much more likely to increase their labor force participation in response to the event. Specifically, it implies larger increases in spousal labor supply among households in which the deceased had provided a larger share of the household's overall income.<sup>15</sup> 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 this estimation to match on pre-event own and spousal income would lead to similar average responses across genders.<sup>16</sup>

*Variation in Social Insurance.* We now ask whether better social insurance attenuates the labor supply responses in our context of spousal mortality. We take advantage of spatial variation in the administration of social survivors benefits through the Social DI program.<sup>17</sup> The program and its effective survivors benefits component are locally administered, where regional councils (in a total of 15 regions) 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

---

<sup>14</sup>Online Appendix Table C.1 additionally provides estimates of these gradients in responses by the deceased's earnings status for each gender with similar conclusions.

<sup>15</sup>Since controlling for additional interactions does not change the results much, the evidence suggests that the heterogeneous responses may indeed be driven by differential income replacement rates.

<sup>16</sup>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.

<sup>17</sup>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, and to survivors under 67, the age at which the program automatically transitions into the Old-Age Pension for the studied population. Since the increase in the take-up of the program following the event is attributable to females within this sample, we focus the analysis on widows. Panel A of online Appendix Figure C.2 displays the aggregate insurance role of Social DI for these widows, whose take-up of the program increases by 49% in the year their husbands die.

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 we assign to a widow in a treated household, who resided in a given municipality prior to the event, the respective mean of that municipality in a given period excluding her own benefits (the “leave-one-out” mean).<sup>18</sup> The results are presented in online Appendix Table C.3. The regression estimate translates to a participation elasticity with respect to social benefits of  $-.26$ .<sup>19</sup> This suggests that formal social insurance provided to survivors substitutes for labor supply increases, which could otherwise provide an informal self-insurance mechanism against loss of income following the death of a spouse.

*Changes in Spousal Labor Disutility.* Besides income losses, there are other important ways in which households can be directly affected by mortality events that can impact 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. We briefly discuss a simple intuitive 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 event, 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 they may be more prone to experience social isolation or loneliness following the event. Similarly, if the deceased spouses in these households did not work prior to the event 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

---

<sup>18</sup>The variation in this instrument is displayed in Panel B of online Appendix Figure C.2. 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.

<sup>19</sup>For a sense of scale, estimates for the net-of-tax participation elasticity are on the order of  $.25$  (see, e.g., Chetty 2012).

husbands die. Overall, survivors in this first sub-set of households (in which both spouses did not work prior to the event) are presumably more likely to experience a decrease in the utility cost of labor supply. 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 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.

We first verify in online Appendix Table C.4 our 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). Then, studying their labor supply, we find that the increase in the labor force participation of survivors in households in which the deceased worked is much larger, with a negligible effect for survivors in households in which the deceased did not work. Hence, in the context of this stylized suggestive test, the evidence is inconsistent with the conjecture that a lower cost of labor following the event drives the 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 could be driven by self-insurance when the event leads to large income losses.

*Ability to Respond.* Next, we wish to consider the possibility that the cross-sectional response heterogeneity may be due to survivors' ability to respond with labor supply increases.<sup>20</sup> To this end, we repeat the analysis by household replacement rates among sub-groups of households with survivors who 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 constrain the sample to spouses who did not work before the event and study participation, for extensive margin responses; and we then constrain the sample to spouses who worked part time before the event and study 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 event, and spouses who had positive earnings before the event. In all these estimations we find strong negative correlations with household replacement rates (see online Appendix Table C.5). Overall, while clearly one cannot rule out this alternative explanation, we think

---

<sup>20</sup>This could be the case, for example, if spouses in households with smaller income losses and higher replacement rates also have higher baseline participation rates, work hours, or rates of full-time employment.

the evidence highlighted here reinforces self-insurance as a likely potential mechanism.

*Presence of Children.* Lastly, we are interested in investigating whether the presence or age of children may attenuate widows’ labor supply increases following their spouse’s death. For example, having younger children may increase the costs of supplying labor and lower mothers’ ability to increase their labor supply due to care-giving; whereas having older children may provide informal insurance possibilities (through, e.g., potential support from adult children or moving in with them). Online Appendix Table C.6 estimates an equation where we interact the effect of spousal death with indicators for the presence of children and for having children in different age ranges. We find that presence of younger children meaningfully hinders labor supply increases following a spousal death. There does not seem to be evidence that women with older children respond differentially in their labor supply.

## B Household Labor Supply Responses to Severe Non-Fatal Health Shocks

In our secondary set of results, we study household labor supply responses to severe non-fatal health shocks. Figure 4 plots these outcomes for both the sick individual and the spouse. Table 2 summarizes the effects by estimating difference-in-differences regressions 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 labor supply responses shown in Figure 4. Columns 2 and 4 in panel A of Table 2 reveal 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 35,467 (\$4,433)—a significant drop of 18%. Dobkin et al. (2018) and Meyer and Mok (2019) find similar-magnitude effects of health shocks on own earnings in their U.S. context.

However, while there is a significant drop in the sick individuals’ earnings, columns 5 and 6 in panel A show that the actual loss of income that these households experience amounts to only 3.4% 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 Table 2).<sup>21</sup> While fatal and non-fatal health events differ in many

---

<sup>21</sup>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 García-Gómez et al. (2013)

aspects indeed (e.g., in their effects on the household’s composition), in terms of the financial aspects of the events 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 summarized in Online Appendix D. First, we study how families differentially respond to non-fatal shocks with divergent degrees of severity, as defined by hospitalization days (see online Appendix Table D.1).<sup>22</sup> 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 care-giving motives or other 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. 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).<sup>23</sup> 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 (in online Appendix Table D.3). 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.

---

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.

<sup>22</sup>Specifically, we divide households by the shock’s severity according to the 75th percentile of the distribution of hospitalization days associated with the shock (10 days in our sample). Consistent with this measure proxying for the morbidity of the shock, we show in online Appendix Table D.2 that individuals with “higher severity” are considerably more likely to go on disability insurance and to utilize primary health care. In addition, using the Health and Retirement Study, we investigate the association between the length of hospital stays (for individuals that have experienced a heart attack or a stroke) and the inability to perform daily activities. We show strong associations that remain similar when we add controls for age, gender, and education.

<sup>23</sup>This variation is in part due to the fact that, although DI benefits are tested against current income, they are flat-rated with respect to the sick spouse’s earnings history and are not a function of it (in contrast to, e.g., the U.S. case).

## IV Discussion and Conclusion

This paper provides new evidence on households' labor supply responses to fatal and severe non-fatal health events. Studying the critical event of the death of a spouse, we find large increases in the surviving spouses' labor supply. The evidence suggests that these responses are driven by households for whom the event imposes significant income losses, as we find that the mean increase in labor supply is attributable to survivors whose deceased spouses had provided a large share of the household's income and who are less formally insured by government transfers. We also provide some suggestive 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. Then, analyzing households in which an individual has experienced a severe health shock but survived, for whom income losses are well-insured (primarily through DI benefits), we find no significant spousal labor supply responses. We also find that, while the mean decline in household income is negligible in the pooled sample, there is still cross-household variation in overall income replacement rates and that there is a strong partial correlation between spousal labor supply responses and the imposed income losses also in the context of non-fatal shocks. Overall, 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.

As such, our findings on family labor supply responses to adverse health events and their patterns have direct implications for positive models of household self-insurance behavior over the life cycle. Additionally, the findings may also point to ways to target government transfers through the existing social insurance programs in Denmark. In Fadlon and Nielsen (2019b) we show that the extent to which households self-insure against realized income losses using spousal labor supply is proportional to the degree to which they lack insurance and would gain from more generous government benefits.<sup>24</sup> Exploiting this logic, it may be welfare improving to provide more generous transfers to surviving spouses in Denmark, given the meaningful labor supply increases we have found among widows. Also, since survivors' labor supply is increasing in the share of the household's income that the deceased had provided, there is an argument for survivors benefits to depend on the deceased spouse's work history.<sup>25</sup> Our findings indicate a similar pattern of heterogeneity in responses to

---

<sup>24</sup>We note that from a life-cycle perspective, the target empirical moment is average spousal labor supply in the periods households spend in each state of the world, whereas by the nature of our empirical design we are able to identify effects up to the medium run (of four years).

<sup>25</sup>This is consistent with Persson (2020) 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

non-fatal spousal health shocks, suggesting that disability benefits in Denmark may also be more efficiently distributed if made dependent on the disabled spouse’s work history (unlike their structure today). Evidently, these features characterize the current large survivors and disability insurance schemes within the Social Security system in the U.S. Of course, a full welfare analysis of any policy change should also consider its welfare costs, as captured by the impact of households’ behavioral responses to the policy, or crowd out effects, on the government’s budget.

Finally, it is useful to discuss the Danish results in an international context. While our analysis is informative for household self-insurance behavior more broadly, the extent of surviving spouses’ labor supply responses may naturally differ across contexts for various reasons. For example, female labor force participation rates and extent of full-time/part-time work differ across economies. This could imply that widows in Denmark, where female participation rates are distinctly high, might be more likely prone to a “ceiling effect” in their ability to respond to negative income shocks; as compared to widows in the U.S., for example, where female participation is lower. Another important dimension of potential differences across countries is their social safety net more broadly and, specific to our application, whether there exists an explicit social survivors insurance program which would determine the need to self-insure. The U.S., for example, has the Social Security survivors benefits program that paid more than \$93 billion to over 4 million surviving spouses in 2014. In contrast, Denmark, which is an outlier in terms of the generosity of social insurance, nevertheless has no explicit survivors insurance program. Still, in practice, widows can apply for DI for social reasons as we discussed above. The different schemes of effective income loss coverage hence also differ in other dimensions, e.g., in their proportionality to the deceased’s earnings history. As part of their work, Fadlon et al. (2019) provide some analysis that offers one useful comparison. They find that, on average, self-insurance through spousal labor supply plays a small role in the U.S. Consistent with our findings, they show based on eligibility-age criterion, that widows who are eligible for survivors benefits exhibit no changes in labor supply following a spousal death whereas just-ineligible widows exhibit a non-negligible increase earnings (in magnitudes similar to what we find here). Ultimately, the quantitative nature of widows’ labor supply responses to spousal death across different countries is naturally context dependent and is subject to empirical investigations.

---

includes such a scheme.



## References

- Ashenfelter, O. (1980). Unemployment as disequilibrium in a model of aggregate labor supply. *Econometrica* 48(3), 547–564.
- Autor, D., A. Kostøl, M. Mogstad, and B. Setzler (2019). Disability benefits, consumption insurance, and household labor supply. *American Economic Review* 109(7), 2613–54.
- 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.
- Bengtsson, S. (2002). Bestemmer forvaltningen om du faar foertidspension? -kommunens forvaltningspraksis og tilkendelse af foertidspension. Technical Report 02:15, Socialforskningsinstituttet.
- 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.
- 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.
- 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.
- 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.
- 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 (2019). Who is screened out? application costs and the targeting of disability programs. *American Economic Journal: Economic Policy* 11(4), 213–48.
- Dobkin, C., A. Finkelstein, R. Kluender, and M. J. Notowidigdo (2018). The economic consequences of hospital admissions. *American Economic Review* 108(2), 308–52.
- 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 (2019a). Family health behaviors. *American Economic Review* 109(9), 3162–91.
- Fadlon, I. and T. H. Nielsen (2019b). Household labor supply and the gains from social insurance. *Journal of Public Economics* 171, 18–28.
- Fadlon, I. and T. H. Nielsen (2020). Family labor supply responses to severe health shocks: Evidence from danish administrative records—dataset. *American Economic Journal: Applied Economics*.
- Fadlon, I., S. P. Ramnath, and P. K. Tong (2019). Market inefficiency and household labor supply: Evidence from social security’s survivors benefits. Working Paper 25586, National Bureau of Economic Research.
- 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* 132(3), 1165–1217.
- 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). 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.
- Jäger, S. (2016). How substitutable are workers? evidence from worker deaths.
- Jensen, A. (2017). Loaded but lonely: Housing and saving responses to spousal death in old age.
- 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.
- 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?
- McClellan, M. B. (1998). Health events, health insurance, and labor supply: Evidence from the health and retirement survey. In *Frontiers in the Economics of Aging*, pp. 301–350. University of Chicago Press.


- Meyer, B. D. and W. K. Mok (2019). Disability, earnings, income and consumption. *Journal of Public Economics* 171, 51–69.
- Nahum, R.-A. (2007). Labour supply response to spousal sickness absence.
- Nekoei, A. and D. Seim (2019). How do inheritances shape wealth inequality? theory and evidence from sweden.
- 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. (2020). Social insurance and the marriage market. *Journal of Political Economy* 128(1), 252–300.
- Persson, P. and M. Rossin-Slater (2018). Family ruptures, stress, and the mental health of the next generation. *American Economic Review* 108(4-5), 1214–52.
- 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. (2002). Worker displacement and the added worker effect. *Journal of Labor Economics* 20(3), 504–537.
- White House (2015). Fiscal year 2016: Historical tables – budget of the US government.

Figure 1: Illustration of the Quasi-Experimental Research Design



Notes: This figure compares the labor force participation of a treatment group of individuals born between 1930 and 1950, who experienced a heart attack or a stroke in 1995, to that of potential control groups. These groups include individuals from the same cohorts who have not experienced this shock within our sample period (1980-2011) and those who have experienced the same shock but in different years. The figure reveals the similarity of households that experience shocks closer in time.

Figure 2: Spousal Labor Supply Responses to Fatal Health Events



Notes: These figures plot spouses' labor supply responses to fatal health events. The sample includes households in which one spouse died between years 1985 and 2011 and was age 45 to 80 in the year of the (actual or placebo) event. Panel A depicts labor force participation and panel B depicts annual earnings. The x-axis denotes time with respect to the event, normalized to period 0. For the treatment group, period 0 is when the actual event occurs; for the control group period 0 is when a "placebo" event occurs (while their actual event occurs in period 5). The gray line plots the behavior of the control group (along with the corresponding 95-percent confidence intervals). 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-event 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 (along with the corresponding 95-percent confidence intervals).

Figure 3: Household Responses to Fatal Health Events by Gender

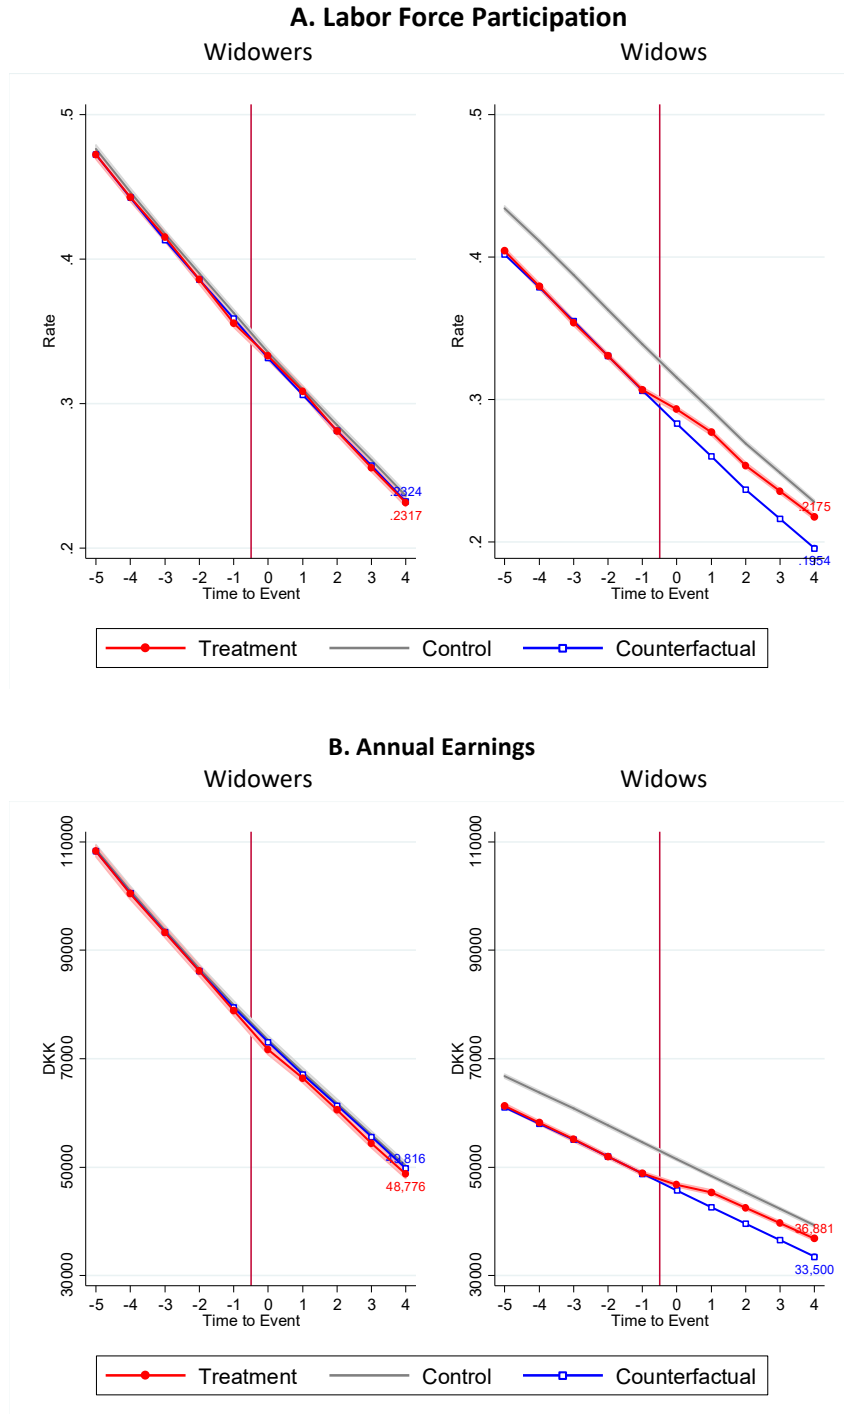
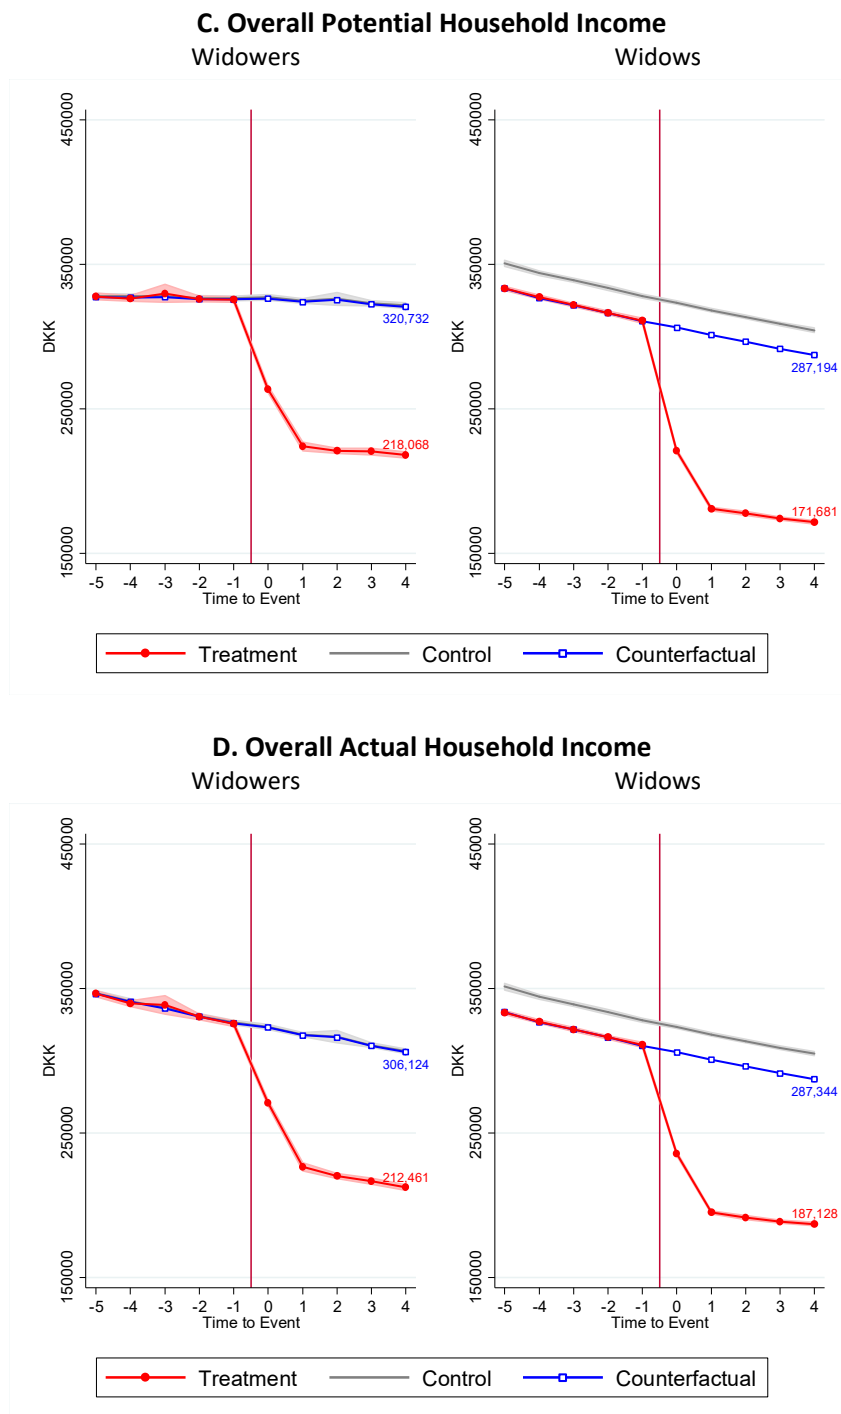


Figure 3: Household Responses to Fatal Health Events by Gender—*Continued*

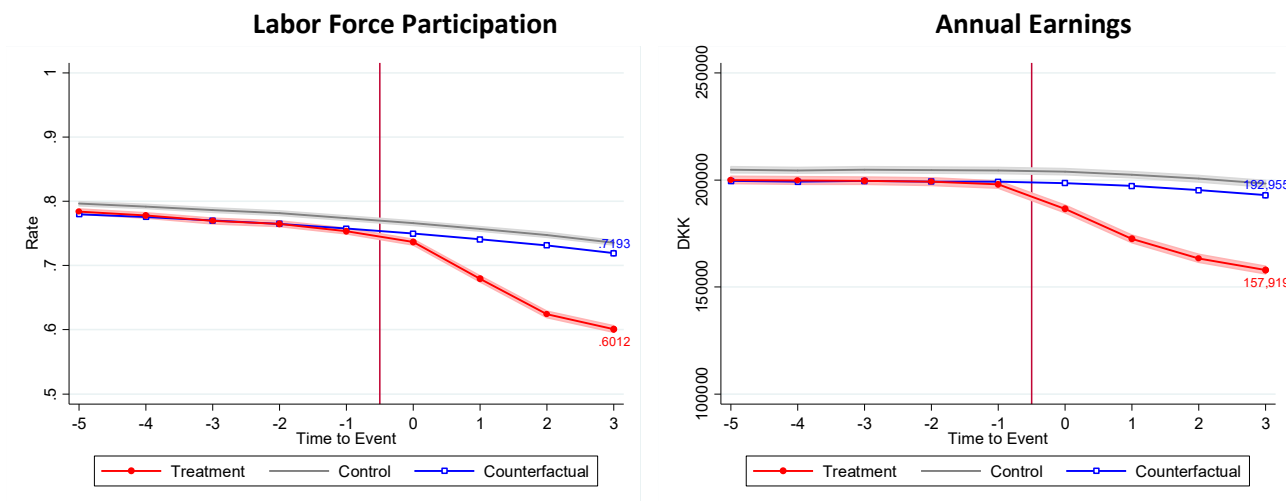


Notes: These figures plot the evolution of different household outcomes in response to fatal health events by the gender of the surviving spouse. The sample includes households in which one spouse died between years 1985 and 2011 and was age 45 to 80 in the year of the (actual or placebo) event. This figure is constructed as described in the notes of Figure 2.




Figure 4: Household Labor Supply Responses to Severe Non-Fatal Health Shocks

**A. Sick Individual**



**B. Spouse**



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. This figure is constructed as described in the notes of Figure 2.

Table 1: Household Responses to Fatal Health Events by Gender

<i>Panel A: Labor Supply</i>								
	Widowers				Widows			
	Participation (1)	Participation (2)	Earnings (3)	Earnings (4)	Participation (5)	Participation (6)	Earnings (7)	Earnings (8)
Treat × Post	-0.0016 (0.0016)	-0.0017 (0.0016)	-939 (460)	-906 (430)	0.0188 (0.0010)	0.0164 (0.0010)	2,957 (190)	2,707 (181)
Year and Age FE		X		X		X		X
Counterfactual	0.2577	0.2578	55,535	55,502	0.2167	0.2191	36,740	36,990
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	166,703	166,703	166,703	166,703	335,670	335,670	335,670	335,670

<i>Panel B: Household Income</i>								
	Widowers				Widows			
	Potential (1)	Potential (2)	Actual (3)	Actual (4)	Potential (5)	Potential (6)	Actual (7)	Actual (8)
Treat × Post	-103,714 (1,989)	-103,257 (1,978)	-94,874 (1,992)	-94,590 (1,979)	-117,485 (780)	-119,006 (782)	-102,526 (779)	-104,860 (780)
Year and Age FE		X		X		X		X
Counterfactual	323,597	323,140	311,379	311,095	292,056	293,577	291,703	294,037
Number of observations	1,396,657	1,396,657	1,396,657	1,396,657	2,919,243	2,919,243	2,919,243	2,919,243
Number of clusters	166,638	166,638	166,638	166,638	335,533	335,533	335,533	335,533

Notes: This table reports estimates of household responses to fatal health events by the gender of the surviving spouse, using the specification of equation (2). 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. All specifications include households fixed effects, and robust standard errors clustered at the household level are reported in parentheses.

Table 2: Household Responses to Severe Non-Fatal Health Shocks

	<i>Panel A:</i>					
	Sick Individual				Household Income	
	Participation		Earnings		Short Run	Medium Run
	Short Run	Medium Run	Short Run	Medium Run		
(1)	(2)	(3)	(4)	(5)	(6)	
Treat × Post	-0.0862 (0.0027)	-0.1199 (0.0031)	-28,903 (884)	-35,467 (995)	-9,342 (2,216)	-17,077 (2,667)
Counterfactual	0.7384	0.7215	197,068	193,729	503,607	505,155
Number of observations	521,609		521,609		521,609	
Number of clusters	67,749		67,749		67,749	

	<i>Panel B:</i>			
	Spouse			
	Participation		Earnings	
	Short Run	Medium Run	Short Run	Medium Run
(1)	(2)	(3)	(4)	
Treat × Post	-0.0012 (0.0024)	-0.0060 (0.0027)	-1,599 (639)	-1,778 (702)
Counterfactual	0.7541	0.7426	168,649	168,355
Number of observations	521,609		521,609	
Number of clusters	67,749		67,749	

Notes: This table reports estimates of household responses to severe non-fatal health shocks. 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 4. All specifications include household fixed effects, and robust standard errors clustered at the household level are reported in parentheses.