HOUSEHOLD RESPONSES TO SEVERE HEALTH SHOCKS AND THE DESIGN OF SOCIAL INSURANCE*

Itzik Fadlon and Torben Heien Nielsen

Abstract

This paper studies how households respond to severe health shocks and the insurance role of spousal labor supply. In the empirical part of the paper, we provide new evidence on individuals' labor supply responses to spousal health and mortality shocks. Analyzing administrative data on over 500,000 Danish households in which a spouse dies, we find that survivors immediately increase their labor supply and that this effect is entirely driven by those who experience significant income losses due to the shock. Notably, widows – who experience large income losses when their husbands die – increase their labor force participation by more than 11%, while widowers – who are significantly more financially stable – decrease their labor supply. In contrast, studying over 70,000 households in which a spouse experiences a severe health shock but survives – for whom income losses are well-insured in our setting – we find no economically significant spousal labor supply responses, suggesting adequate insurance coverage for morbidity (vs. mortality) shocks. In the theoretical part of the paper, we develop a method for welfare analysis of social insurance using only spousal labor supply responses. In particular, we show that the labor supply responses of spouses fully identify the welfare gains from insuring households against health and mortality shocks. Our findings imply large welfare gains from transfers to survivors and identify efficient ways for targeting government transfers. *JEL* codes: H00, I10, J10, J20.

*We thank David Autor, Martin Browning, Raj Chetty, Jeffrey Clemens, Julie Cullen, David Cutler, Gordon Dahl, Ariel Dora Stern, Raluca Dragusanu, Liran Einav, Eric French, Elhanan Helpman, Nathan Hendren, Simon Jäger, Xavier Jaraverl, Larry Katz, Claus Thustrup Kreiner, David Laibson, Jessica Laird, Søren Leth-Petersen, Magne Mogstad, Joana Naritomi, Arash Nekoei, Jann Speiss, Yoram Weiss, Gal Wettstein, Danny Yagan, and Eric Zwick for helpful discussions and comments. We gratefully acknowledge financial support from the Alfred P. Sloan Foundation through the NBER Household Finance working group (under the direction of Brigitte Madrian and Stephen Zeldes) and the NBER Pre-doctoral Fellowship Program on the Economics of an Aging Workforce (under PI David Card), the Boston College Center for Retirement Research, the NBER Roybal Center for Behavior Change in Health and Savings, and EdLabs at Harvard University. The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium. The opinions and conclusions expressed are solely those of the authors and do not represent the opinions or policy of SSA or any agency of the Federal Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States Government or any agency thereof. The research reported in this publication was also supported by the National Institute on Aging of the National Institutes of Health under Award Number P30AG034532. The content is the sole responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. Corresponding author: Itzik Fadlon, email: fadlon@fas.harvard.edu, address: National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138. Total word count: 18,294.

I Introduction

Does the labor supply of household members insure against adverse shocks? The answer to this question is important for our understanding of household behavior and is central to the design of social insurance policies.

This paper studies how households respond to severe health shocks and insure against these shocks through spousal labor supply. In the empirical part of the paper, we provide new evidence on how individuals' labor supply responds to spousal health and mortality shocks. In the theoretical part of the paper, we develop a method for welfare analysis of social insurance that uses only spousal labor supply responses. We show that under plausible conditions the labor supply responses of spouses fully identify the welfare gains of insuring households against adverse shocks, and map our empirical findings on spousal labor supply responses to the welfare implications of providing more generous social insurance.

Spousal labor supply is a potential source of self-insurance when households experience sizable income shocks that are otherwise only partially insured. Therefore, in order to study spousal labor supply behavior as an insurance mechanism, our empirical analysis focuses on an extreme shock that leads to significant and permanent income losses – the death of a spouse. To recover the causal effect of this shock we offer a quasi-experimental design that constructs non-parametric counterfactuals to affected households by using households that experience the same shock a few years in the future, and combines event studies for these two experimental groups. The identification strategy we develop relies on the assumption that the **exact** timing of the shock is as good as random, and is therefore applicable to the analysis of a wide range of other common economic shocks.

Analyzing administrative data on health and labor market outcomes from the years 1980-2011, we study over 500,000 Danish households of married and cohabiting couples in which a spouse has died. We find a large and persistent increase in the surviving spouses' labor supply starting immediately after their spouses die, which amounts to an average increase of 7.6% in labor force participation and 6.8% in annual labor income by the fourth year after the shock.

These effects are driven by households that experience significant income shocks due to the loss of a spouse, and therefore have greater need for self-insurance through labor supply. In particular,

we show that the average increase in labor supply is entirely attributable to survivors whose deceased spouses had earned a large share of the household's income, who have less disposable income at the time of the shock, and who are less formally insured by government transfers. We also find that high-earning survivors, who experience smaller relative income losses and face better financial conditions, decrease their labor supply as their high income is no longer necessary to support two people. Notably, widowers – who tend to be financially stable when losing their wives – decrease their labor supply, while widows – who tend to experience considerably larger income losses when losing their husbands – significantly increase their labor supply. By the fourth year after their husbands die, widows increase their participation by 11.3%, which translates into a 10.1% increase in their annual earnings.

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

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

In the theoretical part of the paper, we map these estimates of spousal labor supply responses to predictions about the welfare gains from providing more generous social insurance. Using a model of efficient household behavior, we show that spousal labor supply responses can fully identify the

benefits of social insurance and develop a new method for welfare analysis that depends only on the spouse's labor supply behavior. This result relies on the observation that within each state of nature the spouse's labor force participation decision reveals the household's valuation of additional consumption (in the form of labor earnings). Hence, the sensitivity of spousal labor supply to shocks and economic incentives reveals the household's preference for consumption across different states of nature, which captures the benefits from insurance. We also consider the welfare implications of potential state dependence of the unaffected spouse's willingness to work.

Applying our welfare method to mortality shocks in our setting, we find substantial gains from benefit increases for elderly widows. Under a benchmark calibration of our model, an additional dollar to widows over 67 is equivalent to an additional \$1.55 to other elderly households, creating a net benefit of \$0.55 per \$1. However, for younger widows who are more attached to the labor force, we find very small gains from additional benefits through the social insurance system (with a net benefit of only \$0.04 per \$1), suggesting that for them the current level of transfers is near optimal. A key implication of our findings within our conceptual framework, driven in part by the differential attachment to the labor force over the life-cycle, is that age-dependence is a feature of the optimal social insurance policy for spousal mortality shocks.

This paper relates to several strands of the literature. First, numerous empirical studies have analyzed spousal labor supply and its responses to shocks in order to uncover the extent to which it is used as insurance. However, while spousal labor supply is commonly modeled as an important self-insurance mechanism against adverse shocks to the household (e.g., Ashenfelter 1980; Heckman and Macurdy 1980; Lundberg 1985), this prior empirical work has been unable to find evidence of significant increases in spousal labor supply in response to shocks (e.g., Heckman and Macurdy 1980, 1982; Lundberg 1985; Maloney 1987, 1991; Gruber and Cullen 1996; Spletzer 1997; Coile 2004; Meyer and Mok 2013). The leading explanation for this lack of evidence has been that within the context of temporary unemployment, on which the empirical literature has focused, income losses are small relative to the household's lifetime income and are already sufficiently insured through formal social insurance (Heckman and Macurdy 1980; Cullen and Gruber 2000). In order to uncover the self-insurance role of spousal labor supply within unemployment shocks, Cullen and Gruber (2000) study whether it is crowded out by unemployment insurance benefits and find a large crowd-out

effect; and, in ongoing work, Autor, Kostøl, and Mogstad (2015) find similar crowd-out effects in the context of disability insurance. We take an alternative empirical approach and directly study the effects of severe health shocks with different degrees of income loss – mortality shocks, which impose large and permanent income losses, and morbidity shocks, which are well-insured.

Second, prior work on estimating welfare gains from insurance has focused on studying its "consumption-smoothing" effects.[1] While it aims at directly identifying the benefits from insurance, this consumption-based method has two limitations. First, it is very sensitive to the value of risk aversion, for which the literature has a wide range of estimates (Chetty and Finkelstein 2013). Second, the choice of the studied consumption measure – most commonly food consumption – is usually driven by data availability rather than theoretical underpinnings. As emphasized by Aguiar and Hurst (2005), focusing on one aspect of expenditure can lead to very misleading conclusions about actual consumption in the presence of home production.[2]

The labor market approach to welfare analysis that we develop addresses these problems by relying solely on directly-observed changes in participation rates and labor supply elasticities and by utilizing labor market data that exactly match the theoretical behaviors of interest (participation rates and earned income). In addition, the wide availability of large-scale accurate data from the labor market renders our approach desirable for empirical applications.

Our method also relates to and builds on recent work on labor market methods for welfare analysis in the context of unemployment. Chetty (2008) recovers gains from social insurance using liquidity and substitution effects in the search effort of the unemployed, and Shimer and Werning (2007) use comparative statics of reservation wages with respect to government benefits.[3] In the

[1] This work includes reduced-form studies in the context of health shocks and the death of a spouse (e.g., Auerbach and Kotlikoff 1991; Stephens 2001; Bernheim et al. 2003; Meyer and Mok 2013; Ball and Low 2014) and studies that rely on structural economic modeling in the context of disability insurance and Social Security (e.g., İmrohoroğlu, İmrohoroğlu, and Joines 1995, 2003; Huang, İmrohoroğlu, and Sargent 1997; Kotlikoff, Smetters, and Walliser 1999; Bound et al. 2004; Benitez-Silva, Buchinsky, and Rust 2006; Nishiyama and Smetters 2007; Chandra and Samwick 2009; Bound, Stinebrickner, and Waidmann 2010; Low and Pistaferri 2012).

[2] Even comprehensive and accurate data on overall expenditure across health states, which are rarely available, would have to be accompanied by time-use data (on home production) and would require strong assumptions on their translation into individual consumption. Among other things, this procedure should take into account consumption flows of durable goods as well as economies of scale in the household's consumption technology (see, e.g., Browning, Chiappori, and Lewbel 2013).

[3] Following Chetty (2008), who uses variations in severance payments, other recent papers estimate the magnitude of the liquidity effects of social insurance programs – LaLumia (2013) uses variations in the timing of EITC refunds and Landais (forthcoming) uses kinks in the schedule of unemployment insurance benefits.

shocks that we consider, these methods cannot be applied because the directly affected individual may be unresponsive to economic incentives (or even deceased) and hence cannot reveal the household's preferences through labor market behavior. Exploiting the interplay between the labor supply decisions of household members, our method uses only the responses of the indirectly affected spouse. As such, our method offers a labor market approach that is also applicable to any economic shock in which the individual who is directly impacted may be unresponsive to economic incentives or at a corner solution.

The paper is organized as follows. Section II sets the conceptual framework for the empirical analysis and theoretically illustrates the self-insurance role of spousal labor supply using a model of household labor force participation. Prior to our empirical analysis, Section III describes the private and social institutional environment in Denmark and the data sources that we use to estimate individuals' labor supply responses to spousal health and mortality shocks. In Section IV we specify the empirical research design that we develop for recovering the causal effect of adverse shocks. Section V presents our estimates for spousal labor supply responses as a self-insurance mechanism. In Section VI we develop our method for welfare analysis of social insurance, and we study the welfare implications of our empirical findings in Section VII. Section VIII concludes.

II Conceptual Framework: A Model of Household Labor Force Participation

We begin with a baseline static model of extensive labor supply decisions.[4] The purpose of this section is to motivate our empirical analysis by formalizing how spousal labor supply can be used as insurance against income shocks to the household. Intuitively, when individuals experience severe health shocks that cause them to decrease their labor supply and earn less income, their spouses can compensate for this income loss by increasing their own labor supply. Moreover, the relative increase in spousal labor force participation in response to shocks increases with the income loss, which can reveal the extent to which the household needs to self-insure. In Section VI.C we discuss important extensions to the simple framework that we present here. Most importantly, we analyze

[4] We focus on the participation margin rather than work intensity since it turns out to be the operative margin of response in our main empirical analysis. We analyze an intensive-margin model in Online Appendix B.

a fully-dynamic life-cycle model that allows for endogenous savings (as well as private and informal insurance arrangements), which can easily incorporate a general class of arbitrary choice variables, such as time investment in home production.[5]

II.A Baseline Model

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

Individual Preferences. Let $U_i(c_i^s, l_i^s)$ represent i's utility as a function of consumption, c_i^s, and labor force participation, l_i^s, in state s (such that $l_i^s = 1$ if i works and $l_i^s = 0$ otherwise). We assume for now that $U_i(c_i^s, l_i^s) = u_i(c_i^s) - v_i \times l_i^s$, where the utility from consumption, $u_i(c_i^s)$, satisfies $u_i'(c_i^s) > 0$ and $u_i''(c_i^s) < 0$, and v_i is i's disutility from labor. The couple's disutilities from labor (v_w, v_h) are drawn from a continuous distribution defined over $[0, \infty) \times [0, \infty)$. We denote the marginal probability density function of v_w by $f(v_w)$ and its cumulative distribution function by $F(v_w)$.

Household Preferences. We follow the collective approach to household behavior (Chiappori 1988, 1992; Apps and Rees 1988) and assume that household decisions are Pareto efficient. Therefore, with equal Pareto weights for both spouses, household decisions can be characterized as solutions to the maximization of $U_w(c_w^s, l_w^s) + U_h(c_h^s, l_h^s)$.[6] It is important to emphasize here that the entire positive and normative analyses that follow do not rely on this particular modeling choice. **Any** model with efficient household behavior (an assumption that we discuss in Section VI.A), such as the widely used unitary model, would provide the same comparative statics that we explore in the empirical part of the paper and the theoretical welfare results that we provide thereafter.

[5] The simple model of this section is most closely related to the collective setting analyzed in Blundell et al. (2007), in which one spouse is on the participation margin while the other is on the intensive margin, as well as to Immervoll et al. (2011) who study optimal tax-and-transfer programs for couples with extensive-margin labor supply responses.

[6] More generally, household decisions can be characterized as solutions to the maximization of $\beta_w U_w(c_w^s, l_w^s) + \beta_h U_h(c_h^s, l_h^s)$, where β_w and β_h are the Pareto weights on w and h, respectively. However, setting $\beta_w = \beta_h = 1$ is without loss of generality as long as the spouses' relative bargaining power is stable across states of nature. Browning, Chiappori, and Weiss (2014) discuss the important distinction in the collective model between ex-post realizations of different states of nature, which do not affect the spouses' relative bargaining power under efficient risk sharing, and ex-ante distributions of income shocks, which may affect the Pareto weights. Similar to Chiappori (1992), baseline weights do not affect our welfare results.

Household's Problem. The household's choices reduce to the allocation of consumption to each spouse i in state s, c_i^s, as well as w's labor force participation in each state, l_w^s. Note that there are no savings decisions involved in the baseline static model (we introduce endogenous savings in the dynamic extension to the model). Each choice of w's employment determines the household's overall income in state s, $y^s(l_w^s)$, such that $y^s(l_w^s) = A + \bar{z}_h^s \times l_h^s + \bar{z}_w^s \times l_w^s + B^s(l_w^s)$, where A is the household's wealth and $\bar{z}_i^s$ is i's net-of-tax labor income in state s. $B^s(l_w^s)$ represents transfers from the government in state s, which we allow to depend on w's participation, so that transfers can be state-dependent as well as earnings-tested at the household level. More generally, the model allows for any type of state-contingent income and assets. These include life insurance and any other source of private insurance, employer-provided insurance, transfers from relatives, social insurance, medical expenses, etc.[7]

At each of w's potential employment statuses, consumption is efficiently allocated across spouses, such that the consumption bundles $c_w^s(l_w^s)$ and $c_h^s(l_w^s)$ are the solutions to

$$\begin{aligned} V(y^s(l_w^s)) \quad \equiv \quad & \max_{c_w^s, c_h^s} u_w(c_w^s) + u_h(c_h^s) \\ & \text{s.t. } c_w^s + c_h^s = y^s(l_w^s), \end{aligned} \tag{1}$$

where $V(y^s(l_w^s))$ is the household's "consumption utility" for any level of household income. We define y_{-w}^s as the household's resources excluding those directly attributed to w's labor supply decision – i.e., $y_{-w}^s \equiv A + \bar{z}_h^s \times l_h^s$.

The unaffected spouse, w, works in state s if and only if

$$v_w < \bar{v}_w^s \equiv V(y^s(1)) - V(y^s(0)). \tag{2}$$

That is, the unaffected spouse works if the household's valuation of the additional consumption of his or her labor income compensates for his or her utility loss from working.[8]

Spousal Labor Supply as Insurance. At this point it is easy to see the self-insurance role of spousal

[7] It is also straightforward to include economies of scale in the household's resource constraint (using a general transformation of income into individual consumption bundles as in Browning, Chiappori, and Lewbel 2013) as well as differential tax rules for joint filing.

[8] This simple decision rule reveals the household's preferences for additional consumption and allows us to map consumption utility to spousal labor force participation. It is the key source for identifying the gains from insurance based on the unaffected spouse's labor supply (as we show below in Section VI).

labor supply responses to shocks, which is our main outcome of interest. Denote w's probability of participation (or the participation rate of unaffected spouses in the population) in state s by $e_w^s \equiv F(\bar{v}_w^s)$, and the income loss from the shock by $d \equiv y_{-w}^g - y_{-w}^b$. In each state the unaffected spouse's probability of participation decreases in his or her unearned income:

$$\frac{\partial e_w^s}{\partial y_{-w}^s} = -f(\bar{v}_w^s)[u_w'(c_w^s(0)) - u_w'(c_w^s(1))] < 0. \tag{3}$$

This implies that we should expect a relative increase in spousal labor force participation, $\frac{e_w^b}{e_w^g} - 1 > 0$, whenever there is some income loss that is not fully insured, $d > 0$. That is, income shocks lead to self-insurance through the unaffected spouse's labor force participation. Furthermore, the unaffected spouse's labor supply response to the shock ($\frac{e_w^b}{e_w^g} - 1$) increases in the income loss d:

$$\frac{\partial \left(\frac{e_w^b}{e_w^g} - 1 \right)}{\partial d} = \frac{f(\bar{v}_w^b)}{F(\bar{v}_w^g)}[u_w'(c_w^b(0)) - u_w'(c_w^b(1))] > 0, \tag{4}$$

so that the extent of spousal self-insurance increases in the degree of income loss imposed by the shock. These comparative statics are no more than simple income effects at the household level and are a direct implication of the concavity of $u_i(c_i^s)$, which translates into the concavity of $V(y^s(l_w^s))$.[9]

II.B State-Dependent Preferences

Besides income losses, there are other important ways in which households can be directly affected by the shocks that we analyze. In particular, individual preferences can change in several dimensions, which can lead to spousal labor supply responses even in the presence of full insurance. In this section, we consider different potential types of such state dependence in preferences.

Let $U_i^s(c_i^s, l_i^s)$ represent i's utility in state s as a function of consumption, c_i^s, and labor force participation, l_i^s, in state s and assume that $U_i^s(c_i^s, l_i^s) = u_i^s(c_i^s, l_i^s) - v_i^s \times l_i^s$. This formulation generalizes

[9] The complete formal description of the household's problem in each state is

$$\max_{l_w^s \in \{0,1\}, c_w^s(l_w^s), c_h^s(l_w^s)} l_w^s(U_w(c_w^s(1), 1) + U_h(c_h^s(1), l_h^s)) + (1 - l_w^s)(U_w(c_w^s(0), 0) + U_h(c_h^s(0), l_h^s))$$
$$\text{s.t. } c_w^s(l_w^s) + c_h^s(l_w^s) = y^s(l_w^s)$$
$$y^s(l_w^s) \equiv A + \bar{z}_h^s \times l_h^s + \bar{z}_w^s \times l_w^s + B(l_w^s).$$

There is another natural approach to modeling the household's decision-making process. One can assert that each individual works if his or her *own* utility from working is higher than his or her *own* utility from not working, and then – conditional on the participation decisions – the couple engages in efficient bargaining that allocates resources according to their respective bargaining power (which in our case implies maximizing $u_w(c_w^s) + u_h(c_h^s)$). The qualitative theoretical results of our analysis (both positive and normative) remain unchanged in this alternative model.

preferences as follows. First, it allows for a completely flexible dependence of consumption utility on the state of nature. Note in particular that this allows us to study the death of h within our framework by setting $u_h^b(c_h^b, l_h^b) = 0$. That is, the bad state can capture either the state of nature in which h is sick or the state in which h is deceased. Second, it allows for flexible consumption-leisure complementarities by allowing the consumption utility to depend freely on participation.[10]

Third, we allow labor disutility, v_i^s, to change across states of nature. For the affected spouse h, this captures the direct effect of health on the ability to work when state b is h's sickness. For the unaffected spouse w, this generalization captures the potential state dependence of labor disutility. For example, when the bad state is h's sickness, v_w^b might be greater than the baseline labor disutility v_w^g if w places greater value on time spent at home – e.g., to take care of his or her sick spouse. When the bad state is h's death, working may become less desirable if the surviving spouse experiences depression and has difficulties working, or conversely, working may become more desirable if the surviving spouse feels lonely and wishes to seek social integration. For simplicity, we model this type of state dependence as $v_w^g = v_w$ and $v_w^b = \theta^b \times v_w^g$, such that θ^b captures the mean percent change in the utility cost of labor compared to the baseline state g.[11]

With these generalized preferences, the comparative statics in equations (3) and (4) still hold. However, potential changes in the unaffected spouse's labor disutility (or willingness to work) can directly lead to spousal labor supply responses (and affect our normative results). Even with complete insurance ($d = 0$), a decrease in spouse w's labor disutility in the transition from state g to state b (i.e., $\theta^b < 1$) will cause an increase in spousal labor force participation (such that $\frac{e_w^b}{e_w^g} - 1 > 0$).[12]

The remainder of the paper proceeds with the empirical analysis of the impact of health and mortality shocks. Our main outcome of interest is spousal labor supply responses to these shocks

[10] One may also want to allow individual i's consumption utility, $u_i^s(c_i^s, l_i^s)$, to depend on the spouse's consumption and participation, which will not change our welfare results. We abstract from these potential dependencies for keeping notation simpler.

[11] In Online Appendix A we show that this is a simplification and that it is not necessary to define such a global parameter for our theoretical results. We illustrate how it can be locally and non-parametrically defined in the more general dynamic search model. In addition, in Online Appendix E we offer an example for allowing heterogeneity in θ^b.

[12] Since this type of response would be driven by w's relative preference for work, we show in the normative analysis that it would not translate into welfare gains from more generous social insurance, in contrast to labor supply responses that are driven by income losses and self-insurance.

$(\frac{c_H^b}{c_H^g} - 1)$, which we show to have important normative implications for the design of social insurance within our conceptual framework. In order to provide empirical support for the insurance role of spousal labor supply we analyze the heterogeneity of these responses by the degree of income loss imposed by the shocks that we study, as emerged from the comparative statics of our model in equation (4). To analyze other potential mechanisms that may underlie the average responses we also develop and empirically apply tests to assess the extent to which spousal labor disutility changes in response to shocks (θ^b). We then return to the model of household labor supply to develop our method of welfare analysis that relies only on spousal responses, and study the welfare implications of our empirical results. We do so by showing that the relative difference in marginal utilities of consumption across health states, which captures the benefits from social insurance, can be fully recovered by the labor supply responses of spouses.

III Data and Institutional Background

To study labor supply responses to severe spousal health shocks we turn to the Danish institutional setting and its rich administrative data on health and labor market outcomes. In this section, we describe the Danish insurance environment, both social and private, as it relates to sick individuals and surviving spouses, and list our data sources. It is useful to distinguish between two types of insurance: health insurance (coverage of medical care) and income insurance (insurance against income losses in different health states). Health insurance in Denmark is a universal scheme in which almost all costs are covered by the government, with a few exceptions such as dental care, chiropractic treatments, and prescription drugs that entail a limited degree of out-of-pocket expenses. Therefore, the Danish setting allows us to concentrate on (social and private) income insurance for losses that go beyond immediate medical expenses, as we describe below. Note, however, that the theory allows for medical expenses (and any other state-contingent expenses) and that our welfare analysis method is robust to any degree of medical coverage.

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

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

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

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

While Social DI and its surviving benefits component are state-wide schemes, they are locally administered. Regional councils (in a total of 15 regions) decide whether to approve or reject an individual's application, and municipal caseworkers (in a total of 270 municipalities) administer the application and handle all aspects of each case – including any contact with the applicant, preparation of the application, and collection of financial and health status records. The local administration of the program has led to differential application behavior across municipalities, which has resulted in substantial variation in rejection rates – ranging from 7% to 30% – and thus

in the mean receipts of the program's benefits across the different municipalities (Bengtsson 2002). We exploit this cross-municipality variation over time in the awarding of the survivors benefits component of the program later in the paper.

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

It is important to emphasize that while the private market for life insurance is large, there is a potentially important role for government interventions as we study in the welfare analysis section of the paper. First, since purchasing life-insurance products in Denmark requires answering a health status and behaviors questionnaire (and even undergoing medical exams) applications by older and unhealthy households are likely to be rejected.[13] Second, it is common that even when the life-insurance product is purchased by younger and healthy households (both in group and nongroup markets) the coverage sharply declines with age.[14] This leaves older and unhealthy Danish households with poorer coverage through the private market.

Lastly, there are old-age social insurance programs that can indirectly protect eligible survivors or households that experience other shocks, who can decide to take them up at different ages according to their financial needs. When crossing into their 60s and until they reach their old-age pension retirement age, individuals who have (voluntarily) been members of an unemployment fund for a sufficiently long period (10 years before 1992 and gradually increasing to 20 years thereafter) are eligible for the Voluntary Early Retirement Pension (VERP). Approximately 80% of the population is eligible for VERP, which provides a flat-rate annual income of roughly DKK 130,000 ($16,250).

[13] These rejections by the insurance companies can be explained by private information that is held by these rejected households (Hendren 2013).

[14] For example, some large white-collar group-market policies guarantee DKK 1,076,000 ($162,050) if the insured 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.

At the "full-retirement" age of 67 (or 65 for those born after July 1, 1939) all residents become eligible for the Old-Age Pension (OAP), which provides income-tested annuities of up to DKK 99,000 ($12,375) per year for singles and DKK 75,000 ($9,375) for coupled individuals (at 2000 rates). Note that the benefits to single survivors who qualify for Social DI sharply reduce at age 67 (from $18,000 to $12,375) when the program transitions into the Old-Age Pension, adding to the financial vulnerability of older households.[15]

Data Sources. We have merged data from several administrative registers to obtain annual information on Danish households of married and cohabiting couples from 1980 to 2011. We use the following registers: (1) the national patient register, which covers all hospitalization records (from both private and public hospitals), and from which we extract information on all the individuals that experienced a heart attack or a stroke; (2) the cause of death register, from which we identify death dates; (3) income registers, which include all sources of household income – e.g., labor income, capital income, annuity payouts, and government benefits from any program – as well as annual measures of gross wealth and liabilities;[16] and (4) the Integrated Database for Labor Market Research, which includes measures from which we construct full-time and part-time labor supply variables and extract demographic variables. All monetary values are reported in nominal Danish Kroner (DKK) deflated to 2000 prices using the consumer price index. In that year the exchange rate was approximately DKK 8 per US $1. We postpone describing the summary statistics of the analysis sample to the next section since they directly relate to the discussion on the advantages of our research design.

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

[16] In our main analysis sample of spousal mortality shocks, the net assets of the median household amount to only DKK 13,236 ($1,655) while the median annual household-level income is DKK 239,922 ($29,990). Therefore, our analysis of labor supply responses focuses on income losses, and we use the wealth data for robustness checks.

IV Research Design

In this section we describe our empirical strategy for identifying the causal effects of spousal health and mortality shocks on individuals' labor supply, $\frac{e_w^b}{e_w^g} - 1$.

The ideal experiment would randomly assign these shocks to households and track labor supply responses over time. Both our baseline theoretical model of Section II and the dynamic life-cycle model (that we discuss in Section VI.C and develop in Online Appendix A) call for comparing the responses to shocks of affected households to the counterfactual behavior of ex-ante similar unaffected households. This requires comparing households with same expectations over the distribution of future paths, but with different realizations. The access to over three decades of administrative panel data on the universe of Danish households allows us to develop a quasi-experimental research design that mimics this ideal experiment, by exploiting the potential randomness of the **exact** timing of a severe health shock or death within a short period of time.

To do so, we look only at households that have experienced the shocks that we consider at some point in our sample period, and identify the treatment effect from the timing at which the shock was realized. We construct non-parametric counterfactuals to affected households using households that experience the same shock a few years in the future, and recover the treatment effect by performing event studies for these two experimental groups. Note that a simple event study, which analyzes the evolution of outcomes of a treated group around the time of a shock, is not appropriate for our application. Pure event studies identify short-run responses, while we are interested in identifying longer-run effects because of potential delays in adjustment (due to, e.g., labor market frictions). This requires a control group, as we construct in our design, that can account for complex life-cycle trends in the counterfactual behavior in the absence of a shock (in, e.g., spousal labor force participation as depicted in Online Appendix Figure I).

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

Illustrative Example. Let us focus on a treatment group of individuals born between 1930 and 1950 who experienced a severe health shock, in particular, a heart attack or a stroke, in 1995. Consider studying the effect of the shock on some economic outcome of these individuals, e.g., their

labor force participation. Panel A of Figure I plots the outcome for these households against the outcomes for households that have not experienced this shock in our sample period, and reveals very different life-cycle patterns across the two groups prior to 1995. This suggests that traditional matching estimators, which use these unaffected households as a control group, are inappropriate for our application as their validity will rely heavily on the set of available controls and on the unconfoundedness assumption.[17] Panel B of Figure I plots the outcome for the treatment group of households that experienced a shock in 1995 as well as for households that experienced the same shock in 2010 (15 years later), in 2005 (10 years later), in 2000 (5 years later), and in 1996 (1 year later). Studying the behavior of households that experienced the shock in different years reveals increasingly comparable patterns to those of the treatment group's behavior – in terms of trends before 1995 – the closer the year in which the individual experienced the shock was to 1995. These patterns confirm our intuition and suggest using households that experienced a shock in $1995+\triangle$ as a control group for households that experienced a shock in 1995. Panel D of Figure I displays a potential control group when we choose $\triangle = 5$.

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

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

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

can identify effects up to four years after the shock. We assessed the robustness of our analysis to this choice and found that local perturbations to $\triangle$ provide very similar results.

Formal Description of the Design and Estimator. Fix a group of cohorts, denoted by Ω, and consider estimating the treatment effect of a shock experienced at some point in the time interval $[\tau_1, \tau_2]$ by individuals who belong to group Ω. We refer to these households as the treatment group and divide them into sub-groups indexed by the year in which they experienced the shock, $\tau \in [\tau_1, \tau_2]$. We normalize the time of observation such that the time period, t, is measured with respect to the year of the shock – that is, $t = year - \tau$, where $year$ is the calendar year of the observation. As a control group, we match to each treated group τ the households among cohorts Ω that experienced the same shock but at $\tau + \triangle$ for a given choice of $\triangle$. For these households we assign a "placebo" shock at $t = 0$ by normalizing time in the same way as we do for the treatment group ($t = year - \tau$).[18]

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

$$\beta_n \equiv \left(y_n^T - y_n^C\right) - \left(y_p^T - y_p^C\right). \tag{5}$$

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

The identifying assumption is that, absent the shock, the outcomes of the treatment and control groups would run parallel. In particular, in accordance with the differences-in-differences research design, there is no requirement regarding the levels of outcomes. The plausibility of this assumption relies on the notion that within the short window of time of length $\triangle$ the **exact** time at which the shock occurs is as good as random. To test the validity of our assumption, we accompany our empirical analysis with the treatment and control groups' behavior in the five years prior to the shock year 0 in order to assess their co-movement in the pre-shock period. By showing that there

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

are virtually no differential changes in the trends of the treatment and control groups before period 0, we alleviate concerns that the two groups may still differ by, e.g., their expectations over the timing of the shock.[19]

Other papers that use similar identifying assumptions include earlier studies in the context of the long-run effects of job displacement (Ruhm 1991) and the effect of arrests on employment and earnings (Grogger 1995), as well as more recent studies such as that by Hilger (2014), who exploits variation in the timing of fathers' layoffs in order to study the effect of parental income on college outcomes. Our quasi-experimental design can be applied to these shocks and any other shock of which the exact timing is random, which can be easily validated in any particular setting by studying the pre-trends of the experimental groups.

IV.A Analysis Sample and Summary Statistics

Table I displays key summary statistics for the analysis sample. The sample of our main analysis includes households in which one spouse died between ages 45 and 80 and is comprised of 310,720 households in the treatment group and 409,190 households in the control group.[20]

The table reveals the advantage of our research design – the comparability of the year of observation and the age of unaffected spouses across experimental groups. The average survivor in the treatment group loses his or her spouse in 1993 at age 62.86 and the average unaffected spouse in the control group experiences the placebo shock in year 1993 at age 62.27. The sub-sample of survivors under age 60, the age at which there is a large drop in labor force participation (due to eligibility for early retirement benefits as shown in Online Appendix Figure I), displays even closer similarities. By construction, the research design nets out calendar year effects non-parametrically. However, due to the randomness of the exact timing of the shock, it also nets out life-cycle effects by comparing groups of very similar ages, so that we effectively compare spouses who experience a

[19] Conceptually, as long as there is no perfect foresight we can use our strategy with the appropriate choice of $\triangle$. This choice is context dependent and requires empirical investigation (where any potential difference across the experimental groups would be included in the bias consideration in the choice of $\triangle$). Comparability is then an empirical question that can be investigated in several ways, such as analyzing sub-samples of shocks that are more likely to come as a surprise and studying the robustness of the results to a rich set of controls, along with testing for parallel trends in the pre-period and investigating the sensitivity of the results to the chosen control group by changing $\triangle$ as we mentioned above. Conducting this set of tests verifies the robustness of our results, supporting our underlying identifying assumption.

[20] Importantly, we conducted additional analysis that constrained the sample to households in which a spouse experienced a heart attack or a stroke for the first time and died within the same year in order to focus on deaths that are more likely to come as a surprise. The qualitative results are similar to those presented here and are available from the authors on request.

shock at age a to same-age spouses who experience a shock at age $a + \triangle$.

The sample for our secondary analysis of severe health shocks includes households in which one spouse experienced a heart attack or a stroke (for the first time) and survived for at least three years. These shocks are among the leading causes of death in the developed world and their exact timing within a short period of time is likely unpredictable. Since the average age of spouses precisely at the time of these health shocks is just over 60 (60.67), we focus on households with both spouses under 60 to ensure that the results we document are driven only by the health shocks and not by eligibility for early retirement benefits.[21] The sample consists of 37,432 households in the treatment group and 54,926 households in the control group. The unaffected spouse is on average 45.7 years old in the treatment group at the time of the shock and 45.3 years old in the control group, where the mean calendar year of the shock is around 1992 for both groups.[22]

V Spousal Labor Supply Responses

V.A Labor Supply Responses to the Death of a Spouse

In this section, we present our main empirical analysis and study survivors' labor supply responses to the death of their spouse. We begin by estimating average labor supply responses. Then, we analyze the heterogeneity of these responses by the degree of income loss imposed by the death of a spouse in order to provide support for the self-insurance role of spousal labor supply (in accordance with the comparative statics of equation (4)).

Mean Responses. Figure II plots the average labor supply response of individuals whose spouse died between ages 45 and 80. Panel A reveals an immediate increase in labor force participation (defined as having any positive level of annual earnings) following the death of a spouse. By the fourth year after the shock, the surviving spouses' participation increases by 7.6% – an increase of 1.6 percentage points (pp) on a base of 20.6 pp. Panel B of Figure II shows that this response translates into a 6.8% increase in annual earnings (including zeros for those who do not work), which represents an annual increase of DKK 2,572 ($322) from a low base of DKK 37,952 ($4,744).

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

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

With significant disparities in baseline participation rates and labor income, men and women may face substantially different financial distress when they lose their spouse and, therefore, may respond differently to the death of their spouse. Indeed, Figure III reveals stark differences in the responses of widowers and widows. While on average widowers do not change their labor force participation when their wife dies, widows immediately and significantly increase their labor force participation when they lose their husband. Four years after the shock, widows' labor force participation increases by 2.2 pp from a baseline participation rate of 19.5 pp, which amounts to a large increase of 11.3% in their labor force participation.

This differential response suggests that female survivors have greater need to self-insure through labor supply and that they experience greater income losses when they lose their spouse as compared to their male counterparts. To test this conjecture, we plot the evolution of overall household income (from any source) around the death of a spouse, including earnings, capital income, annuity payouts, and benefits from social programs. We begin by plotting the household's income in the absence of behavioral responses from the unaffected spouse in order to capture the income loss directly attributable to the loss of an earning spouse. To do so, we plot in Panel A of Figure IV the household's overall income, holding the unaffected spouse's earnings and social benefits at their pre-shock level.[23] The graph shows that widowers experience a 32% loss in household income, while widows lose an additional 8% and experience a significantly larger loss of 40%. Panel B of Figure IV studies the actual change in household income, taking into account the surviving spouses' labor supply responses and any change in the benefits they may receive from social or private insurance. The figure shows that widowers experience an actual loss of 31% and that widows manage to decrease their potential loss (through the increase in labor supply and transfers from private and social insurance) to incur an actual lower loss of 35%.[24]

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

[24] Widows' labor supply responses account for 22% of their (5 pp) decrease in potential income loss (from 40% to 35%). Note that surviving spouses do not fully compensate for a loss in household income (to 100%) since as singles they do not need the full pre-shock level of income. However, potential economies of scale in the household's consumption technology may make half of the pre-shock level of household income insufficient for maintaining the pre-shock level of utility (see, e.g., Nelson 1988; Browning, Chiappori, and Lewbel 2013). The share of household income that keeps consumption utility at its pre-shock level is usually assumed to lie between 0.5 and 1 and is commonly referred to as the adult "equivalence scale". We return to this issue in Section V.A.1 below.

Younger Households. Surviving spouses under 60 have a stronger attachment to the labor force and higher labor earnings and are, therefore, more financially resilient after the loss of an earning spouse.[25] Consistent with the view that their higher participation rates and annual earnings effectively insure them better against losing an earning spouse, Panel A of Figure V reveals that survivors under 60 exhibit a smaller relative increase in labor force participation compared to the universe of survivors – only 2.1% (1.4 pp on a base of 67.2). Similar to the overall treatment effect, this increase is entirely driven by women. As seen in Panel B of Figure V, widows increase their labor force participation by 3.3%, while widowers – who have a higher baseline participation rate (0.78) as compared to widows (0.715) – respond with a small (but statistically significant) decrease of 1.1% in their participation. These responses translate to a 3.2% increase in annual earnings for the lower-earning widows and, interestingly, to a decrease of 4.1% in annual earnings for the higher-earning widowers, who as singles may not need their entire (much) higher pre-shock levels of income (see Panel C of Figure V). As before – and as displayed in Panel C of Figure IV – these differential responses are consistent with the differential financial shock that they experience, with men experiencing a decline of 31% in household income and women experiencing a strikingly larger loss of 44%.

We report estimates for the regression counterparts of these figures in Table II, which replicates our results so far. As we alluded to in Section IV, the treatment effect can be recovered by the simple differences-in-differences estimator of equation (5). The regression specification for this estimator, averaged over the years after the shock, is of the form:

$$l_{w,i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_3 treat_i \times post_{i,t} + \beta_4 X_{i,t} + \alpha_i + \varepsilon_{i,t}. \quad (6)$$

In this regression $l_{w,i,t}$ denotes an indicator for the labor force participation or annual earnings of the unaffected spouse w in household i at time t; $treat_i$ denotes an indicator for whether a household belongs to the treatment group; $post_{i,t}$ denotes an indicator for whether the observation belongs to post-shock periods; $X_{i,t}$ denotes a vector of potential controls; and α_i is a household fixed effect. The parameter β_3 represents the causal effect of the death of a spouse on the labor supply of the unaffected spouse. As we show in Online Appendix Figure II, in periods 0 and 1 there are temporary

[25] Recall that at 60 there is a sharp drop in participation when most of the labor force becomes eligible for early retirement benefits.

transitions to part-time work, consistent with spending time with the dying spouse and mourning his or her loss. These transitions stabilize thereafter such that the active decision margin becomes full-time work vs. non-participation. Throughout the analysis, $post_{i,t}$ therefore assumes the value 1 for periods 2 to 4.

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

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

First, similarly to Panels A and C of Figure IV, we calculate for each household the overall income (from any source) holding the unaffected spouse's earnings and social benefits at their pre-shock level (in $t = -1$). Second, we calculate the ratio of this "potential income" measure in $t = 1$ to the household's income in $t = -1$. Third, we normalize this ratio for the treated households by the mean ratio of the control households in order to purge life-cycle and time effects. This leaves us with a measure of the potential income replacement rate for each treated household, which we denote by rr_i, that captures the change in household income directly attributed to (and only to) the loss of a spouse.

To study the heterogeneity in labor supply responses by the income replacement rate (rr_i) we estimate the following augmented differences-in-differences model

$$l_{w,i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_{3i} treat_i \times post_{i,t} + \beta_4 X_{i,t} + \alpha_i + \varepsilon_{i,t}, \tag{7}$$

where

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

In this regression $l_{w,i,t}$ denotes an indicator for the labor force participation of the unaffected spouse w in household i at time t. We augment the basic differences-in-differences design by allowing the treatment effect, β_{3i}, to vary across households and model it as a function of the household's

potential replacement rate rr_i. Our parameter of interest is β_{31}, which captures the extent to which the surviving spouse's labor supply response correlates with the income loss he or she experiences. Since β_{31} can capture other dimensions of heterogeneity beyond the income replacement rate, we let the treatment effect vary with additional household-level characteristics, $Z_{i,t}$, such that β_{31} further isolates the treatment effect's partial correlation with the loss of household income.[26]

Table III reports the results of estimating (7) separately for each gender, with and without $Z_{i,t}$, for the entire sample of surviving spouses and for only the sub-sample of survivors under age 60. The results consistently show throughout the specifications the strong correlation between labor supply responses and income losses; survivors in households with lower potential income replacement rates (lower rr_i), who experience larger income losses, are much more likely to increase their labor force participation in response to the shock. Since controlling for the additional interactions with $Z_{i,t}$ does not change the results much, the evidence suggests that the heterogeneous responses are indeed driven by differential income replacement rates. In addition, the estimation results reveal quite similar sensitivity to income losses across genders. This verifies that gender differences in preferences do not drive the differential average labor supply responses across female and male survivors (but rather their divergent income losses).

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

We constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than that of their experimental-group-specific 20th percentile. Then, for each household we calculate the pre-shock labor income share

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

of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force. Specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households in which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household.[27]

We divide the remaining sample into five equal-sized groups according to their pre-shock level of earnings, and plot in Panel A of Figure VI the average labor income response (as well as its 95-percent confidence interval) against the pre-shock mean earnings for each group. The figure reveals a strong gradient of labor supply responses with respect to the survivors' own level of earnings when the shock occurs. Survivors at the bottom of the income distribution increase their annual earnings by 7.79% in order to meet their consumption needs, while those at the top decrease their earnings by 2.93% as their high income is no longer necessary to support two people and since they may find lower levels of income sufficient for their consumption needs as singles.

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

[27] These restrictions also imply that the results below are mainly driven by intensive-margin responses.

Spatial Variation in Social Insurance over Time. Lastly, we take advantage of spatial variation in the administration of social survivors benefits to study survivors' labor supply responses by the generosity of social insurance. This allows us to test the self-insurance hypothesis of spousal labor supply by analyzing whether better social insurance crowds out labor supply increases in response to shocks. We find that the increase in survivors' participation due to the shock declines in the formal insurance they receive from the government, which provides further support to the self-insurance hypothesis.

For this analysis, we constrain the sample to survivors under 67 (the age at which the program transitions into the Old-Age Pension) and to the period prior to 1994 due to a data break in the reporting method of survivors benefits received through Social DI. In addition, since for this sample the increase in the take-up of the program following the shock is attributable to females, we focus the analysis on widows. Inclusion of widowers does not change the qualitative results.

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

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

$$l_{w,i,t} = \beta_0 + \beta_1 treat_i + \beta_2 post_{i,t} + \beta_{3i} treat_i \times post_{i,t} + \beta_4 X_{i,t} + \varepsilon_{i,t}, \tag{8}$$

where

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

In this regression, $l_{w,i,t}$ denotes the participation of individual w of household i at time t, and $X_{i,t}$ includes municipality m's unemployment rate and average earnings (and their interaction with

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

The two-stage least squares results are presented in Table V. The estimate for our parameter of interest, $\beta_{31} = \frac{\partial(e_w^b - e_w^g)}{\partial b^b}$, is -.0057. With an average of DKK 23,262 ($2,908) in actual survivors benefits receipts by widows in the analysis sample (including zeros for those not on the program) and a participation rate of 0.5054, this estimate translates to a participation elasticity with respect to social benefits of $\varepsilon(e_w^b, b^b) = -0.26$ for widows under 67. This implies that social insurance crowds out labor supply responses as a self-insurance mechanism against loss of income following the death of a spouse.

In summary, the results reveal a clear pattern: there are significant increases in labor supply in response to losing a spouse, which are entirely driven by households that experience large income losses. The results provide strong evidence of the self-insurance role of spousal labor supply in the extreme case of the death of a spouse, which translates into large and permanent income losses for most households. Within our conceptual framework, the mean effects suggest large welfare benefits from additional transfers to widows due to incomplete insurance of spousal mortality, as we analyze in Section VII.

V.A.1 Assessing Labor Disutility State Dependence for Survivors

In this section, we briefly discuss two strategies to assess the extent to which survivors' labor disutility (or willingness to work) changes in response to the death of their spouse.[28] The evidence

[28] One empirical motivation to account for this sort of state dependence is the striking change in the surviving spouse's health-care utilization following the loss of a spouse. Online Appendix Figure III shows that the overall expenditure on primary medical care (Panel A) as well as the prescription rate for antidepressants (Panel B) exhibit sharp increases in the year of bereavement. While part of these phenomena may be purely driven by changes in take-up of medical care and supply-side responses rather than in actual changes in health, they call for an empirical investigation of labor disutility state dependence.

is inconsistent with the conjecture that the mean increase in the surviving spouses' labor supply is driven by lower cost of labor following the shock ($\theta^b < 1$), e.g., due to loneliness and the desirability of social integration. The analysis consistently supports the view that this increase is driven by self-insurance and large income losses.[29]

The first strategy for studying labor disutility state dependence provides a simple test for the hypothesis that loneliness and seeking social integration – after losing a partner with whom the surviving spouse spent leisure time – may drive the surviving spouses' labor supply responses. Consider widows, for whom we find an increase in participation in response to spousal death, who did not work before the shock in a model where time is divided between labor and leisure. Widows in households in which the deceased spouse did not work before his death experience smaller income losses (taking into account the deceased's income from any source including government transfers), but consumed more joint leisure and hence may be more likely to experience loneliness. In contrast, widows in households in which the deceased spouse worked before his death consumed less joint leisure, but experience larger income losses. The social integration (or "loneliness") hypothesis is consistent with spouses in the first group of households increasing their labor supply more than spouses in the latter group do, while the self-insurance hypothesis is consistent with the opposite pattern. Online Appendix Table I verifies the differential level of income losses across the two groups and shows that the spousal labor supply increase in households in which the deceased worked is much larger (by 4.61 pp), with a negligible effect in households in which the deceased did not work (0.78 pp). Moreover, among households in which the deceased did not work and received low levels of income, there was no increase in the widows' labor supply.

The second strategy, which we develop and prove in detail in Online Appendices C and D, provides a formal method to study general forms of labor disutility state dependence. This method studies whether after the shock survivors fully adjust their income and consumption levels through labor supply responses so that they can achieve their pre-shock level of consumption utility. Intuitively, if survivors **under** adjust their consumption levels after the shock, it implies that supplying labor has become **more** costly. If survivors **over** adjust their income loss through labor supply

[29] The strategies that we develop and apply here also test the hypothesis that work time of survivors may have become less costly since they no longer need to care for the deceased spouses, who had been potentially ill before their death.

responses, it implies that supplying labor has become **less** costly.

As a benchmark for full adjustment we use the adult "equivalence scale", which quantifies the share of the household's overall pre-shock income that survivors need **as singles** in order to achieve the same level of (individual) consumption utility that they enjoyed before the shock (see, e.g., Blundell and Lewbel 1991).[30] Using different common estimates for the adult equivalence scale, we find that the surviving spouses' labor supply responses are consistent with under- (and sometimes close-to-full) adjustment of their consumption following the shock.[31] This implies that survivors' labor disutility is likely to have gone up after the death of their spouse. Notably, the calibrations are in contradiction to the hypothesis that the increases in survivors' labor supply are driven by lower cost of labor.

Overall, our analysis of spouses' potential labor disutility state dependence in the context of death events supports the self-insurance hypothesis, and suggests that the average increase in labor supply is attributable to self-insurance of large income losses rather than to decreases in survivors' labor disutility.

V.B Labor Supply Responses to Spousal Health Shocks

In this section we briefly study individuals' labor supply responses to severe spousal health shocks. The purpose of studying this additional shock is to provide further evidence for the self-insurance hypothesis of spousal labor supply. Recall that our analysis sample for this shock consists of households in which a spouse experienced a heart attack or a stroke (for the first time) and survived for at least four years (until $t = 3$), and in which both spouses were under age 60.

Panel A.1 of Figure VII shows that within three years of the shock, the affected spouses' participation sharply falls, which translates into a large loss of annual earnings as shown in Panel A.2.

[30] We focus on the adult equivalence scale since we study older households. 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.

[31] Specifically, taking into account their labor supply responses, survivors' post-shock income is on average 0.665 of their counterfactual overall household income. Some commonly used equivalence scales (such as the modified OECD equivalence scale of 0.67 and the square-root scale of 0.71) imply close-to-full adjustment. Note that the implicit equivalence scale in the Danish Social DI is approximately 0.65 and is 0.66 in the Old-Age Pension (see Section III). Other model-based estimates for adult equivalence scales, such as those of Browning, Chiappori, and Lewbel (2013), who offer separate scales for men and women, imply under-adjustment. Their estimates for households with equal sharing of income among the two spouses are 0.80 for males and 0.72 for females, while our results (from Panel B of Figure IV) imply that widowers' income is 0.69 and widows' income is 0.65 of their corresponding counterfactual overall household income. See Online Appendix C for more details.

Table VI quantifies these effects by estimating a differences-in-differences regression, in which we allow for differential treatment effects in the "short run" (periods 1 and 2) and the "medium run" (period 3), to account for the gradual responses documented in Panel A of Figure VII.[32] Columns 2 and 4 of Table VI reveal that by the third year after the shock the labor force participation rate of the sick spouses drops by 12 pp – about 17% – and that annual earnings drop by DKK 36,015 ($4,500) – a significant drop of 19%.

However, while there is a significant drop in the sick spouses' earnings, Columns 5 and 6 of Table VI show that the actual loss of income that these households experience is much smaller and amounts to only 3.3% of overall household income. That is, taking into account the entire household income, including any transfers from social or private sources (particularly Disability Insurance), reveals that these shocks are very well-insured in our Danish setting. Consistent with this lack of an income drop, there are no economically significant labor supply responses among unaffected spouses (as shown in Panel B of Figure VII and Columns 7 to 10 of Table VI) as there is no significant need to self-insure.[33]

Note that the rich data-set and our research design allow for a precise estimation of these economically insignificant spousal responses to shocks.[34] In particular, our results imply a small but positive degree of complementarity in spouses' labor supply in response to health shocks, with an estimate of 0.065 for the unaffected spouse's earnings elasticity with respect to the affected spouse's earnings. Since the household's income is **not perfectly** insured, this response implies – in the context of our theoretical framework – health-state dependence of the household's utility. Intuitively, the fact that given a small loss of income due to the shock the unaffected spouses' decrease in labor supply involves an additional (very small) loss (through their lower earnings) is consistent with two main state dependence channels. First, it is consistent with households in the

[32] We estimate the following specification

$$y_{i,t} = \beta_0 + \beta_1 treat_i + \beta_{2a} post^a_{i,t} + \beta_{3a} treat_i \times post^a_{i,t} + \beta_{2b} post^b_{i,t} + \beta_{3b} treat_i \times post^b_{i,t} + \alpha_i + \varepsilon_{i,t}, \quad (9)$$

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

[33] As in spousal mortality shocks, we find a strong correlation between spousal labor supply responses and income losses in the context of health shocks. The analysis is available from the authors on request.

[34] This may explain the survey-based noisy estimates of Coile (2004) and Meyer and Mock (2013), who study responses to health shocks in the US. Note that Meyer and Mock (2013) similarly find that the typical disabled individual in the US loses about 21% in earnings but only 6.75% in post-transfer household income by the fourth year after the shock.

bad state valuing income less than do households in the good state – i.e., a consumption utility state dependence. Second, it is consistent with an increase in the unaffected spouses' utility loss from time spent away from home either because they would like to take care of their sick spouse or due to preferences for joint leisure – i.e., a labor disutility state dependence. With no additional assumptions, we can only reach conclusions about the ratio of these two types of potential state dependence. See Online Appendix F for a formal analysis.

VI Welfare Analysis: Theory

We now return to our conceptual framework and show that spousal labor supply responses can be sufficient for welfare analysis of social insurance. By doing so, we map how individuals respond to spousal shocks into predictions about the welfare implications of providing more generous social insurance. We proceed with analyzing the optimal design of social insurance in our baseline model and then discuss important extensions and generalizations.

VI.A Optimal Social Insurance

Policy Tools. The planner observes the state of nature as well as the employment status of each spouse. Since some spouses work and earn more than others do, the optimal policy is dependent on whether the spouse is employed. We denote the tax on spouse i's labor income in state g by T_i^g and the benefits given to non-working spouses in state g by b^g. In state b, households in which the unaffected spouse, w, works receive transfers of the amount B^b and households in which w does not work receive benefits of the amount b^b. This tax-and-benefit structure allows for the analysis of flexible policy designs and mimics features of existing social insurance programs in most developed countries (such as income-testing that characterizes the Supplemental Security Income program within the Old-Age, Survivors and Disability Insurance in the US and the Social Disability Insurance in Denmark). The exact way in which we model transfers simplifies the analysis and is not necessary for our results. Any system that conditions transfers on the state of nature and employment, as well as on age in the dynamic model, can be analyzed within our framework. We denote taxes by $T \equiv (T_w^g, T_h^g)$ and benefits by $B \equiv (b^g, B^b, b^b)$.

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

$$W^s(v_w) \equiv \begin{cases} V(y^s(1)) - v_h \times l_h^s - v_w & \text{if } v_w < \bar{v}_w^s \\ V(y^s(0)) - v_h \times l_h^s & \text{if } v_w \geq \bar{v}_w^s. \end{cases}$$

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

$$\max_{B,T} J(B,T) \quad \text{s.t. } \mu^g(1-e_w^g)b^g + \mu^b(e_w^b B^b + (1-e_w^b)b^b) = \mu^g(T_h^g + e_w^g T_w^g). \tag{10}$$

To solve the planner's problem we characterize the first-order conditions of the program in (10) by perturbing the tax-and-benefit system. We consider the optimal distribution of benefits to households with non-working spouses across states b and g. To do so, we analyze a small increase in b^b financed by a corresponding balanced-budget decrease in b^g. In the simple model, this captures the efficient distribution of transfers to low-income households across different health states. Any other perturbation of the system will follow the steps of the analysis conducted below, and the complete optimal system can thus be characterized in the same manner. We focus on this particular aspect of the policy since it captures the essence of insuring households against shocks in a simple and policy-relevant way.

The welfare gain from a \$1 (balanced-budget) increase in b^b is $\frac{dJ(T,B)}{db^b} = \mu^b \frac{\partial}{\partial b^b}\left(\int_0^\infty W^b(v_w) f(v_w) dv_w\right) + \mu^g \frac{\partial}{\partial b^g}\left(\int_0^\infty W^g(v_w) f(v_w) dv_w\right) \frac{db^g}{db^b}$. Since this is expressed in utility units with no cardinal interpretation, we follow the recent social insurance literature and normalize it by the welfare gain from a \$1 transfer to households with non-working spouses in the good state, scaled by the targeted population (Chetty and Finkelstein 2013). That is, the normalized net gain that we analyze is $M_W(b^b) \equiv \frac{\frac{dJ(T,B)}{db^b}/\mu^b(1-e_w^b)}{\frac{\partial J(T,B)}{\partial b^g}/\mu^g(1-e_w^g)}$. Differentiating the budget constraint to calculate $\frac{db^g}{db^b}$ and using the household's choices, which imply that $\frac{\partial}{\partial b^b}\left(\int_0^\infty W^b(v_w) f(v_w) dv_w\right) = u_w'(c_w^b(0))(1-e_w^b)$ and $\frac{\partial}{\partial b^g}\left(\int_0^\infty W^g(v_w) f(v_w) dv_w\right) = u_w'(c_w^g(0))(1-e_w^g)$, yield the normalized welfare gain

$$M_W(b^b) = MB(b^b) - MC(b^b), \tag{11}$$

where the marginal benefit is $MB(b^b) \equiv \frac{u'_w(c^b_w(0))-u'_w(c^g_w(0))}{u'_w(c^g_w(0))}$, the marginal cost is $MC(b^b) \equiv \frac{\varepsilon(1-e^b_w,b^b)-\varepsilon(1-e^g_w,b^g)}{1+\varepsilon(1-e^g_w,b^g)}$, and $\varepsilon(1-e^s_w,b^s) = \frac{\partial(1-e^s_w)}{\partial b^s}\frac{b^s}{(1-e^s_w)}$ is the elasticity of the unaffected spouse's non-participation with respect to government benefits. Note that when the consumption of h is positive (e.g., when he or she survives the shock), $MB(b^b)$ is also the gap in his or her marginal utilities due to consumption allocation efficiency in the household, which is determined by the program in (1).

Equation (11) is a simple variant of Baily's (1978) and Chetty's (2006a) formula for the optimal level of social insurance. The marginal benefit from a balanced-budget increase in b^b is captured by the insurance value of transferring resources from the good to the bad state, which is measured by the gap in marginal utilities of consumption across the two states. This "rate of return" on shifting funds, which is zero in the first-best allocation in which marginal utilities are smoothed across states of nature, measures market inefficiency and quantifies the potential benefit from government intervention. The marginal cost of transferring \$1 across states is due to behavioral responses, which capture the fiscal externality that households impose on the government budget when changing their participation decisions. In our case, the government's revenue could decrease since there are more spouses not working in the bad state due to higher benefits but could increase since there are fewer spouses not working in the good state as they receive fewer transfers.

Identifying the Benefits of Social Insurance. While estimating the marginal cost is conceptually straightforward, estimating the marginal benefit is challenging since it requires knowledge of the consumption utility function, particularly of the value of risk aversion, and of each individual's overall consumption. To circumvent the challenges posed by this consumption-based approach, which we discuss below, we use simple but powerful implications of the household's labor supply decisions, which allow us to rewrite the marginal benefit solely in terms of the unaffected spouse's labor supply. The following proposition summarizes this main welfare result and demonstrates the way in which the unaffected spouse's labor supply behavior fully reveals the gap in the marginal utilities of consumption across states of nature. We provide a simple proof and then discuss the intuition behind the formula; namely, that it identifies the gains of insurance by evaluating changes in the consumption of leisure.

Proposition 1. *Under a locally linear approximation of F in the threshold region $(\bar{v}_w^g, \bar{v}_w^b)$, the marginal benefit from raising b^b by \$1 is*

$$MB(b^b) \cong L^b + M^b, \tag{12}$$

where $L^b \equiv \frac{e_w^b - e_w^g}{e_w^g}$ *and* $M^b \equiv \left(\frac{\left|\varepsilon(e_w^b, b^b)\right|/b^b}{\left|\varepsilon(e_w^g, b^g)\right|/b^g} - 1 \right) \frac{e_w^b}{e_w^g}$.[35]

Proof. Recall that the unaffected spouse works when the value of additional consumption from his or her labor income, $\bar{v}_w^s \equiv V(y^s(1)) - V(y^s(0))$, outweighs his or her disutility from labor, v_w. This decision rule reveals the household's consumption value of an additional dollar, $V'(y^s(0))$, through the change in the critical labor-disutility threshold below which the spouse works ($\bar{v}_w^s$) in response to an increase in benefits, since $\left|\frac{\partial \bar{v}_w^s}{\partial b^s}\right| = V'(y^s(0))$. In addition, since (1) implies that $V'(y^s(0)) = u_w'(c_w^s(0))$, we can rewrite the marginal benefit from social insurance using the change in the marginal entrant's disutility of labor – that is, $MB(b^b) = \left(\left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right| - \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right| \right) / \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right|$. The last step to represent $MB(b^b)$ by using labor supply responses of the unaffected spouse is to map this expression onto directly observable participation rates, $e_w^s = F(\bar{v}_w^s)$, and their elasticities, $\varepsilon(e_w^s, b^s)/b^s = \frac{f(\bar{v}_w^s)}{F(\bar{v}_w^s)} \frac{\partial \bar{v}_w^s}{\partial b^s}$, with simple algebra. Together, the equalities $MB(b^b) = \left(\left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right| - \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right| \right) / \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right|$, $e_w^s = F(\bar{v}_w^s)$, $\varepsilon(e_w^s, b^s)/b^s = \frac{f(\bar{v}_w^s)}{F(\bar{v}_w^s)} \frac{\partial \bar{v}_w^s}{\partial b^s}$, and the approximation in the proposition yield the result.

This formula shows that the marginal benefit from social insurance can be fully recovered from two moments of the unaffected spouse's labor supply, which we examine successively.

The first term, L^b, is composed of the unaffected spouse's labor supply response to the shock – or the labor force participation "shock elasticity" – which captures exactly the self-insurance role of the spouse's labor supply. Recall from the comparative statics of our model (equation (4) in Section II) that the relative increase in spousal labor force participation across states of nature increases with the income loss due to the shock. Therefore, the participation response reveals the degree of

[35] This local first-order expansion of $F(v_w)$ is supported by the empirical analysis of the spouse's participation across states of nature, which implies that $\bar{v}_w^g$ and $\bar{v}_w^b$ are within a small region of the support $[0, \infty)$. If one wishes to avoid this approximation, one can accompany the analysis with assumptions regarding the family of distributions to which F belongs, and then calibrate its parameters with the participation rates observed in the data. Note that this approximation is isomorphic to a second-order approximation of the search effort function in a search model of participation that we analyze in Online Appendix A.

the household's income loss and hence the extent to which the household needs additional insurance against the shock. This term is exactly the key moment of our empirical analysis in Section V.

The second term, M^b, captures the gains from the consumption of leisure by the marginal spouses due to behavioral responses to the policy change. When we increase benefits to non-working spouses in the bad state, b^b, we let more spouses meet their consumption needs if they choose not to work and consume more leisure – which is a welfare gain from the individual's and hence from the planner's perspective. The relative share of spouses who are on the labor force participation margin is captured by the semi-elasticity $\left|\varepsilon(e_w^b, b^b)\right| / b^b$, which quantifies the percent change in labor force participation in state b when we increase non-participation transfers b^b by \$1. This is illustrated in Figure VIII: Panel A depicts the pre-perturbation labor force participation in state b, and Panel B depicts the response of the spouses that are on the participation margin in state b. Since we finance the increase in b^b by a decrease in b^g, the marginal spouses in state g who now work as a response – and whose relative share is $\left|\varepsilon(e_w^g, b^g)\right| / b^g$ – represent a welfare loss due to their reduced consumption of leisure. Therefore, the net gain through the change in the consumption of leisure due to the policy change is captured by $\frac{\left|\varepsilon(e_w^b, b^b)\right|/b^b}{\left|\varepsilon(e_w^g, b^g)\right|/b^g} - 1$.[36] To scale these within-state elasticities into cross-state terms (which are relevant for our cross-state perturbation), we multiply this gain by the relative labor supply across states, $\frac{e_w^b}{e_w^g}$ (which we estimated in Section V).[37] This results in the second term of the formula: $M^b \equiv \left(\frac{\left|\varepsilon(e_w^b, b^b)\right|/b^b}{\left|\varepsilon(e_w^g, b^g)\right|/b^g} - 1 \right) \frac{e_w^b}{e_w^g}$. Whenever we transfer resources from the good to the bad state, the formula adjusts through the semi-elasticity ratio that enters this term; it is always the ratio of the responses to the specific policy tools that we consider changing.

Discussion. The alternative method for recovering welfare gains from social insurance is consumption-based and aims at directly identifying the gap in marginal utilities of consumption across states of nature. The reduced-form literature uses the approach that was developed by Baily (1978) and Chetty (2006a) and was first implemented by Gruber (1997) in the context of unemployment in-

[36] Note that **within** a state, marginal spouses are indifferent between working and not working. In the absence of full insurance, this is not the case **across** states, which is the relevant comparison for our policy change and is represented by the semi-elasticity ratio.

[37] Recall that we study the welfare implications of transferring resources from state g to state b. This transfer induces behavioral responses within each state of nature, which are expressed here in terms of semi-elasticities, since it changes the economic incentives within each state. However, to evaluate the implications of these elasticities in "cross-state" rather than "within-state" terms, we need to scale the elasticity ratio by the relative labor supply flow across states, $\frac{e_w^b}{e_w^g}$. The first term of the welfare formula (L^b) is already in cross-state terms so that no scaling is required.

surance. This approach is based on analyzing consumption fluctuations across states, which are transformed to utility losses with estimates for the curvature of the utility function. The structural literature follows a similar approach but with the additional complexity of estimating the full set of the economic model's primitives. Our approach maps the identification problem from the consumption domain to the labor supply domain. By doing so, it does not rely on assumptions regarding the appropriate value of risk aversion about which there is tremendous uncertainty in the literature and to which the consumption-based calculations of gains from insurance are highly sensitive (Chetty and Finkelstein 2013). Additionally, it requires only data from the labor market, which are typically more precise and widely available than consumption data. While consumption measures are usually partial (and cover only a sub-set of goods, such as expenditure on food), and strong assumptions are needed to translate overall expenditure into individuals' consumption bundles, labor market data exactly match the theoretical behaviors of interest, namely, participation and earned income.

Two other labor-market methods have been developed in the context of unemployment in Chetty (2008) and Shimer and Werning (2007), which are based on the labor supply responses of the directly affected individual.[38] However, these methods are not applicable to the case of a severe health shock (or, of course, of a death event) in which the affected individual's labor supply can no longer identify preferences. This is because a non-negligible share of those experiencing severe health shocks (and whose ability to work is directly affected) may be forced out of the labor market and become unresponsive to economic incentives. Their implied small behavioral responses to changes in policy tools and other economic incentives may wrongly imply a low value of additional insurance, while they are actually driven by significant shocks to their ability to work. Our approach falls within this group of labor market approaches but extends the scope of identifying welfare gains from social insurance using labor supply responses. In particular, it can be applied to important cases in which the directly affected individual may be at a corner solution, such as a severe health shock or the extreme case of death.[39]

[38] These methods as well as our own identify directly-estimable moments that are sufficient for welfare analysis. The advantage of the sufficient statistics approach to welfare analysis is that it offers results about optimal policy that do not utilize strong assumptions that are made in structural studies for tractability and identification. The cost is that without extrapolations and additional structure it can only be used to analyze marginal changes in policy. See Chetty (2009) for a more detailed discussion on this issue.

[39] We discuss additional such cases in the Conclusion.

The analysis above has also shown that, in contrast to conventional wisdom, the level of optimal benefits does not necessarily decrease in the degree of crowd-out of self-insurance by social insurance. It is indeed the case that increased benefits to non-working spouses in the bad state impose a fiscal externality on the government's budget through an increase in this group's non-participation rate, which is captured by the non-participation elasticity $\varepsilon(1 - e_w^b, b^b)$ in $MC(b^b)$. However, at the same time, the decreased participation entails a gain from consumption of additional leisure, which is captured by the participation elasticity $\varepsilon(e_w^b, b^b)$ in $MB(b^b)$. Therefore, our analysis formalizes Gruber's (1996) argument that in any assessment of net welfare gains from social insurance both effects have to be taken into account and weighted appropriately.

Identifying Assumption: Efficiency. Before we proceed with extensions to the basic model, it is worth emphasizing the source of identification of the household's preferences by using the unaffected spouse's labor supply responses. The key assumption underlying our analysis is that household decisions are Pareto efficient. This implies that on the margin, all members of the household exhibit the same returns to additional resources; hence any member not at a corner solution can reveal the preferences of each member of the household.

This approach relies on the premise that when spouses have symmetric information about each other's preferences and consumption (because they interact on a regular basis) we would expect them to find ways to exploit any possibilities of Pareto improvements. Importantly, as emphasized by Browning, Chiappori, and Weiss (2014), this does not preclude the possibility of power issues such that the allocation of resources within the household can depend on its members' respective Pareto weights. The approach simply assumes that no resources are left on the table. An additional advantage of the collective model is that it does not require specifying the mechanism that households use, e.g., the bargaining process, but only assumes such a mechanism exists. Note that the unitary model is a special case of our collective framework, and therefore our results readily apply to the unitary assumption that is widely used in models of the household.[40]

[40] There are some cases in which the efficiency assumption fails (see discussion in Browning, Chiappori, and Weiss 2014). To model these cases, one would need to specify the underlying model of household decision making and make additional assumptions in order to identify one spouse's preferences from the other spouse's behavior.

VI.B Accounting for State-Dependent Preferences

To adjust our normative result of Proposition 1 to account for potential changes to preferences in response to shocks we consider the generalized preference structure of Section II.B. In what follows we successively describe how different types of state dependence in the household's preferences may affect the analysis. Since our welfare method identifies gains from the labor supply behavior of the unaffected spouse, the sort of state dependence that affects the normative analysis is confined to potential changes in the unaffected spouse's labor disutility. The adjustment of our welfare formulas is summarized in Proposition 2.

First, recall that the state-dependent preference structure allows for a completely flexible dependence of consumption utility on the state of nature and for flexible consumption-leisure complementarities. These two extensions to the baseline preferences have no effect on the welfare formulas since we mapped the welfare evaluation problem from the consumption domain completely onto the labor force participation domain. This simplifies the analysis tremendously since the estimation of consumption utility state dependence has proven very challenging (Finkelstein, Luttmer, and Notowidigdo 2009), and also allows us to avoid the common practice of assuming consumption-leisure independence.

Next, recall that the labor disutility, v_i^s, can entirely change across states of nature. Since we identify welfare gains from the behavior of the unaffected spouse, allowing the labor disutility of the affected sick spouse, v_h^s, to change completely across states of nature does not affect the analysis. It is indeed the underlying motive for studying the unaffected spouse's behavior in the first place since in the case of health shocks the affected spouse's preferences can change in many unidentifiable ways as a result of the shock. However, the potential state dependence of the unaffected spouse's labor disutility, v_w^s, requires adjusting the welfare formula in the following way.

Proposition 2. *Assume the generalized preference structure of Section II.B. Under a locally linear approximation of F in the threshold region $(\bar{v}_w^g, \bar{v}_w^b)$, the marginal benefit from raising b^b by \$1 is*

$$MB(b^b) \cong L^b + M^b + S^b, \tag{13}$$

where $L^b \equiv \frac{e_w^b - e_w^g}{e_w^g}$, $M^b \equiv \left(\frac{\left|\varepsilon(e_w^b, b^b)\right|/b^b}{\left|\varepsilon(e_w^g, b^g)\right|/b^g} - 1 \right) \frac{e_w^b}{e_w^g}$, *and* $S^b \equiv \left(\theta^b - 1\right)\left(1 + L^b + M^b\right)$.

Proof. With the generalized preferences, it is straightforward to show that $MB(b^b) = \left(\theta^b \left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right| - \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right|\right) / \left|\frac{\partial \bar{v}_w^g}{\partial b^g}\right|$. Combining this equality with $e_w^s = F(\bar{v}_w^s)$, $\varepsilon(e_w^s, b^s)/b^s = \frac{f(\bar{v}_w^s)}{F(\bar{v}_w^s)} \frac{\partial \bar{v}_w^s}{\partial b^s}$, and the approximation in the proposition yields the result.

The additional component, $S^b \equiv \left(\theta^b - 1\right)\left(1 + L^b + M^b\right)$, essentially "prices" in utility terms the cost of the first two labor supply "quantity" expressions, L^b and M^b. The unaffected spouse's labor supply is more costly by $\theta^b - 1$ percent. This additional cost needs to be applied to the overall relative labor supply response across health states, i.e., the sum of the baseline participation rate (normalized to 1) and the two quantity components: $1+L^b+M^b$. Since our welfare method identifies the gains from insurance by evaluating the change in the consumption of leisure, higher valuation of leisure, that is, $\theta^b > 1$, renders leisure more valuable in state b, which makes the transfer of resources from state g to state b more socially desirable in our framework. On the other hand, lower cost of labor supply due to the shock ($\theta^b < 1$) can lead to an increase in labor force participation even if households are well insured. The welfare implications of labor supply responses in this case are different (so that additional transfers to the bad state may even become undesirable) since these responses would be driven by preferences for work and not by under-insurance.

VI.C Additional Generalizations and Extensions

Dynamic Life-Cycle Model. In the context of social insurance over the life-cycle, it is important to consider households' self-insurance through ex-ante mechanisms such as precautionary savings. In Online Appendix A, we analyze life-cycle participation decisions using a dynamic search model that allows for endogenous savings.[41] The general result of this analysis is that our formulas extend to the dynamic case with the adjustment that post-shock responses in the static case are replaced by mean responses when a shock occurs. This is exactly what we recover in our empirical analysis. Hence, our results as well as the welfare analysis we conduct below readily apply to the dynamic case.[42] The intuition behind this theoretical result is that responses of forward-looking households to shocks internalize the full expected path of future consumption and leisure. Therefore, responses in

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

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

periods right after a shock occurs reveal the household's life-time welfare implications of additional transfers.[43]

The dynamics of the life-cycle analysis likewise enter the marginal costs of social insurance. A household in state g not only decreases its labor supply due to higher taxes in the present, but also in response to increased benefits in the hitherto unencountered state b. The prospect of higher benefits in the case that the household experiences a shock lowers its need to save for that scenario, which translates into a decrease in labor supply in state g.

Intensive-Margin Model of Labor Supply. One can construct similar formulas for the case in which the household's intensive labor supply decisions are considered. With no individuals on the participation margin, the formulas consist only of the labor supply changes across states of nature and the potential change in the utility cost of labor. See Online Appendix B for an analysis of this model. Note that the choice of the appropriate model for welfare analysis should depend on the data. For example, studying a sub-population with full employment before a shock occurs calls for the intensive-margin model because in such a case work intensity is the operative margin.[44]

Home Production and Additional Decision Variables. Recall that our propositions for identifying the welfare gains from social insurance rely on optimality conditions that are implied by the household's labor supply choices and are derived using the envelope theorem. Since these optimality conditions must hold in our model when we add arbitrary decision variables, we can still identify the gains from insurance based on the unaffected spouse's labor supply in that case. Therefore, our welfare results hold in the generalized setting that accounts for other decision variables that may be included in the household's optimization problem. Importantly, we can easily incorporate the household's time use decisions, e.g., of how to allocate non-work time between home production and leisure. The robustness of our approach to the inclusion of additional margins of response is a

[43]The setting we analyze in the online appendix also extends the model by allowing for multiple and sequential shocks. In particular, we analyze a model in which h can experience a health shock and may die as a consequence. This illustrates our analysis in a more complex and realistic setting that can be applied to different types of sequential shocks.

[44]Studying the discrete participation decision rather than the intensive-margin decision has several important advantages. First, it allows for flexible consumption-leisure complementarities. Second, it captures additional moral hazard responses that the social insurance literature discusses. By modeling means-tested transfers that can condition on household-level income we can study the welfare effects of the potential crowd-out of spousal labor force participation. Third, labor market frictions (such as hour requirements set by employers) can limit employees' ability to optimize; hence participation decisions may reveal preferences more accurately (since the potential costs of non-optimization are higher).

general feature of the sufficient statistic approach to welfare analysis (see Chetty 2006a).

VII Welfare Analysis: Implications of Empirical Results

We now turn to illustrate our method for welfare analysis and to study the welfare implications of the surviving spouses' labor supply responses. In accordance with the theoretical analysis in the previous section, consider the following policy question: how should we divide a given budget between households of widows (for which we find a large increase in labor supply) and non-widows in our framework? This is essentially a comparison of the social "returns" of two "investment" vehicles – a \$1 transfer to non-working spouses in state b that yields a return of $u'_w(c^b_w(0))$ vs. a \$1 transfer to non-working spouses in state g that yields a return of $u'_w(c^g_w(0))$. This comparison focuses on the marginal benefit from increasing b^b by lowering b^g, and puts aside the associated costs on which the literature has focused. We employ the formula of Proposition 1 (equation (12)) and abstract from labor disutility state dependence by assuming $\theta^b = 1$. Since the evidence is consistent with $\theta^b \geq 1$, this approach delivers a lower bound on the potential welfare gains from this policy change (as implied by Proposition 2).

In order to assess the marginal benefit from this policy perturbation, we need to calibrate the ratio $\frac{\varepsilon(e^b_w,b^b)}{\varepsilon(e^g_w,b^g)}$. Here we make the simplifying assumption of equal elasticities and use the approximation that this ratio is locally constant, which allows us to illustrate our method in the simplest possible way.[45] Assuming that $\frac{\varepsilon(e^b_w,b^b)}{\varepsilon(e^g_w,b^g)} = 1$, the formula for the welfare benefits reduces to

$$MB(b^b) \cong \frac{b^g}{b^b} \times \frac{e^b_w}{e^g_w} - 1,$$

where $\frac{e^b_w}{e^g_w}$ is the moment we estimated in Section V.A.

To study existing social programs in Denmark, we divide the analysis into two sub-populations. First, we consider widows over age 67, who are eligible for the Danish Old-Age Pension (the equivalent of Social Security in the US) and analyze the perturbation within this program. Second, we consider widows younger than 67, who are more attached to the labor force, and analyze changes

[45] In Online Appendix G we add some structure and estimate these elasticities for surviving spouses under age 60. The estimates imply a ratio of 1.375, which suggests that our calibration provides a lower bound for the welfare gains from the policy change.

to social survivors benefits for which they can apply through Social DI .

Old-Age Pension. In Panel A.1 of Figure IX we plot the labor supply responses of widows over 67. The graph reveals that even the elderly need to self-insure and increase their participation by 1.08 pp on a very low base of 1.19 pp. This implies that the participation rate of widows over 67 almost doubles when their husbands die with $\frac{e_w^b}{e_w^g} = \frac{0.0227}{0.0119} = 1.91$. As the Old-Age Pension includes adjustments to the household's composition (as explained in Section III and seen in practice in Panel A.2 of Figure IX), widows during our sample period received on average DKK 87,454 ($10,932) and their non-widow counterparts received DKK 70,684 ($8,836) such that for this population $\frac{b^g}{b^b} = \frac{70{,}684}{87{,}454} = 0.81$. Together, these imply that

$$MB(b^b) \cong \frac{b^g}{b^b} \times \frac{e_w^b}{e_w^g} - 1 = 0.81 \times 1.91 - 1 = 0.55.$$

That is, an additional $1 transferred to widows through the Old-Age Pension creates a net benefit equivalent to 55 cents as compared to transferring $1 to non-widows. This large marginal benefit from an additional dollar to elderly widows is driven by their significant relative increase in participation, which reveals their high valuation of additional insurance. This suggests that in our conceptual framework increasing the relative compensation to **older** widows within the Old-Age Pension (beyond the current household-composition adjustment) entails significant welfare improvement.

Social Survivors Benefits. To focus on the value of survivors benefits through Social DI, we constrain the sample to widows under 67 (the age at which the program transitions into the Old-Age Pension). In addition, we constrain the sample to the period prior to 1994 due to a data break in the reporting method of benefits received through Social DI. Panel A.3 of Figure IX plots the labor force participation behavior of this sample and shows that $\frac{e_w^b}{e_w^g} = \frac{0.4718}{0.4537} = 1.04$, which is smaller than the effect among the elderly as well as among the overall sample of widows as shown in Section V.A. Panel A.4 of Figure IX clearly displays the insurance role of Social DI for widows, whose take-up of the program increases by more than 50% in the year that their husbands die. For this time period, the mean benefits received from Social DI by those on the program are the same for

widows and for non-widows and, therefore, $\frac{b^g}{b^b} \cong 1$.[46] Combining these estimates, it follows that

$$MB(b^b) \cong \frac{b^g}{b^b} \times \frac{e_w^b}{e_w^g} - 1 = 1 \times 1.04 - 1 = 0.04.$$

That is, an additional \$1 transfer to widows through Social DI is worth 4 cents more to each household than is transferring this additional \$1 to non-widows. These small (but positive) welfare gains are a direct result of the smaller relative increase in labor force participation among widows under 67 that are eligible for this program.

Therefore, a key implication of our findings, driven in part by the differential attachment to the labor force over the life-cycle, is that age-dependence is a feature of the optimal social insurance policy for spousal mortality shocks in our setting.

An additional valuable welfare exercise allows us to use our method to assess how far the benefits are from our model's prediction of their optimal level, by evaluating the local rate at which marginal benefits change, $MB'(b^b)$. To evaluate this derivative we take advantage of the spatial variation in the administration of survivors benefits through Social DI, which we analyzed in Section V.A. Recall that our estimate for $\frac{\partial(e_w^b - e_w^g)}{\partial b^b}$ is -.0057. Using this estimate, Panel B of Figure IX plots the behavior of $MB(b^b)$ around the sample mean of DKK 65,000 (\$8,115). The figure shows that in our model additional DKK 1,500 (\$188) in annual benefits decrease the excess benefit to zero. Converting these monetary values into net replacement rates out of the deceased spouse's pre-shock earnings,[47] the current system stands at 0.648 and the optimal allocation of benefits across states stands at 0.663, suggesting that for **younger** widows the current levels are near optimal. To evaluate the overall value of the program, we can approximate the integral $\int_0^{65,000} MB(b^b)db^b$ by using our estimates. This integral answers the question: within the Social DI system in Denmark, what is the welfare gain from the benefits given to widows relative to non-widow beneficiaries of the program? The estimate amounts to DKK 99,942 (\$12,500) annually, which means that transfers to widows

[46] The exact figures are b^b =\$8,115 for widows and b^g =\$8,016 for non-widows (in 2000 dollars). These are averages of the actual benefits received by each group, which are a result of adjusting the statutory benefit amounts to income-testing at the household level.

[47] We calculate the average earnings in $t = -1$ for affected spouses who had positive earnings the year before they passed away. The average is DKK 170,000 (\$21,250), which implies net wage earnings of DKK 100,300 (\$12,538) using an average labor income tax rate of 41% (OECD estimates).

relative to non-widows create a benefit of ($12,500/$8,115-1=) 54%. That is, on average, each dollar given to younger widows through Social DI generates a net benefit equivalent to 54 cents relative to a dollar given to non-widow recipients. This reveals the large value of social survivors benefits within our framework.

VIII Conclusion

This paper provides evidence of household self-insurance through labor supply in response to large and persistent income losses, and develops a new labor market method for welfare analysis of social insurance. Studying the critical event of the death of a spouse, we find large increases in the surviving spouses' labor force participation rate driven by households for whom this event imposes significant income losses. We show that the unaffected spouse's self-insurance response identifies the household's marginal utility from consumption. As the gap in marginal utilities across states of nature captures the value of insurance, we offer a way to recover the gains from social insurance based solely on spousal labor supply responses. Applying this method to spousal mortality shocks, we show that in our conceptual framework allocation of additional resources to elderly widows has significant welfare gains and that the optimal structure of survivors benefits is age-dependent.

We additionally exploit the Danish setting to analyze households in which an individual has experienced a severe health shock but survived, for whom income losses are well-insured. Together, the results point to a potential explanation for the elusiveness of the insurance role of spousal labor supply in previous literature. In support of the hypotheses raised by Heckman and MaCurdy (1980) and Cullen and Gruber (2000), we find that spousal labor supply plays a significant self-insurance role when the income loss incurred by the shock is large relative to the household's lifetime income – as in the death of a spouse – and is irrelevant when the loss is sufficiently insured through formal social insurance – as in spousal health shocks.

Our findings have further implications for potentially improving efficiency in the distribution of government benefits. The significant heterogeneity in responses that we find across different **pre-shock** dimensions of household characteristics suggests that enriching the policy tools to condition transfers on these observable characteristics may be welfare improving. For example, since increases in the surviving spouse's labor supply are strongly correlated with the income shock that he or she

experiences after losing an earning spouse, it may be welfare improving to let survivors benefits increase in the deceased spouse's pre-shock share of annual household earnings.[48]

More broadly, our quasi-experimental design for identifying the effect of shocks as well as our method for welfare analysis can be applied to other important economic questions. Our research design, which relies on comparing households that are affected only a few years apart, can be applied to estimating the effect of a shock in any setting in which its exact timing is likely to be random. Our welfare analysis method, which relies on spousal labor supply, can be applied to evaluating the welfare gains from social insurance in any setting in which the directly affected individual may be at a corner solution. For example, relevant to the debate on the privatization of Social Security, the value of protecting against pension-wealth losses in the 401(k) account of a working individual can be recovered by the labor supply response of his or her spouse. Spousal labor supply can also be used to evaluate the welfare losses caused by the discontinuation of an employee's compensation, such as health insurance, as well as the value of unemployment insurance for the long-term unemployed (whose long durations of unemployment significantly harm their employment prospects).[49]

Itzik Fadlon, National Bureau of Economic Research

Torben Heien Nielsen, University of Copenhagen

[48] A similar feature is implicit in the US system, where survivors are eligible for their deceased spouses' Social Security benefits, which are a function of the deceased's work history.

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

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

The model we analyze in this appendix generalizes our baseline model in two ways. Most importantly, it analyzes life-cycle participation decisions using a dynamic search model, which allows for endogenous savings. Second, we extend the one-shock model we analyzed in the text to include different and potentially sequential shocks. In addition, we use the generalized preference structure of Section II.B in the text and allow for additional extensions to it as we describe below.

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

At the beginning of the planning period, w does not work and searches for a job. When w enters period t in state s without a job he or she chooses search intensity, e^s_{wt}, which we normalize to equal the probability of finding a job in the same period. If w finds a job, the job begins at time t and is assumed to last until the end of the planning period once found.[1]

Individual Preferences. Let $u^s_{it}(c^s_{it}, l^s_{it}, l^s_{jt})$ represent i's flow consumption utility at time t in state s as a function of consumption, c^s_{it}, labor force participation, l^s_{it}, and the other spouse's labor force participation, l^s_{jt}, where $\frac{\partial u^s_{it}}{\partial c^s_{it}} > 0$ and $\frac{\partial^2 u^s_{it}}{\partial (c^s_{it})^2} < 0$. We denote w's cost of search effort at time t in state s by $\kappa^s(e^s_{wt})$, which we assume to be strictly increasing and convex. The relative cost of time invested in search effort across states is captured by $\theta^\sigma(e^\sigma_{wt}) \equiv \kappa^{\sigma\,\prime}(e^\sigma_{wt})/\kappa^{g\,\prime}(e^\sigma_{wt})$, where $\sigma \in \{b, d\}$.

Household Preferences. We assume that the household's per-period utility weights individual utilities according to their respective Pareto weights β_w and β_h, such that the household's flow utility at time t in state s is $\beta_w u^s_{wt} + \beta_h u^s_{ht}$ augmented by w's weighted search cost $\beta_w \kappa^s(e^s_{wt})$ when w is unemployed. We assume equal Pareto weights and normalize $u^d_{ht} = 0$. In the following analysis we suppress the dependence of the consumption utility on participation for ease of notation only.

Policy Tools. The planner observes the state of nature as well as the employment status of each spouse. Since some spouses work and earn more than others do, the optimal policy is dependent on whether the

[1] This simplifies the algebra of the analysis. We later allow for job separations such that employment is absorbing within a health state but not across health states.

spouse is employed. We denote the tax on spouse i's labor income in state g by T_i^g and the benefits given to non-working spouses in state g by b^g. In state $\sigma \in \{b, d\}$, households in which the unaffected spouse, w, works receive transfers of the amount B^σ and households in which w does not work receive benefits of the amount b^σ. We denote taxes by $T \equiv (T_w^g, T_h^g)$ and benefits by $B \equiv (b^g, B^b, b^b, B^d, b^d)$, and let $B^s(l_{wt}^s)$ represent the actual transfers received by a household in state s as a function of w's participation.

Household's Problem. The household's choices include the allocation of consumption to each spouse, c_{it}^s, as well as w's search effort if w is unemployed, e_{wt}^s. In each period, w's employment status, l_{wt}^s, determines the household's income flow, $y_t^s(l_{wt}^s)$, such that $y_t^s(l_{wt}^s) = \bar{z}_{ht}^s \times l_{ht}^s + \bar{z}_{wt}^s \times l_{wt}^s + B^s(l_{wt}^s)$, where z_{it} is i's labor income and $\bar{z}_{it}^s = z_{it} - T_i^s$ is i's labor income net of taxes in state s (with $T_i^\sigma = 0$, $\sigma \in \{b, d\}$). This implies that each period's consumption as well as the next period's wealth – where we denote assets in period t by A_t – are functions of w's participation, which we denote by $c_{it}^s(l_{wt}^s)$ and $A_{t+1}(l_{wt}^s)$, respectively. Therefore, the value function for households in state s who enter period t when w is without a job and with household assets A_t is

$$V_t^{s,0}(B,T,A_t) \equiv \max \left\{ \begin{array}{c} e_{wt}^s \left(u_h^s(c_{ht}^s(1)) + u_w^s(c_{wt}^s(1)) + W_{t+1}^{s,1}(B,T,A_{t+1}(1)) \right) \\ +(1-e_{wt}^s) \left(u_h^s(c_{ht}^s(0)) + u_w^s(c_{wt}^s(0)) + W_{t+1}^{s,0}(B,T,A_{t+1}(0)) \right) - \kappa_w^s(e_{wt}^s) \end{array} \right\},$$

where the budget constraints satisfy

$$c_{ht}^s(l_{wt}^s) + c_{wt}^s(l_{wt}^s) + A_{t+1}(l_{wt}^s) = A_t + y_t^s(l_{wt}^s),$$

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

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

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

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

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

$$V_t^{s,1}(B,T,A_t) \equiv \max \left\{ u_h^s(c_{ht}^s(1)) + u_w^s(c_{wt}^s(1)) + W_{t+1}^{s,1}\left(B,T,A_{t+1}(1)\right) \right\}.$$

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

$$\begin{aligned}&\Big\{\Big(u_h^s(c_{ht}^s(1))+u_w^s(c_{wt}^s(1))+W_{t+1}^{s,1}(B,T,A_{t+1}(1)\Big)\\&-\Big(u_h^s(c_{ht}^s(0))+u_w^s(c_{wt}^s(0))+W_{t+1}^{s,0}(B,T,A_{t+1}(0))\Big)\Big\}\end{aligned}\quad=\quad\kappa_w^{s\,\prime}(e_{wt}^s),\tag{1}$$

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

$$\frac{\partial e_{wt}^s}{\partial b^s}=-\frac{1}{\kappa_w^{s\,\prime\prime}(e_{wt}^s)}\left(u_w^{s\,\prime}(c_{wt}^s(0))+\frac{\partial W_{t+1}^{s,0}}{\partial b^s}\right).\tag{2}$$

Planner's Problem. We define the household's expected utility at the beginning of the planning period by $J_0(B,T)\equiv(1-\rho_0)V_0^{g,0}(B,T,A_0)+\rho_0 V_0^{b,0}(B,T,A_0)$. The social planner's objective is to choose the tax-and-benefit system that maximizes the household's expected utility subject to a balanced-budget constraint. For simplicity, we assume there is some expected revenue collected from each household and study the optimal redistribution of this revenue. We abstract from the specific way in which revenue is collected (or, similarly, assume a lump-sum tax that is determined outside of our problem) since our focus is on the benefits from social insurance and not its fiscal-externality costs. The perturbations we study involve increasing b^σ, $\sigma\in\{b,d\}$, by lowering b^g. Therefore, to further simplify the analysis we assume that $B^b=B^d=0$, as well as that $b^d=0$ when we perturb b^b and that $b^b=0$ when we perturb b^d.

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

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

Optimal Social Insurance

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

$$\frac{dJ_0(T,B)}{db^b}=Q_1^b+Q_2^b\frac{db^g}{db^b},\tag{4}$$

where $Q_1^b \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial b^b} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial b^b}\right)$ and $Q_2^b \equiv \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial b^g} + (1-\rho_0)\frac{V_0^{g,0}}{\partial b^g}\right)$. The following proposition provides an approximated formula for the normalized version of this gain.

Proposition A1.

Under a locally quadratic approximation of the effort function $\kappa_w^g(e_{wt}^g)$ around e_{w0}^g and assuming that the ratio $\theta^b(e_{wt}^b)$ is locally constant at $\hat{e}_{w0}^b$, the marginal benefit from raising b^b by \$1 is

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

with

1. $MB(b^b) \equiv L^b + M^b + S^b$, *where* $L^b \equiv \frac{\hat{e}_{w0}^b - e_{w0}^g}{e_{w0}^g}$, $M^b \equiv \left(\frac{\left|\varepsilon(\hat{e}_{w0}^b, b^b)\right|/b^b}{\left|\varepsilon(e_{w0}^g, b^g)\right|/b^g} - 1\right)\frac{\hat{e}_{w0}^b}{e_{w0}^g}$, $S^b \equiv \left(\theta^b - 1\right)\left(1 + L^b + M^b\right)$, $\varepsilon(x,y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$, $\theta^b \equiv \theta^b(\hat{e}_{w0}^b)$, e_{w0}^g *is w's participation rate at the beginning of the planning period, and* $\hat{e}_{w0}^b$ *is w's mean participation rate in households that transition to state b.*

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

Proof.

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

We begin by providing expressions for $\frac{\partial V_0^{b,0}}{\partial b^b}$ and $\frac{\partial V_0^{g,0}}{\partial b^b}$ in order to characterize Q_1^b. First, we have that $\frac{\partial V_t^{b,0}}{\partial b^b} = (1 - e_{wt}^b)\left(u_w^b{}'(c_{wt}^b(0)) + \frac{\partial W_{t+1}^{b,0}}{\partial b^b}\right)$ and $\frac{\partial W_{t+1}^{b,0}}{\partial b^b} = (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial b^b}$, which imply that $\frac{\partial V_t^{b,0}}{\partial b^b} = (1 - e_{wt}^b)\left(u_w^b{}'(c_{wt}^b(0)) + (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial b^b}\right)$. Working backwards one can show that $\frac{\partial V_t^{b,0}}{\partial b^b} = (1 - e_{wt}^b)\left[u_w^b{}'(c_{wt}^b(0)) + \sum_{i=t+1}^{T-1}\left(\prod_{j=t+1}^{i}(1-e_{wj}^b)(1-\lambda_j)\right)\left(u_w^b{}'(c_{wi}^b(0))\right)\right]$.

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

[2] Specifically, $\beta_0^b \equiv \frac{\sigma^b D^b\left(1-\hat{e}_w^b\right) - D^g\left(1-\hat{e}_w^g\right)}{D^g\left(1-\hat{e}_w^g\right)}$, $\beta_1^b \equiv \sigma^b \frac{b^g}{b^b}$, and $\beta_2^b \equiv \sigma^b \frac{D^b\left(1-\hat{e}_w^b\right)}{D^g\left(1-\hat{e}_w^g\right)}$, where $\sigma^b \equiv (1-p_0)\left(1-e_{w0}^g\right)/\rho(1-\hat{e}_{w0}^b)$ and $\rho \equiv \sum_{i=0}^{T-1}\left(\prod_{j=0}^{i-1}(1-\rho_j)\rho_i\right)$. Note that the elasticities in $MC(b^b)$ consist of the total effect of increasing b^p, which takes into account the effect of lowering the level of the financing tool, b^g. Also note that with forward-looking households, transfers in states not yet encountered can have effects through ex-ante responses. For example, individuals in state g can lower labor supply and savings today in response to larger benefits in state b.

Putting the terms together, it follows that

$$Q_1^b = \left(\rho_0 \frac{\partial V_0^{b,0}}{\partial b^b} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial b^b}\right) = \sum_{i=0}^{T-1}\left[\prod_{j=0}^{i-1}(1-e_{wj}^g)(1-\rho_j)\rho_i\right]\frac{\partial V_i^{b,0}}{\partial b^b}. \tag{5}$$

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

$$Q_1^b = -\sum_{i=0}^{T-1}\left[\prod_{j=0}^{i-1}(1-e_{wj}^g)(1-\rho_j)\rho_i\right](1-e_{wi}^b)\kappa_w^b{}''(e_{wi}^b)\frac{\partial e_{wi}^b}{\partial b^b}. \tag{6}$$

To understand the meaning of this formula let us break it down into its components. First, note that it is a weighted sum of a function of the change in effort (or participation rate), $\frac{\partial e_{wi}^b}{\partial b^b}$. The weight, the term in brackets, is the probability of reaching period i with w unemployed and transitioning to state b exactly in that period. For households that transition to state b in period i when w is employed, the change in effort and participation rates is zero (because they stay employed and do not engage in search effort). Therefore, dividing the probability weights by the chance of transitioning to state b at some point throughout the planning horizon, $\rho \equiv \sum_{i=0}^{T-1}\left(\prod_{j=0}^{i-1}(1-\rho_j)\rho_i\right)$, and rewriting (6) in terms of elasticities (with $\varepsilon(x,y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$) yield $Q_1^b = \rho E_b\left\{(1-\bar{e}_{w0}^b)\kappa_w^b{}''(\bar{e}_{w0}^b)\left|\varepsilon(\bar{e}_{w0}^b, b^b)\right|\frac{\bar{e}_{w0}^b}{b^b}\right\} \equiv \rho E_b(g(\bar{e}_{w0}^b))$, where $\bar{e}_{w0}^p$ denotes participation in the period the household transitions to state b and E_b is the expectation operator conditional on being in state b. By expanding $g(e)$ around w's average participation in households in which h becomes sick – which we denote by $\hat{e}_{w0}^b$ – such that $g(e) \cong g(\hat{e}_{w0}^b) + g'(\hat{e}_{w0}^b)(e-\hat{e}_{w0}^b)$, we approximate $E_b(g(\bar{e}_{w0}^b)) \cong E_b(g(\hat{e}_{w0}^b)) = g(\hat{e}_{w0}^b)$ and obtain the approximation

$$Q_1^b \cong \rho(1-\hat{e}_{w0}^b)\kappa_w^b{}''(\hat{e}_{w0}^b)\left|\varepsilon(\hat{e}_{w0}^b, b^b)\right|\frac{\hat{e}_{w0}^b}{b^b}. \tag{7}$$

We now turn to provide expressions for $\frac{\partial V_0^{b,0}}{\partial b^g}$ and $\frac{V_0^{g,0}}{\partial b^g}$ in order to characterize Q_2^b. Since households that transitioned to state b either stay in state b or transition to state d, we have that $\frac{\partial V_0^{b,0}}{\partial b^g} = 0$. In addition, $\frac{V_t^{g,0}}{\partial b^g} = (1-e_{wt}^g)\left(u_w^g{}'(c_w^g(0)) + \frac{\partial W_{t+1}^{g,0}}{\partial b^g}\right)$, which combined with equation (2) yields $\frac{V_t^{g,0}}{\partial b^g} = -(1-e_{wt}^g)\left(\kappa_w^g{}''(e_{wt}^g)\frac{\partial e_{wt}^g}{\partial b^g}\right)$. Put together, we get that

$$Q_2^b = (1-\rho_0)(1-e_{w0}^g)\kappa_w^g{}''(e_{w0}^g)\left|\varepsilon(e_{w0}^g, b^g)\right|\frac{e_{w0}^g}{b^g}. \tag{8}$$

To complete the proof we need to calculate $\frac{db^g}{db^b}$. Total differentiation of the simplified budget constraint

$D^g\left(1-\hat{e}_w^g\right)b^g + D^b\left(1-\hat{e}_w^b\right)b^b = r$ with respect to b^b gives us

$$\frac{db^g}{db^b} = -\frac{b^g}{b^b}\epsilon(1-\hat{e}_w^g, b^b) - \frac{D^b\left(1-\hat{e}_w^b\right)}{D^g\left(1-\hat{e}_w^g\right)}\epsilon(1-\hat{e}_w^b, b^b) - \frac{D^b\left(1-\hat{e}_w^b\right)}{D^g\left(1-\hat{e}_w^g\right)}, \tag{9}$$

where $\epsilon(x,y) \equiv \frac{dx}{dy}\frac{y}{x}$. Plugging (7), (8), and (9) into (4), using a quadratic approximation of the effort function $\kappa^g(e_{wt}^g)$ around e_{w0}^g and assuming that the ratio $\theta^b(e_{wt}^b)$ is locally constant at $\hat{e}_{w0}^b$, we obtain the approximated formula for the normalized welfare gain $M_w(b^p) \equiv \frac{\frac{dJ_0(T,B)}{db^b}/\rho(1-\hat{e}_{w0}^b)}{\frac{\partial J_0(T,B)}{\partial b^g}/(1-\rho_0)(1-e_{w0}^g)}$ that is stated in the proposition, which completes the proof. ■

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

$$\frac{dJ_0(T,B)}{db^d} = Q_1^d + Q_2^d\frac{db^g}{db^d}, \tag{10}$$

where $Q_1^d \equiv \left(\rho_0\frac{\partial V_0^{b,0}}{\partial b^d} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial b^d}\right)$ and $Q_2^d \equiv \left(\rho_0\frac{\partial V_0^{b,0}}{\partial b^g} + (1-\rho_0)\frac{V_0^{g,0}}{\partial b^g}\right)$. We present the approximated formula in the following proposition.

Proposition A2.

Under a locally quadratic approximation of the effort function $\kappa_w^g(e_{wt}^g)$ around e_{w0}^g and assuming that the ratio $\theta^d(e_{wt}^d)$ is locally constant at $\hat{e}_{w0}^d$, the marginal benefit from raising b^d by \$1 is

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

with

1. $MB(b^d) \equiv L^d + M^d + S^d$, *where* $L^d \equiv \frac{\hat{e}_{w0}^d - e_{w0}^g}{e_{w0}^g}$, $M^d \equiv \left(\frac{\left|\varepsilon(\hat{e}_{w0}^d, b^d)\right|/b^d}{\left|\varepsilon(e_{w0}^g, b^g)\right|/b^g} - 1\right)\frac{\hat{e}_{w0}^d}{e_{w0}^g}$, $S^d \equiv \left(\theta^d - 1\right)\left(1 + L^d + M^d\right)$, $\varepsilon(x,y) \equiv \frac{\partial x}{\partial y}\frac{y}{x}$, $\theta^d \equiv \theta^d(\hat{e}_{w0}^d)$, e_{w0}^g *is w's participation rate at the beginning of the planning period, and* $\hat{e}_{w0}^d$ *is w's mean participation rate in households that transition to state d.*

2. $MC(b^d) \equiv \beta_0^d + \beta_1^d\epsilon(1-\hat{e}_w^g, b^d) + \beta_2^d\epsilon(1-\hat{e}_w^d, b^d)$, *where the coefficients β_0^d, β_1^d, and β_2^d are functions of the transition probabilities, average participation rates, and benefits and* $\epsilon(x,y) \equiv \frac{dx}{dy}\frac{y}{x}$.[3]

[3] Specifically, $\beta_0^d \equiv \frac{\sigma^d D^d\left(1-\hat{e}_w^d\right) - D^g\left(1-\hat{e}_w^g\right)}{D^g\left(1-\hat{e}_w^g\right)}$, $\beta_1^b \equiv \sigma^d\frac{b^g}{b^d}$, and $\beta_2^b \equiv \sigma^d\frac{D^d\left(1-\hat{e}_w^d\right)}{D^g\left(1-\hat{e}_w^g\right)}$, where $\sigma^d \equiv (1-p_0)\left(1-e_{w0}^g\right)/\lambda(1-\hat{e}_{w0}^d)$, $\lambda \equiv \sum_{i=0}^{T-1}\mu_i^d$, and μ_i^d is the probability of transitioning to state d in period i.

Proof.

We first find expressions for $\frac{\partial V_0^{b,0}}{\partial b^d}$ and $\frac{\partial V_0^{g,0}}{\partial b^d}$ in order to characterize Q_1^d. With $\frac{\partial V_t^{b,0}}{\partial b^d} = (1-e_{wt}^b)\left(\frac{\partial W_{t+1}^{b,0}}{\partial b^d}\right)$ and $\frac{\partial W_{t+1}^{b,0}}{\partial b^d} = (1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial b^d} + \lambda_{t+1}\frac{\partial V_{t+1}^{d,0}}{\partial b^d}$ we have that $\frac{\partial V_t^{b,0}}{\partial b^d} = (1-e_{wt}^b)\left((1-\lambda_{t+1})\frac{\partial V_{t+1}^{b,0}}{\partial b^d} + \lambda_{t+1}\frac{\partial V_{t+1}^{d,0}}{\partial b^d}\right)$. Working backwards from period $T-1$ to period 0 one can show that $\frac{\partial V_t^{b,0}}{\partial b^d} = \sum_{i=t+1}^{T-1}\prod_{j=t}^{i-1}(1-e_{wj}^b)\prod_{j=t+1}^{i-1}(1-\lambda_j)\lambda_j\frac{\partial V_i^{d,0}}{\partial b^d}$.

In state g we have $\frac{\partial V_t^{g,0}}{\partial b^d} = (1-e_{wt}^g)\left(\frac{\partial W_{t+1}^{g,0}}{\partial b^d}\right)$ and $\frac{\partial W_{t+1}^{g,0}}{\partial b^d} = \rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial b^d} + (1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial b^d}$, which imply that $\frac{\partial V_t^{g,0}}{\partial b^d} = (1-e_{wt}^g)\left(\rho_{t+1}\frac{\partial V_{t+1}^{b,0}}{\partial b^d} + (1-\rho_{t+1})\frac{\partial V_{t+1}^{g,0}}{\partial b^d}\right)$. Define the probability of transitioning to state d exactly at time i while w is unemployed by $\mu_i^{d,0}$ (which takes into account all the possible transition paths). Then, combining the results so far one can show by working backwards that $Q_1^d = \left(\rho_0\frac{\partial V_0^{b,0}}{\partial b^d} + (1-\rho_0)\frac{\partial V_0^{g,0}}{\partial b^d}\right) = \sum_{i=t}^{T-1}\mu_i^{d,0}E_{\mu_i^{d,u}}\left[\frac{\partial V_i^{d.0}}{\partial b^d}\right]$, where $E_{\mu_i^{d,0}}$ is the expectation operator conditional on arriving at period i with w unemployed and transitioning to state d then (taken over all possible paths).

Since $\frac{\partial V_t^{d.1}}{\partial b^d} = 0$ we have that $\frac{\partial V_t^{d.0}}{\partial b^d} = (1-e_{wt}^d)\left(u_w^{d\,\prime}(c_{wt}^d(0)) + \frac{\partial V_{t+1}^{d.0}}{\partial b^d}\right)$. Combined with (2) it can be expressed as $\frac{\partial V_t^{d.0}}{\partial b^d} = -(1-e_{wt}^d)\kappa_w^{d\,\prime\prime}(e_{wt}^d)\frac{de_{wt}^d}{db^d}$. Putting the terms together we obtain

$$Q_1^d = \sum_{i=t}^{T-1}\mu_i^{d,0}E_{\mu_i^{d,0}}\left[(1-e_{wi}^d)\kappa_{wi}^{d\,\prime\prime}(e_{wi}^d)\left|\varepsilon(e_{wi}^d,b^d)\right|\frac{e_{wi}^d}{b^d}\right]. \tag{11}$$

Define the probability of transitioning to state d in period i by μ_i^d and note that for those households that arrive at this period with w employed the change in participation is zero. Dividing the probabilities in (11) by the chance of transitioning to state d at some point, $\lambda \equiv \sum_{i=0}^{T-1}\mu_i^d$, we can rewrite Q_1^d as $Q_1^d = \lambda E_\lambda\left\{(1-\bar{e}_{w0}^d)\kappa_w^{d\,\prime\prime}(\bar{e}_{w0}^d)\left|\varepsilon(e_{wi}^d,b^d)\right|\frac{\bar{e}_{w0}^d}{b^d}\right\} \equiv \lambda E_\lambda(g(\bar{e}_{w0}^d))$, where $\bar{e}_{w0}^d$ denotes participation in the period the household transitions to state d and E_λ is the expectation operator conditional on being in state d. Expanding $g(e)$ around w's average participation upon the transition to state d – which we denote by $\hat{e}_{w0}^d$ – we can approximate Q_1^d by

$$Q_1^d \cong \lambda(1-\hat{e}_{w0}^d)\kappa_w^{d\,\prime\prime}(\hat{e}_{w0}^d)\left|\varepsilon(\hat{e}_{w0}^d,b^d)\right|\frac{\hat{e}_{w0}^d}{b^d}. \tag{12}$$

In addition, as in the proof of Proposition A1

$$Q_2^d = \left(\rho_0\frac{\partial V_0^{b,0}}{\partial b^g} + (1-\rho_0)\frac{V_0^{g,0}}{\partial b^g}\right) = (1-\rho_0)(1-e_{w0}^g)\kappa_w^{g\,\prime\prime}(e_{w0}^g)\left|\varepsilon(e_{w0}^g,b^g)\right|\frac{e_{w0}^g}{b^g}. \tag{13}$$

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

$$\frac{db^g}{db^d} = -\frac{b^g}{b^d}\epsilon(1-\hat{e}^g_w, b^d) - \frac{D^d\left(1-\hat{e}^d_w\right)}{D^g\left(1-\hat{e}^g_w\right)}\epsilon(1-\hat{e}^d_w, b^d) - \frac{D^d\left(1-\hat{e}^d_w\right)}{D^g\left(1-\hat{e}^g_w\right)}. \tag{14}$$

Plugging (12), (13), and (14) into (10), using a quadratic approximation of the effort function $\kappa^g(e^g_{wt})$ around e^g_{w0} and assuming that the ratio $\theta^d(e^d_{wt})$ is locally constant at $\hat{e}^d_{w0}$, we obtain the approximated formula for the normalized welfare gain $M_w(b^d) \equiv \frac{\frac{dJ_0(T,B)}{db^d}/\lambda(1-\hat{e}^d_{w0})}{\frac{\partial J_0(T,B)}{\partial b^g}/(1-\rho_0)(1-e^g_{w0})}$ that is stated in the proposition, which completes the proof. ■

Extension: Exogenous Separations

One natural extension of our search model is to allow for w's employment status to change at state transitions. For example, a working w is state g may want to decrease his or her labor supply in state b to take care of the ill h and may decide to quit his or her job and start searching for a job again in a year or two after the shock occurs. We can extend the model such that employment is only absorbing within each health state, but can be exogenously terminated at rate δ_t at health-state transitions. To demonstrate how to include this sort of separation, let us reconsider the value of entering period t in state g when w is unemployed. In this case, the household's value function is

$$V^{g,0}_t(B,T,A_t) \equiv \max\left\{\begin{array}{c} e^g_{wt}\left(u^g_h(c^g_{ht}(1)) + u^g_w(c^g_{wt}(1)) + W^{g,1}_{t+1}(B,T,A_{t+1}(1))\right) \\ +(1-e^g_{wt})\left(u^g_h(c^g_h(0)) + u^g_w(c^g_w(0)) + W^{g,0}_{t+1}(B,T,A_{t+1}(0))\right) - \kappa^g_w(e^g_{wt}) \end{array}\right\},$$

where as before

$$W^{g,0}_{t+1}(B,T,A_{t+1}) \equiv (1-\rho_{t+1})V^{g,0}_{t+1}(B,T,A_{t+1}) + \rho_{t+1}V^{b,0}_{t+1}(B,T,A_{t+1}),$$

but with the adjustment that now

$$W^{g,1}_{t+1}(B,T,A_{t+1}) \equiv \rho_{t+1}\left((1-\delta_{t+1})V^{b,1}_{t+1}(B,T,A_{t+1}) + \delta_{t+1}V^{b,0}_{t+1}(B,T,A_{t+1})\right) + (1-\rho_{t+1})V^{g,1}_{t+1}(B,T,A_{t+1}).$$

That is, if h becomes sick when w works, there is a probability of δ_{t+1} that w stops working and then renews his or her search effort. In this case, it is no longer true that $\frac{\partial W^{g,1}_{t+1}}{\partial b^b} = 0$, but rather $\frac{\partial W^{g,1}_{t+1}}{\partial b^b} = \rho_{t+1}\delta_{t+1}\frac{\partial V^{b,0}_{t+1}}{\partial b^b}$. In turn, this implies that in equation (5) one needs to take into account additional paths to reach period i with w unemployed and transition to state b exactly in that period. It is no longer merely the probability of becoming exactly sick in period i and staying unemployed until that period. Rather, it is also the probability of being employed before period i and then transitioning into state b and becoming unemployed in that period (with probability δ_t). However, recall that our final formulas include expected values and

averages. Before, those who were employed contributed a value of zero to the summations. But, now, with a positive probability they contribute a non-zero value (because a fraction δ_t responds on the effort margin as for them employment is not absorbing). Therefore, our formulas remain the same under this extension such that welfare is still identified by the means stated in our formulas. The change is that conceptually these means include additional individuals that respond. The sample moments that one needs to calculate to recover welfare remain unchanged.

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

In this appendix we present a baseline static model that is the intensive-margin counterpart to the participation model in the text. The analysis of the dynamic version of this model follows the logic of the analysis in Online Appendix A and is available from the authors on request. The general conclusion of the dynamic model in the intensive-margin case is similar to that in the extensive-margin case – the labor supply responses that identify the marginal benefits from social insurance are replaced by average labor supply responses. For completeness, we describe the full setup of the model although it has close similarities to the model of Section II.A in the text.

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

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

Household Preferences. We follow the collective approach to household behavior and assume that household decisions are Pareto efficient and can be characterized as solutions to the maximization of $\beta_w U_w(c_w^s, l_w^s) + \beta_h U_h(c_h^s, l_h^s)$, where β_w and β_h are the Pareto weights on w and h, respectively. For simplicity, we assume equal Pareto weights ($\beta_w = \beta_h = 1$), which is without loss of generality as long as the spouses' relative bargaining power is stable across states of nature.

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

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

$$V^s(B,T,A) \equiv \max_{c_i^s, l_i^s} U_h(c_h^s, l_h^s) + U_w(c_w^s, l_w^s)$$

$$\text{s.t. } c_h^s + c_w^s = A + w_h^s\left(1-\tau_h^s\right) l_h^s + w_w\left(1-\tau_w^s\right) l_w^s + B^s,$$

where A is the household's wealth, w_h^s is h's wage rate in state s and w_w is w's wage rate. The household's first-order conditions imply that $\frac{\partial U_h}{\partial c_h^s} = \frac{\partial U_w}{\partial c_w^s} = -\frac{\partial U_w}{\partial l_w^s}\frac{1}{w_w(1-\tau_w^s)}$. Importantly, note that we allow h to be at a corner solution in state b – that is, $l_h^b = 0$ – and use only w's labor supply first-order conditions.

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

$$\max_{B,T} J(B,T) \quad \text{s.t. } \mu^b B = \mu^g(\tau_h^g w_h^g l_h^g + \tau_w^g w_w l_w^g) + \mu^b(\tau_h^b w_h^b l_h^b + \tau_w^b w_w l_w^b). \tag{15}$$

Optimal Social Insurance

Consider a \$1 increase in B financed by an appropriate increase in taxes, e.g., through τ_h^g. To simplify notation we assume that $\tau_w^g = \tau_h^b = \tau_w^b = 0$, which allows us to obtain concise welfare formulas.[4] The welfare gain from this perturbation is $\frac{dJ(B,T)}{dB} = \mu^b \frac{\partial V^b}{\partial B} + \mu^g \frac{\partial V^g}{\partial \tau_h^g}\frac{d\tau_h^g}{dB}$, which we normalize by the welfare gain from raising h's net-of-tax labor income in state g by \$1 (scaled by the targeted population) to gain a cardinal interpretation.[5] Exploiting the envelope theorem (in the differentiation of the household's value functions) and using the household's first-order conditions, we obtain $\frac{\partial V^g}{\partial \tau_h^g} = -w_h^g l_h^g \frac{\partial U_w}{\partial c_w^g}$ and $\frac{\partial V^b}{\partial B} = \frac{\partial U_w}{\partial c_w^b}$. Differentiating the budget constraint with respect to B we get $\frac{d\tau_h^g}{dB} = \frac{\mu^b}{\mu^g z_h^g}\left(1 + \frac{\varepsilon(z_h^g, 1-\tau_h^g)\frac{\tau_h^g}{1-\tau_h^g}}{1-\varepsilon(z_h^g,1-\tau_h^g)\frac{\tau_h^g}{1-\tau_h^g}}\right)$, where $z_h^g \equiv w_h^g l_h^g$ is h's taxable income and $\varepsilon(z_h^g, 1-\tau_h^g) \equiv \frac{\partial z_h^g}{\partial(1-\tau_h^g)}\frac{1-\tau_h^g}{z_h^g}$ is the commonly estimated net-of-tax taxable income elasticity. Put together, it follows that the normalized welfare gain from a marginal increase in B is $M_W(B) = MB(B) - MC(B)$, where $MB(B) \equiv \frac{\frac{\partial U_w}{\partial c_w^b} - \frac{\partial U_w}{\partial c_w^g}}{\frac{\partial U_w}{\partial c_w^g}}$ and $MC(B) \equiv \frac{\varepsilon(z_h^g,1-\tau_h^g)\frac{\tau_h^g}{1-\tau_h^g}}{1-\varepsilon(z_h^g,1-\tau_h^g)\frac{\tau_h^g}{1-\tau_h^g}}$.

Identifying the Benefits of Social Insurance. The identification of the gap in marginal utilities of consumption using the unaffected spouse's labor supply responses in the intensive-margin model is summarized

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

[5] The formula for the normalized gain is $M_W(B) \equiv \frac{\frac{dJ(B,T)}{dB}/\mu^b}{\frac{\partial J(B,T)}{\partial z_h^g(1-\tau_h^g)}/\mu^g}$, where $z_h^g \equiv w_h^g l_h^g$.

in the following proposition.

Proposition B1. *Assuming consumption-leisure separability,*[6] *the marginal benefit from raising B in \$1 is approximately*

$$MB(B) \cong L^b + M^b, \tag{16}$$

where $L^b \equiv \frac{l_w^b - l_w^g}{l_w^g}$, $M^b \equiv (\varphi - 1)\frac{l_w^b - l_w^g}{l_w^g}$, *and* $\varphi \equiv \frac{\partial^2 U_w / \partial (l_w^g)^2}{\partial U_w / \partial l_w^g} l_w^g$.

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

Identification of φ

In this section we derive a relationship between φ and observable labor supply elasticities. The analysis uses a similar strategy as that introduced by Chetty (2006b) to recover risk aversion – i.e., we recover the curvature of the labor disutility function in the same way that Chetty (2006b) recovers the curvature of the consumption utility function. The intuition for the method is that the extent to which an individual responds to changes in economic incentives (wages and income) is directly linked to the rate at which preferences change (over consumption or labor). To conduct the analysis at the individual level, we use the "sharing-rule" interpretation of the collective model as defined by Chiappori (1992). That is, we assume that non-labor income in state s, denoted by y^s, is shared between the members such that $y_w^s \equiv \pi_w^s(w_w, w_h^s, A)$ is the amount received by w and $y_h^s \equiv y^s - \pi_w^s(w_w, w_h^s, A)$ is the amount received by h. With these definitions, one can write w's program in state g as

$$\max_{c_w^g, l_w^g} U_w(c_w^g, l_w^g)$$

$$\text{s.t. } c_w^g = y_w^g + w_w l_w^g.$$

Since we are focusing on the analysis of spouse w in state g, we drop spouse subscripts and state superscripts for convenience.

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

The first-order conditions of this program imply that $wU_c(y+wl,l) = -U_l(y+wl,l)$, where U_x denotes the partial derivative of U with respect to x. Partially differentiating the latter equation with respect to y and w yields $\frac{\partial l}{\partial y} = -\frac{wU_{cc}+U_{cl}}{w^2U_{cc}+U_{ll}+2wU_{cl}}$ and $\frac{\partial l}{\partial w} = -\frac{U_c+wlU_{cc}+lU_{cl}}{w^2U_{cc}+U_{ll}+2wU_{cl}}$.[7] It follows that $\varphi \equiv \frac{U_{ll}}{U_l}l = \frac{1+\varepsilon(l,y)\frac{wl}{y}}{\varepsilon^c(l,w)} + \varepsilon(U_c,l)$, where $\varepsilon(l,y) \equiv \frac{\partial l}{\partial y}\frac{y}{l}$, $\varepsilon(l,w) \equiv \frac{\partial l}{\partial w}\frac{w}{l}$, $\varepsilon(U_c,l) \equiv \frac{U_{cl}}{U_c}l$, and $\varepsilon^c(l,w) \equiv \varepsilon_{l,w} - \varepsilon_{l,y}\frac{wl}{y}$. With consumption-leisure separability the formula reduces to $\varphi = \frac{1+\varepsilon(l,y)\frac{wl}{y}}{\varepsilon_{l,w}-\varepsilon_{l,y}\frac{wl}{y}}$.

Online Appendix C: Calibrating Labor Disutility State Dependence for Survivors

Calibration Method. Consider an insurance environment that fully compensates for consumption losses imposed by income shocks. Normally, this would imply fully compensating for the household's pre-shock level of income. However, in our case, as the composition of the household changes, we need to ask how much income does a surviving spouse need while single to achieve the same level of consumption utility that he or she enjoyed before the shock? The classic answer to this question is the adult "equivalence scale" (see, e.g., Blundell and Lewbel 1991), which is commonly assumed to lie within the interval (0.5,1). It is less than 1 since the household becomes a one-person household and is more than 0.5 due to economies of scale in consumption within a two-person household. We denote the equivalence scale by r^0. A direct implication of its definition is that in the absence of labor disutility state dependence, when $\theta^b = 1$, the labor force participation of the surviving spouses would not change across states of nature if they receive r^0 of their pre-shock household income.

Next, denote the replacement rate that surviving spouses receive in equilibria in which their labor supply does not change after experiencing the shock by r^{eq}. This is the level of compensation they are implicitly willing to accept so that their labor supply would remain the same following the shock.

We show below that the comparison of the two replacement rates, r^0 and r^{eq}, can reveal the degree of state dependence. Intuitively, if $r^{eq} < r^0$, survivors are willing to accept less than full compensation for their consumption after the shock to avoid self-insuring through labor supply, which implies an increase in its utility cost, $\theta^b > 1$. That is, incomplete adjustment of post-shock consumption (captured by r^{eq}) to the consumption level which achieves the pre-shock level of utility (captured by r^0) implies that labor disutility has increased. If $r^{eq} \cong r^0$, then state dependence on average is likely to be negligible. To formalize this intuition we use the following lemma. We denote the household's "consumption value function" for the generalized state-dependent preferences of Section II.B in the main text by $V^s(y^s(l_w^s)) \equiv \max\ u_w^s(c_w^s) + u_h^s(c_h^s)$ s.t. $c_w^s + c_w^s = y^s(l_w^s)$.

[7] Note the subtlety that we focus on *partial* derivatives of the unaffected spouse's behavior with respect to y and w. In particular, y is held fixed when we change w.

Lemma C1. *Let $V^s(y^s(1))$ denote the household's consumption value function when w works in state s, $\theta^u \equiv \frac{V^{b\,\prime}(y^b(1))}{V^{g\,\prime}(y^b(1))}$ denote the change in the marginal value of household income, and $\gamma \equiv -\frac{V^{g\,\prime\prime}(y^g(1))}{V^{g\,\prime}(y^g(1))} \times y^g(1)$ denote the household-level pre-shock relative risk aversion. Then, in equilibria in which w's labor supply is the same in state g and state b the following holds*

$$\theta^u(1+\gamma(1-r^{eq})) \cong \theta^b, \tag{17}$$

where $r^{eq} \equiv y^b(1)/y^g(1)$ is the steady state replacement rate that satisfies this relationship.

Proof. The proof relies on the necessary relationship between household income streams across states of nature in equilibria where labor supply remains unchanged such that $\bar{v}_w^g = \bar{v}_w^b$, where $\bar{v}_w^s \equiv \frac{1}{\theta^s}[V^s(y^s(1)) - V^s(y^s(0))]$. See Online Appendix D for details.

The relationship in (17) has a simple intuition: if labor supply is unchanged when the shock occurs, then the change in the cost of labor must equal the change in the marginal utility from income. The right-hand side of the equation captures the change in the marginal entrant's labor disutility by the definition of θ^b. The left-hand side evaluates the marginal utility from income in the new state. It is the baseline pre-shock marginal utility from income (normalized to one), augmented by the change in the marginal utility due to income changes, $1-r^{eq}$, and the curvature of the consumption value function $V^s(y^s(1))$, γ (which is captured by the term $\gamma(1-r^{eq})$). Then, we multiply the resulting expression by the change in the marginal value of household income across states, θ^u.

When r^{eq} is directly observed (i.e., revealed by individuals' choices) – as in the case of widowers who do not change their mean participation rate when their wife dies – we can recover θ^b with two simple steps, which correspond to the intuitive explanation above. First, since (17) is satisfied when $\theta^b = 1$ and $r^{eq} = r^0$, we can recover θ^u by $\theta^u = 1/(1+\gamma(1-r^0))$, if we borrow estimates for r^0 from the literature as we discuss below. Second, we can use θ^u and the observed r^{eq} to recover θ^b by using (17) again, such that $\theta^b \cong \frac{1+\gamma(1-r^{eq})}{1+\gamma(1-r^0)} = 1 + \frac{\gamma(r^0-r^{eq})}{1+\gamma(1-r^0)}$. This formalizes our intuition: whenever $r^{eq} < r^0$ – that is, whenever survivors avoid self-insurance through labor supply by willing to receive less than the compensation that gets them back to their level of consumption utility before the shock – it follows that self-insurance became more costly when the shock occurred, $\theta^b > 1$.

When r^{eq} is not directly observed by choices – e.g., when participation increases in response to a shock as in the case of widows – we can use the equilibrium responses to construct a bound on θ^b with additional identifying assumptions, such as monotonicity as defined below.

Assumption (monotonicity). *Consider a sub-set of individuals $\mathcal{I}$. For every $i \in \mathcal{I}$ define the potential outcome $Y_i(0)$ to be i's participation decision that would be realized were he or she not to experience a shock and $Y_i(1)$ to be i's participation decision that would be realized if he or she were to experience a shock. If $Y_i(1) \geq Y_i(0)$ for every $i \in \mathcal{I}$, we say that monotonicity is satisfied for $\mathcal{I}$.*

Under monotonicity for, say, widows (or widows who experience large income losses), the mean increase in spousal participation is driven by individuals who switch from working to not working ("compliers"), while the remaining spouses either keep working ("always-takers") or stay out of the labor force ("never-takers"). Given this response, we observe an aggregate income replacement rate in the data, denoted by r', which is composed of the rate among compliers, denoted by r'_c, and a replacement rate for the rest of the sample. Now, assume that we change the environment only by offering the compliers a higher income if they do not work. Since working is costly, there must exist $r''_c < r'_c$ such that compliers prefer receiving r''_c without working to receiving r'_c and working. Therefore, under monotonicity, in an equilibrium in which the mean participation rate does not change when the shock occurs, r^{eq} (which involves r''_c) must be smaller than the r' that we actually observe (which involves r'_c). This imposes a lower bound on θ^b such that $\theta^b \cong 1 + \frac{\gamma(r^0 - r^{eq})}{1+\gamma(1-r^0)} \geq 1 + \frac{\gamma(r^0 - r')}{1+\gamma(1-r^0)}$.

Calibration Results. We begin by studying the implications of a commonly used equivalence scale – the modified OECD equivalence scale which implies $r^0 = 0.67$. Other widely used adult equivalence scales lead to similar conclusions.[8] Combining widowers and widows in Panel B of Figure IV yields an average post-shock replacement rate of $r' = 0.665$. Given the increase in mean labor force participation and using the bound we derived above, these estimates imply that $\theta^b \geq 1$ and suggest that state dependence is negligible.

Next, we consider model-based estimates for adult equivalence scales. In particular, we use recent estimates from Browning, Chiappori, and Lewbel (2013), which offer separate estimates for "indifference scales" for men and women.[9] Since widowers do not change their mean participation rate when their wives die, we can directly observe their r^{eq}. Recall from Panel B of Figure IV that widowers experience an actual loss of 31% in household income and hence for them $r^{eq} = 1 - 0.31 = 0.69$. This implies that they are willing to accept 69% of their pre-shock level of household income to avoid increasing their labor supply. Browning, Chiappori, and Lewbel (2013) find that in households with equal sharing of income among the two spouses,

[8] For example, the square-root scale which implies $r^0 = 0.71$ (see, e.g., Cutler and Katz 1992; OECD 2011). Note that the implicit equivalence scale in the Danish Social DI is approximately 0.65 and is 0.66 in the Old-Age Pension. See Section III in the main text for institutional details.

[9] Their notion of "indifference scales" is an individual-based version of equivalence scales, which aims at identifying the fraction of the household's income a member would need in order to buy a bundle of privately consumed goods at market prices that put him or her on the same indifference curve over goods that he or she attained as a member of the household. Their method relies on recovering the consumption demand functions of individuals within a household based on a collective household model, which they estimate by using the Canadian Survey of Family Expenditures.

the indifference scale for males is about 0.80. This suggests that for widowers $r^0 = 0.80 > r^{eq} = 0.69$ and thus on average their labor disutility increases when they lose their wives – that is, $\theta^b \cong 1 + \frac{0.11\gamma}{1+0.20\gamma} > 1$. For widows, who increase their labor force participation, we can recover a bound for state dependence using Browning, Chiappori, and Lewbel's (2013) indifference ratio of 0.72. Recall from Panel B of Figure IV that for widows $r' = 1 - 0.35 = 0.65$. This implies a lower bound of $\theta^b \geq 1 + \frac{0.07\gamma}{1+0.28\gamma} > 1$, which suggests that on average labor disutility likewise increases for widows when they lose their husbands.

Online Appendix D: Proof of Lemma C1

In this section we provide a proof for Lemma C1 in Online Appendix C. We begin with the baseline model and then provide a proof for the dynamic model. Similar analysis can be conducted for the intensive-margin case and is available from the authors on request.

Static Extensive Margin Model

Recall that $V^s(y^s(l_w^s)) \equiv \max\ u_w^s(c_w^s) + u_h^s(c_h^s)$ s.t. $c_w^s + c_w^s = y^s(l_w^s)$, where $y^s(l_w^s) \equiv A + \bar{z}_h^s \times l_h^s + \bar{z}_w \times l_w^s + B^s(l_w^s)$. Since we are interested in analyzing steady-state equivalence scales we account for transitory labor income shocks and employ conditions under which the scales that we study are not sensitive to these shocks. We decompose w's net labor income, $\bar{z}_w$, into its permanent component, $\tilde{z}_w$, and its transitory component, ς_w, such that $\bar{z}_w = \tilde{z}_w + \varsigma_w$ and $y^s(l_w^s) = A + \bar{z}_h^s \times l_h^s + (\tilde{z}_w + \varsigma_w) \times l_w^s + B^s(l_w^s)$.

Next, recall that w works when $v_w^s < \bar{v}_w^s \equiv V^s(y^s(1)) - V^s(y^s(0))$, where $v_w^g = v_w$ and $v_w^b = \theta^b \times v_w^g$. In equilibria in which w's participation rate in state g and in state b are the same it must be that $\bar{v}_w^g = \bar{v}_w^b/\theta^b$, or: $V^g(y^g(1)) - V^g(y^g(0)) = \frac{1}{\theta^b}\left\{V^b(y^b(1)) - V^b(y^b(0))\right\}$. This implies a necessary condition that the household income flows – $y^g(0)$, $y^g(1)$, $y^b(0)$, and $y^b(1)$ – must satisfy when labor supply is unchanged across states of nature. In a steady state, this equality is not sensitive to local income shocks. By equating the derivative of both sides with respect to the transitory income shock, ς_w, we get the relationship

$$V^{g\,\prime}(y^g(1)) = \frac{1}{\theta^s}\left\{V^{b\,\prime}(y^b(1))\right\}. \tag{18}$$

Let $\theta^u \equiv V^{b\,\prime}(y^b(1))/V^{g\,\prime}(y^b(1))$ denote the change in the marginal value of household income, and let $\gamma \equiv -[V^{g\,\prime\prime}(y^g(1))/V^{g\,\prime}(y^g(1))] \times y^g(1)$ denote the household-level pre-shock relative risk aversion. A second-order expansion of the value function V^b on the right-hand side of (18) around $y^g(1)$ yields the result in Lemma C1

$$\theta^u(1 + \gamma(1 - r^{eq})) \cong \theta^b, \tag{19}$$

where $r^{eq} \equiv y^b(1)/y^g(1)$ is the steady state replacement rate that satisfies this relationship.

Dynamic Search Model

The notation and definitions we use here are described in Online Appendix A. To simplify the analysis we assume two states of nature as in the baseline model, $s \in \{g, b\}$. Recall from Online Appendix A that $c^s_{ht}(l^s_{wt}) + c^s_{wt}(l^s_{wt}) + A_{t+1}(l^s_{wt}) = A_t + y^s_t(l^s_{wt})$ and $y^s_t(l^s_{wt}) \equiv \bar{z}^s_{ht} \times l^s_{ht} + \bar{z}_{wt} \times l^s_{wt} + B^s(l^s_{wt})$. As in the baseline case, we decompose w's net labor income, $\bar{z}_{wt}$, into its permanent component, $\tilde{z}_{wt}$, and its transitory component, ς_{wt}, such that $\bar{z}_{wt} = \tilde{z}_{wt} + \varsigma_{wt}$ and $y^s_t(l^s_{wt}) = \bar{z}^s_{ht} \times l^s_{ht} + (\tilde{z}_{wt} + \varsigma_{wt}) \times l^s_{wt} + B^s(l^s_{wt})$. For each period in which w is not working define the flow consumption utility at the optimal choices as a function of the period's wealth and income by

$$U^s(A_t, y^s_t(l^s_{wt})) \equiv u^s_h(c^{s*}_{ht}(l^s_{wt})) + u^s_w(c^{s*}_{wt}(l^s_{wt})),$$

where

$$(e^{s*}_{wt}, c^{s*}_{ht}(l^s_{wt}), c^{s*}_{wt}(l^s_{wt}), A^*_{t+1}(l^s_{wt})) \equiv \underset{e^s_{wt}, c^s_{ht}, c^s_{wt}, A_{t+1}}{\arg\max} \left\{ \begin{array}{ll} e^s_{wt}\left(u^s_h(c^s_{ht}(1)) + u^s_w(c^s_{wt}(1)) + W^{s,1}_{t+1}(B,T,A_{t+1}(1))\right) + & \\ (1-e^s_{wt})\left(u^s_h(c^s_{ht}(0)) + u^s_w(c^s_{wt}(0)) + W^{s,0}_{t+1}(B,T,A_{t+1}(0))\right) & -\kappa^s_w(e^s_{wt}) \end{array} \right\}.$$

We can, therefore, rewrite the first-order condition for w's effort as

$$\left(U^s(A_t, y^s_t(1)) + W^{s,1}_{t+1}(B,T,A^*_{t+1}(1)\right) - \left(U^s(A_t, y^s_t(0)) + W^{s,0}_{t+1}(B,T,A^*_{t+1}(0))\right) = \kappa^{s\,\prime}_w(e^{s*}_{wt}). \tag{20}$$

In equilibria in which w's participation rate in state g and state b are the same it must be that $e^{g*}_{wt} = e^{b*}_{wt}$. For a given period, which we normalize to 0, define $\theta^b \equiv \kappa^{b\,\prime}(e^{b*}_{w0})/\kappa^{g\,\prime}(e^{b*}_{w0})$, which implies that

$$\begin{array}{c} \left(U^g(A_0, y^g_0(1)) + W^{g,1}_1(B,T,A^*_1(1)\right) - \left(U^g(A_0, y^g_0(0)) + W^{g,0}_1(B,T,A^*_1(0))\right) = \\ \frac{1}{\theta^b}\left\{\left(U^b(A_0, y^b_0(1)) + W^{b,1}_1(B,T,A^*_1(1)\right) - \left(U^b(A_0, y^s_0(0)) + W^{b,0}_1(B,T,A^*_1(0))\right)\right\}. \end{array} \tag{21}$$

Differentiating both sides with respect to the transitory shock ς_{w0} yields $U^g_y(A_0, y^g_0(1)) = \frac{1}{\theta^b} U^b_y(A_0, y^b_0(1))$, where U^s_x is the partial derivative of U^s with respect to x. Let $\theta^u \equiv U^b_y(A_0, y^b_0(1))/U^g_y(A_0, y^b_0(1))$ denote the change in the marginal value of household income, and let $\gamma \equiv -[U^g_{yy}(A_0, y^g_0(1))/U^g_y(A_0, y^g_0(1))] \times y^g_0(1)$ denote the household-level pre-shock relative risk aversion. A second-order expansion of the consumption

flow "value function" U^b around $y_0^g(1)$ yields the result in Lemma C1

$$\theta^u(1+\gamma(1-r^{eq})) \cong \theta^b, \tag{22}$$

where $r^{eq} \equiv y_0^b(1)/y_0^g(1)$ is the steady state replacement rate that satisfies this relationship.

Online Appendix E: Heterogeneity in θ^b

In this section we return to our generalized participation model of Sections II.B and VI.B in the text and provide an approximated formula for the case in which the labor disutility state dependence is heterogeneous.

Denote the joint distribution of the vector of w's labor disutility and labor disutility state dependence, (v_w, θ^b), by $\Gamma(v_w, \theta^b)$, the marginal distribution of θ^b by $K(\theta^b)$, and the marginal distribution of v_w by $F(v_w)$ as before. In addition, denote the distribution of v_w conditional on θ^b by $F_{\theta^b}(v_w)$ and the corresponding probability density function by $f_{\theta^b}(v_w)$. Define $y^s \equiv \theta^s v_w$ (where $\theta^g = 1$ by normalization) and denote its distribution by $G^s(y^s)$ for $s \in \{g, b\}$ with a probability density function $g^s(y^s)$. Using this notation, w works in state s whenever $y^s < \bar{y}^s$ where

$$\bar{y}^s \equiv \left[u_h^s(c_h^s(1)) + u_w^s(c_w^s(1))\right] - \left[u_h^s(c_h^s(0)) + u_w^s(c_w^s(0))\right].$$

It follows that we can rewrite the marginal benefit in labor disutility terms as $MB(b^b) = \frac{\left|\frac{\partial \bar{y}^b}{\partial b^b}\right| - \left|\frac{\partial \bar{y}^g}{\partial b^g}\right|}{\left|\frac{\partial \bar{y}^g}{\partial b^g}\right|}$. Define participation by $e_w^s \equiv G^s(\bar{y}^s)$ and note that $\frac{\partial e_w^b}{\partial b^b} = g^b(\bar{y}^b)\frac{\partial \bar{y}^b}{\partial b^b}$. To continue, we would want to express $g^b(\bar{y}^b)$ in terms of the marginal distribution of v_w. Since $G^b(y^b) = \int_0^\infty k(\theta^b)\left[\int_0^{y^b/\theta^b} f_{\theta^b}(v_w)dv_w\right]d\theta^b$ we have that $g^b(y^b) = \int_0^\infty \left[f_{\theta^b}(\frac{y^\sigma}{\theta^b})\frac{1}{\theta^b}\right]k(\theta^b)d\theta^b \equiv E_{\theta^b}\left[f_{\theta^b}(\frac{y^\sigma}{\theta^b})\frac{1}{\theta^b}\right]$, where E_{θ^b} is the expectation operator with respect to θ^b. Next, consider approximating $g^b(\bar{y}^b)$. Define $\mu(\theta^b) \equiv f_{\theta^b}(\frac{\bar{y}^b}{\theta^b})\frac{1}{\theta^b}$ and take a first-order Taylor expansion around the mean value of θ^b, $E\theta^b$, to get $\mu(\theta^b) \cong \mu(E\theta^b) + \mu'(E\theta^b)(\theta^b - E\theta^b)$. Hence, to a first approximation $g^b(\bar{y}^b) = E_{\theta^b}\left[\mu(\theta^b)\right] \cong \mu(E\theta^b) = f_{E\theta^b}(\frac{\bar{y}^b}{E\theta^b})\frac{1}{E\theta^b}$. Define $\bar{v}^b$ to be the value of v_w which satisfies $v_w E\theta^b = \bar{y}^b$. This implies that $g^b(\bar{y}^b) \cong \frac{1}{E\theta^b}f_{E\theta^b}(\bar{v}^b)$ and hence that $\frac{\partial e_w^b}{\partial b^b} = g^b(\bar{y}^b)\frac{\partial \bar{y}^b}{\partial b^b} \cong \frac{1}{E\theta^b}f_{E\theta^b}(\bar{v}^b)\frac{\partial \bar{y}^b}{\partial b^b}$. If, for example, v_w is distributed independently of θ^b, such that $f_{E\theta^b}(\bar{v}^b) = f(\bar{v}^b)$, a first-order approximation of F in the threshold region $(\bar{v}_w^g, \bar{v}_w^b)$ will yield the same approximated formula for $MB(b^b)$ as in Proposition 2 in the main text where θ^b is replaced by its mean value, $E\theta^b$.

Online Appendix F: Implications for Health-State Dependence of the Household's Preferences

In this section we formalize the discussion in Section V.B in the text on health-state dependence. Since we found the unaffected spouse's labor supply response to spousal health shocks to be on the intensive margin, we refer to the intensive-margin model of the household behavior developed in Online Appendix B. We generalize preferences such that each spouse's preferences in state s can be represented by the utility function $U_i^s(c_i^s, l_i^s)$, where c_i^s and l_i^s are spouse i's consumption and labor supply in state s, respectively. Efficiency requires the marginal utility of h's consumption, $\frac{\partial U_h^s}{\partial c_h^s}$, to equal w's marginal disutility of labor, $-\frac{\partial U_w^s}{\partial l_w^s}\frac{1}{w_w}$. This is the basic logic behind the welfare result for the intensive margin case, which implies that $\left(\frac{\partial U_h^b}{\partial c_i^b} - \frac{\partial U_h^g}{\partial c_i^g}\right)/\frac{\partial U_h^g}{\partial c_i^g} = \left(\frac{\partial U_w^b}{\partial l_w^b} - \frac{\partial U_w^g}{\partial l_w^g}\right)/\frac{\partial U_w^g}{\partial l_w^g}$. Define $\theta^u \equiv \frac{\partial U_h^b}{\partial c_h^b}/\frac{\partial U_h^g}{\partial c_h^g}$ at c_h^b and $\theta^b \equiv \frac{\partial U_w^b}{\partial l_w^b}/\frac{\partial U_w^g}{\partial l_w^g}$ at l_w^b to be the local consumption utility and labor disutility state dependence parameters, respectively. With consumption-leisure separability it follows that $\theta^u\gamma\frac{\triangle c_h}{c_h} + \theta^b\varphi\frac{\triangle l_w}{l_w} \cong \theta^b - \theta^u$, where $\gamma \equiv -\frac{\partial^2 U_h^g/\partial\left(c_h^g\right)^2}{\partial U_h^g/\partial c_h^g}c_h^g$ is h's risk aversion parameter, $\varphi \equiv \frac{\partial^2 U_w^g/\partial(l_w^g)^2}{\partial U_w^g/\partial l_w^g}l_w^g$ is the curvature of w's disutility from labor, and $\frac{\triangle x}{x} \equiv \frac{x^g - x^b}{x^g}$.[10] Since we find that $\frac{\triangle l_w}{l_w} > 0$, since $\theta^u, \theta^b, \gamma, \varphi > 0$, and if $\frac{\triangle c_h}{c_h} > 0$ due to the small income loss the household experiences, it must be that $0 < \theta^u\gamma\frac{\triangle c_h}{c_h} + \theta^b\varphi\frac{\triangle l_w}{l_w} \cong \theta^b - \theta^u$. This implies that $\frac{\theta^b}{\theta^u} > 1$, which includes the extreme cases of $\theta^u = 1$ with $\theta^b > 1$ and $\theta^b = 1$ with $\theta^u < 1$. More generally, our results imply that labor disutility state dependence is greater than the potential state dependence in the sick spouse's consumption utility.

Online Appendix G: An Empirical Model of Labor Force Participation

In this section we estimate for spousal mortality shocks an empirical counterpart to the theoretical model of household labor force participation in order to provide suggestive estimates for $\varepsilon(e_w^b, b^b)$ and $\varepsilon(e_w^g, b^g)$. We model w's participation such that in the years before the event his or her decision is conditional on h's behavior. Specifically, the income h contributes to the household – whether through transfers or through labor income – is perceived as non-labor income in w's decision making.[11] We constrain the sample to individuals who are younger than 60 to avoid retirement transitions that are due to eligibility for early retirement benefits and Social Security.

Labor Force Participation. We let w's labor supply depend on his or her potential wage if deciding to

[10] This is achieved by taking a Taylor expansion of $\theta^u\frac{\partial U_h^g}{\partial c_h^g}$ around c_h^g and of $\theta^b\frac{\partial U_w^g}{\partial l_w^g}$ around l_w^g.

[11] This is a common practice in the empirical literature on married women's labor force participation (see, e.g., a review in Keane, Todd, and Wolpin 2011) and is in-line with the sharing-rule representation of the collective model (in Chiappori 1992).

work, on the potential transfers he or she would receive if deciding not to work, as well as on his or her unearned income. Denote the participation decision and the latent index of spouse w in household i at time t in state s by $l^s_{w,i,t}$ and $I^s_{w,i,t}$, respectively. Then, $l^s_{w,i,t} = 1$ if $I^s_{w,i,t} > 0$ and $l^s_{w,i,t} = 0$ otherwise. We assume the following linear form for the participation latent index

$$I^s_{w,i,t} = \delta_0 + \delta_1 z^s_{w,i,t} + \delta_2 b^s_{w,i,t} + \delta_3 y^s_{i,t} + \delta_4 wealth_{i,t} + controls + \varepsilon^s_{i,t}, \tag{23}$$

where

$$\delta_0 = \delta_{00} + \delta_{01} treat_i + \delta_{02} post_{i,t} + \delta_{03} treat_i \times post_{i,t},$$

$$\delta_k = \delta_{k0} + \delta_{k1} treat_i \times post_{i,t} \;\; k = 1, ..., 4.$$

In this specification $z^s_{w,i,t}$ denotes w's potential labor income in state s, $b^s_{w,i,t}$ denotes his or her potential government transfers if deciding not to work in state s, $y^s_{i,t}$ denotes w's unearned income as well as any income (earned or unearned) that is attributed to h before his or her death, and $wealth_{i,t}$ denotes the household's net wealth. The coefficients are allowed to freely change across states of nature, since $treat_i \times post_{i,t}$ is the differences-in-differences interaction variable. The controls include dummies for w's age, calendar year, and municipality of residence before the shock occurs.

Wage Equations. Following Blundell et al. (2007), we take the standard human capital approach to wages and additionally allow for the relative prices of education to change over time. In particular, we assume[12]

$$\begin{aligned} z^s_{w,i,t} \quad = \quad & \pi_{0t} + \pi_{1t} educ_i + \pi_{2t} educ_i^2 + \pi_3 gender_i + \pi_4 age_{i,t} \\ & + \pi_5 local\, labor\, market_{i,t} + \pi_6 health_{i,t} + \pi_7 X_{i,t} + \kappa^s_{i,t}. \end{aligned}$$

This assumes that wage offers are a function of calendar year, education (and its square), gender, age indicators, local labor market conditions (which include municipality fixed-effects and municipality-level unemployment rate and average labor income), health (current and lagged hospitalization), and additional characteristics $X_{i,t}$ in which we include a dummy variable for whether the person is a native or an immigrant and indicators for the number of children (of any age). The coefficients on education are allowed to vary over time. To account for selection into the labor force in the imputation of wage offers, we employ the (two-stage) Heckman (1979) correction. The analysis is repeated separately for each combination of timing (before/after the shock) and experimental group (treatment/control).

Potential Transfers. In the same manner we need to impute the expected potential government transfers

[12]For expositional reasons we use the notation that whenever the variable is multidimensional (e.g., $age_{i,t}$, which denotes a complete set of age dummies), the corresponding coefficient is a vector of the same dimension (e.g., π_4 has as many entries as the number of unique ages observed in our sample).

in the case an individual chooses not to work. The labor-supply-dependent transfers are Social Disability Insurance (Social DI) benefits, which are awarded in Denmark for medical reasons as well as for social reasons. Recall that Social DI is a state-wide means-tested program that is locally administered (at the municipality level). Hence, we model expected benefits as a function of calendar year dummies (which capture overall national trends in benefits), municipality fixed effects, and interactions of municipality dummies with year dummies. The source of variation we use to identify the effect of potential transfers on participation is within municipalities over time since we include municipality and calendar year fixed effects as controls in the participation equation (23). We also include deciles of gross wealth, liabilities, and home value since some portion of DI is asset-tested, as well as age dummies, gender, and health indicators (hospitalization and lagged hospitalization). We use the following specification

$$\begin{aligned} b^s_{w,i,t} &= \sigma_{0t} + \sigma_1 municipality_i + \sigma_2 municipality_i \times year_{i,t} + \sigma_3 age_{i,t} + \sigma_4 gender_i + \sigma_5 health_{i,t} \\ &\quad + \sigma_6 gross\, wealth_{i,t} + \sigma_7 liabilities_{i,t} + \sigma_8 home\, value_{i,t} + \omega^s_{i,t}. \end{aligned}$$

We estimate this equation using the sample of individuals that do not participate in the labor force, separately for different combinations of timing (before/after the shock) and experimental groups (treatment/control). In this way we construct the transfers an agent who decides not to work expects to receive at time t in state s.

Non-Labor Income and Net-Wealth. We want a measure for non-labor income that is exogenous to other decisions such as take-up of social benefits (beyond direct government transfers that are captured by $b^s_{w,i,t}$), withdrawals from savings accounts, claims from private insurance policies, etc. Therefore, we treat w's component of unearned income $y^s_{i,t}$ as endogenous (following Blundell et al. 2007), and use predictions based on reduced-form projections, which we run for each combination of timing and experimental group for the effective unearned income on a series of pre-shock household economic variables and characteristics.[13] We then construct non-labor income $y^s_{i,t}$ as the sum of h's income and w's predicted non-labor income. To account for potential endogeneity in household-level net wealth (excluding home value), we use pre-shock wealth levels as the right-hand side variable for wealth.

Stochastic Specification and Estimation. We estimate the model as a probit and hence assume that the error in the latent index, $\varepsilon^s_{i,t}$, is normally distributed with unit variance. The participation equation is estimated using the imputed wages, the expected government benefits, the household-level non-labor income, pre-shock net wealth, and the additional controls (age, year, and municipality dummies).

[13] To improve the fit of this reduced-form we included a rich set of predictors. These include age and year dummies as well as their interaction, deciles of pre-shock wealth, liabilities, and home value, pre-shock income flows from different private and social sources available in the register-based data, occupation, employment and earnings history, health indicators, education, cohort dummies, as well as gender and municipality fixed effects.

Elasticity Estimates

The estimation of the model above provides us with the following elasticities for our sub-sample of individuals younger than 60, evaluated at sample means: $\varepsilon(e_w^b, b^b) = -0.1937$ with a confidence interval of [-0.2031,-0.1842] and $\varepsilon(e_w^g, b^g) = -0.1409$ with a confidence interval of [-0.1468,-0.1350]. The estimate for their ratio is $\varepsilon(e_w^b, b^b)/\varepsilon(e_w^g, b^g) = 1.375$ with a confidence interval of [1.292,1.457].

References

Aguiar, M., and E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–948.

Aguila, E., O. Attanasio, and C. Meghir (2011): "Changes in Consumption at Retirement: Evidence from Panel Data," *Review of Economics and Statistics*, 93(3), 1094–1099.

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

Ashenfelter, O. (1980): "Unemployment as Disequilibrium in a Model of Aggregate Labor Supply," *Econometrica*, 48(3), 547–564.

Auerbach, A. J., and L. J. Kotlikoff (1991): "The Adequacy of Life Insurance Purchases," *Journal of Financial Intermediation*, 1(3), 215–241.

Autor, D., A. Kostøl, and M. Mogstad (2015): "Disability Benefits, Consumption Insurance, and Household Labor Supply," Unpublished manuscript.

Baily, M. N. (1978): "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, 10(3), 379–402.

Ball, S., and H. Low (2014): "Do Self-Insurance and Disability Insurance Prevent Consumption Loss on Disability?," *Economica*, 81(323), 468–490.

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

Benitez-Silva, H., M. Buchinsky, and J. Rust (2006): "Induced Entry Effects of a $1 for $2 Offset in SSDI Benefits," Unpublished manuscript.

Bernheim, B. D., K. G. Carman, J. Gokhale, and L. J. Kotlikoff (2003): "Are Life Insurance Holdings Related to Financial Vulnerabilities?," *Economic Inquiry*, 41(4), 531–554.

Blundell, R., P.-A. Chiappori, T. Magnac, and C. Meghir (2007): "Collective Labour Supply: Heterogeneity and Non-Participation," *Review of Economic Studies*, 74(2), 417–445.

Blundell, R., and A. Lewbel (1991): "The Information Content of Equivalence Scales," *Journal of Econometrics*, 50, 49–68.

Bound, J., J. B. Cullen, A. Nichols, and L. Schmidt (2004): "The Welfare Implications of Increasing Disability Insurance Benefit Generosity," *Journal of Public Economics*, 88(12), 2487–2514.

Bound, J., T. Stinebrickner, and T. Waidmann (2010): "Health, Economic Resources and the Work Decisions of Older Men," *Journal of Econometrics*, 156(1), 106–129.

Browning, M., P.-A. Chiappori, and A. Lewbel (2013): "Estimating Consumption Economies of Scale, Adult Equivalence Scales, and Household Bargaining Power," *Review of Economic Studies*, 80(4), 1267–1303.

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

Chandra, A., and A. A. Samwick (2009): "Disability Risk and the Value of Disability Insurance," in *Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly*, pp. 295–336. University of Chicago Press.

Chetty, R. (2006a): "A General Formula for the Optimal Level of Social Insurance," *Journal of Public Economics*, 90(10-11), 1879–1901.

——— (2006b): "A New Method of Estimating Risk Aversion," *American Economic Review*, 96(5), 1821–1834.

——— (2008): "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 116(2), 173–234.

——— (2009): "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1(1), 451–488.

Chetty, R., and A. Finkelstein (2013): "Chapter 3 - Social Insurance: Connecting Theory to Data," vol. 5 of *Handbook of Public Economics*, pp. 111–193. Elsevier.

Chiappori, P.-A. (1988): "Rational Household Labor Supply," *Econometrica*, 56(1), 63–90.

——— (1992): "Collective Labor Supply and Welfare," *Journal of Political Economy*, 100(3), 437–467.

Coile, C. C. (2004): "Health Shocks and Couples' Labor Supply Decisions," Working Paper 10810, National Bureau of Economic Research.

Cullen, J. B., and J. Gruber (2000): "Does Unemployment Insurance Crowd Out Spousal Labor Supply?," *Journal of Labor Economics*, 18(3), 546–572.

Cutler, D. M., and L. F. Katz (1992): "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980's," *American Economic Review*, 82(2), 546–551.

Finkelstein, A., E. F. P. Luttmer, and M. J. Notowidigdo (2009): "Approaches to Estimating the Health State Dependence of the Utility Function," *American Economic Review*, 99(2), 116–121.

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. (1996): "Disability Insurance Benefits and Labor Supply," Working Paper 5866, National Bureau of Economic Research.

——— (1997): "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, 87(1), 192–205.

Gruber, J., and J. B. Cullen (1996): "Spousal Labor Supply as Insurance: Does Unemployment Insurance Crowd Out the Added Worker Effect?," Working Paper 5608, National Bureau of Economic Research.

Heckman, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, pp. 153–161.

Heckman, J. J., and T. E. Macurdy (1980): "A Life Cycle Model of Female Labour Supply," *Review of Economic Studies*, 47(1), 47–74.

——— (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. (2014): "Parental Credit Constraints and College Attendance: Evidence from Timing of Fathers' Layoffs," Unpublished manuscript.

Huang, H., S. İmrohoroğlu, and T. Sargent (1997): "Two Computations to Fund Social Security," *Macroeconomic Dynamics*, 1(1), 7–44.

Imbens, G. W., and J. M. Wooldridge (2009): "Recent Developments in the Econometrics of Program Evaluation," *Journal of Economic Literature*, 47(1), 5–86.

Immervoll, H., H. J. Kleven, C. T. Kreiner, and N. Verdelin (2011): "Optimal Tax and Transfer Programs for Couples with Extensive Labor Supply Responses," *Journal of Public Economics*, 95(11), 1485–1500.

İmrohoroğlu, A., S. İmrohoroğlu, and D. H. Joines (1995): "A Life Cycle Analysis of Social Security," *Economic theory*, 6(1), 83–114.

——— (2003): "Time-Inconsistent Preferences and Social Security," *Quarterly Journal of Economics*, 118(2), 745–784.

Keane, M. P., P. E. Todd, and K. I. Wolpin (2011): "Chapter 4 - The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications," vol. 4, Part A of *Handbook of Labor Economics*, pp. 331–461. Elsevier.

Kotlikoff, L. J., K. Smetters, and J. Walliser (1999): "Privatizing Social Security in the United States – Comparing the Options," *Review of Economic Dynamics*, 2(3), 532–574.

Kroft, K., F. Lange, and M. J. Notowidigdo (2013): "Duration Dependence and Labor Market Conditions: Evidence from a Field Experiment," *Quarterly Journal of Economics*, 128(3), 1123–1167.

LaLumia, S. (2013): "The EITC, Tax Refunds, and Unemployment Spells," *American Economic Journal: Economic Policy*, 5(2), 188–221.

Landais, C. (forthcoming): "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design," *American Economic Journal: Economic Policy*.

Low, H., and L. Pistaferri (2012): "Disability Insurance and the Dynamics of the Incentive-Insurance Tradeoff," Discussion paper.

Lundberg, S. (1985): "The Added Worker Effect," *Journal of Labor Economics*, 3(1), 11–37.

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.

——— (1991): "Unobserved Variables and the Elusive Added Worker Effect," *Economica*, 58(230), 173–187.

Meyer, B. D., and W. K. Mok (2013): "Disability, Earnings, Income and Consumption," Working Paper 18869, National Bureau of Economic Research.

Nelson, J. A. (1988): "Household Economies of Scale in Consumption: Theory and Evidence," *Econometrica*, 56(6), 1301–1314.

Nishiyama, S., and K. Smetters (2007): "Does Social Security Privatization Produce Efficiency Gains?," *Quarterly Journal of Economics*, 122(4), 1677–1719.

OECD (2011): "Divided We Stand: Why Inequality Keeps Rising," Discussion paper, OECD Publishing.

Ruhm, C. J. (1991): "Are Workers Permanently Scarred by Job Displacements?," *American Economic Review*, 81(1), 319–324.

Shimer, R., and I. Werning (2007): "Reservation Wages and Unemployment Insurance," *Quarterly Journal of Economics*, 122(3), 1145–1185.

Spletzer, J. R. (1997): "Reexamining The Added Worker Effect," *Economic Inquiry*, 35(2), 417–427.

Stephens, M. J. (2001): "The Long-Run Consumption Effects of Earnings Shocks," *Review of Economics and Statistics*, 83(1), 28–36.

TABLE I
Summary Statistics of Analysis Sample

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

Notes: This table presents means of key variables in our analysis sample. All monetary values are reported in nominal Danish Kroner (DKK) deflated to 2000 prices using the consumer price index. In this year the exchange rate was approximately DKK 8 per US $1. For each event, the treatment group comprises households that experienced a shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Columns 1 and 2 report statistics for the death event sample of households in which a spouse died of any cause between ages 45 and 80 from 1985 to 2011. Column 1 reports statistics for the entire sample, and Column 2 reports statistics for the sub-sample of surviving spouses under age 60. Column 3 reports statistics for the health event sample. It includes households in which one spouse experienced a heart attack or a stroke between 1985 and 2011 and survived for at least three years, and in which both spouses were under age 60. The values reported in the table are based on data from two periods before the shock occurred (period t = -2).

TABLE II
Survivors' Labor Supply Responses to the Death of Their Spouse

A. Surviving Spouses of All Ages

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

B. Surviving Spouses under 60

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

Notes: This table reports the differences-in-differences estimates of the surviving spouses' labor supply responses (equation (6)). The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Panel A reports the responses of all survivors by gender, where widowers are those who lost their wives and widows are those who lost their husbands. Panel B reports the responses of survivors under 60 by gender. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

TABLE III

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

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

Notes: This table reports the interaction of the treatment effect of the death of a spouse with the household's post-shock income replacement rate (equation (7)). The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011.The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Panel A reports estimates for the sample of all survivors by gender; Panel B reports estimates for the sample of survivors under age 60 by gender. In each panel, we report estimates of two specifications. Specification 1 in each panel estimates a baseline differences-in-differences specification which interacts the treatment effect with the replacement rate variable. This replacement rate is calculated as follows. First, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their pre-shock levels (in period -1). Then, we calculate the ratio of this adjusted household income in period 1 (post-shock) to that in period -1 (pre-shock), and normalize it by the average ratio for the control group in order to account for calendar year trends as well as for life-cycle effects. Specification 2 in each panel extends specification 1 to include interactions of the treatment effect with additional household characteristics: age dummies for the surviving spouse, dummies for the age of the deceased at the year of death, year dummies, indicators for the number of children in the household as well as the surviving spouse's months of education (and its square). The results are also robust to the inclusion of a quadratic in the household's net wealth. All the variables that are interacted with "Treat × Post" are interacted with "Treat" and "Post" and enter the regressions separately as well. Since there are households with missing values for some of the controls (that are therefore included in the estimation of specification 1 but not 2), we show the robustness of our estimate of interest ("Treat × Post × Replacement Rate") to the inclusion of this set of controls by reporting estimates for specification 1 for the sub-sample of households that are included in the estimation of specification 2. All specifications include year, age, and household fixed effects. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

TABLE IV
Survivors' Annual Earnings Responses to the Death of Their Spouse

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

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

Notes: This table reports the differences-in-differences estimates of the surviving spouses' annual earnings by the level of their own earnings when their spouses died. The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011, where we constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than their experimental-group-specific 20th percentile. Then, we calculate for each household the pre-shock labor income share of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force; specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households for which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household. In addition, to guarantee that our results are not driven by outliers, we exclude households with dying spouses whose mean pre-shock earnings did not fall within their group-specific 5th and 95th percentiles or households with unaffected spouses whose mean pre-shock earnings were higher than their group-specific 95th percentile. We divide the remaining sample into five equal-sized groups by their pre-shock level of earnings. Panel A separately estimates a differences-in-differences specification for each surviving spouses' quintile. Column 1 includes all surviving spouses; Column 2 includes households in which the dying spouses' pre-shock labor income fell within the bottom three quintiles of its group-specific distribution, to which we refer as "low-earners"; Column 3 includes households in which the dying spouses' pre-shock labor income fell within the top two quintiles, to which we refer as "high-earners". The gradient of survivors' labor supply responses with respect to their own level of pre-shock earnings is also robust to the inclusion of a quadratic in the household's net wealth. Panel B reports the average treatment effect for this sample. The second row reports the counterfactual outcome based on the differences-in-differences estimation. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

TABLE V

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

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

Notes: This table reports the interaction of the treatment effect of the death of a spouse with the actual survivors benefits widows received through the Social Disability Insurance (Social DI) program (equation (8)). The regression is estimated by two-stage least squares, where the instrument for actual benefits is constructed as follows. In each year we calculate for each municipality the average benefits received by non-working surviving spouses through Social DI. Then, we assign to each widow in the treatment group her respective municipality-year leave-one-out mean. The sample includes widows under age 67 (the age at which the program transitions into the Old-Age Pension) in years prior to 1994 (when there is a data break in the reporting method of survivors benefits received through Social DI). The controls included in the estimation are municipality unemployment rate and average earnings (and their interaction with "Treat", "Post", and "Treat × Post") as well as age, year, and municipality fixed effects. The identifying assumption is that, given our set of controls, the average social survivors benefits transferred to widows in a municipality in a given year affects a widow's participation only through its influence on her own survivors benefits receipts. Note that the source of variation we use is within municipalities over time since we include municipality and calendar year fixed effects as controls. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the municipality level are reported in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

TABLE VI
Household Responses to Severe Health Shocks in which the Affected Spouse Survived

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

Notes: This table reports the differences-in-differences estimates of household labor supply responses to severe health shocks in which the affected spouse survived and the effect of these shocks on overall household income (equation (9) in footnote 32). The sample includes households in which one spouse experienced a heart attack or a stroke and survived for at least three years, and in which both spouses were under age 60. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). We allow for differential treatment effects for the "short run" – periods 1 and 2 – and the "medium run" – period 3, to account for the gradual responses documented in Figure VII. The pre-shock periods include periods -5 to -2. Household income (Columns 5 and 6) includes income from any source – including earnings, capital income, annuity payouts, and benefits from any social program. The third row reports the counterfactual outcome based on the differences-in-differences estimation. Robust standard errors clustered at the household level are reported in parentheses.
*** p<0.01, ** p<0.05, * p<0.1.

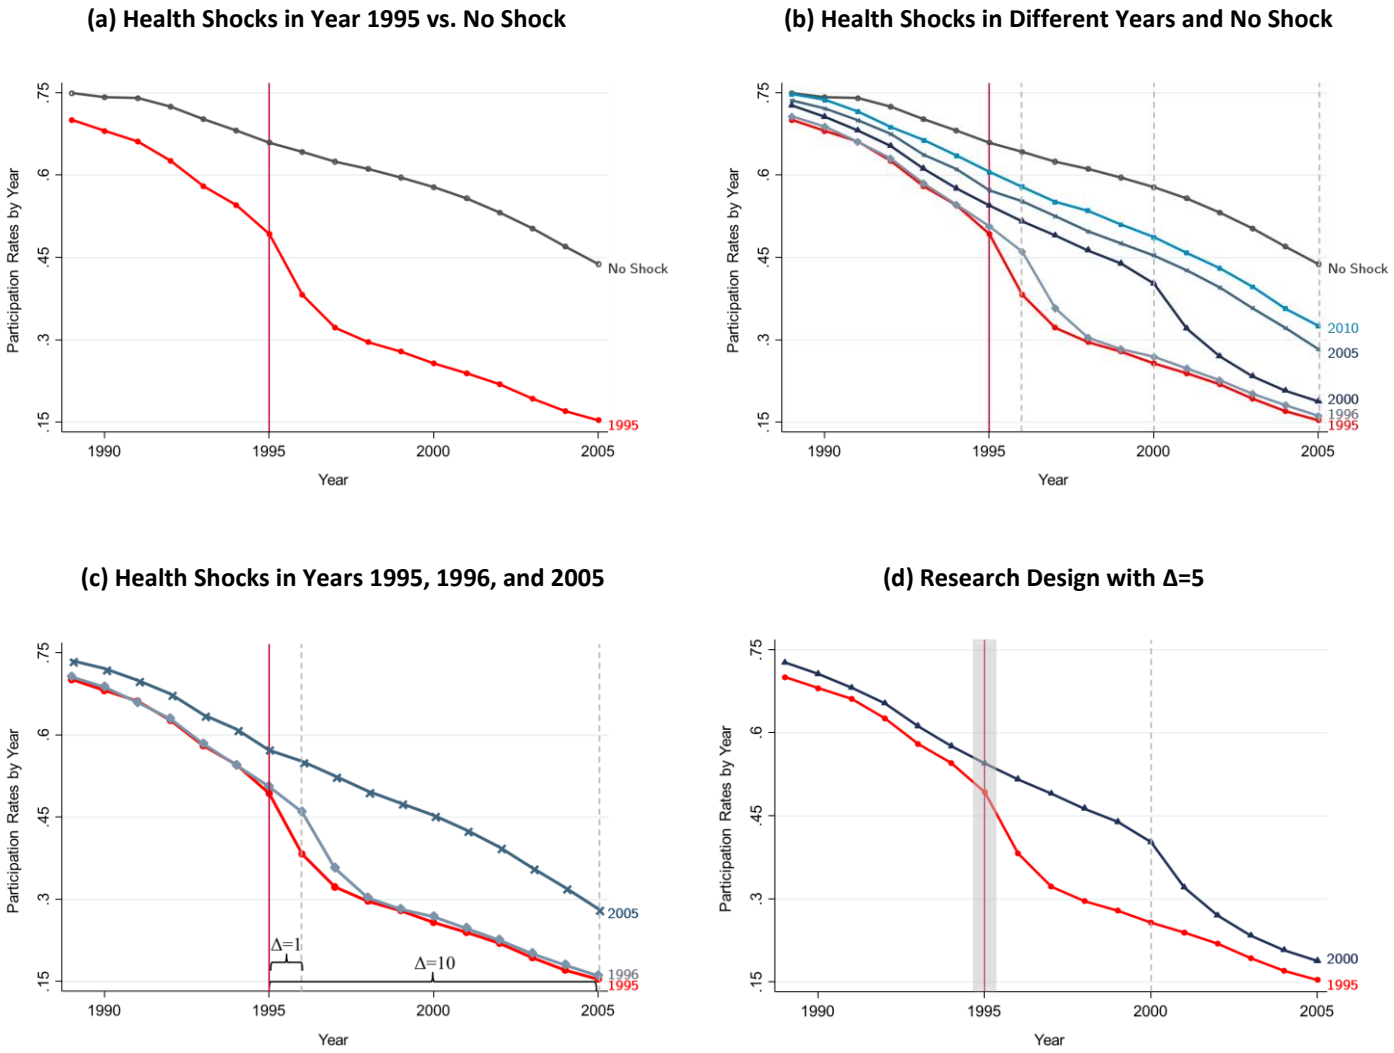
ONLINE APPENDIX TABLE I
Labor Force Participation Responses of Widows Who Did Not Work before the Shock

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

Notes: This table reports the differences-in-differences estimates of the labor force participation responses of widows who did not work during the five-year period preceding their spouse's death. The sample includes households in which the husband died between ages 45 and 80 from 1985 to 2011 and in which he either worked throughout the entire five-year period preceding his death (periods -5 to -1) or did not work altogether during this period. The treatment group comprises households that experienced the shock in different years, to which we match households that experienced the same shock five years later as a control group (Δ=5). Column 1 reports the simple differences-in-differences estimate in a regression in which the outcome variable is spousal labor force participation. Column 2 adds an interaction of the treatment effect with an indicator for whether the husband worked before his death. Column 3 runs the same specification as in Column 2 but where the outcome variable is the household's overall income. Columns 4 and 5 report the spousal labor force participation effect for sub-samples of the households in which the husband did not work before his death. Column 4 reports the treatment effect for households in which the non-working deceased's overall income before the shock (periods -3 to -1), including any transfer from government programs, was lower than the 10th percentile of this sample's income distribution; Column 5 reports the treatment effect for households in which the non-working deceased's overall income before the shock was lower than the 5th percentile. All specifications include year, age, and household fixed effects. The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4. Robust standard errors clustered at the household level are reported in parentheses.

*** p<0.01, ** p<0.05, * p<0.1.

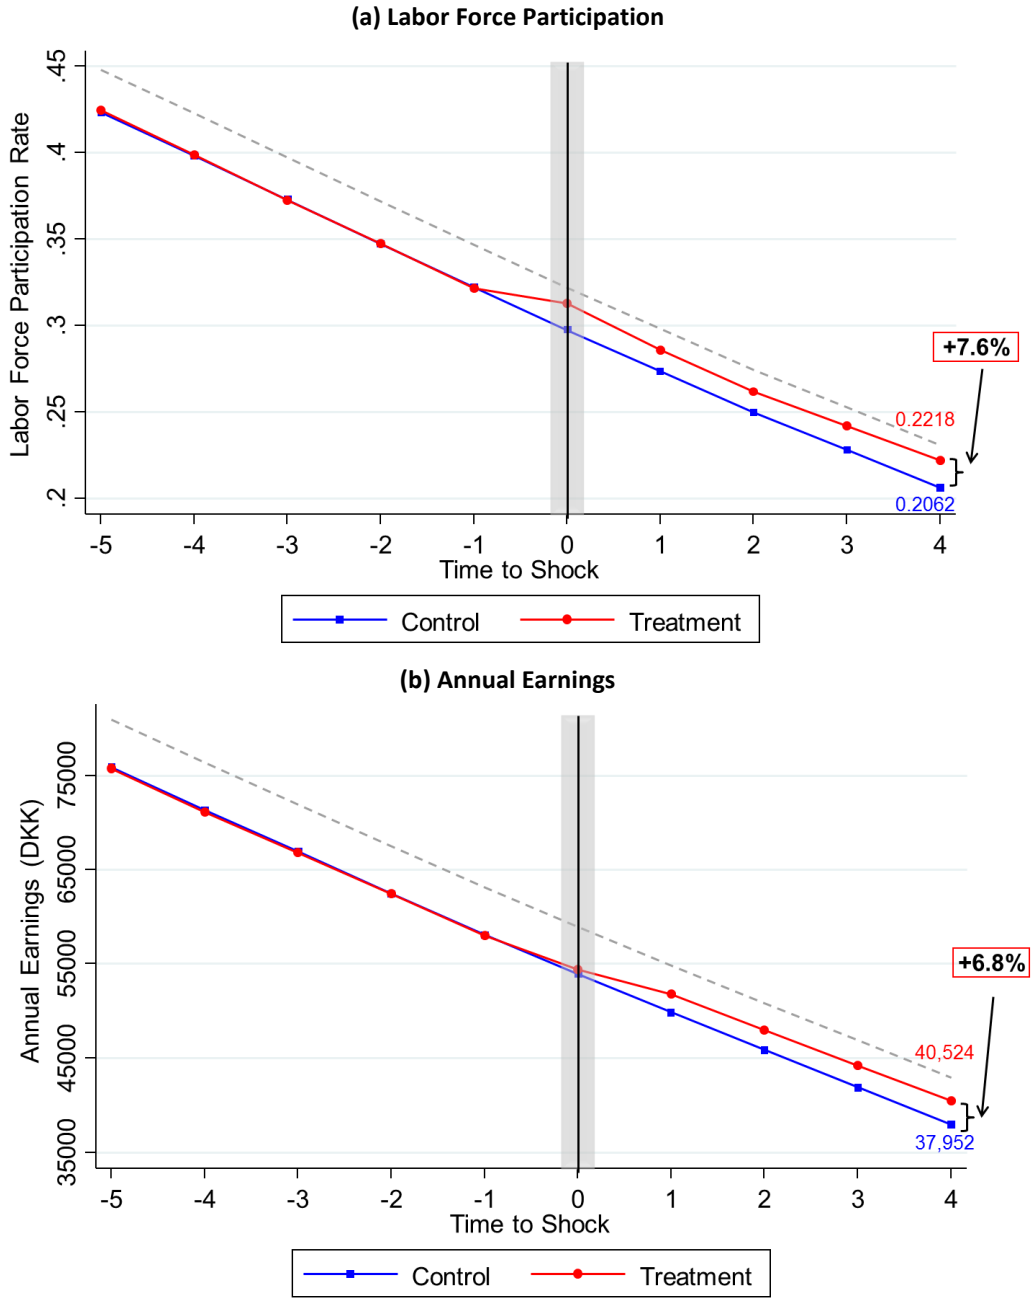
FIGURE I
Illustration of the Empirical Research Design



Notes: These figures compare the labor force participation of a treatment group of individuals who were born between 1930 and 1950 and experienced a heart attack or a stroke in 1995 to that of potential control groups. Panel A compares the treatment group to those who belong to the same cohorts but did not experience a shock in our data window, years between 1985 and 2011, and shows that the pre-1995 patterns of these groups are far from parallel. Panel B adds the behavior of households that experienced the same shock but in different years, and shows that the groups are becoming increasingly comparable to the treatment group – in terms of parallel trends before 1995 – the closer the year in which the individual experienced the shock was to the year the treatment group experienced the shock (1995). The figures suggest using households that experienced a shock in year 1995+Δ as a control group for households that experienced a shock in 1995. The trade-off in the choice of Δ is presented in Panel C. On the one hand, we would want to choose a smaller Δ such that the control group would be more closely comparable to the treatment group, e.g., year 1996 which corresponds to Δ=1. On the other hand, we would want to choose a larger Δ in order to be able to identify longer-run effects of the shock, up to period Δ-1. Using those that experienced a shock in 2005, which corresponds to Δ=10, will allow us to estimate up to the 9-year effect of the shock. However, this entails a potentially greater bias since the trend in the behavior of this group prior to 1995 is not as tightly parallel to that of the treatment group. Panel D displays the potential control group for this example when we choose Δ=5. Our research design generalizes this example by aggregating different calendar years.

FIGURE II

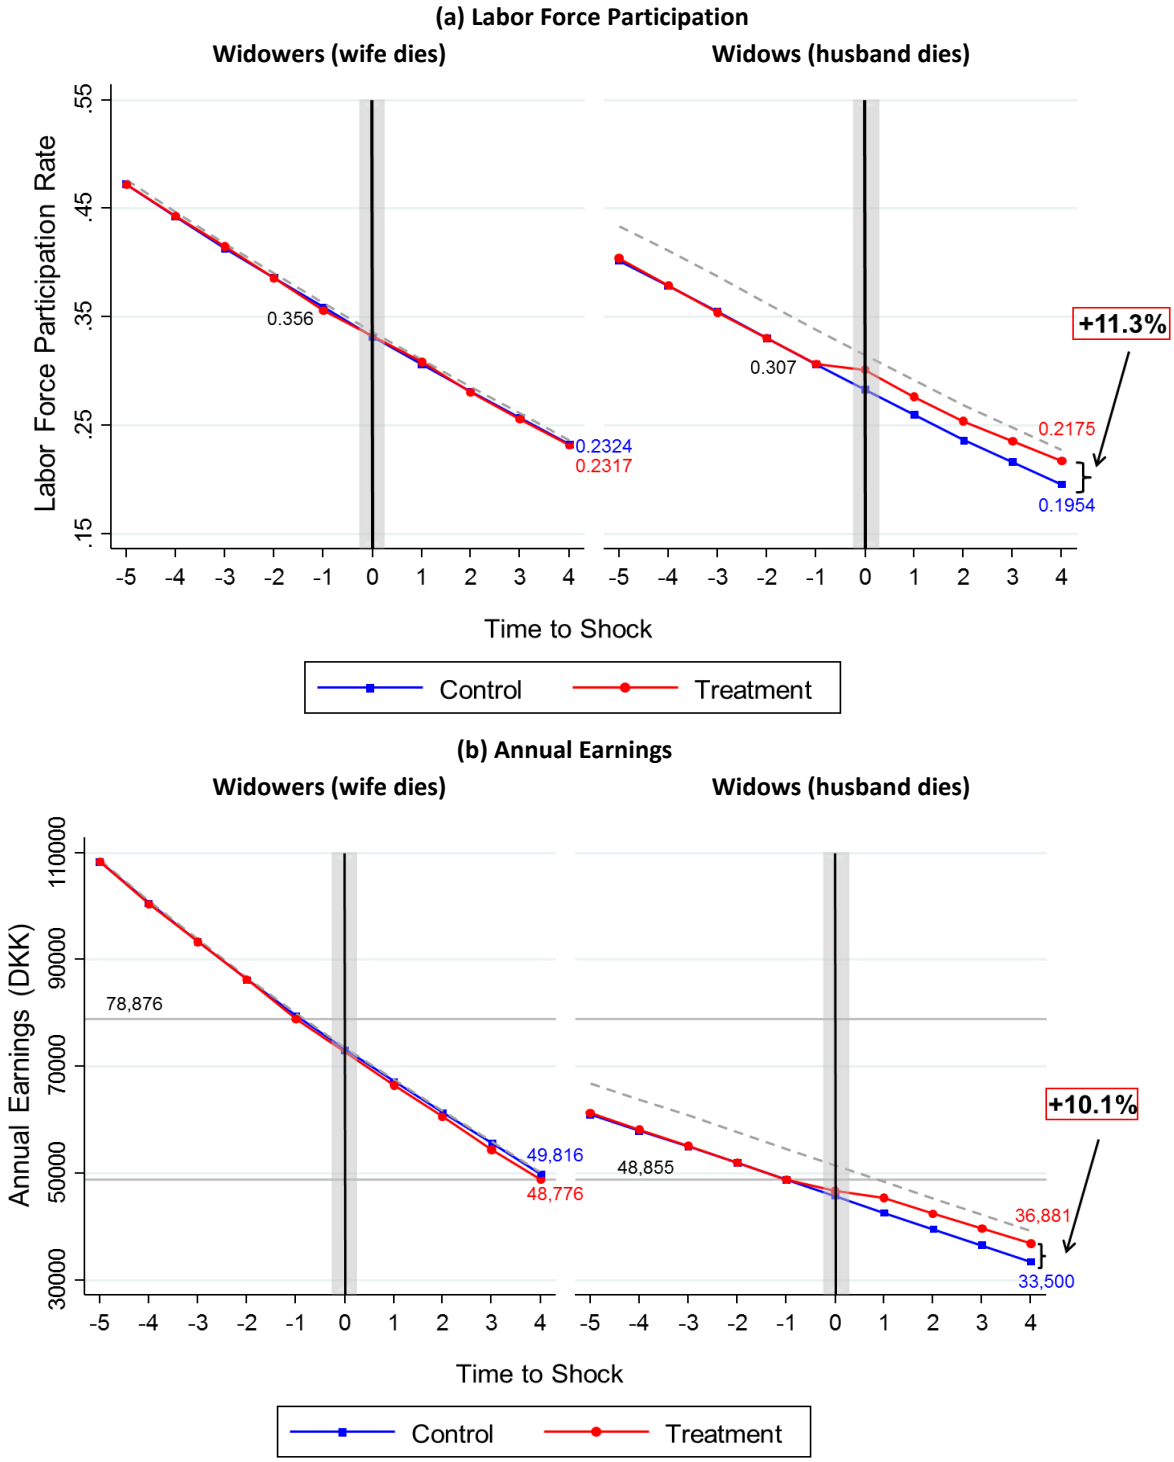
Survivors' Labor Supply Responses to the Death of Their Spouse



Notes: These figures plot the labor supply responses of survivors to the death of their spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A depicts the behavior of labor force participation, and Panel B depicts the behavior of annual earnings. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

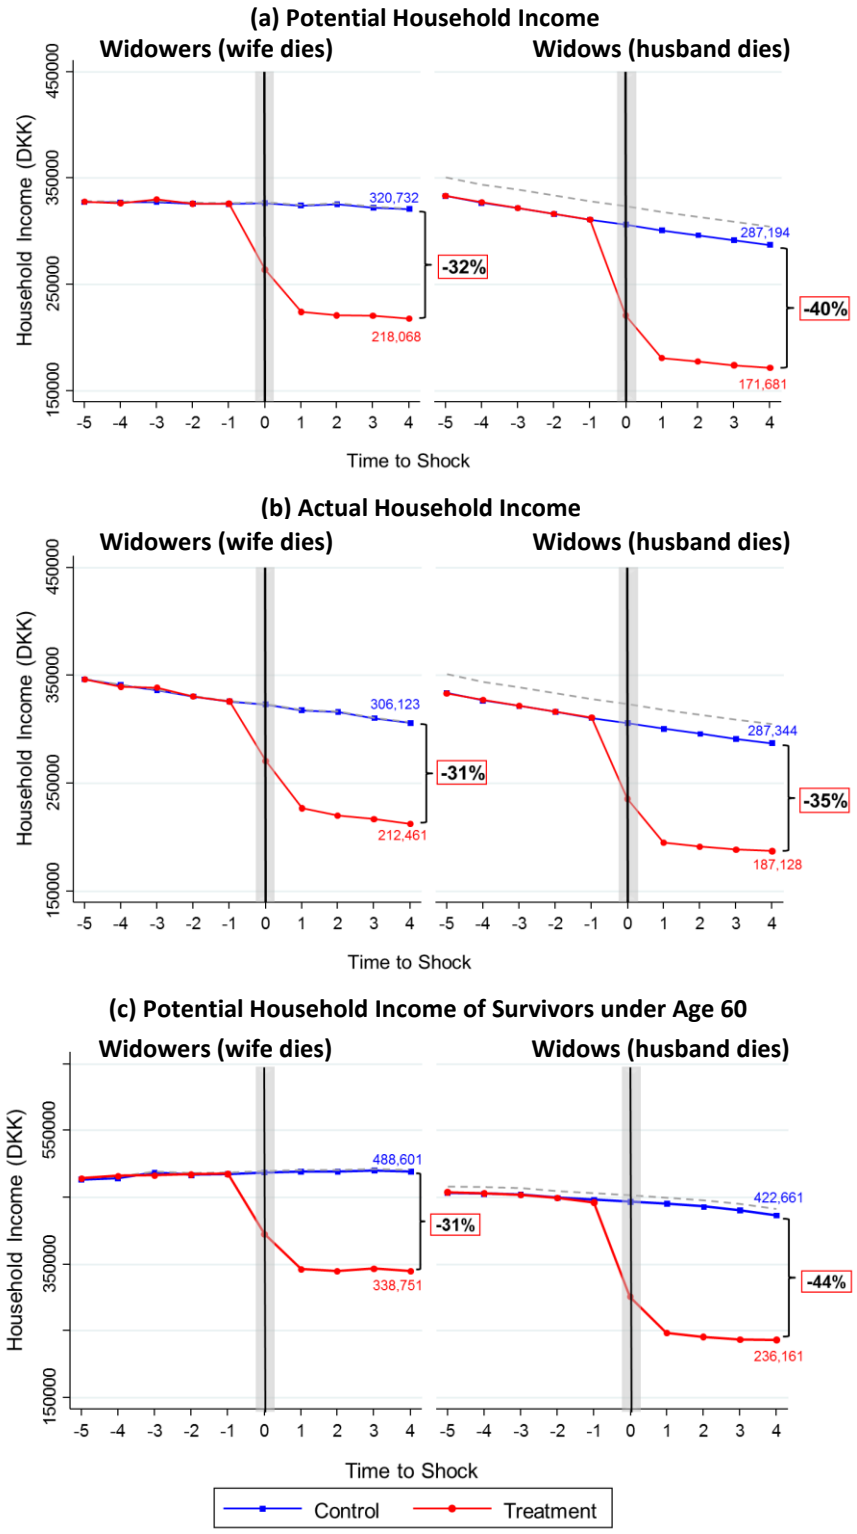
FIGURE III

Survivors' Labor Supply Responses to the Death of Their Spouse by Gender



Notes: These figures plot the labor supply responses of survivors to the death of their spouse by the gender of the surviving spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A depicts the behavior of labor force participation, and Panel B depicts the behavior of annual earnings. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group period 0 is when the actual shock occurs; for the control group period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

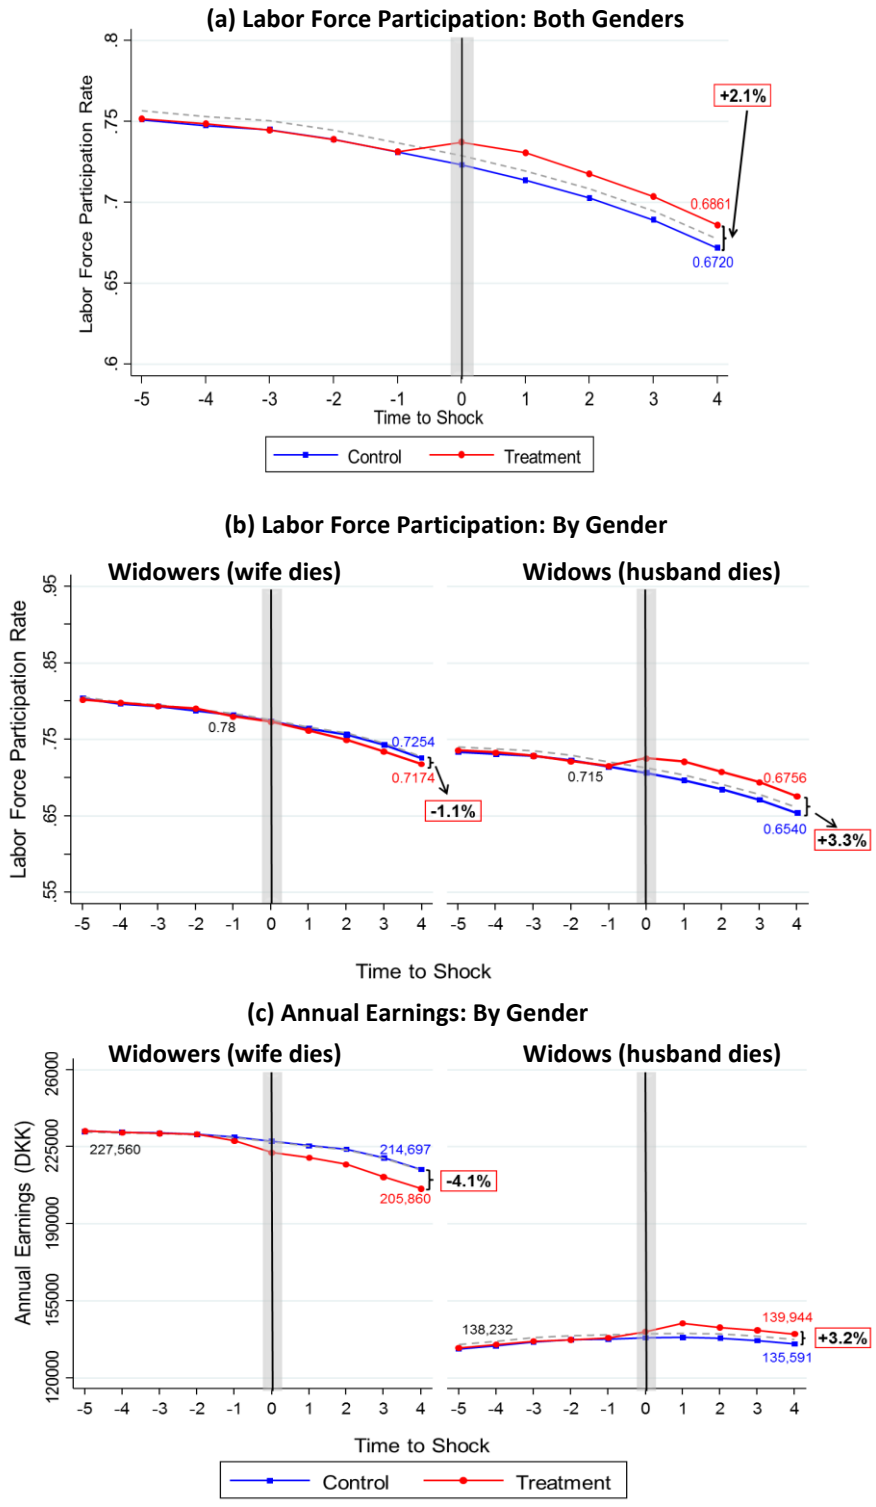
FIGURE IV
Household Income around the Death of a Spouse



Notes: These figures plot different measures of household-level income around the death of a spouse by the gender of the surviving spouse. The sample includes individuals whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A plots an adjusted measure of household income. Specifically, we fix the surviving spouse's labor income, Social Disability and Social Security benefits as well as sick-pay benefits at their pre-shock levels (in period -1). Hence, this measure captures the income loss that is directly attributed to the loss of a spouse. Panel B plots the actual household income that is observed in the data, which takes into account the surviving spouse's behavioral responses. Panel C plots the adjusted measure from Panel A for survivors under age 60. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group period 0 is when the actual shock occurs; for the control group period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

FIGURE V

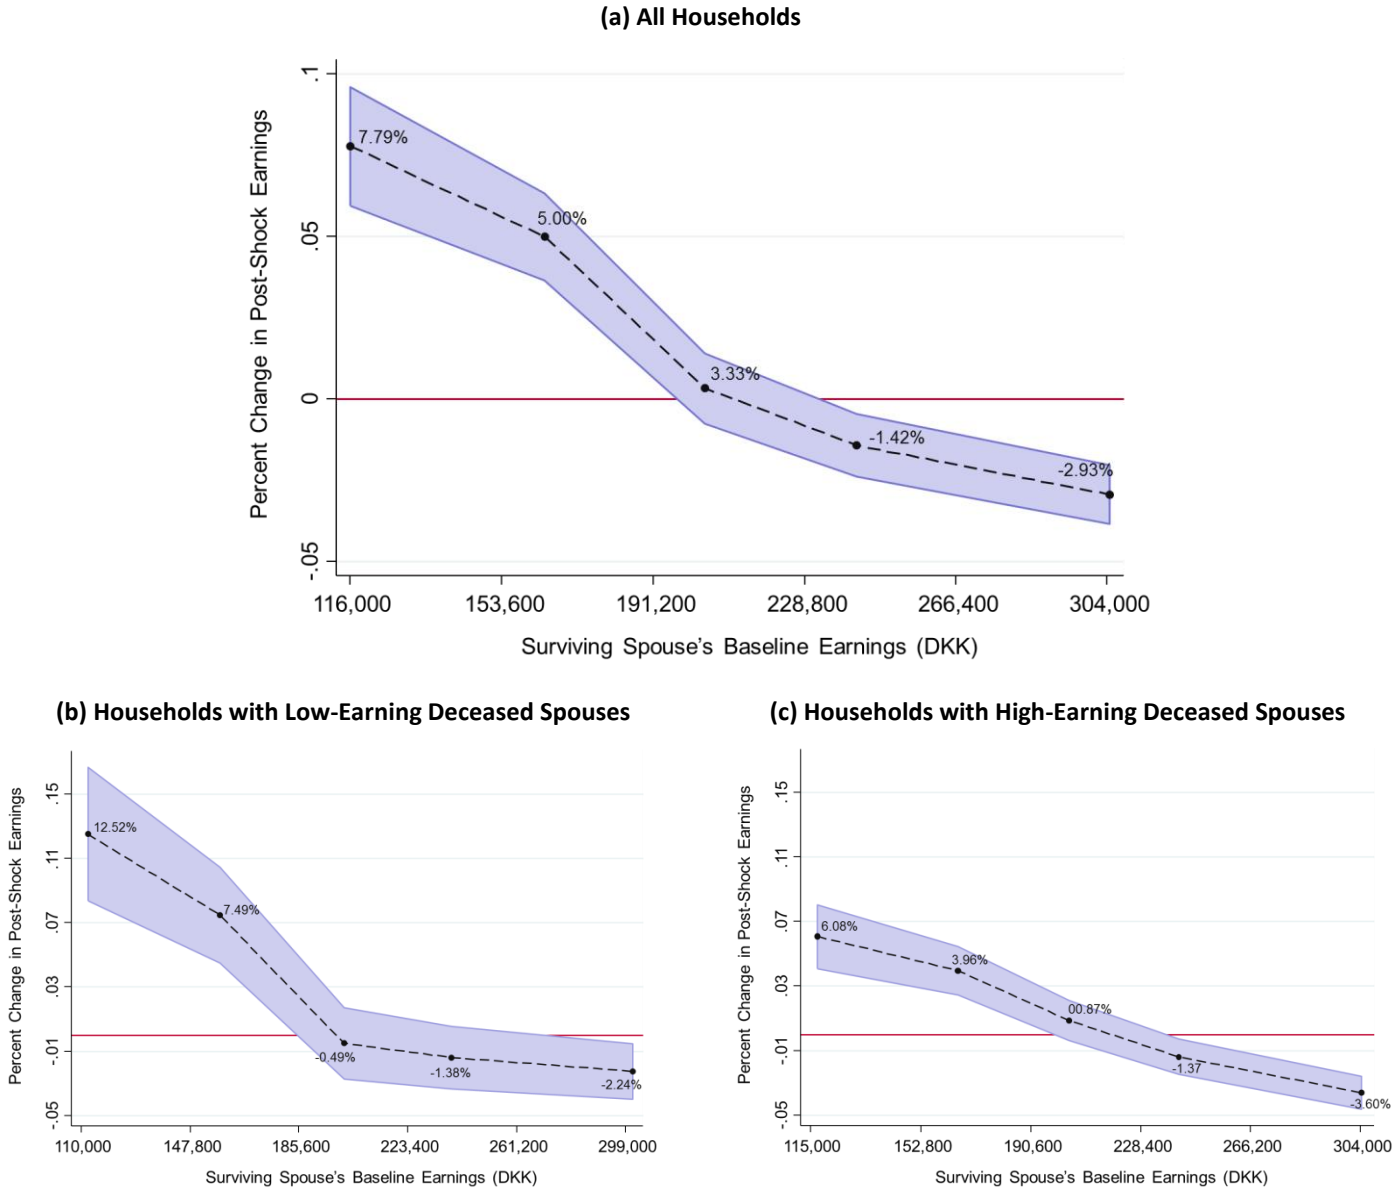
Labor Supply Responses of Survivors under Age 60 to the Death of Their Spouse



Notes: These figures plot the labor supply responses of survivors under age 60 to the death of their spouse. The sample includes individuals under 60 whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A depicts the behavior of labor force participation of the overall sample; Panel B divides the sample by the gender of the surviving spouse; Panel C depicts the behavior of annual earnings by the gender of the surviving spouse. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group period 0 is when the actual shock occurs; for the control group period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

FIGURE VI

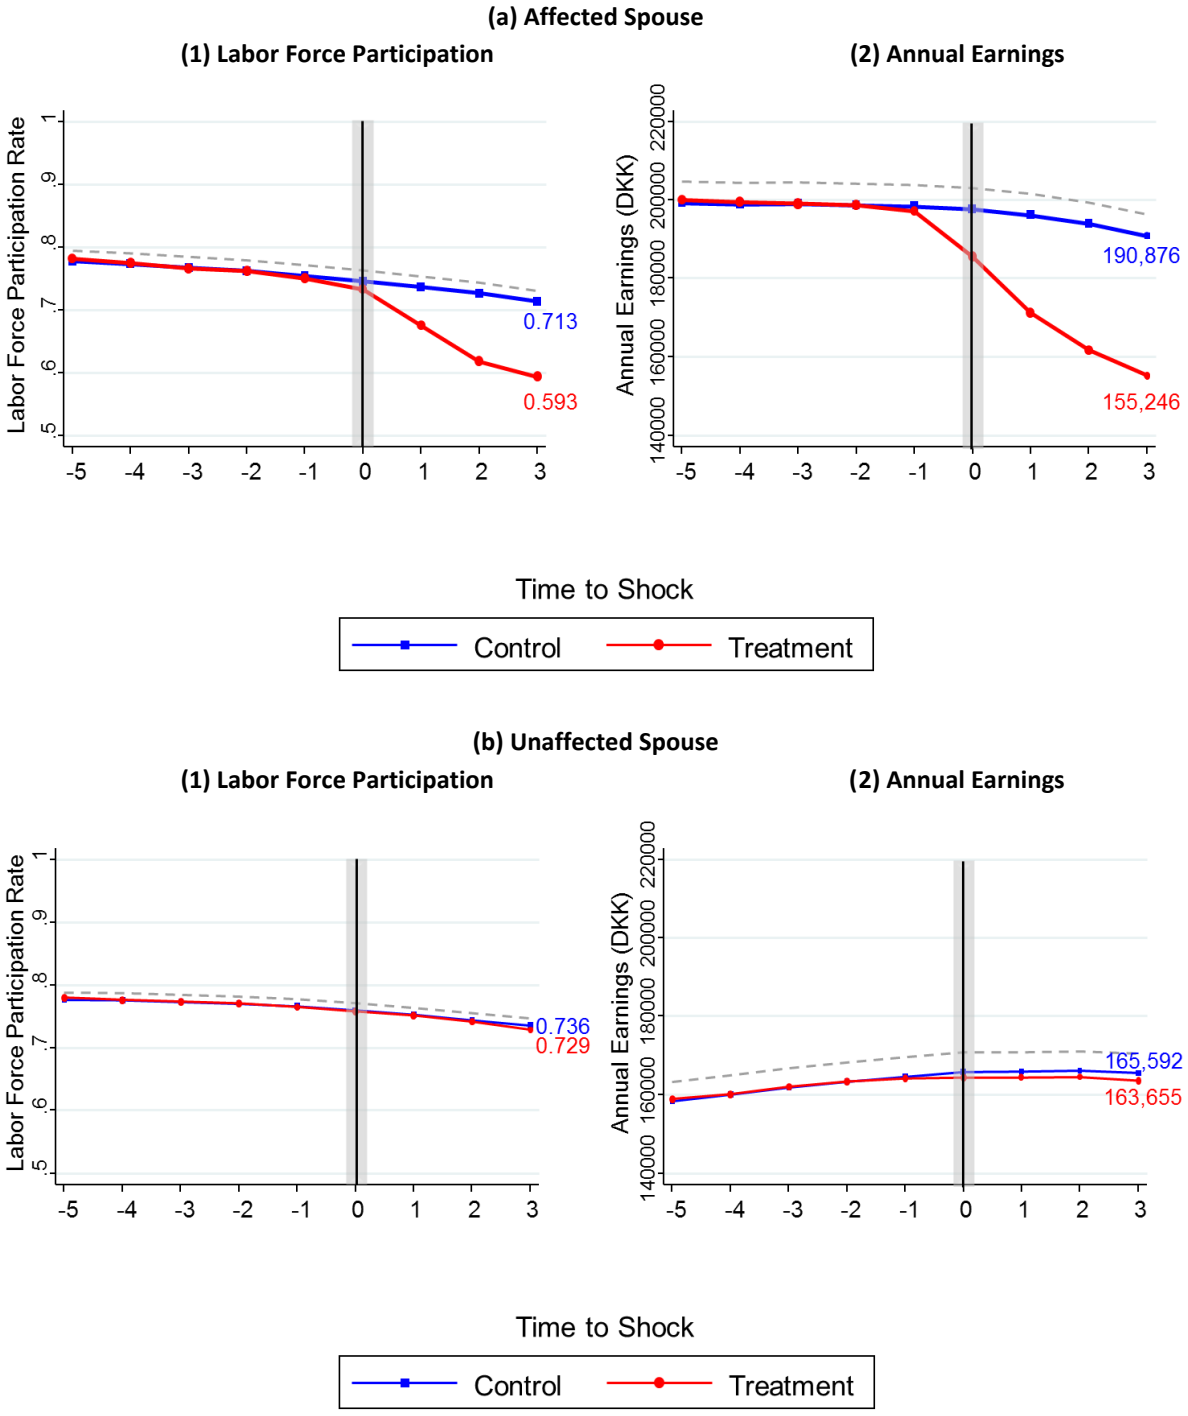
Survivors' Annual Earnings Responses to the Death of Their Spouse
by the Level of their Own Pre-Shock Earnings



Notes: These figures include individuals whose spouses died between ages 45 and 80 from 1985 to 2011, where we constrain the sample in the following way. First, we exclude surviving spouses whose average labor income before the shock (in periods -5 to -2) was lower than their experimental-group-specific 20th percentile. Then, we calculate for each household the pre-shock labor income share of the deceased spouse out of the household's overall labor income and include only households in which both spouses were sufficiently attached to the labor force; specifically, we keep households for whom the average share was between 0.20 and 0.80. These restrictions allow us to focus on households for which there has been some loss of income due to the death of a spouse and in which the surviving spouse earned non-negligible labor income both in levels and as a share within the household. In addition, to guarantee that our results are not driven by outliers, we exclude households with dying spouses whose mean pre-shock earnings did not fall within their group-specific 5th and 95th percentiles as well as households with unaffected spouses whose mean pre-shock earnings were higher than those of their group-specific 95th percentile. We divide the remaining sample into five equal-sized groups by their pre-shock level of earnings and plot the average labor income response as well as its 95-percent confidence interval (where standard errors are calculated using the Delta method) against the pre-shock mean earnings for each group. Panel A includes all households; Panel B includes households in which the dying spouse's pre-shock labor income fell within the bottom three quintiles of its group-specific distribution, to which we refer as "low-earners"; Panel C includes households in which the dying spouse's pre-shock labor income fell within the top two quintiles, to which we refer as "high-earners". The pre-shock periods include periods -5 to -3. The post-shock periods include periods 2 to 4.

FIGURE VII

Household Labor Supply Responses to Severe Health Shocks in Which the Affected Spouse Survived

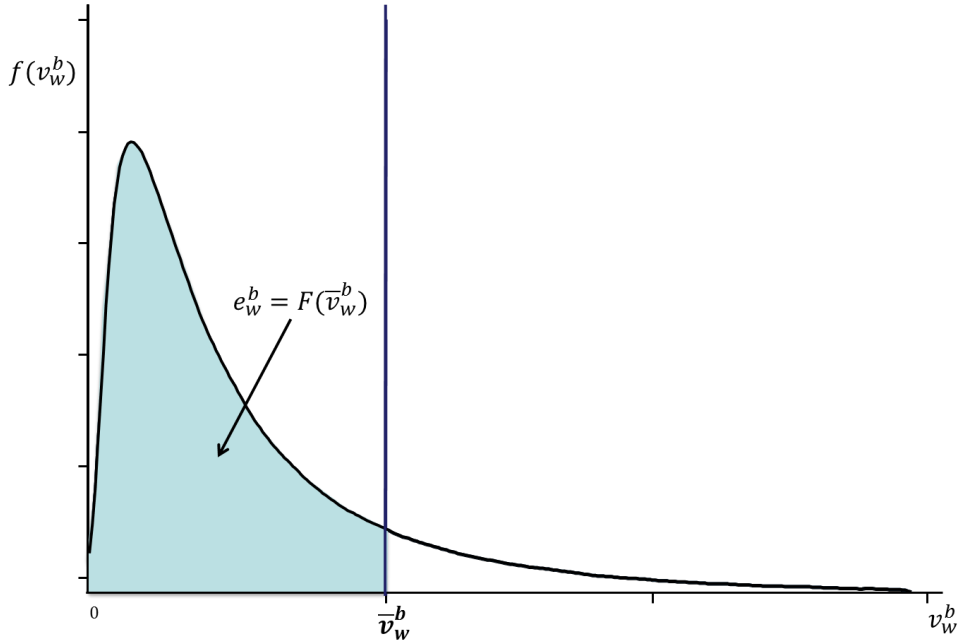


Notes: These figures plot the labor supply responses of households in which an individual experienced a heart attack or a stroke between 1985 and 2011 and survived for at least three years. The sample includes households in which both spouses were under age 60. Panels A.1 and A.2 depict the labor force participation and annual earnings of the individual that experienced the shock, respectively. Panels B.1 and B.2 depict the labor force participation and annual earnings of the unaffected spouse, respectively. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group, period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

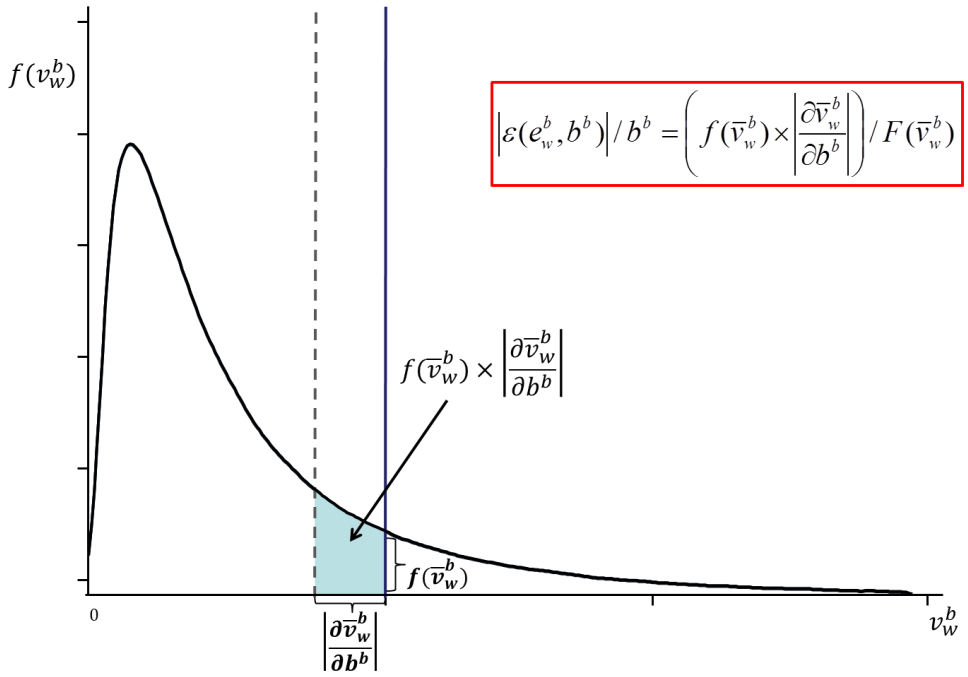
FIGURE VIII

The Unaffected Spouses' Labor Force Participation Responses to Policy Changes

(a) Spousal Labor Force Participation in the Bad State



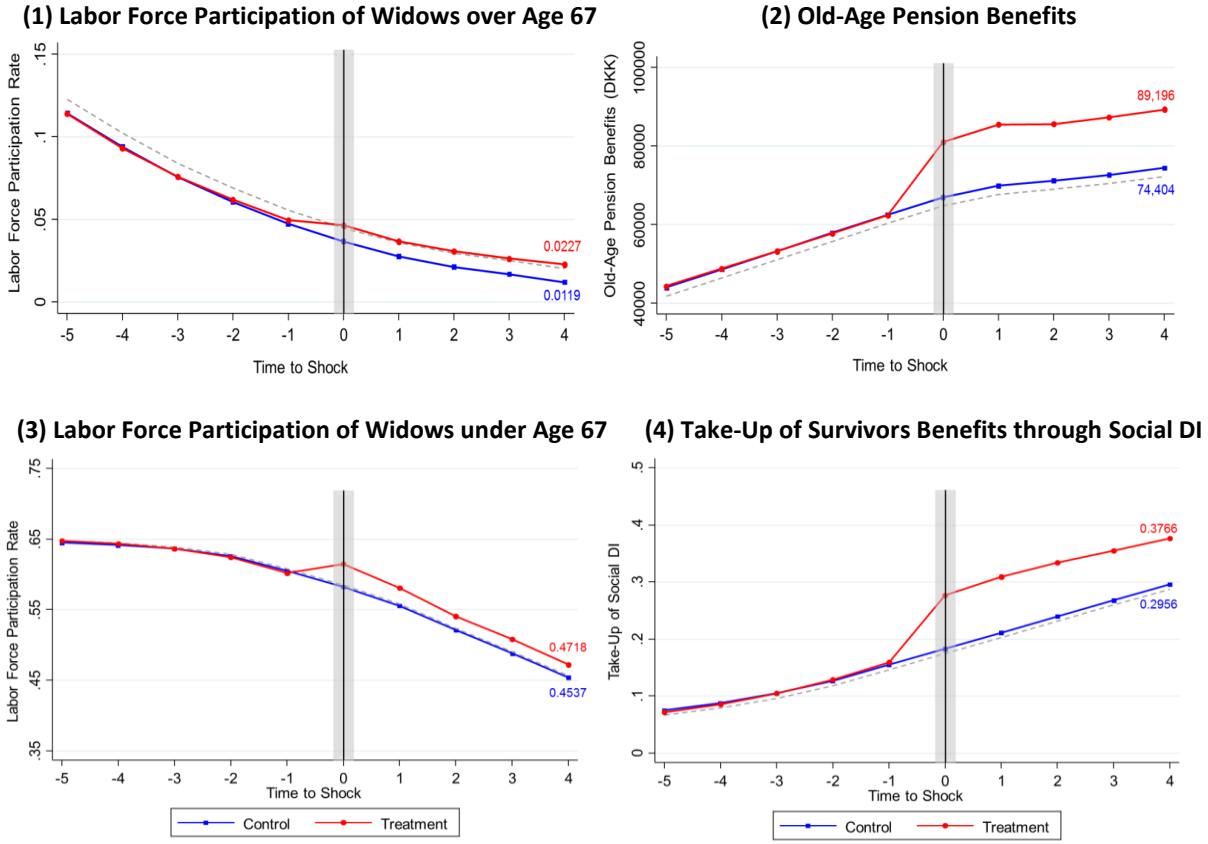
(b) The Change in Spousal Labor Force Participation in the Bad State in Response to the Policy Change



Notes: These figures plot a potential probability density function (pdf) for the labor disutility of the unaffected spouse (spouse w) in state b, v_w^b. The x-axis corresponds to v_w^b and the y-axis corresponds to the pdf, $f(v_w^b)$. In these figures, $\bar{v}_w^b$ is the threshold value below which spouse w chooses to work in state b. Therefore, the area between 0 and $\bar{v}_w^b$ below the pdf is the aggregate labor supply of spouses in state b, $e_w^b \equiv F(\bar{v}_w^s)$. This is the shaded area in Panel A. When government transfers locally change, the threshold changes by $\frac{\partial \bar{v}_w^b}{\partial b^b}$ and the approximated change in w's labor supply is the shaded area in Panel B, $f(\bar{v}_w^b) \times \left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right|$. Hence, the relative within-state change in labor force participation can be approximated by $\left(f(\bar{v}_w^b) \times \left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right|\right)/F(\bar{v}_w^b)$, which is exactly the semi-elasticity of participation, e_w^b, with respect to benefits, b^b. That is, $\varepsilon(e_w^b, b^b)/b^b = \left(f(\bar{v}_w^b) \times \left|\frac{\partial \bar{v}_w^b}{\partial b^b}\right|\right)/F(\bar{v}_w^b)$.

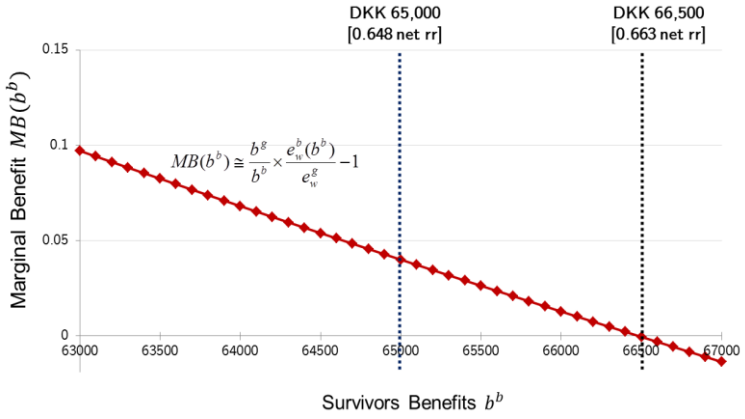
FIGURE IX
Welfare Implications

(a) Widows' Labor Force Participation and Government Transfers around the Death of Their Spouse by Age Group



Notes: These figures plot outcomes for survivors around the death of their spouse by age group. Panels A.1 and A.2 plot outcomes for widows over age 67 whose husbands died between 1985 and 2011. Panel A.1 plots their labor force participation, and Panel A.2 plots the benefits they received from the Old-Age Pension program. Panels A.3 and A.4 plot outcomes for widows under age 67 (the age at which the Social Disability Insurance transitions into the Old-Age Pension) in years prior to 1994 (when there is a data break in the reporting method of benefits received through Social Disability Insurance). Panel A.3 plots their labor force participation, and Panel A.4 plots their take-up of survivors benefits through the Social Disability Insurance program. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group, period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

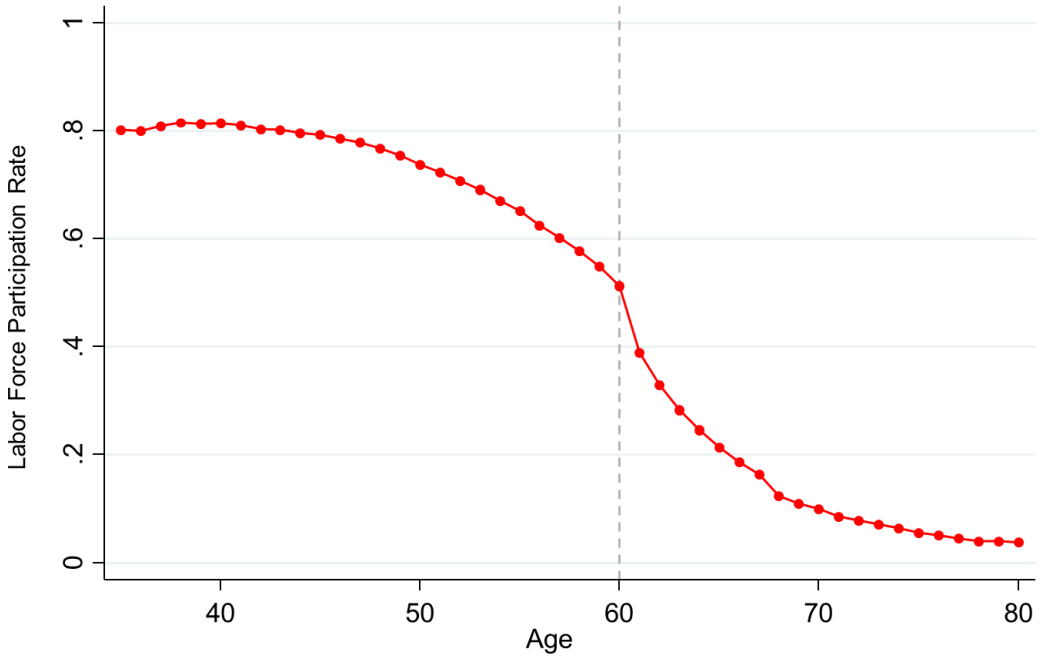
(b) Welfare Gains from Survivors Benefits within the Social Disability Insurance Program



Notes: This figure plots the marginal benefit from transfers to widows within the Social Disability Insurance program. The x-axis denotes the benefit level, b^b, measured in Danish Kroner (DKK), and the y-axis denotes the marginal benefit, $MB(b^b)$. The vertical dashed line at DKK 65,000 ($8,115) denotes the mean benefits transferred to widows who are on the program. It represents a net replacement rate (denoted by "net rr" in the figure) of 0.648 relative to the mean pre-shock annual earnings of deceased spouses who worked before they died. To convert the monetary values into net replacement rates out of the deceased spouse's pre-shock earnings, we calculate the average earnings in t = -1 for deceased spouses who had positive earnings in the year before they died. The average is DKK 170,000 ($21,250), which implies net wage earnings of DKK 100,300 ($12,538) using an average labor income tax rate of 41% (OECD estimates). The vertical dashed line at DKK 66,500 ($8,300) denotes the benefit level that sets the marginal benefit to zero. It represents a net replacement rate (denoted by "net rr" in the figure) of 0.663. Using our model, this suggests that for widows under 67 the current benefit levels are near optimal.

ONLINE APPENDIX FIGURE I

Life-Cycle Labor Force Participation of the Unaffected Spouses in the Death Event Sample

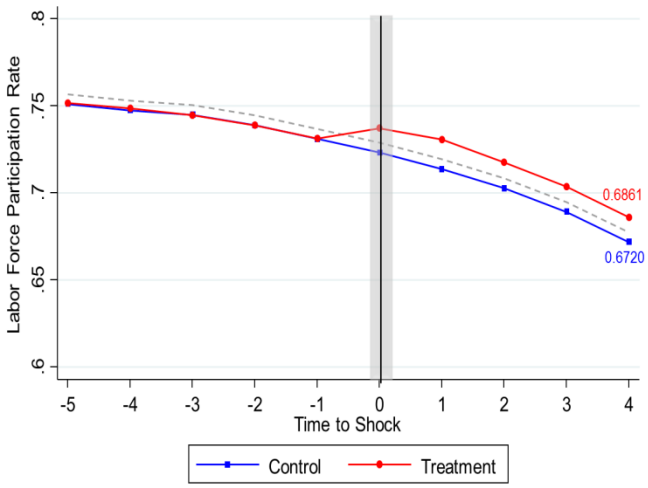


Notes: This figure displays the life-cycle labor force participation of the unaffected spouses that are included in the death event sample (i.e., individuals whose spouses died between ages 45 and 80 from 1985 to 2011). The observations include the pre-shock periods (specifically, periods -5 to -2).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

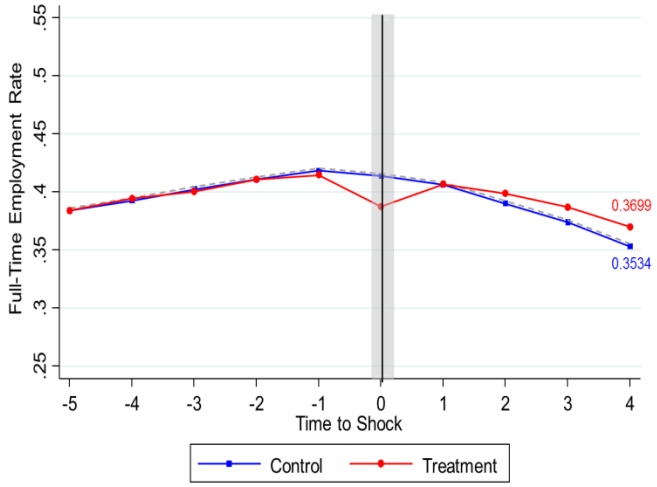
ONLINE APPENDIX FIGURE II

Labor Supply Responses of Survivors under Age 60 to the Death of Their Spouse

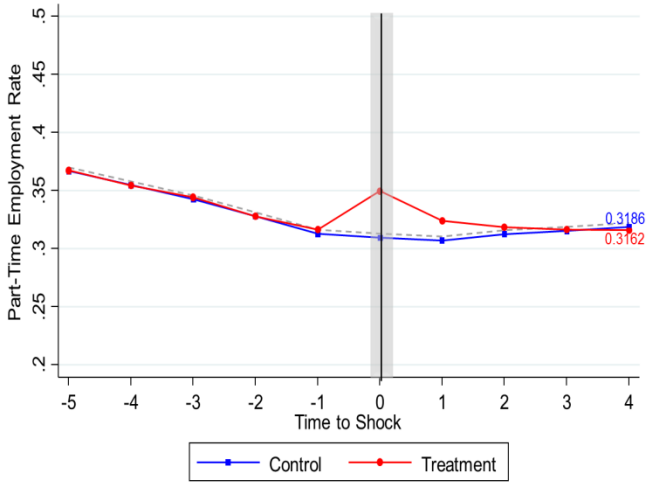
(a) Labor Force Participation



(b) Full-Time Employment



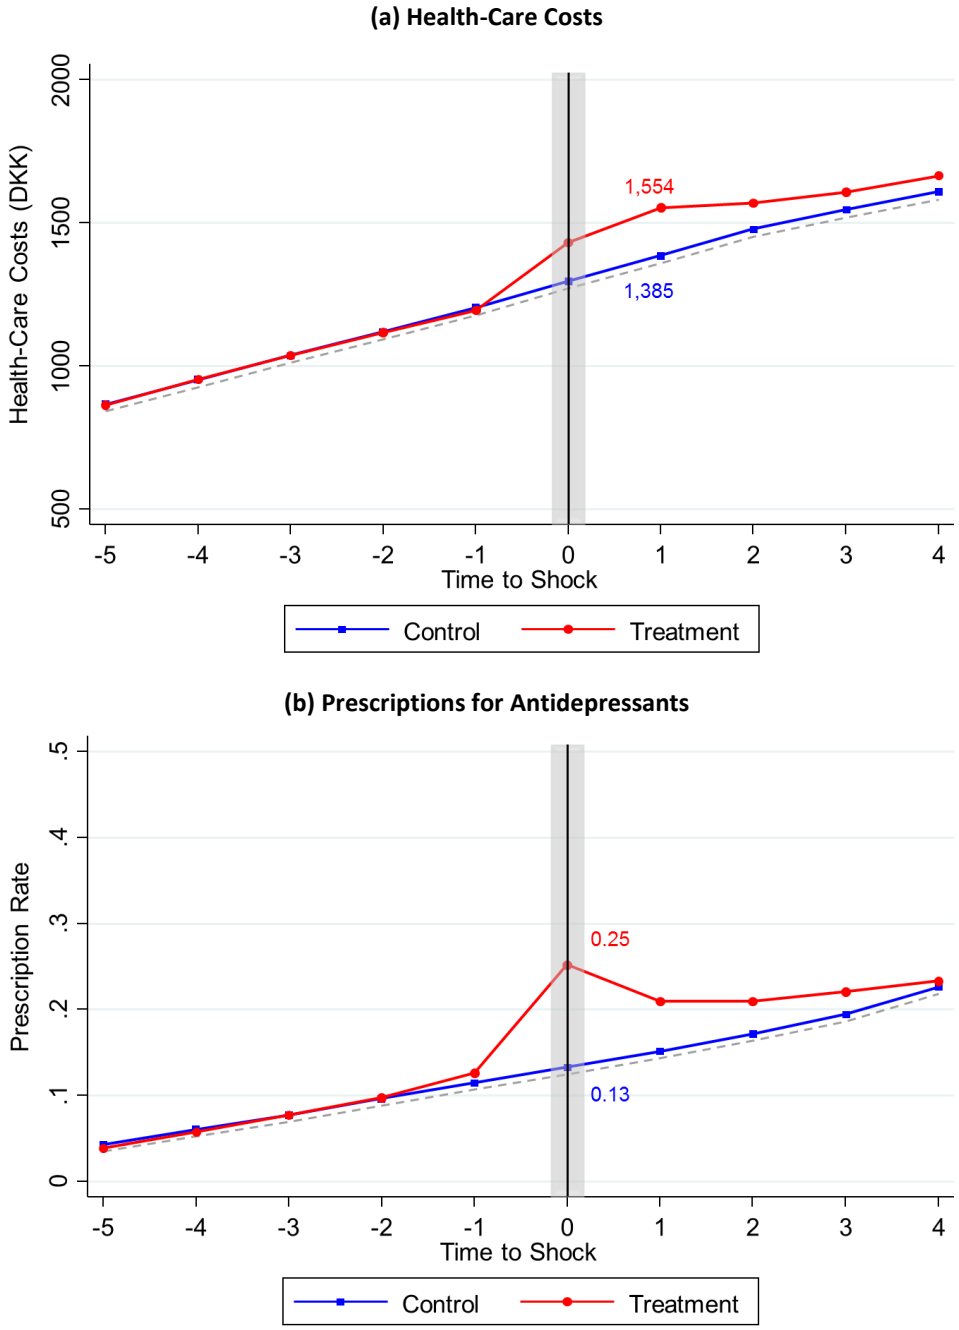
(c) Part-Time Employment



Notes: These figures plot labor supply responses of survivors under age 60 to the death of their spouse. The sample includes individuals under 60 whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A depicts labor force participation; Panels B and C depict the fraction of surviving spouses who are employed full time and part time, respectively. The pictures are constructed from ATP data available for workers under 60. Full-time employment is defined as working at least 30 hours per week all 12 months of the calendar year ("full-time full-year"); part-time employment is defined as working at some point during the year, but either fewer than 30 hours per week or fewer than 12 months within the calendar year. The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group, period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.

ONLINE APPENDIX FIGURE III

Survivors' Health-Care Utilization around the Death of Their Spouse



Notes: These figures plot measures of survivors' health-care use around the death of their spouse. The sample includes individuals born between 1930 and 1950 (for whom we have data on drug prescriptions) whose spouses died between ages 45 and 80 from 1985 to 2011. Panel A depicts overall expenditure on primary medical care, and Panel B depicts the prescription rate for antidepressants (Psycholeptics and Psychoanaleptics). The x-axis denotes time with respect to the shock, normalized to period 0. For the treatment group, period 0 is when the actual shock occurs; for the control group, period 0 is when a "placebo shock" occurs (while their actual shock occurs in period 5). The dashed gray line plots the behavior of the control group. To ease the comparison of trends, we normalize the level of the control group's outcome to the pre-shock level of the treatment group's outcome. This normalized counterfactual is displayed by the blue line and squares. The red line and circles plot the behavior of the treatment group.