Household Labor Supply and the Value of Social Security Survivors Benefits *

David Coyne, *U.S. Department of the Treasury*

Itzik Fadlon, *University of California, San Diego and NBER*

Shanthi P. Ramnath, *Federal Reserve Bank of Chicago*

Patricia K. Tong, *RAND Corporation*

Abstract:

We combine quasi-experimental variation in spousal death and age-eligibility for survivors benefits using U.S. tax records to study the effects on American households' labor supply and the design of Social Security's survivors insurance. Benefit eligibility at the exact age of 60 induces sharp reductions in the labor supply of newly widowed households, highlighting the value of survivors benefits and the liquidity they provide following the shock. Among eligible widows, the spousal death event induces no increases in labor supply, suggesting little residual need to self-insure. Using theory, we underscore the program's protective insurance role and its high valuation among survivors.

* Acknowledgement/disclaimer: This research was supported by the U.S. Social Security Administration through grant #RRC08098400-09 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not necessarily represent the views of the SSA, the U.S. Department of the Treasury, RAND Corporation, any agency of the Federal government, or the NBER. We thank Julie Cullen, Gordon Dahl, Liran Einav, Alex Gelber, Jon Gruber, Henrik Kleven, Amy Finkelstein, Nathan Hendren, Erzo Luttmer, Karthik Muralidharan, Petra Persson, Tommaso Porzio, Kaspar Wüthrich, three anonymous referees, and seminar participants at UCSD, Middlebury College, MIT, University of Michigan, the OTA Workshop, the 2018 NBER Summer Institute, the 2018 NTA Annual Meeting, NBER Public Economics Meeting (Fall 2018), the 2018 Retirement Research Consortium Annual Meeting, the 2018 All-California Labor Economics Annual Meeting, the ASU Annual Empirical Microeconomics Conference, and the Bocconi-CEPR Workshop on Social Security for helpful comments and discussions. Jonathan Leganza provided excellent research assistance.

1. Introduction

The death of a primary earner is among the most devastating shocks that a household can face, and it poses a major source of economic risk for American families. In the U.S., there are approximately 15 million surviving spouses at any given point in time, with 1.4 million newly widowed households each year. The social insurance program that aims to protect against the income losses imposed by this shock—namely, Social Security's survivors benefits—has rapidly grown into one of the largest safety-net programs in the United States. In 2015, the government paid more than $95 billion to 4.2 million surviving spouses (up from $64 billion in 2000), where, by comparison, unemployment benefits and the Earned Income Tax Credit amounted to $35 and $60 billion, respectively (White House 2016; SSA 2018a). Moreover, several proposals to both expand and reform the program are currently under consideration.[1] The significance of social insurance against spousal death in the U.S. is further magnified by evidence of a considerable inadequacy in Americans' life insurance holdings, which have also been declining in recent decades.[2]

Surprisingly, despite the potential importance of the Social Security survivors benefits program to the welfare of vulnerable American households, there is virtually no causal evidence of its economic effects. Specifically, we lack knowledge about the impact of the program's benefits on American families' behavior and economic well-being, the protective role of transfers against spousal death, and the value of liquidity they provide—aspects that are all central for the optimal design of survivors benefits insurance programs.

This paper combines quasi-experimental evidence and theoretical modeling to provide the most comprehensive analysis to date for improving our understanding of the design of survivors benefits insurance programs. We use U.S. tax records from 1999 through 2014 to analyze about a quarter of a million households that have experienced a spousal death. We exploit variation in both the timing of a spousal death and benefit eligibility to study the causal responses of American families along multiple margins. With respect to eligibility for survivors benefits, our research design exploits a sharp discontinuity in the benefit schedule at exactly age 60, providing compelling visual evidence of household responses to the policy. Finally, we combine our empirical analysis with theoretical modeling that utilizes labor supply responses as a well-measured, directly-observable input to individuals' utility that is particularly informative in normative assessments.[3] We map our empirical estimates to various measures of households' willingness to pay for survivors benefits, allowing us to analyze the welfare implications of our findings and the program's design.

[1] E.g., the "Surviving Widow(er) Income Fair Treatment Act of 2018" that was introduced by Senator Robert P. Casey, Jr. on September 18, 2018 as well as various suggested changes at https://www.ssa.gov/oact/solvency/provisions/index.html (Section D).
[2] See Auerbach and Kotlikoff (1987, 1991a,b), Bernheim et al. (2003a,b), Hartley et al. (2018).
[3] For the use of labor supply in normative assessments, see, for example, Shimer and Werning (2007), Chetty (2008), Landais (2015), Hendren (2017), Fadlon and Nielsen (2018), Giupponi (2019), and Wettstein (2020).

Our empirical analysis provides two main sets of findings. First, we identify the causal effect of eligibility for survivors benefits on newly widowed households' behavior. Importantly, conditional on widowhood, eligibility is determined by age, thus creating variation in eligibility status within the widowhood state that allows for shifts in widows' preferences due to the death of their spouse. Using the granularity in birthdate from administrative data and focusing on widows of ages around the sharp eligibility discontinuity at age 60, we provide cleanly identified estimates that are robust to state dependence in preferences. We find clear evidence that eligibility leads to substantial increases in household net income and induces meaningful declines in widows' labor supply. Second, we identify the dynamic causal effects of spousal death by constructing counterfactuals for affected households using households who experience the event in future periods. We show that widows experience small declines in equivalence-scale-adjusted income, with economically negligible increases in labor supply following the death of their spouse.

In the context of our conceptual framework, the welfare implications of these empirical findings can be summarized as follows. The labor supply responses to eligibility point to a high valuation of benefits among newly widowed households. With quadratic labor disutility as one example, the reduction in labor supply maps one-for-one with the excess valuation ineligible widows would assign to $1 of available survivors benefits. These excess valuations amount to 10% among all widows and 37% among low-earning widows, who are potentially more exposed to financial risk. The excess valuation among compliers, who take up benefits upon initial eligibility, amounts to 51%, pointing to a very high valuation of benefits among this policy-relevant population. The age-based variation in eligibility implies these responses are driven by the liquidity effect of immediate cash-on-hand, which has indeed been identified as a key driver of households' valuation of assistance programs (Chetty 2008, Shimer and Werning 2008).[4] As for the coverage level of current benefits, we fail to find a labor supply response to spousal death among age-eligible widows. In the absence of state dependence in preferences, this lack of a response implies that age-eligible households with access to SSA survivors benefits are adequately insured, suggesting limited welfare gains from making the program's benefit levels more generous.

Our findings have two main takeaways. First, we find that households place a high value on the large Social Security survivors benefits program and the liquidity it provides. Second, we find that current benefit levels adequately compensate for the need to self-insure, with potential welfare gains from increased coverage if the cost of supplying labor becomes higher in widowhood. Overall, we highlight the importance of government transfers to vulnerable Americans facing the major economic risk of spousal death.

[4] Some more recent papers, in the context of job displacement and unemployment insurance, study excess sensitivity to cash-on-hand using high-frequency administrative expenditure data and discuss the implications for the design of unemployment benefits—see for example Ganong and Noel (2019) in the U.S. and Gerard and Naritomi (2021) in Brazil.

Relation to Literature. To date, there is a lack of causal evidence on Americans' exposure to financial risk from spousal death and how the associated income loss is compensated. With small-sample survey data, earlier work studied projected losses from hypothetical spousal deaths or assessed changes in financial outcomes between two survey waves of women who became widows.[5] This work aimed to assess Americans' adequacy of life insurance holdings—a key economic question that has remained largely open even though life insurance constitutes one of the largest insurance markets in the U.S. The most related work with respect to widows and survivors insurance in the U.S. is Hurd and Wise (1996), who simulate how widows' poverty would mechanically change (i.e., abstracting from behavioral responses) if they were given higher survivors benefits, and McGarry and Schoeni (2000), who study factors including Social Security benefits that may explain patterns in widows' living arrangements.

The literature is similarly sparse outside of the U.S., with two notable exceptions. Giupponi (2019) estimates the long-run wealth effect of welfare transfers on individual labor supply using benefit cuts in survivors insurance in Italy and a regression discontinuity design. Her findings suggest a high value of additional income among Italian widows of about 50% in excess valuation. Fadlon and Nielsen (2021) exploit variation in event timing to study the effect of spousal death on labor supply in Denmark, where there is no explicit survivors insurance program. They find that widows' earnings increase by about 10% following the shock. This result is similar in magnitude to our findings for age-ineligible widows, suggesting comparable valuations of potential coverage among ineligible households in their study and ours. While our work is related to these two studies, it is also useful to highlight some important differences. Both studies investigate settings in which there is one source of variation, allowing for only one dimension (state of nature or benefit eligibility) to be "held constant" while studying the effect of the other. A clear distinction of our setting relative to past work (on any country) includes quasi-experimental variation that simultaneously exists in the timing of spousal death and in benefit eligibility. As we show, these two margins are key for assessing both the valuation of benefits by vulnerable families (based on responses to benefit eligibility) and the adequacy of the generosity of coverage (based on responses to spousal death). In other words, both assessments are necessary and central for understanding the optimal design of survivors benefits. The simultaneity of our variation also allows us to provide novel statements about state dependence with respect to spousal death. Specifically, we provide evidence that the utility cost of self-insurance through labor supply increases relative to the gain from recovered consumption losses following the event.

[5] See, Holden et al. (1986), Auerbach and Kotlikoff (1987, 1991a,b), Myers et al. (1987), Holden et al. (1988), Hurd (1989), Hurd and Wise (1989), Weir and Willis (2000), Bernheim et al. (2003a), Bernheim et al. (2003b), Sevak et al. (2003), McGarry and Schoeni (2005), Angel et al. (2007).

To summarize, we provide several contributions to past work on survivors benefits in both the United States and the broader international context. First, using administrative IRS tax data, our paper is the first to provide a causal analysis of the expansive Social Security survivors benefits program and widowhood in the U.S. Our research designs allow for clear visual evidence of household responses, leading to novel conclusions about the program's importance to the welfare of American households. Combining our empirical results with theory, we offer a comprehensive analysis of the design of survivors benefits insurance programs. Through the presence of multiple sources of variation, including unique age-based eligibility for benefits, we can provide a rich assessment of the program (in terms of benefit valuation and coverage generosity) that addresses major obstacles facing classic welfare assessments. In particular, we allow for scale economies in household production by relying on individual labor supply, and, notably, we allow for state dependence in preferences in assessing the valuation of benefits among widows. The latter is key since potential changes in preferences across states of nature pose one of the most difficult challenges in assessing insurance inefficiencies and the value of benefits for any type of risk, particularly in the context of spousal death.[6]

The remainder of the paper is organized as follows. Section 2 provides background on widowhood in the U.S. and describes the institutional setting. Section 3 outlines our conceptual framework that allows us to map labor supply moments to expressions of valuation of survivors benefits. In Section 4, we describe the data and lay out our empirical frameworks. Section 5 presents our empirical analysis of the effects of eligibility for Social Security's survivors benefits (in Section 5.1) and the effects of spousal death (in Section 5.2). Section 6 discusses the welfare implications of our findings. Section 7 concludes.

2. Background and Institutional Details

We begin by providing context for our study and outlining the main features of the U.S. Social Security survivors benefits program.

Facts on Widowhood in the U.S. We first briefly describe background information on the prevalence of widowhood in the U.S. using data from the American Community Survey (ACS). The ACS is a representative survey conducted by the U.S. Census Bureau that contains information on households' demographics and economic characteristics.

Appendix Figure A.1 provides a series of facts on widowhood. In any given year, more than 15 million Americans are either a widow or a widower, with a flow of more than 1.4 million newly widowed households each year. In the cross-section of surviving spouses, across ages, surviving female spouses

[6] See general discussions in Finkelstein et al. (2009, 2013), Chetty and Finkelstein (2013), Hendren (2017), Fadlon and Nielsen (2018), and Landais and Spinnewijn (2021). Fadlon and Nielsen (2017, 2018) describe the particular challenges that we overcome in this paper, pointing out that potential changes in preferences from spousal death could occur as a result of lost preference complementarities across spouses and from the significant declines in widows' health (as seen, e.g., in Stroebe et al. 2007).

represent the vast majority of households. In 2016, for example, widows (whose husbands died) accounted for 78% of all surviving spouses, whereas widowers (whose wives died) accounted for the remaining 22%. By age 65, 13.5% of all American females are widows, and this share increases to 21.5% by age 70 (for males, these shares are 4.3% and 6.9%, respectively). Compton and Pollak (2018) additionally estimate that widowhood is expected to constitute a considerable share of one's life cycle. For example, the survivor life expectancy of a 60-year-old wife in a typical American couple in 2010 is more than 12 years if she is the surviving spouse. These numbers underscore that Social Security's survivors insurance program is directly relevant for many households. Indeed, the program pays benefits to more than 4 million surviving spouses every year (SSA 2018a).

U.S. Social Security Survivors Benefits. Surviving spouses become universally eligible for Social Security's survivors benefits at exactly age 60.[7] Historically, the original Social Security Act from 1935 provided only retirement benefits to the worker. The 1939 Amendments made fundamental changes that were designed to extend the provision of a minimum subsistence income for retired workers to their dependents and survivors, where a new category of covered older Americans were widows of age 65 or older. Widows' universal eligibility age declined to 62 in 1956 and then again to the current age of 60 in 1965.[8] The rationale behind providing benefits earlier to widows at age 60 has been that many women were widowed years after having left the labor market to become housewives and mothers. Thus, they may have lacked the skills necessary to qualify for suitable employment. Moreover, widows in their late fifties and sixties were often denied employment because of their age.[9]

To be eligible for survivors benefits, surviving spouses cannot remarry before age 60; otherwise, they lose their entitlement altogether. The benefits that surviving spouses can receive are based on the deceased spouse's potential Social Security retirement benefits, which are determined by the deceased's work history. Specifically, Social Security retirement benefits accrue to individuals whose earnings are subject to Social Security taxes. To become eligible for retirement benefits, individuals are required to accumulate 40 "credits" (where, e.g., in 2016, $1,260 in earnings = 1 credit), which translates to 10 years of work since workers can earn up to 4 credits each year. The retirement benefits aim to reflect lifetime earnings and are based on a worker's Average Income Monthly Earnings (AIME) over the 35 years in which the worker earned the most. The Social Security Administration estimates that the vast majority of workers (about 96.5%) are generally eligible for benefits at retirement age (SSA 2022).

[7] Disabled survivors are eligible for survivors benefits when they reach age 50, and surviving spouses with dependent children under age 16 are eligible for benefits regardless of their own age. The non-zero take-up of benefits by widows prior to age 60 is attributable to these groups.
[8] For details, see https://www.ssa.gov/history/1939amends.htm; https://www.ssa.gov/history/reports/1939no2.html; https://www.ssa.gov/history/benefittypes.html.
[9] See "Social Security Amendments of 1965 Volume 1" at https://www.ssa.gov/history/pdf/Downey%20PDFs/Social%20Security%20Amendments%20of%201965%20Vol%201.pdf

Survivors benefits are then calculated as a percentage of the deceased's potential retirement benefits, and this percentage is determined by the surviving spouse's age at the beginning of benefit claiming. The percentage jumps discontinuously from 0% to 71.5% at exactly age 60, and then continuously rises to 100% at the widow's full retirement age, which represents actuarial adjustments to account for the different lengths of benefit collection when benefits are claimed at different ages.[10]

By design, the present discounted value (PDV) of Social Security's survivors benefits depends only on the deceased's earnings history and does not depend on the survivor's earnings history.[11] This feature is advantageous in that it implies the PDV of survivors benefits is fixed from the point of the husband's death onward. In turn, it implies there are no actual differential substitution effects at eligibility, so the potential impacts of the program on widows' labor supply should operate through a non-distortionary channel. Similar to retirement benefits, however, survivors benefits are subject to an earnings test when claimed prior to full retirement age. If the surviving spouse's labor income exceeds a certain level (e.g., $16,920 in 2017), then benefits are withheld at a specified rate but are later paid back in the form of increased benefits (SSA 2018b). Since research in the context of Social Security retirement benefits has shown that such benefit adjustments may be misperceived as a tax (Liebman and Luttmer 2012, 2015; Brown et al. 2013), the earnings test is a program feature that may create a "substitution" effect. We explore this feature later by studying subsamples of households that are infra-marginal to the earnings test.

It may be useful to discuss other potentially important ages in the vicinity of our age 60 threshold. The first is age 62, which is the early eligibility age for standard Social Security retirement benefits. Note that the claiming of survivors benefits and its timing do not alter widows' schedule of their own retirement benefits, which they can become eligible for and transition to at that age. To account for this threshold, we restrict the analysis to observations of widows younger than 62. Indeed, our monthly-age analysis clearly shows how the effects of survivors benefits kick in at exactly 60 (and we provide figures up to age 70 in the appendix for completeness). A second notable age is 59.5, when withdrawals from private retirement savings accounts are no longer penalized. We note, however, that this has reduced relevance in our context of widows and their overall financial portfolio since the death event itself already allows for non-penalized distributions from the deceased husband's accounts. Still, we can again leverage our high-frequency graphical analysis, which shows that effects kick in promptly at the monthly age of 60. Moreover, we study households who likely do not have access to such retirement accounts and see similar findings.[12] Both of these age considerations (and potential others) are further addressed when we augment the analysis with a

[10] Widows' full retirement ages depend on their year of birth. It is 65 for widows born in 1939 or earlier, and it is increasing by two month increments for every birth year thereafter. Social Security does not notify widows when they become age-eligible for benefits.
[11] This stands in contrast to traditional age-contingent benefit schemes studied in past literature, such as old-age pensions, where own work can directly affect the present discounted value of benefits.
[12] We analyze households who did not make contributions to savings accounts in previous periods in Appendix Table C.3.

control group of future widows and find very similar results. This finding also alleviates concerns that age 60 could be an 'anchor' in the general population (e.g., through norms or workplace retirement policies).

Finally, for our analysis we focus on female surviving spouses, who comprise the vast majority of all widowed households throughout the age distribution (around 80%) and close to 100% of all the program's beneficiaries (e.g., 98% in 2017; SSA 2018c). Using the tax data, we study widows whose husband died between the years 2002-2007. On average, these women were about 5 years younger than their husbands, and, prior to their death (i.e., averaged over years 3 and 2 before the event), these husbands had a 47 percentage points (pp) labor force participation rate, earned about \$4,200 more in annual wages relative to their wives (with an average level of \$24,000), and had a 45 pp likelihood to have started collecting Social Security benefits.

3. Conceptual Framework

In this section, we provide a framework for thinking about how the estimated labor supply responses to benefit eligibility and spousal death together can inform the optimal design of survivors benefits.

3.1. Model of Household Behavior

Setup. We consider the decisions of a two-person household, which consists of member 1 and member 2, in a world with two states of nature: a "good" state, state g, and a "bad" state, state b. In state b, member 1 dies and member 2 becomes a widow.[13] We study the planning problem of households who start in the good state and (similar to the operation of Social Security survivors benefits) would be either age-eligible or age-ineligible for survivors benefits if the bad state occurs. Conditional on starting a period in the good state, the probability of staying in the good state is μ^g and the probability of transitioning to the bad state is μ^b (with $\mu^g + \mu^b = 1$). In what follows, we employ the notation that x_i^s denotes outcome x for household member i, $i \in \{1,2\}$, in state of nature s, $s \in \{g, b\}$. The household's planning problem involves choices over labor supply and consumption allocation (as well as potential insurance purchases and arrangements and use of savings), which can depend on state of nature *and* age-based eligibility for benefits.

[13] For simplicity, we model one period of the absorbing bad state since newly widowed households who differ in benefit eligibility become similarly eligible in subsequent periods. The framework we provide here builds on past literature, and further extensions have been laid out in previous work (such as Chetty 2008 and Landais 2015 in the context of unemployment). Specific to a household-level analysis, Fadlon and Nielsen (2018) provide and discuss extensions and generalizations to the model analyzed here, including a multi-period dynamic life-cycle model (which our results can extend to since they rely on classic Euler conditions that hold more generally), general choice variables and the labor force participation decision, alternative assumptions about the individuals' and the household's preference structure (with an explicit analysis of different types of state dependencies and preference complementarities/non-separabilities), different approaches to modeling the household's behavior (i.e., collective or unitary), and means-testing in government transfers. We can also model borrowing constraints in the dynamic model by setting a lower bound on assets.

Therefore, when relevant, we extend our notation to differentiate between outcomes and choices of eligible households, denoted by $x_i^s(e)$, and ineligible households, denoted by $x_i^s(ne)$.

Household Budget. Denote c_i^s and l_i^s as the individual consumption and labor supply of each member i, respectively. Let $\bar{A}$ represent the household's baseline wealth and non-labor income, and let A^s represent the household's state-contingent wealth and non-labor income. We let state-contingent income be inclusive of premiums to and transfers from any private insurance arrangement, payouts from savings accounts, as well as any informal insurance arrangements across relatives, etc. Denote i's labor income by $z_i^s = w_i l_i^s$, where w_i is the wage rate. We assume that the household pays taxes τ while in the good state and receives a flow of survivors benefits upon transition to the bad state of widowhood. For completeness, we use $b(e)$ to denote the amount received if age-eligible and $b(ne) = 0$ to denote the amount received if age-ineligible, with $b = (b(e), b(ne))$.

Scale Economies. In the presence of scale economies in household consumption (e.g., in utilities such as electricity and heating) a dollar of income in a two-person household can produce more than one unit of the consumption numeraire at the individual level. Since the household's composition changes upon the adverse event in our context of spousal deaths, we need to directly take such consumption technology changes into account in the transition to widowhood. Scale economies essentially create state dependencies that come from differences in technology across the good state and the bad state. We allow for economies of scale in household consumption by introducing a simple consumption technology that transforms household-level income into individual-level consumption (akin to Browning et al. 2013). Specifically, if we denote the household's overall income in state s by y^s, then in the good state we let $c_1^g + c_2^g = \lambda y^g$, where $\lambda \geq 1$ is a scaling factor. This directly maps to the "equivalence scale" measure for scale economies in household consumption, which assesses the share of a household-level income a single individual would need to maintain a given level of consumption as compared to when consuming in a couple. By construction, the equivalence scale is captured by $1/\lambda$.

Household Preferences. Let $U = u_1(c_1) - v_1(l_1) + u_2(c_2) - v_2(l_2)$ represent the household's per-period utility, where $u_i(c_i)$ is member i's utility from consumption and $v_i(l_i)$ is member i's disutility from labor (including the utility loss from direct work costs, labor supply adjustments, and the opportunity costs of lost home production). We later extend the model to allow for state dependence in preferences, i.e., potential changes in spouses' preferences due to the death event. We employ the normalization $u_1(0) = v_1(0) = 0$, which allows us to model the bad state of a spousal death by setting $c_1 = l_1 = 0$ in realizations in which the household is in state b. With these assumptions, the household's preferences in the good state take the form $U = u_1(c_1) - v_1(l_1) + u_2(c_2) - v_2(l_2)$, and the household's preferences reduce in the bad state to the utility of member 2 (the surviving spouse), that is, $U = u_2(c_2) - v_2(l_2)$. We additionally assume that the consumption utility and the labor disutility functions are well-behaved—i.e., that $u_i'(c_i) > 0$, $u_i''(c_i) < 0$, $v_i'(l_i) > 0$, and $v_i''(l_i) > 0$. The household's expected utility at the beginning of the planning period is denoted by $V \equiv E(U)$, with the expectation operator taken over the possible realizations.

Household Problem and Behavior. Starting from the good state, the household maximizes its expected utility, V, subject to the state-contingent budget constraints so that total consumption is bounded by the (potentially scaled) household's net realized income described above. Optimality then implies that the marginal utility from consumption must equate across spouses (when both are alive) and that member 2's marginal utility from consumption must equate to her disutility from supplying labor to produce a unit of consumption (considering wages and the consumption technology). That is,

$$u_1'(c_1^g) = u_2'(c_2^g),\ u_2'(c_2^g) = \frac{v_2'(l_2^g)}{\lambda w_2}, \text{ and } u_2'(c_2^b) = \frac{v_2'(l_2^b)}{w_2}.$$

Otherwise, there are trivial, possible utility-enhancing perturbations.

3.2. Welfare Analysis

Using this framework, we can now derive welfare expressions that map directly to our empirical analysis to assess: (1) the value that widows place on Social Security survivors benefits, and (2) the adequacy of the insurance coverage these benefits provide. These two valuations are necessary for understanding the optimal design of survivors benefits and they require the two different sources of variation we simultaneously have in our setting, namely, both in eligibility for benefits and in the timing of the widowhood event. Accordingly, our first welfare experiment considers transfers across widowed households who differ in benefit eligibility ("holding constant" the state of nature), which maps the labor supply effects of benefit eligibility to the valuation of benefits. Our second welfare experiment considers transfers across households in different states of nature ("holding constant" benefit eligibility), which maps the labor supply effects of spousal death to the degree of insurance coverage.

There are two points worth noting before we proceed. First, using our estimated labor supply responses and under the assumptions we lay out, we can directly speak to several aspects of the value and optimality of the Social Security survivors' benefits program from the point of view of the household. We note, however, that the valuation of benefits from a societal perspective would also need to account for the potential net resource costs.[14] We come back to this point with an illustrative example in Section 6. Second, our revealed-preference approach to welfare relies on households' ability to optimize on the labor supply margin. Hence, potential adjustment frictions may attenuate labor supply responses and lead to an

[14] We emphasize the valuation of benefits, similar to prominent recent studies such as Finkelstein et al. (2019) on Medicaid and Landais and Spinnewijn (2021) on unemployment benefits, since the value added of our analysis lies there. We do not allude to the cost side of changes to the benefit schedule as the Social Security Administration already has mechanisms in place for scoring the cost to the system of various changes to the benefit structure.

underestimation of households' benefit valuation. Such frictions can come, for example, in the form of grief among newly widowed households or the inability to change labor supply or retire at will due to frictions in the labor market. We address this point in the empirical analysis, where the findings suggest that adjustment frictions may have a limited role in our setting. We now turn to our welfare experiments.

Valuation of Survivors Benefits. Consider assessing the utility value of resources to a household that transitions to widowhood. Exploiting variation in benefit eligibility, we can evaluate a reallocation of resources across potential realization paths. Specifically, we consider a household that transitions to widowhood when age-ineligible for survivors benefits. We can assess that household's willingness to trade resources with a counterfactual household that transitions to widowhood when age-eligible for survivors benefits. In the context of our model, this excess valuation of resources from survivors benefits among ineligible widows can be derived through resource-balanced perturbations of b.

Applying the envelope theorem, a widow's willingness to trade resources across eligibility contingencies is captured by the relative gap in the respective marginal utilities from consumption, $\frac{u_2'(c_2^b(ne))-u_2'(c_2^b(e))}{u_2'(c_2^b(e))}$. The intra-temporal optimality conditions imply that this marginal rate of substitution could be alternatively represented by the widow's labor disutility, $\frac{v_2'(l_2^b(ne))-v_2'(l_2^b(e))}{v_2'(l_2^b(e))}$. A first-order expansion then provides us with the following approximation that we define as our first welfare expression:

$$\omega_1 \equiv \varphi(l_2^b(e)) \times \left(\frac{l_2^b(ne)-l_2^b(e)}{l_2^b(e)}\right),$$

where $\varphi(l_2) \equiv \frac{v_2''(l_2)}{v_2'(l_2)} l_2$ is the curvature of labor disutility. The term $\frac{l_2^b(ne)-l_2^b(e)}{l_2^b(e)}$ is the *causal* effect of benefit eligibility on the labor supply of newly widowed households. We investigate the effects of benefit eligibility in Section 5.1 and their welfare implications in Section 6.

To complete the calculation of ω_1, we would also need to recover the preference parameter $\varphi(l_2)$. There are several methods for obtaining the value of φ, such as through calibration or estimation. Most directly, we can assume a particular functional form that possesses desirable features of behaviors related to utility costs. For example, the utility parameter φ equals 1 under the commonly used quadratic labor disutility (of the form $a + bl^2$, $a, b > 0$). For simplicity, we later make this restrictive assumption to derive quantitative normative statements, and we provide various sensitivity analyses.[15] In any case, within our framework, the value of benefits is proportional to the labor supply responses we estimate.

[15] Alternatively, φ could be identified using directly estimable labor supply elasticities (see Fadlon and Nielsen 2018). This approach relies on the idea that the extent to which a household member responds to changes in economic incentives (own wages and income) is directly linked to the rate at which preferences evolve over labor or consumption. This approach is similar to that of Chetty (2006) for estimating risk aversion (i.e., the curvature of consumption utility).

As the welfare experiment shows, our empirical estimate for the causal effect of benefit eligibility upon widowhood carries direct information about the valuation of survivors benefits. Notably, since it "holds constant" the state of nature, our results provide an important contribution to the literature by allowing for welfare evaluations of survivors benefits that are robust to any form of state dependence in preferences.

We note that the age-based benefit eligibility for Social Security survivors benefits lends to a sharp discontinuity in cash-on-hand (i.e., liquidity) among newly widowed women at the precise age 60 cutoff with no discontinuity in their present discounted value of benefits (i.e., income/wealth).[16] Our empirical results of responses to benefit eligibility will therefore offer insights into the value a newly widowed household assigns to the added liquidity from government transfers.

Evaluation of Coverage. Next, consider assessing the insurance value of a dollar in the event that the household experiences widowhood. Exploiting variation in the timing of a spousal death, we can evaluate a reallocation of resources across states of nature. Here, we "hold constant" the status of benefit eligibility. That is, for a given eligibility status, we evaluate the household's willingness to trade resources in the good state for additional resources in widowhood. This evaluation could be achieved in our model by resource-balanced marginal changes in τ and b.

Focusing on age-eligible households, their net willingness to pay for transfers across states of nature is captured by $\frac{u_2'(c_2^b(e))-\lambda u_2'(c_2^g(e))}{\lambda u_2'(c_2^g(e))}$, that is, the relative gap in the corresponding marginal consumption utilities from an additional dollar. We note that in this assessment we need to account for economies of scale in resources since evaluating transfers across states of nature involves changes to the household's composition.[17] The willingness to pay could again be alternatively represented using the widow's labor disutility, $\frac{v_2'(l_2^b(e))-v_2'(l_2^g(e))}{v_2'(l_2^g(e))}$, which has the desirable feature that it is not impacted by scale economies. In the current case of no state dependence in preferences, using a first-order expansion we define our second welfare expression as:

$$\omega_2 \equiv \varphi(\, l_2^g(e)) \times \left(\frac{l_2^b(e)-l_2^g(e)}{l_2^g(e)}\right),$$

[16] See Appendix A.1 for details.

[17] To see how consumption technology plays a role in this assessment we note the following. In applying the envelope theorem, the utility loss from a marginal increase in τ is captured by $\mu^g \lambda u_2'(c_2^g)$, where the monetary loss is scaled by the amount of consumption units it produced (λ). This can finance an increase in survivors benefits of the amount μ^g/μ^b. The household's valuation of an additional dollar of benefits is captured by $\mu^b u_2'(c_2^b)$, so that the overall increase in benefits is evaluated as $\frac{\mu^g}{\mu^b} \times \mu^b u_2'(c_2^b) = \mu^g u_2'(c_2^b)$. The "excess" valuation of making this resources transfer is then $\frac{u_2'(c_2^b)-\lambda u_2'(c_2^g)}{\lambda u_2'(c_2^g)}$. A similar exercise holds among age-ineligible households.

where $\varphi(l_2)$ is the curvature of labor disutility. Here, the term $\frac{l_2^b(e)-l_2^g(e)}{l_2^g(e)}$ is the *causal* effect of spousal death on widows' labor supply. We investigate the effects of spousal death in Section 5.2 and their welfare implications in Section 6.

Estimating the effect of spousal death among eligible households maps to the willingness to pay for higher insurance among potential beneficiaries. Importantly, it offers a direct evaluation of the adequacy of SSA's survivors benefits generosity. The basic intuition for this welfare experiment is that the value of transferring resources from the good state to the bad state can be captured by the degree to which it mitigates the need to self-insure through family labor supply.[18]

State Dependence in Preferences. We can extend the model to allow for state dependence in preferences, where $u_i^s(c_i)$ and $v_i^s(l_i)$ denote member i's utility from individual consumption and disutility from labor is state $s \in \{g, b\}$, respectively. With this preference structure, the evaluation of coverage requires an additional parameter that captures the degree to which marginal assessments of choices (consumption or labor) change across states of nature. In our setting of labor supply, the relevant local parameter that enters the modified welfare formula can be defined as $\theta^v(l) \equiv \frac{v_2^{b\prime}(l)}{v_2^{g\prime}(l)}$, which measures the extent to which the marginal cost of spousal labor supply varies across states of nature.

To see this, consider again age-eligible households' willingness to pay for transfers across states of nature, but now with state-dependent preferences: $\frac{u_2^{b\prime}(c_2^b(e))-\lambda u_2^{g\prime}(c_2^g(e))}{\lambda u_2^{g\prime}(c_2^g(e))}$. When cast in terms of labor disutility, this willingness to pay becomes $\frac{v_2^{b\prime}(l_2^b(e))-v_2^{g\prime}(l_2^g(e))}{v_2^{g\prime}(l_2^g(e))}$, which leads to the following expression using a first-order expansion:

$$\omega_2' \equiv \theta^v\left(l_2^b(e)\right) \times \varphi\left(\, l_2^g(e)\right) \times \left(\frac{l_2^b(e) - l_2^g(e)}{l_2^g(e)}\right) + \theta^v\left(l_2^b(e)\right) - 1$$
$$= \theta^v(l_2^b(e)) \times \omega_2 + \theta^v(l_2^b(e)) - 1.$$

To gain intuition for how state dependence would impact the welfare assessment of coverage, suppose there are no labor supply responses to a spouse's death (i.e., $\frac{l_2^b(e)-l_2^g(e)}{l_2^g(e)} = 0$ so that $\omega_2 = 0$) and assume that widows have an increased marginal disutility from working after their husbands die (i.e., $\theta^v(l_2^b(e)) > 1$). In this case, even though there are no increases in self-insurance via labor supply, there would still be gains from providing more generous social insurance to widows against the preference shock since supplying labor has become more costly.

[18] This within-eligibility evaluation captures the value of additional coverage from both an income effect and higher liquidity.

3.3. Summary

We have shown how the labor supply responses we will estimate in Section 5 can inform the optimal design of survivors benefits. Our unique setting that includes two sources of variation, in benefit eligibility within the bad state and in the state of nature, allows us to provide two key welfare assessments of both widows' valuation of benefits *and* the generosity of SSA's survivors insurance.

4. Data and Empirical Framework

4.1. Data Sources and Variable Definitions

Data Sources. We use administrative tax records on American households for the years 1999 through 2014. The data include both information returns filed by third parties (e.g., Form W-2, Form SSA-1099, and Form 1099-R) and self-filed income tax returns (e.g., Form 1040). We observe exact dates of death (to identify spousal death events) and exact dates of birth (to determine age-eligibility of survivors benefits for widows) using merged Social Security Administration (SSA) records. Spousal linkages are established through jointly filed tax returns in the year prior to the death event.

From the information returns, we extract wage earnings (using Form W-2), Social Security benefits paid from the retirement and the disability trust funds (which are reported separately on Form SSA-1099), unemployment benefits (using Form 1099-G), and distributions from pensions, annuities, retirement plans, individual retirement accounts (IRAs), and insurance contracts (as reported on Form 1099-R). From the income tax returns, we extract Adjusted Gross Income (AGI). Among other sources of income, AGI includes earnings, capital income, retirement income, and taxable Social Security benefits.

Outcomes and Variable Definitions. Our analysis focuses on widows' labor supply behavior. Based on data from Form W-2, we study wage earnings as our primary outcome of interest.[19] Wage earnings comprise an aggregate measure of labor supply that captures all responses on both the intensive and extensive margins. For completeness, we also provide an analysis of labor force participation, where we define participation as having positive earnings in a given period.

In our high-frequency study of the effects of benefit eligibility, we also analyze the rate of change in earnings to capture an overall measure of changes in 'work intensity.' The benefit of using this outcome is that, as a measure of changes, it provides a clear visual illustration of responses precisely at the age of eligibility. We describe in Appendix B our exact method for constructing this variable to avoid division by zeros when analyzing earnings changes from one period to the next. Our definition additionally has the appealing property that it is 0 if there are no changes in earnings, and it is -1 if a person stops working altogether. We also provide figures for retirement behavior (or change in participation) using a conventional

[19] Annual income from self-employment is very low among widows and is therefore not a meaningful margin for responses in our setting. Nevertheless, we report the analysis of self-employment in Appendix E.

measure, where retirement is defined as having positive earnings in the current period and no earnings in the next period.[20] Since this outcome is also capturing changes, it is likewise valuable for illustrating the immediate responses at eligibility.

To further shed light on how the program impacts financial well-being, we analyze the household's overall income. We define this outcome as the net pre-tax family income available from any reported source, which broadly follows the convention in the literature that uses U.S. federal income tax records (see, e.g., Chetty et al. 2014). For income tax filers, this measure includes AGI, tax-exempt interest, and nontaxable Social Security income; for non-filers, this measure includes wages, unemployment benefits, and gross Social Security income, as well as taxable distributions from retirement savings accounts. As such, family income includes labor earnings, capital income, unemployment benefits, and any payments from Social Security (including retirement, survivors, or disability benefits) or retirement accounts.

4.2. Empirical Framework

4.2.1. Effects of Eligibility for Social Security Survivors Benefits

Research Design. To identify the causal effects of eligibility for Social Security survivors benefits on widowed households, we exploit the age discontinuity in the program's benefit-eligibility schedule. Specifically, we estimate sharp breaks in levels and trends of widows' post-shock outcomes at the exact eligibility cutoff age of 60. Importantly, we allow for smooth underlying trends in widows' outcomes to account for any changes that are continuous in age and would, therefore, not affect the interpretation of our results. One specific implication is that the estimation is not confounded by potential changes in preferences occurring from a spousal death as long as those changes are continuous in the widow's age at the event since all households are analyzed after their exposure to spousal death and its direct impacts.

The population-level data allow us to lead our analysis with a high-frequency graphical representation of the results, which we further use to guide our estimation strategy. We take advantage of survivors' exact dates of birth and plot raw means of each outcome variable of interest against the widow's monthly age. To focus on the eligibility cutoff of age 60, we plot outcomes of widows who, at the time of observation, were between ages 55 and 62 (where 62 is the early eligibility age for standard retirement benefits). We provide figures that extend the age range up to age 70 in the appendix (Appendix Figure C.1).

Estimating Equation. Since tax information is observed as of December for a given year, monthly age is defined as a person's age at the end of the calendar year of observation. The data's annual frequency and the utilized variation in monthly age at the end of a calendar year imply that the effect of being "fully exposed" to eligibility for survivors benefits in the year is captured when widows are eligible for benefits

[20] This definition follows the literature on retirement and old-age government transfers (see, e.g., Coile and Gruber 2007, Deshpande et al. 2021).

for the entire calendar year. Widows who turn 60 at the beginning of the year (and whose age in data is just under 61) are eligible for benefits throughout an entire calendar year and would display the full-exposure effect. That is, the effect of full exposure to eligibility for survivors benefits is identified when widows are just below age 61.

Accordingly, we quantify the full-exposure effect of benefit eligibility using the following equation:

$$y_{i,t} = \beta_0 + \beta_1(age_{i,t} - 60) + \beta_2 \mathbb{I}\{age_{i,t} > 60\} + \beta_3 \mathbb{I}\{age_{i,t} > 60\} \times (age_{i,t} - 60) + \varepsilon_{i,t}. \quad (1)$$

In this regression, $y_{i,t}$ denotes an outcome for widow i at time t, $age_{i,t}$ represents the widow's age in months, and $\mathbb{I}\{age_{i,t} > 60\}$ is an indicator variable that assumes the value 1 if the widow is observed at an age older than 60 in terms of monthly age and the value 0 otherwise. We estimate this equation using the sample of widows between ages 55 to 61, and we include two separate linear trends in outcomes: one for observations before and one for observations after the eligibility age of 60. Our choice of the parametric assumptions in equation (1) is strongly guided by the graphical analysis of the raw data. Accordingly, we combine our graphical and regression analyses in figures that plot raw means of outcomes by widow's monthly age, while superimposing the regression lines from equation (1).

In this specification, β_0 captures a baseline level, and β_1 captures an underlying trend. We estimate the treatment effect of benefit eligibility as the full-exposure impact by age 61, which equals:

$$\beta_2 + \beta_3 \times (11/12).$$

That is, the estimator is composed of sharp behavioral changes around the eligibility-age cutoff, which come from both a break in levels (captured by the change to the intercept, β_2) and a break in trend (captured by the change to the slope, β_3), multiplied by $11/12$ to reflect exposure to eligibility for the full year. The treatment effect will be visually represented in the figures. Counterfactual levels, which will also be visually represented in the figures, will be estimated using $\beta_0 + \beta_1 \times (11/12)$. Finally, we augment the design with a control group of future widows as a robustness check.

Analysis Sample. We study all widows in the tax records whose husband died between 2002-2007, comprising 293,857 widows. We analyze the outcomes of newly widowed households in the periods just after the event occurs. In choosing these periods, we must consider that the data are annual and measure values at the end of a calendar year, so that the year of the death event (period 0) is a transitional period since households experience the husband's death at different points during the calendar year. The first full period in which all sample households have been subject to the spousal death event is, therefore, period 1. We also include period 2 in the analysis for increased statistical power and visual clarity, though the results remain the same with only period 1 (see Appendix Table C.2).

As a point of reference, we provide labor supply statistics for American women near retirement age in Appendix Table A.1. Specifically, we provide averages of annual wage earnings, labor force participation, and the share of individuals with earnings lower than the earnings test for ages 55-59 among all women in the U.S. (using the ACS) and among our sample of newly widowed women. The groups show very comparable levels of baseline labor supply.[21]

4.2.2. Effects of Spousal Death

Research Design. The ideal experiment for identifying the economic impacts of spousal death would randomly assign mortality events to families and track the responses of surviving members over time. Therefore, in practice, we need to compare the ex-post responses to the event by affected households to the counterfactual behavior of ex-ante similar unaffected households. To utilize a quasi-experimental research design that approximates this experiment, we restrict the analysis to households that have experienced a spousal death at some point in our sample period, and we identify the treatment effect using variation in the timing of when the event was realized. Specifically, we construct counterfactuals for affected households using households that experience the same event but a few years in the future. As such, our two experimental groups consist of a treatment group, composed of spouses in households that experience the event in year τ, and a matched control group (based solely on event time), composed of spouses from the same cohorts in households that experience the event but in year $\tau + \Delta$. Correspondingly, we assign a placebo event to control households at year τ. We then recover the treatment effect by using a standard dynamic difference-in-differences estimator, identifying the event's impact purely from the change in the differences in outcomes across the two groups over time.

The identifying assumption is that, absent the realization of the event, the outcomes of the treatment and control groups would run parallel. To test its validity based on the necessary condition of parallel pre-event trends, we present the evolution of the treatment and the control groups' outcomes in the periods prior to the (actual or placebo) event year. We show there are virtually no differential changes in the trends of the treatment and control groups before period 0 and that, in the few instances where pre-trends are statistically significant given the large sample size, they are economically negligible.[22]

[21] These patterns can also speak to the generalizability of our findings to the context of retirement benefits. Albeit suggestive, the analysis could be informative for the discussion regarding the possible responses to reforming the Social Security Early Eligibility Age (EEA) since the eligibility for benefits among widows at age 60 is the only source of variation in early eligibility since its introduction. In addition, we find that working widowed women exhibit a considerable increase in retirement rates in response to survivors benefits. With the significant growth in female labor force participation at older ages in the U.S. (Goldin and Katz 2018) and the meaningful share of widows among older American women, the evidence suggests that the Social Security survivors benefits program itself could play an increasingly important role in female retirement behavior.

[22] For more details on the validity of this research design, its identifying assumptions, and the trade-offs in the choice of Δ, see Fadlon and Nielsen (2019, 2021).

Estimating Equation. We normalize the time of observation such that period n is measured with respect to the assigned event year; that is, $n = year - \tau$, where $n = 0$ is when the assigned event occurs. Our estimating equation is of the following dynamic difference-in-differences form:

$$y_{i,t} = \alpha + \beta treat_i + \sum_{n \neq -1; n=-3}^{3} \gamma_n \times I_n + \sum_{n \neq -1; n=-3}^{3} \delta_n \times I_n \times treat_i + \lambda X_{i,t} + \varepsilon_{i,t}, \qquad (2)$$

where $y_{i,t}$ denotes an outcome for household i at time t, $treat_i$ denotes an indicator for whether a household belongs to the treatment group, and I_n are indicators for time relative to the assigned event year. The key parameters of interest are δ_n, which estimate the period n treatment effects ($n > 0$) relative to the pre-period $n = -1$.[23] $X_{i,t}$ is a vector of controls that includes age fixed effects and calendar year fixed effects. We report robust standard errors clustered at the household level. For visualizing our empirical strategy, we provide a graphical analysis of the raw data for the outcomes we study.[24]

To quantify mean treatment effects, we also estimate the standard difference-in-differences equation of the following form, which averages over years before and after the event:

$$y_{i,t} = \alpha + \beta treat_i + \gamma post_{i,t} + \delta treat_i \times post_{i,t} + \lambda X_{i,t} + \varepsilon_{i,t}. \qquad (3)$$

In this regression, $post_{i,t}$ denotes an indicator for whether the observation belongs to post-event periods. The parameter δ represents the average effect of spousal death.

Analysis Sample. We construct our sample as follows. We randomly draw 20% of all men who died between the years 2002 and 2007 and were married in the year prior to their death. Our analysis sample is comprised of their surviving widows. We then restrict this sample of widows to those who were between the ages of 50 and 70 in the assigned event year (the actual event year for the treatment group and the placebo event year for the control group) to pivot around the eligibility age of 60.

Based on the time range of the data and to allow for a sufficient number of years before and after the event, our choice of Δ is four years, i.e., our research design matches households who experience the event 4 years apart. As a result, our treatment group is composed of women whose husband died between 2002-2003, and our control group is composed of women whose husband died between 2006-2007. Setting $\Delta = 4$ implies that we can estimate the effects of spousal death up to three years after the event (since the control group becomes "treated" 4 years out). Our sample includes 63,707 households in the treatment group and 74,214 households in the control group. In the baseline period (the year before the assigned event year), wives are, on average, 60.4 years old in the treatment group and 60.2 years old in the control group, where the mean calendar year is around 2001.5 for both groups.

[23] Testing for parallel pre-trends is based on estimates of δ_n for $n < 0$, which we report in Appendix Table D.1.

[24] In Appendix Figure D.1 we also provide a version of the figures for the dynamic effects of spousal death that split observations around the benefit eligibility age of 60.

5. Empirical Evidence

We now turn to our core empirical analysis. In this section, we first estimate the causal effects of eligibility for Social Security's survivors benefits on widows' labor supply and household income and then analyze the impact of spousal mortality on American families. In the next section, we will analyze the welfare implications of these estimated causal effects and accordingly calculate the moments that enter our welfare formulas.

5.1. Effects of Eligibility for Social Security Survivors Benefits

Benefit Claiming. We begin by looking at the claiming behavior of survivors benefits by newly widowed women as a first stage in the analysis. Panel A of Figure 1 plots the take-up rate of benefits from Social Security. The structure of this and subsequent figures is as follows. The x-axis denotes the age of the widow in months (at the end of the calendar year of the observation), and the y-axis denotes the behavior of the outcome of interest. The circles represent means of the raw data at each monthly-age bin. The solid lines plot the piecewise linear fit using equation (1). The dashed line in the age range 60-61 represents the counterfactual behavior in the absence of eligibility for survivors benefits based on specification (1), which extrapolates the linear relationship estimated on observations prior to age 60. Eligibility for benefits begins at exactly age 60 (which is marked by the vertical dashed line). The full-exposure effect of benefit eligibility is represented by the vertical gap between the solid and the dashed regression lines at age 61 (which is marked by the vertical solid line). The estimates for the full-exposure effects and their standard errors are reported in the figures along with the outcome's counterfactual level. The estimates are also summarized in Appendix Table C.1.

Panel A of Figure 1 clearly shows a jump in the take-up of benefits by just-eligible widows at the cutoff age 60. By age 61, the full-exposure effect amounts to a 51 percentage point increase in claiming. The corresponding pattern in benefit amounts is displayed in panel B of Figure 1. The trend in benefit levels breaks exactly at the cutoff age as the increase in claiming begins and reaches its full effect by age 61. At that point, the average increase in benefits, including zeros for those not claiming, amounts to $5,605.

The overall patterns of benefit claiming can be summarized as follows. Recall that disabled survivors are eligible for benefits when they reach age 50 (as well as surviving spouses with dependent children under age 16) and that survivors benefits are calculated as a percentage of the deceased's potential retirement benefits, where this percentage jumps discontinuously to 71.5% at age 60 and then continuously rises to 100% at the widow's full retirement age. Also, recall that SSA estimates the vast majority of workers (about 96.5%) will be eligible for retirement benefits, so we would generally expect widows of eligible ages to be entitled to some amount (where the amount would depend on the deceased husband's work history). These features together create the patterns we observe: low levels of claiming prior to age 60 of about 11 pp (among widows who are disabled or have dependents), a jump in the discontinuity region of age 60 that amounts to an effect of 51 pp, and then a steady rise as generosity continuously increases when widows age into their full retirement age. By age 62, the take-up of survivors benefits among all widows amounts to about 70 pp, and thereafter take-up can also begin to reflect widows who claim their own retirement benefits.[25] These first-stage patterns establish our research design, and we now turn to our outcomes of interest.

Household Income. To evaluate the impact of eligibility for survivors benefits on widows' overall financial well-being, we analyze our comprehensive measure of net household income. Panel C of Figure 1 reveals a clear break in the trend in overall household income exactly at the point where widows are just-eligible for Social Security's survivors benefits. Benefit-eligible widows' annual income then increases at a rapid rate over the initiated eligibility range until it reaches the full-exposure effect, as displayed by widows of age 61. The net increase in income totals $4,799, which represents an 11.4% increase in family income (on a counterfactual of $42,127). Scaling the effect of eligibility on net income by using the claiming rate, the effect of benefit receipt on the sample of compliers who take up the program upon initial eligibility amounts to $9,345 (=$4,799/0.51351). Appendix F further constructs the counterfactual level for this subsample, which we estimate to be $31,318. Hence, the treatment effect of benefit receipt on net household income among compliers represents an increase of 30%.

Labor Supply. Next, we turn to our main outcome and investigate how benefit eligibility affects the labor supply of widows. For visual evidence, highlighting that households respond immediately when eligible, we first plot our measures of changes in labor supply. Panel D of Figure 1 displays a clear and considerable jump in widows' retirement rate upon benefit eligibility (with a full effect of 1.8 pp on a counterfactual of 5.7 pp). Similarly, panel F of Figure 1 shows a clear break in work intensity patterns exactly at the start of eligibility with substantial follow-up declines, which then flatten back after full exposure. The full-exposure effect on the rate of change in earnings is −4 pp, relative to a counterfactual rate of −2.8 pp. To evaluate the cumulative labor supply effect as our primary outcome of interest, we study the levels of widows' labor force participation and wage earnings. The full exposure effect on labor force participation amounts to a decline of 2.9 percentage points (panel E of Figure 1). Overall, widows' total labor supply responses correspond to an average decrease of $1,751 in annual earnings (panel G of Figure 1).

[25] To further investigate the take-up behavior, we also split the sample by the degree of household labor market "specialization" in the periods prior to spousal death. We calculate the pre-event share of households' overall earnings that is attributable to the husband and study the effects of eligibility on widows in households where this share was low (below 0.25) and those where this share was high (above 0.75). The idea is that, given the structure of the benefit schedule that depends on the deceased's earnings history, widows in households where the husband was a primary earner will have a stronger incentive to claim survivors benefits (since the relative share of benefits out of the overall household income post-event would be higher). Indeed, Appendix Table C.4 shows that these households have significantly higher take-up rates and receive higher average benefits from the program, which is consistent with widowed households responding to financial incentives.

Again, it is valuable to convert these responses into the effect of benefit receipt and focus on the group of compliers. The effect on their overall labor supply, as captured by responses in wage earnings, translates to a decline of $3,410 (=$1,751/0.51351). Given that we estimate their average counterfactual level of earnings to be $10,050 (see Appendix F), the response among compliers represents a decrease of 34% in labor supply.

Low-Earnings Households and the Earnings Test. Under standard preferences, declines in labor supply among those eligible for benefits are always favorable from the point of view of a single household and therefore represent an important component of the gains from government programs. However, the overall net welfare consequences from the social planner's perspective depend on the degree to which our estimated labor supply responses represent a liquidity effect versus a substitution/moral hazard effect (see, e.g., a discussion in Chetty 2008). As discussed earlier, unlike Social Security retirement benefits, survivors benefits are generally decoupled from the beneficiary's own labor supply, so there is a lack of direct distortions in financial incentives to work upon eligibility. In that sense, the estimated effect on widows' labor supply could be attributed to a welfare-beneficial liquidity effect. Nonetheless, since research has suggested that individuals may misperceive earnings tests as distortionary income taxation, we also analyze the subsample of infra-marginal households that earn below the corresponding amount thresholds. Appendix Table C.2 provides the results for widows whose pre-shock earnings were below the annual earnings test thresholds, where we similarly find meaningful declines in overall labor supply. As there is likely no moral hazard component in their responses, it points to a meaningful non-distortionary increase in the consumption of leisure.

We also investigate the effect of benefit eligibility on family income for this subsample of low-earning households. Low-earning spouses are likely more exposed to financial risk since they generate little income on their own, as suggested by Fadlon and Nielsen (2017). We find that, with a claiming rate of 60 pp, they receive $7,258 in annual benefit amounts. The increase in their net income totals $7,085, which represents an increase of 21% (on a counterfactual of $34,061).

Frictions and Liquidity. As noted earlier, frictions causing delays in labor supply adjustments could attenuate the observed responses and affect the degree to which just-widowed households reveal their valuation of benefits. To investigate the presence of such frictions, we study women who transitioned to widowhood well in advance of the eligibility cutoff, giving them more time to respond.[26] Among these widows, we find labor supply responses of similar magnitudes, suggesting that labor market frictions are less of a concern in our setting (see Appendix Table C.5).

[26] Specifically, we use observations from periods 6-10 after the spousal death so that, among all the included observations, the husband had died at least 5 years in the past. This design means, for example, that observations at the critical age of 60 are comprised of 60-year-old widows whose husband had died when they were between the ages of 50 and 54.

Finally, it is worth highlighting the observed lack of anticipation in labor supply responses as widows approach the age-eligibility threshold. Along with discontinuous responses to the anticipated age threshold for benefit eligibility, these patterns are at odds with frictionless benchmark models of intertemporal optimization. This is particularly the case for widows who experienced the spousal death event several years prior to eligibility age who were given time to financially recover from the initial shock of losing their spouse (as theoretically laid out in Appendix C.1). Their observed labor supply patterns imply that at least one of the underlying assumptions of the benchmark model is meaningfully violated, suggesting either that households are extremely myopic and exhibit a lack of financial planning or that households are severely liquidity constrained and may be unable to borrow against their anticipated increase in benefits.[27] Among these widows, we find evidence suggesting that liquidity constraints can rationalize the results. In particular, we proxy for liquidity using lagged unearned income (including capital income), and we study whether widows' labor supply responses vary by this measure. We find a strong gradient of labor supply responses with respect to household liquidity, where the highest-liquidity households do not exhibit labor supply responses and behave as the frictionless model predicts (see Appendix Figure C.2 and Appendix Table C.6).[28] These results are consistent with low-liquidity households being unable to smooth the consumption of leisure prior to the actual receipt of the anticipated benefits and with high-liquidity households using their own resources to smooth behavior in anticipation of future income flows.[29]

Robustness. To address potential confounding changes around our cutoff age, we augment our design with a control group of future widows. We include in the treatment group observations of widowed households from periods 1 and 2, and we include in the control group observations of future-widowed households from periods -2 and -1. To guarantee the comparability of calendar years across the treatment and control groups, we focus the analysis to a subset of our original treatment group, which means that estimates should align across designs albeit imperfectly. Reassuringly, the findings are similar (see Appendix Table C.2).

[27] Such absence of anticipation has been also observed in the context of unemployment insurance (Ganong and Noel 2019, Gerard and Naritomi 2021) in anticipation of a decrease in benefits, where individuals could set aside savings for smoothing consumption.

[28] This is in the spirit of liquidity/borrowing constraints tests in Landais and Spinnewijn (2021), who study variation in consumption responses to unemployment and find that consumption drops are particularly sensitive to the level of liquid wealth.

[29] We find suggestive evidence inconsistent with widows exhibiting complete myopia and lack of financial planning. As described earlier, to be eligible for Social Security's survivors benefits widows cannot remarry before age 60. We find evidence in support of strategic timing of remarriage (in the form of increased marriage rates just after age 60) and that those who potentially strategically remarry themselves exhibit meaningful declines in labor supply at eligibility (see Appendix Figure C.2 and Appendix Table C.6). Moreover, observed patterns of take-up of benefits points to anticipation and planning. Widows who experienced the spousal death years prior to eligibility exhibit high take-up rates despite the fact that they are not notified by Social Security once they turn age-eligible for benefits at age 60. In addition, we provided evidence earlier that is consistent with widows' responsiveness to financial incentives, whereby widows who lose a primary earner and have stronger incentives to claim survivors benefits at eligibility indeed take up the program at higher rates.

5.2. Effects of Spousal Death

Dynamic Impacts. Figure 2 displays the effects of a husband's death on the evolution of several household outcomes. The structure of this figure is as follows. The x-axis denotes time with respect to the event year, normalized to period 0. For the treatment group, period 0 is when the actual event occurs; for the control group, period 0 is when a placebo event occurs (while their actual event occurs in period 4). The dashed gray line plots the behavior of the control group (along with the corresponding 95-percent confidence intervals). To ease the comparison of trends, from which the effect is identified, we normalize the level of the control group's outcome to the pre-event level of the treatment group's outcome (in period $n = -1$). This normalized counterfactual is displayed by the solid gray line with square markers. The purple line with circle markers plots the behavior of the treatment group (along with the corresponding 95-percent confidence intervals). Since period 0 is a transitional year, we consider period 1 as the initial period that captures a "full" impact of the event. In each panel, we report the treatment effect averaged over periods 1-3 based on the δ_n's from equation (2), as well as the corresponding counterfactual level. Appendix Table D.1 reports the full set of dynamic estimates for δ_n.

Household Income. We first assess households' financial well-being by looking at the effect of spousal death on overall net household income. This composite measure captures the direct decline from lost husband income and changes through the compensating income sources reported in the tax data, including transfers from Social Security. Panel A of Figure 2 illustrates a prompt and persistent decline in household net income following the event. We find that households' overall income declines by $25,090 in the post-event years. With a counterfactual household income of $70,803, spousal death leads to a decline in income on the order of 35%. A useful benchmark for comparison is the square-root equivalence scale which implies that income declines on the order of 29% would still suggest full income compensation.[30]

Labor Supply. Next, we proceed to analyze the effect of spousal death on widows' labor supply. In panel B of Figure 2, we analyze the surviving spouses' labor earnings. In the overall sample, we find that, while widows increase their labor supply following their spouse's death, the average increase is negligible. Wage earnings, including zeros for those who do not work, display an average increase of $358 in response to the death of a spouse (on a baseline of $11,583). This finding suggests that, in the overall sample, households have a negligible residual need to self-insure when experiencing a spousal death.

Labor supply responses to spousal death may be again attenuated in the presence of labor market frictions. As an intuitive test to gauge the degree to which such frictions may play a role in our setting, we study heterogeneity in labor supply responses between widows who were employed versus not employed

[30] Full income compensation (equating equivalence-scale-adjusted income levels across states of nature) and full insurance (equating marginal utility across states of nature) are naturally not the same when preferences are state dependent. We later leverage the nature of this relationship to assess state dependencies in preferences upon spousal death.

prior to their spouses' deaths. Intuitively, employed individuals could likely adjust their labor supply more quickly than those who were not employed, as the latter group would need to engage in job search. In fact, in Appendix Table D.2, we find that widows who were not employed prior to their spouse's death exhibit a meaningfully larger labor supply response compared to widows who were employed prior to the shock. In line with our findings from Section 5.1, this pattern is consistent with greater exposure to financial risk among lower-earning spouses and, thus, a greater need to self-insure. For the purpose of our investigation, it again suggests there is limited evidence of adjustment frictions attenuating the results.

Finally, for completeness, the evolution of additional household outcomes around spousal death is provided in panels C-E of Figure 2. These include widows' labor force participation as well as Social Security benefit claiming and average amounts.

6. Welfare Implications

To summarize our main results: first, we find significant labor supply decreases among widows in response to eligibility for Social Security survivors benefits and the liquidity these benefits provide; second, we find economically negligible increases in labor supply in response to spousal death. We now combine our theoretical framework (from Section 3) with our empirical analysis (from Section 5) to study the welfare implications of our findings.

Valuation of Survivors Benefits. To quantify the value of survivors benefits, we return to our welfare measure from Section 3.2 of the excess valuation of benefits among widows. This is captured by ω_1, which is proportional to the labor supply response to benefit eligibility, $\frac{l_2^b(ne)-l_2^b(e)}{l_2^b(e)}$, that we estimated in Section 5.1.

If we begin by calibrating the utility parameter φ to 1 (as in the particular case of quadratic labor disutility), our findings suggest that the excess value of an additional dollar to ineligible widows is $\frac{l_2^b(ne)-l_2^b(e)}{l_2^b(e)} = \frac{\$1{,}751}{(\$18{,}787-\$1{,}751\ )} = 10.3\%$. That is, each additional \$1 made available to a widowed household generates 10 extra cents in welfare gains. Notably, among low-earning households, the response in labor supply is significantly larger and amounts to 36.6% (=\$1,065/(\$3,978-\$1,065)). That is, low-earning widows assign an even greater value to government benefits, consistent with the notion that spouses who had generated little income on their own are more exposed to financial risk. Lastly, among compliers, for whom the difference in benefits received between ineligible and eligible households is largest by construction, the excess valuation of a dollar of benefits by ineligible widows amounts to 51% (=\$3,410/(\$10,050-\$3,410)). This finding points to a very high value of the liquidity that survivors benefits provide to the policy-relevant population of widows who take up the program upon eligibility.

As a sensitivity analysis, panel A of Appendix Table G.1 provides imputations of these welfare effects for a range of values for the labor-disutility curvature parameter, φ, which we vary from 0.5 to 2 in 0.25 increments.

Degree of Coverage. We have shown that estimating the effect of spousal death among eligible households maps to their willingness to pay for more generous survivors insurance. To assess the adequacy of coverage, we now use our welfare measure ω_2 from Section 3.2, which is proportional to widows' labor supply responses to their spouse's death, $\frac{l_2^b(e)-l_2^g(e)}{l_2^g(e)}$.

We estimate an augmented version of equation (3) for a narrow age band around the eligibility cutoff (ages 58-61) and let the average treatment effect depend on benefit eligibility to study the effect of spousal death among age-eligible widows. We define a dummy variable for eligibility as the widow being 60 or older, which we interact with the main term $treat_i \times post_{i,t}$. Finally, we include age dummies and their interactions with $post_{i,t}$ and with $treat_i$ separately. The results, given in Table 1, show no evidence of a labor supply increase among eligible households (if anything, we find a slight decrease). This lack of response suggests that widows who are eligible for survivors benefits upon spousal death are adequately covered in terms of income flows and the liquidity these flows provide, meaning they have no additional need to self-insure. In other words, for all values of φ we have that $\omega_2 \cong 0$ among these households, which implies there would be no meaningful gains to households' welfare from expanding the current generosity levels of Social Security's survivors insurance for age-eligible widows.

Table 1 additionally finds that widows who are age-ineligible for survivors benefits when their husbands die increase their labor earnings by about 9% (=$1,290/$14,909). This differential labor supply response to spousal death between eligible and ineligible households can be used to assess the extent of resources lost due to widowhood that are not covered. Specifically, we see that eligibility induces a benefit transfer of $4,735. The results imply that widows who exogenously receive $4,735 less in resources would increase their labor supply by 13% (=$1,895/$14,909). In our overall sample of widows, panel B of Figure 2 showed that surviving spouses increase their labor earnings by an average of $358 on a baseline of $11,583, that is, by 3%. We can use the sensitivity to benefit eligibility to proportionally scale the response to spousal death in the overall sample to assess the part of widows' income loss that is uninsured by SSA's survivors benefits. The uninsured loss amounts to $1,092 (=$4,735*(3%/13%)), as compared to widows' overall income of the amount $45,713 (=$70,803-$25,090). This suggests that the share of the uninsured loss is rather small (2.4%) and therefore supports the conjecture that there is little need for additional self-insurance among our overall sample of widows.

Finally, we explicitly quantify the potential utility gain among the full sample of widowed households from a marginal transfer of resources across states of nature, which is expected to be low given the discussion above. We calculate ω_2 for this sample, which is a product of their labor supply responses to spousal death ($(l_2^b - l_2^g)/l_2^g$) and the labor-disutility curvature parameter φ. For sensitivity analysis, we look at a range of labor supply responses by taking the point estimate along with the upper and lower bounds of the corresponding 95-percent confidence interval, and we similarly offer a range of values for φ, which we vary from 0.5 to 2 where $\varphi = 1$ corresponds to quadratic labor disutility. Panel B of Appendix Table G.1 summarizes these results. With the point estimate of labor supply responses among the overall sample of widows (of 3.1%), the excess gain to households from a transfer of \$1 across states of nature is small, ranging from 1.5 up to at most 6.2 cents.

Income Scaling and State Dependencies. Beyond changes in consumption technology across states of nature that affect choices, how one responds to an adverse event could also be impacted by changes in preferences that result from the shock. For example, if the cost of using labor supply for self-insurance increases relative to the consumption valuation of the additional resources it provides, then the value of benefits could be higher than what is implied by any observed labor supply increases (and vice versa). Our empirical setting allows us to shed some light on relative changes in preferences upon spousal death.

Specifically, we can use our model from Section 3 to develop a relationship between labor supply responses to spousal death and state dependencies in consumption and labor as revealed by widows' behavior. To proceed, recall our definition for the local labor-disutility state dependence parameter $\theta^v(l) \equiv \frac{v_2^{b\prime}(l)}{v_2^{g\prime}(l)}$. Similarly, we define a local consumption-utility state dependence parameter as $\theta^u(c) \equiv \frac{u_2^{b\prime}(c)}{u_2^{g\prime}(c)}$. We note that $\theta^u(c)$ is a *preference* state dependence parameter that captures the degree to which the value of a unit of *consumption* changes across states of nature. In combination with the *technology* state dependence parameter, λ, which converts dollars to individual-level consumption, $\tilde{\theta}^u(c) \equiv \frac{\theta^u(c)}{\lambda} = \frac{u_2^{b\prime}(c)}{\lambda u_2^{g\prime}(c)}$ is the effective welfare-relevant composite parameter of state dependence in consumption that governs the value of transfers across states of nature. It captures the overall degree to which the consumption value of an extra dollar changes upon spousal death.[31] Putting widows' state-contingent first order conditions together yields the following approximated relationship:[32]

---

[31] To see that the composite consumption parameter $\tilde{\theta}^u(c)$ governs normative statements for the optimal transfers policy, consider the consumption representation of the welfare formula for the willingness to pay for transfers across states of nature, $\frac{u_2^{b\prime}(c_2^b(e)) - \lambda u_2^{g\prime}(c_2^g(e))}{\lambda u_2^{g\prime}(c_2^g(e))}$. Using a first order expansion, it can be approximated as $\frac{u_2^{b\prime}(c_2^b(e)) - \lambda u_2^{g\prime}(c_2^g(e))}{\lambda u_2^{g\prime}(c_2^g(e))} \cong \tilde{\theta}^u(c_2^b(e))\gamma(c_2^g(e))\left(\frac{c_2^g(e) - c_2^b(e)}{c_2^g(e)}\right) + \tilde{\theta}^u(c_2^b(e)) - 1$, where $\gamma(c_2) \equiv -\frac{u_2^{g\prime\prime}(c_2)}{u_2^{g\prime}(c_2)} c_2$, $\theta^u(c) \equiv \frac{u_2^{b\prime}(c)}{u_2^{g\prime}(c)}$, and $\tilde{\theta}^u(c) \equiv \frac{\theta^u(c)}{\lambda}$.

[32] We derive this relationship by dividing the first order conditions $u_2^{b\prime}(c_2^b) = \frac{v_2^{b\prime}(l_2^b)}{w_2}$ and $u_2^{g\prime}(c_2^g) = \frac{v_2^{g\prime}(l_2^g)}{\lambda w_2}$ and by using the first order expansions $\theta^u(c_2^b)\frac{u_2^{g\prime}(c_2^b)}{u_2^{g\prime}(c_2^g)} \cong \theta^u(c_2^b)\left[1 + \gamma(c_2^g)\frac{c_2^g - c_2^b}{c_2^g}\right]$ and $\theta^v(l_2^b)\frac{v_2^{g\prime}(l_2^b)}{v_2^{g\prime}(l_2^g)} \cong \theta^v(l_2^b)\left[1 + \varphi(l_2^g)\frac{l_2^b - l_2^g}{l_2^g}\right]$.

$$\tilde{\theta}^u(c_2^b)\left[1+\gamma(c_2^g)\frac{c_2^g-c_2^b}{c_2^g}\right]=\theta^v(l_2^b)\left[1+\varphi(l_2^g)\frac{l_2^b-l_2^g}{l_2^g}\right],$$

where $\varphi(l_2)\equiv\frac{v_2''(l_2)}{v_2'(l_2)}l_2$ is the curvature of labor disutility and $\gamma(c_2)\equiv-\frac{u_2^{g\prime\prime}(c_2)}{u_2^{g\prime}(c_2)}c_2$ is the curvature of consumption utility.

We use two additional inputs to calibrate this relationship. First, we need an equivalence scale measure to benchmark the share of a household-level income a single individual would need to maintain a given level of consumption as compared to when consuming in a couple. We use the standard square-root scale for income equivalence scaling, which is used by the U.S. government in calculating welfare transfers.[33] This scaling would imply that a one-person household would need 71% of the two-person household's income to maintain the same level of individual consumption. Second, we need to impute the *effective* equivalence scale of household income that equates the marginal value of earning an extra dollar before and after widowhood (taking into account both its consumption value and its labor cost). We use our findings from Figure 2 for this imputation. The net household income at widowhood that is required for this equivalence is captured by the sum of widows' realized income (\$70,803-\$25,090) and the imputed effective uninsured loss (\$1,092); that is, \$46,805. Hence, relative to the counterfactual in the absence of the event (\$70,803), the income equivalence rate that would equate widows' valuation of earning a dollar before and after spousal death is expected to be 66% (=\$46,805/\$70,803). Accordingly, if this amount of income compensation is received following spousal death, then we would expect widows to display no self-insurance through labor supply responses.

Together, this means that a labor supply response of $\frac{l_2^b-l_2^g}{l_2^g}=0$ corresponds to a willingness to incur a consumption loss of $\frac{c_2^g-c_2^b}{c_2^g}=\frac{0.71-0.66}{0.71}=0.07$.[34] Accordingly, we can provide the following calibrated relationship:

$$\frac{\tilde{\theta}^u(c_2^b)}{\theta^v(l_2^b)}\equiv\frac{\theta^u(c_2^b)/\lambda}{\theta^v(l_2^b)}=\frac{1}{1+\gamma\times 0.07}.$$

A few observations are worth noting. First, recall that $\theta^v(l_2^b)$ represents the degree to which earning an extra dollar has become more costly and that $\tilde{\theta}^u(c_2^b)$ represents the degree to which the value of an extra dollar in terms of individual consumption utility has changed. Therefore, the ratio $\frac{\tilde{\theta}^u(c_2^b)}{\theta^v(l_2^b)}$ captures the overall

[33] See, for example: https://www.census.gov/topics/income-poverty/income-inequality/about/metrics/equivalence.html

[34] This assumes that changes in resources reflect changes in consumption. It turns out to be a reasonable assumption in our setting, as evidenced by a calibration exercise we provide in the Appendix on the degree to which widowed households exhibit behaviors that are consistent with a hand-to-mouth benchmark. Specifically, in Appendix C.1 we provide a simple calibration which suggests that the representative widowed household's responses constitute 70% of the hypothetical response under a complete hand-to-mouth benchmark. This finding is in line with findings from Card et al. (2007) of 70% for job searchers in Austria.

change in the consumption utility value of producing an additional dollar through earnings following a spousal death. Since we find this ratio to be smaller than 1 as $\gamma > 0$, it means that, when accounting for the different channels of state dependence (in both preferences and technology), the returns to household self-insurance using earnings declines following a spousal death; or, alternatively, it implies that the cost of self-insuring an additional dollar increases in the transition to widowhood. Second, we can use equivalence scaling to draw implications about relative state dependence in preferences themselves. With $\lambda = \frac{1}{0.71} = 1.41$, we have that $\frac{\theta^u(c_2^b)}{\theta^v(l_2^b)} = \frac{1.41}{1+\gamma\times 0.07}$. Following Landais and Spinnewijn (2021), we consider as standard calibration of γ the range of values from 1 to 4. For this range, we have that $\frac{\theta^u(c_2^b)}{\theta^v(l_2^b)} > 1$. In turn, this implies that, upon widowhood, the utility value of an additional unit of consumption increases by more than the increase in the utility cost of an additional unit of labor. This translates to stronger state dependence in consumption utility than in labor disutility following spousal death.

It is important to note that while these implications are novel, they do not directly map to the preference state-dependence parameters that govern welfare calculations. Specifically, the parameter that enters into our modified welfare formula for assessing the degree of coverage in ω_2' is $\theta^v(l_2^b)$. The state dependence exercises we provide here can only speak to the relative changes in the valuation of consumption and labor, but whether those relative changes are driven by changes to consumption utility or changes to labor disutility is central for making quantitative normative statements. We study the sensitivity of our welfare conclusions regarding coverage adequacy to the value of $\theta^v(l_2^b)$ below in our discussion of the net social returns to a marginal increase in benefits.

Fiscal Externality and Net Social Returns. It is useful to provide a brief discussion of fiscal externalities in the context of marginal increases in transfers to widows financed by an increase in SSA payroll taxes. The fiscal externality would involve two components: the behavioral responses to the increased generosity of benefits among widows and the behavioral responses to the increased tax liability among workers. For the former, we use our empirical estimates, and for the latter, we rely on standard magnitudes of net-of-tax elasticities. In Appendix G.1, we calculate this fiscal externality to be \$0.049 per \$1 of transfers. This calculation assumes a net-of-tax elasticity of 0.20, although research (in settings outside the U.S.) suggests that payroll tax elasticities are much lower, which would reduce the efficiency costs.[35]

[35] Saez et al. (2012) and Lehmann et al. (2013) find no labor supply responses to variation in payroll taxes (in Greece and France, respectively). In the extreme but plausible case of close to zero net-of-tax elasticities in the context of payroll tax, the fiscal externality would reduce to \$0.0225 from widows' behavioral responses to increased benefit generosity.

To assess the net societal return from such transfer expansions, we combine our calibration of the potential utility gain to widowed households (from panel B of Appendix Table G.1) with its fiscal externality (from Appendix G.1). The results are summarized in panel C of Appendix Table G.1, where we consider the value range 0.5-2 for φ and the value range 0-0.30 for the net-of-tax elasticity. We find that the return to an additional \$1 of transfers to widows ranges from -4.7 to 3.9 cents with an average of -0.4 cents, suggesting that benefits hover around the appropriate levels from the planner's perspective.

Finally, we can incorporate potential state dependencies discussed earlier into our evaluation of the net societal returns to a marginal transfer. To do so, we impute the level of labor disutility state dependence, θ^v, that would justify the current levels of transfers to widows. Accordingly, we assess state dependence levels that guarantee non-negative net social returns on the marginal dollar. We choose the highest value of the net-of-tax elasticity from panel C of Appendix Table G.1 to provide a more conservative measure regarding the net social returns. For the marginal dollar of transfers to widows to be justified from a net welfare perspective, we must have that the marginal gain to households (captured by ω_2') is larger or equal to the fiscal externality, that is:

$$\omega_2' = \theta^v \times \varphi \times 0.031 + \theta^v - 1 \geq 0.062.$$

Panel D of Appendix Table G.1 provides these imputations for the range 0.5-2 for φ. The state dependence parameter θ^v that would justify marginal transfers ranges from 1 (no state dependence) to 1.046, where the latter translates to the case in which the marginal cost of labor supply increases by 4.6% upon spousal death. That is, for the parameter ranges we consider for φ, even a moderate degree of state dependence in labor (θ^v) would be enough to justify the generosity of transfers to American widows.

7. Conclusion

Using tax records for the U.S. population and exploiting quasi-experimental variation in the timing of spousal death and survivors benefit eligibility, we study American widows' labor supply and the design of Social Security's survivors insurance. Combined with theory, we offer a novel and comprehensive investigation of survivors benefits insurance programs, which overcomes key challenges in welfare assessments, including scale economies and state dependence in preferences. Our evidence implies that widows highly value the liquidity the program provides and have little residual need to self-insurance. Overall, our analysis underscores the importance of the survivors benefits program and the protection it offers to American households against the significant financial risks associated with spousal death.

References:

Angel, Jacqueline L., Maren A. Jiménez, and Ronald J. Angel. (2007). "The Economic Consequences of Widowhood for Older Minority Women." *The Gerontologist*, 47(2): pp. 224-234.

Auerbach, Alan J., and Laurence J. Kotlikoff (1987). "Life Insurance of the Elderly: Adequacy and Determinants." In *Work, Health and Income Among the Elderly* (pp. 229–267), G. Burtless, ed., Brookings Institution Press.

Auerbach, Alan J., and Laurence J. Kotlikoff (1991a). "The Adequacy of Life Insurance Purchases." *Journal of Financial Intermediation*, 1(3): pp. 215-241.

Auerbach, Alan J., and Laurence J. Kotlikoff (1991b). "Life Insurance Inadequacy-Evidence from a Sample of Older Widows." Working Paper No. w3765, National Bureau of Economic Research.

Bernheim, B. Douglas, Lorenzo Forni, Jagadeesh Gokhale, and Laurence J. Kotlikoff (2003). "The Mismatch between Life Insurance Holdings and Financial Vulnerabilities: Evidence from the Health and Retirement Study." *American Economic Review*, 93(1): pp. 354-365.

Bernheim, B. Douglas, Katherine G. Carman, Jagadeesh Gokhale, and Laurence J. Kotlikoff (2003). "Are Life Insurance Holdings Related to Financial Vulnerabilities?" *Economic Inquiry*, 41(4): pp. 531-554.

Brown, Jeffrey, Arie Kapteyn, Olivia Mitchell, and Teryn Mattox (2013). "Framing the Social Security Earnings Test." Wharton Pension Research Council Working Paper 2013-06.

Card, David, Raj Chetty, and Andrea Weber (2007). "Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market." *Quarterly Journal of Economics*, 122(4): pp. 1511-1560.

Chetty, Raj (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy*, 116(2): pp. 173-234.

Chetty, Raj, and Amy Finkelstein (2013). "Social Insurance: Connecting Theory to Data." In *Handbook of Public Economics*, Vol. 5: pp. 111-193. Elsevier.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014). Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *Quarterly Journal of Economics*, 129(4): pp. 1553-1623.

Coile, Courtney, and Jonathan Gruber (2007). "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics*, 89(2): pp. 234-246.

Compton, Janice, and Robert A. Pollak (2018). "The Life Expectancy of Older Couples and Surviving Spouses."

Deshpande, Manasi, Itzik Fadlon, and Colin Gray (2021). "How Sticky Is Retirement Behavior in the US?" *Review of Economics and Statistics*: 1-55.

Fadlon, Itzik, and Torben Heien Nielsen (2017). "Family Labor Supply Responses to Severe Health Shocks." NBER Working Paper No. 21352.

Fadlon, Itzik, and Torben Heien Nielsen (2018). "Household Labor Supply and the Gains from Social Insurance." *Journal of Public Economics*.

Fadlon, Itzik, and Torben Heien Nielsen (2021). "Family Labor Supply Responses to Severe Health Shocks: Evidence from Danish Administrative Records." *American Economic Journal: Applied Economics*, 13 (3): pp 1-30.

Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo (2009). "Approaches to Estimating the Health State Dependence of the Utility Function." *American Economic Review*, 99 (2): pp. 116-21.

Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo (2013). "What Good is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption." *Journal of the European Economic Association*, 11(s1): pp. 221-258.

Ganong, Peter, and Pascal Noel (2019). "Consumer Spending During Unemployment: Positive and Normative implications." *American Economic Review*, 109(7): 2383-2424.

Gerard, François, and Joana Naritomi (2021). "Job Displacement Insurance and (the Lack of) Consumption-Smoothing." *American Economic Review*, 111(3): 899-942.

Giupponi, Giulia (2019). "When Income Effects are Large: Labor Supply Responses and the Value of Welfare Transfers."

Goldin, Claudia, and Lawrence F. Katz, eds. (2018). "Women Working Longer: Increased Employment at Older Ages." University of Chicago Press.

Hartley, Daniel, Anna Paulson, and Katerina Powers (2018). "What Explains the Decline in Life Insurance Ownership?" *Economic Perspectives*, 8: pp. 1-20.

Hendren, Nathaniel (2017). "Knowledge of Future Job Loss and Implications for Unemployment Insurance." *American Economic Review*, 107(7): pp. 1778-1823.

Holden, Karen C., Richard V. Burkhauser, and Daniel J. Feaster (1988). "The timing of Falls into Poverty after Retirement and Widowhood." *Demography,* 25(3): pp. 405-414.

Holden, Karen C., Richard V. Burkhauser, and Daniel A. Myers (1986). "Pensioner's Annuity Choice: Is the Well-being of Their Widows Considered?" University of Wisconsin Institute for Research on Poverty Discussion Paper.

Hurd, Michael D. (1989). "The Poverty of Widows: Future Prospects." In *The Economics of Aging* (pp. 201-230), Wise, ed., University of Chicago Press.

Hurd, Michael D., and David A. Wise (1989). "The Wealth and Poverty of Widows: Assets before and after the Husband's Death." In *The Economics of Aging* (pp. 177-200), Wise, ed., University of Chicago Press.

Hurd, Michael D., and David A. Wise (1996). "Changing Social Security Survivorship Benefits and the Poverty of Widows." In *The Economic Effects of Aging in the United States and Japan* (pp. 319 - 332), Michael D. Hurd and Naohiro Yashiro, eds., University of Chicago Press.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006). "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review,* 96(5): pp. 1589-1610.

Kolsrud, Jonas, Camille Landais, Daniel Reck, and Johannes Spinnewijn (2021). "Retirement Consumption and Pension Design."

Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn (2018). "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden." *American Economic Review,* 108(4-5), pp. 985-1033.

Landais, Camille (2015). "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design." *American Economic Journal: Economic Policy,* 7 (4): pp. 243-78.

Landais, Camille, and Johannes Spinnewijn (2021). "The Value of Unemployment Insurance." *Review of Economic Studies*, 88(6): 3041-3085.

Lehmann, Etienne, François Marical, and Laurence Rioux (2013). "Labor Income Responds Differently to Income-Tax and Payroll-Tax Reforms." *Journal of Public Economics*, 99: 66-84.

Liebman, Jeffrey B., and Erzo F.P. Luttmer (2012). "The Perception of Social Security Incentives for Labor Supply and Retirement: The Median Voter Knows More Than You'd Think." *Tax Policy and the Economy*, 26(1): pp. 1-42.

Liebman, Jeffrey B., and Erzo F.P. Luttmer (2015). "Would People Behave Differently if They Better Understood Social Security? Evidence from a Field Experiment." *American Economic Journal: Economic* Policy, 7(1): pp. 275-99.

MaCurdy, Thomas E. (1981). "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy,* 89(6): pp. 1059-85.

McGarry, Kathleen, and Robert F. Schoeni (2000). "Social Security, Economic Growth, and the Rise in Elderly Widows' Independence in the Twentieth Century." *Demography,* 37(2): pp. 221-236.

McGarry, Kathleen, and Robert F. Schoeni (2005). "Widow(er) Poverty and Out-of-Pocket Medical Expenditures near the End of Life." *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 60(3): pp. S160-S168.

Myers, Daniel A., Richard V. Burkhauser, and Karen C. Holden (1987). "The Transition from Wife to Widow: The Importance of Survivor Benefits to Widows." *Journal of Risk and Insurance*: pp. 752-759.

Parker, Jonathan A. (1999). "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review*, 89(4): pp. 959-973.

Saez, Emmanuel, Manos Matsaganis, and Panos Tsakloglou (2012). "Earnings Determination and Taxes: Evidence from a Cohort-Based Payroll Tax Reform in Greece." *Quarterly Journal of Economics*, 127(1): 493-533.

Sevak, Purvi, David R. Weir, and Robert J. Willis (2003). "The Economic Consequences of a Husband's Death: Evidence from the HRS and AHEAD." Social Security Bulletin, 65(3): pp. 31-44.

Shimer, Robert, and Ivan Werning (2007). "Reservation Wages and Unemployment Insurance." *Quarterly Journal of Economics,* 122(3): pp. 1145-1185.

Social Security Administration (2018a). Social Security Beneficiary Statistics.

Social Security Administration (2018b). Exempt Amounts under the Earnings Test.

Social Security Administration (2018c). Benefits Paid by Type of Beneficiary.

Social Security Administration (2020). Never Beneficiaries, Aged 60 or Older, 2020.

Souleles, Nicholas S. (1999). "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*, 89(4): pp. 947-958.

Stroebe, Margaret, Henk Schut, and Wolfgang Stroebe (2007). "Health Outcomes of Bereavement." *The Lancet,* 370 (9603): pp. 1960-1973.

Weir, David R., and Robert J. Willis (2000). “Prospects for Widow Poverty in the Finances of Married Couples in the HRS.” In *Forecasting Retirement Needs and Retirement Wealth* (pp. 208-234), Olivia Mitchell, P. Brett Hammond, and Anna Rappaport, eds., Philadelphia: University of Pennsylvania Press.

Wettstein, Gal (2020). “Retirement Lock and Prescription Drug Insurance: Evidence from Medicare Part D.” *American Economic Journal: Economic* Policy, 12 (1): pp. 389-417.

White House (2016). Fiscal Year 2017: Historical Tables – Budget of the U.S. Government.

Zeldes, Stephen P. (1989). “Consumption and Liquidity Constraints: An Empirical Investigation.” *Journal of Political Economy*, 97 (2): pp. 305-346.

Figure 1: Effects of Eligibility for Social Security's Survivors Benefits

A. Social Security Benefit Claiming

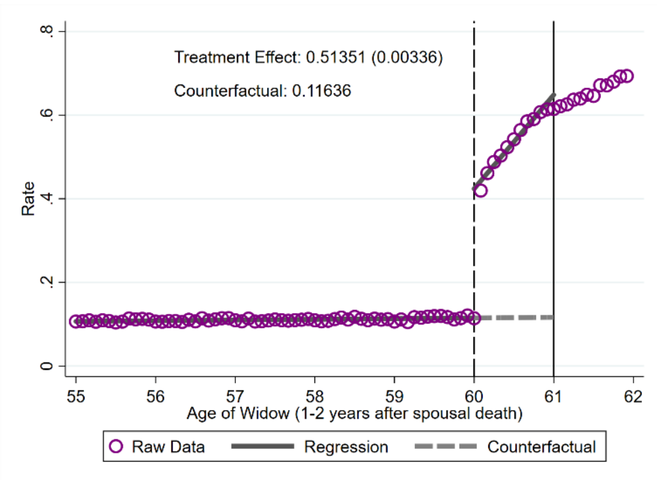

B. Social Security Benefit Amounts

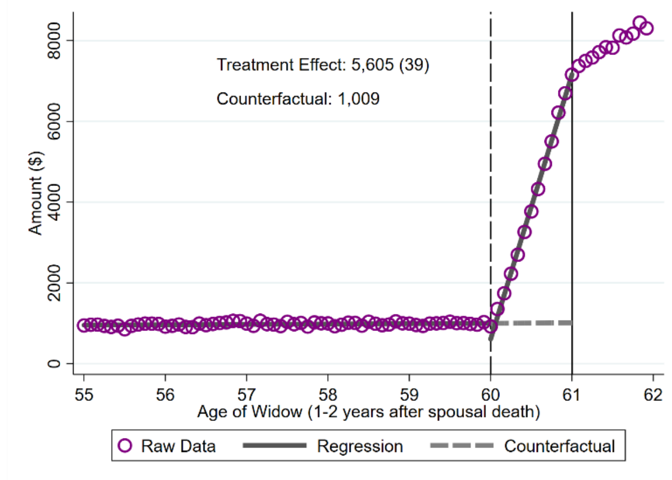

C. Overall Net Household Income

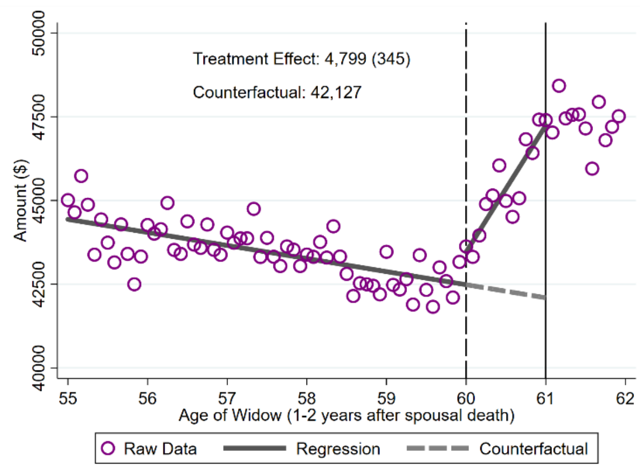

Figure 1: Effects of Eligibility for Social Security's Survivors Benefits (*continued*)

D. Retirement Rate

E. Labor Force Participation

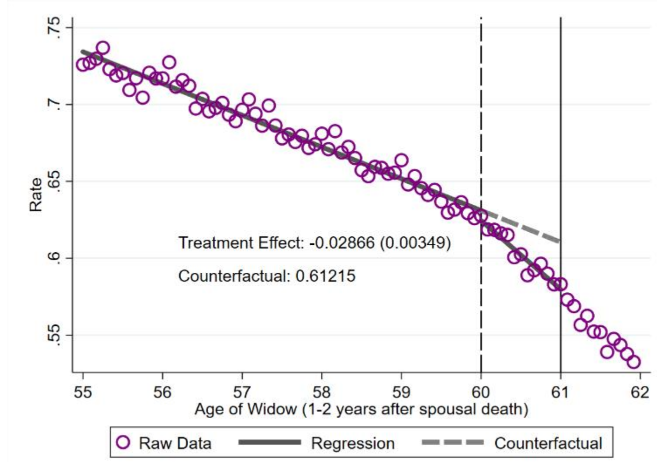

F. Work Intensity

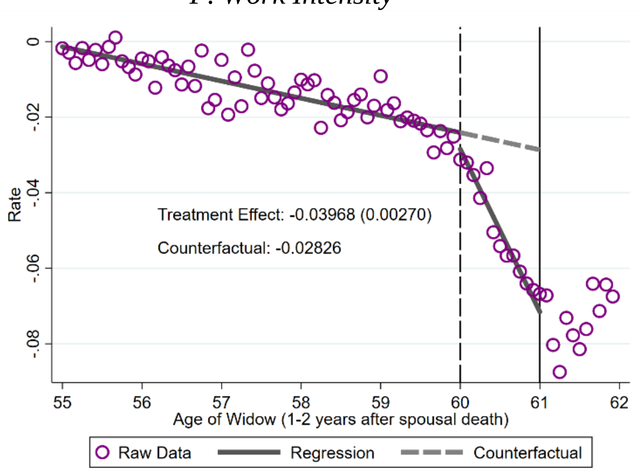

G. Wage Earnings

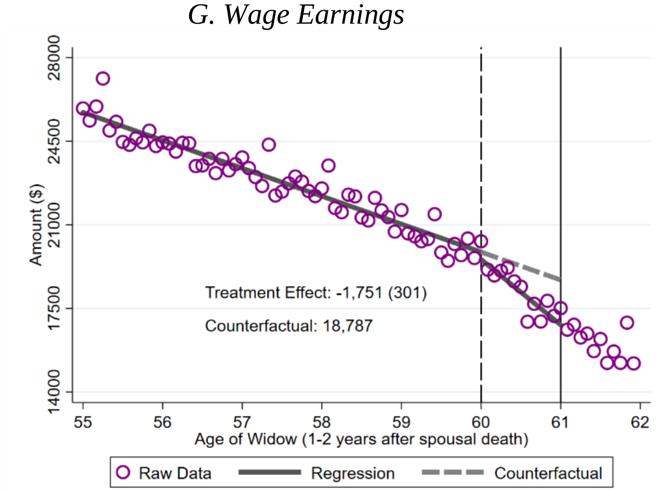

Notes: These figures plot various household outcomes in the years just after a husband's death event (years 1 and 2 following the event) as a function of the surviving spouse's age in months. The purple circles represent the means of the raw data for each "monthly age" bin. The solid gray lines plot the piecewise linear fit using equation (1). The dashed gray line in the age range 60-61 represents the counterfactual behavior in the absence of eligibility for survivors benefits based on specification (1), which extrapolates the linear relationship estimated on observations prior to age 60. Eligibility for benefits begins at exactly age 60 (which is marked by the vertical dashed black line). The full-exposure effect of benefit eligibility is represented by the vertical gap between the solid and the dashed gray regression lines at age 61 (which is marked by the vertical solid black line). Each panel reports the full-exposure effect (with standard errors in parentheses) as well as counterfactual levels.

Figure 2: Dynamic Effects of Spousal Death

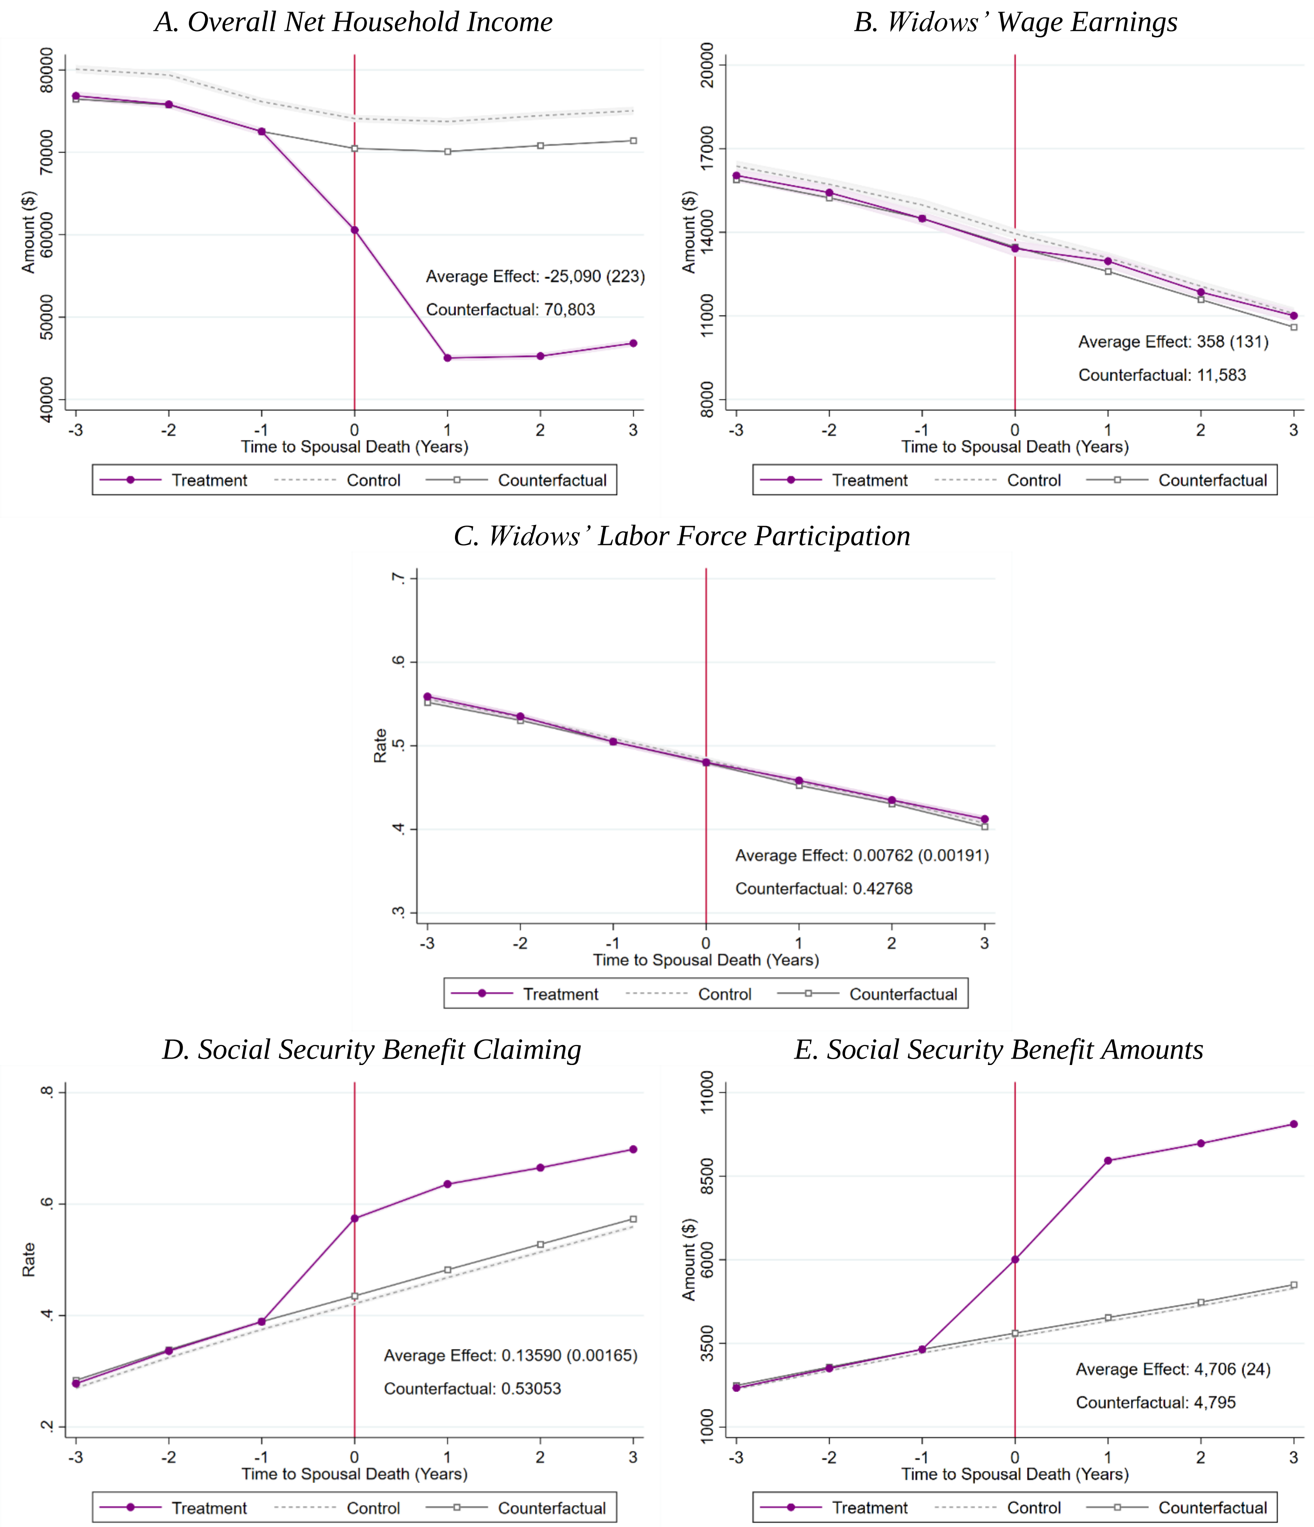

Notes: These figures plot the evolution of different household outcomes in response to spousal death. The x-axis denotes time with respect to the event year, normalized to period 0. For the treatment group, period 0 is when the actual event occurs; for the control group, period 0 is when a placebo event occurs (while their actual event occurs in period 4). The dashed gray line plots the behavior of the control group (along with the corresponding 95-percent confidence intervals). To ease the comparison of trends, from which the effect is identified, we normalize the level of the control group's outcome to the pre-event level of the treatment group's outcome (in period -1). This normalized counterfactual is displayed by the solid gray line and squares. The purple lines and circles plot the behavior of the treatment group (along with the corresponding 95-percent confidence intervals). In each panel, we report the average treatment effect over periods 1-3, based on δ_n from equation (2), as well as the corresponding counterfactual level.

Table 1: Effects of Spousal Death by Age-Eligibility for Survivors Benefits

	Widows' Labor Supply		Overall Net Household Income	Social Security Benefits	
	Wage Earnings	Participation		Claiming Rate	Benefit Amount
	(1)	(2)	(3)	(4)	(5)
Average Treatment Effect	312	0.00112	-26,373	0.36528	3,662
	(380)	(0.00668)	(838)	(0.00333)	(39)
Effect Among Age Eligible	-605	-0.01690	-23,578	0.59143	5,919
	(461)	(0.00866)	(1,067)	(0.00471)	(56)
Effect Among Age Ineligible	1,290	0.02048	-29,412	0.11975	1,184
	(545)	(0.00874)	(1,109)	(0.00346)	(41)
Difference: Ineligible − Eligible	1,895	0.03737	-5,834	-0.47168	-4,735
	(661)	(0.01112)	(1,384)	(0.00545)	(64)
Counterfactual	14,909	0.54281	69,090	0.06207	1,535.27
Number of Obs.	161,073	161,073	161,073	161,073	161,073
Number of Clusters	67,140	67,140	67,140	67,140	67,140

Notes: This table studies the effects of spousal death by age-eligibility for survivors benefits. We estimate a specification of equation (3), where we let the average treatment effect vary by age eligibility. We estimate this equation on a narrow age band around the eligibility cutoff (ages 58-61). We define a dummy variable for eligibility as the widow being 60 or older, which we interact with the main term $treat_i \times post_{i,t}$, and we flexibly include age controls in the regression: age dummies and their interactions with $post_{i,t}$ and with $treat_i$ separately. Counterfactuals are calculated based on the estimated equation for age 61 in the median post-shock calendar year (2004). All specifications include as controls age indicators and calendar year fixed effects, and we report robust standard errors clustered at the household level.