NBER WORKING PAPER SERIES

HOUSEHOLD RESPONSES TO TRANSFERS AND LIQUIDITY: EVIDENCE FROM SOCIAL SECURITY'S SURVIVORS BENEFITS

Itzik Fadlon
Shanthi P. Ramnath
Patricia K. Tong

Working Paper 25586
http://www.nber.org/papers/w25586

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019

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 represent the views of 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, Jon Gruber, Henrik Kleven, Amy Finkelstein, Nathan Hendren, Erzo Luttmer, Petra Persson, Tommaso Porzio 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, and the 2018 All-California Labor Economics Annual Meeting for helpful comments and discussions. Jonathan Leganza provided excellent research assistance.

Household Responses to Transfers and Liquidity: Evidence from Social Security's Survivors Benefits

Itzik Fadlon, Shanthi P. Ramnath, and Patricia K. Tong

NBER Working Paper No. 25586
February 2019
JEL No. D1,D61,G22,H0,H55,I1,I38,J2

ABSTRACT

We use administrative tax data that cover the U.S. population to identify the causal effects of Social Security's survivors benefit receipt on American families' behavior and financial well-being. We analyze over a quarter of a million widowed households in which the husband died between 2002-2007, and we exploit a sharp age discontinuity in benefit eligibility to study the responses of financially vulnerable households to government transfers. We first study how households respond to unanticipated benefit receipt in the immediate periods following a large financial shock to investigate the protective role of transfers. We find significant impacts of the program on newly-widowed families' net income and labor supply behavior, which points to considerable allocative inefficiencies in the life insurance market and to a high valuation of survivors benefits in protecting Americans against mortality shocks. Second, to investigate the particular role of liquidity and benefit timing, we then study how already-widowed women's labor supply responds to anticipated survivors benefit receipt. We find considerable responses to cash-on-hand via benefit availability that underscore allocative inefficiencies in the credit market and the value of liquidity itself provided by government transfers. These responses and their heterogeneity highlight mechanisms that underlie the labor supply behavior of older vulnerable households, and they point to liquidity constraints, rather than myopia or benefit-schedule misperceptions, as the likely operative channel. Our results have implications for survivors benefits in the U.S., and, more generally, for retirement behavior and response mechanisms to transfers among older vulnerable populations.

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

Patricia K. Tong
RAND Corporation
1776 Main Street
P.O. Box 2138
Santa Monica, CA 90407-2138
ptong@rand.org

Shanthi P. Ramnath
Office of Tax Analysis
US Department of the Treasury
1500 Pennsylvania Ave
Washington DC, 20220
shanthi.ramnath@treasury.gov

1. Introduction

The death of a primary earner is among the most devastating shocks that households face, and it poses a major source of economic risk for American families. Fadlon et al. (2019) show that a husband's death leads to significant declines in equivalence-scale adjusted household income and to considerable increases in financial insolvency, which are both immediate and persistent.[1] In the United States, there are approximately 15 million surviving spouses at any given time, with 1.4 million newly-widowed households each year, among which 80% are women.[2] 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 U.S. government paid more than $95 billion to 4.2 million surviving spouses (up from $64 billion in 2000); where, by comparison, unemployment and Earned Income Tax Credit benefits amounted to $35 and $60 billion, respectively (White House 2016; SSA 2018a). Moreover, several proposals to both expand and reform this large program are currently under consideration.[3] In addition, the literature has pointed to considerable inadequacies in Americans' life insurance holdings (Auerbach and Kotlikoff 1987, 1991a,b; Bernheim et al. 2003a,b), which have been declining in recent decades (Hartley et al. 2018), suggesting an even larger and increasing role for social insurance against spousal death. Yet, despite all these facts and to the best of our knowledge, there is virtually no direct causal evidence on the economic effects of Social Security's survivors benefits or on the financial protection they provide to vulnerable American households.[4]

In this paper, using tax records that cover the U.S. population, we estimate the causal effects of Social Security's survivors benefit receipt on households' behavior and economic well-being. From a panel that spans the years 1999 through 2014, we analyze close to a quarter of a million households that have experienced the event of a husband's death. The data include a rich set of financial outcomes. This provides us with the distinct advantage to investigate aggregates of effective financial well-being, accounting for potential substitution of alternative sources of income such as private savings; as well as with the ability to analyze households' behavioral responses, most importantly, in their labor supply.

[1] The declines in income amount to 12% among all households with widows of ages 50-70 and to 23% among the subset of households in which the husband continuously worked in the years prior to his death.

[2] Compton and Pollak (2018) also show 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.

[3] See, e.g., the "Surviving Widow(er) Income Fair Treatment Act of 2018" that was introduced on September 18, 2018 by Senator Robert P. Casey, Jr., as well as various suggested changes at https://www.ssa.gov/oact/solvency/provisions/index.html (Section D).

[4] Relatedly, Hurd and Wise (1996) simulated how widows' poverty would mechanically change (i.e., abstracting from behavioral responses) if Social Security's survivors benefits were to increase in a balanced-budget way through declines in spousal benefits to couples; and McGarry and Schoeni (2000) study factors that can explain the long-run changes in widows' living arrangements toward independent living and point to income growth through increased Social Security benefits as a potential explanation. Studies on government transfers to survivors in the context of countries other than the U.S. include Fadlon and Nielsen (2017) in the case of Denmark and Giupponi (2018) in the case of Italy. Aizer, Eli, Ferrie, and Lleras-Muney (2016) study the effects on children of a welfare program in the U.S. for poor mothers with dependent children whose large share of recipient mothers were widows.

To estimate the economic impacts of Social Security's survivors benefits, we exploit a sharp discontinuity in the age at which surviving spouses become eligible for transfers. In particular, widowed individuals become eligible at exactly age 60 for their deceased spouses' potential Social Security retirement benefits. We leverage the scope and detailed nature of the data to investigate financial outcomes and behaviors of households that fall just below and just above the age eligibility cutoff at different time horizons following the spousal death. Based on this design, we utilize the advantages of our setting—a considerable financial shock with well-defined timing, a large social program relevant for a significant share of households, and a clean quasi-experiment—to study two related broader questions on the responses of financially vulnerable households to government transfers.

The first question that we study is how households respond to unanticipated benefit receipt in the immediate periods following a large financial shock. In our first empirical exercise, we analyze the effects of being eligible for benefit receipt in the immediate run among newly-widowed households. Studying widows' responses just after the death event occurs allows us to test for allocative inefficiencies in the life insurance market and to evaluate the protective value of transfers from the government. This analysis is based on the notion that while, ex-ante, the benefit schedule itself can be fully anticipated, ex-post eligibility for benefit receipt at the particular period of the transition to widowhood is unanticipated to extent that the exact timing of death is unpredictable. The second broader question that we study is how vulnerable households respond to anticipated benefit receipt. In our second empirical exercise, we analyze the effects of becoming eligible for benefits in the longer run among already-widowed households, whose benefit entitlement has been determined so that benefit availability is completely predictable. Studying these longer run responses allows us to test for allocative inefficiencies in the credit market and to cleanly isolate the value of liquidity itself provided by government transfers.

Overall, a key takeaway that emerges from the combination of our analyses is the important dual role of liquidity over the course of Americans' life-cycle. The first is the insurance role of liquidity in smoothing consumption across states of nature, and the second is the intertemporal role of liquidity in smoothing consumption across time periods; both of which we show to be qualitatively and quantitatively important.

We begin by addressing the first question and we find large impacts of eligibility for survivors benefit receipt in the periods that just follow households' transition to widowhood. First, we show that eligibility leads to an average increase of $4,804 in net annual income, accounting for adjustments in alternative income sources, which represents a substantial improvement in widows' overall financial circumstances. Second, eligibility for survivors benefits leads to meaningful declines in widows' labor supply. Together, these findings point to a high valuation of Social Security's survivors benefits by American families in terms of the immediate protection they provide against the adverse financial impact

of mortality shocks: the benefits generate gains to newly-widowed households both through significant increases in income available for consumption, and through greater consumption of leisure as the social program mitigates the need for income compensation via family labor supply.

Our setting lends itself to an intuitive test for inefficiencies in the life insurance market, which constitutes one of the largest private insurance markets in the United States. Specifically, we show that efficient markets imply no labor supply responses to benefit eligibility by newly-widowed households. Our findings reveal significant deviations from this benchmark, indicating considerable allocative inefficiencies. By relying on labor supply responses across widows with differential benefit eligibility status, a key advantage of this test is that it does not rely on estimates for scale economies within the household or on assumptions about state dependence in preferences. With this analysis, our paper contributes to the important earlier work that aimed to assess the adequacy of households' life insurance holdings in the U.S. (Auerbach and Kotlikoff 1987, 1991a,b; Bernheim et al. 2003a,b). In light of the deviation from efficient insurance, the evidence and the features of the program indicate that these labor market responses are driven by the effect of cash-on-hand (rather than a wealth effect). Moreover, they suggest that these responses are likely socially desirable, as we document notable non-distortionary effects (vs. substitution effects) in family labor supply behavior.

We then proceed to address our second question and we study the labor supply and retirement responses of households who had already transitioned to widowhood several years back. Since the present discounted value of these households' benefit entitlement is completely fixed, benefit receipt at the eligibility age cutoff represents a discontinuous predictable increase in cash-on-hand. As such, it identifies labor supply responses of vulnerable older Americans that are driven purely by the availability of anticipated liquidity. With this analysis, we contribute to the literature on household behavior in the context of liquidity constraints and credit market inefficiencies, which spans various fields including public finance and macroeconomics (see, e.g., Zeldes 1989, Johnson, Parker, and Souleles 2006, and Card, Chetty, and Weber 2007). Our analysis relies on the observation that, within our setting, the frictionless world implies there should be no labor supply responses when the anticipated benefits become available—that is, the efficient benchmark of responses is zero. Hence, a key advantage of our setting is that it provides a test for credit inefficiencies and assessments of the value of liquidity that do not require calibrations of benchmarks for evaluating excess sensitivity to income. Even more, we are also able to take advantage of the program's design and our data to devise tests for the behavioral channels that underlie household responses to predictable changes in cash-on-hand, focusing on the degree of forward-looking behavior and liquidity or borrowing constraints.

The findings point to significant deviations from the frictionless first-best benchmark in which labor supply should be smoothed and exhibit no sensitivity to anticipated benefit timing. Widows'

retirement rate displays a clear jump at survivors benefit availability, which translates to sizable declines in labor force participation and earnings in response to anticipated benefit receipt. These responses first point to notable liquidity and credit allocation inefficiencies and imply that liquidity plays an important role in the labor supply decisions of vulnerable American families.[5] We also show that the value from injecting liquidity earlier in the life-cycle is proportional to the considerable labor supply responses that we estimate.

Our results underscore that at least one of the underlying assumptions of the frictionless model is violated, consistent with either myopia or liquidity constraints as potential driving channels. Providing empirical tests that investigate which response channels are likely operative, we first show that complete myopia among widowed households is unlikely to explain the results. We find that widows clearly exhibit strategic timing of remarriage for ensuring benefit entitlement, which requires forward-looking planning. We then show a considerable response gradient in liquidity, which can also account for the patterns that we find: the labor supply declines in response to anticipated benefit receipt are attributable to lower-liquidity households who exhibit large effects, whereas highest-liquidity households exhibit labor supply smoothing as predicted by the frictionless model. In line with these findings, we also illustrate that presumably myopia-free households who are still potentially subject to liquidity constraints—specifically, either those who display strategic remarriage or those with private retirement savings but low balances—exhibit considerable declines in labor supply when benefits become available. The results therefore point to liquidity constraints, rather than a lack of forward-looking behavior, as a key operative mechanism driving the estimated responses to anticipated benefit receipt. These findings can also have broader positive implications for Americans' retirement behavior and response mechanisms to transfers and liquidity at older ages. Widows represent a considerable share of older households, as more than 13.5% of all American women are widowed by age 65 (Fadlon et al. 2019).

In addition to the main contributions mentioned so far, our results have direct implications for Social Security's survivors insurance program. First, in regards to the program's generosity, the evidence is consistent with close-to-full income compensation for the immediate losses from spousal deaths among eligible families, and it indicates that ineligible households are exposed to substantial risk and would well value insurance through government transfers. Second, the revealed role of liquidity implies that merely changing the timing of benefits and smoothing the benefits' profile, keeping their current level of present discounted value unchanged, has potential for generating considerable value to households. In addition, albeit more suggestively, the analysis could be informative for the discussion of the possible responses to reforming the Social Security Early Eligibility Age (EEA), since the eligibility for anticipated benefits among widows at age 60 is the only source of variation in early eligibility since its introduction. Lastly, we

[5] Somewhat relatedly, Ganong and Noel (2018) find that liquidity matters in the context of mortgage debt obligations and the default behavior of mortgage borrowers in the U.S., although in a setting where provision of liquidity was unanticipated.

find that working widowed women exhibit a considerable spike in retirement rates in response to benefit receipt. With the significant trend of increases in female labor force participation at older ages in the U.S. (Goldin and Katz 2018), and as widows comprise a meaningful share of older American women, the evidence suggests that the Social Security survivors benefits program itself could play an increasing role in female retirement behavior.

The remainder of the paper is organized as follows. Section 2 outlines a simple conceptual framework to provide benchmarks for the household responses that we estimate. In Section 3 we describe the data and lay out our empirical framework. Section 4 presents our analysis of the effects of eligibility for Social Security's survivors benefits and discusses the implications of our findings. Section 5 concludes.

2. Conceptual Framework and Useful Benchmarks

We leverage our empirical setting to conduct two related exercises that differ by the degree of the anticipation for actual benefit receipt. Our first exercise investigates the effects of unanticipated benefit receipt at the eligibility cutoff on newly-widowed households in the period that immediately follows the death event, and it focuses on the insurance market. Whereas the benefit schedule itself can be fully anticipated, eligibility for benefits at the exact transition to widowhood is unanticipated to degree that the particular timing of the husband's death is unpredictable. Our second exercise investigates the effects of anticipated benefit availability on already-widowed households several years after the event—so their receipt of benefits is entirely predictable—and it focuses on liquidity constraints and the credit market. In this section we describe first-best benchmarks for each exercise to provide useful anchors for the effects that we estimate in the empirical analysis. These benchmarks also provide natural tests for market allocation inefficiencies, which are valuable for interpreting the implications of our findings. This section describes the setup of simple models of household behavior in the context of survivors benefits and provides the economic intuition for the optimality results. Complete details appear in Appendix A.

(i) Responses by Newly-Widowed Households to Unanticipated Benefits. Consider the decisions of a two-person household, which consists of individuals 1 and 2, in a world with two states: a "good" state, state g, and a "bad" state, state b, in which member 1 dies and member 2 becomes a widow. Households begin their planning problem in the good state, and they can transition to widowhood when they are just-ineligible for Social Security survivors benefits (i.e., when the wife is just below 60) or when they are just-eligible for benefits (i.e., when the wife is just over 60). As in our empirical setting, households that are just ineligible in the first period of widowhood become eligible in the periods that follow. That is, just-eligible and just-ineligible households differ in benefit eligibility during the period just after spousal death, and they are equally eligible for benefits in future periods. In our analysis, we investigate the behaviors of newly-

widowed households (at the first period of widowhood), comparing those who are just-eligible and just-ineligible for benefits at that stage.

As our natural benchmark, we consider a first-best world in which households can purchase actuarially-fair life insurance policies. The eligibility schedule for Social Security's survivors benefits is deterministic in age and thus fully predictable at the beginning of the analysis horizon. Hence, the household's optimal choices follow eligibility-contingent consumption bundles and insurance purchases through age-contingent plans.

With this setup, optimality leads to the classic result of full insurance in the presence of actuarially-fair insurance markets, extended to a setting of eligibility-contingent purchases and plans when the eligibility schedule is fully anticipated by age. That is, the wife's marginal utility from consumption is equated both across states of the world and across eligibility statuses upon the transition to widowhood. We derive this result in Appendix A.1. The household's standard optimal behavior also implies that at each contingency the wife chooses her labor supply so that the marginal utility from consumption equals her wage-weighted marginal disutility from labor. Together, the optimality conditions imply the following necessary condition in a first-best world: the marginal disutility from labor of a newly-widowed wife who is just-ineligible for benefits should equal the marginal disutility from labor of a newly-widowed wife who is just-eligible for benefits.

This equality forms a useful benchmark for our empirical analysis. Since it compares households that transitioned to the bad state who only differ by whether they are just-eligible or just-ineligible for benefits, equality of marginal disutility from labor is equivalent to equality of labor supply. Hence, we would expect no labor supply responses to eligibility by newly-widowed households in the presence of perfect insurance markets. As such, this necessary condition immediately provides a test for life insurance market inefficiency by estimating the degree of deviations from labor supply smoothing around our eligibility-age cutoff.

This test provides two advantages over classical tests that are based on smoothness across states of nature—in either marginal utility from consumption (see Chetty and Finkelstein 2013 for a review) or marginal disutility from labor (see, e.g., Fadlon and Nielsen 2018). First, by analyzing labor supply, the test does not rely on estimates for economies of scale within the household, for which there is no consensus in the literature. Second, by exploiting variation across widowed households (i.e., within the widowhood state), the test is not confounded by state dependence in preferences. Our analysis and conclusions are robust to potential changes in spouses' utility when the event of widowhood occurs.

Nonetheless, it is important to note that the test also has a clear limitation, in that it is "asymmetric" with respect to zero. That is, while deviations from zero would imply market inefficiencies, equality to zero would not imply market efficiency. Equality of marginal disutility from labor across just-eligible and just-

ineligible widows is necessary but not sufficient, as it can still be the case that these marginal disutilities from labor in the bad state differ from that in the good state. Such a test would naturally be susceptible to state dependence.

In the presence of insurance inefficiencies, we can further infer the underlying economic forces that lead to newly-widowed households' labor supply responses to benefit eligibility. Specifically, declines in labor supply would imply that the responses are driven by the sharp change in cash-on-hand and would point to the value of immediate unanticipated liquidity upon an adverse household event. To see this, we need to consider the degree to which benefit eligibility at the transition to widowhood induces a cash-on-hand (or liquidity) effect as compared to an income (or wealth) effect, which we now discuss.

Our empirical analysis compares the behaviors of women at the first year of widowhood as a function of their age in months. As such, per the exact structure of Social Security's survivors benefits, there is a sharp discontinuity in cash-on-hand among newly-widowed women at the precise age-60 eligibility cutoff. This sharp increase in unanticipated liquidity through government transfers would induce a discontinuous decline in labor supply at the age cutoff. On the other hand, there is no discontinuity in newly-widowed households' present discounted value (PDV) of benefits at the eligibility cutoff. Among widows younger than the full retirement age (which is 65 or older for all cohorts), the program is designed to provide an entitlement for the same benefit PDV for a given history of a husband's earnings. Claiming benefits at different ages over 60 involves actuarial adjustments to account for the different length of benefit collection; and widows younger than 60 are entitled for the same PDV of benefits as those who are older, which they can collect starting age 60. Thus, the entitlement formula for benefit PDV is approximately flat around our threshold. Since the PDV of survivors benefits weakly increases in the husband's earnings history, which is weakly increasing the older the household transitions to widowhood, the PDV of survivors benefits at widowhood may display moderate increases in our running variable of the widows' monthly age—but such potential increases are smooth as per the benefit calculation formula.[6] This potential continuous underlying evolution of life-time income from benefits in our running variable—which does not confound the identification of responses but can complicate their interpretation—is the key motivation for our second exercise in which life-time income is pre-determined and constant throughout.

[6] Specifically, as we describe in the institutional details, survivors benefits are a percentage of the husband's potential Social Security retirement benefits. Retirement benefits aim to reflect life-time earnings and are based on a worker's Average Income Monthly Earnings (AIME) over 35 years (in which the worker earned the most). As such, potential increases in benefit PDV over our running variable from an additional month of a husband's earnings are smooth and muted. Note that with less-than-full insurance, one may think up ways in which households could respond in the good state to benefit availability in the bad state. If households reduce their ex-ante self-insurance through savings, due to a lesser need for cash-on-hand if the event occurs, husbands may continuously reduce their labor supply as their wife approaches 60. Such potential responses, which would already be muted through the event's small probability and the 35-year averaging of earnings, may further flatten the potential increase in benefit PDV. This could, if anything, induce an income effect that would push toward an increase in labor supply at the region of the eligibility cutoff, mitigating the observed labor supply reductions.

(ii) Responses by Already-Widowed Households to Anticipated Benefits. Consider an already-widowed household and let us analyze the case in which the following assumptions hold: (1) households are forward-looking and understand the Social Security benefit schedule and rules; (2) there are no liquidity constraints. We analyze a two-period model, where periods are indexed by $t \in \{1,2\}$, to capture a period of benefit ineligibility followed by a period of benefit eligibility. The results extend, of course, to multi-period dynamic models, since they rely on classical Euler conditions that hold more generally. We consider the planning problem of a household that had already transitioned to widowhood prior to the beginning of period 1. We assume that benefit eligibility comes into effect only in period 2, such that $B_1 = 0$ and $B_2 > 0$, and that this benefit schedule is deterministic, and hence can be fully anticipated, at the beginning of the planning period. The model is described in detail in Appendix A.2.

The household maximizes its life-time utility subject to the within-period budget constraints, where the choice of saving or borrowing is unconstrained beyond guaranteeing that consumption is non-negative. At the optimum, widows smooth consumption and leisure across time periods, and the whole planning problem can be rewritten in terms of the present discounted value of life-time unearned income. Hence, the main prediction of this familiar model, which we use as our benchmark, is that of labor supply smoothing: there should be no discontinuity in labor market choices when the anticipated benefits become available. For a given level of the present discounted value of benefits, the household's behavior should not depend on their timing (following a similar logic as that in MaCurdy 1981). It is straightforward to also explicitly incorporate an earnings test similar to that of the Social Security survivors insurance, whereby benefits increase permanently to account for the months in which benefits are withheld if widows' earned income crosses a given threshold (SSA 2018b). If households correctly perceive the earnings test, the qualitative results of our model remain the same.

This benchmark provides us with a clean test for credit market imperfections and the role of anticipated liquidity in our application: comparing household responses to anticipated changes in cash-on-hand against the prediction of identically zero responses in a frictionless world. This is due to a key advantage of our empirical setting, in which the present discounted value of life-time unearned income is completely unchanged and does not depend on the widow's behavior, whereas there is a sharp discontinuity in cash-on-hand.[7] As such, our analysis does not require calibrations of behavioral models against which one should test for excess sensitivity to increases in income as a measure for the effects of liquidity (to isolate them from wealth effects), which are needed in commonly-studied settings that involve some degree of changes in life-time wealth/permanent income. In addition, recall that in this second exercise we analyze the behavior of households that are several years into widowhood, which provides them with the necessary

[7] This advantageous feature also stands in contrast to traditionally-studied age-contingent benefit schemes, such as old-age pensions, in which own continued work can directly affect the PDV of benefits.

time to respond to the shock (e.g., in self-insurance through earnings), to anticipate the benefit receipt (and accordingly plan for the future and for retirement), as well as to make needed financial arrangements (e.g., borrowing). Hence, our setting allows for general forms of delays in adjustment to the household shock and in anticipation of benefits, so that our assessment of the role and value of anticipated liquidity is not confounded by their presence.

Deviations from this benchmark—i.e., from insensitivity of widows' labor supply profile to anticipated benefit timing—would imply that at least one of the underlying assumptions of the model is violated. In our empirical analysis we offer several ways to distinguish between the candidates for the underlying channels: lack of planning, liquidity constraints, and benefit-schedule misperceptions.

3. Data and Empirical Framework

3.1. 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 income tax returns (e.g., Form 1040) and information returns filed by third parties (e.g., Form W-2, Form SSA-1099, and Form 1099-R). We observe exact dates of birth (to determine age-eligibility for survivors benefits by widows) and exact dates of death (to identify spousal death events) using the Social Security Administration (SSA) records. Spousal linkages are established through filing a joint tax return in the year prior to the death event.

Outcomes and Variable Definitions. 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. 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).

We analyze two sets of outcomes: household income and labor supply. We define overall household income as the net pre-tax family income available from any reported source, which broadly follows the recent convention in the literature that uses 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.

For labor supply outcomes, based on data from Form W-2, we study as our primary outcomes of interest wage earnings and labor force participation.[8] We define participation as having positive earnings in a given period. When discussing our findings we emphasize wage earnings, as they comprise an aggregate measure that captures responses on both the intensive margin and the extensive margin. Within our main analysis we also provide complementary figures for retirement behavior, where retiring is defined as having positive earnings in the current period and no earnings in the next period. Since this is a flow outcome that captures changes, responses in it are less informative quantitatively (e.g., in comparisons across settings or across subsamples) as they do not represent the full aggregate effect of eligibility, unlike the cumulative labor supply outcomes that we focus on. Nonetheless, their nature makes them qualitatively valuable in that they can illustrate in a visually clear way the promptness of responses to eligibility.

3.2. Empirical Framework

To estimate the causal effects of Social Security's survivors benefits receipt on households' behavior and economic outcomes, we exploit a sharp discontinuity in the program's benefit-eligibility schedule. In this subsection, we outline the main features of Social Security's survivors benefits program, and we then describe our research design for estimating the effects of the program on American households.

Institutional Details. Surviving spouses become universally eligible for Social Security's survivors benefits at exactly age 60.[9] The benefit amounts that surviving spouses can receive are based on the deceased spouse's potential Social Security retirement benefits, which are themselves determined by the deceased's work history. Specifically, Social Security retirement benefits accrue to individuals whose earnings are subject to Social Security taxes. Generally, to become eligible for retirement benefits, individuals are required to accumulate 40 "credits," which translate to 10 years of work since workers can earn up to 4 credits each year (where, e.g., in 2016, $1,260 in earnings = 1 credit). Survivors benefits are then calculated as a percentage of the deceased's retirement benefits, and this percentage is determined by the surviving spouse's age at the beginning of benefit claiming. The percentage ranges from 71.5% at age 60 to 100% at the widow's full retirement age.[10] Social Security does not notify widows when they become age-eligible for benefits. Rather, widows who had claimed spousal retirement benefits are contacted by the Social Security Administration upon the beneficiary husband's death with a notification of the eligibility rules and their potential entitlement.[11]

[8] Annual income from self-employment is very low among widows (with baselines of $593 in our sample of newly-widowed households and $404 in our sample of already-widowed households) and is therefore not a meaningful margin for responses in our setting. We report the analysis of self-employment in Appendix Table 7.
[9] To be eligible, surviving spouses cannot remarry before age 60. 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.
[10] Claiming of survivors benefits and its timing do not alter widows' schedule of own retirement benefits, which they can become eligible for and transition to at age 62 according to the standard Social Security rules.
[11] We have learned of these practices through former SSA field officers.

By design, the PDV of Social Security's survivors benefits depends on the deceased's earnings history and does not depend on the survivor's earnings history. This feature provides two notable advantages. First, it implies that the PDV of benefit entitlement is fixed from the point of the husband's death onward. Second, it implies there are no actual substitution effects upon eligibility, so the potential impacts of the program on widows' labor supply are likely to operate primarily through a non-distortionary liquidity effect. Still, similar to retirement benefits, 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), benefits are withheld at a specified rate, but are later paid back in the form of increased benefits. Since research has shown that such benefit adjustments may be misperceived as a tax, the earnings test is a program feature that may create a "substitution" effect. We utilize this feature of the program later in the analysis by studying subsamples of households that are infra-marginal to the earnings test.

Research Design. We exploit the sharp eligibility-age discontinuity in Social Security's survivors benefits to identify their causal effects on widowed households. Specifically, we study the patterns of widows' outcomes in the post-shock years as a function of their age, and we conduct causal inference by estimating sharp breaks in trends at the exact eligibility cutoff age of 60.

Importantly, we allow for smooth underlying trends in widows' outcomes. These trends account for any changes that are continuous in age, and would therefore not affect the interpretation of our results. Such changes could be directly attributed to the evolution of the running variable of widows' age, or indirectly attributed to factors correlated with a widow's age in the post-shock years, such as the husband's age at death and the associated time of exposure to the event (with its impacts on the household's realized life-time wealth). One specific implication is that the estimation is not confounded by potential changes to preferences as a result of a spousal death, as long as those are continuous in the widow's age upon the event, since all households are analyzed after having been exposed to spousal death and its direct impacts.

Estimation. We study all widows in the tax records whose husband died in the years 2002-2007. We focus on female surviving spouses since they comprise the vast majority of all widowed households—approximately 80%—throughout the age distribution (see Fadlon et al. 2019). The population-level data allow us to lead our analysis with a graphical representation of the results. We take advantage of survivors' exact dates of birth, and plot means of each outcome variable of interest against the widow's "monthly age." Since tax information is observed as of December in a given year, age is defined as a person's age at the end of the calendar year of observation. 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 (the early eligibility age for standard retirement benefits).

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 Social Security's survivors benefits is

captured when widows are eligible for benefits for the entire calendar year. Specifically, widows who turn 60 in January are eligible for benefits throughout an entire calendar year, as they just turned 60 at its beginning, whereas widows who turn 60 in December are eligible for only one month at most. Hence, it is the behavior of widows in the former group that displays the full-exposure effect. Technically, as these widows turn 60 at the beginning of a year and since age is defined at the end of a calendar year, the effect of being fully exposed to eligibility for survivors benefits is identified by widows whose age at the end of the year is just below 61.

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

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

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 $\{age_{i,t} > 60\}$ 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 closely guided by the graphical analysis of the raw data. For visual clarity in presenting our findings and assessing these assumptions, we combine the graphical analysis and the regression analysis. In particular, we present figures that plot raw means of outcomes by widow's monthly age, and we superimpose the regression lines from the corresponding estimation of 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 trends (captured by the change to the slope, β_3), accounting for the time period that captures full exposure. We also augment the design with a control group of future widows as a robustness check, which we describe in the empirical analysis.

Implementation. We implement our two empirical exercises by analyzing different horizons of post-shock years in the following way. First, we study outcomes of newly-widowed households in the periods just after the event occurs. In choosing these periods we must consider the property that the data are annual and measure values at the end of a calendar year; so that the year of the event (which we index by $t = 0$) is a transitional period since households experience the husband's death at different points during the calendar year. The first period in which all sample households have been fully exposed to the spousal death event—and involves a minimal degree of benefit anticipation to the extent that the exact timing of death is unpredictable—is therefore $t = 1$. We also include $t = 2$ in the analysis for increased statistical

power and visual clarity, though the results remain the same when only $t = 1$ is considered (see Appendix Table 1).

Second, we study the responses of already-widowed households. Specifically, we analyze the behavior of widows using observations from periods $t = 6 - 10$ after the spousal death, so that among all the included observations the husband had died at least 5 years in the past. This 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 50 and 54.

4. The Effects of Social Security's Survivors Benefit Receipt

We now turn to estimate the effects of Social Security's survivors benefits on widows' labor supply and household income. We first analyze their protective role on newly-widowed households in the immediate years following a spousal death, and we then study responses by already-widowed households to anticipated benefit receipt.

4.1. Responses in Immediate Post-Shock Periods to Unanticipated Benefits

Benefit Claiming. We begin by looking at the claiming behavior of survivors benefits by newly-widowed women, which constitutes a first stage in our analysis. Panel A of Figure 1 first 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 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). 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 (see column 1 of Table 1).[12] 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 increased claiming begins, with the average amount of benefits transferred to survivors

[12] The non-zero take-up rate prior to age 60 is attributable to disabled survivors who are eligible for benefits when they reach age 50 and surviving spouses with dependent children under age 16.

reaching its full effect by age 61. At that point, the average increase in benefits, including zeros for those not claiming, amounts to $5,605 (see column 2 of Table 1).[13]

Household Income. To evaluate the impact of eligibility for survivors benefits on newly-widowed women's 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 fast rate over the eligibility range, until it reaches the full-exposure effect as displayed by widows of age 61. The net increase in income totals to $4,804 (see column 3 of Table 1), which represents an increase in family income of 11.4%.[14]

Labor Supply. Next, we investigate how benefit eligibility affects the labor supply of widows, which constitutes an important dimension of household gains from social insurance through changes in the consumption of leisure. For visual clarity of response promptness, we first plot our supplementary flow outcome of widows' retirement behavior. Panel D of Figure 1 displays a clear and considerable jump in widows' retirement rate at benefit eligibility.[15] To evaluate how these changes in flows translate to the cumulative labor supply effects which are the focus of our interest, we study widows' labor force participation rates and wage earnings. The full exposure effect on labor force participation amounts to a decline of 2.87 percentage points (see panel E of Figure 1 and column 4 of Table 1). Overall, widows' labor supply responses amount to an average decrease of $1,751 in annual earnings (see panel F of Figure 1 and column 5 of Table 1).

4.1.1. Implications

By design, this analysis identifies the effects of benefit receipt in the immediate post-shock years. Both households just below and just above the threshold would be eligible for benefits in future periods, but only those above the threshold are eligible for receiving benefits right after the shock's realization. Therefore, these effects capture the protective insurance role of survivors benefit receipt against the immediate adverse financial consequences of a spousal death. As such, the results underscore eligible households' valuation of the protection provided by the federal government against mortality events. The Social Security survivors benefits program generates gains to newly-affected households both through significant increases in household income and through meaningful increases in the consumption of leisure due to a mitigated need to self-insure.

---

[13] It is useful to note that by construction and due to the identification strategy of exposure to eligibility, binary outcomes tend to exhibit immediate breaks in both levels and slope and continuous outcomes exhibit an immediate break in the slope. Also, binary flow variables exhibit a tendency to revert back to pre-eligibility levels once the incentives of full exposure have come into effect, as these outcomes essentially capture the "derivative" of a given behavior.

[14] The counterfactual level is visually represented in the figure, and using equation (1) it is estimated to be $\beta_0 + \beta_1 \times (11/12) = 42{,}456 + (-388) \times (11/12) = 42{,}100$, so that the effect on income is 11.4% compared to the counterfactual.

[15] The estimate for the full exposure effect on retirement is 0.01829 (with s.e. 0.00188) on a counterfactual of 0.05704.

Insurance Inefficiencies and Value of Benefits to Ineligible Households. The results point to a clear deviation from the first-best benchmark described in Section 2, indicating notable allocative inefficiencies in the large life insurance market. Even more, the degree of this deviation in labor supply responses has direct implications for the excess value ineligible newly-widowed households would assign to a dollar of benefits through survivors insurance relative to the benchmark of eligible newly-widowed households.

To see this, consider our conceptual framework from Section 2 (which is presented in more detail in Appendix A.1). Let $u_2^b(c_2)$ represent the wife's flow utility from consumption in the bad state, let $v_2^b(l_2)$ represent her disutility from labor; and, for any variable x, define $x(0)$ to be the outcome for a just-ineligible newly-widowed household, and $x(1)$ to be the outcome for a just-eligible newly-widowed household. The value of a dollar is exactly given by the marginal utility from consumption, $u_2^{b\prime} = v_2^{b\prime}/w_2$ (where w_2 is the widow's wage rate). Hence, the excess value of transferring benefits to ineligible newly-widowed households on the margin is captured by the relative gap in the marginal disutility from labor, $\frac{v_2^{b\prime}(l_2^b(0))-v_2^{b\prime}(l_2^b(1))}{v_2^{b\prime}(l_2^b(1))}$, which can be approximated by $\varphi\frac{l_2^b(0)-l_2^b(1)}{l_2^b(1)}$, where $\varphi \equiv \frac{v_2^{b\prime\prime}(l_2^b(1))}{v_2^{b\prime}(l_2^b(1))} l_2^b(1)$ is the curvature of labor disutility. As this gain is proportional to our estimated causal effect of benefit eligibility on labor supply, $\frac{l_2^b(0)-l_2^b(1)}{l_2^b(1)}$, our results point to potentially meaningful valuation of benefits by ineligible widows.[16] For example, calibrating the utility parameter φ to equal 1 as is the case under quadratic labor disutility, the findings suggest that the excess value of an additional dollar to ineligible widows is approximately 9.32% ($1,751 on a counterfactual of $18,787).[17]

Recall from the discussion of the benefit structure in Section 2 that, in the presence of life insurance inefficiencies evidenced by labor supply responses, the effects of eligibility are driven by discontinuities in cash-on-hand from benefit availability whereas there is no discontinuous change in life-time income. Hence, the results point to an important role of the immediate liquidity provided by transfers following the realization of the adverse household event, which allows under-insured households to smooth consumption across states of nature.[18]

Assessing the Degree of Income Flow Coverage. To understand the scope of the program, we gauge the extent to which households eligible for Social Security's survivors benefits are protected against the

[16] We note that we only point to gross gains from any consideration of changes to the benefit schedule since the value added of our analysis lies there. We do not allude to the cost side as the Social Security Administration already has mechanisms in place for scoring the cost to the system of various changes to the benefit structure.

[17] Fadlon and Nielsen (2018) show how φ can alternatively be estimated using labor supply elasticities. The analysis here is merely an application of their analysis across states of nature to an analysis across states of eligibility.

[18] There are two additional pieces of evidence that favor the liquidity interpretation. First, we have found there are no lingering effects of eligibility for benefits upon spousal death. We show in Appendix Table 3 that longer-run outcomes of widows, e.g., at ages 67-69 (given the range of our data), do not depend on eligibility for benefits when the shock occurs. Second, we split the sample into high-liquidity and low-liquidity households based on the median level of lagged unearned income, and we find to some degree larger labor supply responses among lower-liquidity households. See Appendix Table 4.

financial burden imposed by spousal death. For this exercise, we need to assess the effects of the spousal death event itself on eligible households. To do so, we utilize an event-study approach that exploits the potential randomness of the particular timing at which a death event was realized within a short period. Specifically, we construct counterfactuals for affected households using households that experience the death event at a later period, and we correspondingly assign a placebo shock for control households in the year at which the treatment group experience their actual shock.[19] Full details on the design and its identifying assumptions appear in Fadlon and Nielsen (2017) and investigation of its validity within our setting (in terms of comparability and pre-trends) is provided in Fadlon et al. (2019).

We assess the impact of spousal death on overall annual household income among women who at the year of observation were of the eligible ages 60-61. We accompany the analysis with a similar exercise for women of the ineligible ages 58-59. We quantify the degree of income coverage by estimating the standard difference-in-differences equation of the following form:

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

In this regression, $treat_i$ denotes an indicator for whether a household belongs to the treatment group, $post_{i,t}$ denotes an indicator for whether the observation belongs to post-shock periods ($t = 1,2$) or pre-shock periods ($t = -2,-1$), and $X_{i,t}$ is a vector of controls that includes age indicators and calendar year fixed effects. The parameter δ represents the average effect of the event on households' income.

The results in panel A of Table 2 indicate that eligible households experience a decline of $22,803 in household income, which represents a decline of 33.5%. We interpret this finding through the lens of commonly used adult equivalence scales to account for the household's compositional change. Specifically, the modified OECD equivalence scale of 0.67 and the square-root scale of 0.71 suggest that declines in household income following a spousal death on the order of 29-33 percentage points could be interpreted as full compensation.[20] Hence, the evidence is consistent with close-to-full compensation for income losses from spousal deaths among eligible families. However, this assessment relies on the accuracy of equivalence scales in capturing economies of scale within the household. To avoid this issue, we assess the degree of coverage further by additionally analyzing widows' labor supply behavior, which is an input that directly enters an individual's utility and does not require scaling. Fadlon and Nielsen (2018) demonstrate that, under certain conditions, labor supply responses to a spousal death which act as self-insurance can capture the extent to which households lack formal insurance and would value government benefits. Panel

[19] Specifically, for this illustration, we draw a 20% random sample of men who died between the years 2002 and 2007 and who were married in the year prior to their death, and we study the effects on their surviving widows. Based on the time range of the data, our treatment group is composed of women whose husband died in years 2002-2003 and our control group is composed of women whose husband died in years 2006-2007.

[20] Of course, full income compensation (equating equivalence-scale adjusted income levels across states) and full insurance (equating marginal utility across states) are not the same, specifically when preferences are state dependent.

A of Table 2 indicates that, in response to a spousal death, eligible widows exhibit no changes in labor supply following the shock. This suggests that self-insurance through labor supply is not required for those with access to Social Security's survivors benefits, which further supports the view that the program provides close-to-full compensation to eligible households. Note that this labor supply approach is not assumption free either, and it requires taking a stand on whether and how labor disutility may change as a result of the death of a spouse. In contrast to eligible households, panel B of Table 2 points to the financial vulnerability of ineligible households: they experience a significantly larger income decline and exhibit a non-negligible increase in labor supply (of 7.5%), consistent with a need to self-insure against the income loss imposed by the mortality event.

Liquidity versus Substitution Effects. 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 they 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. This is because substitution effects are socially suboptimal responses to distortionary wedges between private and social marginal costs, while liquidity effects are socially beneficial responses to the correction of insurance and credit market failures (Chetty 2008).

Unlike Social Security retirement benefits, survivors benefits are generally decoupled from own labor supply, so there are presumably no direct distortions in the incentives to work. In that sense, the estimated effect on widows' labor supply could therefore be attributed to a welfare-beneficial liquidity effect. Intuitively, the liquidity provided by the social insurance program attenuates the need for costly self-insurance through labor supply, leading to efficient increases in the consumption of leisure toward the optimal allocation in the absence of a market failure.

Nonetheless, research in the context of Social Security retirement benefits has suggested that individuals may misperceive earnings tests as distortionary income taxation, even though transfer reductions due to the earnings test are paid back to beneficiaries after they reach full retirement age (Liebman and Luttmer 2012, 2015; Brown et al. 2013). We therefore proceed by analyzing a subsample of households for whom only a non-distortionary effect is likely operative in their responses. To do so, we study the labor supply of widows whose pre-shock earnings were below the annual earnings test thresholds.

The results are reported in Appendix Table 2 (and Appendix Figure 1). Similar to the analysis of the full sample, we find meaningful declines in labor force participation among the current subsample of households. Widows with labor income below the earnings test thresholds exhibit a decline of 2.42

percentage points on a counterfactual baseline of approximately 30.[21] As there is likely no moral hazard component involved in their responses, this points to a meaningful non-distortionary (or corrective) increase in the consumption of leisure. Notably, among these households, which represent a large share of our sample (43%), the relative response in earnings is significantly larger. It amounts to a decline of $1,065 on a baseline of $3,978, which translates to 27% and points to even greater valuation of insurance benefits among the current subsample. This finding is consistent with the notion that low-earnings spouses are more exposed to financial risk and are effectively less well insured against spousal death, as they generate little income on their own (Fadlon and Nielsen 2017).

4.1.2. Robustness

We provide several robustness checks for the estimated effects. Appendix Table 5 summarizes these results, and complete details are provided in the notes of the table. First, to account for potential confounding changes of a general source 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' observations, the treatment group narrows to a (majority) subset of our original treatment group, so that estimations should naturally not perfectly align across designs. Still, the findings are very similar (see panel A). Second, although our high-frequency graphical analysis shows that effects kick in promptly at the monthly age of 60, we may want to pay particular attention to age 59.5 after which withdrawals from private savings accounts are no longer penalized. However, the households that we analyze may be less likely affected by this change in incentives from the point of view of their overall portfolio, since the death event allows for non-penalized distributions from the deceased husband's accounts. Indeed, we find that less than 4% of all analyzed households make any withdrawals that are indexed with reasons other than "death," so that constraining the analysis to the remaining 96% of households provides similar results (see panel B). To alleviate remaining concerns, we additionally analyze households who from an ex-ante standpoint are less likely to be affected by this change in withdrawal incentives, as they did not make contributions to savings accounts in the pre-shock periods. This analysis also provides similar results (see panel C).

[21] It may be useful to compare this response to that of the overall sample as a natural benchmark. To provide a comparison across more similar moments, we convert the labor supply effects into elasticities. Specifically, we estimate the percent change in participation divided by the percent change in household income that is attributed to government benefits. The overall sample and the current subsample display very similar elasticities. In the full sample, the elasticity of labor force participation with respect to government-provided income is $\frac{-0.02866/0.61215}{5{,}605/42{,}100} = -0.35$; and it is $\frac{-0.02424/0.30}{7{,}258/34{,}043} = -0.38$ in the sample of widows whose earnings were below the earnings test thresholds in the pre-shock years. We note that this exercise is only suggestive due to potential heterogeneity in labor supply responses along the earnings distribution.

4.2. Responses to Anticipated Benefits by Already-Widowed Households

We now proceed to study whether and to what extent widowed households' labor supply responds to cash-on-hand via anticipated benefit receipt, as compared to the frictionless benchmark of labor supply smoothing. Recall that we analyze already-widowed households, who had time to adjust to the shock (e.g., in self-insurance through earnings or assets), to anticipate the benefit receipt, and to make necessary financial arrangements (e.g., borrowing). This exercise hence inherently focuses on the particular role of predictable changes in cash-on-hand and benefit timing for given life-time wealth, and it isolates the effects and value of liquidity provided by government transfers to vulnerable older families.

Results. Table 3 first verifies the existence of a first stage, indicating that the take-up of survivors benefits amounts to 34 percentage points which translates to an increase of $3,655 in annual transfers from the government (see columns 1 and 2). Then, studying labor supply outcomes, we find clear deviations from the frictionless benchmark. Notably, there is a pronounced jump in widows' retirement rate at benefit availability (see panel A of Figure 2).[22] The full exposure effect on labor force participation totals to a decline of 3.12 percentage points, with considerable overall labor supply decreases that amount to a decline of $1,938 in annual labor earnings (see panels B-C of Figure 2 and columns 4-5 of Table 3).

The evidence is inconsistent with the conjecture that these responses can be explained by misperceptions of the earnings test. We repeat the analysis for widows whose lagged earnings were below the earnings test thresholds, and we find large labor supply responses among them as well: decreases of 10% in participation (2.8 pp on a baseline of 27.6) and 20% in earnings ($497 on a baseline of $2,472). See Appendix Table 6 (and Appendix Figure 2).

4.2.1. Implications

The results point to a meaningful reduction in labor supply in response to predictable increases in cash-on-hand at the benefit eligibility age of 60. In fact, a simple calibration suggests that the representative household's responses constitute about 70% of the hypothetical response under a complete hand-to-mouth benchmark (see Appendix B).[23] This significant deviation from the frictionless benchmark of labor supply smoothing from Section 2 has two sets of implications which we now discuss.

Normative Implications. First, the results indicate considerable allocative inefficiencies in credit and liquidity among U.S. households. The findings underscore that the timing of benefits and liquidity play a considerable role and can have direct value in allowing households to smooth consumption across time periods. The evidence suggests there are potential gains from changing the benefits' timing to inject liquidity earlier and smooth their distribution over the course of widowhood (for a given program size).

[22] This effect averages to 0.01997 (with s.e. 0.00165) on a counterfactual baseline of 0.04411.
[23] This is in line with findings from Card et al. (2007) for job searchers in Austria.

That is, holding the present discounted value of benefits unchanged, merely transferring benefits from later periods to earlier periods could get widowed households closer to first-best smoothing.

The gross marginal gains from such budget-neutral retiming of benefits are exactly captured by the extent to which households fail to smooth their behavior, in either consumption or leisure. Based on our conceptual framework from Section 2 detailed in Appendix A.2, let x_t represent the value of any variable x in period $t \in \{1,2\}$, where period 1 captures a period of benefit ineligibility followed by period 2 of benefit eligibility; and let $u(c)$ and $v(l)$ represent the widow's flow utility from consumption and disutility from labor, respectively. The household's gains from benefit retiming are captured by the relative gap in $u'(c_t)$, or equivalently in $v'(l_t)$, across periods. In the context of our model, this would translate to $\frac{v'(l_1)-v'(l_2)}{v'(l_1)}$, which can again be approximated by $\varphi \left|\frac{l_2-l_1}{l_1}\right|$ where $\varphi \equiv \frac{v''(l_1)}{v'(l_1)} l_1$ is the curvature of the labor disutility function.[24]

That is, the gain from incrementally smoothing the distribution of benefits across periods is captured by gain from incrementally smoothing labor supply across periods. The latter is proportional to the meaningful labor supply responses to benefit availability that we find, which are on the order of 9.4% ($1,938 on a counterfactual of $20,566). The evidence points to even greater gains among low-earnings households (those with earnings that fall below the earnings test). For these households, who comprise a large share of 52% of all households in our sample, we have found that the relative responses amount to 20%. Importantly, the gains from a smoother benefit profile are similar irrespective of the reason households fail to smooth their behavior; in particular, whether it is driven by lack of forward-looking behavior or by liquidity constraints.

Positive Implications and Response Channels. Second, our findings have implications for the mechanisms that underlie widows' labor supply responses to transfers, which can be also informative more generally for the channels that govern vulnerable older Americans' retirement decisions. Since households meaningfully deviate from the frictionless benchmark, the results are consistent with either myopia and lack of forward-looking behavior or with liquidity and borrowing constraints (and credit market imperfections) as driving channels.[25] To further investigate the source of this deviation, we offer tests that aim to distinguish between these potential channels.

Forward Looking. We first examine whether the responses can be explained by complete myopia and lack of forward-looking behavior among our sample of households. To do so, we exploit the unique feature of Social Security's survivors benefits program that, to be eligible, surviving spouses cannot remarry

[24] This is similar to Fadlon and Nielsen's (2018) analysis of transferring resources across states of nature but applied to transferring resources across periods.

[25] Recall from Section 3 that these already-widowed households are not notified by Social Security once they become eligible for benefits at age 60. Hence, their benefit take-up rate itself exactly at the cutoff points to knowledge of the program and to anticipation of benefit receipt prior to actual eligibility.

before age 60; if they do remarry before reaching the eligibility age, they lose their entitlement for survivors benefits altogether. This gives rise to an empirical test for the presence of planning. Specifically, we study whether there is strategic timing of remarriage in the form of increased rates just after age 60. Evidence of such responses would be generally inconsistent with myopia since it requires planning.

Figure 3 shows clear evidence in support of strategic timing of remarriage. The break in the trends is visible exactly at age 60, where the full-exposure effect amounts to an increased remarriage hazard rate of 0.893 percentage points (see column 6 of Table 3) on a counterfactual of 0.819. Consistent with optimal responses to incentives (and with economic theory; see, e.g., Persson 2017), we also show that the sample of widows who likely strategically time their remarriage, takes up benefits at a higher rate and receives higher average benefits from the program as compared to the overall sample (see columns 1-2 of Table 5 compared to columns 1-2 of Table 3).

Liquidity Constraints. Next, we investigate if there is evidence that liquidity constraints play a role. To this end, we study whether household labor supply responses vary by the degree of liquidity as proxied by lagged unearned income. We split households by the sample median, and we analyze our labor supply outcomes for each subsample.

Panels A and B of Table 4 summarize the labor supply responses, in both participation and earnings, among households with liquidity levels below and above the median. The results show considerable differences in the effects of benefit availability across the two subsamples. Despite receiving economically similar levels of benefits from the government (see column 1), lower-liquidity households display meaningfully larger labor supply reductions in response to availability of liquidity (see columns 2 and 3). Moreover, we also find that the highest-liquidity households (within the top quartile) do not respond in labor supply and behave as the frictionless model predicts (see panel C).[26] These results are consistent with low-liquidity households' inability to smooth the consumption of leisure prior to the actual receipt of the anticipated benefits, and with high-liquidity households' ability to use their own resources to smooth consumption and behavior. The findings therefore not only indicate that liquidity can play a meaningful role, but they also show that liquidity provides an explanation that can account for the estimated responses to anticipated cash-on-hand.

Lastly, we further take advantage of the richness and scope of the data to focus on subsamples of households who are likely forward-looking, but are still potentially subject to liquidity constraints. We study responses of two such subsamples using specification (1), albeit with naturally lower precision due to sample sizes. The first is the sample of widows who remarry at or just after the year they turn 60. The

[26] This is not the case in the context of our first exercise in which households are also subject to the immediate adverse financial effects of the death event, whereas the current exercise isolates the effect of liquidity. Notice that these highest-liquidity households do have "room to respond" in labor supply reductions, as their baseline participation rate is 45 pp.

second is the sample of households who have Individual Retirement Accounts (IRAs), but whose (lagged) account balances are low (i.e., below the median).[27] For both subsamples, we find evidence of meaningful declines in labor supply in response to the availability of anticipated survivors benefits from Social Security (see columns 3-6 of Table 5). Consistent with the recent literature on the "wealthy hand-to-mouth" (Kaplan and Violante 2014), we also find meaningful responses to anticipated benefit receipt by households with low liquidity (below median lagged unearned income) who hold illiquid assets as proxied by homeownership (see columns 7-8 of Table 5).

Overall, the evidence points to liquidity constraints, rather than benefit misperception or myopia, as the likely operative mechanism that underlies the large estimated responses to predictable increases in cash-on-hand from the receipt of anticipated benefits.

5. Conclusion

Exploiting the age discontinuity in eligibility for Social Security's survivors benefits, we find important impacts of this large program on the behavior and financial security of American families. Our findings highlight considerable allocative inefficiencies in both the life insurance market and the credit market. The evidence first points to a high valuation of survivors benefits in protecting Americans against the immediate adverse economic impact of spousal mortality events, which comes in the form of considerable increases both in net household income and in the consumption of leisure by widowed households. Second, the evidence firmly points to households' particular valuation of liquidity provided by the federal government, both in terms of its insurance value in smoothing consumption across states of nature and in terms of its intertemporal value in smoothing consumption over time. The results also indicate that the effects of anticipated changes in cash-on-hand on household behavior are likely due to liquidity constraints, rather than the lack of planning. With these conclusions, our findings suggest potentially important gains from providing coverage to ineligible families and from a smoother benefit profile. More broadly, the results shed light on vulnerable American households' retirement behavior and response mechanisms to receiving transfers at older ages.

[27] Among those with below-median IRA balances the mean level of balances is $8,596. Note that initiation of IRA accounts is generally more likely to involve some degree of active savings choices (whereas initiation of 401k accounts could be subject to passive behavior to a higher degree due to employer-based policies such as defaults).

References:

Aizer, A., Eli, S., Ferrie, J., and Lleras-Muney, A. (2016). "The Long-Run Impact of Cash Transfers to Poor Families." *American Economic Review*, 106(4): pp. 935-971.

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 Grace 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, Chetty, Raj, and Weber, Andrea (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 Finkelstein, Amy (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.

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

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, Shanthi P. Ramnath, and Patricia K. Tong (2019). "Mortality Risk and Financial Consequences in the U.S."

Ganong, Peter, and Pascal Noel (2018). "Liquidity vs. Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession." NBER Working Paper No. 24964.

Giupponi, Giulia (2018). "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, D., Paulson, A., and Powers, K. (2018). "What Explains the Decline in Life Insurance Ownership?" *Economic Perspectives*, 8: pp. 1-20.

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

Kaplan, Greg, and Giovanni L. Violante (2014). "A Model of The Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): pp. 1199-1239.

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

Persson, Petra (2017). "Social Insurance and the Marriage Market."

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

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

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 Social Security's Survivors Benefit Receipt on Newly-Widowed Households

A. Social Security Benefit Claiming

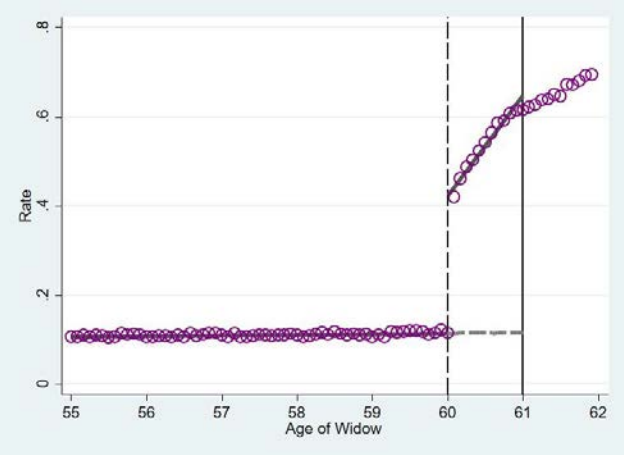

B. Social Security Benefit Amounts

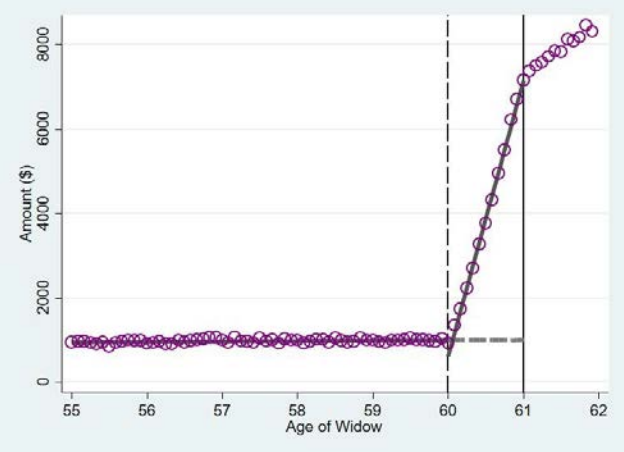

C. Overall Net Household Income

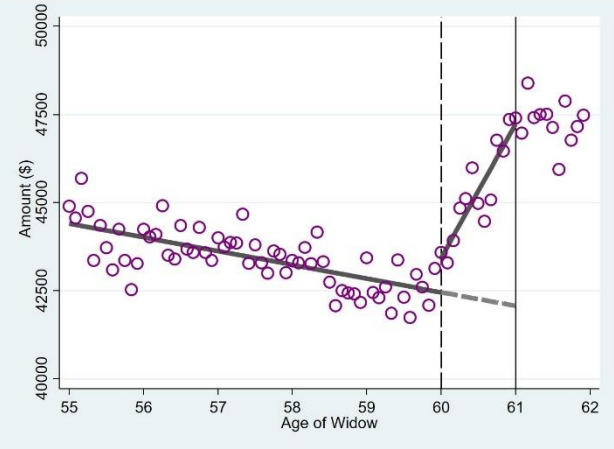

Figure 1: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Households (*continued*)

D. Retirement Rate

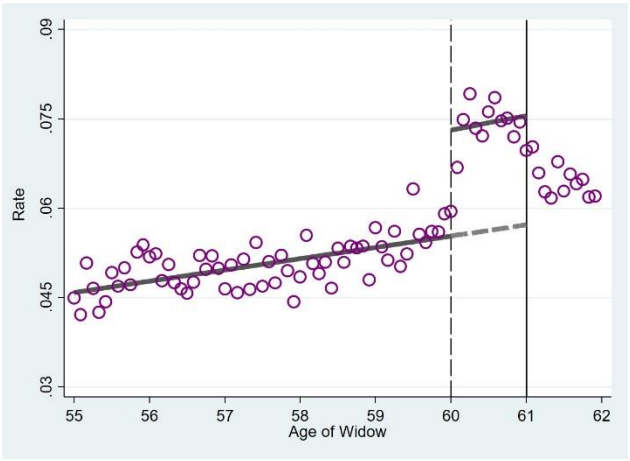

E. Labor Force Participation

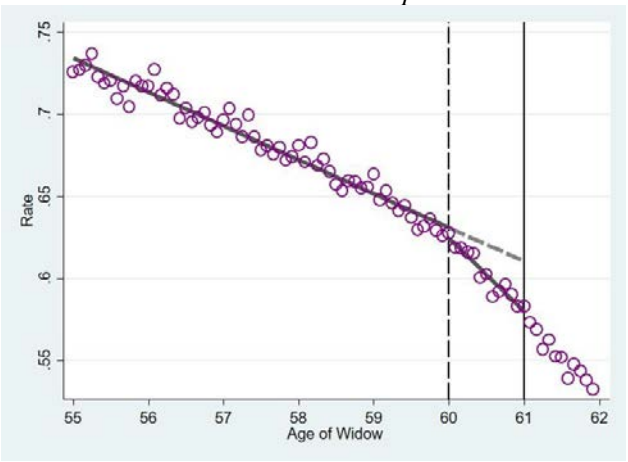

F. Wage Earnings

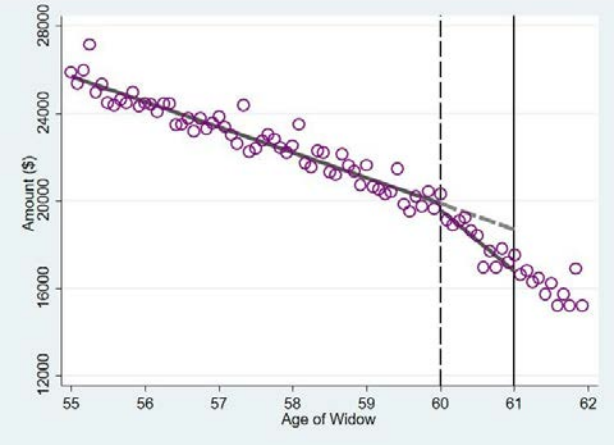

Notes: These figures plot various household outcomes in the years just after a husband's death event ($t = 1,2$) as a function of the surviving spouse's age in months. The purple circles represent means of raw data for each "monthly age" bin. The solid gray lines plot the piecewise linear fit using equation (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 (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 (marked by the vertical solid black line).

Figure 2: Labor Supply Responses to Anticipated Benefit Receipt by Already-Widowed Households

A. Retirement Rate

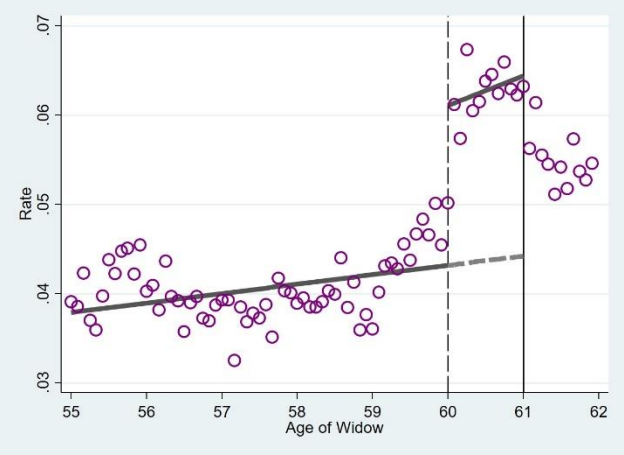

B. Labor Force Participation

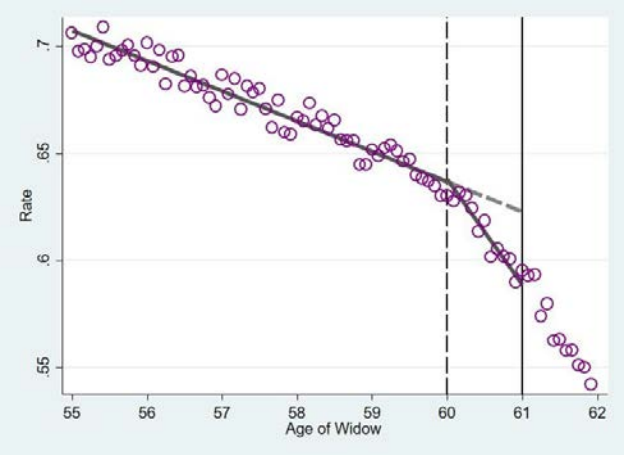

C. Wage Earnings

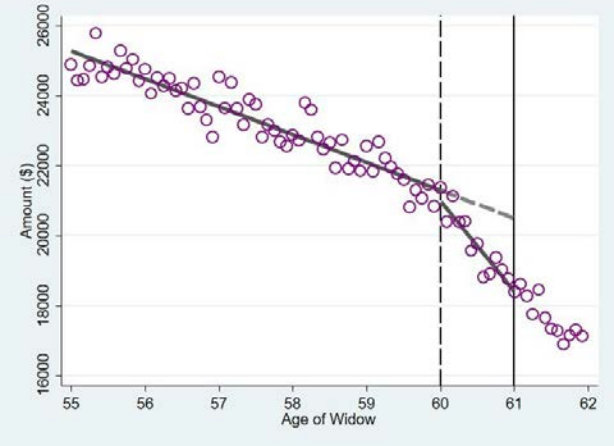

Notes: These figures plot various labor supply outcomes of already-widowed households (using observations from periods $t = 6 - 10$ following the spousal death) as a function of the surviving spouse's age in months, to display responses to anticipated benefit receipt. The purple circles represent means of 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 (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 (marked by the vertical solid black line).

Figure 3: Responses to Anticipated Benefits—Timing of Remarriage

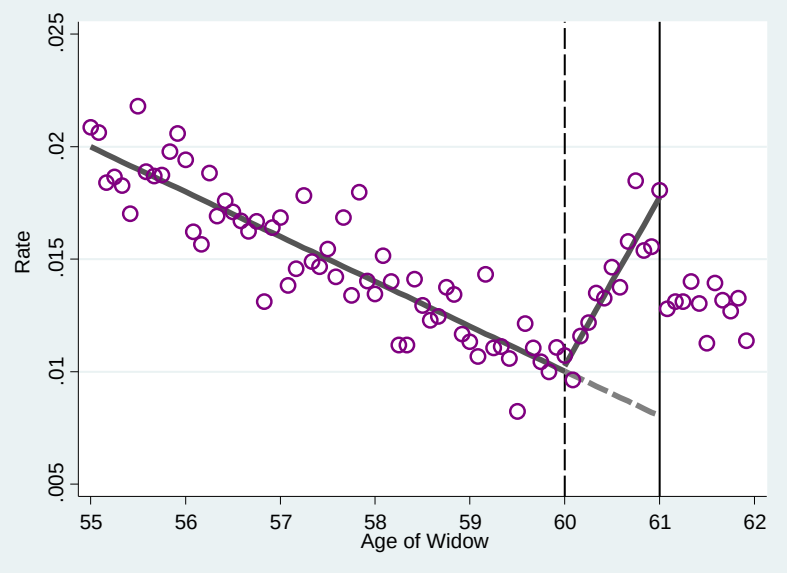

Notes: This figure investigates the presence of forward-looking behavior based on the feature of Social Security's survivors insurance that, to be eligible, surviving spouses cannot remarry before age 60. We study widows' remarriage rate as a function of age among widows who were single in the lagged period. The sample includes already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. The purple circles represent means of 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 (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 (marked by the vertical solid black line).

Table 1: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Households

	Social Security Benefits		Overall Net Household Income	Labor Supply	
	Claiming Rate	Benefit Amounts		Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
Full-Exposure Effect	0.51351***	5,605***	4,804***	-0.02866***	-1,751***
	(0.00336)	(39)	(343)	(0.00349)	(301)
Number of Obs.	504,104	504,104	504,104	504,104	504,104
Number of Clusters	293,857	293,857	293,857	293,857	293,857

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1). It provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations from the immediate post-shock years ($t = 1,2$). We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2: Effects of Spousal Death

	Household Income	Wage Earnings
	(1)	(2)
Panel A: Widows of Ages 60-61		
Treat x Post	-22,803***	-652
	(1,213)	(531)
Counterfactual	68,072	15,927
Percent Change	-33.5	-4.1
Number of Obs.	55,478	55,478
Number of Clusters	41,626	41,626
Panel B: Widows of Ages 58-59		
Treat x Post	-29,951***	1,354**
	(1,265)	(653)
Counterfactual	72,253	18,046
Percent Change	-41.5	+7.5
Number of Obs.	52,324	52,324
Number of Clusters	39,273	39,273

Notes: This table reports difference-in-differences estimates for changes in household outcomes in response to mortality events using specification (2). It is based on an event-study design that exploits the potential randomness of the particular timing at which a death event was realized within a short period, so that we construct counterfactuals for affected households using households that experience the death event at a later period. Panel A includes observations of women who at the year of observation were of the eligible ages 60-61. Panel B includes observations of women who at the year of observation were of the ineligible ages 58-59. We include as controls age indicators and calendar year fixed effects, and we report robust standard errors clustered at the household level. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 3: Responses to Anticipated Benefit Receipt by Already-Widowed Households

	Social Security Benefits		Overall Net Household Income	Labor Supply		Remarriage
	Claiming Rate	Benefit Amounts		Participation	Wage Earnings	
	(1)	(2)	(3)	(4)	(5)	(6)
Full-Exposure Effect	0.33598***	3,655***	2,911***	-0.03118***	-1,938***	0.00893***
	(0.00308)	(36)	(250)	(0.00302)	(197)	(0.00089)
Number of Obs.	544,223	544,223	544,223	544,223	544,223	485,798
Number of Clusters	226,701	226,701	226,701	226,701	226,701	206,903

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1). It provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations of already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. Column 6 additionally constrains the sample to estimate remarriage rates among those who were single in the previous period. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Responses to Anticipated Benefit Receipt by Degree of Liquidity

	Social Security Benefit Amounts	Labor Supply	
		Participation	Wage Earnings
	(1)	(2)	(3)
Panel A: Low Liquidity			
Full-Exposure Effect	3,835***	-0.04004***	-2,884***
	(46)	(0.00394)	(339)
Number of Obs.	270,990	270,990	270,990
Number of Clusters	129,023	129,023	129,023
Panel B: High Liquidity			
Full-Exposure Effect	3,442***	-0.00991**	-598***
	(53)	(0.004319)	(226)
Number of Obs.	273,233	273,233	273,233
Number of Clusters	131,382	131,382	131,382
Difference Low-High	393***	-0.03013***	-2,286***
	(71)	(0.00595)	(418)
Panel C: Highest Liquidity			
Full-Exposure Effect	3,390***	-0.00418	269
	(75)	(0.00614)	(342)
Number of Obs.	136,765	136,765	136,765
Number of Clusters	71,149	71,149	71,149

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1) for subsamples with varying degrees of liquidity. We proxy for the degree of liquidity using lagged unearned income. Panel A and panel B split the observations into high liquidity and low liquidity based on the sample median, and panel C includes observations from the top quartile of liquidity. For each subsample, the table provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations of already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Table 5: Labor Supply Responses to Anticipated Benefit Receipt— Various Subsamples

	Remarried at or Just after 60				Positive but Low IRA		Low-Liquidity Homeowners	
	Social Security Benefits		Participation	Wage Earnings	Participation	Wage Earnings	Participation	Wage Earnings
	Claiming Rate	Benefit Amounts						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Full-Exposure Effect	0.50079***	5,714***	-0.05111*	-5,292***	-0.04499***	-2,893***	-0.02446***	-2,386***
	(0.02623)	(329)	(0.02874)	(1,470)	(0.00879)	(527)	(0.00500)	(628)
Number of Obs.	7,042	7,042	7,042	7,042	61,976	61,976	118,849	118,849
Number of Clusters	6,980	6,980	6,980	6,980	30,164	30,164	59,648	59,648

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1) for different subsamples of households. Columns 1-4 analyze the sample of widows who remarry at or just after the year they turn 60. Columns 5-6 analyze the sample of households who have Individual Retirement Accounts (IRAs), but whose (lagged) account balances are low (i.e., below the median). Columns 7-8 analyze households with low liquidity (below median lagged unearned income) who hold illiquid assets as proxied by homeownership. Homeownership is defined based on whether the widow received Form 1098 indicating payment of mortgage interest. The table provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations of already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix

A. Models of Household Behavior

In this section we describe first-best benchmarks for each of our exercises to provide anchors for the effects that we estimate in the empirical analysis, as well as to provide tests for market allocation inefficiencies. We use simplified models of household behavior in the context of survivors benefits to make our points, where some main extensions are straightforward as we describe.[1]

A.1. Responses by Newly-Widowed Households to Unanticipated Benefits

Setup. Consider the decisions of a two-person household, which consists of individuals 1 and 2, in a world with two states: a "good" state, state g, and a "bad" state, state b, in which member 1 dies and member 2 becomes a widow. There are two time periods: period 0 in which households are just-ineligible for benefits if the bad state occurs (i.e., the wife is below age 60), and period 1 in which households are just-eligible for benefits if the bad state occurs (i.e., the wife is over 60). We set the discount rate and the interest rate to zero for simplicity and since we analyze periods that are minimally far apart.[2] For any variable x, we define $x(0)$ to be the outcome for a just-ineligible household (i.e., in period 0), and $x(1)$ to be the outcome for a just-eligible household (i.e., in period 1). We let the subscript $i \in \{1,2\}$ refer to the household member and the superscript $s \in \{g, b\}$ refer to the state of nature.

We consider the planning problem of households at the beginning of period 0, starting from the good state. Conditional on arriving at 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 $\mu^g + \mu^b = 1$). Choice variables require tracking the household's history in terms of state of nature realizations to capture state-contingent plans. Hence, for household choices we use the history vector superscript h_t where $t \in \{0,1\}$, so that $h_0 \in \{g; b\}$ and $h_1 \in \{(g, g); (g, b); (b, b)\}$. With this notation, $c_2^{b,g}(1)$, for example, would denote the wife's consumption in period 1 for a household that transitioned to widowhood at that period.

Household Budget. Denote by $c_i^{h_t}$ and $l_i^{h_t}$ the individual consumption and labor supply of member i in history h_t, respectively. Let $\bar{A}$ represent the household's baseline wealth and non-labor income, and denote by A^{h_t} the household's state-contingent wealth and non-labor income inclusive of premiums to and transfers from any private insurance arrangement, as well as any informal insurance arrangements across

[1] Additionally, Fadlon and Nielsen (2018) provide and discuss other extensions and generalizations to the simple model analyzed here, including the labor force participation decision counterpart, a dynamic life-cycle model, general choice variables, alternative assumptions about the household's preference structure (with an explicit analysis of different types of state dependence and preference complementarities/non-separabilities), different approaches to modeling the household's behavior (i.e., collective or unitary), means-testing in government transfers, and the presence of household public goods and economies of scale in the household's consumption technology.

[2] Within our empirical application we account for these dynamic aspects by allowing for an underlying trend.

relatives, etc. We denote i's labor income by $z_i^{h_t} = w_i l_i^{h_t}$, where w_i is the (net-of-tax) wage rate. Additionally, let $B^{h_t}(t)$ represent benefits from the government; so that it is positive for households who are in the bad state in period 1, and it equals 0 otherwise. This age/contingency-dependent schedule is predictable at the beginning of the planning problem. Finally, the household can also make savings decisions at the beginning of each period (and for any contingency), which we denote by $s(0)$ at the beginning of the analysis horizon and $s^{h_0}(1)$ at the beginning of period 1 (which can be contingent on the state realization in period 0).

Efficient Insurance Market Benchmark. As a natural benchmark, we study a first-best world in which households can purchase actuarially-fair life insurance policies. Since eligibility if the bad state occurs is fully deterministic in age, and hence fully predictable, households make age/eligibility-contingent insurance purchases at the ex-ante stage of the planning problem. In period 0, the household would pay $p(0)$ if the good state occurs in that period; and if the bad state occurs in that period, the household would receive a series of payments $b(0)$ and $b(1)$. A household that arrives at period 1 in the good state, would pay $p(1)$ if the good state occurs in that period, and would receive a payment $\bar{b}(1)$ if the bad state occurs in that period. Actuarially-fair pricing implies that $\mu^g p(0) = \mu^b (b(0) + b(1))$ and $\mu^g p(1) = \mu^b \bar{b}(1)$.

Household Preferences. Let $U^s = u_1^s(c_1) - v_1^s(l_1) + u_2^s(c_2) - v_2^s(l_2)$ represent the household's per-period state-dependent utility, where $u_i^s(c_i)$ is member i's utility from consumption and $v_i^s(l_i)$ represents member i's disutility from labor (including the utility loss from direct work costs and the opportunity costs of lost home production). The model freely allows for potential changes in spouses' utility when the shock occurs. Importantly, all our conclusions are robust to any such dependencies. We employ the normalization $u_1^b(0) = v_1^b(0) = 0$ which allows us to model the bad state of a spousal death by setting $c_1 = l_1 = 0$ in periods/contingencies in which the household is in state b. With these assumptions, the household's preferences in the good state take the form: $U^g = u_1^g(c_1) - v_1^g(l_1) + u_2^g(c_2) - v_2^g(l_2)$, and the household's preferences reduce to the utility from member 2's allocation in the bad state: $U^b = u_2^b(c_2) - v_2^b(l_2)$. We additionally assume that the consumption utility and the labor disutility functions are well-behaved—i.e., that $u_i^{s\prime}(c_i) > 0$, $u_i^{s\prime\prime}(c_i) < 0$, $v_i^{s\prime}(l_i) > 0$, and $v_i^{s\prime\prime}(l_i) > 0$ for $i = 1$ and $s = g$ and for $i = 2$ and $s \in \{g, b\}$. The household's expected flow utility in period t is denoted by $V(t) \equiv E_{h_t}(U^s)$, where the expectation operator is taken over the possible realization paths h_t.

Household Behavior. The household's choices involve the labor supply and consumption allocation decisions, as well as insurance purchases, which are time/age and state contingent. The household maximizes it expected life-time utility at the beginning of period 0, $V(1) + V(2)$, subject to the time-state contingent budget constraints so that total consumption at each contingency is bounded by the household's realized income (net of savings) described above. Recall that in our first empirical exercise we compare the responses of households in the immediate period following a spousal death as a function of benefit

eligibility. Hence, in characterizing the household's optimal choices under the efficient-insurance benchmark, our goal is to compare the behavior of households that are just-eligible for benefits and households that are just-ineligible for benefits in the first period they transitioned to the bad state. To derive these equilibrium results, we make a series of perturbation arguments that must hold at the optimal allocation:

i) Consumption and labor supply: Within each period/state the marginal utility from consumption must equate across spouses (when both are alive) and the wife's marginal utility from consumption must equate to her wage-weighted marginal disutility from labor. That is, $u_1^{g\prime}(c_1) = u_2^{g\prime}(c_2)$ and $u_2^{s\prime}(c_2) = \frac{v_2^{s\prime}(l_2)}{w_2}$. Otherwise, there are trivial possible utility-enhancing perturbations.

ii) Insurance purchase for period 0: First, consider the allocation of insurance payouts in case the bad state occurs in period 0. A marginal decrease in $b(0)$ accompanied by a similar increase in $b(1)$ would generate a loss of $u_2^{b\prime}\left(c_2^b(0)\right)$ and a gain of $u_2^{b\prime}\left(c_2^{b,b}(1)\right)$ which must equate at the optimum: $u_2^{b\prime}\left(c_2^b(0)\right) = u_2^{b\prime}\left(c_2^{b,b}(1)\right)$. Second, consider a marginal increase in $p(0)$ with an increase in $b(0)$ that should amount to μ^g/μ^b with actuarially-fair pricing. The utility loss $\mu^g u_2^{g\prime}\left(c_2^g(0)\right)$ and the utility gain $\left(\mu^g/\mu^b\right)\mu^b u_2^{b\prime}\left(c_2^b(0)\right)$ must equate, so that $u_2^{g\prime}\left(c_2^g(0)\right) = u_2^{b\prime}\left(c_2^b(0)\right)$.

iii) Insurance purchase for period 1: Similar to above, consider a marginal increase in $p(1)$ with an increase of μ^g/μ^b in $\bar{b}(1)$. The utility loss $\mu^g u_2^{g\prime}\left(c_2^{g,g}(1)\right)$ and the utility gain $\left(\mu^g/\mu^b\right)\mu^b u_2^{b\prime}\left(c_2^{g,b}(1)\right)$ must equate, so that $u_2^{g\prime}\left(c_2^{g,g}(1)\right) = u_2^{b\prime}\left(c_2^{g,b}(1)\right)$.

iv) Savings decision: At the beginning of period 0, an additional dollar of savings will incur a loss of $\mu^g u_2^{g\prime}\left(c_2^g(0)\right) + \mu^b u_2^{b\prime}\left(c_2^b(0)\right)$, which equals $u_2^{g\prime}\left(c_2^g(0)\right)$ from ii). The expected gains from its consumption in period 1 would amount to: $\mu^b u_2^{b\prime}\left(c_2^{b,b}(1)\right) + \mu^g\left[\mu^g u_2^{g\prime}\left(c_2^{g,g}(1)\right) + \mu^b u_2^{b\prime}\left(c_2^{g,b}(1)\right)\right]$, which equals $\mu^b u_2^{g\prime}\left(c_2^g(0)\right) + \mu^g u_2^{g\prime}\left(c_2^{g,g}(1)\right)$ from ii) and iii). Equality of marginal gains and losses implies that $u_2^{g\prime}\left(c_2^g(0)\right) = \mu^b u_2^{g\prime}\left(c_2^g(0)\right) + \mu^g u_2^{g\prime}\left(c_2^{g,g}(1)\right)$, so that $u_2^{g\prime}\left(c_2^g(0)\right) = u_2^{g\prime}\left(c_2^{g,g}(1)\right)$. A similar condition can be derived if savings decisions are made at the end of period 0 in the transition from period 0 to 1.

The combination of conditions i)-iv) implies that age-contingent life insurance purchases result in equality of the wife's marginal utility from consumption across states of nature and the eligibility status, for any possible history/contingency. Specific for our purposes, it implies about the immediate period following widowhood that $u_2^{b\prime}\left(c_2^b(0)\right) = u_2^{b\prime}\left(c_2^{g,b}(1)\right)$. Then i) also implies that the same equality holds for the wife's marginal disutility from labor, so that $v_2^{b\prime}\left(l_2^b(0)\right) = v_2^{b\prime}\left(l_2^{g,b}(1)\right)$. This condition implies that the

labor supply of newly-widowed households that are just-eligible and just-ineligible should be similar, which forms the benchmark for our empirical analysis.

A.2. Responses by Already-Widowed Households to Anticipated Benefits

Consider an already-widowed household and let us analyze the case in which the following assumptions hold: (1) households are forward-looking and understand the Social Security benefit schedule and rules; (2) there are no liquidity constraints. We analyze a two-period model to capture a period of benefit ineligibility followed by a period of benefit eligibility. The results extend to multi-period dynamic models, since they rely on classical Euler conditions that hold more generally. Since we focus on spouses who are already in the bad state, we suppress any indexes for the household member or the state of nature.

For any variable x, let x_t represent the value of x in period $t \in \{1,2\}$, and consider the planning problem of a household that transitioned to widowhood prior to the beginning of period 1. For simplicity, we again set the discount rate and the interest rate to zero. We assume that benefit eligibility comes into effect only in period 2, so that $B_1 = 0$ and $B_2 > 0$. This deterministic benefit schedule can be fully anticipated at the beginning of the planning problem. We carry the term B_1 rather than setting it to 0 to demonstrate behavior in the presence of a more general form of benefit timing. Formally, the household solves the problem: max $U_1 + U_2 = u(c_1) - v(l_1) + u(c_2) - v(l_2)$, subject to the within-period budget constraints: $c_1 = A + wl_1 + B_1 - s$ and $c_2 = A + wl_2 + B_2 + s$, where A is a baseline level of wealth. The choice of saving or borrowing, s, is unconstrained beyond guaranteeing that $c_1, c_2 \geq 0$.

At the optimum, widows smooth consumption and leisure, so that $c_1 = c_2$ and $l_1 = l_2$, and the whole planning problem can be rewritten in terms of the present discounted value of life-time unearned income, $I \equiv A + B_1 + B_2$.[3] Hence, the main prediction of this familiar model, which we use as our benchmark, is that of labor supply smoothing: there should be no discontinuity in labor market choices when the anticipated benefits become available. That is, for a given level of the present discounted value (PDV) of benefits, the household's behavior should not depend on their timing. It is straightforward to also explicitly incorporate an earnings test similar to that of the Social Security survivors insurance, whereby benefits increase permanently to account for the months in which benefits are withheld if widows' earned income crosses a given threshold. If households correctly perceive the earnings test, the qualitative results of our model remain the same.

[3] The saving/borrowing decision follows $s = \frac{1}{2}(A + B_1 - B_2)$, and the planning problem reduces to maximizing $U_1 + U_2$ subject to $c_t = wl_t + \frac{1}{2}I$ for $t \in \{1,2\}$.

B. Benchmark Calibration of Full Hand-to-Mouth

To provide this calibration in the context of the model from Appendix A.2, we first employ the simple within-period first-order condition $u'(c_t) = \frac{v'(l_t)}{w}$ in the absence of benefits ($t = 1$) to calibrate parameters; and we then use the same equation at benefit eligibility ($t = 2$) to impute the responses if households were to display complete hand-to-mouth (HtM) behavior (so that current income equals consumption). We make the following parametric assumptions and calibrations: $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, $\gamma = 2$; and $v(l) = a + bl^2$. The first-order condition is then: $c^{-\gamma} = \beta z$, where $\beta \equiv \frac{2b}{w^2}$ and $z \equiv wl$. At the imputed counterfactual in the absence of benefits, income equals \$40,341 and earnings equal \$20,567, which implies that $\beta = \frac{40{,}341^{-2}}{20{,}567}$. Among eligible households, income equals \$43,252. To satisfy the first-order condition under complete hand-to-mouth, their earnings should be: $z = \frac{43{,}252^{-2}}{\beta} = \$17{,}891$. To measure the degree to which households display hand-to-mouth behavior, we divide the gap between the actual earnings treatment effect and the treatment effect under no liquidity constraints (our benchmark of zero) by the treatment effect under complete credit constraints (that is, \$17,891-\$20,567 = -\$2,675). This measure has the properties that it equals zero under the permanent income hypothesis (PIH) and it equals 1 under HtM. In practice, the measure equals $\frac{-1{,}938-0}{-2{,}675}$ =0.72, so that the representative widowed household displays behavior that is 72% away from PIH and 28% away from HtM.

Appendix Figure 1: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Women with Pre-Shock Labor Income below the Earnings Test

A. Social Security Benefit Amounts

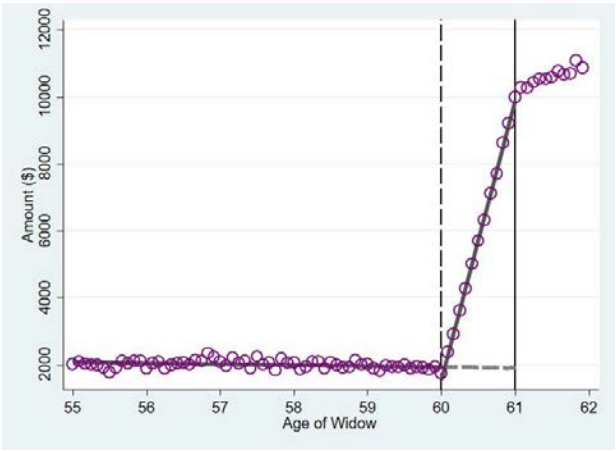

B. Overall Net Household Income

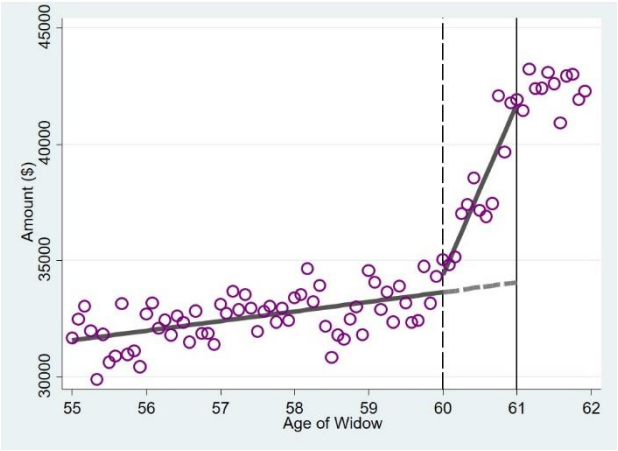

C. Labor Force Participation

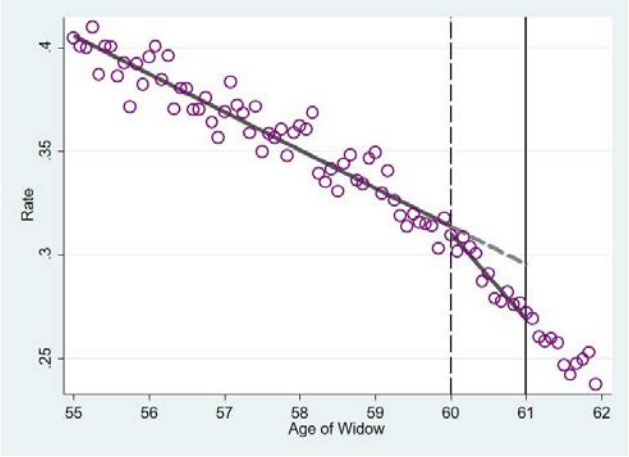

D. Wage Earnings

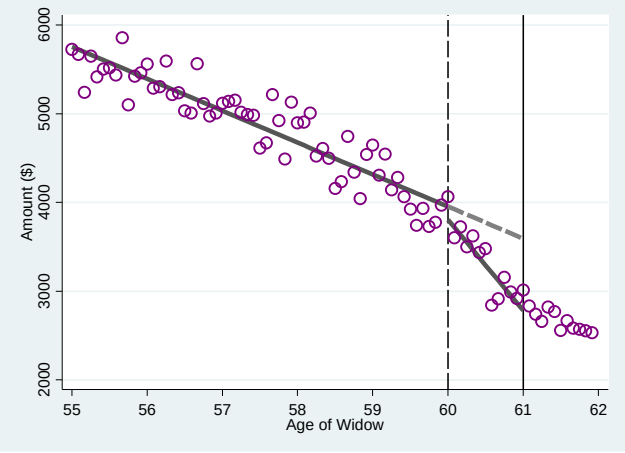

Notes: These figures plot various household outcomes in the years just after a husband's death event ($t = 1,2$) as a function of the surviving spouse's age in months. They include the sample of widows whose pre-shock earnings were below the earnings test thresholds. The purple circles represent means of 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 (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 (marked by the vertical solid black line).

Appendix Figure 2: Responses to Anticipated Benefit Receipt by Already-Widowed Women with Lagged Labor Income below the Earnings Test

A. Social Security Benefit Amounts

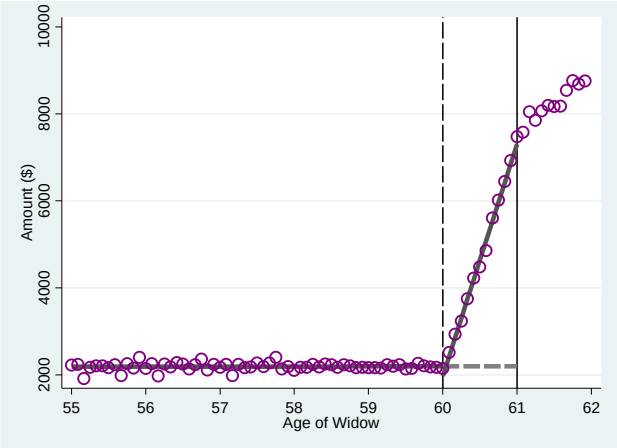

B. Labor Force Participation

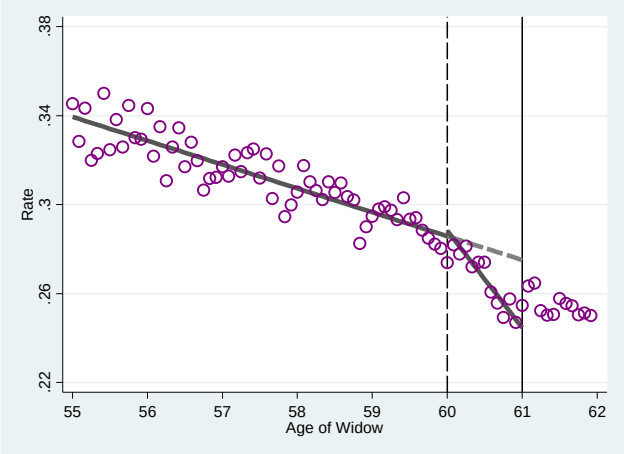

C. Wage Earnings

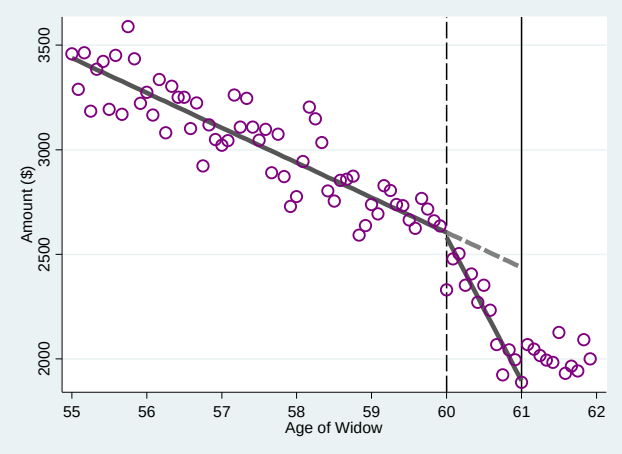

Notes: These figures plot various outcomes of already-widowed households (using observations from periods $t = 6 - 10$ following the spousal death) as a function of the surviving spouse's age in months. They include the sample of widows whose lagged earnings were below the earnings test thresholds. The purple circles represent means of 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 (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 (marked by the vertical solid black line).

Appendix Table 1: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Households in Post-Shock Year $t = 1$

	Social Security Benefits		Overall Net Household Income	Labor Supply	
	Claiming Rate	Benefit Amounts		Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
Full-Exposure Effect	0.54086***	5,959***	4,912***	-0.03514***	-1,784***
	(0.00461)	(53)	(501)	(0.00504)	(448)
Number of Obs.	259,407	259,407	259,407	259,407	259,407
Number of Clusters	259,407	259,407	259,407	259,407	259,407

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1). It provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations from the immediate post-shock year $t = 1$. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 2: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Women with Pre-Shock Labor Income below the Earnings Test

	Social Security Benefits		Overall Net Household Income	Labor Supply	
	Claiming Rate	Benefit Amounts		Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
Full-Exposure Effect	0.60301***	7,258***	7,074***	-0.02424***	-1,065***
	(0.00461)	(60)	(522)	(0.00482)	(214)
Number of Obs.	216,167	216,167	216,167	216,167	216,167
Number of Clusters	126,635	126,635	126,635	126,635	126,635

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1) for the sample of widows whose pre-shock earnings were below the earnings test thresholds. It provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations from the immediate post-shock years ($t = 1,2$). We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 3: Effects of Eligibility for Social Security's Survivors Benefit Receipt upon Spousal Death on Widow's Longer-Run Outcomes

	Overall Net Household Income	Labor Supply	
		Participation	Wage Earnings
	(1)	(2)	(3)
Full-Exposure Effect	-126	-0.00395	-167
	(476)	(0.00528)	(160)
Number of Obs.	309,539	309,539	309,539
Number of Clusters	151,022	151,022	151,022

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits at the occurrence of spousal death on widows' later-life outcomes. Specifically, the estimation includes observations of widows of ages 67-69 (based on the range of our data), and it studies how they may differ by the widow's age at the year of the husband's death using the full-exposure effect based on specification (1). We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 4: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Households by Degree of Liquidity

	Social Security Benefit Amounts	Labor Supply	
		Participation	Wage Earnings
	(1)	(2)	(3)
Panel A: Low Liquidity			
Full-Exposure Effect	5,057***	-0.02714***	-2,219***
	(51)	(0.00454)	(514)
Number of Obs.	278,860	278,860	278,860
Number of Clusters	191,347	191,347	191,347
Panel B: High Liquidity			
Full-Exposure Effect	6,140***	-0.02053***	-916***
	(59)	(0.00518)	(296)
Number of Obs.	225,244	225,244	225,244
Number of Clusters	162,387	162,387	162,387
Difference Low-High	-1,083***	-0.00661	-1,303**
	(77)	(0.00696)	(595)

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1) for high-liquidity and low-liquidity households. We proxy for the degree of liquidity using lagged unearned income and split the observations by the sample median. For each subsample, the table provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations from the immediate post-shock years ($t = 1,2$). We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 5: Effects of Social Security's Survivors Benefit Receipt on Newly-Widowed Households—Robustness

	Social Security Benefits		Overall Net Household Income	Labor Supply	
	Claiming Rate	Benefit Amounts		Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
Panel A: Augmented Design—Control Group					
Full-Exposure Effect	0.53745***	5,759***	5,412***	-0.02866***	-1,363***
	(0.00489)	(56)	(767)	(0.00696)	(465)
Number of Obs.	547,279	547,279	547,279	547,279	547,279
Number of Clusters	317,332	317,332	317,332	317,332	317,332
Panel B: No Non-Death Withdrawals					
Full-Exposure Effect	0.51162***	5,553***	4,642***	-0.02890***	-1,770***
	(0.00343)	(40)	(348)	(0.00356)	(306)
Number of Obs.	484,395	484,395	484,395	484,395	484,395
Number of Clusters	282,336	282,336	282,336	282,336	282,336
Panel C: No Pre-Shock Contributions					
Full-Exposure Effect	0.55553***	5,905***	5,095***	-0.02447***	-1,597***
	(0.00458)	(54)	(432)	(0.00496)	(272)
Number of Obs.	252,192	252,192	252,192	252,192	252,192
Number of Clusters	147,400	147,400	147,400	147,400	147,400

Notes: This table reports various estimations of the full-exposure effects of eligibility for Social Security's survivors benefits as robustness checks. Panel A augments 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' observations and due to the horizon of our data, the control group is based on households that experience a spousal death in years 2005-2007, and the treatment group is based on households that experience a spousal death in years 2002-2004, so that all included observations are from years 2003-2006. We estimate a specification that fully interacts the terms in equation (1) with an indicator for whether a household belongs to the treatment group, denoted by $treat_i$; that is, we estimate: $y_{i,t} = \beta_0 + \beta_1(age_{i,t} - 60) + \beta_2\{age_{i,t} > 60\} + \beta_3\{age_{i,t} > 60\} \times (age_{i,t} - 60) + treat_i \times [\gamma_0 + \gamma_1(age_{i,t} - 60) + \gamma_2\{age_{i,t} > 60\} + \gamma_3(age_{i,t} - 60) \times \{age_{i,t} > 60\}] + \varepsilon_{i,t}$. The full-exposure effect is then assessed by $\gamma_2 + \gamma_3 \times (11/12)$. Panels B and C estimate equation (1) for different subsamples, and they report the full-exposure effect based on the estimate for $\beta_2 + \beta_3 \times (11/12)$ using observations from the immediate post-shock years ($t = 1,2$). Panel B includes all households that do not make any withdrawals indexed with reasons other than "death," and panel C includes all households who did not make contributions to savings accounts in the pre-shock periods. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 6: Responses to Anticipated Benefit Receipt by Already-Widowed Women with Lagged Labor Income below the Earnings Test

	Social Security Benefits		Overall Net Household Income	Labor Supply	
	Claiming	Benefit Amounts		Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
Full-Exposure Effect	0.38572***	4,666***	4,201***	-0.02772***	-497***
	(0.00454)	(58)	(334)	(0.00404)	(59)
Number of Obs.	255,325	255,325	255,325	255,325	255,325
Number of Clusters	117,735	117,735	117,735	117,735	117,735

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits based on specification (1) for widows whose lagged earnings were below the earnings test thresholds. It provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation includes observations of already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.

Appendix Table 7: Responses to Benefit Receipt—Self-Employment

	Newly-Widowed		Already-Widowed	
	Indicator for Self-Employment	Self-Employment Income Amount	Indicator for Self-Employment	Self-Employment Income Amount
	(1)	(2)	(3)	(4)
Full-Exposure Effect	-0.00552**	-45	-0.00622***	-44***
	(0.00219)	(32)	(0.00193)	(13)
Number of Obs.	504,104	504,104	544,223	544,223
Number of Clusters	293,857	293,857	226,701	226,701

Notes: This table reports estimates for the impact of eligibility for Social Security's survivors benefits on self-employment based on Schedule C. Using specification (1), it provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. The estimation in columns 1-2 includes observations from the immediate post-shock years ($t = 1,2$). The estimation in columns 3-4 includes observations of already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. We report robust standard errors clustered at the household level. *** p<0.01, ** p<0.05, * p<0.1.