

NBER WORKING PAPER SERIES

MARKET INEFFICIENCY AND HOUSEHOLD LABOR SUPPLY:
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, Revised June 2019

Previously circulated as "Household Responses to Transfers and Liquidity: Evidence from Social Security's Survivors Benefits." 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, Alex Gelber, Jon Gruber, Henrik Kleven, Amy Finkelstein, Nathan Hendren, Erzo Luttmer, Karthik Muralidharan, Petra Persson, Tommaso Porzio, Kaspar Wüthrich, and seminar participants at UCSD, Middlebury College, MIT, University of Michigan, the OTA Workshop, the 2018 NBER Summer Institute, the 2018 NTA Annual Meeting, NBER Public Economics Meeting (Fall 2018), the 2018 Retirement Research Consortium Annual Meeting, the 2018 All-California Labor Economics Annual Meeting, and the ASU Annual Empirical Microeconomics Conference for helpful comments and discussions. Jonathan Leganza provided excellent research assistance.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Itzik Fadlon, Shanthi P. Ramnath, and Patricia K. Tong. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Market Inefficiency and Household Labor Supply: Evidence from Social Security's Survivors Benefits

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

NBER Working Paper No. 25586

February 2019, Revised June 2019

JEL No. D1,D61,G22,H0,H55,I1,I38,J2

ABSTRACT

We study the effects of the Social Security survivors benefits program on household labor supply and the efficiency implications for insurance and credit markets. We use U.S. population tax records and exploit a sharp age discontinuity in benefit eligibility for identification. We find that eligibility induces considerable reductions in labor supply both among newly-widowed households in the immediate post-shock periods and among already-widowed households whose benefit receipt is entirely predictable. The evidence points to liquidity constraints, rather than myopia, as a leading operative mechanism underlying household responses to anticipated benefits. Our findings identify important inefficiencies in the life insurance market and in the allocation of credit. Our results further highlight the protective insurance role of the social program and the importance of liquidity provided by the government, and they suggest potential gains from expanding and smoothing the program's benefit schedule.

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 a household can face, and it poses a major source of economic risk for American families. Fadlon et al. (2019) show that, among American households, a husband’s death leads to significant declines in equivalence-scale adjusted income and to considerable increases in financial insolvency, which are both immediate and persistent. In the U.S., there are approximately 15 million surviving spouses at any given point in time, with 1.4 million newly-widowed households each year.¹ The social insurance program that aims to protect against the income losses imposed by this shock—namely, Social Security’s survivors benefits—has rapidly grown into one of the largest safety-net programs in the United States. In 2015, the government paid more than \$95 billion to 4.2 million surviving spouses (up from \$64 billion in 2000); where, by comparison, unemployment 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 the program with respect to its benefit generosity, eligibility ages, and benefit timing are currently under consideration.² The significance of social insurance against spousal death in the U.S. is further magnified by the evidence of a considerable inadequacy in Americans’ life insurance holdings, which have also been declining in recent decades.³

Surprisingly, despite the importance of the program to the welfare of vulnerable American households, there is virtually no causal evidence on the economic effects of Social Security’s survivors benefits. Specifically, we lack knowledge about two central policy questions on the impact of the program’s benefit eligibility that directly concern the operation of key markets in the U.S. The first question—which directly pertains to the efficiency of the life insurance market and the protective role of government transfers—is how newly-widowed households respond to benefit eligibility in the immediate periods following the spousal death. The second question—which directly pertains to credit market efficiency and the role of liquidity provided by the transfers—is how becoming age eligible affects the behavior of already-widowed households whose benefit availability is completely predictable.

In this paper, we address these two questions by studying the behavior of American households at different stages of widowhood. Our analysis has three main attributes. First, we use tax records that cover the U.S. population from 1999 through 2014, and we analyze close to a quarter of a million households who experience a spousal death. Second, our research design exploits a sharp discontinuity in the benefit schedule at exactly age 60, when widows become eligible to receive survivors benefits. We leverage the scope and detailed nature of our data to plot raw means of widows’ outcomes against their age in months

¹ Moreover, more than 13.5% of all American women are widowed by age 65 (Fadlon et al. 2019) and widowhood is expected to constitute a considerable share of one’s life-cycle (Compton and Pollak 2018).

² 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).

³ See Auerbach and Kotlikoff (1987, 1991a,b), Bernheim et al. (2003a,b), Hartley et al. (2018).

around the eligibility cutoff. These figures provide compelling visual evidence of household responses to the policy. Third, methodologically, we can utilize labor supply responses as a well-measured, directly-observable input of individuals' utility that is particularly informative in identifying market efficiency.⁴ We show that households who have access to perfect markets should not respond to survivors benefit availability since they would be able to smooth their behavior across eligibility statuses. Even more, we illustrate that the extent of responses through labor supply is directly proportional to the marginal value widowed households would assign to the correction of any existing market failure.

We provide two main sets of results. First, we find large effects of eligibility for survivors benefits among newly-widowed households in the periods that immediately follow the spousal death event. Eligibility leads to substantial increases in our measure of household net income, which accounts for adjustments in alternative sources such as private savings and other government transfers. These increases amount to 11.4% among all widows and to 20.1% among low-earning widows. Notably, we find that benefit eligibility induces meaningful declines in widows' labor supply. In particular, the decreases in wage earnings are on the order of 9.32% among the overall sample and of 27% among low-earning widows, who comprise 43% of all widowed households in our analysis. These findings first highlight the protection the program provides. That is, eligibility generates gains both through considerable increases in income available for consumption and through greater consumption of leisure, as the social insurance program mitigates the need for income compensation via family labor supply.

Importantly, these labor supply responses in the immediate post-shock periods reveal considerable allocative inefficiencies in the life insurance market. With quadratic labor disutility as one example, the meaningful reductions in labor supply map one-for-one to the excess value ineligible newly-widowed households would assign to a dollar of survivors benefits. The evidence implies that these responses are driven by the effect of immediate cash-on-hand rather than a wealth effect. The evidence also suggests that the responses are socially desirable toward the optimal allocation, as we document notable non-distortionary effects (vs. substitution effects) in family labor supply.

Second, we find considerable labor supply declines in response to benefit eligibility among households who had already transitioned to widowhood several years earlier. For these widows, the present discounted value of survivors benefit entitlement is fixed and benefit receipt at the eligibility age cutoff represents only a discontinuous predictable increase in cash-on-hand. Our findings imply clear deviations from the frictionless first-best benchmark, in which labor supply should display no sensitivity to anticipated benefit timing. A simple calibration suggests that the responses represent 70% of the hypothetical response under the full hand-to-mouth benchmark. This points to notable credit allocation inefficiencies among

⁴ For the use of labor supply in normative assessments, see, for example, Shimer and Werning (2007), Chetty (2008), Landais (2015), Hendren (2017), Fadlon and Nielsen (2018), Giupponi (2018), Wettstein (2019).

American households, and it underscores the value from injecting liquidity earlier in widows' life-cycle, which is proportional to the sizable labor supply responses to anticipated benefit availability.

We next devise several tests to provide suggestive evidence of the behavioral channels underlying these responses, focusing on liquidity constraints and myopia. We show a considerable response gradient in household liquidity (as proxied by unearned income): the labor supply declines are attributable to lower-liquidity households who exhibit large effects, whereas the highest-liquidity households are able to fully smooth their labor supply behavior as predicted by the frictionless model. We also show that complete myopia among widowed households is unlikely to explain the results; widows clearly exhibit strategic timing of remarriage for ensuring benefit entitlement, which requires forward-looking planning.⁵ Overall, the results point to liquidity constraints as a leading operative mechanism for the estimated responses to anticipated benefits.

Our paper contributes to several existing literatures. Our primary contribution is to the influential earlier work that aimed to assess the adequacy of households' life insurance holdings in the U.S.—a major economic question that has remained largely open in recent decades despite the fact that life insurance constitutes one of the largest insurance markets in the United States. The past work (e.g., Auerbach and Kotlikoff 1987, 1991a,b; Bernheim et al. 2003a,b) has used small-sample survey data to project the living standard of surviving spouses, by analyzing the household state-contingent income in hypothetical cases of a spousal death. Based on administrative population-level data, we provide an alternative test and new evidence for allocative inefficiency in the life insurance market and for the value of transfers to survivors. Our approach offers two major advantages over existing work. First, by relying on labor supply that is directly assignable to individuals, our analysis does not require estimates for scale economies within the household for which there is no consensus in the literature. This is in contrast to evaluations of state-contingent income, which rely on and are sensitive to these estimates. Second, by relying on variation in eligibility status within the widowhood state alone, our analysis is robust to any form of state dependence in preferences. This is key since potential changes in preferences across states of nature pose one of the most difficult challenges in assessing insurance efficiency and value of benefits for any type of risk and, in particular, in the context of spousal death.⁶

⁵ 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.

⁶ See general discussions in Finkelstein, Luttmer, and Notowidigdo (2009, 2013), Chetty and Finkelstein (2013), Hendren (2017), Fadlon and Nielsen (2018), and Landais and Spinnewijn (2019). Fadlon and Nielsen (2017, 2018) describe the particular important challenges that we overcome in this paper, and Giupponi (2018) addresses in concurrent work similar issues in the valuation of survivors benefits in the Italian context. Fadlon and Nielsen (2017) specifically point out that potential changes in preferences from spousal death could occur as a result of lost preference complementarities across spouses and from the significant declines in widows' health following the event (see, e.g., Stroebe et al. 2007).

We also contribute to the large literature on household behavior in the context of liquidity constraints and credit inefficiencies, which spans various fields including public finance and macroeconomics.⁷ In our setting, the efficient frictionless benchmark for the responses to anticipated benefit availability is (identically) zero. Hence, an advantage of our analysis is that we provide a wide-scope setting with 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 able to offer suggestive analysis for the potential mechanisms that underlie older American households' responses to predictable changes in cash-on-hand.

The broad takeaway from the combination of our results is the important dual role of liquidity via government transfers over the course of the 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. We show both to be qualitatively and quantitatively important. Our results, which provide the first causal estimates for the effects of the Social Security survivors insurance program,⁸ also have implications for its benefit scheme. With 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 highly value insurance through government transfers. Furthermore, the revealed role of anticipated liquidity implies that merely changing the timing of benefits and smoothing the survivors benefits' profile, keeping their current level of present discounted value unchanged, has potential for generating considerable value to households.⁹

The remainder of the paper is organized as follows. Section 2 describes the institutional setting. Section 3 outlines a simple conceptual framework to provide benchmarks for the household responses that we estimate. In Section 4 we describe the data and lay out our empirical framework. Section 5 presents our analysis of the effects of eligibility for Social Security's survivors benefits and discusses the implications of our findings. Section 6 concludes.

⁷ See, for example, Zeldes (1989), Parker (1999), Souleles (1999), Johnson et al. (2006), and Card, Chetty, and Weber (2007).

⁸ To the best of our knowledge, the most related existing work is by Hurd and Wise (1996) who simulated how widows' poverty would mechanically change (i.e., abstracting from behavioral responses) if survivors benefits increased through declines in spousal benefits, and by McGarry and Schoeni (2000) who study factors that can explain the changes in widows' living arrangements and point to Social Security benefits as a potential explanation.

⁹ 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 find that working widowed women exhibit a considerable increase in retirement rates in response to survivors benefits. With the significant growth in female labor force participation at older ages in the U.S. (Goldin and Katz 2018) and the meaningful share of widows among older American women, the evidence suggests that the Social Security survivors benefits program itself could play an increasing role in female retirement behavior.

2. Background and Institutional Details

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

Surviving spouses become universally eligible for Social Security's survivors benefits at exactly age 60.¹⁰ To be eligible, surviving spouses cannot remarry before age 60, otherwise they lose their entitlement altogether. 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 translates 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). The retirement benefits aim to reflect life-time earnings and are based on a worker's Average Income Monthly Earnings (AIME) over the 35 years in which the worker earned the most.

Survivors benefits are then calculated as a percentage of the deceased's potential retirement benefits, and this percentage is determined by the surviving spouse's age at the beginning of benefit claiming. The percentage ranges from 71.5% at age 60 to 100% at the widow's full retirement age (65 or older for all cohorts), which represents actuarial adjustments to account for the different length of benefit collection when benefits are claimed at different ages. 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 the widow's potential entitlement.¹¹

By design, the present discounted value (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 advantages. First, it implies that the PDV of survivors benefit entitlement is fixed from the point of the husband's death onward. Second, it implies there are no actual differential substitution effects at eligibility, so the potential impacts of the program on widows' labor supply should operate through a non-distortionary 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 (SSA 2018b). 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

¹⁰ 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.

¹¹ We have learned of these practices through former SSA field officers.

utilize this feature of the program later in the analysis by studying subsamples of households that are infra-marginal to the earnings test. We further analyze these subsamples to gain additional insights on the program’s effects on low-earnings households.

It may be useful to discuss other potentially important ages in the vicinity of our age 60 threshold. The first is age 62 which is the early eligibility age for the standard Social Security retirement benefits. Note that 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 that age. To account for this threshold, we restrict the analysis to observations of widows younger than 62. 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 retirement savings accounts are no longer penalized. However, this is effectively not a relevant margin for our context of widows and their overall financial portfolio, since the death event itself already allows for non-penalized distributions from the deceased husband’s accounts. Indeed, only 4% and 1% of all analyzed newly-widowed households and already-widowed households, respectively, ever make any withdrawals that are indexed with reasons other than “death.”¹²

Finally, due to the nature of our setting, we focus on female surviving spouses. Women comprise the vast majority of all widowed households throughout the age distribution (around 80%; Fadlon et al. 2019) and close to 100% of all the program’s beneficiaries (e.g., 98% in 2017; SSA 2018c).

3. Conceptual Framework and Estimation Benchmarks

We utilize our empirical setting to conduct two related empirical analyses that differ by the time horizon following the spousal death. Our first analysis investigates the effects of eligibility for benefit receipt on newly-widowed households in the periods that immediately follow the death event, and it focuses on the insurance market. Our second analysis 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 and the credit market. In this section we describe first-best benchmarks for each of the empirical analyses to provide useful anchors for the effects that we estimate. These benchmarks form natural tests for market allocation inefficiencies and will be used to interpret our findings and to draw their normative implications. 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.

¹² Naturally, excluding these households does not change the results (see panel A of Appendix Table 6). To alleviate remaining concerns, we also analyze households who are less likely to be affected by this change in withdrawal incentives (from an ex-ante standpoint), as they did not make contributions to savings accounts in previous periods (see panel B of Appendix Table 6).

(i) *Responses by Newly-Widowed Households to Eligibility for Benefit Receipt.* Consider the decisions of a two-person household, which consists of member 1, the husband, and member 2, the wife, 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. Hence, in our model analysis that follows, 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. We later show that the degree of deviation further has direct implications for the value of transferring insurance benefits to ineligible households.

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, which is directly assignable to individuals (so it does not require scaling), the test does not rely on estimates for economies of scale within the household. 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 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 liquidity upon the 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 first empirical analysis compares the behaviors of women at the initial stage 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. If liquidity matters, this sharp increase in 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 of benefits at the eligibility cutoff for the following reasons. First, among widows older than 60 but younger than the full retirement age, the program is designed to provide an entitlement for the same benefit PDV for a given history of a husband’s earnings, whereby claiming benefits at different ages involves actuarial adjustments as we mentioned above. Second, 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 the widows' monthly age—but such potential increases are smooth as per the benefit calculation formula.¹³

This potential continuous underlying evolution of life-time income (through benefits or husband's earnings) in the new widow's monthly age—which does not confound the identification of responses but can complicate their interpretation—is the key motivation for our second analysis in which life-time income is pre-determined and constant throughout.

(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, and that this benefit schedule is deterministic, and hence can be fully anticipated, at the beginning of the planning period.

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. The model analysis is described in detail in Appendix A.2. 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. 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

¹³ Specifically, potential increases in benefit PDV over a newly-widowed wife's monthly age from an additional month of a husband's earnings are smoothed and muted by the averaging of the husband's Average Income Monthly Earnings (AIME) over 35 years. 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 per-period (of a month) 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.

advantage of our empirical setting, in which the present discounted value of life-time unearned income through survivors benefits is unchanged after the husband's death and does not depend on the widow's behavior, whereas there is a sharp discontinuity in cash-on-hand.¹⁴ 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 analysis we study the behavior of households that are several years into widowhood. This provides them with the necessary time to respond to the event (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 event 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 of insensitivity of widows' labor supply 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 suggestive ways to distinguish between the candidates for the underlying channels: lack of planning, liquidity constraints, and benefit-schedule misperceptions. We also illustrate later that the degree of deviation maps into the value of injecting liquidity earlier by smoothing the benefit profile.

4. Data and Empirical Framework

4.1. Data Sources and Variable Definitions

Data Sources. We use administrative tax records on American households for the years 1999 through 2014. The data include both information returns filed by third parties (e.g., Form W-2, Form SSA-1099, and Form 1099-R) and income tax returns (e.g., Form 1040). 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.

From the information returns, we extract wage earnings (using Form W-2), Social Security benefits paid from the retirement and the disability trust funds (which are reported separately on Form SSA-1099), unemployment benefits (using Form 1099-G), and distributions from pensions, annuities, retirement plans, individual retirement accounts (IRAs), and insurance contracts (as reported on Form 1099-R). From the

¹⁴ This advantageous feature stands in contrast to traditionally-studied age-contingent benefit schemes, such as old-age pensions, in which own work can directly affect the PDV of benefits.

income tax returns, we extract Adjusted Gross Income (AGI). Among other sources of income, AGI includes earnings, capital income, retirement income, and taxable Social Security benefits.

Outcomes and Variable Definitions. Our analysis focuses on widows' labor supply behavior. Based on data from Form W-2, we study as our primary outcomes of interest wage earnings and labor force participation.¹⁵ 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.¹⁶ 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.

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

4.2. Empirical Framework

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

¹⁵ 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. Nevertheless, we report the analysis of self-employment in Appendix Table 5.

¹⁶ This definition follows the literature on retirement behavior in response to old-age government transfers (see, e.g., Coile and Gruber 2007).

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 a widow's age in the post-shock years or indirectly attributed to factors correlated with it, 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. The population-level data allow us to lead our analysis with a graphical representation of the results, which we further use to guide our estimation strategy. We take advantage of survivors' exact dates of birth, and plot raw means of each outcome variable of interest against the widow's monthly age. To focus on the eligibility cutoff of age 60, we plot outcomes of widows who at the time of observation were between ages 55 and 62 (the early eligibility age for standard retirement benefits).

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. 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:

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

In this regression, $y_{i,t}$ denotes an outcome for widow i at time t , $\text{age}_{i,t}$ represents the widow's age in months, and $\{\text{age}_{i,t} > 60\}$ 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.¹⁷

Implementation. We implement our two empirical analyses by studying 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 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 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.

5. The Effects of Social Security’s Survivors Benefits Eligibility

We now turn to estimating the effects of eligibility for Social Security’s survivors benefits on widows’ labor supply and household income. We first analyze the impact on newly-widowed households in the immediate years following a spousal death, and we then study the responses by already-widowed households to anticipated survivors benefits.

5.1. Responses in Immediate Post-Shock Periods to Eligibility for Benefit Receipt

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

¹⁷ Later, we augment the design with a control group of future widows as a robustness check.

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).¹⁸ 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 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).

Household Income. As an initial evaluation of 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 rapid 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%.¹⁹ Scaling the effect of eligibility on net income by using the claiming rate, the effect of benefit receipt on the sample of compliers amounts to \$9,355 ($=\$4,804/0.51351$). Appendix C further constructs the counterfactual level for this subsample, which we estimate to be \$31,307. Hence, the treatment effect of benefit receipt on net household income among compliers represents an increase of 29.9%.

Finally, we investigate the effect of benefit eligibility on our sample of low-earning households—that is, widows with pre-shock earnings lower than the earnings test thresholds. These widows’ labor supply responses, which we analyze below, guarantee the isolation of non-distortionary effects. But, analyzing the effect on their income is also valuable in that low-earnings spouses are likely more exposed to financial risk since they generate little income on their own, as suggested by Fadlon and Nielsen (2017). Among these low-earnings households, who represent a large share of our sample (43%), the claiming rate is 60 pp

¹⁸ 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.

¹⁹ 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.

and they receive \$7,258 in annual benefit amounts. The increase in their net income totals to \$7,074, which represents an increase of 20.1% (on a counterfactual of \$34,043). See Appendix Table 1 and Appendix Figure 1.

Labor Supply. Next, we turn to our core analysis and investigate how benefit eligibility affects the labor supply of widows, which constitutes an important dimension of household gains from social insurance through the consumption of leisure and has direct implications for insurance efficiency. 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.²⁰ To evaluate the cumulative labor supply effect, which is our primary 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).

Again, it is useful to convert these responses to the effect of benefit receipt by focusing on the group of compliers. The effect on their overall labor supply, as captured by responses in wage earnings, translates to a decline of \$3,410 ($=\$1,751/0.51351$). Given that we estimate their average counterfactual level of earnings to be \$10,050 (see Appendix C), these responses represent a decrease of 33.9% in labor supply among compliers as a result of being eligible for survivors benefit receipt.

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 market imperfections (see, e.g., a discussion in Chetty 2008).

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

²⁰ The estimate for the full exposure effect on retirement is 0.01829 (with s.e. 0.00188) on a counterfactual of 0.05704.

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 our subsample of households for whom only a non-distortionary effect is likely operative in their responses. Specifically, 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 1 (and Appendix Figure 1). Similar to the analysis of the full sample, we find meaningful declines in overall labor supply among the current subsample of households. Widows with labor income below the earnings test thresholds exhibit a decline of 2.42 percentage points in labor force participation on a counterfactual baseline of approximately 30.²¹ Their decline in wage earnings amounts to \$1,065 on a baseline of \$3,978. 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.

Robustness. Lastly, to address 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 similar (see Appendix Table 1).

5.1.1. Implications

Recall that, by design, this first analysis identifies the effects of eligibility for benefits 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 event's realization. Therefore, these effects capture and underscore the protective insurance role of survivors benefits against the immediate adverse financial consequences of a spousal death. In particular, the Social Security survivors benefits program generates gains to newly-affected households both through

²¹ It may be useful to compare this response to that of the overall sample as a 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 period. We note that this exercise is only suggestive due to potential heterogeneity in labor supply responses along the earnings distribution.

significant increases in household income flow and through meaningful increases in the consumption of leisure due to a mitigated need to self-insure.

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, $u_2^{b'} = v_2^{b'}/w_2$ (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, $\frac{v_2^{b'}(l_2^b(0)) - v_2^{b'}(l_2^b(1))}{v_2^{b'}(l_2^b(0))}$.²² This expression can be approximated by $\varphi \left| \frac{l_2^b(1) - l_2^b(0)}{l_2^b(0)} \right|$, where $\varphi \equiv \frac{v_2^{b''}(l_2^b(0))}{v_2^{b'}(l_2^b(0))} l_2^b(0)$ is the curvature of labor disutility.

As this gain is proportional to our estimated causal effect of benefit eligibility on labor supply, $\left| \frac{l_2^b(1) - l_2^b(0)}{l_2^b(0)} \right|$, our results point to potentially meaningful valuation of benefits by ineligible widows.²³ For example, calibrating the utility parameter φ to equal 1 as is the case under quadratic labor disutility (of the form $a + bl^2$, $a, b > 0$), 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).²⁴ Notably, among low-earning households, the overall relative response in labor supply as captured by wage earnings is significantly larger and amounts to 27% (\$1,065 on a counterfactual of \$3,978). This points to even greater valuation of insurance benefits among low-earnings widows, and is consistent with the notion that spouses who generate little income on their own are more exposed to financial risk and are effectively less well insured against spousal death (Fadlon and Nielsen 2017). Lastly, among compliers, for whom the difference in benefits received between ineligible and eligible households is largest by construction, the excess valuation of a dollar of benefits by ineligible widows would amount to 33.9%.

²² 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.

²³ The valuation of benefits can be represented relative to $v_2^{b'}(l_2^b(1))$ instead, in which case the gain would be proportional to the term $\left| (l_2^b(1) - l_2^b(0))/l_2^b(1) \right|$ which is larger.

²⁴ Fadlon and Nielsen (2018) show how, alternative to calibration, φ can 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.

Recall from the discussion of the program’s benefit structure in Section 2 that, in the presence of life insurance inefficiencies evidenced by labor supply reductions, 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.²⁵

Assessing the Degree of Income Flow Coverage relative to Pre-Shock Levels. To additionally understand the scope of the program, we complement our main analysis by gauging the extent to which households eligible for Social Security’s survivors benefits are protected against the financial burden imposed by spousal death. For this assessment, we need to evaluate the average 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 event for control households in the year at which the treatment group experience their actual event.²⁶ Full details on this 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 assessment 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’ overall 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

²⁵ 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 2 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 event 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 3.

²⁶ 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 the years 2002-2003 and our control group is composed of women whose husband died in the years 2006-2007.

commonly used adult equivalence scales to account for the household's compositional change. 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.²⁷ 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 evaluate 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 income insurance coverage. Panel A of Table 2 indicates that, in response to a spousal death, eligible widows exhibit no changes in labor supply following the event. 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.²⁸ 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, consistent with a need to self-insure against the income loss imposed by the mortality event.

5.2. 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 event (e.g., in self-insurance through earnings or assets), to anticipate the benefit receipt, and to make necessary financial arrangements (e.g., borrowing). This analysis hence inherently focuses on the impact of predictable changes in cash-on-hand and benefit timing for given life-time wealth, and it identifies the particular role of liquidity provided by government transfers to our sample of vulnerable older families.

Results. Figure 2 (panels A-B) and Table 3 (columns 1-2) first verify the existence of a first stage, indicating that the take-up of survivors benefits amounts to 34 percentage points which translates to an

²⁷ 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.

²⁸ It is worth highlighting again that our primary analysis of the effects of the Social Security survivors benefits program does not suffer from the disadvantages of the complementary assessment that we provide here, which re-emphasizes the key advantages of our design and setting.

increase of \$3,655 in annual transfers from the government. Then, studying labor supply outcomes, we find significant deviations from the frictionless benchmark. There is a break around benefit availability in the pattern of widows' retirement rate with a local spike at the eligibility region (see panel C of Figure 2).²⁹ 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 \$1,938 in annual labor earnings (see panels D-E of Figure 2 and columns 4-5 of Table 3). As for the effect of benefit receipt, these responses imply a decrease of \$5,768 ($=\$1,938/0.33598$) in annual earnings among compliers (on a counterfactual of \$13,203; see Appendix C).

The evidence is inconsistent with the conjecture that these responses may be explained away by misperceptions of the earnings test. We repeat the analysis for low-earning 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 4 (and Appendix Figure 2).

5.2.1. 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. 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).³⁰ This significant deviation from the frictionless benchmark of labor supply smoothing from Section 2 has two sets of implications.

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. That is, when holding the present discounted value of benefits unchanged, 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

²⁹ This effect averages to 0.01997 (with s.e. 0.00165) on a counterfactual baseline of 0.04411.

³⁰ This is in line with findings from Card et al. (2007) for job searchers in Austria.

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.³¹

That is, the (gross) gain from incrementally smoothing the distribution of benefits across periods is captured by the gain from incrementally smoothing labor supply across periods. The latter is proportional to the meaningful labor supply responses to benefit availability that we find. For the overall sample, the responses are on the order of 9.4% (\$1,938 on a counterfactual of \$20,566). For low-earning widows (with earnings that fall below the earnings test), who comprise a large share of 52% of all households in our current sample, we have shown that the relative responses amount to 20%. Hence, the evidence points to even greater gains from benefit retiming among low-earning households. Finally, among compliers, for whom the change in liquidity flows is largest by construction, the gains are proportional to a response of 43.7% (= \$5,768/\$13,203). It is important to note that these 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 Mechanisms. Second, our analysis has implications for the mechanisms that underlie widows' labor supply responses to anticipated government transfers, which can be also informative more generally for the channels that govern vulnerable older Americans' retirement decisions.³² 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. To further investigate the source of this deviation, we offer suggestive 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 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.

³¹ This is similar to Fadlon and Nielsen's (2018) analysis of transferring resources across states of nature but this time applied to transferring resources across periods.

³² Recall from Section 2 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.

Panel A of 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 2019), 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 (of any source). We split households by the sample median, and we analyze labor supply outcomes for each subsample.³³

Panels A and B of Table 4 summarize the labor supply responses 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). The gradient of labor supply responses with respect to household liquidity becomes even clearer when we split households by quartiles, as evident by their differential degree of response in both participation and earnings in panel B of Figure 3. Even more, we 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 of Table 4 and panel B of Figure 3).³⁴ 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. Thus, the findings suggest that liquidity constraints can play a meaningful role, and they further show that household liquidity can provide an explanation for the overall responses to anticipated cash-on-hand and their patterns.

Lastly, we additionally 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 second is the sample of households who have Individual Retirement Accounts (IRAs), but whose (lagged)

³³ Of course, this heterogeneity analysis should not be assigned a causal interpretation, as households with differential levels of liquidity could differ in many other ways that relate to their responses to changes in cash-on-hand.

³⁴ This is not the case in the context of our first analysis in which households are also subject to the immediate adverse financial effects of the death event, whereas the current analysis isolates the effect of anticipated liquidity. Notice that these highest-liquidity households do have "room to respond" in labor supply reductions, as their baseline participation rate is 45 pp.

account balances are low (i.e., below the median).³⁵ 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 benefits 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 government benefits.

6. Conclusion

Using tax records for the U.S. population and exploiting an 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 inefficiencies in the large U.S. life insurance market, as well as in the allocation of credit among American households. The evidence points to significant protection the program provides against the immediate adverse economic impacts of spousal mortality, which comes in the form of considerable increases in net household income and in the consumption of leisure by widowed households. What is more, the evidence underscores the importance of liquidity provided to widowed households by the federal government, both in terms of its insurance role of smoothing consumption across states of nature and in terms of its intertemporal role of 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 analysis further suggests potentially important gains from providing coverage to ineligible families and from a smoother benefit profile.

³⁵ 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:

- Auerbach, Alan J., and Laurence J. Kotlikoff (1987). "Life Insurance of the Elderly: Adequacy and Determinants." In *Work, Health and Income Among the Elderly* (pp. 229–267), G. Burtless, ed., Brookings Institution Press.
- Auerbach, Alan J., and Laurence J. Kotlikoff (1991a). "The Adequacy of Life Insurance Purchases." *Journal of Financial Intermediation*, 1(3): pp. 215-241.
- Auerbach, Alan J., and Laurence J. Kotlikoff (1991b). "Life Insurance Inadequacy-Evidence from a Sample of Older Widows." Working Paper No. w3765, National Bureau of Economic Research.
- Bernheim, B. Douglas, Lorenzo Forni, Jagadeesh Gokhale, and Laurence J. Kotlikoff (2003). "The Mismatch between Life Insurance Holdings and Financial Vulnerabilities: Evidence from the Health and Retirement Study." *American Economic Review*, 93(1): pp. 354-365.
- Bernheim, B. Douglas, Katherine G. Carman, Jagadeesh Gokhale, and Laurence J. Kotlikoff (2003). "Are Life Insurance Holdings Related to Financial Vulnerabilities?" *Economic Inquiry*, 41(4): pp. 531-554.
- Brown, Jeffrey, Arie Kapteyn, Olivia Mitchell, and Teryn Mattox (2013). "Framing the Social Security Earnings Test." Wharton Pension Research Council Working Paper 2013-06.
- Card, David, Raj Chetty, and Andrea Weber (2007). "Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market." *Quarterly Journal of Economics*, 122(4): pp. 1511-1560.
- Chetty, Raj (2008). Moral Hazard versus Liquidity and Optimal Unemployment Insurance. *Journal of Political Economy*, 116(2): pp. 173-234.
- Chetty, Raj, and Amy Finkelstein (2013). "Social Insurance: Connecting Theory to Data." In *Handbook of Public Economics*, Vol. 5: pp. 111-193. Elsevier.
- Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez (2014). Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States. *Quarterly Journal of Economics*, 129(4): pp. 1553-1623.
- Coile, Courtney, and Jonathan Gruber (2007). "Future Social Security Entitlements and the Retirement Decision." *Review of Economics and Statistics*, 89(2): pp. 234-246.
- Compton, Janice, and Robert A. Pollak (2018). "The Life Expectancy of Older Couples and Surviving Spouses."
- 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 Economic Consequences in the U.S.: Evidence from Population Tax Records."
- Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo (2009). "Approaches to Estimating the Health State Dependence of the Utility Function." *American Economic Review*, 99 (2): pp. 116-21.
- Finkelstein, Amy, Erzo F. P. Luttmer, and Matthew J. Notowidigdo (2013). "What Good is Wealth Without Health? The Effect of Health on the Marginal Utility of Consumption." *Journal of the European Economic Association*, 11(s1): pp. 221-258.
- Giupponi, Giulia (2018). "When Income Effects are Large: Labor Supply Responses and the Value of Welfare Transfers."

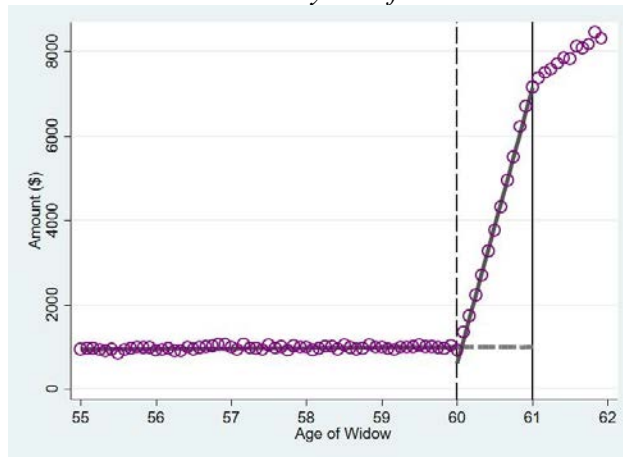
- Goldin, Claudia, and Lawrence F. Katz, eds. (2018). "Women Working Longer: Increased Employment at Older Ages." University of Chicago Press.
- Hartley, Daniel, Anna Paulson, and Katerina Powers (2018). "What Explains the Decline in Life Insurance Ownership?" *Economic Perspectives*, 8: pp. 1-20.
- Hendren, Nathaniel (2017). "Knowledge of Future Job Loss and Implications for Unemployment Insurance." *American Economic Review*, 107(7): pp. 1778-1823.
- Hurd, Michael D., and David A. Wise (1996). "Changing Social Security Survivorship Benefits and the Poverty of Widows." In *The Economic Effects of Aging in the United States and Japan* (pp. 319 - 332), Michael D. Hurd and Naohiro Yashiro, eds., University of Chicago Press.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles (2006). "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review*, 96(5): pp. 1589-1610.
- Kaplan, Greg, and Giovanni L. Violante (2014). "A Model of The Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): pp. 1199-1239.
- Landais, Camille, and Johannes Spinnewijn (2019). "The Value of Unemployment Insurance."
- Liebman, Jeffrey B., and Erzo F.P. Luttmer (2012). "The Perception of Social Security Incentives for Labor Supply and Retirement: The Median Voter Knows More Than You'd Think." *Tax Policy and the Economy*, 26(1): pp.1-42.
- Liebman, Jeffrey B., and Erzo F.P. Luttmer (2015). "Would People Behave Differently if They Better Understood Social Security? Evidence from a Field Experiment." *American Economic Journal: Economic Policy*, 7(1): pp. 275-99.
- MaCurdy, Thomas E. (1981). "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy*, 89(6): pp. 1059-85.
- McGarry, Kathleen, and Robert F. Schoeni (2000). "Social Security, Economic Growth, and the Rise in Elderly Widows' Independence in the Twentieth Century." *Demography*, 37(2): pp.221-236.
- Parker, Jonathan A. (1999). "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review*, 89(4): pp. 959-973.
- Persson, Petra (2019). "Social Insurance and the Marriage Market." *Journal of Political Economy* (Forthcoming).
- Shimer, Robert, and Ivan Werning (2007). "Reservation Wages and Unemployment Insurance." *Quarterly Journal of Economics*, 122(3): pp. 1145-1185.
- Social Security Administration (2018a). Social Security Beneficiary Statistics.
- Social Security Administration (2018b). Exempt Amounts under the Earnings Test.
- Social Security Administration (2018c). Benefits Paid by Type of Beneficiary.
- Souleles, Nicholas S. (1999). "The Response of Household Consumption to Income Tax Refunds." *American Economic Review*, 89(4): pp. 947-958.
- Stroebe, Margaret, Henk Schut, and Wolfgang Stroebe (2007). "Health Outcomes of Bereavement." *The Lancet*, 370(9603): pp. 1960-1973.
- Wettstein, Gal (2019). "Retirement Lock and Prescription Drug Insurance: Evidence from Medicare Part D." *American Economic Journal: Economic Policy* (Forthcoming).
- White House (2016). Fiscal Year 2017: Historical Tables – Budget of the U.S. Government.
- Zeldes, Stephen P. (1989). "Consumption and Liquidity Constraints: An Empirical Investigation." *Journal of Political Economy*, 97(2): pp. 305-346.

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

A. Social Security Benefit Claiming



B. Social Security Benefit Amounts



C. Overall Net Household Income


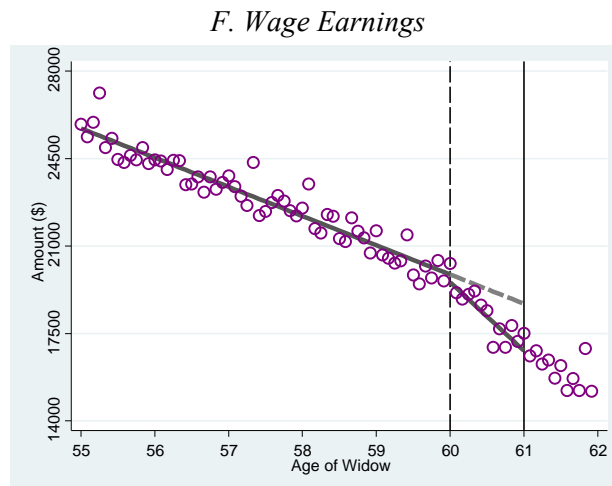
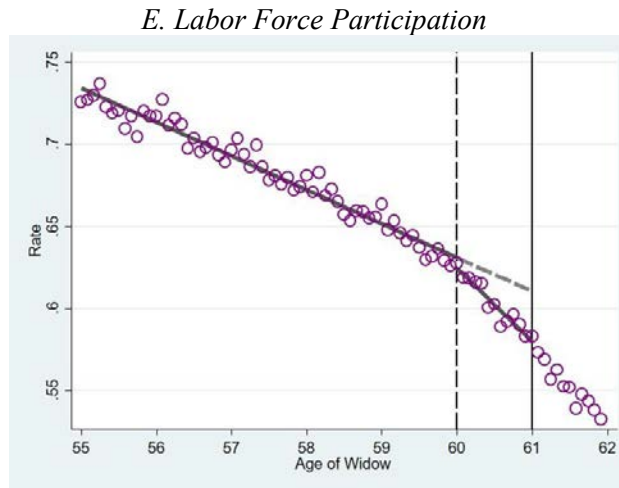




Figure 1: Effects of Eligibility for Social Security’s Survivors Benefit Receipt on Newly-Widowed Households (continued)



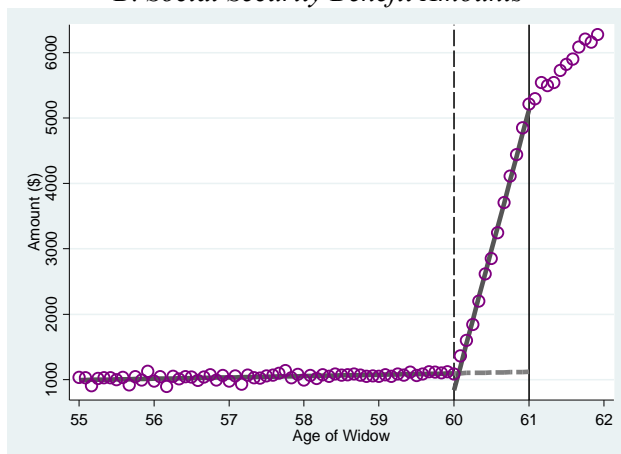
Notes: These figures plot various household outcomes in the years just after a husband’s death event ($t = 1,2$) as a function of the surviving spouse’s age in months. The purple circles represent means of raw data for each “monthly age” bin. The solid gray lines plot the piecewise linear fit using equation (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: Responses to Anticipated Benefits by Already-Widowed Households

A. Social Security Benefit Claiming



B. Social Security Benefit Amounts



C. Overall Net Household Income






Figure 2: Responses to Anticipated Benefits by Already-Widowed Households (*continued*)

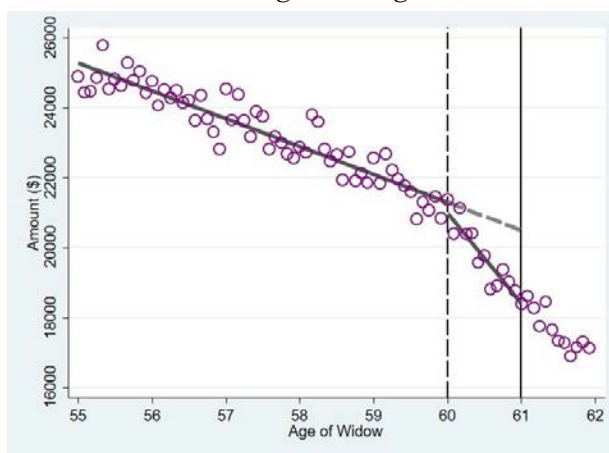
C. Retirement Rate



D. Labor Force Participation




E. Wage Earnings



Notes: These figures plot various household 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 benefits. 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: Mechanisms of Responses to Anticipated Benefits

A. Timing of Remarriage




B. Labor Supply Responses by Household Liquidity

Labor Force Participation



Wage Earnings



Notes: These figures offer different tests to investigate the mechanisms that underlie widows' labor supply responses to anticipated survivors benefits. The analysis sample includes already-widowed households using observations from periods $t = 6 - 10$ following the spousal death. Panel A 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 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). Panel B investigates whether there is evidence that liquidity constraints may play a role. We study whether household labor supply responses vary by the degree of liquidity as proxied by lagged unearned income. We split households by quartiles, and we analyze labor supply outcomes for each subsample. The figures plot full-exposure effects (using equation (1)) on both participation and wage earnings by liquidity quartiles, along with the corresponding 95-percent confidence intervals.

Table 1: Effects of Eligibility for 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*** (0.00336)	5,605*** (39)	4,804*** (343)	-0.02866*** (0.00349)	-1,751*** (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)
<i>Panel A: Widows of Ages 60-61</i>		
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
<i>Panel B: Widows of Ages 58-59</i>		
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 spousal 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. See footnote 26 for more details. 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 Benefits 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*** (0.00308)	3,655*** (36)	2,911*** (250)	-0.03118*** (0.00302)	-1,938*** (197)	0.00893*** (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 Benefits by Degree of Household Liquidity

	Social Security	Labor Supply	
	Benefit Amounts (1)	Participation (2)	Wage Earnings (3)
<i>Panel A: Low Liquidity</i>			
Full-Exposure Effect	3,835*** (46)	-0.04004*** (0.00394)	-2,884*** (339)
Number of Obs.	270,990	270,990	270,990
Number of Clusters	129,023	129,023	129,023
<i>Panel B: High Liquidity</i>			
Full-Exposure Effect	3,442*** (53)	-0.00991** (0.004319)	-598*** (226)
Number of Obs.	273,233	273,233	273,233
Number of Clusters	131,382	131,382	131,382
Difference Low-High	393*** (71)	-0.03013*** (0.00595)	-2,286*** (418)
<i>Panel C: Highest Liquidity</i>			
Full-Exposure Effect	3,390*** (75)	-0.00418 (0.00614)	269 (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 Benefits—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*** (0.02623)	5,714*** (329)	-0.05111* (0.02874)	-5,292*** (1,470)	-0.04499*** (0.00879)	-2,893*** (527)	-0.02446*** (0.00500)	-2,386*** (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 appendix we describe first-best benchmarks for each of our two empirical analyses to provide anchors for the effects that we estimate, 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 Eligibility for Benefit Receipt

Setup. Consider the decisions of a two-person household, which consists of member 1, the husband, and member 2, the wife, 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 prior to 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 tracing the household’s history in terms of state of nature realizations to capture state-contingent plans. Hence, for household choice variables 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, $c_2^{g,b}(1)$, 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 $u_i^{s'}(c_i) > 0$, $u_i^{s''}(c_i) < 0$, $v_i^{s'}(l_i) > 0$, and $v_i^{s''}(l_i) < 0$ 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 its 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 analysis 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, $u_1^{g'}(c_1) = u_2^{g'}(c_2)$ and $u_2^{s'}(c_2) = \frac{v_2^{s'}(l_2)}{w_2}$.

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 $u_2^{b'}(c_2^b(0))$ and a gain of $u_2^{b'}(c_2^{b,b}(1))$ which must equate at the optimum: $u_2^{b'}(c_2^b(0)) = u_2^{b'}(c_2^{b,b}(1))$. 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 $\mu^g u_2^{g'}(c_2^g(0))$ and the utility gain $(\mu^g/\mu^b)\mu^b u_2^{b'}(c_2^b(0))$ must equate, so that $u_2^{g'}(c_2^g(0)) = u_2^{b'}(c_2^b(0))$.

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 $\mu^g u_2^{g'}(c_2^{g,g}(1))$ and the utility gain $(\mu^g/\mu^b)\mu^b u_2^{b'}(c_2^{g,b}(1))$ must equate, so that $u_2^{g'}(c_2^{g,g}(1)) = u_2^{b'}(c_2^{g,b}(1))$.

iv) Savings decision: At the beginning of period 0, an additional dollar of savings will incur a loss of $\mu^g u_2^{g'}(c_2^g(0)) + \mu^b u_2^{b'}(c_2^b(0))$, which equals $u_2^{g'}(c_2^g(0))$ from ii). The expected gains from its consumption in period 1 would amount to: $\mu^b u_2^{b'}(c_2^{b,b}(1)) + \mu^g [\mu^g u_2^{g'}(c_2^{g,g}(1)) + \mu^b u_2^{b'}(c_2^{g,b}(1))]$, which equals $\mu^b u_2^{g'}(c_2^g(0)) + \mu^g u_2^{g'}(c_2^{g,g}(1))$ from ii) and iii). Equality of marginal gains and losses implies that $u_2^{g'}(c_2^g(0)) = \mu^b u_2^{g'}(c_2^g(0)) + \mu^g u_2^{g'}(c_2^{g,g}(1))$, so that $u_2^{g'}(c_2^g(0)) = u_2^{g'}(c_2^{g,g}(1))$. 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 $u_2^{b'}(c_2^b(0)) = u_2^{b'}(c_2^{g,b}(1))$. Then i) also implies that the same equality holds for the wife's marginal disutility from labor, so that $v_2^{b'}(l_2^b(0)) = v_2^{b'}(l_2^{g,b}(1))$. It follows from this condition

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

C. Constructing Counterfactual Outcomes for Compliers

In this appendix we describe how we calculate counterfactuals for the sample of compliers. It follows the classical Local Average Treatment Effect (LATE) framework and assumptions as applied to our setting.

Let Y_d be the potential outcome as a function of potential benefit take-up, $d \in \{0,1\}$, and let d_z be the potential take-up for each possible value of the eligibility instrument, $z \in \{0,1\}$. Denote the share of never-takers (for whom $d_z = 0$) by π_n , the share of always-takers (for whom $d_z = 1$) by π_a , and the share of compliers (for whom $d_0 = 0$ and $d_1 = 1$) by π_c . Additionally, we let $D \in \{0,1\}$ and $Z \in \{0,1\}$ denote the actual benefit take-up and benefit eligibility status, respectively.

We recover counterfactual outcomes for compliers, $E[Y_0|complier]$, using the following relationships:

$$E[D|Z = 0] = \pi_a; E[D|Z = 1] = \pi_a + \pi_c; \pi_n = 1 - E[D|Z = 1]$$

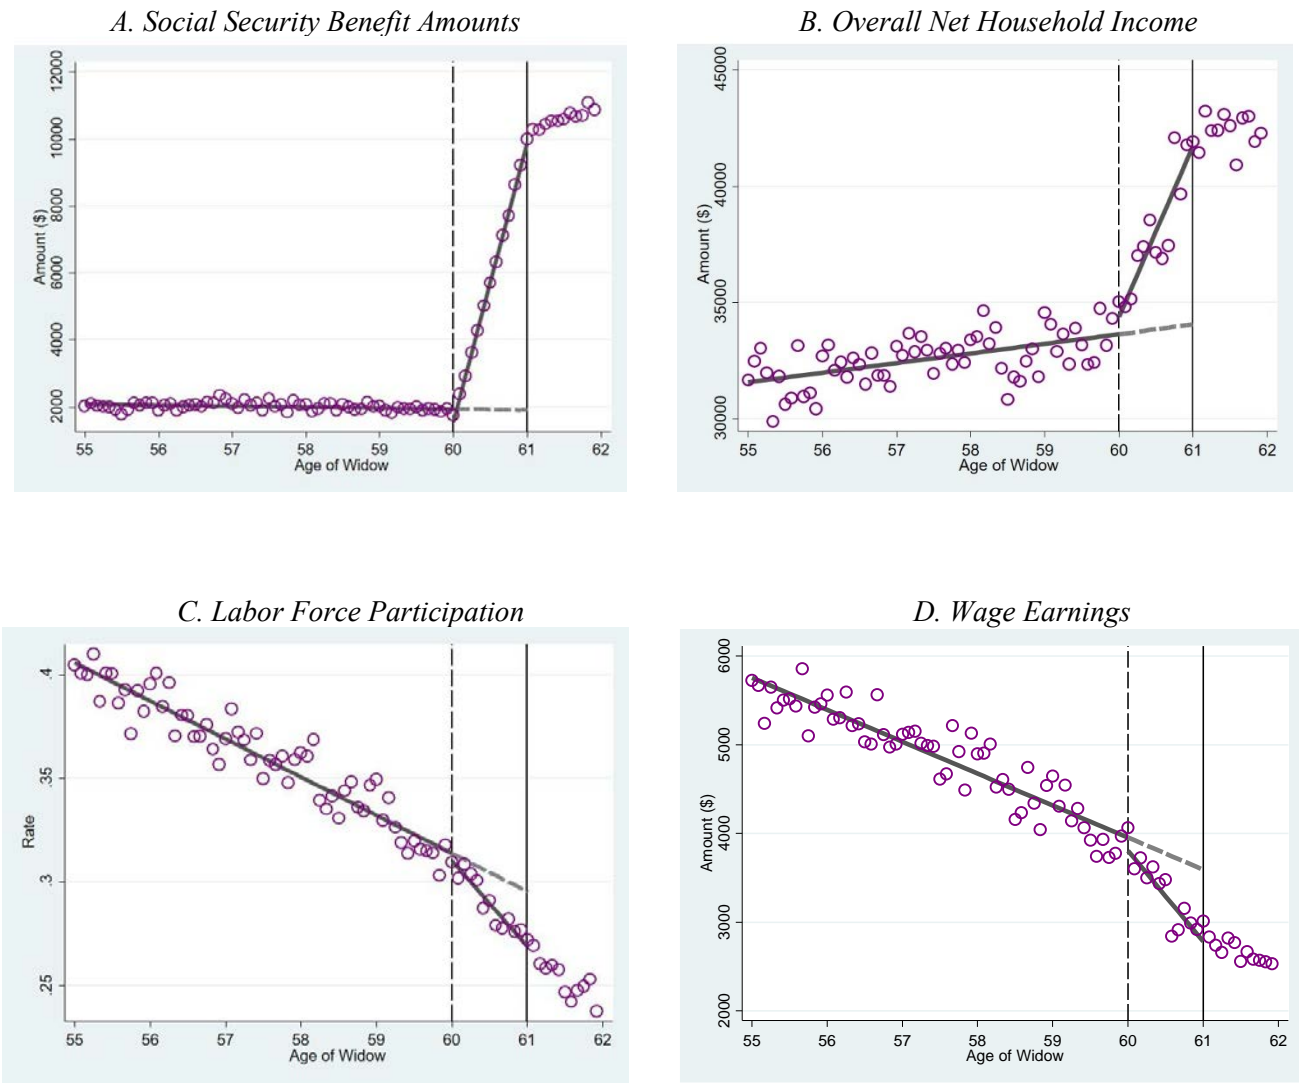
$$E[Y|D = 0, Z = 0] = \frac{\pi_n}{\pi_n + \pi_c} E[Y_0|never-taker] + \frac{\pi_c}{\pi_n + \pi_c} E[Y_0|complier]$$

$$E[Y|D = 0, Z = 1] = E[Y_0|never-taker]$$

Lastly, we estimate the observable moments as follows. First, for moments that pertain to ineligible households ($Z = 0$), we use predictions from a linear specification estimated based on households below the eligibility age (younger than 60): $y_{i,t} = \beta_0 + \beta_1(\text{age}_{i,t} - 60) + \varepsilon_{i,t}$. The estimator for the counterfactual is then $\beta_0 + \beta_1 \times (11/12)$. Second, for moments that pertain to eligible households ($Z = 1$), we calculate average outcomes based on households that are fully exposed, specifically using widows of the monthly ages $60\frac{11}{12}$ and 61. When the moments are conditional on takeup (D), we add that as a restriction to the estimation sample.

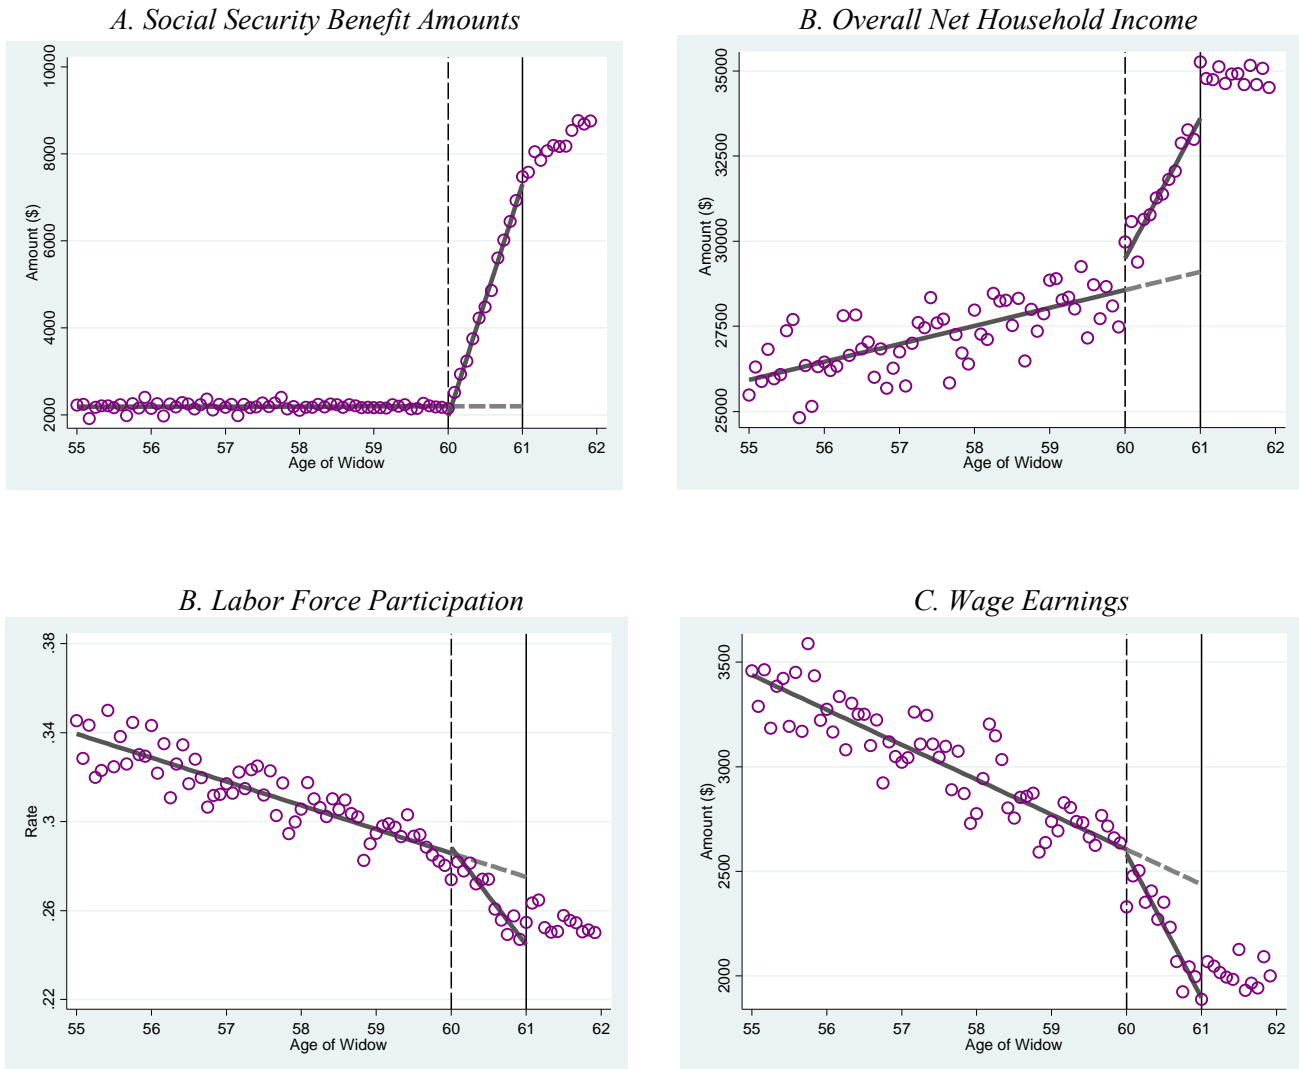
The results are as follows. For newly-widowed households: $\pi_a = 0.11643$; $\pi_c = 0.49792$; $\pi_n = 0.38565$. For household income: $E[Y_0|\text{never-taker}] = 59,244$ and $E[Y_0|\text{complier}] = 31,307$. Since the treatment effect for compliers is 9,355 it represents an increase of 29.9%. For earnings: $E[Y_0|\text{never-taker}] = 34,696$ and $E[Y_0|\text{complier}] = 10,050$. Since the treatment effect for compliers is -3,410 it represents a decrease of 33.9%. For already-widowed households: $\pi_a = 0.12053$; $\pi_c = 0.32421$; $\pi_n = 0.55525$. For earnings: $E[Y_0|\text{never-taker}] = 29,291$ and $E[Y_0|\text{complier}] = 13,203$. Since the treatment effect for compliers is -5,768 it represents a decrease of 43.7%.

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



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 Benefits by Already-Widowed Women with Lagged Labor Income below the Earnings Test



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 Eligibility for Social Security’s Survivors Benefit Receipt on Newly-Widowed Households

	Social Security Benefits Claiming Rate	Benefit Amounts	Overall Net Household Income	Labor Supply Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: Post-Shock Year $t = 1$</i>					
Full-Exposure Effect	0.54086*** (0.00461)	5,959*** (53)	4,912*** (501)	-0.03514*** (0.00504)	-1,784*** (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
<i>Panel B: Labor Income below Earnings Test</i>					
Full-Exposure Effect	0.60301*** (0.00461)	7,258*** (60)	7,074*** (522)	-0.02424*** (0.00482)	-1,065*** (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
<i>Panel C: Augmented Design—Control Group</i>					
Full-Exposure Effect	0.53745*** (0.00489)	5,759*** (56)	5,412*** (767)	-0.02866*** (0.00696)	-1,363*** (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

Notes: This table reports various estimations of the impact of eligibility for Social Security’s survivors benefits on newly-widowed households. Panels A and B estimate equation (1) and report the full-exposure effect based on the estimate for $\beta_2 + \beta_3 \times (11/12)$. Panel A uses observations from the immediate post-shock year $t = 1$. Panel B includes the sample of widows whose pre-shock earnings were below the earnings test thresholds and uses observations from the post-shock years $t = 1, 2$. Panel C augments our main 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 the years 2005-2007, and the treatment group is based on households that experience a spousal death in the years 2002-2004, so that all included observations are from the 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 $\gamma_2 + \gamma_3 \times (11/12)$. We report robust standard errors clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix Table 2: Effects of Eligibility for Social Security’s Survivors Benefit Receipt upon Spousal Death on Widow’s Longer-Run Outcomes

	Overall Net Household Income (1)	Labor Supply	
		Participation (2)	Wage Earnings (3)
Full-Exposure Effect	-126 (476)	-0.00395 (0.00528)	-167 (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 3: Effects of Eligibility for Social Security’s Survivors Benefit Receipt on Newly-Widowed Households by Degree of Liquidity

	Social Security Benefit Amounts (1)	Labor Supply	
		Participation (2)	Wage Earnings (3)
<i>Panel A: Low Liquidity</i>			
Full-Exposure Effect	5,057*** (51)	-0.02714*** (0.00454)	-2,219*** (514)
Number of Obs.	278,860	278,860	278,860
Number of Clusters	191,347	191,347	191,347
<i>Panel B: High Liquidity</i>			
Full-Exposure Effect	6,140*** (59)	-0.02053*** (0.00518)	-916*** (296)
Number of Obs.	225,244	225,244	225,244
Number of Clusters	162,387	162,387	162,387
Difference Low-High	-1,083*** (77)	-0.00661 (0.00696)	-1,303** (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 4: Responses to Anticipated Benefits 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*** (0.00454)	4,666*** (58)	4,201*** (334)	-0.02772*** (0.00404)	-497*** (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 5: Responses to Eligibility for Survivors 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** (0.00219)	-45 (32)	-0.00622*** (0.00193)	-44*** (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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Appendix Table 6: Effects of Eligibility for Social Security’s Survivors Benefit Receipt

	Social Security Benefits Claiming Rate	Benefit Amounts	Overall Net Household Income	Labor Supply Participation	Wage Earnings
	(1)	(2)	(3)	(4)	(5)
<i>Panel A: No Non-Death Withdrawals</i>					
<i>Newly-Widowed Households</i>					
Full-Exposure Effect	0.51162*** (0.00343)	5,553*** (40)	4,642*** (348)	-0.02890*** (0.00356)	-1,770*** (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
<i>Already-Widowed Households</i>					
Full-Exposure Effect	0.33599*** (0.00310)	3,650*** (36)	2,841*** (251)	-0.03156*** (0.00304)	-1,979*** (198)
Number of Obs.	538,247	538,247	538,247	538,247	538,247
Number of Clusters	224,443	224,443	224,443	224,443	224,443
<i>Panel B: No Pre-Period Contributions</i>					
<i>Newly-Widowed Households</i>					
Full-Exposure Effect	0.55553*** (0.00458)	5,905*** (54)	5,095*** (432)	-0.02447*** (0.00496)	-1,597*** (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
<i>Already-Widowed Households</i>					
Full-Exposure Effect	0.37308*** (0.00371)	4,191*** (45)	3,300*** (278)	-0.03140*** (0.00367)	-1,517*** (173)
Number of Obs.	387,944	387,944	387,944	387,944	387,944
Number of Clusters	170,683	170,683	170,683	170,683	170,683

Notes: This table reports various estimations of the impact of eligibility for Social Security’s survivors benefits. The analysis of newly-widowed households includes observations from the immediate post-shock years ($t = 1, 2$). The analysis of already-widowed households includes observations from periods $t = 6 - 10$ following the spousal death. Using specification (1), the table provides the effect of full exposure to eligibility for the program, which is captured by $\beta_2 + \beta_3 \times (11/12)$. Panel A uses households that do not make any withdrawals indexed with reasons other than “death.” Panel B includes all households who did not make contributions to savings accounts in previous periods (specifically, the pre-shock periods for newly-widowed households and the lagged period for already-widowed households). We report robust standard errors clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.