



The Effect of Social Security Information on the Labor Supply and Savings of Older Americans

Philip Armour and Michael F. Lovenheim



The Effect of Social Security Information on the Labor Supply and Savings of Older Americans

Philip Armour
RAND Corporation

Michael F. Lovenheim
Cornell University and NBER

September 2016

Michigan Retirement Research Center
University of Michigan
P.O. Box 1248
Ann Arbor, MI 48104
www.mrrc.isr.umich.edu
(734) 615-0422

Acknowledgements

The research reported herein was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement Research Consortium through the University of Michigan Retirement Research Center Award RRC08098401. The opinions and conclusions expressed are solely those of the author(s) and do not represent the opinions or policy of SSA or any agency of the federal government. Neither the United States government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of the contents of this report. Reference herein to any specific commercial product, process or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply endorsement, recommendation or favoring by the United States government or any agency thereof.

Regents of the University of Michigan

Michael J. Behm, Grand Blanc; Mark J. Bernstein, Ann Arbor; Laurence B. Deitch, Bloomfield Hills; Shauna Ryder Diggs, Grosse Pointe; Denise Ilitch, Bingham Farms; Andrea Fischer Newman, Ann Arbor; Andrew C. Richner, Grosse Pointe Park; Katherine E. White, Ann Arbor; Mark S. Schlissel, *ex officio*

The Effect of Social Security Information on the Labor Supply and Savings of Older Americans

Abstract

This paper examines how older workers adjust their labor supply in response to information they receive about their retirement wealth from the provision of the Social Security Statement. We find that older male workers' labor supply is highly responsive to receiving personalized information about future Social Security benefits, leading to a reduction of 119 hours worked per year, on average. However, our estimates point to significant heterogeneity in this response, with workers at the lower end of the hours-worked distribution increasing their labor supply and those at the high end decreasing their labor supply. We argue differences in knowledge about Social Security benefits across the labor supply distribution can explain much of this heterogeneity. We additionally explore the extent to which the information on the Statement may have led some workers to mistakenly reduce their labor supply by too much due to a lack of understanding of the dynamic nature of the Statement's benefit projections with respect to earnings. Receipt of a second Statement led all but the lowest hour workers to increase their labor supply relative to workers who did not receive a second Statement. This is consistent with workers misunderstanding the information provided as accumulated rather than projected wealth. Our results point to older workers being very responsive to Social Security information, which highlights the need to accurately convey information about both pension wealth and its sensitivity to changes in earnings.

Citation

Armour, Philip and Michael F. Lovenheim. 2016. "The Effect of Social Security Information on the Labor Supply and Savings of Older Americans." Ann Arbor, MI. University of Michigan Retirement Research Center (MRRC) Working Paper, WP 2016-361.
<http://www.mrrc.isr.umich.edu/publications/papers/pdf/wp361.pdf>

Authors' acknowledgements

We gratefully acknowledge funding for this research through a grant from the Michigan Retirement Research Center in the Social Security Administration's Retirement Research Consortium. This paper used the Health and Retirement Study, Respondent Cross-Year Summary and Detailed Earnings, and Respondent Cross-Year Benefits restricted use datasets produced by the University of Michigan with funding from the National Institute on Aging (grant number NIA U01AG009740), Ann Arbor, MI (2012) and distributed to authorized users only. We thank seminar participants at the MRRC Researcher Workshop, the NBER Summer Institute Economics of Aging Working Group, the RAND Corporation, Syracuse University, and Tulane University. We also have benefitted from discussions of this work with Richard Burkhauser, Gopi Goda, John Laitner, Nicole Maestas, Bruce Meyer, Kathleen Mullen, Ted O'Donoghue, Maria Fitzpatrick and Zhuan Pei.

1. Introduction

A central question in economics is how information affects decisions, especially when this information is imperfect. Older workers' retirement behavior is a particularly relevant area in which partial information may lead to sub-optimal decisions, as the incentives embedded in pension plans often are complex and difficult for workers to understand. Indeed, American workers have rather poor knowledge of their pension and Social Security wealth levels (Gustman and Steinmeier, 2001; Mastrobuoni, 2011). This lack of information provides a clear role for information-based interventions that can inform workers about their retirement wealth in order to help them make better intertemporal labor supply and savings decisions. The complexity of many pension systems in general and of Social Security more specifically makes it difficult to structure an intervention that provides information that workers will understand. As a result, there is a possibility of giving workers information that is misleading, which can cause optimization errors that render them worse off.

There currently is a poor understanding of whether workers' knowledge about their future pension benefits affects their labor supply decisions. It therefore is critical to study whether and how workers respond to interventions aimed at providing them with better information about their pension wealth. In particular, we are interested in whether labor supply is affected by future pension benefit information as well as whether workers use such information to make working decisions that render them better off. Pension systems typically are difficult to understand and embed complex incentives for workers that vary at different points in the age-experience profile, and few systems are more complex than Social Security's retirement benefits. Pension information for younger workers necessarily needs to be based on projections of future earnings, which leads to a tradeoff between providing workers with accurate information and providing them with information that is easy to understand. Giving workers a simple projection based on strong assumptions about future labor supply may maximize salience of the information, but it also may be misleading because it does not contain information about how current labor supply decisions map into changes in pension benefits. Conversely, if pension information consists of too much detail about how changes in earnings affect pension wealth, workers may ignore it. How to provide accurate and salient retirement information to workers is a question of central importance that has received little attention in the previous literature.

In this paper, we study the effects of the largest retirement information program in the US, the Social Security Statement, on the labor supply of older workers. The Statement was phased in from 1994 to 2000 according to worker age; that different-aged workers received the Statement in different years generates exogenous cross-cohort differences in the timing of information receipt. Because the exact timing of Statement receipt is based on birth month, there is additional within-birth cohort variation. Furthermore, workers receive multiple statements that are staggered over several years depending on their birth cohort, which allows us to examine how workers respond to updated Social Security benefit information. Prior research on the Statement, conducted variously by independent researchers (Greenwald et al. 2010, Mastrobuoni 2011), the GAO (T-HEHS-96-210, HEHS-97-19, HEHS-98-228, T-HEHS-00-101, and GAO-05-192), and SSA itself (Smith and Couch 2014), indicates a remarkably high recall of Statement receipt, with over two-thirds of *intended* recipients remembering receiving a Statement from SSA, even up to three years after being sent one. Over 90% of this population remembered that the Statement included personalized benefit projections.

A central motivation of this paper is that the Statement itself provided highly salient but limited information to workers: it informed them of their *projected* Social Security monthly benefit at ages 62, Full Retirement Age (FRA), and age 70, assuming constant earnings. That is, it assumed no change in earnings from the current age until the claiming age, nor any change in national price or wage indexes.¹ Given the complex relationship between lifetime earnings and Social Security benefits, such assumptions are necessary to provide any projection. However, the Statement did not include any measures of how sensitive these projections are to *changes in future earnings*. It hence was difficult if not impossible for workers to use the information contained in the Statement to forecast how labor supply changes would impact their future benefits. For example, the paper Statement does not provide an estimate of accumulated Social Security benefits at the time of calculation (i.e., their monthly benefit they would have if they stopped working today even if they did not claim benefits until 62 or the FRA). This feature highlights the importance of observing worker reactions to subsequent Statements, when they

¹ The projections in the Statement are based on the program parameters in place in the year of the Statement's production, and all benefit dollar amounts are in the current year's dollars. An alternative, weaker interpretation of these assumptions is that average price and wage growth will be equal, and the resulting benefit projection at each age is deflated to the current year's dollars.

would be able to see how their labor supply changes affected their projected benefits. We estimate how workers respond to initial receipt of the Statement and to repeated information about projected retirement benefit levels. The dynamic response provides insight into whether workers may have made errors in their labor supply decisions when they received the initial information.

The introduction of the Statement previously has been used to study the effect of retirement benefit information on worker benefit knowledge and retirement timing (Mastrobuoni 2011, Smith and Couch 2014). The findings indicate that although the Statement increased the accuracy of Social Security benefit predictions, it had no average effect on the timing of Social Security claiming or on the timing of self-reported retirement. The analysis of binary claiming and retirement decisions, however, can miss many of the ways in which workers' labor supply responds to information. For one, the transition to retirement is not binary. Many older workers reduce their labor supply quite dramatically on the intensive margin before leaving the labor force altogether, and they also frequently re-enter the labor force after they first leave (Ruhm 1990, Rust and Phelan 1997). Both of these behaviors are not captured by a retirement indicator variable. In addition, there are large spikes at the early and full Social Security retirement ages. These spikes likely are due to the incentives embedded in the Social Security system as well as rule-of-thumb behavior and interactions with other government programs and work rules. The large retirement spikes at these ages make it difficult to observe any impact of an intervention on a binary retirement measure, since so many individuals are not on the decision-making margin.

The first contribution of this paper is to estimate the effect of Social Security benefit information on a continuous measure of labor supply that captures both the intensive and extensive margins. Because of the often slow (and non-monotonic) transition from full-time work to full-time retirement, examining the direct labor supply measure of hours worked allows us to analyze in far more detail how labor supply decisions among older workers are influenced by this information intervention. We show extensive evidence that using a continuous measure of labor supply leads to dramatically different conclusions about the role of information than simply examining retirement behavior.

The second contribution of our paper is to examine the dynamic responses of workers to partial Social Security benefit information, i.e. benefit projections based on *constant* real

earnings. Does the partial nature of the information provided cause workers to make “mistakes” that are then corrected when the information is updated? We present evidence that indicates this is indeed what occurs. To our knowledge, this question has not been addressed by prior research.

We combine restricted-access Health and Retirement Study (HRS) data that include Social Security earnings histories on workers born between 1932 and 1947 with the timing of the rollout of the Statement across these birth cohorts. Our primary focus is on older male workers who have sufficient earnings to qualify them for Social Security when they become age eligible. We exclude women throughout most of the analysis because in the cohorts we consider, they are much less likely to have earnings that would qualify them for Old Age Insurance amounts greater than their spousal benefits. Additionally, we show that female labor supply responds to receipt of the husband’s Statement (but not vice versa), suggesting that female labor supply in our sample is subject to more complex intra-household bargaining behavior that is beyond the scope of our analysis. We do show estimates that include women, however, which are very similar to those for men only.

We first estimate the effect of Statement receipt on hours worked. Our results indicate that receiving the Statement reduced annual hours worked by 119 hours, which is an 11% reduction relative to the mean. We find much evidence of heterogeneity, most notably across the distribution of pre-Statement hours worked. Workers who were not working or who worked few hours substantially increase their labor supply, while there are declines in hours worked among those who were working full-time prior to Statement receipt. We present evidence that these findings are not driven by mean reversion but rather reflect the causal effect of the Social Security Statement on labor supply decisions. In addition, we show using HRS questions about expected Social Security benefits that these effects are driven by heterogeneity in pre-Statement worker knowledge. Workers whose prior Social Security expectations understate their benefit levels decrease hours worked when they receive a Statement, while workers who have little knowledge of their benefits or overstate them increase labor supply due to the information treatment. Thus, our results point to large labor supply responses to receiving a Social Security Statement that are driven in part by how the Statement affects workers’ knowledge about their pension wealth.

Given the evidence that Statement receipt leads most workers to reduce their labor

supply, we next examine the impact of receiving a second Statement that provides workers with updated information on their projected benefit levels. We hypothesize that some workers may misinterpret the information they receive such that they think the projected benefit on the Statement represents accumulated wealth, causing them to reduce their labor supply by more than if they had had full knowledge of the degree to which changing earnings affects future benefits. Receiving the second Statement provides them with information that their Social Security wealth has declined. If this decline was unintended, we then should see these workers partially or wholly reverse these prior labor supply decisions. We examine how receipt of the second Statement impacts hours of work among respondents with the same pre-Statement labor supply and who have had the first Statement for the same amount of time. Our approach thus identifies how second Statement receipt impacts the response to the first Statement.

The results point to marked labor supply increases among all but the lowest hours workers when they receive a second Statement relative to workers who have not yet received the second Statement. Receiving the second Statement attenuates the effect of the first Statement by between 53 and 65 percent; the effects are most prevalent for previously full-time workers. This evidence suggests receipt of the second Statement induces workers who are reducing their hours worked over time on the path to retirement to either increase hours worked or slow their reductions in labor supply. We argue these results are consistent with workers misunderstanding the information they received in the initial Statement, leading to an unintended reduction in Social Security benefits. As supporting evidence, we show using self-reported expected benefit levels that receiving a Statement leads workers to report that reduced earnings will not lower their Social Security benefits, contrary to what actual accrual rates are for most of these workers. Although we do not have direct feedback from individuals as to how sensitive they considered their benefit projections were to changes in earnings, our results show that workers attempted to correct decisions made based on first Statement information when subsequent information became available. For some workers, these intervening years of labor supply decisions that were subsequently at least partially reversed may have led to suboptimal earnings trajectories, insofar as updated Statements led to labor supply changes in the opposite direction of the first Statement effect.

Taken together, the results from this analysis suggest that information about public

retirement benefits, when provided by the public agency in question, has substantial effects on the labor supply of older, male Americans, whether this information is well-understood or not. Although Social Security Statements ceased being automatically sent out starting in April 2011, these automatic mailings restarted in September 2014. Our analysis sheds light on the essential difficulty of providing clear information without distorting knowledge of the dynamic qualities of pension programs. Specifically, providing a particular point estimate increased accuracy of expected benefits but it also decreased knowledge of how this benefit can vary as a function of labor supply. Given how responsive we find workers are to pension wealth information, much care needs to be taken to ensure its accuracy and transparency.

The structure of the paper is as follows: Section 2 describes essential components of Social Security benefits; Section 3 describes the Social Security Statement and its implementation; Section 4 discusses the data used in this analysis; Section 5 presents our analysis of the effect of the first Statement receipt on labor supply; Section 6 shows the dynamic labor supply analysis from second Statement receipt; and Section 7 concludes.

2. Old Age, Survivors, and Disability Insurance

Social Security, officially known as Old Age, Survivors, and Disability Insurance (OASDI), provides a suite of potential benefits to individuals who have a sufficient work history in the US. This program is large: in 2014, total expenditures were \$785 billion. Chief among these programs is the Old Age Insurance (OAI) portion. Because OASDI is a social insurance program, eligibility for benefits and benefit level are both based on one's history of covered earnings. OAI in particular requires individuals to have paid into the Social Security system with about 10 years of work for eligibility.²

For OAI benefit calculation, the highest 35 years of an individual's national-wage-level-indexed calendar year earnings are used. An Average Indexed Monthly Earnings (AIME) amount then is calculated. To determine one's Primary Insurance Amount (PIA), or monthly benefit available upon retirement at the Full Retirement Age, the Social Security Administration (SSA) applies a progressive benefit formula to one's AIME. As of 2015, this formula provides a 90% marginal replacement rate for the first \$826 of an AIME, a 32% marginal replacement rate for the next \$4,154 of the AIME, and a 15% marginal replacement rate for any remaining

² Specifically, the requirement to be insured is 40 Quarters of Coverage (QC), where in 2014 a QC is earned for every \$1,200 of earnings, up to 4 per year.

earnings up to a family cap. Hence, benefits are increasing in previous earnings, although at a decreasing rate. This PIA is then reduced if one opts for early retirement, available starting at age 62, or is increased if one delays collecting benefits after the Full Retirement Age, currently at 66.

Although a full discussion of program details is outside of the scope of this paper, a few points are relevant to the analysis below: if a potential retiree does not have 35 years of earnings in his work history, then his AIME will contain zero earnings years. Because most individuals are earning at their highest levels late in their careers, there can be large returns to work among older workers when these higher earnings years replace zero or low earnings years (Coile et al. 2002). The extent to which these high accrual rates apply depends on the specific pattern of a worker's earnings history, and thus workers with similar current or recent income levels may have vastly different returns to remaining in the labor force as they approach Social Security claiming ages.

Additionally, individuals can collect benefits based on their spouse's work history, generally limited to 50% of their spouse's PIA. Since we focus on older Americans in the 1990s in this sample, we limit the majority of our analysis to men largely to avoid the complex incentives facing women who may be deciding whether to collect benefits based on their husband's work history or their own. Men have been shown to be largely unresponsive to the impact of their own claiming behavior on spousal benefits (Sass et al. 2007), so our sample represents individuals responding to their own retirement benefits for whom receipt of a Social Security Statement based on their own earnings history would be directly relevant. However, we show that our estimates are robust to including the subset of women who are covered by Social Security and have a PIA of more than 50% that of their husband.

A large literature measures the effects that the various components of the Social Security system have on labor supply, largely through changes in the parameters or scope of these components.³ Most papers in this literature either implicitly or explicitly assume that workers know their future benefits and can accurately weigh alternative income streams when making labor supply and benefit collection decisions. Survey-based evidence, however, suggests that such sophisticated decision-making is rare. In the HRS, only about 50% of respondents are able to provide any estimate of their expected Social Security benefits and fewer than 30% of

³ For example, see Krueger and Pischke (1992), Friedberg (2000), Duggan et al. (2007), Mastrobuoni (2009) and Gelber, Jones and Sachs (2017). Krueger and Meyer (2002) provide a survey of studies that model retirement.

respondents are able to estimate their future benefits to within \$1,500 (in 2000 dollars) per year (Gustman and Steinmeier 2001). These results underscore that respondents do not have high awareness of the range of complex retirement incentives they face, and it is unlikely that they incorporate these incentives into their decision-making years in advance.

Chan and Stevens (2008) estimate that studies finding workers respond to pension incentives are based entirely off of the 20% of workers who correctly perceive these incentives. This highlights the critical role of information in driving worker behavior to pension incentives. Consequently, behavior inconsistent with these incentives obtains among a substantial portion of the population. For example, family members for whom it is more advantageous to delay collecting spousal benefits until after their own labor force exit are more likely instead to immediately collect benefits. Unmarried men who should immediately claim after exiting the labor force are more likely to delay collection as well (Gustman and Steinmeier 2000a).

More recent research has found that a majority of 50- to 70-year-olds understand that future Social Security benefits are linked to one's participation in the labor force on the extensive margin. These individuals also largely understand the incentives behind the delayed retirement credits and widow benefits (Liebman and Luttmer 2012).⁴ However, there are aspects of the Social Security system that impact intensive margin incentives about which individuals have a poor understanding, such as which and how many years of earnings are used in benefit calculations. Moreover, individuals' ability to operationalize this knowledge is unclear, or at least incomplete. In a field experiment designed to increase knowledge about Social Security benefits and the incentives embedded in the benefit formula, Liebman and Luttmer (2015) find that labor force participation increased by 4 percentage points, or over 5%. However, the effects are limited to women, and there is no evidence of an impact on intensive margin labor supply. The intervention we study differs from theirs in two notable ways: first, the Social Security Statement was a letter clearly sent from the Social Security Administration, which may have drawn more attention than a similar letter from a lesser-known agency, and second, Liebman and Luttmer (2015) did not provide any information about participants' Social Security wealth to them. They only provided information about Social Security program provisions, not individual-

⁴ The evidence in Liebman and Luttmer (2012) comes from a survey they conducted in 2008, when their sample would have been comprised of individuals who had received the Statement for at least 8 consecutive years. This may have increased their knowledge about their benefits.

specific benefit projections. In contrast, the Social Security Statement showed both personalized benefit projections as well as the earnings history of the recipient. Thus, responses to the two types of information may be quite different, as our results indicate.

Beyond this field experiment, research focused on understanding the effect of improving knowledge of Social Security incentives has been stymied by the endogeneity of information. Cross-sectional variation in program knowledge can be highly correlated with the benefits themselves and/or with labor force attachment. The staggered introduction of the Social Security Statement across birth cohorts combined with within-cohort differences in the timing of receipt based on birth month produced the type of exogenous variation in knowledge needed to analyze labor supply responses to projected Social Security benefit information.

3. The Social Security Statement

Starting in 1990, the Social Security Administration began providing standardized benefit statements for all individuals who requested them, and starting in late 1994, Statements were automatically sent out. These Social Security Statements eventually were sent annually to all individuals 25 and older between 2000 and 2011 who ever paid payroll tax. They contained personalized information about OASDI benefits upon retirement, disability, or death. The Online Appendix contains a fictional example Statement provided by the SSA.

Consistent with the scale of the mailings, the accuracy of recent addresses reported on tax forms, and the salience of receiving a document from the Social Security Administration with personalized benefit information, prior research has found high rates of recall of Statement receipt. Greenwald et al. (2010) found that over two thirds of individuals sent Statements recall having received one. Of those recalling receipt, 83%-90% report having read it carefully, with over 90% remembering that it contained personalized benefit calculations. A series of GAO reports finds results consistent with Greenwald et al. (2010).⁵ Smith and Couch (2014) show that personalized knowledge of the Full Retirement Age, own monthly benefit, and individual Social Security incentives is lacking before Statement receipt. This is despite there being rather high general knowledge about the Social Security program among Americans approaching retirement. Receipt of the Statement improves these areas of Social Security knowledge, as evidenced by Mastrobuoni (2011)'s findings of fewer errors in HRS recipients' estimated future benefits

⁵ These GAO reports are T-HEHS-96-210, HEHS-97-19, HEHS-98-228, T-HEHS-00-101, GAO-05-192 and can be accessed at www.gao.gov.

compared to projected benefits after they receive a Statement. These results indicate that the Statement was a highly effective outreach effort in terms of being read by the targeted recipients, with strong recall, both qualitatively and quantitatively, of the information presented therein. That Americans lack detailed knowledge about their own Social Security wealth and the incentives they face highlights the importance of providing *personalized* Social Security information to them, which is what the Statement did.

Prominent among the information contained in the Statement are projected retirement benefit levels if a retiree elects to receive benefits at the Early Eligibility Age (62), the Full Retirement Age (between 65 and 67, depending on birth cohort), and age 70. To construct the benefit information, the SSA uses each individual's earnings history up to the calendar year before the Statement's release. The SSA also includes expected future earnings up to the three ages listed on the Statement. These expected future earnings assume the individual will earn the last calendar year's earnings until collecting retirement benefits, with zero real wage growth at either the individual or national level.

Although there can be much debate over whether these assumptions are realistic or individually applicable, more concerning is whether individuals even understand that these retirement benefit levels are based on an assumption of continued similar earnings. Critical for our study, it is not possible to use the information on the Statement to project what might happen to benefits due to a change in labor supply. It is unclear whether individuals knew that any large change in labor supply could alter their projected benefits. The strongest manifestation of this confusion would be if workers believed the benefit levels shown were already accrued, in which case they might think reductions in labor supply would not affect their benefit level. Indeed, some researchers have expressed concern that the static nature of these estimates is misleading and argue conveying information on Social Security wealth accrual rates by different earnings trajectories would be more relevant to the decision-making of potential beneficiaries (Jackson 2006). This concern over the way in which this information was provided is a central motivation for our study.

The Statement was phased in across different age groups in the late 1990s. Automatic mailings began in Fiscal Year (FY) 1995 to those age 60 and over only. In FY 1996, Statements were automatically sent to those age 58 to 60; in FY 1997, 53 to 58; in FY 1998, 47 to 53; and in

FY 1999, 40 to 47. Beginning in FY 2000, Statements were mailed to all individuals 25 and over, including for the first time those who already had received a Statement.⁶ These FY 2000 mailings constitute second Statement receipt for all workers. Figure 1 illustrates which age groups received the Statement in which fiscal year as well as the total number of Statements sent out. This phase-in schedule provides a natural experiment in the provision of information about OASDI benefits in the late 1990s. As evident in Figure 1, there is variation by year and age in first Statement receipt, and consequently, in the time between first and second Statement receipt. The lack of information on how different earnings trajectories might affect benefit levels makes the second Statement receipt the main way workers could determine how their labor supply responses to the first Statement affected their Social Security wealth.

The cross-cohort and age variation in Statement receipt timing that we exploit in our analysis is shown in Figure 2. Each cell is filled in with the age of Statement receipt if that birth cohort was sent a Statement in that year.⁷ For example, the 1936-1938 birth cohorts (ages 58-60) received their first Statement in 1996, while the 1939-1944 cohorts (ages 53-58) had to wait until 1997 and the 1945-1947 birth cohorts (ages 51-53) received their first Statement in 1998. Beginning in 2000, the SSA began sending yearly Statements out to all individuals. As a result of this rollout pattern, there is cross-cohort variation in when individuals first received a Statement and variation in the time between first and second Statement receipt. This rollout also generates variation in age at both first and second Statement receipt, which is driven entirely by timing differences in initial Statement receipt. It is this variation in Statement receipt that allows for the identification of the effect of the Statement separate from age and year effects. We exploit the fact that otherwise similar cohorts have different Statement receipt patterns at different ages and in different years to identify the causal effect of the Statement information on labor supply of older workers.

The cross-cohort Statement patterns shown in Figures 1 and 2 understate the amount of variation in our data, since the actual timing of the Statement mailings depended on one's birth month in the year. Although Figure 2 shows the year in which the majority of each birth cohort

⁶ The years described here correspond to SSA fiscal years, which start in October. The exact timing of Statement receipt depends on one's birth month, but approximately one third of those 60 and over received a Statement in 1994.

⁷ Note that because Statements were sent out approximately three months before the birthday corresponding to the age listed, actual receipt would occur before turning this age; this receipt timing is coded into our analysis.

received a Statement, individuals were sent their Statements approximately three months before their birthday, so individuals with different birth months in the same birth year are sent Statements at different times throughout the year. When combined with variation in the timing of the HRS interview, the result is additional within-birth-year variation in Statement receipt. For example, if an individual born in 1937 is interviewed *before* his first Statement receipt, then we can use his labor supply decision in that survey wave as a control observation, in contrast to another individual born in 1937 who, because of some combination of his birth month and interview month, we observe after having received a Statement. Similarly, if an individual is interviewed in 2000 before his second Statement receipt, his labor supply in 2000 helps form the control group in our analysis of how individuals respond to multiple Statement doses. Because all second Statement receipt occurred in 2000, the within-cohort variation in timing provides much of the identifying variation in our second Statement receipt analysis. A main strength of our data is that it includes information on the month of birth of each respondent, which allows us to closely match the timing of statement receipt with respect to the interview date.

Previous research has shown that once one controls for age and year, no other factors influence Statement receipt, and that after having received these Statements, individuals are much more likely to be able to provide any estimate of their OAI benefits (Biggs 2010; Mastrobuoni 2011). Among those who already provided estimates, the accuracy also improves. However, in the only prior analysis of worker retirement effects of the Statement, Mastrobuoni (2011) found no average change in timing of collecting Old Age Insurance benefits. He also did not find any evidence that the Statement led benefit claiming to become more sensitive to variation in Social Security wealth. To date, there has been no direct analysis of the Statement's effect on labor supply of older Americans, though, which is the focus of this paper.

Theoretically, the effect of Statement receipt on labor supply is ambiguous. Essentially, the Statement provides individuals with *projected* information about their monthly retirement benefits. Prior work has shown that there is a sizable Social Security wealth effect on labor supply (Friedberg 2000; Liebman, Luttmer, and Sief 2009) and that the Statement led workers to more accurately predict their projected benefits (Mastrobuoni 2011). We therefore expect labor supply to decrease among those whose beliefs led them to understate their projected benefits prior to Statement receipt and who thus experienced a positive informational shock to retirement

wealth. We expect the opposite effect among those who over-state their benefits levels. Below, we show that most of the sample believes their benefits are lower than they actually are, which implies we should find an overall negative effect on labor supply. Since error in beliefs is also strongly correlated with pre-Statement labor supply, such that those at the high end of the hours-worked distribution systematically underestimate their future benefits, we examine heterogeneous effects by pre-Statement labor supply and by pre-Statement beliefs over future Social Security retirement benefits.

Predicting how workers will respond to projected Social Security benefit information is further complicated by the fact that the benefits shown are projections rather than accumulated wealth. Indeed, it is virtually impossible to use the information given in the Statement to predict how a given labor supply change will impact your benefits. Nonetheless, if people behave as if the benefit projections represent accumulated wealth, we should see a negative labor supply effect after initial Statement receipt as long as workers believe their benefits are lower than they actually are on average. Then, when they receive a second Statement, some workers who were induced to reduce hours worked from the first Statement should increase labor supply. We test these theoretical predictions directly in the remainder of the paper.

4. Data

We use restricted-access Health and Retirement Study (HRS) panels from 1992-2002 that are matched to Social Security earnings and benefits records. The HRS is a nationally-representative panel survey of individuals over age 50 and their spouses. The survey elicits information about demographics, income, assets, health, cognition, job status and history, expectations, and insurance. We focus on men born between 1932 and 1947. Our analysis sample excludes those in the baby boom cohorts (1948-1959), since they enter the HRS after the Statement had been universally provided to those 25 and older. These respondents are then matched to Social Security Respondent Cross-Year Summary Earnings, for which the match rate is approximately 72% among the cohorts we use (Mitchell et al. 1996). These records provide earnings from 1951 to the year of the match. The match is imperfect due to two factors: approximately a quarter of respondents do not grant permission to have their administrative records matched, and several individuals provided erroneous Social Security Numbers. Previous research using these matched data shows that for the Initial HRS Cohort, the matched subset is

an unbiased subsample (Kapteyn et al. 2006).

The largest problem when using the matched data is that the Social Security records are matched only up until a permission year, and for the vast majority of respondents in our sample there are only three permission years: 1992, 2004, and 2008. In a permission year, an HRS respondent is asked again whether the survey administrators can match his SSA records up until that year. Therefore, an individual must stay in the HRS until 2004 for researchers to observe his or her records past 1992. Although more recent cohorts in the HRS have provided permissions for prospective SSA record matching, our identifying variation occurs only during the retrospective match period. These individuals thus represent a skewed sample of younger and healthier respondents, and we therefore primarily use self-reported measures of earnings and hours worked instead of relying on administrative records post-1992.

We focus our analysis on men for two reasons. First, for this population of older workers, labor force participation rates of men are much higher than among women, and men represent the primary earners in their families. Second, because of their higher lifetime earnings, their Social Security Statement will be informative as to their retirement benefits, while their wives will be much more likely to collect spousal benefits. However, we also show results that include women; these estimates are very similar to those that only include men. This similarity is driven by the small number of women who have benefits of more than 50% of those of their spouse. We further limit our analysis to men under age 62, thereby avoiding the complex incentives facing someone who can choose to receive benefits immediately and for whom the Statement has different informational content. In effect, we are focusing only on men who can change their labor supply in *anticipation* of future Social Security benefits.

Using the SSA-matched data, we calculate whether individuals have earned the 40 Quarters of Coverage in their lifetime to be fully insured for OAI. We drop individuals who are not fully insured by 1991. Although they may subsequently work enough to gain OAI eligibility, their benefits will be very low and they represent an unusual sample of workers. Additionally, we drop those cohorts that were included in the HRS after they received their first Statement because for these workers we cannot measure pre-Statement labor supply. Table 1 shows the effect of these sample restrictions on the size of our primary sample. Ultimately, there are 21,904 observations corresponding to 4,038 unique respondents in our analysis sample.

For variable construction, we draw from the RAND Corporation’s pre-cleaned version of the HRS for self-reported earnings, hours worked, self-employment status, analytic weights, health status, IRA wealth, general assets not including IRAs, and pension information. All labor supply measures are contemporaneous to the time of the interview rather than retrospective.⁸ We use the HRS Tracker File for marriage status, birth and death information, and education as well as the HRS modules for expected OAI benefits at age 62 or 65. We calculate whether an individual had a second job before any Statement receipt, as well as the number of hours they worked in the year before the first Statement receipt. Tables 2 and 3 provide descriptive statistics of the variables we use in our analysis.

Our primary analysis uses the HRS as a natural sample, as is common practice in the Social Security program analysis literature using the HRS (Burkhauser et al. 2004, Li and Maestas 2008, Mastrobuoni 2011). The primary reason for this decision is that the weights are not available in all years for our entire analytic sample, and thus using them distorts the age composition of the sample.⁹ As a check on our results, weighted descriptive statistics and regression results are included in Online Appendix Tables A-3 and A-4 and show our conclusions are robust to using sample weights.

5. The Effect of Statement Receipt on Labor Supply

5.1. Empirical Methodology

Our goal is to estimate the effect of Statement receipt on the labor supply of older male workers. We employ difference-in-difference models that examine how labor supply of men in different cohorts changes when they receive a Statement. The baseline model is:

$$H_{iat} = \alpha + \beta First\ Statement_{iat} + X_{it}\theta + \sum_{j=1}^6 \gamma_j Hours\ Group_{ij}^{pre} + \delta_a + \rho_t + \varepsilon_{it}, \quad (1)$$

where H_{iat} represents annual hours worked of respondent i , of age a in year t . The variable $First\ Statement_{iat}$ is an indicator for whether an individual has received a Social Security Statement before the HRS interview took place in year t . The vector X_{it} is a set of demographic variables

⁸ Our labor supply questions derive from the RAND HRS variables `r#jhours`, `r#jhour2`, `r#jweeks`, and `r#jweek2`, which in turn are based on the “Usual hours worked per week” and “Usual weeks worked per year, counting paid vacation” questions in the HRS. These questions are identical across HRS panels.

⁹ Namely, the expansion of the HRS sample in 1998 included younger spouses of previous respondents in its sampling frame to reduce recruiting costs. However, these spouses have pre-1998 weights of zero, while our pre-Statement labor supply measures for this population derive from these earlier years. While we prefer unweighted estimates, our results are robust to dropping these individuals or including weights.

shown in Table 2, and the model includes both age (δ_a) and year (ρ_t) fixed effects.¹⁰

We control for pre-Statement hours worked by including a set of six indicator variables ($Hours\ Group_j^{pre}$) for whether an individual worked 1-9, 10-19, 20-29, 30-39, exactly 40, or over 40 hours per week in the survey wave immediately prior to first Statement receipt. The omitted category is workers who had zero hours of work. These pre-treatment labor supply controls serve two functions. First, they account for any heterogeneity across workers in pre-existing labor supply levels that may be correlated with the timing of the Statement rollout. Second, changes in labor supply can be influenced by mean reversion, since both low-hours workers and high-hours workers will naturally tend to revert to the mean. Controlling for pre-treatment labor supply helps account for this mean reversion, so we can identify whether workers in each hours group exhibit differential changes in labor supply when they receive the Statement relative to workers who work the same number of hours and who did not receive the Statement.

The main parameter of interest in equation (1) is β , which conditional on the controls in the model estimates the average change in labor supply from the pre-Statement level when a worker receives a Statement compared to a worker who has not yet received a Statement.¹¹ There are two main assumptions under which equation (1) allows us to identify the causal effect of Statement receipt on hours worked. First, the timing of Statement rollout must be unrelated to cross-cohort secular trends in labor supply, conditional on age and calendar year. If there are cohort-specific trends in hours worked that happen to be correlated with the timing of Statement rollout, this would bias our estimates. We believe such a situation is unlikely given the idiosyncratic variation in Statement receipt timing illustrated in Figure 2. Indeed, Mastrobuoni (2011) shows that conditional on controlling for age and year, no other observable factor predicts Statement receipt. Consistent with this assumption, we show below that our estimates remain unchanged if we only use within-cohort variation in receipt based on birth and survey month.

¹⁰ We do not control for the existence or value of private pension plans because this information is missing or unreliably volatile for a large proportion of the sample. Below, we examine heterogeneity by private defined contribution plan wealth, but we note that private pension wealth is unlikely to be a confounder in this setup because it is very unlikely that it is correlated with the rollout of the Statement. We performed additional sub-analyses using the pension wealth variables constructed by Gustman et al. (2014) and subsequently imputed for respondents with missing pension information, but we found no heterogeneity by pension holdings once controlling for pre-Statement labor supply.

¹¹ All workers in our sample eventually receive a Statement, so the control group in our model is comprised entirely of individuals who have not yet received a Statement but who will receive one in the future.

Figure 3 provides direct evidence that our estimates are not affected by differential cross-cohort trends correlated with the timing of Statement receipt. The figure presents event study estimates in which we re-scale calendar time to be relative to the year in which we first see each cohort post Statement receipt. We plot labor supply means and the 95% confidence interval by relative year and cohort that have been residualized to marital status, education, race, pre-Statement hours of work, and age and year fixed effects. Additionally, we include the difference-in-differences estimate from our preferred specification in Column 1 of Table 4. All respondents receive the Statement between relative time -2 and 0. There is no pre-receipt trend in labor supply, but there is a large decline in hours worked after the first Statement is received of over 100 hours that persists in the subsequent waves. It is clear from this figure that the timing of Statement receipt is associated with a sizable decline in labor supply.

Our final test for the existence of pre-treatment trends is to include a lead of Statement receipt in our empirical models. If Statement rollout is correlated with cohort-specific trends in hours worked, this lead variable will be large and statistically significant. However, consistent with Figure 3, we estimate a precise coefficient close to zero. These estimates are shown in Section 5.2 and support our main identification assumption.

The second identifying assumption is that there are no cohort-specific shocks that are correlated with the timing of Statement rollout. We do not find it very plausible that such systematic shocks would exist, as the Statement rollout allows us to control for year effects and age effects separately from any effect of the Statement. In short, the time-varying nature of Statement receipt makes it unlikely there were other factors that influenced the relative labor supply of cohorts in a way that was correlated with Statement receipt.¹² In particular, we are aware of no labor market policies that would have differentially affected these cohorts and that was rolled out contemporaneously with the Social Security Statement.¹³

¹² In addition, the fact that we find heterogeneous effects across the pre-Statement hours worked distribution suggests that any cohort-specific shocks that are correlated with the timing of Statement rollout would have to differentially affect low- versus high-hours workers.

¹³ There is one major Social Security policy change in our sample window: the rising Full Retirement Age. Starting in 2000, cohorts in our sample face a rising FRA for those born after 1937. However, specifications that include birth cohort-by-year fixed effects, shown in Table 6, provide nearly identical results to the baseline estimates. These estimates use only within-cohort variation in the timing of Statement receipt driven by differences in month of birth and interview month. In this specification, variation in Statement receipt thus is uncorrelated with differences in the FRA across individuals. Additionally, we conduct permutation tests that indicate our findings are highly dependent on following the exact cohort-phase-in pattern of the Statement and are thus not capturing labor supply responses to

A related concern is mean reversion: because we control for pre-Statement hours of work, our estimates could reflect reversion to the mean that would generate increases among low-hours workers and decreases among high-hours workers. We test directly for mean reversion using a set of randomization tests. This is done by randomly assigning the year of first Statement receipt to whole cohorts or to individuals. Using the randomly-assigned Statement year, we then estimate equation (1), including the controls for labor supply in the year prior to the randomly-assigned treatment year. If our estimates are being driven by mean reversion, we should get similar estimates when randomly assigning treatment years. In contrast, as discussed in Section 5.2.2, the results indicate no effect of the randomly-assigned Statement. These findings are inconsistent with mean reversion biasing our estimates.

Equation (1) estimates the average effect of having received a Statement on labor supply. This average effect, however, might mask substantial heterogeneity across subgroups in responsiveness to the Statement that is of high interest. Thus, we examine heterogeneity along several dimensions: age, educational attainment, whether the respondent has a second job, and pre-Statement hours worked. Because each of these sources of heterogeneity either are determined prior to Statement receipt or are not malleable (e.g., age), these interactions do not pose additional identification concerns.

The possibility for pre-treatment trends by these sources of heterogeneity raises a potential source of bias. To provide some evidence on the existence of such trends, Figure 4 shows raw hours worked by year relative to treatment broken down by pre-Statement hours of work bins, as well as the 95% confidence interval and the difference-in-differences estimates from column (6) of Table 4. There is no evidence for any group of differential pre-Statement labor supply trends. These figures also demonstrate one of our central findings: Statement receipt leads to increases in work hours among low-hours workers and decreases among high-hours workers.¹⁴

Other potentially important sources of heterogeneity in worker responses to information are the extent to which they have large or small pension returns to working and/or non-Social

the falling net present value of Social Security retirement benefits of later cohorts or other uncontrolled-for trends or changes that are not identical to the Statement's phase-in schedule.

¹⁴ The drop in hours worked among full-time workers in the years after initial Statement receipt is likely driven by a continuing path of decreasing labor supply, which we will explore further in our dynamic analysis in Section 6.

Security wealth. In order to account for heterogeneity in Social Security incentives for continuing to work, for each person-year, we construct a measure of the one-year Social Security retirement benefit return to work. This is done by calculating the present discounted value of the older worker's PIA (i.e. Old Age Insurance benefit collected at the Full Retirement Age) if he stopped working now compared to working one additional year at the same earnings level, with this benefit collected from the FRA onward.¹⁵ Although much research in the pension literature (Stock and Wise 1990, Samwick and Wise 2003; Coile and Gruber 2007) has emphasized that the one-year accrual rate does not measure optimal retirement timing, in this context it provides a straightforward measure of the returns to additional work among future Social Security beneficiaries. Its accuracy stems from the relatively monotonic nature of OASDI, especially when looking at workers under the age of 62, and the assumption of an additional year of earnings. Given our analytic design – examining the labor supply change from one interview to the next – this accrual rate provides a strong measure of the relative gains from continued labor supply for men under the age of 62 embedded in the Social Security system.¹⁶

In addition to the Social Security returns to work embedded in the one-year accrual rate, we create four measures of wealth: (1) self-reported non-pension wealth from all sources, including housing, liquid assets, vehicles, IRAs, etc.; (2) the prior definition plus current self-reported balances in defined contribution pension plans; (3) the prior definition plus the value of Social Security retirement and private pension benefits based on work up until the current year, discounted to the current year; and (4) Social Security wealth as measured by the AIME each worker has accrued up until the current calendar year.¹⁷ We examine how workers respond to Statement receipt by each of these wealth levels as well as by the one-year OAI accrual rate.

5.2. Results

5.2.1. Baseline Estimates

¹⁵ Survival probabilities were taken from the 1995 OASDI Trustees Report Lifetables, and future payments were discounted to the current age with a discount rate of 3% (Coile and Gruber 2001).

¹⁶ Coile and Gruber (2001) include auxiliary benefits (spousal and survivors) in their approach to finding optimal retirement timing. However, recent empirical evidence suggests that male workers incorporate only their own retirement benefits into their decision-making vis-à-vis Social Security (Sass et al. 2007, Henriques 2012), and if they do include spousal or survivors benefits, they place little weight on it in their decisions (Knapp 2013).

¹⁷ Measures (2) and (3) come directly from the latest HRS Imputations for Pension Wealth (V2.0) release set (ImpPenW), as variables PV_DB and DBXP.

The main results from estimation of equation (1) are shown in Table 4.¹⁸ In the table, each column presents results from a separate regression that does not use sample weights,¹⁹ and all estimates are accompanied by standard errors that are two-way clustered at the survey year and birth-year levels (Cameron, Gelbach and Miller 2011). In the first column, we show estimates that include all demographic controls, age and year fixed effects, and controls for pre-Statement labor supply. Column (1) demonstrates that Statement receipt reduces the amount of hours worked by 118.9 hours. This is an 11% decline relative to the mean hours worked of 1065.3 shown in Table 2. One way to put the size of this effect into context is relative to the Social Security wealth elasticity of 0.42 found by Liebman, Luttmer and Sief (2009): the effect of Statement receipt on labor supply is equivalent to increasing Social Security wealth by 26%. While large, this effect is consistent with the sizable increase in individual knowledge about projected Social Security benefits driven by the Statement (Mastrobuoni 2011).

As discussed in Section 5.1, a core concern with our difference-in-difference design is that the Statement rollout is correlated with cohort-specific trends in labor supply. To test for selection on fixed trends, we include an indicator for whether the respondent will receive a Statement by the next survey wave in column (2). The coefficient on the *I-Wave Lead of Statement* variable is small, precisely estimated, and is not statistically significantly different from zero at conventional levels. Furthermore, including this variable does not substantively change the estimate on the Statement variable or the R^2 . Thus, there is no evidence of cohort-specific labor supply trends that are correlated with Statement rollout.

There is much reason to believe that the effects of Statement receipt will differ across age groups. In fact, if workers correctly understand the Statement information, younger workers should not react at all to the information, as their PIA will be highly sensitive to hours worked over the remainder of their careers. If anything, we would expect there to be a positive effect among younger workers who are worried that their current PIA is insufficient for their expected

¹⁸ Online Appendix Table A-1 shows estimates using self-employment hours worked. These results are correspondingly noisier due to the paucity of self-employment relative to all-employment but indicate a similar pattern among these individuals who have greater control over their intensive-margin labor supply.

¹⁹ Online Appendix Table A-4 contains results that are akin to the results in Table 4 but use the 1994 sampling weights; descriptive statistics using the weights are shown in Table A-3. This restricts us to using the 1932-1941 birth cohorts, so we cannot examine effects among workers under 50. The weighted estimates match closely with the unweighted results. They are somewhat larger in absolute value, but that is to be expected since the younger workers with positive responses to the Statement now are excluded from the regressions. Overall, we see little evidence that using the larger, unweighted sample influences the conclusions one can draw from our results

retirement plans. Older workers, however, are more likely to reduce their labor supply if the Statement provides information that their Social Security wealth is high enough to fund their retirement. This is exactly the pattern we observe in Column (3) of Table 4. The estimates for workers in their 40s and early 50s are positive, although they are not statistically different from zero. Workers aged 55-61, however, significantly reduce their labor supply when they receive a Statement, and the estimates are much larger in absolute value for the 60-61 year old workers. These results suggest that older workers are more responsive to retirement wealth information.

In columns (4) and (5), we examine whether there are heterogeneous responses by worker education level and by whether a worker has a second job before the first statement receipt, respectively. We find no strong evidence of heterogeneous treatment effects with respect to education: those with graduate training are the most responsive, but high school graduates also reduce labor supply when they receive a Statement. In column (5), the results point to the largest effects of the Statement on those with a second job: such workers experience a 549 hour reduction in hours worked.²⁰ This heterogeneity likely is driven by the fact that those with two jobs have much more flexibility in hours than those with one primary job.

Finally in Table 4, we estimate whether workers respond differentially to receiving a Social Security Statement according to their pre-receipt hours of work. The results, shown in column (6), indicate a large amount of heterogeneity. Those who were previously not working and with low hours of work increase their hours worked, and those working full time (i.e., 40 hours per week or more) reduce their labor supply substantially. Among those who were previously not working (31% of the sample), there is an increase in hours worked of 474 hours, and for those who worked under 10 hours (1.3% of the sample) labor supply increased by 584 hours due to Statement receipt. The effect then decreases monotonically as pre-Statement work hours increase. For those working exactly 40 hours a week (23% of the sample), hours worked decreases by 402 hours per year. This is 19% relative to the mean hours worked for this group (Table 3). Among workers with more than 40 hours of work (34% of the sample), labor supply

²⁰ We also have examined whether those who work part time or those who report being in more hours flexible jobs are more responsive to Statement receipt. These estimates are largely uninformative due to small sample sizes. Additionally, we have examined effects for wage versus salaried workers. We find a negative effect on wage workers of -91 hours (se=40.87) and a positive effect on salaried workers of 244 hours (se=43.30). These estimates do not average to the effect in column (1) of Table 4 because we also allow the fixed effects and the coefficients on the observables to vary across these groups. All of these estimates are available from the authors upon request.

declines by 683 hours, or almost 25% of the baseline mean. These results clearly demonstrate that there is significant heterogeneity in the response to information receipt across the distribution of hours worked. One explanation for these results is that low-hours workers increase labor supply due to receiving information about low projected PIA levels, and vice versa for high-hours workers.²¹

As we discussed previously, our preferred estimates focus on men due to their strong attachment to the labor market, their likelihood of being the main earner in the household, and prior research suggesting that their labor supply behavior is less sensitive to spousal labor supply. Women were also sent Social Security Statements; however, for the birth cohorts during the Statement's phase-in, there is a substantially smaller number of women who have worked enough to qualify for OAI benefits based on their own earnings histories. As well, a large portion of the women in these birth cohorts are eligible for larger spousal benefits (50% of their spouse's PIAs) than own benefits. Labor force attachment for these cohorts of women is also substantially lower than for men.²²

It is instructive to consider how including women in the analysis affects the estimates and whether women respond differently to Statement receipt than men. We examine these questions in Table 5, which contains estimates of equation (1) for men and women together in column (1) and then separately for women in columns (2) and (3). As with men, we analyze those who are likely to receive their own Social Security retirement benefits: unmarried women and women with PIA levels greater than 50% of those of their spouse.

Column (1) of Table 5 shows that adding women into the estimating sample attenuates the results slightly but leads to extremely similar estimates as those in column (6) of Table 4. Given that including this sample of women adds fewer than 5,000 observations to the sample, this is not a surprising result. In column (2), we show that the estimates for women have a similar

²¹ Hours worked is our preferred measure of labor supply since it is a direct measure of worker behavior. Earnings estimates are somewhat hard to interpret in this setting, as workers may be switching between jobs with different wage rates, and there is substantially more variability in earnings than in hours worked from year to year. It thus is not obvious how the hours effects will translate into earnings changes. We show the results from estimation of equation (1) using earnings as the dependent variable in Appendix Table A-2. Earnings in this table are self-reported by respondents in the HRS. The earnings estimates are much noisier than the hours estimates in Table 4, but the signs and magnitudes of the estimates tell a very similar, though somewhat muted, story about worker responses to Statement receipt. On average, earnings decline by about \$599 per year (or 1.7%), and as with hours worked there is a strong age pattern.

²² We present counts of these OAI-eligible women and their pre-Statement labor supply in Appendix Table A-5.

pattern to those of men but are significantly smaller in absolute value. Nonetheless, women with low levels of pre-Statement labor supply increase hours worked upon Statement receipt and full-time female workers reduce labor supply when they receive the Statement.

In the final column of Table 5, we examine how women respond to own- and spousal Statement receipt.²³ We find similar reactions when women receive their own Statement or their husband's Statement. Those who receive both exhibit the opposite pattern across the pre-Statement labor supply distribution, but all of these effects are small. These results are suggestive of more complex intra-household bargaining decisions governing older female labor supply decisions that are not present for men. This is another reason why we focus on men for the remainder of the analysis. We leave for future work the study of gender differences in reactions to Social Security information among more recent cohorts given the growing rates of OAI-eligibility among women over time.

5.2.2. Robustness Checks

Table 6 provides a series of robustness checks of our main labor supply results that strongly support the validity of our identification strategy. In column (1), we reproduce the baseline estimates from Table 4. Column (2) excludes all government workers, who often receive an alternative pension plan in lieu of Social Security. The next three columns test the robustness of our estimates to the existence of cross-cohort shocks or trends. Column (3) includes separate linear year trends for each pre-Statement hours of work bin, column (4) shows results that control for cohort-by-year fixed effects and the column (5) results contain month-by-year fixed effects. In each of these cases, the results are extremely similar to baseline. The estimates in column (4) are of particular interest because they only use within-birth year cohort variation that is driven by differences in birth dates across respondents. Notably, they should be unaffected by the change in the full retirement age that occurred across the birth cohorts we study. These estimates are actually slightly larger than those in Table 4, although they are qualitatively similar.

In column (6) of Table 6, we estimate equation (1) only for those men who are initially ineligible for OAI. This is an imperfect sensitivity test because some of these individuals later become eligible and because the sample sizes are small. Taking the point estimates at face value

²³ Note that for unmarried women, these husband-Statement-receipt variables are always zero. In results available upon request, we find that men are completely unresponsive to their spouse's Statement receipt.

suggests OAI-ineligible respondents are much less influenced by Statement receipt than are OAI-eligible respondents. The final column of Table 6 shows extensive margin estimates, where the dependent variable equals 1 if the respondent works. Consistent with the positive hours effect among those who previously were not working, there is a large extensive margin response for the zero hours group. Most other groups do not experience an extensive margin effect from Statement receipt, and full-time workers slightly reduce employment (consistent with their negative labor supply response). On average, there is no extensive margin effect, which is what Mastrobuoni (2011) finds in his analysis of Statement receipt. However, this average masks some heterogeneity across the pre-Statement hours distribution.

As discussed in Section 5.1, a central identification concern in this analysis is mean reversion: If controlling for lagged hours worked induces mean reversion, it would appear as if low-hours workers were increasing labor supply and high-hours workers were decreasing their labor supply due to Statement receipt. In order to provide evidence on whether mean reversion is driving our results, we conduct randomization tests in which we randomly assign either entire cohorts or individuals to a treatment year. For each randomization algorithm, we conduct 10,000 replications, with replacement, in which treatment year is randomly assigned. We then construct lagged labor supply using the number of hours worked in the survey wave prior to the randomly-assigned first Statement year. We estimate equation (1) on each of the 10,000 replications and compare the Statement estimates to those in Table 4.

Table 7 shows where the overall estimates and estimates by hours worked fall in the CDF of treatment estimates from the randomized datasets. The first column of results shows results from randomizing cohort treatment years, and the second column shows results from randomizing individual treatment years. Focusing on the top row for the full sample, the actual estimate in Table 4 (column 1) is smaller than over 99.8% of the randomized estimates. If mean reversion were generating our results, than randomly assigning treatment years and controlling for lagged labor supply should produce the same results, as the same mean reversion should be at work. In contrast, decreases in hours worked are much larger at the actual time of treatment, which is inconsistent with a large role for mean reversion.

A similar conclusion holds when examining the results by pre-Statement hours of work. For low-hours workers (0-9 hours) for whom the effects are positive, the actual estimates are

larger than over 99% of the estimates from the randomized datasets. Among full-time workers (40 hours and over) for whom the effects in Table 4 are the most negative, the actual estimates are smaller than virtually all of the randomized effects. In Table 4 we find a much more muted effect for workers working between 10 and 29 hours, which is reflected in the fact that the randomized estimates are similar to the actual estimates. Among 30-39 hours workers, we also find that the actual estimates are smaller than most of the randomized ones, but the magnitudes differ somewhat across specifications. Overall, these results suggest that the heterogeneity by hours worked we show in Table 4 is not due to mean reversion but rather reflects how the causal effect of the Statement varies by pre-treatment labor supply.

5.2.3. Estimates Accounting for Wealth

One potentially important source of heterogeneity in how respondents react to Statement receipt is non-Social Security wealth. One might expect those with more wealth to be less responsive to receiving a Statement because Social Security is less important to their overall asset portfolio. However, wealthier people may reduce labor supply more with the receipt of information if they realize their Social Security benefits are small relative to their other wealth. Furthermore, those with higher OAI accrual rates may be less responsive to the Statement information, as they have a higher benefit return to remaining in the labor market.

In Table 8, we test for heterogeneity along these dimensions. Column (1) shows the baseline estimate from Table 4. In column (2) of Table 8, we estimate a version of equation (1) that includes a control for total non-pension wealth (in \$1,000s) and an interaction of wealth with Statement receipt. There is clear evidence that those with more non-pension wealth are more responsive to receiving a Statement. For each \$1,000 of wealth, the effect of receiving the Statement is reduced by 0.057 hours. However, in the next column when we include defined contribution pension plans as a part of wealth, there no longer is evidence of heterogeneity as a function of wealth. The same null result holds if we include the value of Social Security retirement benefits as described in Section 4. It is important to use caution when interpreting these results as there are few observations with defined contribution pensions in this cohort, but these estimates provide suggestive evidence that wealthier workers respond more to the Social Security Statement only when their wealth is not tied up in retirement savings.

In the final column of Table 8, we test whether there are heterogeneous effects of the

Statement as a function of the OAI one-year accrual rate (in percentage points) and of Social Security wealth levels.²⁴ The average accrual rate is 1%, and such an increase in the accrual rate would increase the effect of Statement receipt by 8.3 hours. Thus, those with a higher return to working respond to the information receipt by working more. In contrast, those with higher OAI wealth levels work less when they receive the Statement. Since these are individuals who tend to have a lower return to work with respect to their Social Security benefits, column (5) of Table 8 shows that the responsiveness to Statement receipt is sensitive to the Social Security returns to working. These results are consistent with previous findings of retirement responsiveness to Social Security benefits (Coile and Gruber 2007) and indicate that the information provision of the Statement contributes to this responsiveness.

5.2.4. The Role of Information

The main mechanism through which the Statement affects labor supply is to provide workers with more information about their future Social Security benefits. Prior work has shown the Statement allows individuals to make more accurate PIA predictions (Mastrobuoni 2011). We further explore this relationship by linking worker responses to the breadth and accuracy of their pre-Statement information set. We split workers into 5 groups based on their PIA predictions in the survey directly prior to receiving the first Statement: workers who do not expect benefits, workers who expect benefits but do not report an expected PIA, workers whose expected PIA is at least \$300 too high, workers whose expected PIA is at least \$300 too low, and workers whose expected PIA is within \$300 of the actual level.²⁵

Figure 5 shows the percentage of workers in each group and how this percentage varies across the hours of work distribution. Very few workers can provide an accurate PIA prediction prior to receiving the Statement; the majority of workers either cannot guess what their PIA is or guess an amount that is too low. As hours of work increase, the proportion understating their PIA increases, while the proportion that cannot provide an expected amount decreases. Thus, we hypothesize that a main explanation for the heterogeneity in results by pre-Statement hours is

²⁴ Here, Social Security wealth is calculated as the accrued present discounted value of Old Age Insurance benefits if an individual immediately left the labor force and collected benefits at the Full Retirement Age. It therefore represents the “stock” of Social Security wealth already accumulated, while the one-year accrual rate represents the “flow” to this wealth of working an additional year.

²⁵ We generate information on the accuracy of the PIA prediction using the administrative-linked SSA earnings data. We calculate the PIA each worker has accrued as of the survey year prior to receiving the first Statement and compare this amount to the worker-reported expected PIA in that year.

that those with high labor supply are more likely to under-estimate their benefits. When they receive the Statement, there is an income effect that leads to lower labor supply.

Table 9 provides further evidence supporting this hypothesis. We estimate a version of equation (1) in which we allow the effect of Statement receipt to vary by one's pre-Statement expectations category. We categorize people into 4 groups that are similar to those used in Figure 5, except combining the small number of workers who do not expect benefits with those whose expected benefits are more than \$300 lower than actual. We combine these respondents together because both groups experience a positive shock to predicted Social Security wealth. The aggregate hours declines we document are driven primarily by those whose PIA expectations are too low (or who do not expect benefits). As Figure 5 shows, these are more likely to be higher-hours workers. Furthermore, most workers expect benefits that are too low, which is why the overall effect is negative despite the heterogeneity shown in Table 9. Those whose expectations are too high strongly increase their labor supply, while those with accurate expectations do not alter their labor supply meaningfully. Workers who do not report an expected PIA increase their hours worked, but the increase is smaller than among those whose expectations are too high. This smaller increase likely reflects the more heterogeneous nature of this group. Taken together, these results suggest that the heterogeneity we find across the hours distribution is driven by workers having different beliefs about their retirement wealth prior to receiving the Statement.

6. Dynamic Effect of Second Statement Receipt on Labor Supply

6.1. Empirical Methodology

In Section 5 we established that receiving the Social Security Statement leads to sizable reductions in labor supply overall but has heterogeneous effects based on pre-treatment labor supply levels and accuracy of Social Security benefit knowledge. We now turn to the second question motivating our analysis: how do workers respond to multiple doses of information? The thought experiment underlying our approach is to consider two, otherwise identical, workers who have received one Statement the same number of years in the past. One of these workers then receives a second Statement, while the other worker does not. Our goal is to understand how the labor supply of the worker who receives the second Statement changes relative to the other worker.

Specifically, we estimate event study models in which event time is relative to the first

Statement receipt. We then allow for hours to change differentially by whether an individual has received a second Statement. Essentially, we augment the event study model to include interactions between the post-Statement periods and second Statement receipt, separately by pre-Statement hours of work categories. Our identifying variation thus comes from the relative timing of first versus second Statement receipt.

To make the model more parsimonious, we collapse the pre-Statement hours bins into 5 groups: 0-9 hours, 10-29 hours, 30-39 hours, 40 hours and >40 hours. The estimates in Table 4 suggest these groups are similar in their responses to first Statement receipt. We then estimate the following event study model, which we index with respect to survey waves rather than calendar years for ease of exposition:

$$\begin{aligned}
 H_{iaw} = & \alpha + X_{iw}\theta + \delta_a + \sum_{j=1}^5 \sum_{\tau=-2}^2 \gamma_{j\tau} \text{Hours Group}_{ij}^{pre} * I(w - w_0 = \tau)_{iw\tau} \\
 & + \sum_{j=1}^5 \sum_{\tau=1}^2 \phi_{j\tau} \text{Hours Group}_{ij}^{pre} * I(w - w_0 = \tau)_{iw\tau} * \text{Second Statement}_{iw} \\
 & + \sum_{j=1}^5 \psi_{jw} \text{Hours Group}_j^{pre} * I(w)_w + \epsilon_{iaw}
 \end{aligned} \tag{2}$$

In equation (2), w represents the given survey wave and ψ_{jw} are pre-Statement hours of work bin by survey wave fixed effects. Note that because each survey wave occurs in a specific calendar year, these fixed effects are akin to year by hours group fixed effects. The term w_0 represents the survey wave in which the respondent first received the Statement, and so $w-w_0$ is event time with respect to the first Statement receipt.²⁶ The $\gamma_{j\tau}$ coefficients thus show the effect of first Statement receipt and test for differential pre-trends separately for each pre-Statement hours of work bin. These estimates will be similar to those shown in Figure 4.

The coefficients of interest in equation (2) are $\phi_{j\tau}$, which show how the effect of the first Statement changes when respondents receive the second Statement. No respondent receives the first and second Statements in the same survey wave, so the earliest event time in which we can identify the effect of second Statement receipt is 1. Furthermore, by wave 3 in event time, almost all respondents have received a second Statement. We therefore restrict our analysis to three

²⁶ Recall that the HRS is conducted biennially, so that one wave in event time is equal to two years.

waves after first Statement receipt (relative waves 0-2).²⁷ Because we include pre-Statement work hours bins by survey wave fixed effects, we have to drop one of the $\gamma_{j\tau}$ coefficients. We exclude those working 0-9 hours 2 waves prior to first Statement receipt, so all estimates are relative to hours of work among this group. All other variables in equation (2) are as previously defined in equation (1).

The identification assumptions underlying this model are virtually identical to those for equation (1): the rollout of the Statements needs to be uncorrelated with cross-cohort secular trends and shocks in labor supply. Here, we invoke this assumption both for the first and second Statements. The variation in first Statement receipt is driven by the across- and within-birth cohort rollout pattern discussed previously. As long as the timing of the initial Statement rollout is conditionally exogenous, both the timing of when the initial Statement is received and the number of survey waves between first and second Statement receipt will be exogenous. That the first Statement is uncorrelated with secular trends in labor supply makes it very likely that the same is true for the timing of the second Statement rollout.²⁸

As with equation (1), a core identification concern with equation (2) is mean reversion. If changes induced by first Statement receipt revert back to the mean over time, we may attribute some of that mean reversion to the effect of the second Statement. However, the fact that there is variation in the time between first and second Statement receipt allows us to account for mean reversion. That in each relative time period after first Statement receipt only some respondents have received the second Statement allows us to disentangle the time pattern of responses to the first Statement from the second Statement. Any mean reversion thus should be accounted for by the $\gamma_{j\tau}$ estimates.

Another potential concern with our approach is that the change in projected Social Security benefits is highly dependent on one's work history. Social Security retirement benefits are based on the highest 35 years of a worker's earnings, and therefore there is a natural limit to

²⁷ We also exclude observations for relative waves before -2 for ease of comparison across our pre-Statement hours bins, since there are insufficient observations to estimate coefficients for all of these bins in the -3 wave. However, inclusion of these observations has no substantive impact on the other coefficients estimated or on the overall fit of the model. Estimates from this alternative specification are available upon request.

²⁸ We again note that although all workers in our sample were sent a second Statement in 2000, heterogeneity in interview timing and birth month in our sample results in a portion of our sample being interviewed in 2000 before their second Statement receipt, allowing us to control for general labor market shocks in that year separate from second Statement receipt.

the amount these benefits can change by just one additional year of work. However, the Social Security Statement uses the most recent full year's earnings as the basis of projected earnings for *every year* from the Statement's construction to the listed benefit collection age. Therefore, changes in current earnings can have immediate and large impacts on the Statement's benefit projection. For example, if a 50-year-old earning at the highest level of his career and who plans on retiring at the Late Retirement Age of 70 decides to go to part time work earning half as much after he receives the first Statement, then a second Statement issued the very next year will use 19 years (of 35 computation years) with this new lower earnings rate. This new Statement will show a substantially lower projected benefit, even after just one year of lower labor supply, and the now 51-year-old may be surprised at how much his benefit has dropped.

In order to demonstrate the degree and rapidity with which these projected benefits can fall, we construct a counterfactual Statement in each year with benefit projections not based on current earnings but instead assuming either no earnings or half the current earnings level. Table 10 shows the percent decline in the projected Social Security retirement benefit at the FRA if an individual in our analysis sample who just received his first Statement either stopped working or halved his earnings. As is clear, even among an older population, most workers would experience a decline in projected benefits. Thus, there is much scope for declines in hours worked to affect projected PIA levels and worker behavior, which is the motivation behind estimating equation (2).

6.2. Results

Table 11 shows estimates of equation (2). All estimates in the table come from one regression, with standard errors that are two-way clustered at the survey wave and birth year levels. The first five rows show the effects of first Statement receipt and accord closely with the estimates in Figure 4: there are no differential trends in hours worked prior to Statement receipt, and post-receipt hours increase for previously low hours workers and decline among previously full time workers. Among full-time workers, the effects grow somewhat over event time, suggesting that the effect of first Statement receipt is to immediately reduce hours worked as well as to put workers on a path of declining labor supply over time. Thus, rather than mean reversion, the effect of Statement receipt is to put full time workers on a quicker path to retirement.

The estimates in the last five columns indicate that receiving the second Statement substantially mitigates this effect. Among previously full-time workers, there is a statistically significant relative increase in hours worked in event waves 1 and 2 when workers receive the second Statement. This effect grows over time: by the second event wave, the second Statement attenuates the reduced hours worked effect of the first Statement by between 53 (495.1/928.1) and 65 (348.7/537.5) percent. That is, relative to those who do not receive a second Statement, hours worked increase by 348.7 hours per week among those who previously worked 40 hours and increase by 495.1 hours per week among those who worked over 40 hours. These results clearly indicate that receiving the second Statement induces many previously full-time workers who reduced their labor supply when they received the first Statement to reverse course and either increase hours worked or reduce their labor supply at a slower rate.

Results among the other pre-Statement hours of work groups are qualitatively similar in indicating the opposite effect from that of first Statement receipt, but the estimates are not statistically significantly different from zero at even the 10% level. Among the 10-29 and 30-39 hours groups, the increase in labor supply in event year 2 is almost identical to the effect sizes among full time workers in percent terms. The relative decline in hours worked among low-hours respondents is smaller in magnitude at 26%, but it is still substantial. Taken together, the results in Table 11 suggest that receiving the second Statement leads workers to alter their labor supply in the opposite direction from their response to receiving the first Statement.

This evidence is consistent with a set of workers making sub-optimal labor supply decisions in response to the information on the first Statement due to a misinterpretation of the information provided. The first Statement caused many workers to reduce hours, yet once they received updated information, they responded by increasing their labor supply relative to those who had not yet received such information. While our preferred interpretation of the results in Table 11 is that workers make “mistakes” from misinterpreting the Statement information, it is possible we are instead picking up planned intertemporal changes in labor supply that happen to be correlated with Statement rollouts. If so, the labor supply changes we analyze should show up in respondent-reported expected PIA. That is, if workers reduced their labor supply in response to the Statement knowing full well it would reduce their Social Security wealth, we should observe reported PIAs declining.

We explore this question in Table 12. Here, we use all workers who report an expected PIA as well as workers with an accrual rate of at least 1% to focus on those whose Social Security benefits are strongly tied to their continued labor force participation. For workers with an accrual rate of at least 1%,²⁹ we can be sure that their Social Security wealth will decline if they reduce their labor supply. The dependent variable in columns (1)-(4) of this table is an indicator for reporting that one's expected PIA declined between this survey and the prior survey wave.³⁰ In columns (1) and (3), there is no evidence that workers who decrease their labor supply believe this will reduce their PIA, on average. This is direct evidence that even older workers do not have a firm grasp of the fact that reducing their hours will negatively influence their Social Security wealth. These estimates support our contention that the changes in labor supply shown in Table 11 reflect workers misinterpreting the information provided on the Statement. In columns (5) and (7), we repeat the same exercise with increased expected PIA. Here, there is a strong link between increased hours worked and workers reporting increases in expected PIA.³¹ This asymmetry in beliefs, with workers understanding that working more will increase benefits but not understanding that reducing work will lower benefits, is a novel finding in this literature.

In Columns (2) and (4) of Table 12, we examine how the likelihood a respondent reports a decline in expected PIA responds to Statement receipt. Interestingly, for those who have not received a statement, some workers correctly report that reductions in their labor supply will reduce their PIA. However, having received the Statement eliminates this effect, especially for those individuals who will actually experience a decline in their projected PIA due to lower earnings. Columns (6) and (8), however, show that the reverse situation does not hold: receiving a Statement does not affect workers' understanding of the positive link between working more and Social Security benefits. These results support our contention that workers interpreted the Social Security Statement benefits as accumulated wealth that only could grow over time if

²⁹ This accrual rate is defined as the change in the present discounted value of future Social Security retirement benefits that arises from retiring immediately compared to working an additional year.

³⁰ Prior reductions in labor supply are potentially endogenous. For this exercise to be informative, we require that there are no changes in the underlying endogeneity between prior changes in labor supply and changes in PIA expectations that are correlated with the timing of Statement receipt. Because of the exogenous rollout of the Statement, this assumption is plausible.

³¹ Leibman and Luttmer (2012) find that about 57% of respondents to their survey report that having higher earnings will increase future benefits. Notably, they do not ask about how they think declines in earnings will affect benefits. Our results suggest asymmetric beliefs on the part of older workers, which is not inconsistent with the findings in Leibman and Luttmer (2012).

added to through additional earnings.

On the whole, this evidence suggests that many older workers misinterpreted the information they were given. While there is some knowledge that reducing hours worked will lower one's PIA, a subset of workers did not realize that their PIA would not change, or not change substantially, when they reduced their labor supply, despite the fact that the Statement made clear that the estimated monthly benefits were projections based on continued constant earnings. As a result, many workers reduced their hours worked. When they received the updated Statement, however, they increased their hours worked relative to those who did not receive a second Statement. The second Statement essentially led these workers to either increase labor supply or slow their continuing path to retirement. Our findings thus are consistent with workers being highly responsive to retirement wealth information. But, these findings also highlight the need for this information to be easy to understand and convey how benefits may change if workers change their retirement plans, or workers will respond to the information by making optimization errors.

The behavior observed after second receipt and the change in how individuals perceive the dynamics of their OASDI benefits confirms prior results that the Statement itself gives workers a better point estimate of their projected retirement benefit (Mastrobuoni 2011). Yet, it renders them less likely to understand how this amount is linked to their future labor supply decisions. In effect, it provides a shock to knowledge of income, thus inducing an income effect, while diminishing knowledge about the tradeoff between work and benefits, thus diminishing the substitution effect. These results underscore the importance of giving older workers accurate information about the work incentives embedded in their retirement benefits programs that they can use to make more informed labor supply decisions towards the end of their careers.

7. Conclusion

This paper examines how older workers respond to information they receive about their retirement wealth from the Social Security Statement. We exploit the rollout of the Statement to different birth cohorts and different years that we argue was uncorrelated with the underlying labor force attachment and retirement preferences of each cohort. One of our main contributions is to examine labor supply responses that include both the intensive and extensive margins, and we find that older workers' labor supply is highly responsive to receiving information about

future Social Security benefits. Statement receipt leads to a reduction of 119 hours worked, on average. However, our estimates point to significant heterogeneity in this response across the distribution of pre-treatment hours worked, with low-hours workers increasing their labor supply post-treatment and high-hours workers significantly reducing their labor supply in response to receiving a Statement. We provide evidence that this heterogeneity can be explained in part by the different knowledge and beliefs workers have about their Social Security benefits. Given the high rate of recall among intended recipients of the Statement and the corresponding shock to expected Social Security benefits estimated by Mastrobuoni (2011), our average effect sizes are largely consistent with prior estimates of labor supply responsiveness to changes in Social Security wealth.

In the second part of our analysis, we explore the extent to which the information on the Statement may have led some workers to mistakenly reduce their labor supply. This would happen if workers thought the information they were receiving was accumulated rather than projected wealth. We find evidence that this is the case using the rollout of the second Statement workers received that allowed them to update their information. In the absence of receiving the second Statement, all but very low hours workers continue to reduce their hours worked over time after receiving the first Statement. Second Statement receipt, however, causes workers to increase their labor supply relative to those who have not yet received this Statement. For all but the lowest hours workers, the effect of the second Statement attenuates the first Statement effect by 53-65%. Thus, workers who reduce their labor supply because of first Statement receipt further reduce their hours worked along the path to retirement over time, but receiving the second Statement causes them either to increase their labor supply or slow their path to retirement. These effects are particularly pronounced among full time workers. Furthermore, we show evidence that the Statement decoupled the link between labor supply declines and projected Social Security benefits, which is consistent with workers interpreting the information on the Statement as accumulated benefits. These results highlight the importance of providing clear information to workers, as it is likely at least some people were left worse off between receipt of their first and second Statements due to their misunderstanding of how benefits would respond to labor supply changes.

Taken together, our results suggest that information older workers have about their

retirement benefit levels is an important driver of their labor supply. However, our estimates also underscore the difficulty in providing this information: conveying just point estimates decreases the complexity of the information but also may crowd-out knowledge of the dynamic structure of these projections. That workers are so responsive to this information provides much scope for information-based policy interventions to help them make more informed labor supply decisions as they approach retirement. An important question for further research is how to provide this information in the most salient and accurate way possible to avoid the misunderstandings that can arise from complex information.

References

- Armour, Philip, 2014 "The Role of Information in Disability Insurance Application: An Analysis of the Social Security Statement Phase-in." Unpublished manuscript.
- Biggs, Andrew G, 2010. "Improving the Social Security statement." Financial Literacy Center Working Paper.
- Burkhauser, Richard V., John S. Butler, and Gulcin Gumus, 2004. "Dynamic programming model estimates of Social Security Disability Insurance application timing." *Journal of Applied Econometrics* 19.6 (2004): 671-685.
- Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, 2011. "Robust inference with multiway clustering." *Journal of Business & Economic Statistics* 29.2 (2011).
- Carlson, Barbara Lepidus, and Stephen Williams, 2001. "A comparison of two methods to adjust weights for non-response: propensity modeling and weighting class adjustments." *Proceedings of the Annual Meeting of the American Statistical Association*.
- Chan, Sewin, and Ann Huff Stevens, 2008. "What You Don't Know Can't Help You: Pension Knowledge and Retirement Decision-making." *The Review of Economics and Statistics* 90.2 (2008): 253-266.
- Coile, Courtney, Peter Diamond, Jonathan Gruber, and Alain Jouten, 2002. "Delays in claiming Social Security benefits." *Journal of Public Economics* 84, no. 3: 357-385.
- Coile, Courtney, and Jonathan Gruber, 2001. "Social Security incentives for retirement." In *Themes in the Economics of Aging*, pp. 311-354. University of Chicago Press.
- Coile, Courtney, and Jonathan Gruber, 2007. "Future Social Security Entitlements and the Retirement Decision." *The Review of Economics and Statistics*, 89(2): 234-246.
- Duggan, Mark, Perry Singleton, and Jae Song, 2007. "Aching to Retire? The Rise in the Full Retirement Age and its Impact on the Social Security Disability Rolls." *Journal of Public Economics* 91.7: 1327-1350.
- Friedberg, Leora, 2000. "The Labor Supply Effects of the Social Security Earnings Test." *Review of Economics and Statistics* 82.1: 48-63.
- Gelber, Alexander, Damon Jones and Daniel W. Sachs. 2017. "Estimating Earnings Adjustment Frictions: Method and Evidence from the Social Security Earnings Test." Mimeo
- Greenwald, Mathew, Arie Kapteyn, Olivia S. Mitchell, and Lisa Schneider, 2010. "What Do People Know About Social Security?." *Financial Literacy Consortium Report to the SSA*.
- Gustman, Alan L., and Thomas L. Steinmeier, 2000a. "Research on Pensions and Social Security," NBER Summer Reporter.
- Gustman, Alan L., and Thomas L. Steinmeier, 2000b. "Retirement Outcomes in the Health and Retirement Study," NBER Working Paper no. w8506.
- Gustman, Alan L., and Thomas L. Steinmeier, 2001. "Imperfect Knowledge, Retirement and Saving." NBER Working Paper no. w8406.
- Gustman, Alan L., Thomas L. Steinmeier, and Nahid Tabatabai, 2014. "Updated Pension Wealth Data Files in the HRS Panel: 1992 to 2010 Part III." Health and Retirement Study researcher contributed dataset.
- Henriques, Alice, 2012. "How does Social Security claiming respond to incentives? Considering husbands' and wives' benefits separately." Board of Governors of the Federal Reserve WP.
- Jackson, Howell E., 2006. "Accounting for Social Security Benefits." *Behavioral Public Finance* (2006): 261-303.
- Kapteyn, Arie, Pierre-Carl Michaud, James Smith, and Arthur van Soest, 2006. "Effects of Attrition and Non-Response in the Health and Retirement Study." No. 2246. Institute for the Study of Labor (IZA).

- Knapp, David, 2013. "The Effect of Social Security Auxiliary Spouse and Survivor's Benefits on the Household Retirement Decision." Unpublished manuscript.
- Krueger, Alan B., and Bruce D. Meyer, 2002. "Labor Supply Effects of Social Insurance." *Handbook of public economics* 4: 2327-2392.
- Krueger, Alan B., and Jörn-Steffen Pischke, 1992. "The Effect of Social Security on Labor Supply: A Cohort Analysis of the Notch Generation." *Journal of Labor Economics* 412-437.
- Li, Xiaoyan, and Nicole Maestas, 2008. "Does the Rise in the Full Retirement Age Encourage Disability Benefits Applications? Evidence from the Health and Retirement Study." MRRC Working Paper 2008-198.
- Liebman, Jeffrey B., and Erzo FP 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, Volume 26*. University of Chicago Press. 1-42.
- Liebman, Jeffrey B., and Erzo FP Luttmer, 2015. "Would people behave differently if they better understood Social Security? Evidence from a field experiment." *American Economic Journal: Economic Policy*: 7(1), 275-99.
- Liebman, Jeffrey B., Erzo FP Luttmer, and David G. Sief, 2009. "Labor Supply Responses to Marginal Social Security Benefits: Evidence from Discontinuities." *Journal of Public Economics* 93(11-12): 1208-1223.
- Mastrobuoni, Giovanni, 2009. "Labor Supply Effects of the Recent Social Security Benefit Cuts: Empirical Estimates using Cohort Discontinuities." *Journal of Public Economics* 93.11: 1224-1233.
- Mastrobuoni, Giovanni, 2011. "The Role of Information for Retirement Behavior: Evidence Based on the Stepwise Introduction of the Social Security Statement," *Journal of Public Economics*.
- Mitchell, Olivia, Jan Olson and Thomas Steinmeier, 1996. "Construction of the Earnings and Benefits File (EBF) for Use with the Health and Retirement Survey," NBER WP no. w5707.
- Ruhm, Christopher J., 1997. "Bridge jobs and partial retirement." *Journal of Labor Economics*: 482-501.
- Rust, John, and Christopher Phelan, 1997. "How Social Security and medicare affect retirement behavior in a world of incomplete markets." *Econometrica: Journal of the Econometric Society*: 781-831.
- Samwick, Andrew, and David A. Wise, 2003. "Option value estimation with Health and Retirement Study data." In *Labor markets and firm benefit policies in Japan and the United States*, pp. 205-228. University of Chicago Press.
- Sass, Steven A., Wei Sun, and Anthony Webb, 2007. "Why Do Married Men Claim Social Security Benefits So Early? Ignorance or Caddishness?." *Center for Retirement Research at Boston College Working Paper* 17.
- Smith, Barbara A., and Kenneth A. Couch, 2014. "How Effective Is the Social Security Statement: Informing Younger Workers about Social Security." *Soc. Sec. Bull.* 74: 1.
- Social Security Administration, 2013. "The Online Statement and MySSA Portal: SSA Should Take Additional Steps to More Effectively Communicate with the Public, December 2013," http://ssab.gov/PublicationViewOptions.aspx?ssab_pub=122.
- Stock, James H. and David A. Wise, 1990. "Pensions, the Option Value of Work, and Retirement." *Econometrica* 58, no. 5: 1151-1180.

Figure 1. Social Security Statement Phase-In Schedule

	SSA Fiscal Year							
	1994	1995	1996	1997	1998	1999	2000	2001
60 and Older		X					X	X
58-60			X				X	X
53-58				X			X	X
47-53					X		X	X
40-47						X	X	X
25-40							X	X
Total Statements Sent (millions)	0	7	5.5	12.4	20.7	26.6	134.7	135.6

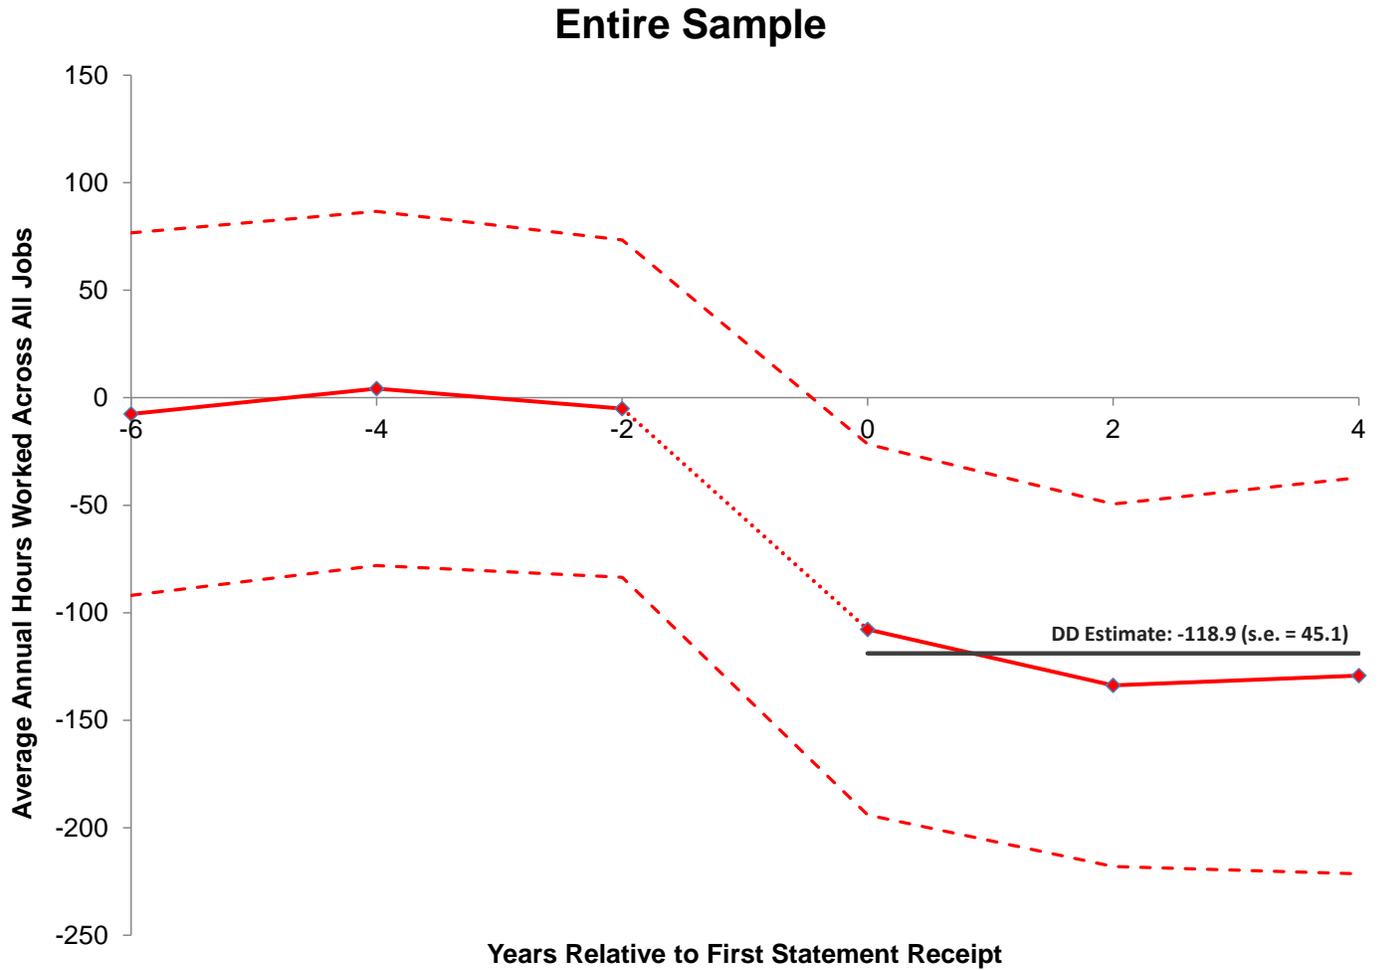
Notes: SSA Fiscal Years are October of the preceding calendar year to September of the stated year. No Statements were sent out before Fiscal Year 1995, and all individuals with Social Security Numbers age 25 and over received a Statement yearly from 2000 to 2011.

Figure 2. Social Security Statement Phase-in, by Birth Cohort and Age at Statement Mailing

Birth Year	SSA Fiscal Year							
	1994	1995	1996	1997	1998	1999	2000	2001
1934		61					66	67
1935		60					65	66
1936			60				64	65
1937			59				63	64
1938			58				62	63
1939				58			61	62
1940				57			60	61
1941				56			59	60
1942				55			58	59
1943				54			57	58
1944				53			56	57
1945					53		55	56
1946					52		54	55
1947					51		53	54

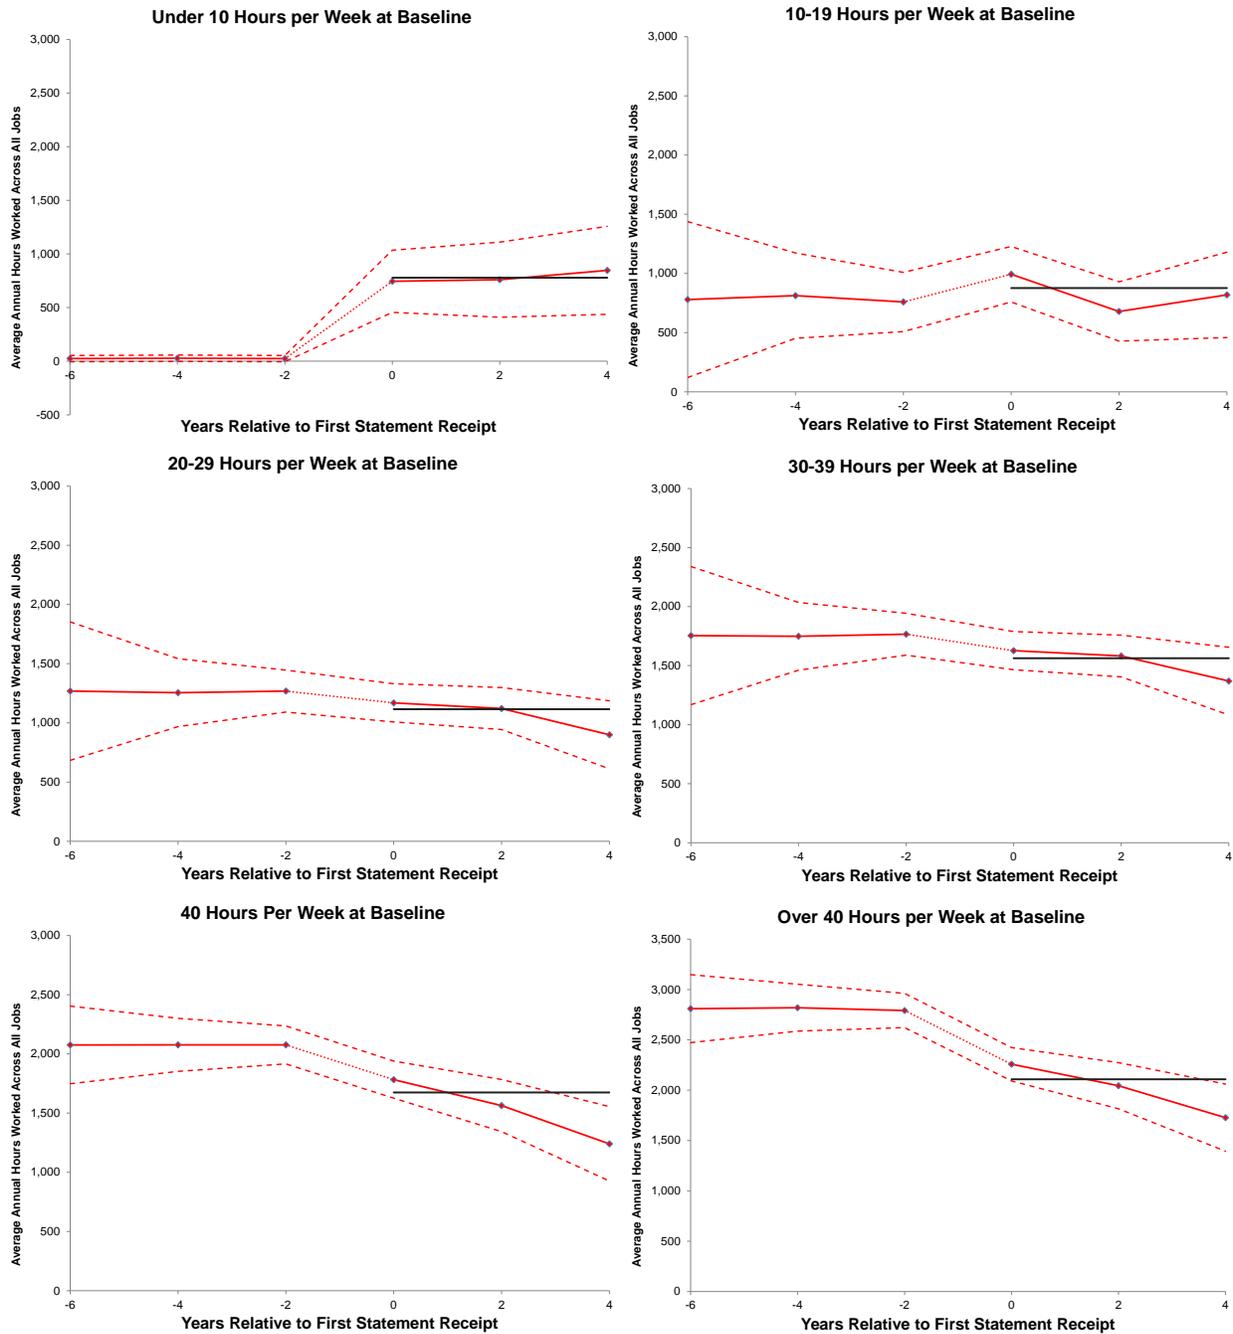
Notes: Each cell shows the age of those in a given birth year when they were mailed a Statement in a given SSA Fiscal Year. Blank space indicates no Statement was sent to that birth cohort in that year. SSA Fiscal Years are October of the preceding calendar year to September of the stated year. No Statements were automatically sent out before Fiscal Year 1995. Differences between fiscal years and calendar years lead to those of the same age being mailed a Statement in different fiscal years.

Figure 3. Trends in Hours Worked Across Cohorts in Years Relative to Statement Receipt



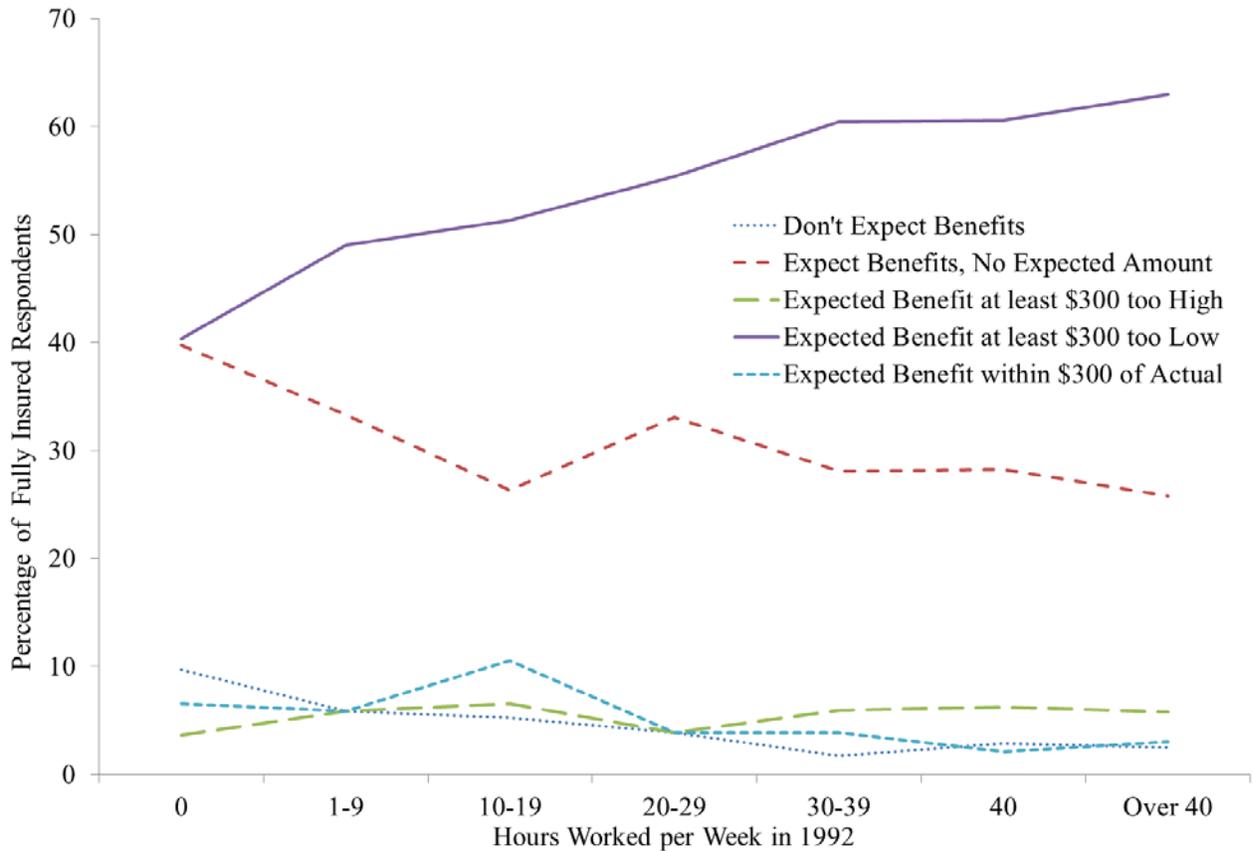
Notes: Unit of observation is person-wave. The figure depicts trends in self-reported hours worked across all jobs in years relative to first Statement Receipt. Respondents receive Statements between relative years -2 and 0. Hours worked are residual to marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. The sample includes men under age 62 who are fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). Dashed lines show 95% confidence interval bounds calculated from standard errors that are two-way clustered by birth year and calendar year.

Figure 4. Trends in Hours Worked Across Cohorts in Years Relative to Statement Receipt, by Pre-Statement Hours of Work



Notes: Unit of observation is person-wave. The figures depict trends in self-reported hours worked across all jobs in years relative to first Statement receipt. Respondents receive Statements between relative years -2 and 0. The sample includes men under age 62 who are fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). The black horizontal lines show the difference-in-difference estimates from column (6) of Table 4. Dashed lines show 95% confidence interval bounds.

Figure 5. PIA Expectations as a Function of Pre-Statement Hours of Work



Notes: Unit of observation is a person included in the HRS in 1992. Categories are mutually exclusive and exhaustive: individuals who expect benefits are then prompted as to what their expected benefits would be at their expected claiming age, to which they can respond “Don’t Know.” Categories are defined from responses to most recent wave before any Statement receipt. Expected benefits are recoded to monthly payment frequency. The sample includes men under age 62 who are fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992).

Table 1. Sample Restrictions

Sample Restriction	Person-year Observations	Unique Individuals
Total	332,721	27,224
Between 1992 and 2002	268,102	
Matched to Earnings Records	237,928	
Under Age 62	112,785	
Male	46,288	
Qualified for OAI	43,260	
In the HRS Before First Statement Receipt	21,904	4,038

Source: Health and Retirement Study respondents from the 1932-1947 birth cohorts.

Table 2. Descriptive Statistics of Analysis Variables

	Mean	SD
<u>Dependent Variables:</u>		
Annual Hours Worked	1065.28	1,252.12
Self-Employment Annual Hours Worked	227.79	764.36
Self-Reported Annual Earnings	34,836.79	58,514.62
Decreased Expected PIA	0.29	0.45
<u>Independent Variables:</u>		
Age	52.4	6.56
White	0.803	0.398
Married	0.491	0.500
High School	0.493	0.500
College	0.137	0.343
More than College	0.133	0.340
Total Wealth	292,313	1,096,645
Wealth + Current Pension Value	515,538	1,194,929
Wealth + Projected Pension Value	564,968	1,213,758
OAI Wealth Accrual Rate	0.01	0.01
OAI Wealth	7,730.06	3,232.21
Ever Received Statement	0.389	0.488
Ever Received 2 Statements	0.176	0.381
Had Second Job Before 1st Statement	0.104	0.305

Note: Unweighted descriptive statistics from 1992-2002 Health and Retirement Study matched to SSA administrative records. Sample restrictions are outlined in Table 1.

Table 3. Descriptive Statistics of Pre-Statement Hours Worked Categories

Hours Category (per week)	Unique Individuals	Person- Years	Pre-Statement Annual Hours	
			Mean	Median
0	1268	6251	0	0
1-9	54	277	261.2	260
10-19	87	387	776.4	780
20-29	127	620	1266.5	1280
30-39	209	1117	1770.8	1820
40	921	5075	2076.3	2080
Over 40	1372	7367	2759.9	2600

Note: Unweighted descriptive statistics from 1992-2002 Health and Retirement Study matched to SSA administrative records. Sample restrictions are outlined in Table 1. The 40 hours per week category includes those working between 50 and 52 weeks, inclusive, per year.

Table 4. Effect of Statement Receipt on Annual Hours Worked, 1992-2002

Independent Variable	Dependent Variable: Annual Hours Worked					
	(1)	(2)	(3)	(4)	(5)	(6)
Ever Received Statement	-118.9 ^{***} (45.1)	-135.5 ^{**} (39.7)		54.1 (40.9)	-79.4 [*] (45.9)	474.1 ^{***} (42.9)
1-Wave Lead of Statement		0.32 (17.41)				
Age Category X Statement						
40-44			87.3 (117.4)			
45-49			34.1 (93.9)			
50-54			15.3 (80.1)			
55-59			-125.3 ^{***} (40.4)			
60-61			-272.7 ^{***} (44.6)			
High School X Statement				-95.2 ^{**} (37.6)		
College X Statement				-27.7 (81.3)		
More than College X Statement				-161.1 ^{**} (70.8)		
2 nd Job Before 1 st Statement* Statement					-469.6 ^{***} (102.6)	
Pre-Statement Hours Worked Category X Statement						
1-9						110.3 (141.6)
10-19						-356.2 ^{***} (111.8)
20-29						-550.7 ^{***} (70.3)
30-39						-609.8 ^{***} (70.3)
40						-876.0 ^{***} (67.3)
>40						-1157.0 ^{***} (72.8)
R ²	0.478	0.478	0.478	0.478	0.463	0.754

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Whether a respondent had a 2nd job prior to 1st Statement receipt is controlled for in column (5). Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table 5. The Effect of Statement Receipt on Annual Hours Worked, Including Women

		Dependent Variable: Annual Hours Worked			
Independent Variable	(1)	(2)	(3)		
Ever Received Own Statement	436.4 ^{***} (65.3)	54.3 ^{**} (39.7)	50.7 ^{**} (24.4)		
Ever Received Husband's Statement			26.0 ^{**} (12.7)		
Ever Received Both Statements			-24.6 [*] (13.0)		
Pre-Statement Hours Worked Category Interacted With:					
Pre-Statement Hours Worked Category	Own Statement	Own Statement	Own Statement	Husband's Statement	Both Statements
1-9	84.2 (150.3)	15.8 (52.5)	24.4 (47.4)	64.6 (70.8)	-40.4 (24.3)
10-19	-361.8 ^{***} (141.1)	-25.3 (116.2)	-28.5 (78.2)	47.1 (84.1)	12.4 (73.8)
20-29	-403.2 ^{***} (54.9)	-56.6 (131.8)	-50.7 (166.8)	-17.0 (122.5)	16.0 (128.6)
30-39	-545.9 ^{***} (65.0)	-71.3 (156.7)	-65.9 (104.9)	-35.3 (139.2)	26.6 (149.8)
40	-661.9 ^{***} (76.4)	-75.4 ^{**} (37.1)	-80.6 ^{**} (36.1)	-44.0 ^{**} (21.2)	38.5 [*] (23.1)
>40	-942.9 ^{***} (99.4)	-107.0 ^{**} (51.4)	-103.2 ^{**} (46.8)	-67.6 ^{**} (29.5)	56.2 [*] (29.2)
Sample	Men and Women; Unmarried or with Own PIA \geq 0.5 Spouse's PIA	Women; Unmarried or with Own PIA \geq 0.5 Spouse's PIA	Women; Unmarried or with Own PIA \geq 0.5 Spouse's PIA		
N	26,856	4,952	4,952		
R ²	0.675	0.522	0.527		

Notes: Unit of observation is person-wave. Each numbered column is a separate regression. The sample is limited to unmarried, OAI-eligible individuals or married individuals with potential PIA at least 50% of spousal PIA. Dependent variable is self-reported hours worked across all jobs in the reference year. Regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table 6. Effect of Statement Receipt on Annual Hours Worked – Robustness Checks

Dependent Variable: Annual Hours Worked							
Independent Variable	Baseline	Excluding Government Workers	Linear Pre-Statement Hours Trends	Cohort x Year Fixed Effects	Month x Year Fixed Effects	Initially Ineligible for OAI	Extensive Margin
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Ever Received Statement	474.1*** (42.9)	472.3*** (41.1)	392.8*** (70.7)	484.8*** (52.4)	448.7*** (51.0)	63.6 (259.7)	0.507*** (0.030)
Pre-Statement Hours Worked Category X Statement							
1-9	110.3 (141.6)	132.1 (124.0)	196.8 (111.9)	162.5 (129.6)	108.2 (124.7)	161.4 (198.2)	-0.393*** (0.065)
10-19	-356.2*** (111.8)	-339.5*** (125.4)	-274.8*** (180.5)	-341.0*** (127.5)	-353.4*** (120.8)	-98.2 (589.8)	-0.469*** (0.090)
20-29	-550.7*** (70.3)	-534.5*** (102.8)	-458.3*** (144.6)	-585.6*** (101.5)	-506.0*** (104.3)	-52.4 (315.2)	-0.449*** (0.037)
30-39	-609.8*** (70.3)	-649.2*** (109.0)	-425.4*** (195.5)	-689.5*** (204.5)	-516.0*** (165.9)	-132.3 (412.3)	-0.455*** (0.046)
40	-876.0*** (67.3)	-899.7*** (54.2)	-765.9*** (76.4)	-821.4*** (50.1)	-934.6*** (45.9)	-101.1 (299.6)	-0.550*** (0.032)
>40	-1157.0*** (72.8)	-1155.2*** (63.0)	-993.5*** (92.4)	-1124.7*** (56.5)	-1191.2*** (60.7)	-131.8 (245.1)	-0.578*** (0.034)
Observations	21,904	20,182	21,904	21,904	21,904	862	21,904
R ²	0.754	0.751	0.755	0.765	0.764	0.693	0.526

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year in columns (1)-(6) and is an indicator for any working hours in the year in column (7). The sample is limited to under age 62 men; outside of column (6) the sample consists of those who are fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table 7. Randomization Tests

Sample	Percentage of Random Draws with Smaller Estimates	
	Randomized Cohort Treatment Assignment (1)	Randomized Individual Treatment Assignment (2)
Full Sample	0.002	0.001
Pre-Statement Hours Worked Category X Statement		
0-9	0.991	1.000
10-19	0.371	0.285
20-29	0.526	0.526
30-39	0.179	0.001
40	0.018	0.000
>40	0.001	0.000

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All randomization tests come from 10,000 replications done with replacement. In column (1), we randomly assign cohorts to a treatment year and in column (2) we randomly assign individuals to a treatment year. All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Pre-Statement hours of work are calculated using labor supply in the survey wave prior to the randomly assigned treatment year.

Table 8. Effect of Statement Receipt on Annual Hours Worked, by Wealth (1992-2002)

Independent Variable	Dependent Variable: Annual Hours Worked				
	(1)	(2)	(3)	(4)	(5)
Ever Received Statement	-118.9*** (45.3)	-31.05 (52.36)	-168.7*** (48.8)	-164.4*** (49.2)	129.7*** (58.6)
Total Wealth (\$1,000)		0.055* (0.026)			
Total Wealth (\$1,000) X Statement		-0.057** (0.025)			
Wealth + Current Pension Value			3.19e-05 (2.74e-05)		
Wealth + Current Pension Value X Statement			-5.56e-05 (4.17e-05)		
Wealth + Projected Pension Value				3.24e-05 (3.03e-05)	
Wealth + Projected Pension Value X Statement				-5.87e-05 (4.27e-05)	
OAI Wealth Accrual Rate					3009** (1500)
OAI Wealth Accrual Rate X Statement					8325** (1230)
OAI Wealth					0.018 (0.014)
OAI Wealth X Statement					-0.037*** (0.010)
Observations	21904	19103	4771	4771	20926
R ²	0.478	0.478	0.478	0.478	0.463

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992) for whom the given wealth measure is available. All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table 9. The Effect of Statement Receipt by Accuracy of PIA Prediction

Expectation Category X Statement	
Expect Benefits, but No Amount	84.2* (48.0)
Expected Benefit at least \$300 too High	114.6** (57.8)
Expected Benefit at least \$300 too Low	-166.5*** (46.8)
Expected Benefit within \$300 of Actual	57.3 (129.5)

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, uninteracted expectation category, and age and year fixed effects. The “Expected Benefit at least \$300 too Low” category includes those who provide an expected benefit as well as those who do not expect any benefits.

Table 10. Percent Change in Projected PIA at Full Retirement Age after First Statement Receipt, by Simulated Change in Earnings

Permanent Labor Force Exit	% Change in Projected PIA
Max	-78.2%
95 th Percentile	-44.8%
Median	-5.9%
5 th Percentile	0.0%
Min	0.0%
Mean	-12.8%
Permanent 50% Decrease in Earnings	% Change in Projected PIA
Max	-38.1%
95 th Percentile	-22.5%
Median	-4.1%
5 th Percentile	0.0%
Min	0.0%
Mean	-6.9%

Note: Change in projected PIA using Social Security Statement methodology by simulated change in earnings at first Statement receipt among the HRS sample outlined in Table 1.

Table 11. The Effect of First and Second Statement Receipt on Labor Supply

Waves Since First Receipt	Waves Since First Receipt X Pre-Statement Hours Category					Waves Since First Receipt X Pre-Statement Hours Category X Second Statement Receipt				
	0-9	10-29	30-39	40	>40	0-9	10-29	30-39	40	>40
-2		-50.4 (36.1)	25.1 (45.9)	17.7 (36.3)	15.4 (55.4)					
-1	69.1 (71.5)	-27.7 (42.5)	69.9 (91.3)	59.4 (65.4)	26.4 (71.6)					
0	428.6*** (79.6)	94.6 (156.5)	-293.6* (158.6)	-415.7*** (137.2)	-562.0*** (183.8)					
1	442.2*** (77.8)	24.8 (154.5)	-199.4* (107.7)	-473.8*** (153.9)	-686.3*** (212.5)	-58.2 (102.3)	3.4 (269.0)	51.8 (316.9)	144.3*** (56.3)	304.1** (131.9)
2	428.8*** (128.6)	-307.6* (167.0)	-398.2** (168.8)	-537.5*** (150.0)	-928.1*** (157.6)	-110.2 (173.6)	162.7 (184.0)	252.1 (301.6)	348.7*** (131.6)	495.1*** (119.4)

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all jobs in the reference year. All estimates come from a single regression that includes reporting wave-by-pre-Statement-hours category fixed effects as well as controls for marital status, education, race, and age. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table 12. Effect of Statement Receipt on Decrease in Expected PIA, 1992-2002

Independent Variable	Decreased Expected PIA from w-1 to w				Increased Expected PIA from w-1 to w			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Decreased Earnings from w-1 to w	0.017 (0.013)	0.040 ^{***} (0.017)	-0.002 (0.023)	0.061 ^{**} (0.027)				
Ever Received Statement		0.049 (0.035)		0.089 [*] (0.050)				
Decreased earnings X Received Statement		-0.027 (0.024)		-0.077 ^{***} (0.028)				
Increased Earnings from w-1 to w					0.059 ^{***} (0.023)	0.063 ^{***} (0.020)	0.076 ^{***} (0.026)	0.096 ^{***} (0.040)
Ever Received Statement						0.0001 (0.0450)		0.026 (0.047)
Increased Earnings X Received Statement						-0.004 (0.028)		-0.024 (0.041)
OAI Wealth Accrual Rate at least 1%	No	No	Yes	Yes	No	No	Yes	Yes
Observations	2046	2046	1171	1171	2046	2046	1171	1171
R ²	0.014	0.015	0.019	0.020	0.027	0.027	0.041	0.041

Notes: Unit of observation is person-wave. Dependent variable is whether individual decreased (increased) expected PIA from last interview to current interview. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992) who had a predicted PIA last interview and current interview and did not change the expected retirement date between the two interviews. PIA Accrual rate calculated as percentage change between PIA based on working one more year at current earnings level and PIA based on only earnings up until current year. All regressions are unweighted and control for marital status, education, race, and age and year fixed effects. Standard errors in parentheses are two-way clustered at year and birth-year levels: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

**Online Appendix
Not for Publication**

Table A-1. Effect of Statement Receipt on Annual Self-Employment Hours Worked, 1992-2002

Independent Variable	Dependent Variable: Annual Hours Worked in Self-Employed Jobs				
	(1)	(2)	(3)	(4)	(5)
Ever Received Statement	12.74 (14.01)		38.47** (19.28)	52.15*** (17.28)	121.40*** (16.58)
Age Category X Statement					
40-44		48.16** (24.40)			
45-49		100.00** (47.38)			
50-54		19.69 (40.91)			
55-59		0.02 (17.56)			
60-61		-13.34 (15.69)			
High School X Statement			-45.58** (17.65)		
College X Statement			3.91 (39.02)		
More than College X Statement			-40.66** (16.58)		
2 nd Job Before 1 st Statement* Statement				-108.90*** (73.55)	
Pre-Statement Hours Worked Category X Statement					
1-9				20.62 (129.50)	
10-19				-215.2 (169.6)	
20-29				-173.3 (124.7)	
30-39				-195.6*** (90.5)	
40				-798.0*** (161.3)	
>40				-1115.0*** (99.1)	
R ²	0.438	0.453	0.439	0.450	0.849

Notes: Unit of observation is person-wave. Dependent variable is self-reported hours worked across all self-employment jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Whether a respondent had a 2nd job prior to 1st Statement receipt is controlled for in column (4). Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table A-2. Effect of Statement Receipt on Annual Self-Reported Earnings, 1992-2002

Independent Variable	Dependent Variable: Yearly Self-Reported Earnings			
	(1)	(2)	(3)	(4)
Ever Received Statement	-599 (869)		627 (1001)	-746 (1235)
Age Category X Statement				
40-44		12130*** (4374)		
45-49		4781 (4331)		
50-54		-512 (1248)		
55-59		-2036 (1399)		
60-61		-974 (2391)		
High School X Statement			-1218 (926)	
College X Statement			-3228 (5065)	
More than College X Statement			-2591 (4523)	
2 nd Job Before 1 st Statement* Statement				8895* (5281)
R ²	0.088	0.087	0.089	0.123

Notes: Unit of observation is person-wave; N=19103. Dependent variable is self-reported earnings across all jobs in the reference year, price adjusted to 2000 dollars with the CPI-U-RS. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are unweighted and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Whether a respondent had a 2nd job prior to 1st Statement receipt is controlled for in column (4). Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table A-3. Descriptive Statistics of Analysis Variables, Weighted

	Mean	SD
<u>Dependent Variables:</u>		
Annual Hours Worked	1760.91	1,179.49
Self-Employment Annual Hours Worked	419.31	1,010.79
Self-Reported Annual Earnings	37,354.99	69,343.22
Administrative Earnings	28,021.21	24,754.33
Decreased Expected PIA	0.28	0.45
<u>Independent Variables:</u>		
Age	57.2	2.90
White	0.908	0.289
Married	0.848	0.359
High School	0.505	0.500
College	0.123	0.329
More than College	0.121	0.326
Ever Received Statement	0.390	0.488
Ever Received 2 Statements	0.135	0.342
Had Second Job Before 1st Statement	0.122	0.348

Hours Worked Category Pre-Statement among those with Positive Hours Worked Pre-Statement

	All Hours Worked	Self-Employed Hours Worked
1-9	12.95%	80.91%
10-19	1.92%	0.96%
20-29	3.95%	1.78%
30-39	7.11%	2.35%
40	27.01%	2.08%
>40	47.07%	11.91%

Note: Weighted descriptive statistics from 1992-2002 Health and Retirement Study matched to SSA administrative records. Sample restrictions are outlined in Table 1. Weights are the 1994 HRS person-level weights.

Table A-4. Effect of Statement Receipt on Annual Hours Worked, Weighted (1992-2002)

Independent Variable	Dependent Variable: Annual Hours Worked					
	(1)	(2)	(3)	(4)	(5)	(6)
Ever Received Statement	-241.6*** (40.5)	-282.2** (14.9)		-144.1*** (45.2)	-207.7*** (47.1)	433.4*** (69.7)
Leading Statement Variable		-7.03 (16.13)				
Age Category X Statement						
50-54			-144.5* (75.2)			
55-59			-204.4*** (34.8)			
60-61			-389.8*** (55.2)			
High School X Statement				-143.1** (54.3)		
College X Statement				-109.8 (127.8)		
More than College X Statement				-139.8 (98.3)		
2 nd Job Before 1 st Statement* Statement					-319.7*** (94.6)	
Pre-Statement Hours Worked Category X Statement						
1-9						92.3 (121.9)
10-19						-405.8** (194.4)
20-29						374.7*** (65.4)
30-39						-500.5*** (102.2)
40						-833.1*** (98.2)
>40						-1098.0*** (97.2)
R ²	0.703	0.704	0.715	0.703	0.703	0.754

Notes: Unit of observation is person-wave; N=16810. Dependent variable is self-reported hours worked across all jobs in the reference year. Sample limited to under age 62 men fully insured for Old Age Insurance in 1991 (those with at least 40 Quarters of Coverage by 1992). All regressions are weighted by HRS sample weights and control for marital status, education, race, pre-Statement hours of work category, and age and year fixed effects. Whether a respondent had a 2nd job prior to 1st Statement receipt is controlled for in column (5). Standard errors two-way clustered on year and birth-year are shown in parentheses: * indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level.

Table A-5. Descriptive Statistics of Pre-Statement Hours Worked Categories – Women

Hours Category (per week)	Unique Individuals	Percent
0	580	43.6%
1-9	25	1.9%
10-19	44	3.3%
20-29	75	5.6%
30-39	160	12.0%
40	315	23.6%
Over 40	133	10.0%

Note: Unweighted descriptive statistics from 1992-2002 Health and Retirement Study matched to SSA administrative records. The 40 hours per week category includes those working between 50 and 52 weeks, inclusive, per year.