NBER WORKING PAPER SERIES

THE WELFARE COST OF PERCEIVED POLICY UNCERTAINTY: EVIDENCE FROM SOCIAL SECURITY

Erzo F.P. Luttmer
Andrew A. Samwick

Working Paper 21818
http://www.nber.org/papers/w21818

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2015

We thank Alan Gustman, Jonathan Skinner, Steve Venti, and Niels Vermeer for helpful comments and Poom Nukulkij and Wan Yan of Knowledge Networks for their work on our survey. We also thank seminar audiences at the Western Economic Association International, NBER Summer Institute, University of Wisconsin, CPB Netherlands Bureau for Economic Policy Analysis, Tinbergen Institute, Groningen University, Erasmus University, Tilburg University, University of Oxford, IZA Bonn, LSE/IFS joint seminar, the Paris School of Economics, the University of Maastricht, the University of Connecticut, the 13th International Pensions Workshop, Texas A&M University, University of Bergen, Loyola Marymount University, the University of Utah, and Middlebury College for helpful comments. We are grateful to Mathew Greenwald for providing us with tabulations of data from his research, to Steve Goss for guidance on the present discounted value of accrued future Social Security benefits, and to Ben Chuchla for exceptional research assistance. This research was supported by the U.S. Social Security Administration through grant #5 RRC08098400-03-00 to the National Bureau of Economic Research as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the author(s) and do not represent the views of SSA, any agency of the Federal Government, or the National Bureau of Economic Research.

The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security
Erzo F.P. Luttmer and Andrew A. Samwick
NBER Working Paper No. 21818
December 2015
JEL No. D89,H55

ABSTRACT

Policy uncertainty can reduce individual welfare when individuals have limited opportunities to mitigate or insure against consumption fluctuations induced by the policy uncertainty. For this reason, policy uncertainty surrounding future Social Security benefits may have important welfare costs. We field an original survey to measure the degree of policy uncertainty in Social Security and to estimate the impact of this uncertainty on individual welfare. On average, our survey respondents expect to receive only about 60 percent of the benefits they are supposed to get under current law. We document the wide variation around the expectation for most respondents and the heterogeneity in the perceived distributions of future benefits across respondents. This uncertainty has real costs. Our central estimates show that on average individuals would be willing to forego around 6 percent of the benefits they are supposed to get under current law to remove the policy uncertainty associated with their future benefits. This translates to a risk premium from policy uncertainty equal to 10 percent of expected benefits.

Erzo F.P. Luttmer
6106 Rockefeller Center, Room 305
Department of Economics
Dartmouth College
Hanover, NH 03755
and NBER
Erzo.FP.Luttmer@Dartmouth.Edu

Andrew A. Samwick
6106 Rockefeller Hall
Department of Economics
Dartmouth College
Hanover, NH 03755-3514
and NBER
andrew.samwick@dartmouth.edu

A data appendix is available at http://www.nber.org/data-appendix/w21818

1. Introduction

Relative to the extensive literature that values risk in insurance and financial markets, economists have paid surprisingly little attention to the welfare consequences of policy uncertainty. The welfare effects of policy uncertainty are likely to be especially pronounced when the policy has a potentially large impact on consumption and the risk associated with the policy is not diversifiable or insurable. For example, uncertainty about future taxes is costly to individuals because it hampers their ability to consumption smooth over the lifecycle and because their investments in human capital will not be privately optimal for the actual realization of the future tax rate.[1] Uncertainty about generosity of expenditure programs for the elderly, such as Social Security and Medicare, also reduces individuals' ability to consumption smooth. Given that Social Security is mandatory, non-diversifiable, and accounts for more than a third of total income among the elderly, we suspect that policy uncertainty regarding its generosity is likely to be one of the major sources of welfare cost of policy uncertainty more generally. This paper's objective, therefore, is to estimate the welfare cost to individuals of policy uncertainty regarding Social Security benefits.[2] In other words, we estimate the risk premium for policy uncertainty in Social Security wealth.

The traditional method of valuing uncertainty by comparing an asset's market value to its expected value is generally not feasible in the case of policy uncertainty. The effect of policy uncertainty is hardly ever fully captured by a publicly traded asset, and even if it were, other sources of uncertainty might also affect the asset's value.[3] To overcome this challenge, the empirical literature on policy uncertainty proceeds in two steps. The first step is to measure the

---

[1] As noted by Weiss (1976) and Stiglitz (1982) in the case of income taxes, policy uncertainty can induce behavioral changes that may counteract existing distortions. Such behavior changes thus yield a positive effect on welfare and this positive effect could potentially more than offset the negative welfare effect of the consumption risk induced by the policy uncertainty. Alm (1988) and Kim, Snow, and Warren (1995) provide further theoretical results regarding the welfare effects of tax policy uncertainty in a second-best world.

[2] Our empirical approach does not allow us to ascertain whether some perceived policy uncertainty is optimal from an intergenerational risk-sharing perspective (see, e.g., Gordon and Varian (1988)). To the extent some component of the perceived policy uncertainty is optimal for intergenerational risk sharing, our estimates of the welfare cost to current individuals of policy uncertainty are an overestimate of the total welfare effect of policy uncertainty. Similarly, we are not able to evaluate any welfare effects of policy uncertainty that stem from the uncertainty reducing existing distortions.

[3] Geanakoplos and Zeldes (2010) estimate the market value of accrued Social Security benefits by adjusting the actuarial value of accrued Social Security benefits for the uncertainty in Social Security benefits that stems from wage indexing. Hence, their paper uses an asset price model to estimate the market risk premium for the main non-policy related source of uncertainty in Social Security benefits whereas we use survey techniques to estimate the individual risk premium for policy uncertainty.

degree of policy uncertainty. This can be done retrospectively by measuring uncertainty as the residuals in a vector-autoregression model, as Skinner (1988) does, or by estimating the variability in past policy changes, which is the approach taken by McHale (2001), Nataraj and Shoven (2003), Shoven and Slavov (2006), Borgmann and Heidler (2007), Dušek (2007), and Blake (2008). Because past variability may not necessarily provide a good estimate of uncertainty about future policy (e.g., if the process is non-ergodic or there is a so-called peso-problem), other studies, including Van der Wiel (2008), Guiso, Jappelli, and Padula (2013), and Giavazzi and McMahon (2012) have measured perceived policy uncertainty using survey questions about future policy. Baker, Bloom, and Davis (2015) take yet another approach and create an index of policy uncertainty based on the frequency that the word triplet "uncertain", "economic," and "policy" (or variants/synonyms of these words) appears in newspaper articles.

Not all of these papers proceed to the second step, but those that do either relate the policy uncertainty estimated in the first step to observed behavior or use the estimated policy uncertainty to calibrate a model that yields a welfare estimate. Papers that relate estimated policy uncertainty to observed individual-level behavior include Giavazzi and McMahon (2012), who analyze its effects on household saving; Guiso, Jappelli, and Padula (2013), who study the effects on enrollment in private pensions and health insurance; and Van der Wiel (2008), who examines the effects on private pension participation. Baker, Bloom, and Davis (2015) relate their indices of policy uncertainty to industry and macro outcomes including stock-price volatility, investment, employment, and output.

These papers, however, do not estimate the welfare cost of the policy uncertainty. In contrast, Skinner (1988) and Dušek (2007) evaluate the estimated uncertainty using a model to calculate the welfare cost of the uncertainty. Skinner estimates that the welfare cost of uncertain taxes is 0.4% of national income, and Dušek finds that the risk premium for the uncertainty around the indexing of Social Security benefits in the Czech Republic is 1.3% when the coefficient of relative risk aversion is assumed to equal 3. Alternatively, it is possible to calculate the welfare cost of policy uncertainty using a calibrated (rather than estimated) measure of policy uncertainty, which is the approach taken by Gomes, Kotlikoff, and Viceira (2012). They focus on a slightly different question, namely the welfare gain from resolving uncertainty about future Social Security benefits earlier holding constant the variance in future Social Security benefits, and find that early resolution can lead to welfare gains that are equivalent to

0.5% of lifetime consumption. Caliendo, Gorry, and Slavov (2015) also use a calibrated measure of policy uncertainty but allow both uncertainty in the timing of the resolution of uncertainty and uncertainty in the structure of the Social Security reform. Their model shows that the welfare cost of Social Security policy uncertainty is just a few basis points of lifetime consumption for individuals who make optimal savings decisions, but that it can exceed 1 percent of lifetime consumption for those who do not save.

In this paper, we take an alternative and, to the best of our knowledge, novel approach to valuing the cost of policy uncertainty: we elicit both the expected policy and the certainty equivalent of uncertain future policy and use the difference between these two measures as the welfare cost to the individual of policy uncertainty. Our approach is thus similar to the asset-pricing approach of valuing uncertainty except that we elicit the certainty equivalent by asking individuals how they value a hypothetical asset that has no policy uncertainty rather than using a market price to observe this certainty equivalent. The chief concern about our approach is that some individuals may have trouble giving a meaningful valuation of a hypothetical asset. Because we believe this is an important concern, we include various forms of randomized variation in the way we elicit expectations and certainty equivalents, and the responses to this randomized variation allow us to evaluate the quality of the responses. The benefit of our approach is that our estimate of the cost of uncertainty does not rely on model specification, parameter assumptions, or estimates of the correlation between policy uncertainty and other sources of uncertainty that affect consumption. This means that our estimate does not rely on any assumptions on, or estimates of, the types of behaviors people may undertake to mitigate the policy risk. Moreover, our estimates capture any direct effects (such as disutility from stress or worrying) related to the policy uncertainty that might not be captured by a standard expected utility model.

We estimate the cost of policy uncertainty for Social Security benefits because this is one of the largest sources of policy uncertainty for individuals. To address the solvency of Social Security, some combination of benefit cuts and tax increases will likely occur at some point in the future.[4] The need for reform to restore the program to long-term financial stability has been an active topic of policy discussion since at least the report of the 1994-1996 Advisory Council

[4] In their most recent report, Social Security's Board of Trustees (2015) projected that the program's trust funds would be exhausted in 2035, at which point annual costs are projected to exceed annual income by 28 percent or 3.2 percentage points of taxable payroll.

(Advisory Council, 1997). Since then, each of the last three presidents has made the reform of Social Security an important part of his policy agenda.[5] With the status of reform still in doubt, individuals can expect some form of policy change but may be uncertain of its timing, size, and composition. To illustrate the role of this policy uncertainty, consider two scenarios in a stylized example. In the first, individuals know for sure that their Social Security benefit will be cut by 20 percent. In the second, they have a 20 percent chance that their benefits will be cut completely and an 80 percent chance that their benefits will not be cut at all. While the expected benefits (and thus the expected cost to the government) are the same in both scenarios, individuals only face policy uncertainty in the second scenario. Because of the uncertainty in the second scenario, risk averse individuals value their benefits less than what they cost in expectation. In particular, they would likely be willing to trade the second scenario for a sure benefit cut, even if that sure benefit cut is somewhat greater than 20 percent. The difference between the expected benefit cut and the largest sure benefit cut people would be willing to accept is an estimate of the cost to individuals of policy uncertainty surrounding Social Security benefits.

We implement our methodology by fielding an original, internet-based survey of 3,000 individuals between the ages 25 and 59 who are broadly representative of the U.S. population in that age range. We focus on this age range because this is the prime age range in which individuals need to prepare for retirement and because older individuals will likely be grandfathered into the existing rules if there is a major Social Security reform. An important innovation relative to the literature that examines perceptions of future Social Security benefits is that we ask about future benefits *relative* to the benefits scheduled under current law.[6] This allows us to filter out any uncertainty (or misperceptions) regarding the current benefit rules as well as uncertainty about benefits that is related to uncertain inputs (such as own future earnings or aggregate future wage growth) to the benefit formula. The key part of the survey consists of two sets of questions about these benefits. In the first, respondents are asked to describe the

[5] The Social Security Administration keeps an archive of presidential statements on Social Security at http://www.ssa.gov/history/presstmts.html. President Bush spent much of 2005 advocating for reform, and the need for reform figured prominently in President Obama's call for a bipartisan fiscal commission in 2010 and negotiations over the debt ceiling increase in the summer of 2011.

[6] There is an extensive literature examining perceptions of expected Social Security benefits. An early example focusing on the relationship between Social Security expectations and private saving is Bernheim and Levin (1989). More recent examples include Gustman and Steinmeier (2005), Dominitz and Manski (2006), Delavande and Rohwedder (2008), and Liebman and Luttmer (2012).

likelihood of receiving benefits in specific ranges relative to "the benefits they are supposed to get under current law." They fill in a histogram of this distribution by putting balls into bins on their computer screens. This histogram allows us to calculate their expected benefits. In the second part, respondents are asked to make a sequence of choices as to whether they would prefer a guaranteed contract at a hypothetical percentage of the benefits they are supposed to get under current law to the distribution of benefits they think they will get. This sequence of questions allows us to bracket their certainty equivalent benefit level. Subtracting the certainty equivalent from the expected benefits yields the respondent's risk premium against policy uncertainty.

Our main results indicate that individuals perceive the risk to which policy uncertainty exposes them and that the welfare cost of that risk is statistically and economically significant. Across respondents, the average expected benefits are 59.4 percent of the benefits the respondents are supposed to get under current law and the average standard deviation is 22.5 percent. The average certainty equivalent is 53.7 percent, yielding an average risk premium of 5.8 percent. At 7.0 percent, the median risk premium is close to the average risk premium. These risk premia are expressed as percent of benefits under current law, but would become 9.7 percent and 11.8 percent, respectively, if expressed as a percent of expected benefits.

Regression results show that the risk premium increases with age and decreases with income. Expected benefits as a fraction of benefits under current law rise with age and the standard deviation of benefits decreases with age. This implies that the increase in the risk premium with age is driven by the fact that it is costlier for older people to bear policy risk in Social Security, for example, because they have fewer means to mitigate this uncertainty by changing their labor supply or savings rate.

Because we recognize that some of the questions may be challenging for a broadly representative subject pool, we build into the survey randomizations that can alert us to respondents giving non-meaningful answers. One of the key randomizations that we insert is the starting value to the series of questions that brackets the value of the certainty equivalent. This starting value should not affect the final valuation of the certainty equivalent for a respondent who can report a stable underlying valuation of the certainty equivalent. We find that the starting value has a moderate, but statistically significant, effect on the reported certainty equivalent. The randomization of the starting value enables us to correct the estimated certainty equivalent for the effect of the starting value since the underlying distribution of certainty equivalents is invariant to the starting value. We obtain an average risk premium of 5.1 percent based on this corrected value of the certainty equivalent. We also examine how reported risk premia vary with indicators of response quality based on other questions asked in the survey (e.g., respondents should not give a lower probability of a policy change by a certain date if the date is further in the future). If we further adjust the risk premia for these indicators of response quality, we obtain an average risk premium of 6.7 percent.

As an additional check, we calculate the risk premium using the methodology that the existing literature has taken, namely applying a model and an assumed coefficient of risk aversion to our estimates of the degree of policy uncertainty as given by the reported histogram of future benefits. The resulting "simulated" risk premium has a median of 4.0 percent and an average of 9.4 percent if we assume a coefficient of relative risk aversion of 3. The simulated risk premium is based on an admittedly very simple model and sensitive to various assumptions including the value of the coefficient of relative risk aversion and the absence of a correlation between policy uncertainty in Social Security and other sources of uncertainty affecting consumption. We nevertheless find it reassuring that the resulting estimate is broadly similar to our main estimate of the risk premium of policy uncertainty in Social Security benefits.

The remainder of the paper is organized as follows. In section 2, we describe our sampling frame and survey instrument and provide summary statistics for the demographic and other control variables used in our analysis. In section 3, we discuss the particular design features of the survey that enable us to elicit information on the distribution of future benefits and its certainty equivalent. We present our main results and sensitivity tests in Section 4. Section 5 provides evidence on the validity of survey responses to questions about benefit distributions. Section 6 considers possible adjustments that could be made to the distribution of risk premia. Section 7 concludes.

2. Data

Our survey is conducted as a module of the KnowledgePanel, created by the survey firm Knowledge Networks. The KnowledgePanel is an address-based sample drawn from the U.S.

Postal Service's Delivery Sequence File.[7] When households without Internet access are recruited, they are provided with a laptop computer and free Internet service so they may participate in the panel. The KnowledgePanel consists of about 50,000 participants over the age of 18 and includes persons living in cell phone only households. Knowledge Networks collects basic demographic characteristics for all its panelists, and its panelists are roughly representative of the adult U.S. population according to these characteristics. Active members of the panel are invited to take specific surveys, with subsamples drawn using probability weighted sampling methods. The burden of panel membership is kept low by having members selected for no more than one survey per week.

We contracted with Knowledge Networks to obtain survey responses from approximately 3,000 KnowledgePanel participants who were between the ages of 25 and 59 in June 2011. Our sample contains the results for 3,053 completed interviews conducted between June 10 and July 1, 2011. The median duration of the interview was 20 minutes, and we paid respondents a $5 cash-equivalent incentive to enhance survey completion. Table 1 contains summary statistics for the demographic and other control variables that we use in our empirical analysis.[8] Online Appendix Table A1 compares summary statistics for these and other demographic variables to the Current Population Survey from March 2010. While for many demographic characteristics we can reject the hypothesis that the mean is the same in the CPS and our sample, the differences are limited in terms of economic magnitude. We therefore consider our sample as broadly representative of the U.S. population between the ages of 25 and 59.

In our regressions, we control for these demographic characteristics, along with MSA residency, homeownership, and employment status, as shown in Table 1. In some specifications, we also include a set of additional control variables that are relevant to perceptions of policy uncertainty in general and the Social Security program in particular.[9] We ask about risk

[7] As discussed in Knowledge Networks (2010), randomly sampled addresses are invited to join the KnowledgePanel through a series of mailings (English and Spanish materials) and by telephone follow-up to non-responders when a telephone number can be matched to the sampled address. Invited households can join the panel by one of several means: completing and mailing back an acceptance form in a postage-paid envelope; calling a toll-free hotline staffed by bilingual recruitment agents; or going to a dedicated Knowledge Networks recruitment Web site and completing the recruitment information online.

[8] We defer the discussion of the first four rows, which summarize the distribution of perceived Social Security benefits, until Section 4 below.

[9] We ask these questions at the end of the survey. The full survey instrument is included as Online Appendix B. As these control variables are not the focus of our analysis, we do not eliminate observations for which they are missing. Instead, we create a dummy variable for whether the response is missing, recode the missing values to zero,

preferences, life expectancy, the importance of Social Security in retirement, optimism, trust in the political system, and financial literacy. Summary statistics are presented in the last six rows of Table 1.

We measure risk preference through a sequence of questions in which respondents can choose a job that offers a certain lifetime income or a job that offers varying degrees of risk, such as a 50-50 chance of doubling lifetime income and a 50-50 chance of reducing it by some percentage. The sequence varies the reduction to bracket the respondent's point of indifference, from which we can infer risk aversion. In a constant relative risk aversion scenario, the brackets are coefficients of less than 0.5, 0.5 – 1, 1 – 2, 2 – 4, 4 – 8, and greater than 8. The median response is consistent with risk aversion of 4 – 8.

Two factors are very important to the role of Social Security in retirement. The first is how long the beneficiary will live. We ask respondents for a subjective probability of surviving to age 75. The mean probability is 67.9 percent and the median is 71 percent. The second is how important Social Security will be as a share of retirement income. We ask this question directly, with possible responses, coded 1 – 4, in the form of ranges of less than 25 percent of spending, 25 – 50 percent, 50 – 75 percent, and greater than 75 percent. There is considerable variation around a mean of 2.8 and a median of 3 (50 – 75 percent).

To measure optimism, we ask six questions about how the respondent perceives the outcomes of uncertain events (e.g., "In uncertain times, I usually expect the best."). The respondent can pick from five choices – strongly disagree, somewhat disagree, neither agree or disagree, somewhat agree, strongly agree – which are given numerical values of 1 – 5, with higher values indicating more optimism. We average the numerical responses across the six questions and standardize the variable to have zero mean and unit standard deviation.

Trust in the political system is measured as the response to the statement, "Most elected officials are trustworthy." As with the optimism question, the five choices range from strongly disagree to strongly agree, with numerical values ranging from 1 – 5. The average response is 2.2 and the median response is 2.0, indicating that most respondents lack trust in the political system.

---

and then include both the recoded variable and the dummy for whether the response was originally missing in our regressions.

Finally, we measure financial literacy as the number of correct answers given by the respondent to four simple questions about a lottery, money illusion, compound interest, and mutual funds. The average score is 2.4, with a median of 3.

3. Methodology

The main part of our survey is designed to gather information from the respondents sufficient to calculate the costs of policy uncertainty. As this is not an everyday topic of conversation for most people, the survey itself needs to guide them through the steps of the process. Moreover, we include randomizations in the survey that allow us to gauge whether respondents are able to give meaningful answers. This section discusses and illustrates three important design features of the survey.

3.1 Choice of Baseline Benefits

The first feature, which to the best of our knowledge has not been implemented before, is to use the respondent's own perception of current law benefits as the baseline. Throughout the survey, respondents are asked to compare expected or hypothetical benefits to "the benefits you are supposed to get under current law." We do not seek to measure whether the respondent has an accurate projection of what those current law benefits would be or whether the respondent is uncertain about benefits under current law because those benefits depend on variables that themselves are uncertain, such as future own earnings or future aggregate earnings. By keeping whatever uncertainty or misconceptions respondents may have about benefits under current law in the baseline, the survey responses will pertain only to the policy uncertainty regarding how current law benefits will be changed by policy makers.

3.2 Constructing the Perceived Distribution of Social Security Benefits

The second feature is to use the visual aspect of the online survey to facilitate the answer to the general question of how uncertain the respondent believes future Social Security benefits to be. This feature was developed in Delavande and Rohwedder (2008) and subsequently used in Liebman and Luttmer (2015). We measure uncertainty in the form of a histogram of where the respondents believe their benefits will be. This allows us to estimate the cumulative distribution function (CDF) of benefits for each respondent as a percent of what he or she is supposed to get under current law. The survey first asks the respondent to allocate 20 balls across four bins reflecting different benefit amounts, where each ball is explained to represent a 1 in 20 chance of that benefit amount occurring. One category is "no benefits whatsoever." The other three categories are lower, the same, and higher benefits relative to the benefits that the respondent is supposed to get under current law. An example of what the survey screen might look like when the respondent has allocated the 20 balls to the 4 bins is:

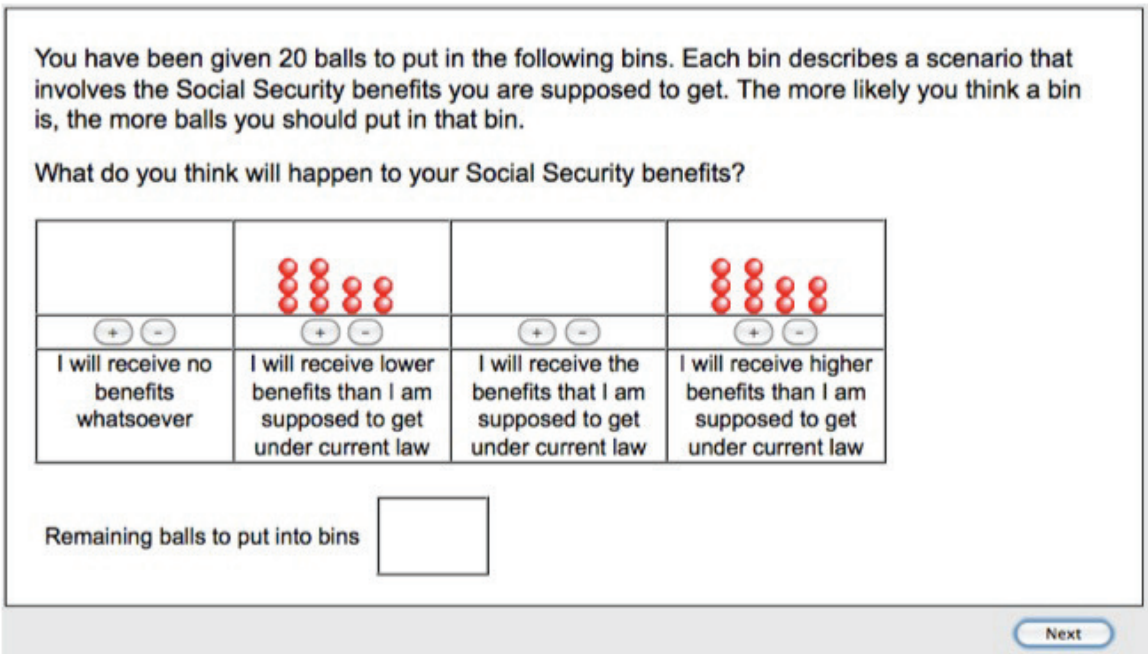

Respondents who put any of these balls in the "lower" or "higher" bins are then asked to further specify which 20-percentage-point bins between 1 and 99% or 101 and 200% should contain these balls. An example of the next screen this respondent will see is:

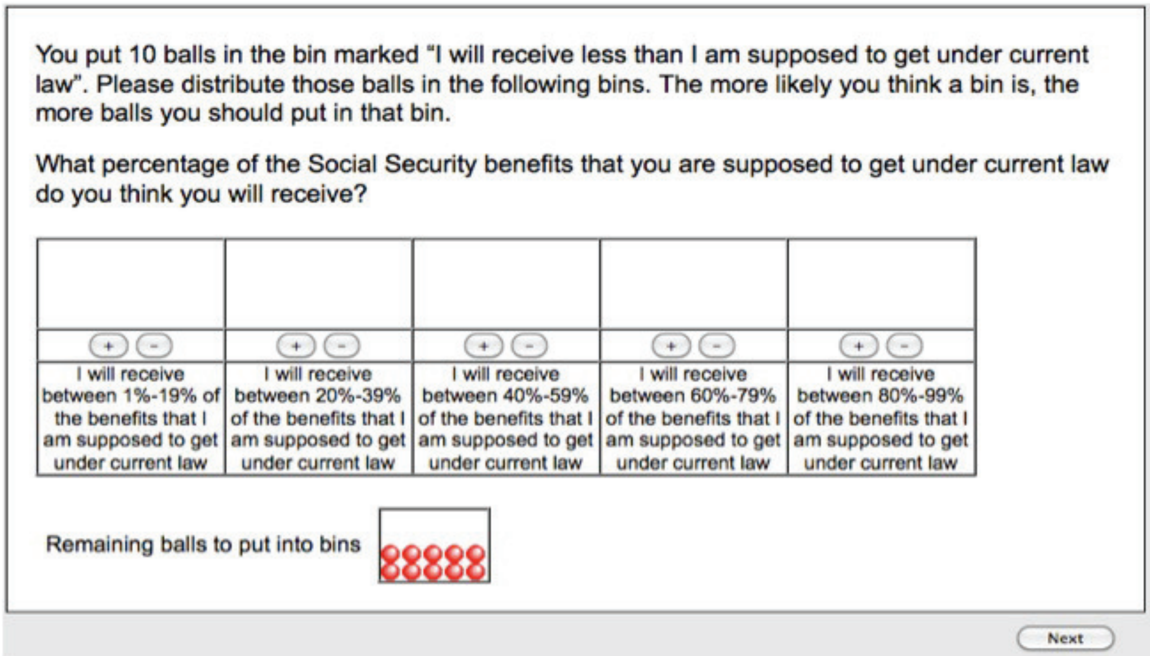

Finally, any bin into which 11 or more balls are placed is further broken down into five smaller bins, and respondents are asked to allocate the balls from the larger bin into the smaller bins. An example of the screen that the respondent would have seen in that case is:

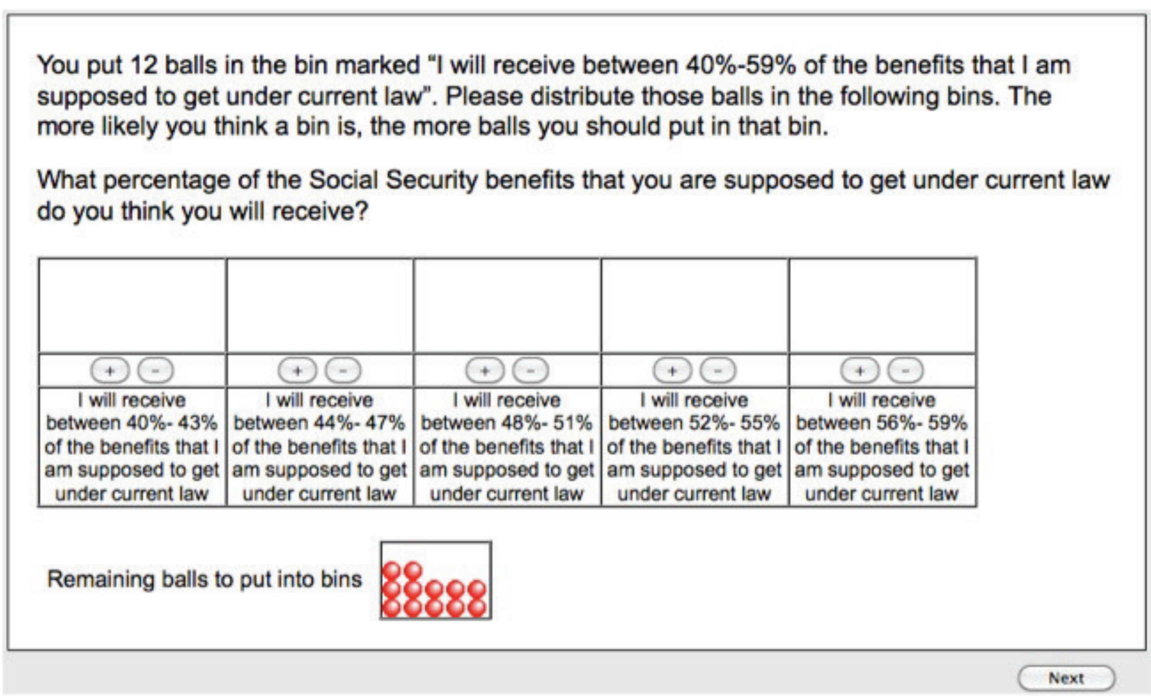

By this three-step process, we obtain the CDF of expected future benefits for each respondent. In order to have greater confidence that respondents will know how to use this tool to express their preferences, we first give an illustration using the weather in Boston.

Recognizing that the shape of the distribution that we show them to illustrate the method might influence the way they fill in the distribution of perceived benefits, we choose two different illustrations and assign them to respondents at random. For example, the wide distribution is:

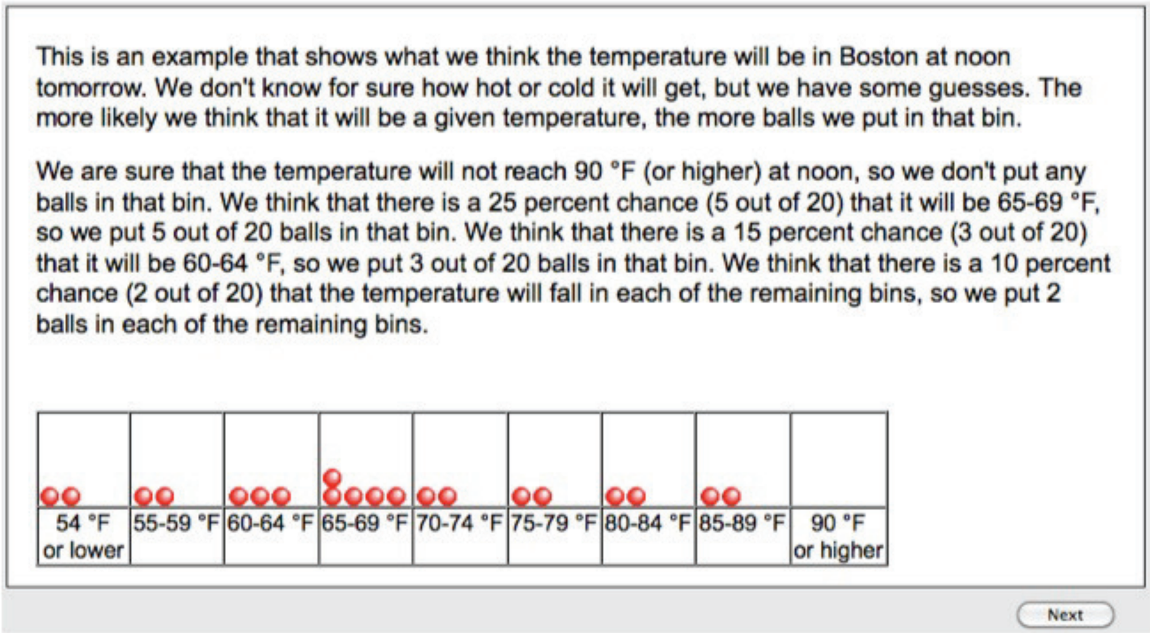

And the narrow distribution is:

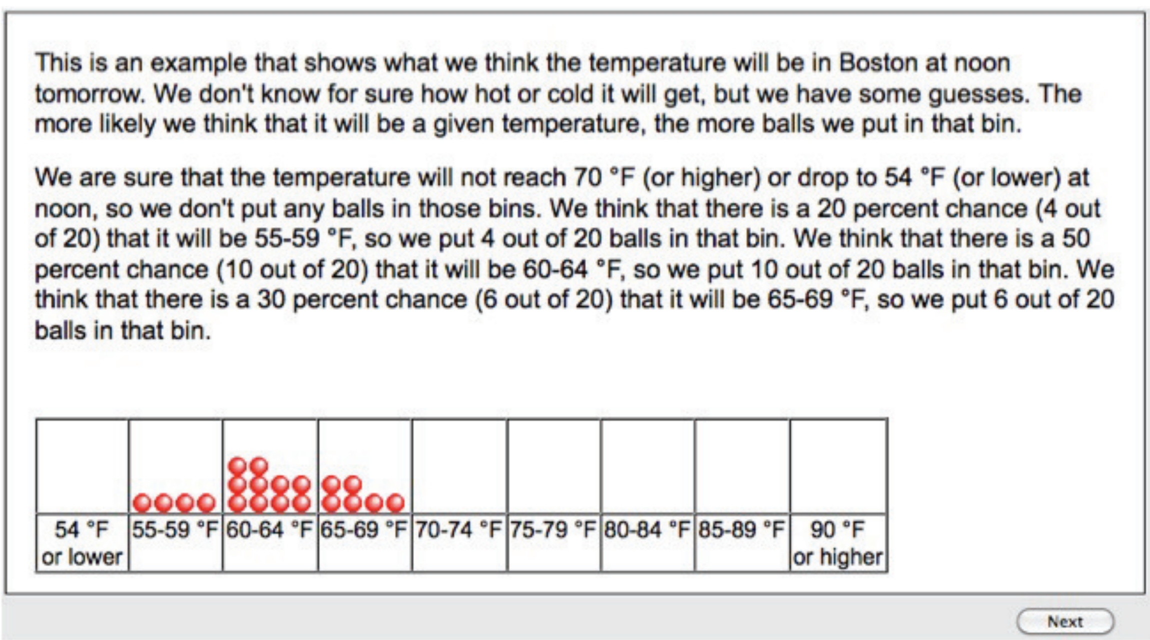

If we had shown no illustration, we could not be sure that respondents would understand the tool well enough to answer the subsequent question. If we had only shown one illustration, then we would have had no way to gauge the size of any bias that our particular choice of illustration may have had on the subsequent question. By choosing two illustrations, we can estimate the impact of the characteristics of the illustration – wide or narrow – on the responses to the subsequent question.

3.3 Obtaining the Certainty Equivalent Benefit

The natural metrics to quantify just how much the uncertainty in the perceived distribution of Social Security benefits matters to respondents are how much they would pay to insure themselves against it or at what discount they would be willing to sell their claim to future Social Security benefits. Even in a more straightforward context, respondents could be expected to have trouble coming up with a sensible answer if we asked for it directly. This concern led us to develop a third important feature of our survey, which is the sequence of binary choices that the survey presents to the respondent that allow us to bracket the respondent's certainty equivalent to the perceived distribution of benefits described in Section 3.2. The survey calculates the expected value (denoted below by the variable *X*) of the benefit distribution each respondent constructed by putting balls into bins and presents the respondent with the following choice:

> The way you put balls into various bins shows that you expect to receive [***X***]**%** of the Social Security benefits you are supposed to get under current law. It also shows that you could receive more or less than this **[*X*]%**. Let's call this distribution of possible benefits, as described by you using the bins and balls, your "uncertain benefits." So, your uncertain benefits are whatever level of benefits you get when you claim benefits.
>
> Imagine a contract that instead guarantees you a certain percentage of the Social Security benefits you are supposed to get under current law. This is like having all 20 balls on this certain percentage. This contract is unbreakable and cannot be changed by anybody, even the United States government.
>
> Would you rather have:
>
> (1) Guaranteed benefits equal to [***Y***]**%** of the Social Security benefits you are supposed to get under current law
>
> (2) Uncertain benefits around [***X***]**%** of the Social Security benefits you are supposed to get under current law

Respondents are prompted with a starting value of Y_1 equal to 30 or 70, chosen randomly, so that we can assess the impact of the starting value on the ultimate results. (Whether the guaranteed benefits are the first or second choice is also randomized, for the same reason.) A respondent who chooses the guaranteed (uncertain) benefits at a given Y_1 is then offered a lower (higher) value of Y_2 and asked the same question. The questioning continues, with the differences between Y_n and Y_{n+1} narrowing, until the respondent has answered that he would take the uncertain benefits if offered the lower of Y_n and Y_m, and the guaranteed benefits if offered the higher of them, where the interval between them is 5.[10]

One problem in generating the certainty equivalent using the question above is that 7.5% of respondents provide distributions that show no uncertainty. For these respondents, we ask a slightly different version of the question:

> Imagine that you were offered a contract that guaranteed you a certain percent of the Social Security benefits you are supposed to get under current law. This contract is unbreakable and cannot be changed by anybody, even the United States government.
>
> Would you rather have:
>
> (1) Benefits as determined by an unbreakable contract that offers you [***Y***]% of the Social Security benefits you are supposed to get under current law
>
> (2) Benefits as determined by Social Security when you claim benefits

The sequencing of the offers of *Y*% is the same as in the original question. This question simply makes no mention of a distribution that shows no uncertainty.

The answers to these questions provide us with upper and lower bounds on a certainty equivalent to the distribution of possible Social Security benefits. Subtracting this certainty equivalent from the distribution's expected value yields the risk premium that the respondent would pay to insure against the policy uncertainty in Social Security. In order to make more precise estimates of this risk premium, we ask a follow-up question of respondents whose range for the certainty equivalent is close to the expected value of their distribution of benefits. Specifically, a respondent whose upper bound for the certainty equivalent is within 5 percentage

[10] The full sequence of offers that the respondents receive is shown in Question 4.3 of the survey instrument in Online Appendix B.

points of the expected value will be asked the question again, with a value of *Y* close to *X* that will ensure that we can ascertain whether or not the risk premium exceeds two percent.

4. Results

4.1 General Expectations about Social Security

The survey begins by soliciting respondents' views on the financial condition of the Social Security program in order to get a qualitative understanding of their views about policy risk as well as the nature of the risk that they perceive. Table 2 aggregates the responses to these general questions. About 91 percent of respondents are aware that Social Security faces a projected financial shortfall. When asked how confident they are that Social Security will be able to provide them with the benefits they are supposed to get under current law, only 3.3 percent were very confident, with another 22.3 percent somewhat confident. Thus, only a quarter expressed any confidence in the program's finances, while 45 percent are not too confident and 29 percent are not at all confident.

The wording of our question about confidence in Social Security matches that of Greenwald et al. (2010), who conducted a nationally representative, random-digit telephone survey. Online Appendix Table A2 provides comparisons of the responses to this question in our sample and the subsample of their respondents age 25 – 59.[11] In their sample, 10.5 percent were very confident and 34.0 percent were somewhat confident. Together, about 45 percent express confidence in Social Security in the Greenwald et al. sample, compared to 25 percent in the Knowledge Networks panel. Of the remaining 55 percent, 36.3 percent are not too confident and 19.2 percent are not at all confident. Thus, our sample respondents show less confidence than those in the Greenwald et al. sample. In both samples, confidence tends to rise with age and is similar across men and women.

The survey then asks respondents how they expect the projected shortfall will be closed. As shown in Table 2, more than half, about 58 percent, expect the shortfall to be addressed by a combination of tax increases and benefit reductions. Nearly a quarter believe the shortfall will be addressed mostly or entirely through tax increases, while 18 percent believe the shortfall will

---

[11] We are indebted to Mathew Greenwald for providing these tabulations. The tabulations of the Knowledge Networks panel in Table A2 pertain to the respondents who answered both the ball/bins questions and the certainty equivalent questions, as described in Section 3 above.

be addressed mostly or entirely through benefit cuts. We focus on benefit cuts in the next several tables and report the results of analogous questions about tax increases in Table A3 in the Online Appendix.

When asked about the chance that the general level of benefits (as distinct from the benefits they expect to get individually) will decline over the next decade, the mean and median probabilities shown are 61 percent. The same question asked about a decline by the time the respondent reaches age 65 yields mean and median probabilities of 66.6 and 71 percent, respectively. This pessimism regarding future benefits is also reflected in expected benefit levels. Compared to the benefits they are supposed to get under current law, only 3 percent of respondents expect to get greater benefits, with 24 percent expecting the same benefits and 73 percent expecting lower benefits. When respondents are asked for a point estimate of benefits they expect to get relative to what they are supposed to get under current law, the mean and median responses for the point estimate of their benefits are 65.9 and 70 percent, respectively.

4.2 The Perceived Distribution of Future Benefits

The responses to the general questions presented in Table 2 show that respondents by and large expect to not receive all of the benefits they are supposed to get under current law. By themselves, they do not indicate whether individuals face uncertainty about the benefits they will get. Respondents could have a firm belief that they will receive, say, 70 percent of their current-law benefits, no more and no less.

Figure 1 graphs the aggregate CDF of perceived future Social Security benefits for all respondents to the survey. Looking at the probability mass at 0 and 100 percent, in aggregate, respondents perceive about a one in six chance of receiving no benefits whatsoever and about a one in four chance of receiving exactly the benefits they are supposed to get under current law. The perceived probability of outcomes strictly above current-law benefits is less than four percent. The remaining 54 percent of the probability mass lies strictly between 0 and 100, with an overall median at 69.5 percent.

The aggregate CDF shown in Figure 1 incorporates both the variation in possible outcomes within individual respondents' CDFs and the variation across respondents' CDFs. Figures 2 and 3 demonstrate that both sources of variation are important. Figure 2 shows the CDF of the mean perceived benefit across respondents. There is very little probability mass at zero, at 100 percent, or above 100 percent. Almost all of the respondents have mean perceived benefits between 0 and 100 percent of the benefits they are supposed to get under current law. The graph shows wide variation across respondents, with summary statistics provided in the first row of Table 1. The 25th and 75th percentiles are 37.1 and 83.4 percent, respectively.

We can use two other questions that we asked about the expectations of future benefits to assess the validity of the subjective probability distribution using our ball/bin question. In the first, we compute the correlation of the mean of the subjective distribution with the straightforward multiple-choice question about confidence in Social Security that we presented in Panel A of Table 2. This correlation is 0.54, indicating that those with more confidence tended to construct distributions with higher expected benefits. In the second, we compute the correlation of the mean of the subjective distribution with the point estimate of future benefits as a fraction of benefits under current law that we presented in Panel F of Table 2. This correlation is 0.69, and like the first, is highly statistically significant.

We use the expectation of the subjective probability distribution of future Social Security benefits, rather the point estimate, as our baseline measure of expected future benefits, for two reasons. First, we are not sure whether the point estimate offered by respondents is an expectation, a median, or a mode, whereas by construction the expectation of subjective benefits is an expectation. Second, the expectation of subjective benefits better predicts confidence in Social Security (as measured by the multiple-choice question) than the point estimate is able to predict confidence in Social Security. This suggests that the subjective expectation has less measurement error than the point estimate.

Figure 3 shows the CDF of the standard deviations of respondent CDFs. Only 7.5 percent have a standard deviation of zero. The second row of Table 1 provides summary statistics, indicating mean and median values of about 23 percent, with a quarter of the standard deviations at 33 percent or higher. These figures and statistics show that respondents perceive uncertainty in the possible benefits they will receive from Social Security and that the perceived distribution of possible benefits varies across respondents.

4.3 The Certainty Equivalent Social Security Benefit

It could be that respondents perceive an uncertain distribution of future benefits but that due to risk-neutrality or indifference, the uncertainty has little impact on their welfare. As a first measure of the importance of uncertain benefits, the survey asks, "How much does it matter to you that you do not know exactly how much you will get in Social Security benefits?" Panel G of Table 2 reports the results. Only 20.5 percent respond that the uncertainty matters little or does not matter, compared to 32 percent who respond that it matters somewhat and 47.5 percent who respond that it matters very much.

Figure 4 shows the distributions of the upper and lower bounds for the certainty equivalents across respondents. In the rest of the paper, we compute the certainty equivalent as the midpoint of the interval between them. Summary statistics for the certainty equivalents are shown in the third row of Table 1, denominated as a percentage of the benefits the respondents are supposed to get under current law. The mean certainty equivalent is 53.7 percent and the median is 57.5 percent.

4.4 Risk Premia for Policy Uncertainty

With the responses for the expected benefit from the elicited benefit distribution and for the certainty equivalent from the sequence of choices between guaranteed and uncertain benefits, we can subtract the average of the upper and lower bounds shown in Figure 4 from the expected value of benefits to obtain our key results: the risk premia that respondents would pay in the form of lower benefits to avoid the policy uncertainty surrounding Social Security.

Summary statistics for the distribution of risk premia are shown in the fourth row of Table 1. The mean risk premium is 5.8 percent and the median risk premium is 7.0 percent. About 25 percent of respondents have a risk premium of zero or less – there is no requirement imposed on their responses that the certainty equivalent obtained through the sequence of choices of guaranteed versus uncertain benefits yields a certainty equivalent below the expected value. The full distribution of risk premia is shown in Figure 5. About 11 percent of respondents have risk premia less than negative 20 percent. At the other end of the distribution, 25 percent of respondents have risk premia of 16.5 percent or more, with 4 percent having risk premia in excess of 50 percent. Given the challenging nature of our questions, we are not surprised to find that the tails of the distribution correspond to risk premia that may seem unreasonably high or low. The estimated risk premia rise moderately if we truncate or ignore observations in the tails. For example, if we ignore all observations below the 10^{th} percentile or above the 90^{th} percentile, the mean risk premium becomes 6.9%. Similarly, winsorizing at the 10^{th} and 90^{th} percentiles yields a mean risk premium of 6.3%. Truncating at the 25th and 75th percentiles increases the mean to 7.3%, while winsorizing at these percentiles yields a mean risk premium of 7.7%.[12]

Our main estimate of the risk premium is based on our novel method of eliciting a certainty equivalent and comparing that to the expected value. The benefits of this method are that the estimate does not rely on modeling or parameter assumptions, that it captures any responses that mitigate the impact of the uncertainty, and that it does not require estimates of the correlation between policy uncertainty and other sources of uncertainty affecting future consumption. Yet, our estimate is based on a question in which respondents are asked to value a hypothetical contract, and some respondents may have found it challenging to answer this question. We therefore compare our main estimate with an estimate of the risk premium that uses the methodology that prior papers have used; namely, to use a model to simulate the risk premium of the policy uncertainty. For each respondent, we calculate the risk premium that would be implied by the self-reported distribution of possible Social Security benefits, assuming constant relative risk aversion preferences with coefficients of relative risk aversion equal to 1, 3, and 5. These simulated risk premia also incorporate the information from the variable that captures how important the respondent expects Social Security to be in financing retirement spending.[13]

By construction, the distributions of simulated risk premia cannot have negative values and will show a zero premium for any respondent who did not indicate variation in the self-reported distribution of future Social Security benefits. Figure 6 shows the CDFs for the risk premia calculated in this manner, along with the CDF from Figure 5 based on self-reported certainty equivalents. The graph shows that for the 75 percent of respondents who reported positive risk premia, the CDF of those risk premia is intermediate between the hypothetical CDFs that would obtain if all respondents had coefficients of relative risk aversion between 3

[12] Recall from Section 3.2 that respondents who have missing benefit expectations or distributions that have no uncertainty are asked an alternative version of the certainty equivalence questions. The latter group tends to have lower risk premia, as would be expected based on the lack of perceived uncertainty. However, relative to Figure 5, which includes all respondents, the difference in the CDF when these respondents are excluded is minimal. We therefore use the full sample of respondents in the analyses below.

[13] Specifically, suppose that the respondent's Social Security benefits will be 100. Recall that the four responses to the survey question for the importance of Social Security are less than 25 percent, 25 – 50 percent, 50 – 75 percent, and more than 75 percent. If Social Security financed 25 percent of spending, that would require other income of 300. For 50 and 75 percent, the other income would have to be 100 and 33, respectively. Thus, we can assign other income of 200, 67, and 17 for the 25 – 50, 50 – 75, and 75 – 100 intervals, respectively. For the interval that is 0 – 25, we choose a value of 500 (consistent with Social Security funding 17 percent).

and 5. This indicates that our main estimate of the risk premium is consistent with the risk premium that would be obtained using a basic model and a reasonable assumption of the coefficient of relative risk aversion.[14]

4.5 Correlates of the Perceived Distribution of Benefits

We next consider the empirical relationships between the characteristics of the perceived distribution of Social Security benefits and the demographic and other control variables included in the survey. The most important of these is the age of the respondent. Figure 7a shows the expected benefits with a 95% confidence interval for 5-year age groups in our sample. The overall pattern is that the expected benefits, as a share of what respondents believe they are supposed to get under current law, are an increasing function of age. This pattern is evident at ages above 40 and even more so above 50. The point estimates for the average expected benefits rise from about 50 percent for the youngest age groups to about 80 percent for the oldest age group. Figure 7b shows the analogous graph of average risk premia by 5-year age group. There is a clear difference between those over 50 and those under 50. The former have risk premia around 11 percent while the latter have risk premia around 4 percent.

We consider age and other factors in regressions in Table 3. Estimates are shown using expected benefits, the standard deviation of benefits, and the risk premium as dependent variables. Each regression includes both the demographic variables from the Knowledge Networks panel and the other control variables about preferences and beliefs that we ask in our survey.[15] Focusing on the regression for expected benefits, an additional year of age leads to a 0.94 percentage point increase in expected benefits and a decrease in the standard deviation of 0.21 percentage points. These estimates are statistically significant at the 1 percent level. They are consistent with political rhetoric on Social Security reform – the older people get, the less likely they are to get a benefit cut, and the less variable they will expect that cut to be.

Table 3 also shows that some demographic and other control variables have significant effects on the expected benefits and the standard deviation of benefits. The effect of being retired on expected benefits is large and significant – equivalent to the effect of 10 years of age.

---

[14] For a 50-50 chance of gaining or losing 25 percent of one's wealth, the risk premia are 3.2, 9.0, and 13.5 percent for coefficients of relative risk aversion of 1, 3, and 5, respectively.

[15] Estimates that exclude the other control variables about preferences are similar and shown in Online Appendix Table A4.

This is consistent with Benítez-Silva et al. (2007), who find that early retirement can be partly explained by individuals retiring early in order to reduce exposure to policy risk. The point estimates for the effect on the standard deviation are negative but significant only at the 10 percent level. A 10 percent increase in income leads to a 0.31 percentage point decline in expected benefits and a 0.09 percentage point reduction in the standard deviation of benefits. This result is also consistent with political rhetoric surrounding Social Security reform, in which potential benefit cuts relative to current law are conjectured to be "progressive."[16] Race and education also matter – being Black or Hispanic or having less than a high school diploma all predict higher standard deviations. Black and Hispanic also predict higher expected benefits. Being female or having kids predicts lower expected benefits, equivalent to being about 3 or 6 years younger, respectively. Among the other control variables, living in the Northeast, greater longevity, greater importance of Social Security to retirement spending, greater trust in the political system, greater optimism and higher financial literacy all predict higher expected benefits.

The third column of Table 3 presents the regression in which the dependent variable is the risk premium calculated based on the certainty equivalent. The coefficient estimates generally follow those in the expected benefits regression in terms of sign and significance. For example, the effect of age is positive and significant while that of income is negative and significant. An additional decade of age increases the risk premium by 3 percentage points. A 10 percent increase income leads to a 0.18 percentage point reduction in the risk premium. This is exactly what theory would predict: a given amount of uncertainty is more costly to people near retirement as they have fewer opportunities to mitigate the benefit uncertainty by adjusting future labor supply or savings. Those with higher incomes tend to rely less on Social Security in retirement and would thus have lower risk premia regarding policy uncertainty. Other significant effects on the risk premia include the positive effects of being retired (equal to 30 years of age), being Black or Hispanic, being more risk averse, having a higher chance of living to age 75, and having more trust in the political system.

5. Cross-validation of Responses

[16] See, for example, Mermin (2005).

Recognizing that our survey asks questions that may be challenging for some respondents to answer, we incorporated a number of design features to enable us to determine how valid the answers to the key questions are. In Table 4, we present regressions that examine the effects of these design features. In Panel A, we construct a dummy variable for whether the respondent took the option of the guaranteed benefits (rather than the uncertain benefits) in the first round of questioning to determine the certainty equivalent. Recall that the first offer of guaranteed benefits was randomized at either 30 or 70 percent of the benefits the respondent is supposed to get under current law. We also randomized whether the guaranteed benefits were listed as the first or second option. If respondents are making reasonable choices, then we would expect that the guaranteed benefits are more likely to be chosen when they are higher and that the results should be insensitive to whether the guaranteed benefits are the first or second choice. The regressions in Panel A show this to be the case. Focusing on the second regression, which includes the demographic and other controls listed in Table 1, the respondent is 34 percentage points more likely to accept the guaranteed benefits when they are at 70 percent rather than 30 percent, an effect which is both large and statistically significant at the 1 percent level. Whether the guaranteed benefits are shown as the first or second option does not statistically significantly affect the probability that the respondent chooses them.

In Panel B of Table 4, we regress the respondent's certainty equivalent on three key variables that should predict it as well as three variables that should not predict it. The three key variables that should predict the certainty equivalent are: the respondent's expected benefits, the respondent's perceived standard deviation of benefits, and the measure of the respondent's risk aversion derived from separate questions about hypothetical gambles described in Section 2. Recall that the expected benefits and standard deviation are derived solely from the distribution of benefits presented by the respondent before questions are asked about the certainty equivalent. All three coefficients have the predicted signs and are statistically significant at the 1 percent level. The regression in column 2 shows that a 1 percentage point increase in expected benefits is associated with a 0.47 percentage point increase in the certainty equivalent, while a 1 percentage point increase in the standard deviation is associated with a reduction in the certainty equivalent by 0.38 percentage points. An increase of 1 unit in the measure of risk aversion (e.g., from a coefficient between 1 – 2 to one between 2 – 4) predicts a reduction of the certainty equivalent by 1.76 percentage points.

The next three regressors in Panel B should not affect the certainty equivalent if the respondent is able to report a stable underlying value: a dummy for the starting value being 70 percent, a dummy for the order in which the guaranteed benefits option is presented, and a dummy for whether the respondent sees the wide (rather than narrow) "weather" example. As in Panel A, the order in which the guaranteed benefits are presented has no statistically significant effect on the result. However, the starting value affects the certainty equivalent in a statistically significant way. The regressions show that if the respondent is first presented with guaranteed benefits of 70 percent rather than 30 percent, then the certainty equivalent that obtains from the sequence of questions is about 7 percentage points higher. This effect is statistically significant at the 1 percent level but should be zero – a fully rational respondent would give the same certainty equivalent regardless of the starting point. We explore possible explanations for this bias, along with one suggested correction, in Section 6 below.

The regressions in Panel B also show that illustrating how to put balls into bins with a wide distribution has a statistically insignificant effect on the certainty equivalent. In Panel C of Table 4, we consider the impact of the weather example on the standard deviation of the distribution of benefits reported by the respondent. The regressions indicate that respondents who are shown the wider distribution of temperatures subsequently report distributions with more uncertainty. A respondent shown the wider distribution has on average a standard deviation 2.88 percentage points greater than one shown the narrower distribution. Regressions not shown indicate that there is a negative but insignificant effect of being shown the wider distribution on the mean of the distribution of expected benefits. Because the estimate of the risk premium based on our methodology relies on the mean but not the reported dispersion of this distribution, our estimate is not affected by the sensitivity of the standard deviation to the weather example. In unreported regressions that are analogous to those in Panel C, but that instead use our main estimate of the risk premium as the dependent variable, we find no effect of the weather example on the risk premium – the point estimates are less than 0.1 in absolute value and statistically insignificant. This would not be the case if our measure of the risk premium were based on simulations from the reported benefit distribution rather than the certainty equivalent. Given the need to illustrate the balls and bins framework, the sensitivity of simulated risk premia to the dispersion in that illustration argues for our methodology based on directly eliciting the certainty equivalent to measure the risk premia.

6. Possible Adjustments to Risk Premia

In this section, we consider two possible adjustments to our estimated risk premia. The first pertains to the starting value, based on the results in Panel B of Table 4 indicating that the starting value in the sequence of questions that determine the respondent's certainty equivalent has an effect on the resulting value. The second considers the possibility that some respondents are not able to fully understand the questions and therefore give unreliable answers.

6.1 Starting Value Adjustment to the Certainty Equivalent

The first type of adjustment is based on a simple linear model in which the reported value (R_i) of the certainty equivalent for respondent i is a weighted average of the respondent's true underlying value (V_i) and the starting value (S):

$$R_i = (1 - \beta)V_i + \beta S_i$$

The parameter, β, can be interpreted as a bias that affects all respondents uniformly when they see a starting value for the certainty equivalent questions that is away from their true underlying value. Rearranging terms, we can write the true value as a function of the reported value and the starting value:

$$V_i = R_i - \frac{\beta}{1 - \beta}(S_i - R_i)$$

The second term on the right-hand side is the adjustment required to remove any starting-value bias. The higher is the starting value, the more negative is the adjustment. Given the linear model, and the fact that we randomly varied the starting values across respondents, we can estimate β by running a regression of the reported values of the certainty equivalent on the starting value, all the covariates from Table 1, and a set of variables related to response quality described below. Our estimate (to be more fully described below) of β is 0.17, suggesting that

0.17/0.83 or about 20 percent of the difference between the starting value and the reported value of the certainty equivalent reflects starting-value bias.

6.2 Response Quality Adjustment

To address the possibility that some respondents are not able to understand the questions fully and therefore give unreliable answers, we construct several indicator variables of response quality using data from the survey, estimate their relationship to the reported risk premia, and use this estimated relationship to adjust the reported risk premia of respondents who have one or more indicators showing low response quality. Specifically, we regress the risk premium on the response quality dummies controlling for the starting value and our standard set of demographic and other control variables. The coefficient on the response quality dummy is an estimate of the effect of giving a high-quality response on the reported risk premium. We create the risk premium adjusted for response quality by adding to the reported risk premium the coefficient of each response quality dummy that took on a value of zero for the respondent in question (i.e., for which the respondent gave a low-quality response). This adjustment yields the distribution of risk premia that would be obtained if everyone were able to give high-quality responses under the assumption that any differences in true underlying risk premia between high- and low-quality respondents are captured by our demographic and other control variables rather than by the quality response indicators.

Our measures of response quality are dummy variables for whether the survey respondents provided answers to questions that are correct, consistent, or otherwise in accordance with what we would expect from someone who fully understood the questions. The full list of measures, along with sample means, is shown in Online Appendix Table A5. For example, the first measure is whether the respondent reported a positive probability of not living to age 75. A respondent who instead gave 100% as the likelihood of being alive at age 75 may not fully understand the concept of a probability. Other measures for response quality focus on whether reported probabilities for an event (e.g., a change in Social Security policy) did not decrease as the time interval increased; on correct answers to the four questions testing financial literacy; on whether the point estimate for expected benefits given directly was close to the expectation of the distribution of benefits collected through the balls/bins method; and whether the survey duration was neither too short nor too long.

6.3 Adjusted Risk Premia

Table 5 shows the impact of the two adjustments to the mean and median risk premium, along with expected benefits, the certainty equivalent, and the standard deviation of benefits. The first row shows the unadjusted estimates. The second row shows the adjustments for the starting value, following the procedure described above and the coefficient on the starting value from the regressions in Online Appendix Table A6, which also include the demographic and other control variables from Table 3 and the response quality measures. Note that because the expected benefits and the standard deviation of benefits are elicited without using a starting value, no adjustment is made for those two outcomes. The starting value adjustment for the certainty equivalent is 0.7 percentage points at the mean. Standard errors (based on 10,000 bootstrap replications) indicate that the adjustment is statistically significant. The mean adjustment turns out to be positive because, on balance, the starting value of 30 percent is farther below the mean reported value 53.7 percent than the starting value of 70 percent is above the mean value. This 0.7 percentage point adjustment (≈ 20% of 53.7 – (30+70)/2) reduces the mean risk premium by the same amount, yielding a mean risk premium of 5.1 percent, shown in the third row of the table. The adjustment to the median risk premium is -1.0 percentage point, resulting in an adjusted median risk premium of 6.0 percent.

The fourth row of the table shows the adjustments to all four outcome measures based on response quality. The adjustments are negative for the standard deviation of benefits and positive for both expected benefits and the certainty equivalent of the distribution of benefits. Respondents who give higher quality answers tend to have less uncertainty about the distribution of future benefits, distributions that center closer to the reductions that would be consistent with official projections of the benefit cuts implied by current funding shortfalls, and higher certainty equivalents given those higher expected benefits. All of these adjustments, at both the mean and the median, are statistically significantly different from zero.

Importantly, the upward adjustments to the expected benefits are larger than those to the certainty equivalent, generating net upward adjustments to the mean and median risk premium of 1.6 and 1.4 percentage points, respectively. Though not statistically significant, the point estimates for these adjustments more than offset the negative adjustment due to the starting value. Thus, the mean and median adjusted risk premiums in the last row of the table are higher than the unadjusted estimates in the first row, suggesting that while starting value bias and response quality do affect the reported benefit distributions and certainty equivalents, the net impact on the risk premium is negative and the unadjusted estimates are a lower bound for the risk premium. Using the adjusted estimates, respondents on average expect about two thirds of the benefits they are supposed to get under current law and view the distribution of future benefits as equivalent to guaranteed benefits of about 60 percent of what they are supposed to get under current law, yielding a risk premium of about 7 percent.

7. Conclusion

While it has been long recognized that policy uncertainty can have welfare consequences, the empirical literature trying to estimate the size of such welfare losses is relatively sparse. This paper contributes to this literature by providing the first empirical estimate of the size of the welfare loss to individuals of policy uncertainty in U.S. Social Security benefits. Relative to the literature on policy uncertainty, we take a novel approach to estimating this welfare loss – we elicit from a nationally representative sample of survey respondents both the expected value and the certainty equivalent of future Social Security benefits as a fraction of the benefits they are scheduled to receive under current law. Our approach mimics the traditional approach of measuring risk premia in the finance literature, except that, by necessity, we measure the certainty equivalent using survey methods rather than from market data. We are keenly aware of the challenges of getting survey respondents to give meaningful answers to hard questions, and we introduce randomizations in our survey instrument that allow us to detect potential biases, and in some cases, correct for them.

We apply our methodology to policy uncertainty surrounding Social Security benefits because this is one of the largest sources of unavoidable and uninsurable economic policy uncertainty to U.S. residents. The projected financial shortfalls in the Social Security program have been the subject of active policy discussion for over 15 years. During that period, no clear policy direction has emerged for how the projected shortfalls will be closed, and, as a result, households are exposed to considerable policy uncertainty. We find that on average respondents would be willing to forego around 6 percent of the benefits they are supposed to get under current law to remove the policy uncertainty associated with their future benefits. Because respondents only expect to receive 60 percent of the benefits they are suppose to get under current law, this risk premium is equivalent to about 10 percent of expected benefits. The informal estimate of the accrued obligation, under current law, for individuals 25-59 by the Office of the Chief Actuary at the Social Security Administration is about $12 trillion.[17] So, in dollar terms, the welfare cost is 6 percent of $12 trillion, which equals about $700 billion. In other words, the government could cut future Social Security benefits by a total of about $700 billion in present value terms without making individuals worse off on average if it (somehow) could remove all policy uncertainty surrounding future benefits.

A promising avenue for further research, in economics as well as political science, is the study of mechanisms by which policy uncertainty could be reduced. From the perspective of the current generation, which we adopted in this paper, all policy uncertainty is costly, but from an intergenerational risk-sharing perspective some degree of policy uncertainty may be optimal. Yet the current degree of policy uncertainty almost surely includes more than just the uncertainty that could potentially be justified by intergenerational risk sharing because it also includes uncertainty that is due to the uncertain behavior of political actors. So, one avenue is to try to eliminate the political component of the policy risk by specifying time-invariant rules that specify benefit levels as functions of macroeconomic parameters. Sweden and Germany have adopted such systems, and their operation and performance are analyzed by Auerbach and Lee (2011). Another avenue for reducing policy risk, but one that has not been implemented anywhere, is to create government securities that pay out a benefit stream that has the same time-profile as Social Security benefits as proposed by Geanakoplos and Zeldes (2009) – in other words, people would be granted ownership of some kind of wage- and/or inflation indexed deferred annuity in return for contributions to the Social Security system. This will not completely eliminate policy risk – for example, the government could renege on annuity payments just like it could default on its Treasury bills, but this would be politically much more difficult than changing benefit rules. Transforming Social Security into a system with personal accounts may be an alternative way of reducing policy risk but could expose individuals to other risks depending on the types of assets that individual could hold in such accounts.[18] Indeed, Smetters and Theseira (2011) find that fundamental reforms away from traditional pay-as-you-go

[17] Steve C. Goss, Personal Communication, April 5, 2012.

[18] Diamond (1997) assesses the insulation against political risk that Chile's privatized mandatory pension system achieves.

Social Security systems to systems with funded accounts can be partly explained as a response to political uncertainty, either coming from a lack of intergenerational trust or from a lack of trust in the government to save.

References

Advisory Council on Social Security (1994 – 1996). 1997. *Report of the 1994 – 1996 Advisory Council on Social Security.* Washington: U.S. Government Printing Office, January 7. http://www.ssa.gov/history/reports/adcouncil/

Alm, James. 1988. "Uncertain Tax Policies, Individual Behavior, and Welfare," *American Economic Review*, 78(1), pp. 237–245.

Auerbach, Alan J. and Ronald Lee. 2011. "Welfare and Generational Equity in Sustainable Unfunded Pension Systems," *Journal of Public Economics*, 95(1-2), pp. 16–27.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2015. "Measuring Economic Policy Uncertainty," unpublished manuscript, Stanford.

Benítez-Silva, Hugo, Debra S. Dwyer, Frank Heiland, and Warren C. Sanderson. 2007. "Retirement and Social Security Reform Expectations: A Solution to the New Early Retirement Puzzle," unpublished manuscript, SUNY-Stony Brook.

Bernheim, B. Douglas and Lawrence Levin. 1989. "Social Security and Personal Saving: An Analysis of Expectations," *American Economic Review*, 79(2), pp. 97–102.

Blake, David. 2008. "What is a Promise from the Government Worth? Quantifying Political Risk in State and Personal Pension Schemes in the United Kingdom, *Economica*, 75(298), pp. 342–361.

Borgmann, Christoph, and Matthias Heidler. 2007. "Volatility of Social Security Wealth: Political Risks of Benefit-Rule Changes in Germany," *FinanzArchiv: Public Finance Analysis*, 63(1), pp. 83–106.

Board of Trustees, for the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. 2015. *Annual Report*. http://www.ssa.gov/oact/tr/2015/tr2015.pdf.

Caliendo, Frank N., Aspen Gorry, and Sita Slavov. 2015. "The Cost of Uncertainty about the Timing of Social Security Reform," unpublished manuscript, Utah State University.

Delavande, Adeline, and Susann Rohwedder. 2008. "Eliciting Subjective Expectations in Internet Surveys," *Public Opinion Quarterly*, 72(5), 866–891.

Diamond, Peter. 1997. "Insulation of Pensions from Political Risk," in Salvador Valdés-Prieto (ed.), *The Economics of Pensions: Principles, Policies and International Experience*. Cambridge: Cambridge University Press, pp. 33–57.

Dominitz, Jeff, and Charles F. Manski. 2006. "Measuring Pension-benefit Expectation Probabilistically," *Labour*, 20(2), pp. 201–236.

Dušek, Libor. 2007. "Political Risk of Social Security: The Case of the Indexation of Benefits in the Czech Republic," CERGE-EI working paper 318, Prague.

Geanakoplos, John, and Stephen P. Zeldes. 2009. "Reforming Social Security with Progressive Personal Accounts," in Jeffrey R. Brown, Jeffrey B. Liebman, and David A. Wise (eds.), *Social Security Policy in a Changing Environment*. Chicago: University of Chicago Press, pp. 73–128.

Geanakoplos, John, and Stephen P. Zeldes. 2010. "Market Valuation of Accrued Social Security Benefits," in Deborah Lucas (ed.), *Measuring and Managing Federal Financial Risk*. Chicago: University of Chicago Press, pp. 213–233.

Giavazzi, Francesco, and Michael McMahon. 2012. "Policy Uncertainty and Household Savings," *Review of Economics and Statistics*, 94(2).

Gomes, Francisco J., Laurence J. Kotlikoff, and Luis M. Viceira. 2012. "The Excess Burden of Government Indecision," in Jeffrey R. Brown (ed.) *Tax Policy and the Economy, Volume 26*, University of Chicago Press.

Gordon, Roger H., and Hal R. Varian. 1988. "Intergenerational Risk Sharing," *Journal of Public Economics*, 37(2), pp. 185–202.

Greenwald, Mathew, Arie Kapteyn, Olivia S. Mitchell, and Lisa Schneider. 2010. "What Do People Know About Social Security?" Pension Research Council Working Paper, September.

Guiso, Luigi, Tullio Jappelli, and Mario Padula. 2013. "Pension Wealth Uncertainty," *Journal of Risk and Insurance*, 80(4), pp. 1057–1085.

Gustman, Alan L., and Thomas L. Steinmeier. 2005. "Imperfect Knowledge of Social Security and Pensions," *Industrial Relations*, 44(2), pp. 373–97.

Kim, Iltae, Arthur Snow, and Ronald S. Warren Jr. 1995. "Tax-Rate Uncertainty, Factor Supplies, and Welfare," *Economic Inquiry,* 33(1), pp. 159–169.

Knowledge Networks. 2010. "KnowledgePanel® Design Summary," Manuscript, May. http://www.knowledgenetworks.com/knpanel/docs/KnowledgePanel(R)-Design-Summary-Description.pdf

Liebman, Jeffrey B. and Erzo F.P. Luttmer. 2012. "The Perception of Social Security Incentives for Labor Supply and Retirement: The Median Voter Knows More Than You'd Think." in Jeffrey R. Brown (ed.) *Tax Policy and the Economy,* Volume 26, University of Chicago Press. 1–42.

Liebman, Jeffrey B. and Erzo F.P. Luttmer. 2015. "Would People Behave Differently If They Better Understood Social Security? Evidence From a Field Experiment," *American Economic Journal: Economic Policy,* 7(1), pp. 275–229.

McHale, John. 2001. "The Risk of Social Security Benefit-Rule Changes: Some International Evidence," in John Y. Campbell and Martin S. Feldstein (eds.) *Risk Aspects of Investment-Based Social Security Reform*. Chicago: University of Chicago Press, 247–290.

Mermin, Gordon B.T. 2005. "Reforming Social Security through Price and Progressive Price Indexing," Older Americans' Economic Security Brief No. 6, The Urban Institute, December. http://www.urban.org/publications/900903.html

Nataraj, Sita and John B. Shoven. 2003. "Comparing the Risks of Social Security with and without Individual Accounts," *American Economic Review*, pp. 348-53.

Shoven, John B., and Sita N. Slavov. 2006. "Political Risk versus Market Risk in Social Security," NBER Working Paper 12135.

Skinner, Jonathan. 1988. "The Welfare Cost of Uncertain Tax Policy," *Journal of Public Economics*, 37(2), 129–145.

Smetters, Kent, and Walter E. Theseira. 2011. "A Matter of Trust: Understanding Worldwide Public Pension Conversions," NBER Working Paper 17015.

Stiglitz, Joseph E. 1982. "Utilitarianism and Horizontal Equity: The Case for Random Taxation," *Journal of Public Economics*, 18(1), pp. 1–33.

Weiss, Laurence. 1976. "The Desirability of Cheating Incentives and Randomness in the Optimal Income," *Journal of Political Economy*, 84(6), pp. 1343–1352.

Van der Wiel, Karen. 2008. "Preparing for Policy Changes: Social Security Expectations and Pension Scheme Participation," Institute for the Study of Labor (IZA) Discussion Paper No. 3623. http://ftp.iza.org/dp3623.pdf

Figure 1: Average Perceived CDF of Future Social Security Benefits

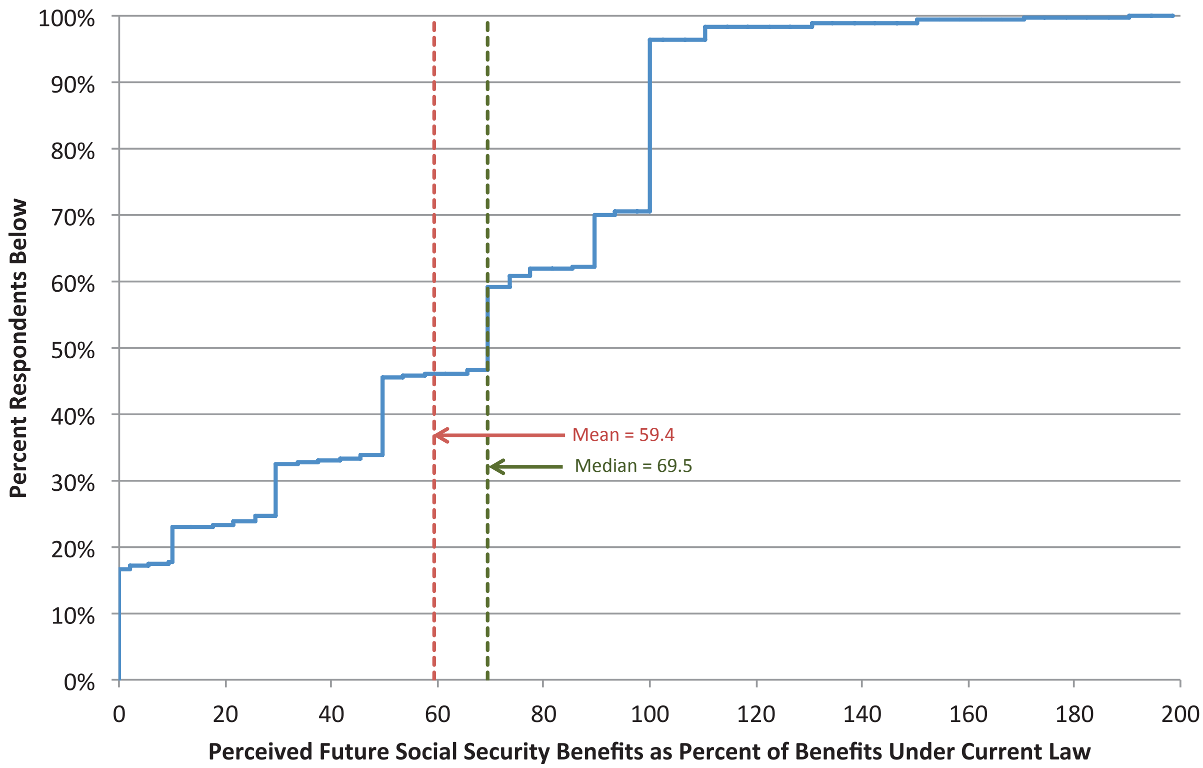

Figure 2: CDF of Perceived Mean Future Social Security Benefits

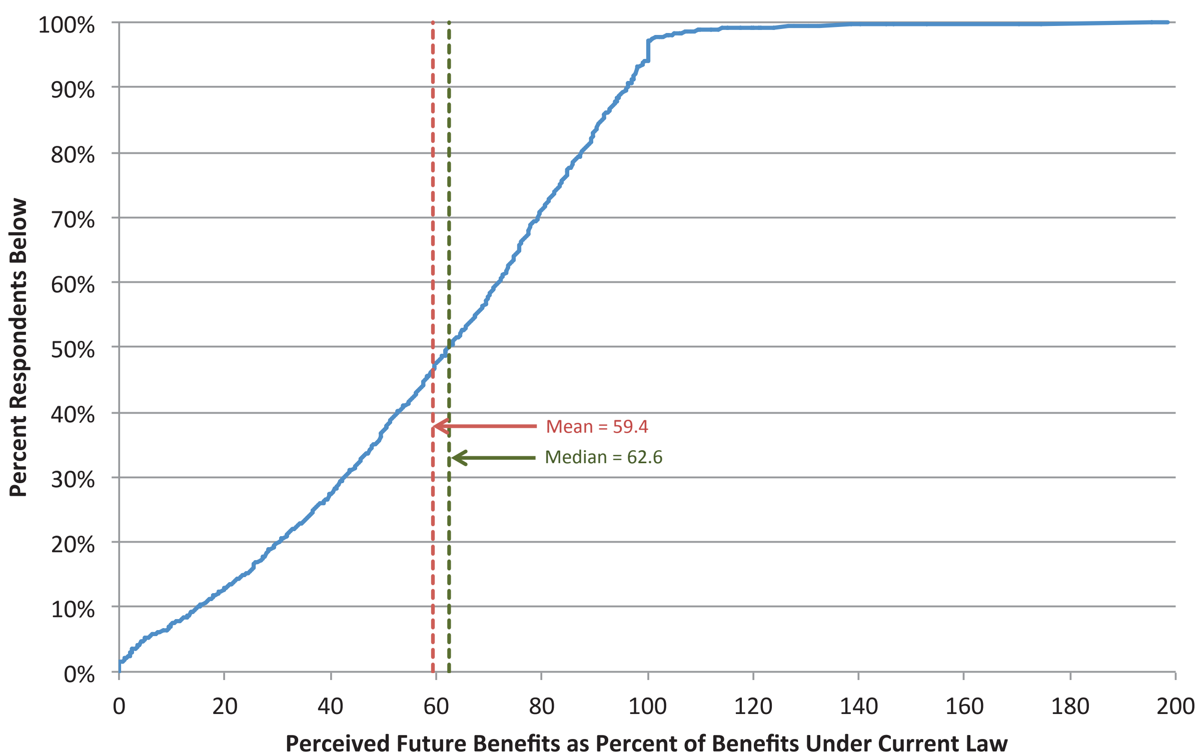

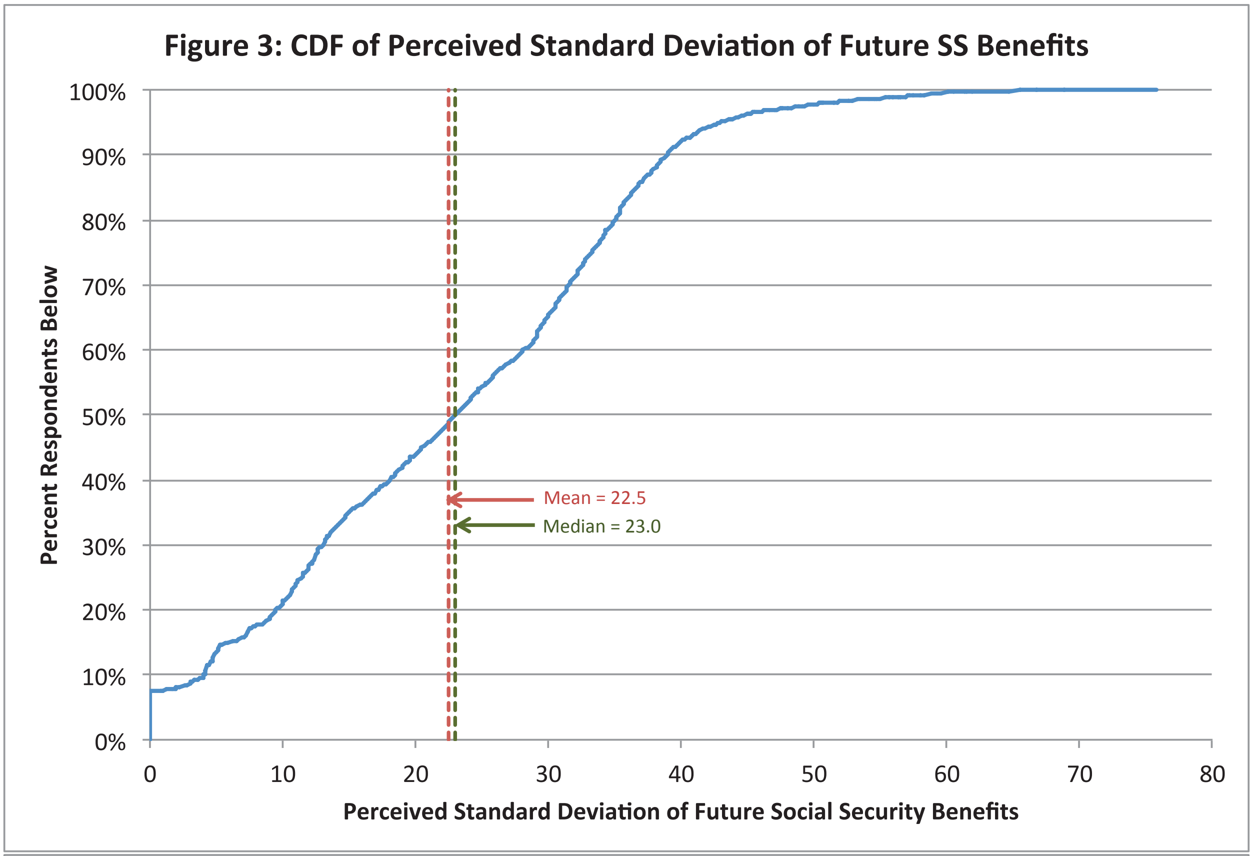
Figure 3: CDF of Perceived Standard Deviation of Future SS Benefits
Percent Respondents Below
Mean = 22.5
Median = 23.0
Perceived Standard Deviation of Future Social Security Benefits

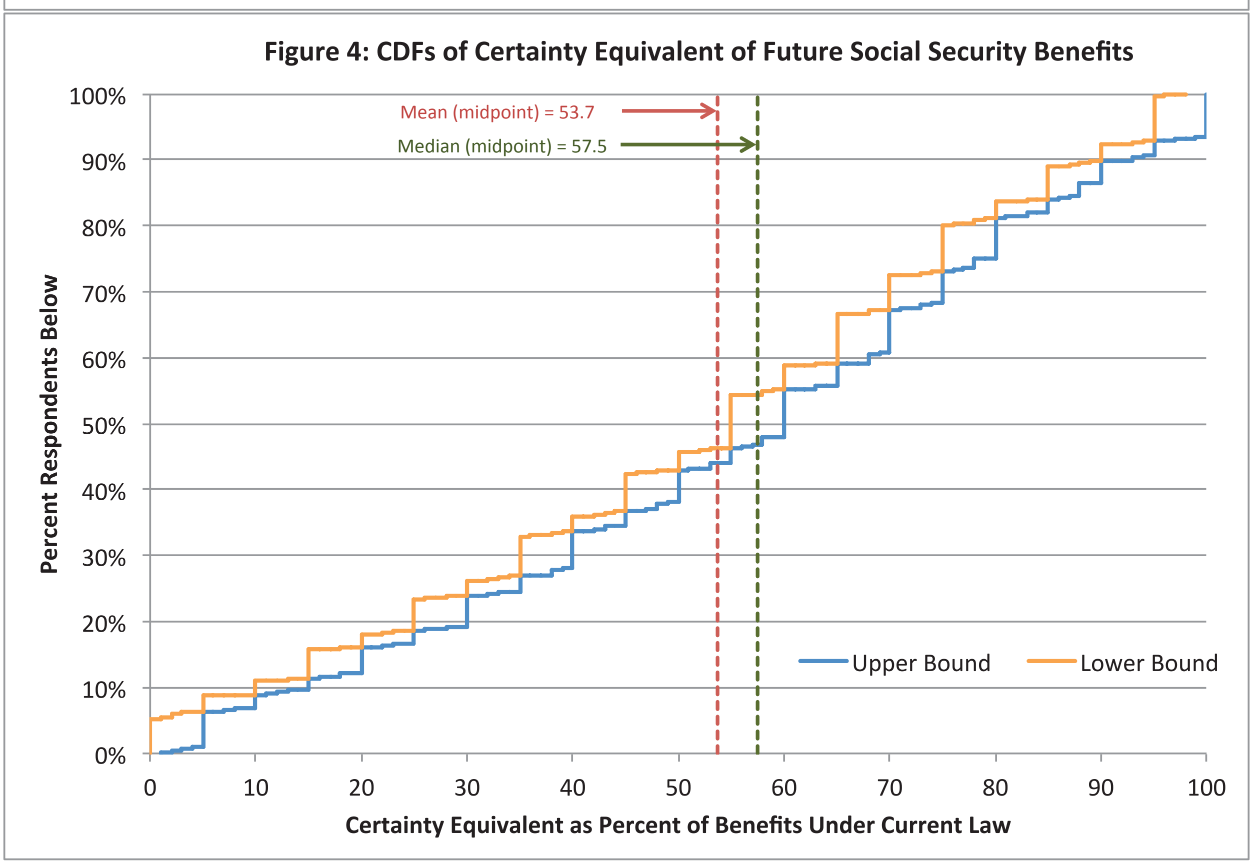
Figure 4: CDFs of Certainty Equivalent of Future Social Security Benefits
Percent Respondents Below
Mean (midpoint) = 53.7
Median (midpoint) = 57.5
Upper Bound
Lower Bound
Certainty Equivalent as Percent of Benefits Under Current Law

Figure 5: Distribution of Risk Premium

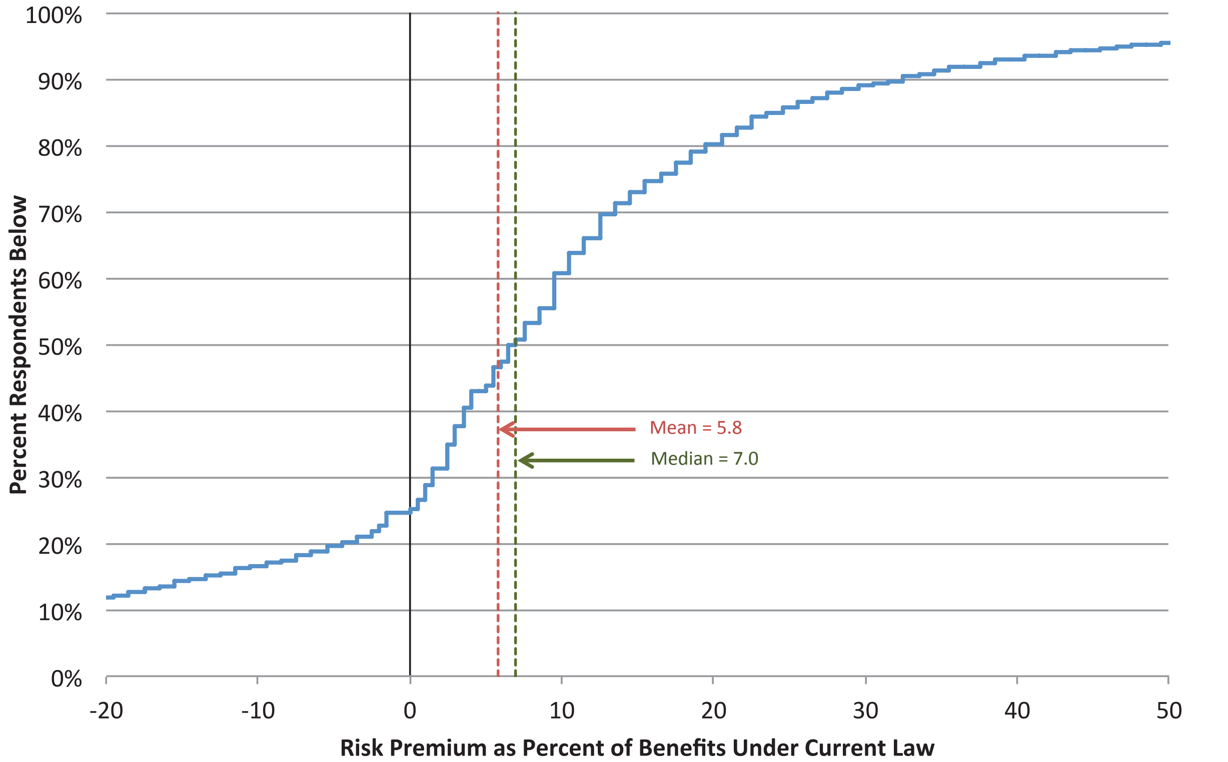

Figure 6: Distribution of Risk Premia Based on Simulated Risk Aversion (CRRA)

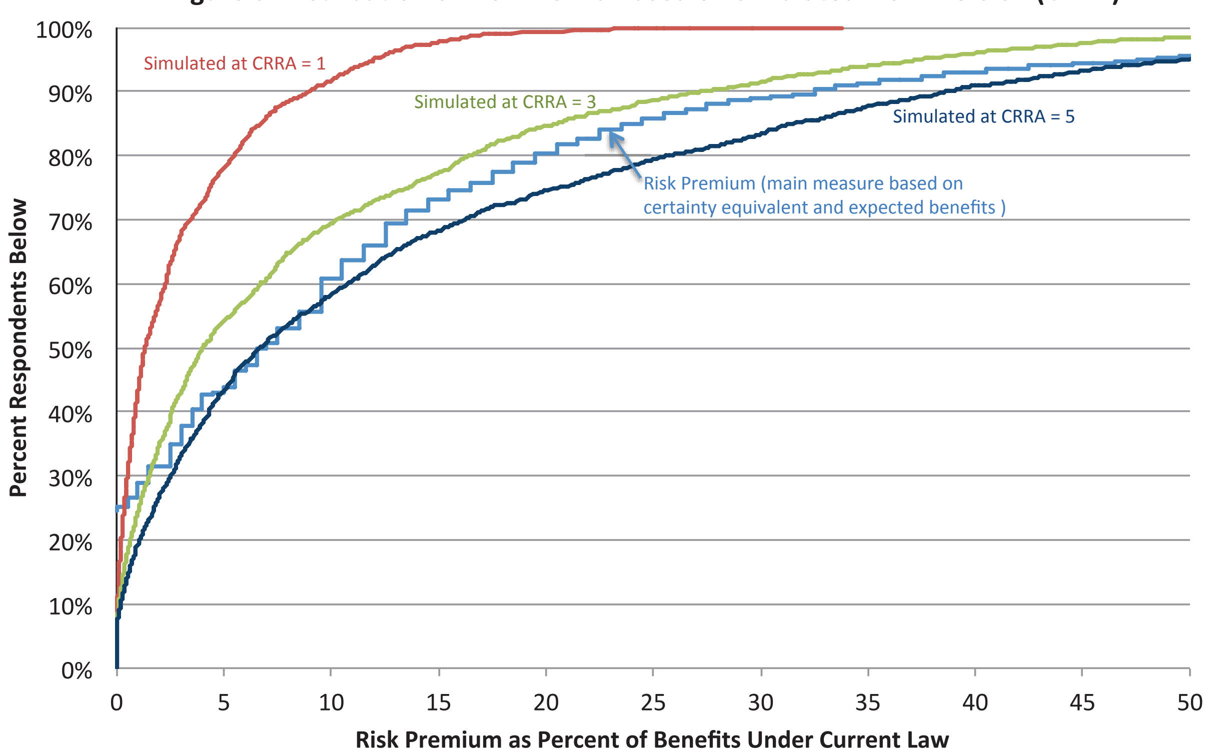

Figure 7a: Expected Benefits by Age Group (Mean and 95% Confidence Interval)

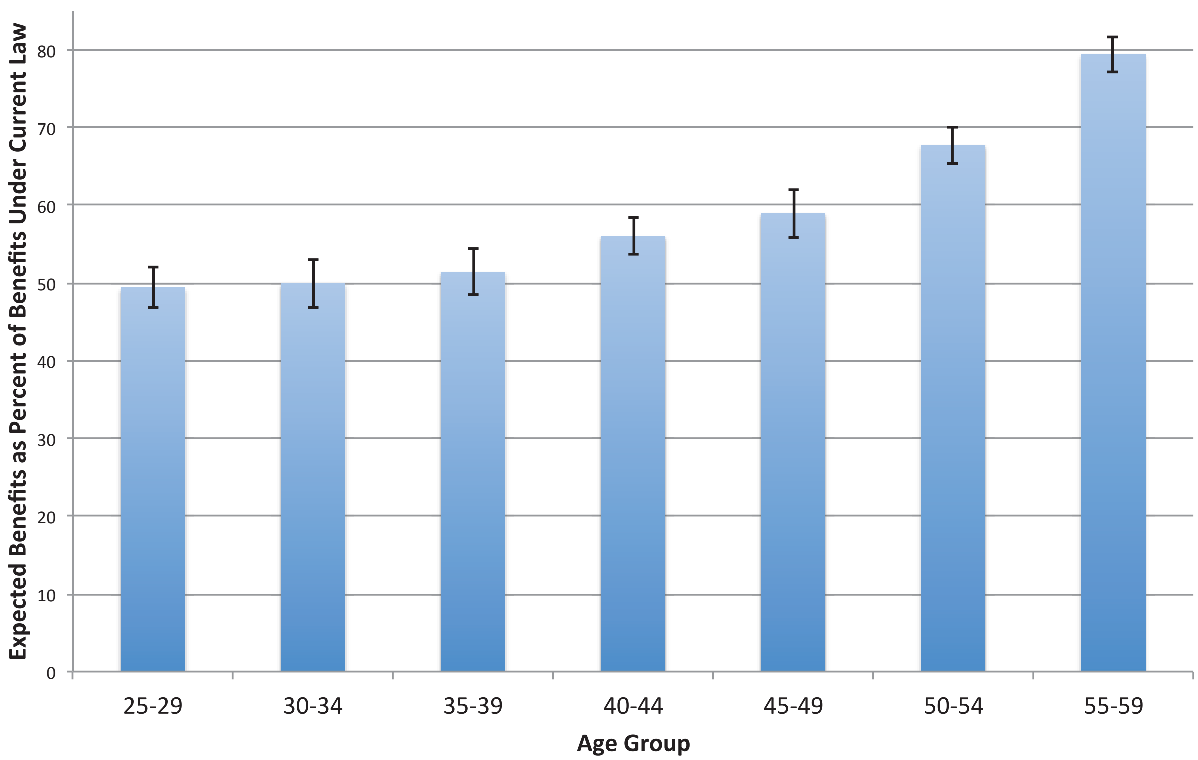

Figure 7b: Perceived Risk Premium by Age Group (Mean and 95% Confidence Interval)

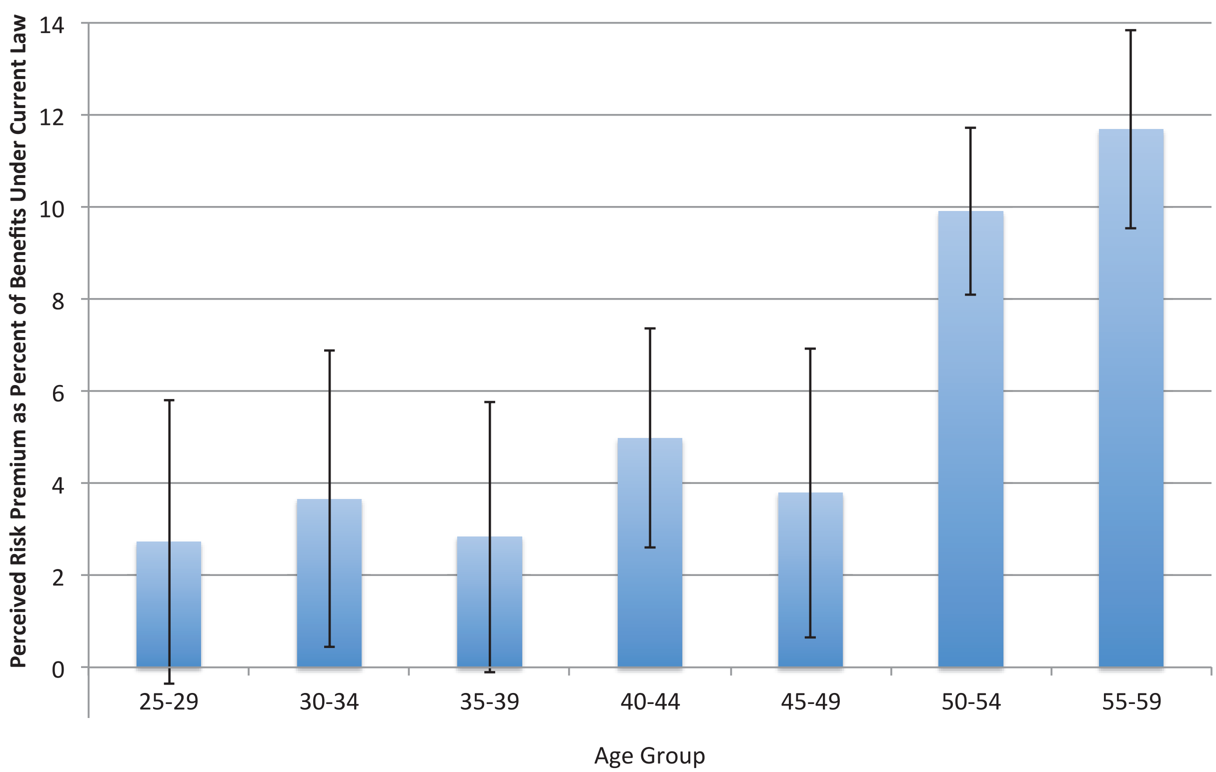

Table 1: Summary Statistics

Variable	Mean	(1) Standard Deviation	(2) 25th Percentile	(3) Median	(4) 75th Percentile	(5) Number of Observations
Key Outcome Variables						
(all in Percent of Benefits Under Current Law)						
Expected Benefits	59.4	30.1	37.1	62.6	83.4	2,960
Standard Deviation of Expected Benefits	22.5	13.7	11.4	23.0	33.3	2,960
Midpoint of Certainty Equivalent	53.7	27.8	32.5	57.5	76.5	2,939
Midpoint of Risk Premium	5.8	28.0	0.0	7.0	16.5	2,939
Demographic Control Variables						
Age	42.5	10.0	34.0	43.0	51.0	3,053
Ethnicity						
White	0.702					3,053
Black	0.103					3,053
Hispanic	0.154					3,053
Other	0.041					3,053
Education						
High School Dropout	0.088					3,053
High School Education	0.286					3,053
Some College	0.229					3,053
Bachelor's or Above	0.397					3,053
ln(Household Size)	1.00	0.52	0.69	1.10	1.39	3,053
ln(Household Income)	10.97	0.89	10.53	11.12	11.63	3,053
Marital Status						
Married	0.643					3,053
Widowed	0.013					3,053
Divorced	0.076					3,053
Separated	0.018					3,053
Never Married	0.157					3,053
Living with partner	0.092					3,053
Female	0.464					3,053
Homeowner	0.726					3,053
Region						
Northeast	0.174					3,053
Midwest	0.237					3,053
South	0.354					3,053
West	0.235					3,053
Lives in MSA	0.843					3,053
Kids in Household	0.467					3,053
Employment Status						
Currently Working	0.788					3,053
Retired	0.019					3,053
Disabled	0.021					3,053
Unemployed	0.086					3,053
Not Working	0.085					3,053
Other Control Variables						
Risk Aversion Index (Using Lifetime Income Gambles, 1-6 scale)	4.6	1.4	4.0	5.0	6.0	2,997
Subjective Probability of Surviving To Age 75 (percent)	67.9	22.5	51.0	71.0	85.0	2,935
Importance of Social Security Funds during Retirement (1-4 scale)	2.8	1.0	2.0	3.0	4.0	2,982
Trust in Elected Federal Officials (1-5 scale)	2.2	1.0	1.0	2.0	3.0	3,018
Optimism Indicator (standarized variable)	0.0	1.0	-0.6	0.0	0.7	2,955
Financial Literacy (0-4 scale)	2.4	1.2	2.0	3.0	3.0	3,053

Notes: Key outcome variables are measured in the June 2011 Social Security Policy Risk Survey, designed by the authors and fielded by Knowledge Networks. The baseline demographics are the values in the standard demographic profile variables at the time of the baseline survey (June 2010). The standard demographic profile is collected by Knowledge Networks. The sample is restricted to individuals between the ages of 25 and 59 as of May 2011. See the text and Appendix B for a discussion of and definitions of the key outcome variables. The risk-aversion variable is an index that runs from 1 to 6 and is based on five questions about hypothetical choices between a riskless and a risky job (Q6.1-Q6.5). The index corresponds respectively to the following six CRRA ranges: [<0.5], [0.5-1],[1-2],[2-4],[4-8],[>8]. Importance of Social Security Funds during Retirement is measured on a 4-point scale from "not so important" to "extremely important" (Q6.10). Trust in Elected Federal Officials is a on a five-point scale, with higher values indicating more trust (Q6.11). The Optimism Indicator is the standardized average of the non-missing responses to the six items (reverse coded when appropriate) of Q6.12. The financial literacy index is the number of correct responses to the four questions on financial literacy (Q6.13-Q6.13).

Table 2: Expectations about Social Security

	(1) Mean		(2) Median		(3) Number of Observations
Panel A: Respondent Confidence in Social Security					
Very Confident	0.033	(0.003)	..	..	3,045
Somewhat Confident	0.223	(0.008)	..	..	3,045
Not too Confident	0.453	(0.009)	..	..	3,045
Not at all Confident	0.291	(0.008)	..	..	3,045
Panel B: Does Social Security Face a Financial Shortfall?					
Yes	0.914	(0.005)	..	..	3,036
No	0.086	(0.005)	..	..	3,036
Panel C: How Will the Government Address the Shortfall?					
Mostly or Entirely through Benefit Cuts	0.183	(0.007)	..	..	3,028
Balanced Mix of Benefit Cuts and Tax Increases	0.576	(0.009)	..	..	3,028
Mostly or Entirely through Tax Increases	0.241	(0.008)	..	..	3,028
Panel D: Chance of a Decline in General Level of Benefits					
Within 10 Years	61.0	(0.5)	61.0	(0.8)	2,937
By Age 65	66.6	(0.5)	71.0	(0.5)	2,840
Panel E: Do you Expect More, the Same, or Less Benefits than you are Supposed to Get Under Current Law?					
More	0.028	(0.003)	..	..	3,026
The Same	0.241	(0.008)	..	..	3,026
Less	0.731	(0.008)	..	..	3,026
Panel F: Point Estimate of Future Benefits as % of Current Benefits (measured using a slider)	65.9	(0.6)	70.0	(0.5)	2,956
Panel G: Importance of Benefit Amount Uncertainty					
Matters Very Much	0.475	(0.009)	..	..	3,038
Matters Somewhat	0.320	(0.008)	..	..	3,038
Matters Little	0.148	(0.006)	..	..	3,038
Does Not Matter at All	0.057	(0.004)	..	..	3,038

Notes: Robust standard errors in parentheses. Data from the June 2011 Social Security Policy Risk Survey, designed by the authors and fielded by Knowledge Networks. The sample is restricted to individuals between the ages of 25 and 59 as of May 2011. See Appendix B for exact question definitions: Q1.2 for Panel A, Q2.1 for Panel B, Q2.2 for Panel C, Q2.11 for Panel D, Q3.1 for Panel E, Q3.2 for Panel F, Q4.1 for Panel G.

Table 3: Correlates of Expected Benefits, Standard Deviation of Benefits, and the Risk Premium

	(1)		(2)		(3)	
	Dep. Variable: Expected Benefits		Dep. Variable: Standard Deviation of Benefits		Dep. Variable: Risk Premium	
Basic Demographics						
Age	0.94***	(0.06)	-0.21***	(0.03)	0.29***	(0.06)
Black	5.1***	(1.9)	3.1***	(1.0)	9.5***	(2.1)
Hispanic	3.9**	(1.6)	1.8**	(0.8)	4.8***	(1.7)
Other	-0.8	(2.8)	1.9	(1.2)	-4.1	(2.6)
Highschool Dropout	0.7	(2.2)	3.7***	(1.1)	3.4	(2.4)
Some College	-0.3	(1.5)	-0.7	(0.7)	-1.6	(1.5)
Bachelor's Degree or Higher	0.6	(1.5)	0.1	(0.7)	-0.5	(1.5)
Ln Household Size	0.8	(1.6)	0.6	(0.8)	0.3	(1.6)
Ln Household Income	-3.1***	(0.8)	-0.9**	(0.4)	-1.8**	(0.9)
Widowed	6.9*	(4.0)	-0.8	(2.9)	5.3	(4.0)
Divorced	-0.5	(2.0)	0.0	(1.1)	-0.5	(2.1)
Separated	-0.3	(3.7)	-0.1	(1.8)	6.2*	(3.3)
Never Married	2.9	(1.7)	-1.0	(0.9)	1.9	(1.7)
Lives With Partner	0.6	(1.9)	0.7	(0.9)	1.4	(2.1)
Female	-3.0***	(1.1)	0.3	(0.5)	2.0*	(1.1)
Owns House	-1.1	(1.3)	-1.2*	(0.7)	-1.9	(1.4)
Lives in Northeast	4.0***	(1.5)	0.1	(0.7)	-1.1	(1.5)
Lives in Midwest	2.4*	(1.3)	0.1	(0.7)	-1.0	(1.3)
Lives in West	0.0	(1.4)	1.2*	(0.7)	-2.3	(1.4)
Lives in MSA	2.3	(1.4)	-0.1	(0.7)	0.2	(1.5)
Kids in Household	-5.4***	(1.5)	0.1	(0.8)	-1.1	(1.5)
Retired	9.3**	(3.8)	-3.6*	(1.9)	10.6***	(3.4)
Disabled	-2.5	(3.8)	-1.5	(2.2)	-4.9	(4.1)
Unemployed	-2.0	(2.0)	-1.4	(1.0)	-0.7	(2.2)
Not Working	-0.5	(2.0)	1.9**	(0.9)	2.0	(2.0)
Other Control Variables						
Risk Aversion Index	0.0	(0.4)	-0.2	(0.2)	1.6***	(0.5)
Subjective Probability of Surviving To Age 75	0.13***	(0.03)	-0.03**	(0.01)	0.05**	(0.03)
Importance of SS to Retirement Spending	2.3***	(0.6)	-0.4	(0.3)	0.9	(0.6)
Trust in Elected Federal Officials	6.0***	(0.5)	0.0	(0.3)	2.3***	(0.5)
Optimism Index	1.7***	(0.6)	-0.4	(0.3)	0.6	(0.5)
Financial Literacy	1.5***	(0.5)	0.0	(0.3)	0.7	(0.6)
R^2	0.202		0.069		0.075	
N	2,960		2,960		2,939	

Notes: Robust standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%. Missing values of explanatory variables are dummied out. Expected Benefits and Standard Deviation of Expected Benefits are based on the bin/ball question that elicites the subjective distribution of future Social Security benefits (Q3.3-Q3.6). The risk premium is the percent of benefits under current law that respondents are willing to sacrifice in order to receive their expected benefits with certainty. All dependent variables are expressed as a percentage of benefits under current law. See the note to Table 1 for more details on the explanatory variables. Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks. The sample is restricted to individuals between the ages of 25 and 59 as of May 2011.

Table 4: Are Responses on Certainty Equivalence Meaningful?

	(1)	(2)
Panel A: Effect of Starting Value on First Choice	Dependent Variable: Dummy for Respondent Choosing Guaranteed Benefits	
Starting Value is 70%	0.34*** (0.02)	0.34*** (0.02)
Guaranteed Benefits is Second Option	-0.01 (0.02)	-0.02 (0.02)
Demographic and Other Controls	No	Yes
R^2	0.113	0.189
N	2,939	2,939
Panel B: Effects of Randomizations, Perceived Uncertainty, and Risk-Aversion on Certainty Equivalent	Dependent Variable: Certainty Equivalent of Social Security Benefits	
Expected Social Security Benefits	0.49*** (0.02)	0.47*** (0.02)
Perceived Standard Deviation of own Social Security Benefits	-0.42*** (0.04)	-0.38*** (0.04)
Risk Aversion Index	-1.99*** (0.34)	-1.76*** (0.37)
Starting Value is 70%	7.68*** (0.82)	7.44*** (0.82)
Guaranteed Benefits is Second Option	0.94 (0.83)	1.18 (0.82)
Respondent Sees High SD Weather Example	0.81 (0.83)	0.56 (0.83)
Demographic and Other Controls	No	Yes
R^2	0.352	0.375
N	2,939	2,939
Panel C: Effect of Weather Example On Perceived Uncertainty	Dependent Variable: Std Dev. of Perceived Social Security Benefits	
Respondent Sees High SD Weather Example	2.73*** (0.50)	2.88*** (0.49)
Demographic and Other Controls	No	Yes
R^2	0.010	0.080
N	2,960	2,960

Notes: Robust standard errors in parentheses. * significant at 10%, ** significant at 5%, *** significant at 1%. Demographic and other controls is the set of controls used in Table 3. The Certainty Equivalent is the percent of benefits under current law that the respondent is just willing to accept in place of benefits under current law if the certainty equivalent is guaranteed in an unbreakable contract. Expected Social Security Benefits and Perceived Standard Deviation of Social Security Benefits are based on the respondent's subjective probability distribution of future Social Security Benefits as elicited by the Bin/Ball question (Q3.3-Q3.6). The Risk Aversion Index is defined in the note to Table 1. The weather example is an example of a probability distribution using the Bin/Ball format that was presented to the respondent prior to Q3.3. The variable Respondent Sees High SD Weather Example is a dummy that equals 1 if the variance of the distribution in the example was high. Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks. The sample is restricted to individuals between the ages of 25 and 59 as of May 2011.

Table 5: Adjustments for Starting Value and Response Quality

	(1)		(2)		(3)		(4)	
	Risk Premium		Expected Benefits		Certainty Equivalent		Standard Deviation of Benefits	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median
(1) Unadjusted Estimate	5.8	7.0	59.5	62.7	53.7	57.5	22.6	23.0
	(0.5)	(0.4)	(0.5)	(0.9)	(0.5)	(0.1)	(0.3)	(0.4)
(2) Starting Value Adjustment	-0.7	-1.0	0	0	0.7	-0.1	0	0
	(0.1)	(0.4)	--	--	(0.1)	(1.6)	--	--
(3) Estimate Adjusted for Starting Value	5.1	6.0	59.5	62.7	54.4	57.4	22.6	23.0
	(0.5)	(0.4)	(0.5)	(0.9)	(0.6)	(1.6)	(0.3)	(0.4)
(4) Response Quality Adjustment	1.6	1.4	5.6	5.7	4.0	4.6	-3.3	-3.5
	(1.2)	(1.1)	(1.1)	(1.1)	(1.1)	(1.4)	(0.6)	(0.6)
(5) Estimate Adjusted for Starting Value and Response Quality	6.7	7.4	65.1	68.4	58.4	62.0	19.2	19.5
	(1.2)	(1.1)	(1.2)	(1.4)	(1.2)	(1.3)	(0.6)	(0.7)
Number of observations	2939		2939		2939		2939	

Notes: Bootstrapped standard errors in parentheses (based on 10,000 replications). The sample is limited to observations for which both the expected benefits and the certainty equivalent are nonmissing. See text for a description of the methodology of the starting value adjustment. The starting value adjustment is not done for expected benefits and standard deviation of benefits because these two variables were elicited without using a starting value. See text for a description of the response quality adjustment. The variable definitions and means of the variables that are proxies for response quality are listed in Appendix Table A5. The coefficients of the regression of the outcome variables on these measures of response quality are shown in Table A6.

Online Appendices can be found at: http://users.nber.org/~luttmer/polrisk_apx.pdf