American Economic Review 2018, 108(2): 275–307
https://doi.org/10.1257/aer.20151703

The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security†

By Erzo F. P. Luttmer and Andrew A. Samwick*

Policy uncertainty reduces individual welfare when individuals have limited opportunities to mitigate or insure against the resulting consumption fluctuations. We field an original survey to measure the degree of perceived policy uncertainty in Social Security benefits and to estimate the impact of this uncertainty on individual welfare. Our central estimates show that on average individuals are willing to forgo 6 percent of the benefits they are supposed to get under current law to remove the policy uncertainty associated with their future Social Security benefits. This translates to a risk premium from policy uncertainty equal to 10 percent of expected benefits. (*JEL D14, D81, H55*)

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 life cycle and because their investments in human capital will not be privately optimal for the actual realization of the future tax rate.[1] Uncertainty about the generosity of expenditure programs for the elderly, such

*Luttmer: Dartmouth College, 6106 Rockefeller Hall, Hanover, NH 03755, and NBER (email: Erzo.FP.Luttmer@dartmouth.edu); Samwick: Dartmouth College, 6106 Rockefeller Hall, Hanover, NH 03755, and NBER (email: Andrew.Samwick@dartmouth.edu). This paper was accepted to the *AER* under the guidance of Hilary Hoynes, Coeditor. 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 numerous institutions 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. We are grateful to three anonymous referees for thoughtful and detailed comments. We obtained approval for this project from the Institutional Review Boards of the NBER and of Dartmouth College. This research was supported by the US Social Security Administration through grant 5 RRC08098400-03-00 to the NBER as part of the SSA Retirement Research Consortium. The findings and conclusions expressed are solely those of the authors and do not represent the views of SSA, any agency of the Federal Government, or the NBER. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.
† Go to https://doi.org/10.1257/aer.20151703 to visit the article page for additional materials and author disclosure statement(s).

[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. The resulting welfare gain could potentially more than offset the welfare cost of the consumption risk induced by 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.

as Social Security and Medicare, also reduces individuals' ability to consumption smooth. Given that Social Security is mandatory, nondiversifiable, and accounts for more than one-third of total income among the elderly, policy uncertainty regarding its generosity is likely to be one of the major sources of the 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. Throughout the analysis, we focus on future Social Security benefits relative to the benefits scheduled under current law, which abstracts from uncertainty regarding both current benefit rules and the factors that determine benefits under those rules.

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 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, other studies, including Van der Wiel (2008), Giavazzi and McMahon (2012), and Guiso, Jappelli, and Padula (2013), have measured policy uncertainty using survey questions about future policy. Baker, Bloom, and Davis (2016) take yet another approach and create an index of policy uncertainty based on the frequency that the word triplet "policy," "economic," and "uncertain" (or synonyms of these words) appears in newspaper articles.

The papers that proceed to the second step then either relate the policy uncertainty estimated in the first step to observed behavior or use it to calibrate a model that yields a welfare estimate.[4] Skinner (1988) and Dušek (2007) use models to calculate the welfare cost of the estimated uncertainty. Skinner (1988) estimates that the welfare cost of uncertain taxes is 0.4 percent of national income, and Dušek (2007) finds that the risk premium for the uncertainty around the indexing of Social

[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 this is the case, 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 or exacerbating 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.

[4] 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 (2016) relate their indices of policy uncertainty to industry and macro outcomes including stock-price volatility, investment, employment, and output.

Security benefits in the Czech Republic is 1.3 percent when the coefficient of relative risk aversion is assumed to equal 3. Alternatively, a calibrated (rather than estimated) measure of policy uncertainty can be used to estimate welfare, which is the approach taken by Gomes, Kotlikoff, and Viceira (2012). They calculate that early resolution of the uncertainty about future Social Security benefits, holding constant the variance in future Social Security benefits, can lead to welfare gains that are equivalent to 0.5 percent of lifetime consumption. Caliendo, Gorry, and Slavov (2015) also use a calibrated measure of policy uncertainty but allow uncertainty in both the timing and structure of 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.[5]

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 individual's risk premium of the 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 differential 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 risk premium 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. The estimate of the risk premium is the welfare cost to the individual of the policy uncertainty if the individual is rational, accurately answers our survey questions, and does not misperceive the policy uncertainty.[6]

We estimate the risk premium of 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, and the need for reform to restore the program to long-term financial stability has been an ongoing topic of policy discussion.[7] With the status of reform still in doubt, individuals can expect some form of policy

[5] See Kitao (2016) and Nelson (2017) for further analyses of the welfare implications of policy uncertainty around the timing of reform in general equilibrium, overlapping generation frameworks.

[6] Under the additional assumptions that the concerns noted in footnote 2 are not present, the welfare cost to the individual is equal to the social welfare cost of policy uncertainty.

[7] The Social Security's Board of Trustees (2017) projected that the program's trust funds would be exhausted in 2034, at which point annual costs would exceed annual income by 27 percent or 3.61 percentage points of taxable payroll.

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 US 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.[8] 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) to the benefit formula. The key part of the survey consists of two sets of questions. In the first, respondents are asked to describe the 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 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

[8] 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).

percent of benefits under current law, but would become 9.8 percent and 11.8 percent, respectively, if expressed as a percent of expected benefits. In addition, we find 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 randomizations into the survey that can alert us to respondents giving nonmeaningful 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 7.4 percent.

The remainder of the paper is organized as follows. In Section I, we describe our sampling frame and survey instrument and provide summary statistics. In Section II, 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 III. Section IV provides evidence on survey response quality and possible corrections to the distribution of risk premia in light of measurement error. Section V concludes.

I. 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 US Postal Service's Delivery Sequence File.[9] 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. Knowledge Networks collects basic demographic characteristics for all its panelists, and its panelists are

[9] 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 nonresponders 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 website and completing the recruitment information online.

roughly representative of the adult US population according to these characteristics. The burden of panel membership is kept reasonably 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. Our sample contains the results for 3,053 completed interviews conducted between June 10 and July 1, 2011.[10] The median duration of the survey was 20 minutes, and respondents were paid a $5 cash-equivalent incentive to enhance survey completion. Online Appendix Table A1 compares summary statistics for age, gender, race, education, marital status, region, income, and household size from our survey to the Current Population Survey (CPS) from March 2010 (King et al. 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 thus consider our sample as broadly representative of the US population between the ages of 25 and 59.

In the regressions, we control for these demographic characteristics, along with MSA residency, homeownership, employment status, and whether there are kids in the household. Summary statistics are shown in online Appendix Table A2. 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.[11] We ask about risk preferences, life expectancy, the importance of Social Security in retirement, optimism, trust in the political system, and financial literacy. Summary statistics for these control variables are also presented in online Appendix Table A2.

II. Methodology

The main part of our survey gathers information needed to calculate the costs of policy uncertainty. Following a long-standing practice in the field of survey design, we include randomizations in the survey that allow us to gauge whether respondents are able to give meaningful answers. This section discusses three important design features of the survey.

A. *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." Respondents

[10] The invitation to our survey did not mention the topic of the survey. Of the Knowledge Networks panelists receiving the invitation, 63 percent opened the link to the survey. Of those opening the link, 18 percent were disqualified early on because they already were receiving Social Security or because they said they were ineligible to ever receive Social Security. Of the qualified respondents opening the link, 94 percent completed our survey.

[11] We ask these questions at the end of the survey. The full survey instrument is included in online Appendix C. 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, and then include both the recoded variable and the dummy for whether the response was originally missing in our regressions.

may not have a very accurate understanding of how current law determines benefits or may be uncertain about variables, such as future own earnings or future aggregate earnings, that are an input to the benefit calculation. 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 policymakers.

B. *Constructing the Perceived Distribution of Social Security Benefits*

The second feature is to use the visual aspect of the online survey to elicit a probability distribution. This feature was developed in Delavande and Rohwedder (2008) and subsequently used in Liebman and Luttmer (2015). Specifically, to elicit the probability distribution of future benefits, we start by asking the respondent to allocate 20 balls across 4 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.[12] An example of what the survey screen might look like when the respondent has allocated the 20 balls to the 4 bins is

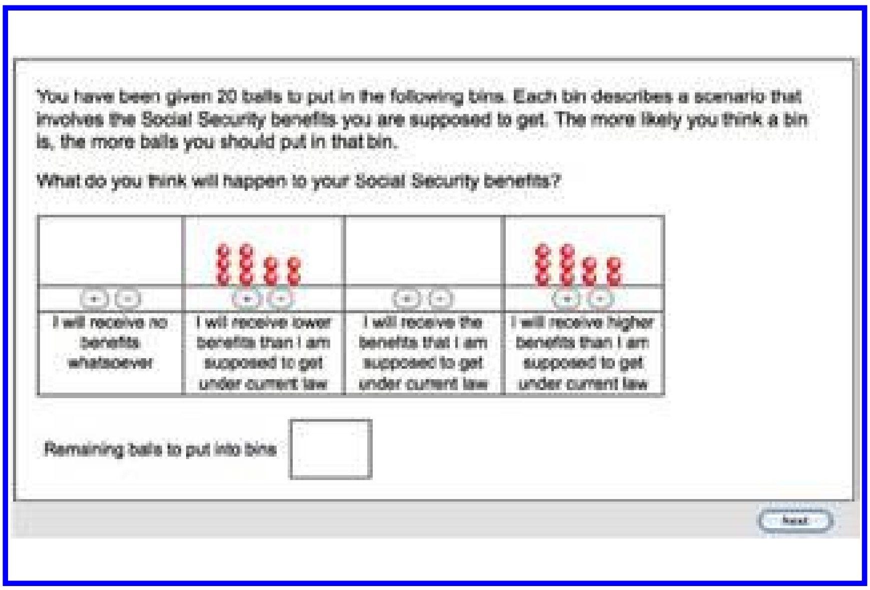

Respondents who put any of these balls in the "lower" or "higher" bins are then asked to further specify which 20 percentage point bins between 1 and 99 percent or

[12] As we discuss in Section III, we elicit this distribution of benefits after respondents answered a series of questions on the financial condition of the Social Security program and the likelihood of various policy changes to the system. The questions on the distribution of future benefits are thus asked in the context of the policy-related reasons why they may get lower benefits than they are supposed to under current law. It is therefore reasonable to assume respondents generally had future policy changes in mind as the reason why they might receive benefits that differ from what they are supposed to get under current law, though we cannot rule out that some respondents interpreted the question differently. See Sections 2 and 3 of the full survey instrument in online Appendix C.

101 and 200 percent should contain these balls. An example of the next screen this respondent will see is

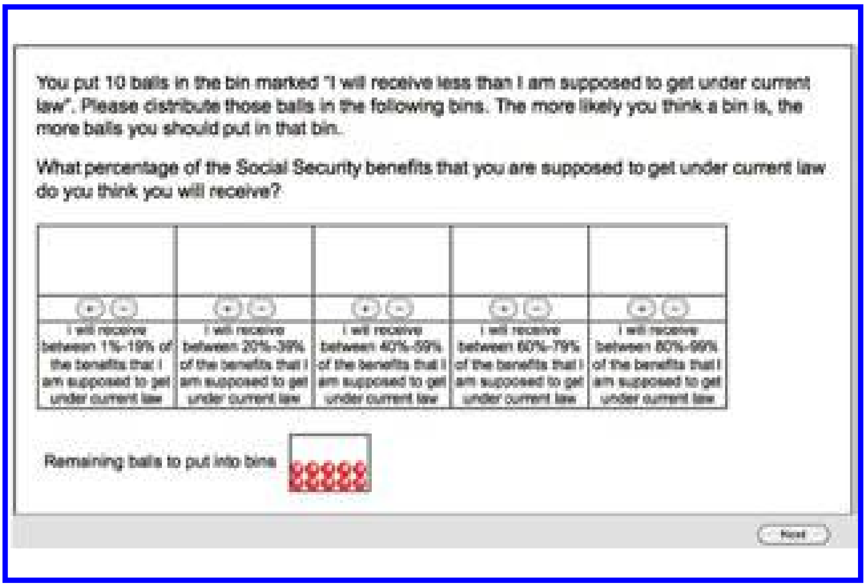

Finally, any bin into which 11 or more balls are placed is further broken down into 5 smaller bins, and respondents are asked to allocate the balls from the larger bin into the smaller bins. An example of the screen that the respondent would have seen in that case is

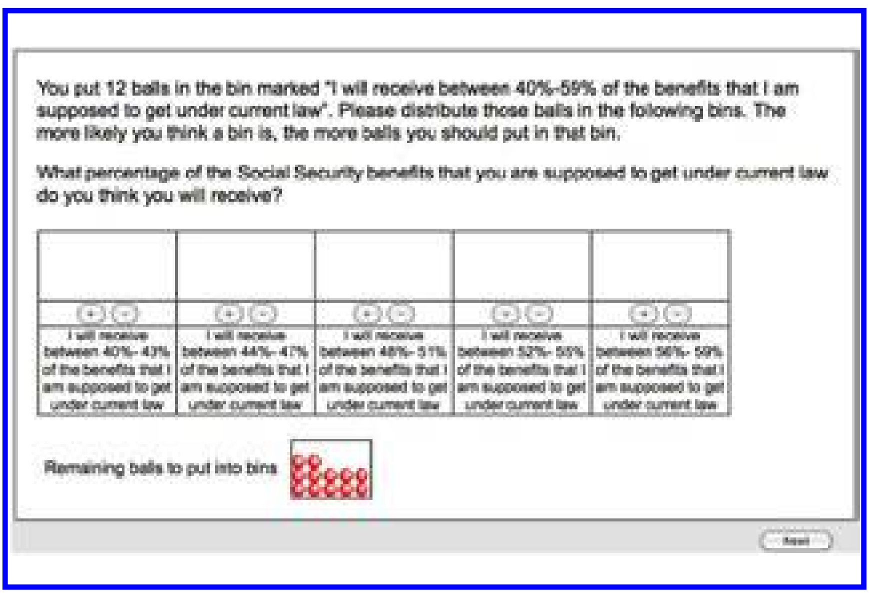

By this three-step process, we obtain the cumulative distribution function (CDF) of uncertain future benefits for each respondent. In order to have greater confidence that respondents will know how to use this tool to express their preferences, we first give an illustration using the weather in Boston.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

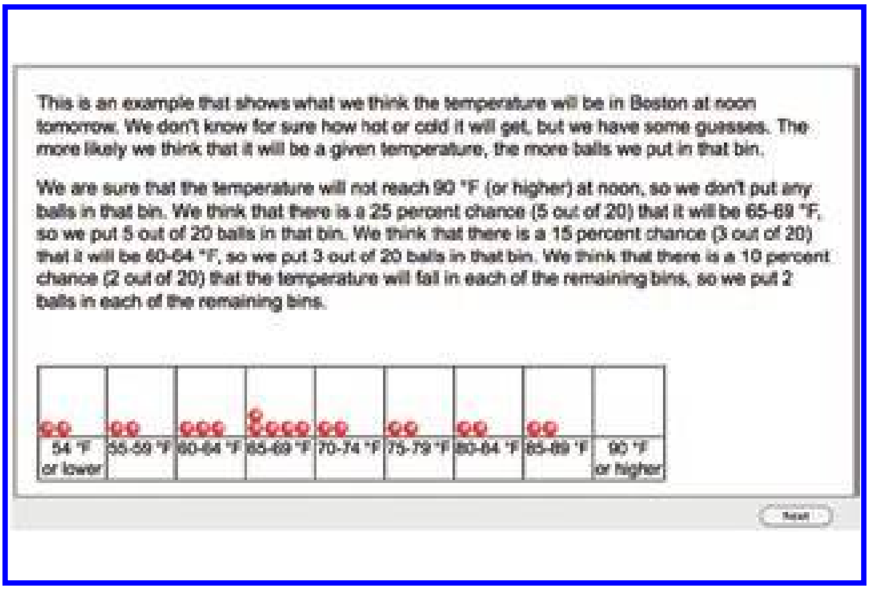

And the narrow distribution is

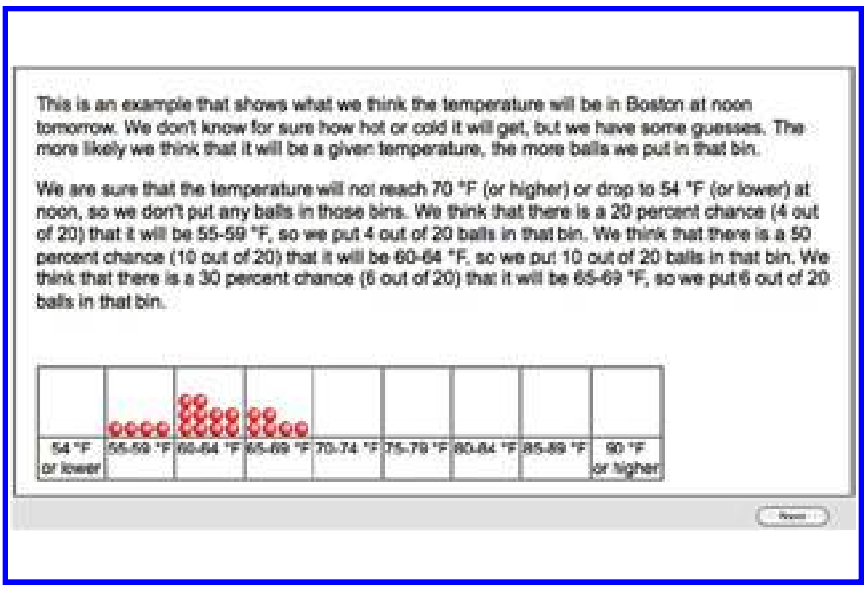

If we had shown no illustration, we could not be sure that respondents would understand the tool well enough to answer the subsequent question. If we had only shown one illustration, then we would have had no way to gauge the size of any bias that our particular choice of illustration may have had on 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.

C. *Obtaining the Certainty-Equivalent Benefit*

The natural metrics to quantify 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. Because respondents might have trouble coming up with a sensible answer if we asked for it directly, we adopt a common method in the contingent-valuation literature, which is to offer the respondent a sequence of binary choices that allow us to bracket the respondent's certainty equivalent to the perceived distribution of benefits described in Section IIB. The survey calculates the expected value (denoted 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 five.[13]

[13] The full sequence of offers that the respondents receive is shown in Question 4.3 of the survey instrument in online Appendix C. For respondents whose certainty equivalent was close to their expected value, we asked the certainty equivalent question an additional time to narrow the interval between Y_n and Y_m below 5. For the 7.4 percent of respondents who provide distributions that show no uncertainty, we ask a slightly different version of Question 4.3, also detailed in online Appendix C, in which the sequencing of the offers of Y% is the same as in the original question, but the question makes no mention of an uncertain distribution.

Table 1—Summary Statistics of Respondents' Estimates of Future Social Security Benefits

	Mean (1)	Standard deviation (2)	25th percentile (3)	Median (4)	75th percentile (5)	Number of observations (6)
Panel A. Respondent confidence in Social Security						
Very confident	0.033	—	—	—	—	3,045
Somewhat confident	0.223	—	—	—	—	3,045
Not too confident	0.453	—	—	—	—	3,045
Not at all confident	0.291	—	—	—	—	3,045
Panel B. Point estimate of future benefits						
Future benefits measured using a slider	65.9	31.8	41.0	70.0	100.0	2,956
Panel C. Key outcome variables						
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

Notes: The sample is restricted to individuals between the ages of 25 and 59 as of May 2011. See online Appendix C for exact question definitions. Panel A reports responses to Q1.2 and panel B reports responses to Q3.2. The construction of the variables in panel C is discussed in the text and based on responses to the balls/bins question that elicits the subjective distribution of future Social Security benefits (Q3.3–Q3.6) and the certainty equivalent question (Q4.3). All values in panels B and C are in percent of benefits under current law.

Source: Data from the June 2011 Social Security Policy Risk Survey, designed by the authors and fielded by Knowledge Networks

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.

III. Results

A. *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 perceived risk. As shown in panel A of Table 1, 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 one-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 the 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 A3 provides comparisons of the responses to this question in our sample and the subsample of their respondents aged 25–59.[14]

[14] We are indebted to Mathew Greenwald for providing these tabulations. The tabulations of the Knowledge Networks panel in online Appendix Table A3 pertain to the respondents who answered both the balls/bins questions and the certainty equivalent questions, as described in Section II.

In their sample, 9.8 percent were very confident and 33.4 percent were somewhat confident. Together, 43 percent express confidence in Social Security in the Greenwald et al. sample, compared to 25 percent in the Knowledge Networks panel. Of the remaining 57 percent, 37.4 percent are not too confident and 19.3 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.

About 91 percent of respondents in our sample are aware that Social Security faces a projected financial shortfall. About 58 percent expect the shortfall to be addressed by a combination of tax increases and benefit reductions. About three-quarters believe the shortfall will be addressed at least in part through benefit cuts. 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.[15] When respondents are asked to use a slider to provide 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, as shown in panel B of Table 1.

B. *The Perceived Distribution of Future Benefits*

The responses to the general questions above 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 1-in-6 chance of receiving no benefits whatsoever and about a 1-in-4 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 4 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. On average, respondents expect to receive 59.4 percent of the benefits they are supposed to get under current law. There is very little probability mass at 0, 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 panel C of Table 1.

[15] These figures are shown in online Appendix Table A4, which also tabulates the responses to more detailed survey questions about expected benefit and tax changes.

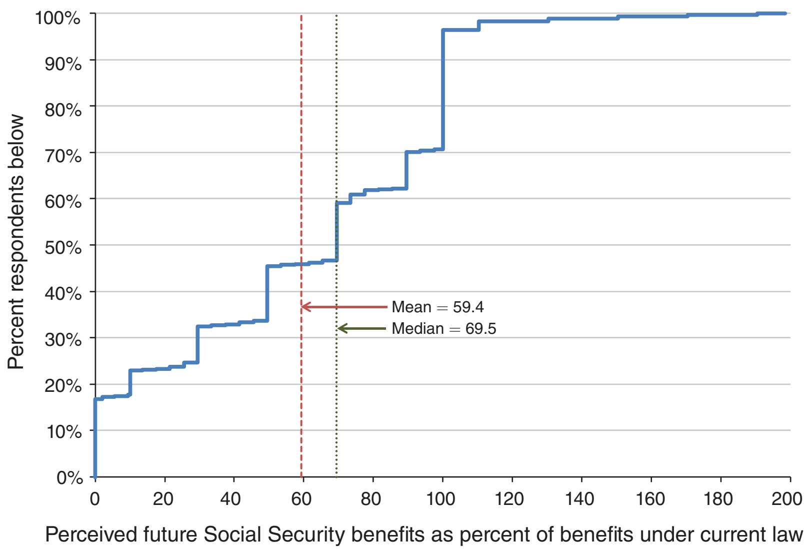

Figure 1. Average Perceived CDF of Future Social Security Benefits

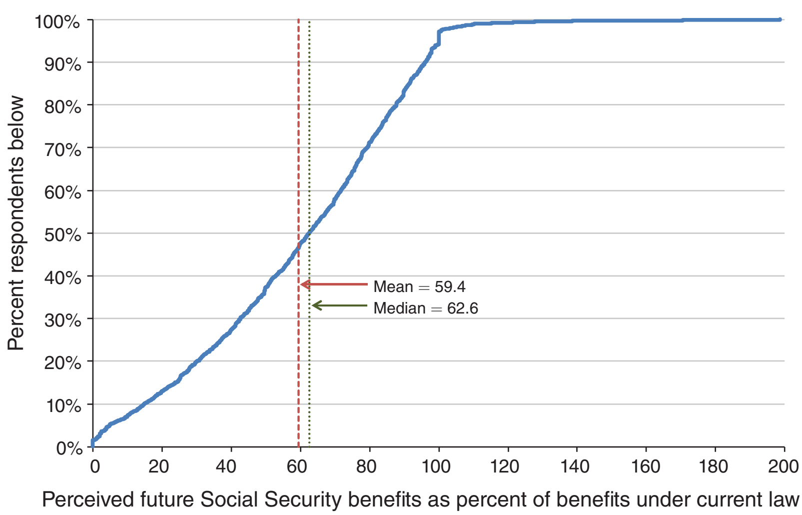

Figure 2. CDF of Perceived Mean Future Social Security Benefits

We use the expectation of the subjective probability distribution of future Social Security benefits in panel C, rather than the point estimate in panel B, 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 a mean. Second, the expectation of subjective benefits better predicts confidence in Social Security (as measured by the multiple-choice question in panel A) than the point estimate

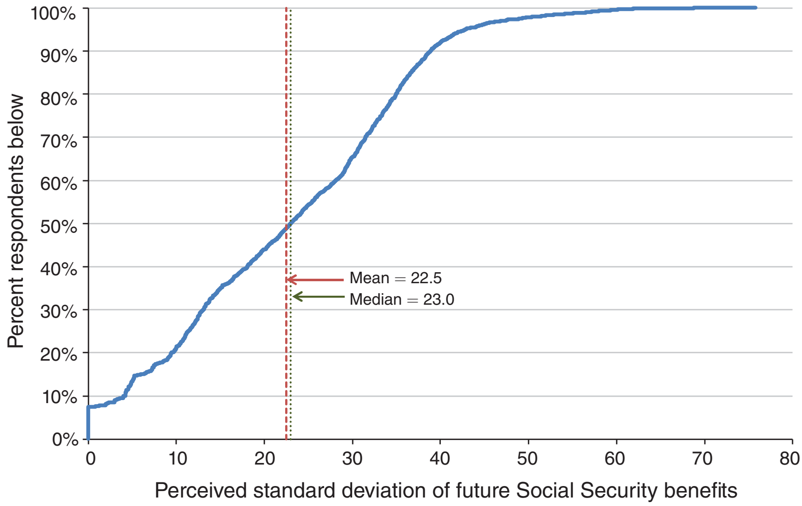

Figure 3. CDF of Perceived Standard Deviation of Future Social Security Benefits

is able to predict confidence in Social Security. This suggests that the subjective expectation has less measurement error than the point estimate.[16]

Figure 3 shows the CDF of the standard deviations of respondent CDFs. Only 7.4 percent have a standard deviation of zero. The second row of panel C of Table 1 provides summary statistics, indicating mean and median values of about 23 percent, with one-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.

C. *The Certainty Equivalent of Social Security Benefits*

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?" Only 20.5 percent respond that the uncertainty matters little or does not matter, whereas 32 percent respond that it matters somewhat and 47.5 percent respond that it matters very much.[17]

More formally, 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 panel C of Table 1, denominated as a percentage of the benefits the respondents are supposed to get

[16] The correlations between the expectation of the subjective probability distribution and the point estimate and multiple-choice question on confidence are 0.69 and 0.54, respectively. That the correlations are not unity suggests the presence of measurement error. We estimate a model to correct for mean-zero measurement error in Section IVC.

[17] These figures are shown in panel F of online Appendix Table A4.

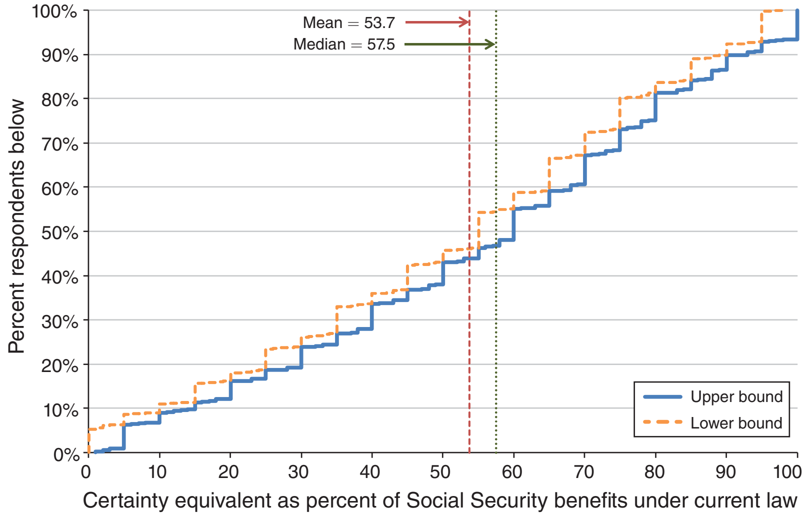

Figure 4. CDFs of Certainty Equivalent of Future Social Security Benefits

under current law. The mean certainty equivalent is 53.7 percent and the median is 57.5 percent.[18]

D. *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 panel C 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 measured 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 measured risk premia is shown in Figure 5.[19] About 11 percent of respondents have measured risk premia less than negative 20 percent. At the other

[18] The certainty-equivalent question was framed as one where the individual makes a personal choice between regular future Social Security benefits (which are uncertain) and the guaranteed contract. If the respondent took the question as a personal choice only applying to herself, then there are negligible aggregate budget implications of this choice, and it would not affect her future tax rates. However, some respondents seem to have interpreted the question as them "voting" on a choice between regular benefits and guaranteed benefits. In other words, they may have assumed that their choice would be applied to all people (if their vote is pivotal), in which case the aggregate budget implications of the choice matter. In online Appendix D, we discuss how this interpretation could lead to bias in our estimates of the average certainty equivalent and risk premium. We also provide a rough estimate of this bias, which indicates that the bias in the risk premium would be less than one-half of a percentage point.

[19] If there is classical measurement error, the distribution of risk premia in Figure 5 is more spread out than the true underlying distribution of risk premia. We implement a correction for classical measurement error in online Appendix A and show the measurement-error corrected distribution of risk premia in online Appendix Figure A2.

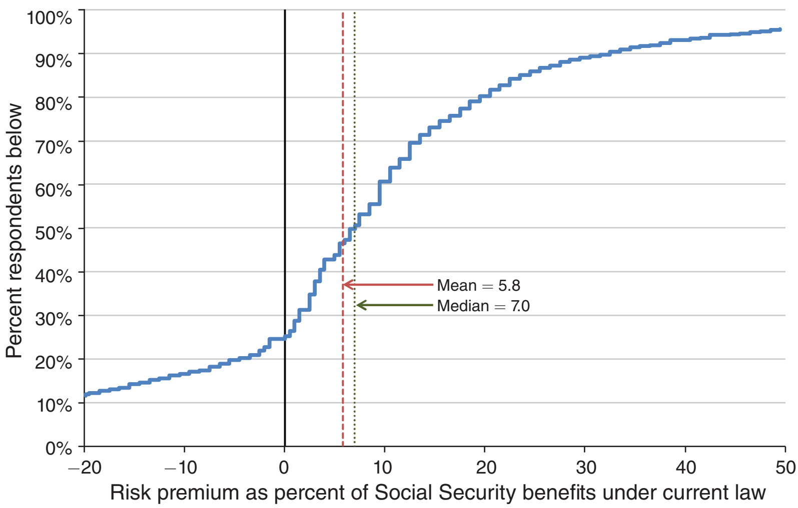

Figure 5. Distribution of Risk Premia

end of the distribution, 25 percent of respondents have measured risk premia of 16.5 percent or more, with 4 percent having one 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 tenth percentile or above the ninetieth percentile, the mean risk premium becomes 6.9 percent. Similarly, winsorizing at the tenth and ninetieth percentiles yields a mean risk premium of 6.3 percent. Truncating at the twenty-fifth and seventy-fifth percentiles increases the mean to 7.3 percent, while winsorizing at these percentiles yields a mean risk premium of 7.7 percent.[20]

The estimate of the risk premium is the welfare cost to the individual of the policy uncertainty if the individual is rational, accurately answers our survey questions, and does not misperceive policy uncertainty. In Section IV, we examine response quality, test implications of rationality, and adjust the estimates for a particular deviation from rationality (anchoring) and for response quality. We have no direct way of testing whether perceptions of policy uncertainty are accurate. Predictions about possible future Social Security policies and their likelihoods are necessarily subjective. However, we can provide evidence on the accuracy of two perceptions for which objective estimates exist, namely life expectancy and benefits under current law. In our sample, the median and mean subjective probabilities of surviving to age 75 are 71 and 67.9 percent, compared to actuarial probabilities of 71.5 and 74.1 percent, respectively. The mean and median estimates among our respondents

[20] Respondents who have missing benefit expectations or distributions that have no uncertainty were asked an alternative version of the certainty equivalence questions. These respondents tend 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. If we impose a risk premium of zero on those who reported perceiving no uncertainty (rather than calculate their risk premium based on the alternative version of the certainty-equivalent question), we obtain a mean risk premium of 6.1 percent and the median risk premium of 5.5 percent.

of the benefits they are supposed to get under current law are \$1,237 and \$1,100, respectively. These are each within 10 percent of the mean and median value of \$1,175 in Social Security administrative data for retired workers at the time of the survey.[21] The accuracy, on average, of perceptions on which we have objective measures increases our confidence in the reported perceptions of policy risk.

Our main estimate of the risk premium is based on our method of eliciting a certainty equivalent through a series of pairwise choices regarding a guaranteed contract 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. The drawback to the method is that it requires respondents 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 calculate the risk premium for each respondent from the self-reported distribution of possible Social Security benefits by making assumptions about risk aversion and other elements of the budget set. Specifically, we assume constant relative risk aversion preferences with coefficients of relative risk aversion equal to 1, 3, and 5 and incorporate the information from the variable that captures how important the respondent expects Social Security to be in financing retirement spending.[22]

By construction, the distributions of these 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 the risk premia is intermediate between the hypothetical CDFs that would obtain if all respondents had coefficients of relative risk aversion between 3 and 5. Thus, the two different methodologies for estimating the risk premium obtain broadly similar estimates when using values for the coefficient of relative risk aversion that lie within the range of values typically used in the public economics literature.

Independent of functional form and exact modeling assumptions, utility-based models of the risk premium would predict that, all else equal, the risk premium is increasing in the amount of uncertainty; the risk premium is increasing in risk aversion; and the risk premium is increasing in the fraction of retirement consumption that depends on Social Security. We test these predictions in Table 2. The regression

[21] Our respondents report their subjective probability of being alive at age 75 in Question 6.9, listed in online Appendix C. We assign to each respondent the gender-, age-, and cohort-specific actuarial probability of surviving to age 75 using the life tables presented in Bell and Miller (2005). Our respondents report their estimate of the benefits they are supposed to get under current law in Question 6.8, listed in online Appendix C. We compare these estimates to benefit payouts for all retired workers over age 62 reported in Table 5.B9 of the Annual Statistical Supplement to the Social Security Bulletin, 2011. See Social Security Administration (2012).

[22] Specifically, suppose that the respondent's Social Security benefits will be 100. The four responses to the survey question on the importance of Social Security (see online Appendix C, Question 6.10) 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 assign other income of 200, 67, and 17 for the 25–50, 50–75, and 75–100 intervals, respectively. For the 0–25 interval, we choose a value of 500 (consistent with Social Security funding 17 percent).

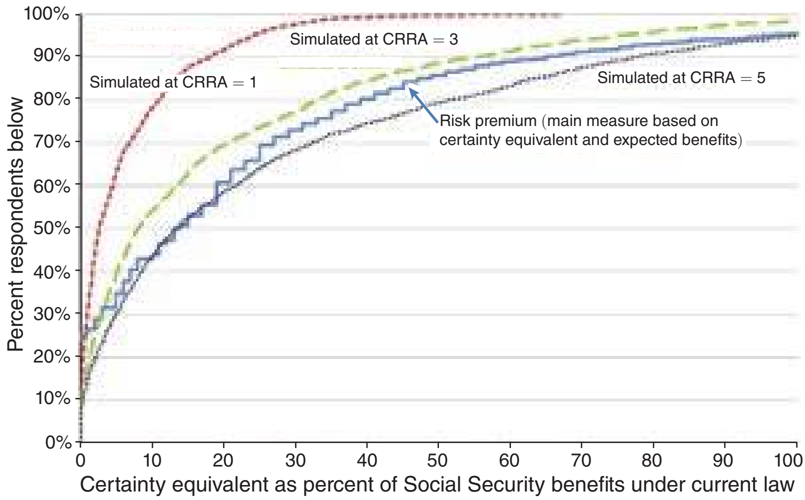

Figure 6. Distribution of Risk Premia Based on Simulated Risk Aversion (CRRA)

Table 2—Theoretical Predictors of the Risk Premium (*Percentage Points*)

	Risk premium of Social Security benefits					
	(1)		(2)		(3)	
Perceived standard deviation of Social Security benefits (percentage points)	0.35	(0.05)	0.35	(0.05)	0.38	(0.05)
Risk aversion index (1–6)			2.07	(0.43)	1.58	(0.44)
Indicator variables for relative risk aversion						
Relative risk aversion < 0.5	Omitted					
0.5 ≤ Relative risk aversion < 1	1.43	(4.06)				
1 ≤ Relative risk aversion < 2	4.80	(3.60)				
2 ≤ Relative risk aversion < 4	10.14	(3.47)				
4 ≤ Relative risk aversion < 8	10.29	(3.42)				
8 ≤ Relative risk aversion	10.78	(3.48)				
Index of importance of Social Security to retirement spending (1–4)			1.66	(0.53)	0.79	(0.60)
Indicator variables for importance of Social Security in retirement spending						
Not so important: Social Security would pay for less than 25% of spending	Omitted					
Important: Social Security would pay for 25% to 50% of spending	4.00	(1.74)				
Very important: Social Security would pay for 50% to 75% of spending	6.31	(1.76)				
Extremely important: Social Security would pay for more than 75% of spending	5.91	(1.82)				
Demographic controls	No		No		Yes	
R^2	0.0499		0.0457		0.0906	
Observations	2,939		2,939		2,939	

Notes: Robust standard errors in parentheses. 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. The standard deviation of Social Security benefits is based on the balls/bins question that elicits the subjective distribution of future Social Security benefits (Q3.3–Q3.6) and is expressed as a percent of benefits under current law. 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 four-point scale from "not so important" to "extremely important" (Q6.10). The demographic control variables in column 3 are listed under the heading "demographic control variables" in Table 4.

Source: Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks

in column 1 shows that a higher standard deviation of perceived future benefits is strongly predictive of a higher risk premium, with each percentage point increase in the standard deviation increasing the risk premium by 0.35 percentage points. The risk premium increases monotonically in the respondent's degree of risk aversion as measured by the respondent's choices over hypothetical job offers with different degrees of risk. Finally, the risk premium increases in the importance of Social Security in retirement spending. In column 2, risk aversion and the importance of Social Security are entered as indices rather than as a set of dummies in order to formally show that both variables positively and statistically significantly predict the risk premium, as predicted by theory. Column 3 shows that the standard deviation of Social Security benefits and the risk aversion index remain significant predictors of the risk premium after we add demographic control variables but that the importance of Social Security benefits is no longer significant.

E. *Heterogeneity in the Perceived Distribution of Benefits and the Risk Premium*

Table 3 shows the heterogeneity across demographic groups in the characteristics of the perceived distributions of Social Security benefits and the risk premia. The age of the respondent is an important source of heterogeneity. 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. This pattern fits with a perception that a reform would likely grandfather older people. The second column shows risk premia by demographic characteristics. There is a clear difference in risk premia between those over 50 and those under 50. The former have risk premia of 10 percent or more while the latter have risk premia of 5 percent or less. An increase in risk premia is consistent with older people having less opportunity to adjust future labor supply or savings in response to a change in Social Security benefits, which increases the cost of uncertainty to them. For both the expected benefits and perceived risk premia, the differences across age groups are jointly significant.

The remaining panels of Table 3 show the analogous comparisons for groups based on income, race/ethnicity, gender, education, and marital status. In terms of expected benefits, the general pattern that emerges in these univariate comparisons is that those for whom Social Security is likely to be more important report expecting to receive relatively high fractions of the benefits they are supposed to get under current law. In particular, benefit expectations are relatively high for lower-income individuals, for Blacks and Hispanics, for those with lower levels of education, and for widowed individuals. The main exception to this general pattern is that men have higher benefit expectations than women. For the risk premia, a similar general pattern emerges. Risk premia are markedly higher for lower-income groups, for Blacks and Hispanics, for women, for those with less education, and for widowed and separated individuals. These groups tend to have fewer other sources of retirement wealth and are therefore likely more vulnerable to changes in Social Security, which would be consistent with their relatively high risk premia. The last panel in the table shows the differences across terciles of response quality (to be described in more detail in Section IV). Respondents who give the highest quality responses tend to have significantly higher expected benefits but not significantly higher risk premia.

Table 3—Expected Benefits and Perceived Risk Premium by Demographic Groups

	Expected benefits			Perceived risk premium		
Demographic group	Mean		*p*-value	Mean		*p*-value
Age groups			0.000			0.000
25–29	49.4	(1.4)		2.7	(1.6)	
30–34	49.9	(1.6)		3.7	(1.6)	
35–39	51.4	(1.5)		2.8	(1.5)	
40–44	56.0	(1.3)		5.0	(1.2)	
45–49	59.0	(1.6)		3.8	(1.6)	
50–54	67.7	(1.2)		9.9	(0.9)	
55–59	79.4	(1.2)		11.7	(1.1)	
Income groups			0.029			0.000
< \$25k	63.9	(1.6)		9.8	(1.7)	
\$25k–\$50k	59.1	(1.2)		9.2	(1.1)	
\$50k–\$75k	58.3	(1.2)		3.5	(1.1)	
\$75k–\$100k	57.4	(1.4)		5.1	(1.2)	
\$100k or more	59.5	(1.0)		3.1	(0.9)	
Race/ethnic groups			0.001			0.000
Non-Hispanic white	58.4	(0.6)		4.0	(0.6)	
Non-Hispanic black	66.4	(1.9)		16.7	(2.0)	
Hispanic	60.4	(1.5)		9.0	(1.6)	
Other	56.3	(2.9)		–0.9	(2.7)	
Gender groups			0.022			0.000
Female	58.0	(0.8)		7.8	(0.8)	
Male	60.6	(0.7)		4.1	(0.7)	
Education groups			0.308			0.001
High school dropout	62.2	(2.1)		12.4	(2.2)	
High school graduate	60.2	(1.1)		7.2	(1.1)	
Some college	58.8	(1.2)		4.7	(1.1)	
College graduate	58.6	(0.8)		4.1	(0.7)	
Marital status groups			0.000			0.001
Married	58.4	(0.7)		4.6	(0.6)	
Widowed	77.3	(4.1)		16.1	(4.0)	
Divorced	65.2	(1.9)		7.0	(1.8)	
Separated	64.7	(4.0)		15.9	(3.5)	
Never married	60.3	(1.4)		7.3	(1.4)	
Lives with partner	56.9	(1.8)		7.2	(2.0)	
Response-quality terciles			0.000			0.517
Lowest tercile	57.7	(1.1)		6.8	(1.2)	
Middle tercile	56.2	(0.9)		5.1	(0.9)	
Highest tercile	64.0	(0.9)		5.6	(0.6)	

Notes: Robust standard errors in parentheses. Expected benefits are based on the balls/bins question that elicits the subjective distribution of future Social Security benefits (Q3.3–Q3.6), and are expressed as a percentage of benefits under current law. 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. The *p*-value refers to the test that outcome variable is the same in the demographic subgroups listed. Response-quality terciles are based on the first principal component of the response-quality indicators listed in online Appendix Table A7.

Source: Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks

In Table 4, we use a regression framework to describe heterogeneity by demographic and other respondent characteristics. The columns show regressions with expected benefits, the standard deviation of benefits, and the risk premium as the dependent variable. Each regression includes both the demographic variables from

Table 4—Correlates of Expected Benefits, Standard Deviation of Benefits, and the Risk Premium

	Expected benefits		Standard deviation of benefits		Risk premium	
Demographic control variables						
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)
High school 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	
Observations	2,960		2,960		2,939	

Notes: Robust standard errors in parentheses. Each column reports a separate OLS regression. Expected benefits and standard deviation of benefits are based on the balls/bins question that elicits 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. Summary statistics of the explanatory variables are provided in online Appendix Table A2. The risk aversion index is defined in the note to Table 2. Importance of Social Security funds during retirement is measured on a four-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 nonmissing 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–Q6.16). Missing values of explanatory variables are dummied out.

Source: Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks

the Knowledge Networks panel and the other control variables for preferences and beliefs that we ask in our survey.[23]

[23] Estimates that exclude the other control variables for preferences and beliefs are similar and shown in online Appendix Table A5.

Focusing on the regression for expected benefits, we find the same general pattern as we found in the univariate comparisons. As before, age is a strong predictor of expected benefits, with each additional year of age leading to a 0.94 percentage point increase in expected benefits. Perceived uncertainty decreases with age, with each year of age decreasing the standard deviation of benefits by 0.21 percentage points. These estimates are consistent with political rhetoric on grandfathering in 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. Another group that is more likely to be grandfathered consists of those who are already retired. The effect of being retired on expected benefits is large and significant: equivalent to the effect of ten years of age. 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 estimate for the effect on the standard deviation is negative but significant only at the 10 percent level. Another result that is consistent with political rhetoric surrounding Social Security reform, in which potential benefit cuts relative to current law are conjectured to be "progressive," is the relation with income.[24] A 10 percent increase in income is associated with a 0.31 percentage point decline in expected benefits and a 0.09 percentage point reduction in the standard deviation of benefits.

The third column of Table 4 presents a regression in which the dependent variable is the risk premium. 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. Like the univariate comparisons, the risk premium is decreasing in educational attainment, but this effect is no longer significant in the regression framework, which includes other proxies for lifetime income. We would expect to see higher risk premia among those who are more vulnerable to risk, either because they have fewer other sources of retirement wealth or they have fewer opportunities to mitigate the benefit uncertainty by adjusting future labor supply or savings. In line with this prediction, we find age and being retired, both of which indicate fewer opportunities to mitigate the benefit uncertainty, are associated with higher risk premia. Apparently, the greater difficulty of mitigating benefit uncertainty outweighs the smaller perceived uncertainty for these groups. Also in line with the prediction above, we find larger risk premia for those that are likely to rely more on Social Security as a source of retirement income: those with lower levels of income, Blacks, Hispanics, and those with greater life expectancy. The self-reported indicator for reliance on Social Security is positive but not significant. As theory predicts and as we already saw in Table 2, the risk premium is increasing in risk aversion. There is no clear theoretical prediction for our finding of a positive association between trust in elected officials and the risk premium.

IV. Response Quality and Measurement Error Corrections

Recognizing that our survey asks questions that may be challenging for some respondents to answer, we expect measurement error in the reported answers. The

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

large tails in the distribution of measured risk premia, as shown in Figure 5, could very well reflect measurement error. We address the presence of measurement error in two steps. First, in the next subsection, we show that the survey responses to the certainty equivalent questions are strongly correlated with the variables that should predict them and, with one exception, unaffected by the design features of the survey needed to obtain meaningful responses. The exception is that the starting value in the sequence of questions that determine the respondent's certainty equivalent has an effect on the resulting value.

Second, in the remaining subsections, we consider three possible adjustments to our estimated risk premia to correct for measurement error. The first employs a measurement-error model that corrects the distribution of risk premia for mean-zero measurement error. The other two corrections deal with measurement error that has a nonzero mean. One of these pertains to the starting value concern noted above. The other considers the possibility that some respondents give systematically biased answers because they are not able to fully understand the questions. This method assesses respondents' response quality based on their responses to questions unrelated to estimating the risk premium. 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.[25] The full list of measures, along with sample means, is shown in online Appendix Table A7. We create an index of response quality by taking the first principal component of the 15 response-quality indicators.

A. *Cross-Validation of the Certainty Equivalent*

In Table 5, 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. 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, which includes the demographic and other control variables from Table 4, 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

[25] For example, the first measure is whether the respondent reported a positive probability of not living to age 75, an age 15 years in the future for our oldest respondent. A respondent who instead gave 100 percent as the likelihood of being alive at age 75 may not fully understand the concept of a probability. The second measure is whether the absolute difference between the individual's subjective probability of living to age 75 and the actuarial probability from Bell and Miller (2005) based on the respondent's age, cohort, and gender is below the median of 15 percentage points. 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.

Table 5—Are Responses on the Certainty Equivalent Meaningful?

Dependent variable: certainty equivalent of Social Security benefits	(1)		(2)	
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 percent	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	
Observations	2,939		2,939	

Notes: Robust standard errors in parentheses. Demographic and other controls are the set of controls used in Table 4. 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 balls/bins question (Q3.3–Q3.6). The risk aversion index is defined in the note to Table 2. Starting value is 70 percent is a dummy variable for the respondent initially being offered a guaranteed contract of 70 percent (rather than 30 percent) of benefits under current law. Guaranteed benefits is second option: a dummy variable for the guaranteed contract being listed second rather than first. The weather example is an example of a probability distribution using the balls/bins format that was presented to the respondent prior to Q3.3. Respondent sees high standard deviation weather example is a dummy variable that equals 1 if the variance of the distribution in the example was high.

Source: Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks

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 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. The order in which the guaranteed benefits are presented and whether the respondent sees the wide weather example have no statistically significant effects on the certainty equivalent.[26] 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 results from the sequence of questions is about 8 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 present a correction for this bias in Section IVD.

[26] In online Appendix Table A6, we present regression results that show that the "wide" weather example has a statistically significant effect on the standard deviation of the perceived distribution of benefits but not its expected value. Because the estimate of the risk premium based on our methodology relies on the mean but not the reported dispersion of this distribution, and because the certainty equivalent is not significantly affected by the weather example, our estimate of the risk premium is not affected by the sensitivity of the standard deviation to the weather example. In panel C of online Appendix Table A6, we find no effect of the weather example on the risk premium—the point estimates are less than 0.2 in absolute value and statistically insignificant.

B. *Measurement Error and the Identification of Risk Premia*

In this subsection, we first formally spell out conditions on the types of measurement error under which the mean risk premium is identified and conditions under which the distribution of risk premia is identified. We define the reported answer to the balls/bins question B to equal the underlying benefit expectation X plus measurement error ε_b:

$$B = X + \varepsilon_b \,. \tag{1}$$

Similarly, we define the reported certainty equivalent C to be the underlying certainty equivalent Y plus measurement error ε_c:

$$C = Y + \varepsilon_c \,. \tag{2}$$

By definition, the risk premium is the difference between the expected benefits and the certainty equivalent. Thus, the reported risk premium, R, is equal to the reported expected benefits minus the reported certainty equivalent:

$$R = B - C = (X - Y) + (\varepsilon_b - \varepsilon_c) = Z + (\varepsilon_b - \varepsilon_c), \tag{3}$$

where $Z \equiv X - Y$ denotes the underlying risk premium.

Equation (3) makes clear that a sufficient condition for the distribution of R and Z to be equal is for the distributions of measurement error in ε_b and ε_c to be equal and perfectly correlated.[27] Clearly, these assumptions on the distributions of ε_b and ε_c are unrealistically strong. However, as we show in Section IVC, we can recover the distribution of underlying risk premia under weaker assumptions on measurement error.

Taking the expectation of equation (3) shows that the mean reported risk premium is equal to the mean underlying risk premium if the mean of measurement error in the balls/bins question is equal to the mean in measurement error in the certainty equivalent question, i.e., if $E[\varepsilon_b] = E[\varepsilon_c]$. In Sections IVD and IVE, we will correct our estimate of the mean risk premium for violations of this condition.

C. *Correction for Mean-Zero Measurement Error*

In this subsection, we estimate the variance of measurement error in expected benefits, as measured by the balls/bins question and use this estimate to remove mean-zero measurement error from our estimated distribution of risk premia. This measurement-error correction model is an adaptation of standard latent factor models to our setting and its details are described in online Appendix A. The model uses the three measures of expected benefits from Table 1: one based on the balls/bins question, another one based on the slider question, and a third one based

[27] Note that this is not a necessary condition. For example, if distributions are Normal, then any distribution of $\varepsilon_b - \varepsilon_c$ that is negatively correlated with Z with correlation coefficient θ and has a variance equal to $4\theta^2 \operatorname{Var}(Z)$ also leads to a distribution of R that is identical to the distribution of Z.

on a qualitative question about expected benefits (the question from Greenwald et al. 2010, Q1.2). Under the assumption that the measurement error in these three questions is uncorrelated, we solve for the variance of measurement error in each of these three questions using the observed variances and covariances of these three questions. Under the additional assumption that the fraction of measurement error in the certainty-equivalent question is the same as the fraction of measurement error in the balls/bins question, we calculate the fraction of measurement error in the risk premium for any given correlation between measurement errors in the balls/bins question and the certainty-equivalent question.

The estimates from the measurement-error correction model are shown in online Appendix Table A8. The model estimates that 23.9 percent of the variance in the balls/bins question comes from measurement error and that the corresponding fraction is 35.4 percent for the slider question. In short, the balls/bins question provides a more precise elicitation of expected benefits. To correct the risk premium for measurement error, we need to know the correlation between measurement error in expected benefits (from the balls/bins question) and the measurement error in the certainty equivalent. Absent information to estimate this correlation, we make what we consider the most natural assumption, namely that this measurement error is uncorrelated.[28] Under this assumption, the variance of the risk premium net of measurement error is 48 percent of the uncorrected variance of the risk premium, indicating that about one-half of the variance in the observed risk premium is estimated to stem from measurement error. Removing the measurement error reduces the standard deviation in the risk premium from 28 percentage points to 19 percentage points, and the fraction of negative risk premia drops from 24.6 percent to 21.0 percent. These results indicate that, while measurement error is substantial, there is significant variability in underlying risk premia and the fraction of negative risk premia only drops moderately.

To see how far we could drive down the fraction of negative risk premia, we also estimated the model setting the correlation between measurement error in the certainty equivalent and measurement error in expected benefits to its lowest possible value (-0.91) that is consistent with all variances in the model remaining nonnegative. Under this assumption, almost all of the variance in the risk premium is attributed to measurement error and only 1.5 percent of observations have a negative measurement-error corrected risk premium. Thus, this 1.5 percent of negative risk premia is a lower bound on the fraction of people with a negative risk premium, though we emphasize this bound is realized only under assumptions (a correlation of -0.91) that are quite stylized. Nevertheless, it serves as an upper bound on what we can explain by measurement error using this model.

We also estimate the measurement-error correction model separately for three terciles of response quality. While the uncorrected variance of the risk premium clearly falls as we move to the highest response-quality terciles, the measurement-error corrected variance of the risk premium exhibits a nonmonotonic pattern. The

[28] To be clear, the underlying values of the expected benefits and the certainty equivalent are strongly positively correlated, both as a matter of theory and in our estimates. But the mean-zero measurement error is the noise that is introduced because our survey may not perfectly elicit these underlying values. Given the very different ways that we elicit those values (a balls/bins distribution for one and a series of bracketing questions for the other), there is no particular reason why the errors would be correlated.

fraction of observations with a negative measurement-error corrected risk premium lies between 20.9 percent and 24.9 percent for all three response-quality terciles.[29]

For our baseline assumption of uncorrelated measurement error in expected benefits and the certainty equivalent, we estimate that a substantial fraction of respondents has a negative risk premium, even after correcting for measurement error and even after limiting the sample to the third of respondents with the highest indicators for response quality. To get a sense of whether some respondents may truly have a negative risk premium, we examined responses to an open-ended question that asked respondents for the reason for their decision if they rejected a guaranteed benefit amount that was at least 5 percentage points higher than their uncertain expected benefits. Details on the exact question asked and the classification of responses are provided in online Appendix Tables A11 and A12. The responses indicate that 35 percent of respondents to the open-ended question gave a motivation that indicates they indeed prefer the risky option over guaranteed benefits. In short, their negative risk premium does not appear to be measurement error. About 6 percent of respondents to the open-ended question indicated that they do not trust the guarantee. The lack of trust in the guarantee means that in order to be chosen, the guarantee must be larger than one that is trusted. This biases the certainty equivalent upward, and, hence, the risk premium is biased down. For these people, the true risk premium is less negative than what we inferred. Finally, 59 percent of respondents to the open-ended question gave a motivation that indicated their answer was based on a different motivation than just a personal trade-off between risk and return. We think of these answers as measurement error. In short, the open-ended question indicates that a substantial fraction (around one-third) of respondents with a negative measured risk premium truly does not appear to be risk averse in this context, which is consistent with Prospect Theory's prediction of risk-loving behavior in a loss domain (Kahneman and Tversky 1979). However, the majority of negative measured risk premia appears to be due to some form of measurement error.

D. *Corrections to the Mean Risk Premium*

We next consider adjustments that affect the mean risk premium. The first such adjustment is based on the simple linear model used by Herriges and Shogren (1996) in which the reported or measured value (M_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_i):

$$M_i = (1 - \beta) V_i + \beta S_i \, . \tag{4}$$

[29] These estimates are presented in online Appendix Table A8. In addition, we also examined whether demographic groups can account for negative risk premia. While the fraction of negative risk premia varies somewhat across demographic groups, the fraction with a negative risk premium lies between 15 and 30 percent for all demographic groups that make up at least 5 percent of the sample. Moreover, demographics and the response-quality terciles together have an R^2 of only 0.043 in explaining which respondents have a negative risk premium. The findings are presented in online Appendix Tables A9 and A10.

The parameter β can be interpreted as a bias that affects all respondents uniformly when they see a starting value for the certainty equivalent question 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:

$$V_i = M_i - \frac{\beta}{1-\beta}(S_i - M_i) . \tag{5}$$

The second term on the right-hand side is the adjustment that removes measurement error induced by the starting value. This adjustment requires that equation (4) correctly specifies the structural relationship between the starting value and the reported certainty equivalent.

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 4, and the set of variables related to response quality described above. Our estimate (shown in online Appendix Table A13 and 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.

The second adjustment to the mean risk premium estimates the effects of response-quality indicators on the risk premium and uses these estimates to construct a response-quality correction to the risk premium. Specifically, we regress the risk premium on the response-quality indicators controlling for the starting value and our standard set of demographic and other control variables. The coefficient on a response-quality dummy provides 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 an unbiased estimate of the average risk premium under two conditions: (i) 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, and (ii) measurement error in the risk premium averages out to zero for those who give high-quality responses on all indicators.

E. *Adjusted Risk Premia*

Table 6 shows the impact of the two adjustments to the risk premium, 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 using the coefficient on the starting value from the regressions in online Appendix Table A13, which also include the demographic and other control variables from Table 4 and the response-quality indicators. 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

Table 6—Adjustments for Starting Value and Response Quality

	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	2.3	2.0	6.4	6.5	4.1	4.8	−2.9	−3.3
	(1.3)	(1.2)	(1.2)	(1.2)	(1.2)	(1.5)	(0.6)	(0.6)
5. Estimate adjusted for starting value and response quality	7.4	8.0	65.9	69.2	58.5	62.2	19.7	19.7
	(1.2)	(1.2)	(1.3)	(1.5)	(1.3)	(1.4)	(0.6)	(0.7)
Observations	2,939		2,939		2,939		2,939	

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 the 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 online Appendix Table A7. The coefficients of the regression of the outcome variables on these measures of response quality are shown in online Appendix Table A13.

Source: Data from the June 2011 Social Security Political Risk Survey, designed by the authors and fielded by Knowledge Networks

0.7 percentage points at the mean.[30] 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 further 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 ($\approx$ 20 percent 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. Respondents who have higher response-quality indicators 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. All of these adjustments, at both the mean and the median, are statistically significantly different from zero.

[30] Allowing the sensitivity to the starting value to depend on the response-quality indicators (by including a full set of interactions between the starting value and the response-quality indicators in the regressions reported in online Appendix Table A13) results in a starting value adjustment that is somewhat smaller (0.4) but estimated much less precisely. This estimate is presented in online Appendix Table A14. Given that these interactions are jointly insignificant (p-value 0.82), we do not include them in Table 6.

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 2.3 and 2.0 percentage points, respectively. Though only marginally 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 premia 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 biased downward. 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.

The response-quality adjustment for the risk premium involves a trade-off. Under the assumptions for the adjustment (spelled out in Section IVD), the adjustment results in an unbiased estimate for the average risk premium in the population. However, this estimate has a higher standard error than the estimate that only adjusts for the starting value. Moreover, the response-quality adjustment may add bias if its underlying assumptions, which are not testable, do not hold. Hence, the estimate that only adjusts for the starting value (but not for response quality) may be the preferred estimate for those who value the precision of the estimate and are not as confident in the assumptions required for the response-quality adjustment.

V. 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 US 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 US households. We find that on average respondents would be willing to forgo 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 supposed to get under current law, this risk premium is equivalent to about 10 percent of expected benefits. The informal estimate of the accrued obligation, at the time of our survey under current law, for individuals 25–59, by the Office of the Chief Actuary at the Social Security Administration is about $12 trillion.[31] 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 the welfare cost of perceived policy uncertainty could be reduced. If policy uncertainty is misperceived, policies to correct misperceptions deserve consideration. There are also avenues to reduce actual policy uncertainty. 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 variables. 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.[32] Indeed, Smetters and Theseira (2011) find that fundamental reforms away from traditional pay-as-you-go 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

Alm, James. 1988. "Uncertain Tax Policies, Individual Behavior, and Welfare." *American Economic Review* 78 (1): 237–45.

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

[31] Steve C. Goss, personal communication, April 5, 2012.

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

Baker, Scott, Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–636.

Bell, Felicitie C., and Michael L. Miller. 2005. *Life Tables for the United States Social Security Area 1900–2100*. Actuarial Study No. 120. Social Security Administration SSA Pub. No. 11-1153. Cleveland: Social Security Administration.

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." https://core.ac.uk/download/pdf/6455038.pdf (accessed December 12, 2017).

Bernheim, B. Douglas, and Lawrence Levin. 1989. "Social Security and Personal Saving: An Analysis of Expectations." *American Economic Review* 79 (2): 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): 342–61.

Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. 2017. *The 2017 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds*. Washington, DC: Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds.

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): 83–106.

Caliendo, Frank N., Aspen Gorry, and Sita Slavov. 2015. "The Cost of Uncertainty about the Timing of Social Security Reform." National Bureau of Economic Research Working Paper 21585.

Delavande, Adeline, and Susann Rohwedder. 2008. "Eliciting Subjective Expectations in Internet Surveys." *Public Opinion Quarterly* 72 (5): 866–91.

Diamond, Peter A. 1997. "Insulation of Pension from Political Risk." In *The Economics of Pensions: Principles, Policies and International Experience*, edited by Salvador Valdés-Prieto, 33–57. Cambridge, UK: Cambridge University Press.

Dominitz, Jeff, and Charles F. Manski. 2006. "Measuring Pension-Benefit Expectations Probabilistically." *Labour* 20 (2): 201–36.

Dušek, Libor. 2007. "Political Risk of Social Security: The Case of the Indexation of Benefits in the Czech Republic." Center for Economic Research and Graduate Education Economics Institute Working Paper 318.

Geanakoplos, John, and Stephen P. Zeldes. 2009. "Reforming Social Security with Progressive Personal Accounts." In *Social Security Policy in a Changing Environment,* edited by Jeffrey Brown, Jeffrey Liebman, and David Wise, 73–121. Chicago: University of Chicago Press.

Geanakoplos, John, and Stephen P. Zeldes. 2010. "Market Valuation of Accrued Social Security Benefits." In *Measuring and Managing Federal Financial Risk*, edited by Deborah Lucas, 213–33. Chicago: University of Chicago Press.

Giavazzi, Francesco, and Michael McMahon. 2012. "Policy Uncertainty and Household Savings." *Review of Economics and Statistics* 94 (2): 517–31.

Gomes, Francisco, Laurence Kotlikoff, and Luis Viceira. 2012. "The Excess Burden of Government Indecision." In *Tax Policy and the Economy*, Vol. 26, edited by Jeffrey Brown, 125–63. Chicago: University of Chicago Press.

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

Greenwald, Mathew, Arie Kapteyn, Olivia S. Mitchell, and Lisa Schneider. 2010. "What Do People Know about Social Security?" Social Security Administration Working Paper 792-SSA.

Guiso, Luigi, Tullio Jappelli, and Mario Padula. 2013. "Pension Wealth Uncertainty." *Journal of Risk and Insurance* 80 (4): 1057–85.

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

Herriges, Joseph A., and Jason F. Shogren. 1996. "Starting Point Bias in Dichotomous Choice Valuation with Follow-Up Questioning." *Journal of Environmental Economics and Management* 30 (1): 112–31.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47 (2): 263–91.

Kim, Iltae, Arthur Snow, and Ronald S. Warren, Jr. 1995. "Tax-Rate Uncertainty, Factor Supplies, and Welfare." *Economic Inquiry* 33 (1): 159–69.

King, Miriam L., Steven Ruggles, J. Trent Alexander, and Rebecca Vick. 2010. *Integrated Public Use Microdata Series, Current Population Survey: Version 3.0.* Minneapolis: University of Minnesota.

Kitao, Sagiri. 2016. "Policy Uncertainty and the Cost of Delaying Reform: A Case of Aging Japan." Research Institute of Economy, Trade and Industry Discussion Paper 16-E-013.

Knowledge Networks. 2010. *KnowledgePanel Design Summary*. 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 *Tax Policy and the Economy*, Vol. 26, edited by Jeffrey R. Brown, 1–42. Chicago: University of Chicago Press.

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): 275–99.

Luttmer, Erzo F. P., and Andrew A. Samwick. 2018. "The Welfare Cost of Perceived Policy Uncertainty: Evidence from Social Security: Dataset." *American Economic Review.* https://doi.org/10.1257/aer.20151703.

McHale, John. 2001. "The Risk of Social Security Benefit-Rule Changes: Some International Evidence." In *Risk Aspects of Investment-Based Social Security Reform,* edited by John Y. Campbell and Martin Feldstein, 247–90. Chicago: University of Chicago Press.

Mermin, Gordon B. T. 2005. *Reforming Social Security through Price and Progressive Price Indexing.* Washington, DC: Urban Institute. Chicago: University of Chicago Press.

Nataraj, Sita, and John B. Shoven. 2003. "Comparing the Risks of Social Security with and without Individual Accounts." *American Economic Review* 93 (2): 348–53.

Nelson, Jaeger L. 2017. "Welfare Implications of Uncertain Social Security Reform." https://economics.indiana.edu/home/conferences/2016-jordan-river-economics-conference/files/2016-04-29-14.pdf.

Shoven, John B., and Sita N. Slavov. 2006. "Political Risk Versus Market Risk in Social Security." National Bureau of Economic Research Working Paper 12135.

Skinner, Jonathan. 1988. "The Welfare Cost of Uncertain Tax Policy." *Journal of Public Economics* 37 (2): 129–45.

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

Social Security Administration. 2012. *Annual Statistical Supplement to the Social Security Bulletin*. SSA Publication 13-11700. Washington, DC: Social Security Administration.

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

Van der Wiel, Karen. 2008. "Preparing for Policy Changes: Social Security Expectations and Pension Scheme Participation." Institute for the Study of Labor Discussion Paper 3623.

Weiss, Laurence. 1976. "The Desirability of Cheating Incentives and Randomness in the Optimal Income Tax." *Journal of Political Economy* 84 (6): 1343–52.