

The Welfare Cost of Retirement Uncertainty*

Frank N. Caliendo[†] Maria Casanova[‡] Aspen Gorry[§] Sita Slavov[¶]

May 20, 2016

Abstract

Uncertainty about the date of retirement is a major financial risk with implications for decision making and welfare over the life cycle. Conservative estimates of the standard deviation of the difference between retirement expectations and actual retirement dates range from 4.28 to 6.92 years. This uncertainty implies large fluctuations in total wage income. Our baseline individual would give up about 4% of total lifetime consumption to fully insure this risk and about 3% of lifetime consumption simply to know the retirement date. Uncertainty about the date of retirement helps to explain consumption spending near retirement and precautionary saving behavior. We study how social insurance programs could be designed to hedge this risk and find that current programs in the U.S. (OASI and SSDI) do little to mitigate the risk.

*We thank seminar audiences at the MRRC research workshop and at Utah State.

[†]Utah State University. frank.caliendo@usu.edu.

[‡]Cal State Fullerton. mcasanova@fullerton.edu.

[§]Utah State University. aspen.gorry@usu.edu.

[¶]George Mason University. sslavov@gmu.edu.

1. Introduction

The date of retirement is one of the most important financial events in the life of an individual. It determines the number of years of wage earnings and the expected length of time over which the individual must survive on accumulated savings, both of which are crucial for lifetime budgeting decisions. If the exact date of retirement were known with certainty, then financial planning for retirement would be a relatively easy task.¹

Unfortunately, life is not that simple. Individuals cannot know for sure at age 25 when they will ultimately retire because the transition into retirement is the result of multiple factors that are hard to predict decades in advance. Among many others, these include health status, the retirement and health status of a spouse, changes in working conditions, and the degree of skill obsolescence.

Uncertainty about the timing of retirement is a major financial risk that affects consumption and saving decisions and welfare over the life cycle. We find that retirement timing uncertainty leads to substantial variation in lifetime income. The associated welfare cost to individuals is at least as large as that of other income shocks such as aggregate business cycle risk and idiosyncratic wage shocks.

Our analysis helps to explain some consumption and saving behaviors that often appear puzzling through the lens of traditional economic theory, such as why consumption spending drops discretely upon retirement and why many individuals accumulate such large precautionary savings balances (Scholz, Seshadri and Khitatrakun (2006)). Our analysis also provides insights into how retirement insurance programs can be designed to hedge retirement timing risk, and we characterize first-best insurance arrangements as well as more modest second-best policy adjustments that partially insure retirement timing risk.

This paper proceeds in three steps. First, we provide empirical evidence about the degree of retirement timing uncertainty. Second, we compute the welfare cost to individuals. Third, we assess how well existing social insurance programs mitigate retirement uncertainty and we explore policy adjustments that improve individual welfare.

The convention in the labor supply literature has been to treat the retirement date either as an exogenous, deterministic event or as a completely voluntary, endogenous choice (French (2005), Rogerson

¹Of course, there are other considerations such as uncertainty over asset returns and other risks, as well as limitations on financial literacy that present challenges to the household budgeting and planning process (Lusardi and Mitchell (2007), Lusardi and Mitchell (2008), van Rooij, Lusardi and Alessie (2012), Lusardi, Michaud and Mitchell (2011), Ameriks, Caplin and Leahy (2003), Campbell (2006)).

and Wallenius (2009) and others). By recognizing retirement as an uncertain event, we clearly depart from studies that treat it as fixed. However, our approach is not inconsistent with the modern labor supply literature that treats the retirement date as an endogenous choice. An ideal model would allow individuals to optimally update the retirement date in response to the many different shocks to life circumstances that contribute to retirement. These would include the often-modeled shocks to health, employment, and wages that account for 30% of retirement decisions (Casanova (2013), Szinovacz and Davey (2005)). But if we were to add the many other (potentially) stochastic events that determine the remaining 70%, the model would quickly become unmanageable.² We take a reduced-form approach that measures how individuals optimally update their retirement date in response to the arrival of new information by comparing their expected and actual retirement dates. We use the standard deviation of this variable to assign to the individual a distribution of retirement dates that accounts for endogenous responses to shocks.

Rather than simply using the dispersion in retirement ages in the population as a measure of uncertainty—which could confound uncertainty with heterogeneity because individuals have private information about their expected retirement age—we use the Health and Retirement Study to measure retirement timing uncertainty directly as the standard deviation of the difference between self-reported retirement expectations and actual retirement dates. We estimate standard deviations for a number of subsamples and we make conservative assumptions to obtain a lower bound on the degree of retirement timing uncertainty that individuals face.³ Our estimates range from 4.28 to 6.92 years, depending on the sample. As a baseline we will use 6 years. This uncertainty about the date of retirement is a major financial risk that can deeply affect the financial preparedness of individuals. For instance, an individual who draws a retirement shock at age 59 instead of age 65 would lose about one-sixth of his total (non-discounted) pre-retirement wage income. The loss of multiple years of earnings puts a significant dent in the individual’s lifetime budget, and this loss is amplified by the need to spread available assets over a longer retirement period.

We assess how costly this timing risk is to individuals. We use a variety of data sources to calibrate a quantitative life-cycle model in which individuals make consumption and saving choices in the face of retirement timing uncertainty.⁴ We calculate the welfare cost of retirement timing uncertainty by

²The many reasons for retirement cited in the HRS include changes in family circumstances such as illness and the retirement of a spouse, changes in working conditions, separation and divorce, and unexpected financial incentives to quit among many others.

³Because the HRS samples people above age 50, our estimates of timing uncertainty are likely conservative relative to the timing uncertainty facing younger individuals.

⁴In this paper we deal only with known probabilities and we therefore use the words risk and uncertainty interchangeably

determining what fraction of total lifetime consumption an individual would be willing to give up in order to live in a safe world with comparable expected wealth but no retirement timing uncertainty. We interpret this as the value of *full insurance* against timing risk, because the benchmark is a world where decision making is not distorted and wealth is fully insured. We find that the welfare cost is 4.26% under *laissez faire* with no Social Security.

We also calculate the value of simply knowing the retirement date, which allows the individual to optimize with full information but does not insure the individual’s wealth across realizations of the retirement date. We refer to this as the *timing premium* because it captures the value of early resolution of uncertainty as in Epstein, Farhi and Strzalecki (2014). Again under *laissez faire* with no Social Security, the timing premium is 2.95% of total lifetime consumption. To put the magnitude of these costs into context, they are larger than estimates of the cost of business cycle fluctuations as in Lucas (2003) and the cost of idiosyncratic fluctuations in wage income as in Vidangos (2009).

Given the magnitude of this welfare cost, a natural question to ask is whether existing social insurance programs help to mitigate the cost. We find that a Social Security retirement program, with a benefit-earning rule that is calibrated to match current U.S. policy, does not mitigate much of the large welfare cost. When we include Social Security, the full welfare cost of retirement timing uncertainty drops slightly from 4.26% to 4.05%, and the timing premium also drops slightly from 2.95% to 2.80%.

This small improvement is not because Social Security provides timing *insurance*, but because Social Security boosts the individual’s expected wealth which makes him less sensitive to earnings shocks. To truly insure against timing risk, a program would need to provide individuals with a big payment if they unexpectedly retire early and a small payment if they retire late. However, Social Security does just the opposite because of the positive relationship between benefits and wage earnings: individuals who suffer early retirement shocks must potentially include some zeros in the calculation of their average earnings, while individuals who draw late shocks have the highest possible average earnings. In this sense, Social Security is anti-insurance because it pays good in good states and it pays bad in bad states.

In some public pension systems such as Japan, the UK, Spain, and other countries in Europe, part

throughout. Our theory extends the recursive method in Caliendo, Gorry and Slavov (2015) and Stokey (2014), which is a technique for solving regime switching optimal control problems where the timing and structure of the new regime are uncertain. Technically speaking, the current paper has the added complication that the timing p.d.f. is truncated, which renders the Pontryagin first-order conditions for optimality insufficient to produce a unique solution. We derive a “stochastic continuity” condition as the limiting case of an otherwise redundant transversality condition, in order to identify the unique solution. Our method works for any generic control problem with a stochastic stopping time and a free endpoint on the state variable.

of retirement benefits are completely independent of the individual's earnings history. In other words, a component of retirement benefits is the same regardless of when the retirement shock occurs. We show that this feature can mitigate up to one-third of the welfare costs of retirement timing uncertainty. The largest insurance gain comes from completely breaking the link between benefits and earnings. However, the benefit-earning link encourages labor force participation, and if this is a politically desirable goal then having a component of benefits that is earnings based can preserve some of this incentive effect, while having a component that is unrelated to earnings can at the same time significantly increase the insurance value of Social Security.⁵

To provide a more comprehensive evaluation of the Social Security program's overall role in mitigating timing uncertainty, we extend the model to include disability risk and a disability component within the Social Security program. In the extended model, individuals not only face uncertainty about the timing of retirement, they also face uncertainty about their disability status upon retirement. We find that disability insurance almost perfectly offsets the disability risk that the individual faces, but it does *not* offset the timing risk at all. That is, disability insurance does a nice job of replacing lost post-retirement (part-time) income if the individual is unable to work at all, but it does not solve the problem that the individual doesn't know when such a shock might strike. The joint welfare cost of timing risk and disability risk, in a model with a Social Security program that features both retirement and disability benefits, is almost the same as when disability risk and disability insurance are excluded from the model.

In sum, retirement timing uncertainty is a major financial risk that has not received much attention even though its welfare consequences are large. While there are a few social insurance programs that may appear to offer partial protection against this risk, in fact they do not. Social Security retirement benefits do not provide any insurance against retirement timing risk; and, while Social Security disability benefits might provide some protection, it is still very far from complete insurance.

At a very basic level, the objective of Social Security is to prevent poverty in old age by helping retirees maintain a minimum standard of living. Because benefits are paid out as a life annuity that lasts as long the individual lives, and because replacement rates are more generous for the poor than for the rich, Social Security is commonly thought to meet its objectives. However, retirement timing risk is a major source of volatility in lifetime earnings and retirement wellbeing, and Social Security retirement benefits compound this risk by paying a small life annuity in bad states and a large annuity in good

⁵The Supplemental Security Income (SSI) program in the U.S. has a flavor of a fixed component that is unrelated to earnings. However, only individuals with little or no income qualify. Insuring against timing uncertainty requires a policy that has a fixed component above and beyond SSI that is available to all retirees.

states. Hence, while Social Security may be meeting some of its objectives, it does not pool risk across one of the most important challenges to the financial wellbeing of retirees.

Our paper is related to a large literature that documents a discrete drop in consumption at the date of retirement.⁶ While a variety of explanations have been proposed, our paper clarifies the role that retirement uncertainty could play in explaining the drop. Timing uncertainty causes a reduction in consumption at retirement no matter when the shock is realized, because the retirement shock leaves the individual poorer than expected from the perspective of a moment before the shock occurred. This causes an abrupt adjustment in consumption irrespective of whether the shock happens at an early age or at a late age. Adding disability risk amplifies the size of the drop even further. Therefore, timing uncertainty is a powerful force that can help to explain the observed drop in consumption near retirement.

In addition, retirement timing uncertainty is a powerful channel that may help to explain precautionary savings balances that otherwise seem large. For instance, Scholz, Seshadri and Khitatrakun (2006) estimate that as much as 80% of Americans in the HRS have asset balances that *exceed* the optimal amount of savings from a life-cycle optimization perspective. Individuals in our model not only save for retirement but they also save because they don't know when retirement will strike, and we find that a large portion of observed savings may be due to uncertainty about the date of retirement. Models without retirement timing uncertainty will tend to understate the precautionary motive for saving.

Perhaps the closest paper to ours is Grochulski and Zhang (2013). They also study consumption and saving decisions over the life cycle with uncertainty about the timing of retirement. Like our setting, uncertain retirement leads to precautionary savings and consumption drops discretely when individuals lose their jobs. We extend their analysis by providing empirical evidence on retirement timing uncertainty, by computing the welfare cost of this uncertainty, and by evaluating the role of social insurance programs in mitigating this risk and considering alternative arrangements that improve insurance coverage.⁷

⁶For instance, see Hamermesh (1984), Mariger (1987), Bernheim, Skinner and Weinberg (2001), Hurd and Rohwedder (2006), Scholz, Seshadri and Khitatrakun (2006), Hurst (2006), Haider and Stephens (2007), and Ameriks, Caplin and Leahy (2007) among others.

⁷On the technical side, Grochulski and Zhang (2013) assume stationarity of the timing risk (constant hazard rate of job loss) in an infinite horizon model. We solve a non-stationary problem in which the hazard rate is allowed to depend on age as in the data and we assume individuals face mortality risk over a finite maximum lifespan. In some parameterizations, we also include uncertainty over the individual's disability status, and we allow this second risk to be non-stationary with respect to age. While allowing for non-stationary risk departs from standard dynamic programming, it allows us to more fully calibrate both risks (timing and disability) to the available data.

2. Measuring retirement uncertainty

When thinking about retirement uncertainty, the distinction between voluntary and involuntary retirements, which is at the forefront of the literature studying retirement patterns, comes to mind. Involuntary retirements are the result of employment constraints—due, for example, to the onset of disability or job loss—while voluntary retirees leave the labor force even though the option to remain employed remains available, usually to enjoy more leisure or spend more time with their families (Casanova (2013)).

The distinction between voluntary and involuntary retirement is often interpreted as a distinction between *expected* and *unexpected* retirement. This interpretation owes much to the retirement-consumption literature, which has focused on the Euler equation for the periods right before and after retirement takes place. Several papers have found that the consumption drop at retirement is considerably larger for individuals who retire involuntarily, suggesting that voluntary retirements are anticipated, and allow individuals to better smooth consumption around that event (Banks, Blundell and Tanner (1998), Bernheim, Skinner and Weinberg (2001), Hurd and Rohwedder (2008), Smith (2006)).

While this distinction may be appropriate when considering individuals that are one period away from retirement, it is no longer helpful from the perspective of a model that focuses on the full life cycle profile of consumption. For a worker just entering the labor force, the degree of uncertainty about the likelihood of retiring for involuntary reasons is not necessarily larger than that of retiring voluntarily. For example, a 25 year old may not be better able to predict the probability of becoming disabled before reaching retirement age than that of getting married to a spouse who will retire early, and who will lead him to anticipate his retirement in order to spend time together. The concept of retirement timing uncertainty we use in this paper is hence not limited to the negative employment shocks that cause the one third of involuntary retirements observed in the data (Casanova (2013), Szinovacz and Davey (2005)), but rather covers all life events that may trigger an exit from the labor force which cannot be perfectly foreseen from a young age, including the retirement of a spouse, the birth of a grandchild, a dislike for the work environment in the pre-retirement years, etc.

In order to measure retirement timing uncertainty, we must first make an assumption on how individuals form expectations regarding their retirement age. A straightforward approach would be to assume that the subjective distribution of retirement probabilities coincides with the actual retirement distribution estimated from the data. In particular, if the expected retirement age is assumed to coincide with the average retirement age in the population, deviations of actual retirements from that expectation would be informative about the degree of uncertainty. This assumption of unconditional rational expectations is

likely to yield biased estimates of retirement uncertainty, given that individuals have private information about, e.g., their health status or taste for work, allowing them to predict whether they will retire earlier or later than average.

We follow an alternative approach that makes use of self-reported retirement expectations, and is consistent in the presence of private information.⁸ The implicit assumption is that individuals use all private information at their disposal when reporting their expected retirement age. The degree of uncertainty is given by the size of the deviations between expected and eventual retirement ages. In particular, we estimate the standard deviation of the following variable:

$$X = (Eret - Ret),$$

where $Eret$ is an individual’s expected retirement age, and Ret is the actual age at which retirement takes place.⁹

2.1. Data and empirical evidence

The data come from the Health and Retirement Study (HRS), a nationally representative longitudinal survey of 7,700 households headed by an individual aged 51 to 61 in the first survey wave. Interviews are conducted every two years, and we follow individuals for a maximum of 11 waves, from 1992 to 2012. We measure retirement expectations on wave 1, and then follow individuals up until the end of the panel in order to establish their retirement age.

The variable $Eret$ is constructed from questions that ask individuals when they “plan to stop work altogether” and when they “think [they] will stop work or retire”.¹⁰ We include observations for males who are aged 51 to 61 in wave 1. We exclude those who are not employed or do not report retirement

⁸The use of expectation variables, and retirement expectations in particular, has become commonplace in the literature in recent years. There is a growing number of papers studying the validity of retirement expectations elicited from individuals, and showing that they are strong predictors of actual retirement dates (Bernheim (1989), Dwyer and Hu (1999), Haider and Stephens (2007)), consistent with rational expectations (Benítez-Silva and Dwyer (2005), Benítez-Silva et al. (2008)), and updated upon arrival of new information (Benítez-Silva and Dwyer (2005), McGarry (2004)).

⁹In addition to computing the standard deviation of X , $\sqrt{E[(X - E(X))^2]}$, we have also computed an alternative measure of the amount of uncertainty about the timing of retirement that individuals face, $\sqrt{E(X^2)}$. This alternative measure may be a little more intuitive because it gives the typical gap between $Eret$ and Ret . However, we focus on the first measure because it is mathematically less than (or equal to) the second, making our estimates of timing uncertainty as conservative as possible. In any case, the difference between the two measures is practically insignificant in our samples.

¹⁰We combine the variables $Rwrplnyr$ and $Rwrplnya$ from the RAND-HRS dataset.

plans, which results in a sample of 3,251 individuals. To be consistent with the wording of the retirement expectations questions, we define retirement as working zero hours. The variable Ret is constructed combining information on the first wave in which a respondent is observed to be retired, with the month and year in which he left his last job. In cases where the retirement age is not observed—either because of attrition or the end of the sample period—and for those individuals who say they will never retire, we make assumptions that allow us to get a conservative value for the variable X . These assumptions, together with the strategy used to control for measurement error in retirement expectations, and further details on sample selection and the construction of the variables $Eret$ and Ret , are described in Appendix A.

The major strength of the HRS for our purposes is the fact that it both elicits retirement expectations and then follows workers over time so that their retirement age can be established. The dataset, however, is not without drawbacks. The main disadvantage is that it samples older individuals, so we measure retirement timing uncertainty for a sample of workers who are close to retirement age. Since this likely understates the degree of retirement timing uncertainty facing young individuals, our welfare estimates will be conservative.¹¹

The first column of Table 1 displays the distribution of retirement expectations in our sample. Close to 15% of individuals report that they will never retire, and another 10% state that they do not know when retirement will take place. For individuals who provide a specific retirement date, two peaks are apparent at the Social Security retirement ages of 62 and 65. The last two columns of the table compare reported retirement expectations with actual retirement ages. To do so, we restrict the sample to individuals for whom both the date at which they expect to retire and their eventual retirement date fall within the sample period. Expected retirement ages for this subsample, shown in column 2, display the same peaks at ages 62 and 65. Two facts are striking when comparing the distribution of expected retirements with that of actual retirements, shown in column 3. First, just slightly over 50% of individuals who planned to

¹¹We also likely overstate the degree of uncertainty facing the oldest workers, although this likely has a small effect on our welfare estimates. First, while the degree of retirement timing uncertainty decreases as retirement approaches and more information becomes available, the evidence indicates that it remains high until very close to retirement age. Haider and Stephens (2007) estimate that less than 70% of HRS respondents who expect to retire within one year are in fact retired by the next survey wave. Our own estimates show that we are not missing a sharp drop in uncertainty as retirement nears. Robustness checks presented in the appendix show that the standard deviation of X decreases by only half a year to one year when comparing the sample of individuals aged 51 to 55 to those aged 56 to 61. Second, uncertainty at younger ages sets the shape of the saving profile in all subsequent years, and hence has a larger effect on welfare than uncertainty in the years right before retirement.

retire at the Social Security ages actually do so. Second, a large share of retirements takes place before age 55 and after age 66, suggesting that a significant number of individuals end up retiring considerably earlier or later than they anticipated.

Table 2 shows estimates of the standard deviation of X for different samples. The most conservative estimate, presented in row 1, equals 4.28. It is obtained from the sample of individuals for whom both $Eret$ and Ret are observed. Because this subsample excludes individuals likely to face the highest degree of uncertainty—those whose actual retirement date is censored, who say they will never retire, or who do not know when they will retire—the resulting estimate yields a lower bound on retirement timing uncertainty. Subsequent rows use larger samples, adding individuals for whom either $Eret$ or Ret are not observed, but can be assigned a value by making a conservative assumption, as discussed in Appendix A. It is important to point out that the estimate shown in the last row (6.82) is not intended to represent an upper bound on uncertainty, as it is still obtained from a sample of individuals close to retirement age.

In the baseline simulations of the model, we use a value of 6 for the standard deviation of uncertainty, implying that an individual who draws a one-standard-deviation shock will stop working 6 years earlier or later than expected. While this value likely understates the true degree of retirement timing uncertainty, for the reasons stated above, we also consider smaller standard deviations in the robustness checks to illustrate the relationship between the degree of uncertainty and welfare.¹²

¹²Instead of using self-reported retirement expectations in the construction of retirement timing uncertainty, suppose we had taken the simple approach of assuming that the subjective distribution of retirement probabilities coincides with the actual retirement distribution estimated from the data. This simple exercise leads to a standard deviation in retirement uncertainty that is a little less than 6 years, and so in the end we would calibrate our theoretical model the same way.

Table 1. Distribution of Expected and Actual Retirement Ages

	All	Both <i>Eret</i> and <i>Ret</i> during sample period	
	<i>Eret</i>	<i>Eret</i>	<i>Ret</i>
Age < 55	0.52	0.74	4.59
Age = 55	1.91	2.69	2.64
Age = 56	1.23	1.85	2.75
Age = 57	1.02	1.37	3.43
Age = 58	1.41	2.22	4.44
Age = 59	1.29	1.69	5.02
Age = 60	4.46	6.39	7.98
Age = 61	2.77	3.70	8.29
Age = 62	18.33	25.30	16.96
Age = 63	8.74	12.15	7.40
Age = 64	1.48	1.85	6.29
Age = 65	16.98	21.45	8.40
Age = 66	7.72	9.93	4.23
Age > 66	8.00	8.66	17.59
Never	14.61		
Do not know	9.54		
<i>N</i>	3,251	1,893	1,893

Table 2. Standard Deviation of X for Different Subsamples

	Sample	Standard Deviation	N
1	Ret observed	4.28	1,903
2	1 + Work past $Eret$, Ret not observed	5.05	2,147
3	2 + $Eret$ after sample period, Ret not observed	5.04	2,152
4	3 + Will never retire, Ret observed	6.54	2,476
5	4 + Will never retire, Ret not observed	6.35	2,627
6	5 + DK when they will retire, Ret observed	6.92	2,840
7	6 + DK when they will retire, Ret not observed	6.82	2,937

3. A model of retirement uncertainty

In this section we construct a dynamic stochastic model of individual consumption and saving decisions over the life cycle in the face of uncertainty about the timing of retirement and uncertainty about disability status after retirement. By not imposing any specific assumptions about the distribution of timing uncertainty—for instance, we do not assume the distribution of timing uncertainty is stationary—our model is flexible enough to conform to the moments of uncertainty observed in the data. We also do not impose any restrictions on the distribution of disability risk, which allows us to calibrate this second layer of uncertainty to estimates of the probability of becoming disabled conditional on each retirement age.

3.1. Notation

Age is continuous and is indexed by t . Individuals start work at $t = 0$ and pass away no later than $t = T$. The probability of surviving to age t is $\Psi(t)$. A given individual collects wages at rate $(1 - \tau)w(t)$ as long as retirement has not yet occurred, where τ is the Social Security tax rate. The retirement date is a continuous random variable with continuously differentiable p.d.f. $\phi(t)$ and c.d.f. $\Phi(t)$, with support $[0, t']$, where $t' < T$ so that everyone draws a retirement shock before some specified age. We truncate the p.d.f. for two practical reasons. First, truncation prevents us from needing to estimate the $w(t)$ profile deep into old age when data are not reliable. Second, truncation prevents us from having an extremely

thin right tail on $\phi(t)$, which creates technical difficulty as the computer is unable to distinguish between $1 - \Phi(t)$ and 0, and the term $1 - \Phi(t)$ appears in the denominator of first-order conditions for optimality.

At the moment retirement strikes at age t , the individual collects a lump sum $B(t, d) = SS(t|d) + Y(t) \times (1 - d)$ where $SS(t|d)$ is the present discounted value (as of shock date t) of Social Security retirement and disability benefits, d is an indicator variable that equals 1 if the individual has become disabled and 0 if he is still able to work part time after retirement, and $Y(t)$ is the present discounted value (as of shock date t) of post-retirement earnings.¹³ Let d be a random variable with conditional p.d.f. $\theta(d|t)$, hence $\theta(0|t) + \theta(1|t) = 1$ for all t . Note that d may be correlated with the retirement shock t , and we assume that $\theta(d|t)$ is continuously differentiable in t .¹⁴ Hence, $\theta(1|t)$ should be interpreted as the probability that the individual will qualify for disability benefits if retirement strikes at date t . We abstract from policy uncertainty about future Social Security reform (Caliendo, Gorry and Slavov (2015)).

Consumption spending is $c(t)$ and private savings in a riskless asset is $k(t)$, which earns interest at rate r . Annuity markets are closed, and capital markets are perfect in the sense that the individual can borrow and lend freely at rate r . The individual starts with no assets, has no bequest motive, and is not allowed to leave debt behind at $t = T$. Hence, $k(0) = k(T) = 0$.

3.2. Individual problem

Period utility is CRRA over consumption with relative risk aversion σ , and utils are discounted at the rate of time preference ρ .¹⁵ The individual takes as given factor prices and government taxes and transfers, while treating the retirement date as a continuous random variable and disability as a binary random variable. We extend the recursive method in Caliendo, Gorry and Slavov (2015) and Stokey (2014) to the current setting and we relegate lengthy proofs and derivations to Appendix B.

As long as retirement has not yet occurred, the individual follows a contingent plan $(c_1^*(t), k_1^*(t))_{t \in [0, t]}$,

¹³Income from asset holdings is not included in $Y(t)$ because asset holdings are modeled separately.

¹⁴We assume continuous differentiability in t for notational convenience. We could easily allow for a finite number of discontinuities in the t dimension, but then we would need to break the p.d.f. apart at each discontinuity and allow for a unique maximum condition for each continuous segment. This would complicate notation without adding much economic content.

¹⁵We abstract from leisure in the period utility function. As we discuss later in the paper, under common assumptions this simplification has no impact on our welfare calculations.

which solves the following dynamic stochastic control problem (where t and d are random variables)

$$\max_{c(t)_{t \in [0, t']}} : \int_0^{t'} \left\{ [1 - \Phi(t)] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \sum_d \theta(d|t) \phi(t) \mathcal{S}(t, k(t), d) \right\} dt$$

subject to

$$\mathcal{S}(t, k(t), d) = \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k(t), d)^{1-\sigma}}{1-\sigma} dz,$$

$$\frac{dk(t)}{dt} = rk(t) + (1 - \tau)w(t) - c(t),$$

$$k(0) = 0, k(t') \text{ free},$$

where $c_2^*(z|t, k(t), d)$ solves the post-retirement deterministic problem for given $k(t)$ and given realizations of t and d

$$\max_{c(z)_{z \in [t, T]}} : \int_t^T e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz,$$

subject to

$$\frac{dK(z)}{dz} = rK(z) - c(z), \text{ for } z \in [t, T],$$

$$t \text{ and } d \text{ given, } K(t) = k(t) + B(t, d) \text{ given, } K(T) = 0,$$

where $K(t)$ is total financial assets at retirement, which includes accumulated savings $k(t)$ plus the lump-sum payment $B(t, d)$.

The pre-retirement solution $(c_1^*(t), k_1^*(t))_{t \in [0, t']}$ obeys the following system of differential equations and boundary condition

$$\frac{dc(t)}{dt} = \left(\frac{c(t)^\sigma e^{(\rho-r)t}}{\sigma \Psi(t)} \sum_d \theta(d|t) \left[\frac{(k(t) + B(t, d)) e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} - \frac{1}{\sigma} \right) \left(\frac{c(t) \phi(t)}{1 - \Phi(t)} \right) + \left[\frac{\Psi'(t)}{\Psi(t)} + r - \rho \right] \frac{c(t)}{\sigma},$$

$$\frac{dk(t)}{dt} = rk(t) + (1 - \tau)w(t) - c(t),$$

$$k(0) = 0,$$

where the remaining boundary condition $c(0)$ is chosen optimally (explained in Appendix B). And the optimal consumption path for $z \in [t, T]$ after the retirement shock has hit at date t with optimal savings

$k_1^*(t)$ is

$$c_2^*(z|t, k_1^*(t), d) = \frac{(k_1^*(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}, \text{ for } z \in [t, T].$$

3.3. Welfare

In this section we introduce two welfare costs. Each is a measure of willingness to pay to avoid retirement uncertainty. The first is our baseline welfare cost, which captures the value of fully insuring against retirement uncertainty. The second captures just the value of early resolution of uncertainty. We refer to the baseline welfare cost as the value of *full insurance*, and we refer to the second welfare cost as the *timing premium*.

We begin with the value of full insurance. As a point of reference, consider the case where the individual faces no risk (NR) about retirement. Instead, the individual is endowed at $t = 0$ with the same expected future income (as in the world with retirement uncertainty) and solves

$$\max_{c(t)_{t \in [0, T]}} : \int_0^T e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} dt,$$

subject to

$$\frac{dk(t)}{dt} = rk(t) - c(t),$$

$$k(0) = \int_0^{t'} \left(\sum_d \theta(d|t) \phi(t) \left(\int_0^t e^{-rv} (1-\tau) w(v) dv + B(t, d) e^{-rt} \right) \right) dt, \quad k(T) = 0.$$

The solution is

$$c^{NR}(t) = \frac{k(0) e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}}{\int_0^T e^{-rv+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv}, \text{ for } t \in [0, T].$$

The baseline welfare cost of living with retirement uncertainty (value of full insurance) Δ is the solution to the following equation

$$\int_0^T e^{-\rho t} \Psi(t) \frac{[c^{NR}(t)(1-\Delta)]^{1-\sigma}}{1-\sigma} dt$$

$$= \int_0^{t'} \left(\sum_d \theta(d|t) \phi(t) \left(\int_0^t e^{-\rho z} \Psi(z) \frac{c_1^*(z)^{1-\sigma}}{1-\sigma} dz + \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k_1^*(t), d)^{1-\sigma}}{1-\sigma} dz \right) \right) dt.$$

By equating utility from expected wealth to expected utility, our baseline welfare cost Δ measures the individual's willingness-to-pay to have one's expected wealth. This captures the value of *full insurance* because the individual is paying to have his expected wealth with certainty, rather than paying merely for information about retirement.

While our baseline welfare cost Δ follows in the tradition of calculating willingness-to-pay to avoid uncertainty by equating utility from expected wealth to expected utility, there are other sensible ways to calculate the welfare cost of retirement uncertainty. For example, rather than using utility from expected wealth as the welfare benchmark, we could instead use as a benchmark a world in which the individual learns at time 0 when and how retirement uncertainty will be resolved so that the individual follows the optimal deterministic consumption path conditional on that information. To compute the welfare cost of retirement uncertainty, we would then compare the ex ante expected utility of this world (expected utility just before the time 0 information is released) to the expected utility of living with retirement uncertainty.

Following this alternative approach, we now formally define the *timing premium*. Now our point of comparison is a world where at time 0 the individual learns both the retirement date t as well as the disability indicator d . Upon learning these things, the individual solves a deterministic problem:

$$\max_{c(z)_{z \in [0, T]}} : \int_0^T e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz,$$

subject to

$$\frac{dk(z)}{dz} = rk(z) - c(z),$$

$$k(0|t, d) = \int_0^t e^{-rv}(1-\tau)w(v)dv + B(t, d)e^{-rt}, \quad k(T) = 0.$$

The solution is

$$c(z|t, d) = \frac{k(0|t, d)e^{(r-\rho)z/\sigma}\Psi(z)^{1/\sigma}}{\int_0^T e^{-rv+(r-\rho)v/\sigma}\Psi(v)^{1/\sigma}dv}, \text{ for } z \in [0, T].$$

The timing premium Δ_0 is the solution to the following equation

$$\int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^T e^{-\rho z} \Psi(z) \frac{[c(z|t, d)(1-\Delta_0)]^{1-\sigma}}{1-\sigma} dz \right) \right) dt$$

$$= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^t e^{-\rho z} \Psi(z) \frac{c_1^*(z)^{1-\sigma}}{1-\sigma} dz + \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k_1^*(t), d)^{1-\sigma}}{1-\sigma} dz \right) \right) dt.$$

In other words, we are calculating how much an individual would pay at time 0 to know his retirement date t and his future disability status upon retirement d ? This exercise is guaranteed by Jensen's inequality to yield a smaller welfare cost from retirement uncertainty than what is generated by our baseline method (see the proof in Appendix C). The individual would always pay more to have his expected wealth with certainty (Δ) than he would pay for retirement information (Δ_0), because simply knowing

one's wealth is not as good as insuring one's wealth.

Our timing premium is related to the timing premium in Epstein, Farhi and Strzalecki (2014). In both cases, it is the amount individuals would pay for early resolution of uncertainty. However, their premium is the result of Epstein-Zin recursive preferences, which carry a taste for early resolution of uncertainty even if early information is not used to reoptimize. Indeed, in their setting individuals do not reoptimize if information is released early. In contrast, in our setting with CRRA utility the timing premium is the result of better decision making in the face of early information. Including a taste for early information would only enhance the magnitude of the welfare cost of retirement uncertainty.¹⁶

Finally, one may be concerned that we have abstracted from leisure in the period utility function. That is, it may seem that the negative consequences of an early retirement shock are partly mitigated if early retirement brings more leisure. However, at least for the common case in which consumption and leisure are additively separable, this is not the case. In fact, if we include leisure in the period utility function, then the baseline welfare cost will strictly increase. This is because retirement timing uncertainty now imposes an additional cost on the individual in the form of uncertainty about leisure time, and he would pay an additional premium to fully insure this risk. On the other hand, adding leisure to the period utility function leaves the timing premium unchanged; the individual would not pay an additional premium for early resolution of uncertainty about his fixed leisure endowment. We prove these points in Appendix D.¹⁷

4. Calibration

The parameters to be chosen are the maximum lifespan T , the survival probability $\Psi(t)$ as a function of age t , the individual discount rate ρ , the coefficient of relative risk aversion σ , the real return on assets r , the age-earnings profile $w(t)$, the p.d.f. over timing risk $\phi(t)$ and its upper support t' , the present discounted value of post-retirement earnings $Y(t)$ as a function of retirement date t , the Social Security tax rate τ , the present discounted value of Social Security retirement and disability benefits $SS(t|d)$ as

¹⁶There are at least two other ways in which our modeling of the welfare cost is conservative. First, we endow the individual with full information about the distributions of the random variables over both timing risk and disability risk. Second, we assume the individual saves optimally in the face of these risks and therefore accumulates optimal precautionary savings balances to buffer the shocks.

¹⁷If consumption and leisure are complements, then we presume retirement timing uncertainty would become even more costly than in our baseline model without leisure, because in this case an early retirement shock would leave the individual with reduced wealth and with a reduced ability to enjoy that wealth. In this way, the stakes are amplified and the welfare cost would naturally increase.

a function of retirement date t and disability state d , as well as the conditional p.d.f. over disability risk $\theta(d|t)$. Table 3 provides a comprehensive summary of our calibration.

4.1. Lifespan, preferences, and wages

The individual starts work at age 25 (model age $t = 0$) and passes away no later than age 100 (model age $t = 1$). Hence we set the maximum lifespan to $T = 1$.

Our survival data come from the Social Security Administration’s cohort mortality tables. These tables contain the mortality assumptions underlying the intermediate projections in the 2013 Trustees Report. The mortality table for each cohort provides the number of survivors at each age $\{1, 2, \dots, 119\}$, starting with a cohort of 10,000 newborns. However, we truncate the mortality data at age 100, assuming that everyone who survives to age 99 dies within the next year. In the baseline results, we assume individuals enter the labor market at age 25, giving them a 75-year potential lifespan within the model. In our baseline parameterization, we use the mortality profile for males born in 1990, who are assumed to enter the labor market in 2015. For this cohort, we construct the survival probabilities at all subsequent ages conditional on surviving to age 25.

We fit a continuous survival function that has the following form:

$$\Psi(t) = 1 - t^x.$$

After transforming the survival data to correspond to model time, with dates on $[0, 1]$, $x = 3.28$ provides a close fit to the data (see Figure 1).

The utility parameters ρ and σ vary somewhat in the literature. We will consider common values, $\rho = 0$ and $\sigma = 3$. We assume a risk-free real interest rate of 2.9% per year, which is the long-run real interest rate assumed by the Social Security Trustees. In our model, this implies a value of $r = 75 * 0.029 = 2.175$.

We truncate wages $w(t)$ at model time $t' = 2/3$ or actual age 75 because of our concern with the reliability of wage data beyond 75.¹⁸ Using data for workers between 16 and 75 years of age, we fit a fifth-order polynomial to simulated wage income (which is described in detail in the next paragraph) and then we normalize the result such that maximum wages are unity. Although we include observations before age 25 with the view that more observations are better, model time zero corresponds to age 25 and therefore we feed just the post-25 segment of the fitted wage profile (model time $[0, t']$) into the

¹⁸For instance, the data show an upward trend in wages for most education groups between ages 75 and 85, which would seem to reflect selection problems rather than the true wage profile of a particular worker.

individual’s optimization problem (see Figure 2)

$$w(t)_{t \in [0, t']} = 0.3863 + 2.5479t - 2.2727t^2 - 8.4466t^3 + 29.9410t^4 - 29.0828t^5.$$

Our simulated wage income is based on data from the 2014 Current Population Survey (CPS) Merged Outgoing Rotation Group (MORG) file created by the National Bureau of Economic Research. Households that enter the CPS are initially interviewed for 4 months. After a break of 8 months, they are then interviewed again for another 4 months before being dropped from the sample. Questions about earnings are asked in the 4th and 8th interviews, and these outgoing interviews are included in the MORG file. We restrict the sample to men and calculate, at each age, the ratio of average annual earnings¹⁹ to the 2014 Social Security average wage index (AWI).²⁰ Next, we project the AWI forward starting in 2015, assuming that it grows at 3.88% per year in nominal terms. This is consistent with the 2015 Social Security Trustees Report’s intermediate assumptions about nominal wage growth. Multiplying this series by the previously calculated age-specific ratios produces a nominal wage profile for a hypothetical worker who is aged 25 in 2015. This series is deflated to 2015 dollars assuming inflation of 2.7% per year, again consistent with the Social Security Trustees’ intermediate assumptions for 2015.

4.2. Retirement timing

We use a truncated beta density to capture uncertainty over the timing of retirement,

$$\phi(t) = \frac{t^{\gamma-1}(t' - t)^{\beta-1}}{\int_0^{t'} t^{\gamma-1}(t' - t)^{\beta-1} dt}, \text{ for } t \in [0, t']$$

with mean and variance

$$\mathbb{E}(t) = t' \frac{\gamma}{\gamma + \beta}$$

$$\text{var}(t) = \frac{t' \beta \mathbb{E}(t)}{(\gamma + \beta)(\gamma + \beta + 1)}.$$

We truncate the density function at age 75 for consistency with the truncation of wages at age 75, or model time $t' = 2/3$. We set the mean retirement age to 65 which corresponds to model time $\mathbb{E}(t) = 40/75$ and

¹⁹Average weekly earnings are provided for non-self employed workers. We multiply these by 52 to obtain annual earnings. We use the CPS earnings weights to calculate average annual earnings by age. Since CPS earnings data are topcoded, our average earnings estimates are likely to be biased downward.

²⁰We project this based on the 2013 AWI of \$44,888.16 by inflating it by the long-term average nominal wage growth assumption in the 2015 Social Security Trustees Report (3.88% per year).

the standard deviation to 6 years (as explained earlier) which corresponds to model time $\sqrt{\text{var}(t)} = 0.08$. Then, from the mean and variance equations we can calculate the remaining parameters²¹

$$\gamma = \frac{[t' - \mathbb{E}(t)] (\mathbb{E}(t))^2}{t' \text{var}(t)} - \frac{\mathbb{E}(t)}{t'} = 8.0889$$

$$\beta = \gamma \left(\frac{t'}{\mathbb{E}(t)} - 1 \right) = 2.0222.$$

See Figure 3 for a graph of the p.d.f.

4.3. Retirement income and insurance

Beginning with the 1,798 individuals included in the calculations of the standard deviation of retirement uncertainty, we drop individuals who do not have a retirement date (either observed or imputed as described above), individuals with a zero respondent-level analysis weight, and individuals who are only observed in a single wave (thus providing no within-person variation for our fixed effects models). This leaves us with 1,300 individuals and 12,132 person-wave observations over the 11 waves of the HRS. To check robustness, we also re-do all of our analysis using the sample of 1,001 individuals (9,437 person-year observations) who are observed to retire within the sample period.

The RAND version of the HRS includes information about several categories of income, including earnings from work, capital income, pension and annuity income, Supplemental Security Income (SSI) and Social Security Disability Insurance (SSDI) income, Social Security retirement income, unemployment insurance and worker’s compensation, other government transfers (including veteran’s benefits, welfare, and food stamps), and other income (including alimony, lump sums from pensions and insurance, inheritances, and any other income). Except for capital income and other income, which are provided at the household level, all income categories are measured at the individual level. We focus on income in two categories: earnings from work and income from non-Social Security transfers (in which we combine unemployment insurance, worker’s compensation, and other government transfers). Since we explicitly

²¹Truncating the timing density at age 75 works well for two reasons. First, if we truncate much earlier, then we are unable to match both the desired mean (65) and desired standard deviation (6 years), because if the mean is too close to the truncation date then it is impossible to deliver a large enough variance. Of course, we are working with a specific distribution, and perhaps other distributions (with fatter tails) could allow truncation at earlier ages while still hitting our targets for the mean and variance. Second, if we truncate too much later than 75, then we end up with an extremely thin right tail, which ultimately creates “division by zero” errors in our computational procedures as the computer is unable to distinguish between zero and the area in the right tail.

model post-retirement SSDI, Social Security retirement benefits, and asset income (which could include income from pensions and annuities, as well as interest, rent, dividends, and other such income) we exclude these components of income from our analysis.²² We also ignore the “other income” category, as pension lump sums would be classified as capital income, and alimony and inheritances are unlikely to be correlated with retirement. All income figures are converted to July 2015 dollars using the Consumer Price Index for all urban consumers (CPI-U).

To determine how income changes after retirement, we regress each component of income on a set of indicators for time since/before retirement, a set of age dummies, a set of wave dummies, and a set of individual fixed effects. We use respondent-level analysis weights in our regressions and cluster standard errors by individual. The results from these regressions are shown in Table 4. The first three columns show results for the full sample, and the last three show results for the subset of individuals who are observed to retire within the sample period. We only report coefficients for the time since/before retirement indicators; full results are available upon request. The omitted category is the wave immediately before retirement; thus, all coefficients show the change in income relative to this benchmark. The wave of retirement denotes the first wave in which the individual is observed to be retired. Since we do not know how long before the interview the respondent retired, and since income amounts are provided for the previous calendar year, the change in earnings in the wave of retirement is relatively small. However, in subsequent waves, earnings from work decline by between \$41,639 and \$44,821. Relative to their mean in the wave just before retirement (shown in the table), earnings drop by around 80 percent in the first wave after retirement. The point estimates suggest that non-Social Security transfers rise slightly in the wave of retirement and possibly one wave after retirement, but fall slightly in subsequent waves. All but one of these coefficients are statistically insignificant at the 5 percent level.

Based on these estimates, we endow the individual with a lump sum at the date of retirement t , that reflects the present value (as of the retirement date) of post-retirement earnings

$$Y(t) = 0.2w(t) \int_t^T e^{-r(v-t)} dv.$$

That is, post-retirement earnings are equal to 20% of what they were at the time of retirement. We ignore non-Social Security transfers since these appear to be small and do not show statistically significant changes upon retirement. Recall that this endowment is collected only if the individual does not draw

²²The capital income category in the HRS also includes self-employment, business, and farm income. Thus, we are also excluding these components of income from our analysis.

the disability shock.²³

The Social Security program $(\tau, SS(t|d))$ is modeled after the current U.S. program with a tax of $\tau = 10.6\% + 1.8\%$ on wage earnings (which includes the retirement and disability parts of the program). We adopt a simplified Social Security arrangement that captures the most important channels through which the stochastic retirement timing mechanism can influence the level of Social Security benefits. First, the date of the retirement shock affects the individual's average wage income, which in turn influences the individual's benefits through the benefit-earning rule. Second, for those who become disabled, the Social Security disability program acts as a bridge between wage income and retirement benefits.

The total level of Social Security benefits collected is state dependent. For those who do not become disabled but instead retire for other reasons, we compute the individual's average wage income corresponding to the last 35 years of earnings (which is virtually equivalent to the top 35 years of earnings for the wage profile that we are using). If retirement strikes before reaching 35 years in the workforce, then some of these years will be zeros in the calculation. Conversely, as the individual works beyond 35 years, average earnings will increase because a low-wage early year drops out of the calculation while a high-wage later year is added to the calculation. Then, we use a piecewise linear benefit-earning rule that is concave in the individual's average earnings, reflecting realistic slopes and bend points. Finally, we calculate benefits based on collection at age 65, and then we make actuarial adjustments to accommodate early and late retirement dates.

On the other hand, for those who become disabled we compute average wage income corresponding to the last 35 years of earnings, and *no* zeros are included in the average if the individual draws a timing shock that leaves him with fewer than 35 years of work experience. Moreover, he begins collecting full benefits at the moment he retires (rather than waiting until age 65).²⁴ See Appendix E for a full explanation of the state-dependent Social Security program.

Finally, to find the probability of becoming disabled conditional on retirement at t , $\theta(1|t)$, we fit a fifth-order polynomial to the joint probability of becoming disabled and retired at age t (which comes from 2009 disability awards for males between the ages of 17 and 67, reported in 5-year bins, Zayatz (2011)), and then we divide the result by our p.d.f. over timing risk $\phi(t)$ to come up with the probability

²³In reality, non-disabled retirees may or may not collect income from part-time work, whereas in our model we are endowing them with post-retirement earnings that reflect the average life-cycle experience. In doing this, we are suppressing another layer of risk that could make our welfare cost even larger: in reality, non-disabled individuals face uncertainty about post-retirement earnings (their skills may or may not become obsolete, for example).

²⁴In the U.S., it takes a few months for a worker to begin collecting disability benefits after becoming disabled. We have simplified so that benefits commence upon disability.

of disability conditional on retirement age. If the resulting ratio is greater than 1, we assign a value of 1; if the resulting ratio is less than 0, we assign a value of 0. Figure 4 is a graph of our estimated $\theta(1|t)$ profile²⁵

$$\theta(1|t) = \frac{0.0022 + 0.0187t - 0.0721t^2 - 0.6021t^3 + 4.4001t^4 - 5.5785t^5}{\phi(t)}.$$

²⁵In making these calculations, we are assuming that recovery doesn't occur once someone is disabled; that is, disability always implies retirement. In reality, some fraction of people do recover, but it's less than 1% per year (Autor (2011)).

Table 3. Summary of Baseline Calibration of Paramaters

Lifespan, preferences, and wages:

$T = 1$	Normalized maximum lifespan
$\Psi(t) = 1 - t^{3.28}$	Survival probabilities from SS mortality files
$\rho = 0$	common discount rate in the literature
$\sigma = 3$	common CRRA value in the literature
$r = 0.029 \times 75 = 2.175$	Real interest rate from Trustees Report
$w(t) = \sum_{i=0}^5 w_i t^i$	pre-retirement wages (w_i estimated from CPS MORG 2014)

Retirement timing:

$\phi(t) = \frac{t^{\gamma-1}(t'-t)^{\beta-1}}{\int_0^{t'} t^{\gamma-1}(t'-t)^{\beta-1} dt}$, for $t \in [0, t']$	truncated beta p.d.f. over retirement date
$t' = 2/3$	truncation at age 75 (max retirement age)
$\mathbb{E}(t) = 40/75$ (age 65)	mean retirement age
$\sqrt{\text{var}(t)} = 0.08$ (6 years)	standard deviation of retirement age (HRS)
$\gamma = \frac{[t' - \mathbb{E}(t)](\mathbb{E}(t))^2}{t' \text{var}(t)} - \frac{\mathbb{E}(t)}{t'} = 8.0889$	calibrated value
$\beta = \gamma \left(\frac{t'}{\mathbb{E}(t)} - 1 \right) = 2.0222$	calibrated value

Retirement income and insurance:

$Y(t) = 0.2w(t) \int_t^T e^{-r(v-t)} dv$	pdv of post-retirement earnings (HRS)
$\theta(1 t)$	prob disability cond. on retirement (Zayatz (2011) and HRS)
$\tau = 10.6\% + 1.8\%$	Statutory rates for SS ret. and SS disability (U.S. system)
$SS(t d)$	state-dependent pdv of SS benefits (U.S. system)

Table 4. Post-Retirement Income

VARIABLES	Full Sample			Observed to Retire		
	(1)	(2)	(3)	(4)	(5)	(6)
	Earnings	Transfers	Total	Earnings	Transfers	Total
2+ Waves Pre-Retirement	3,437 (2,543)	117.4 (202.9)	3,555 (2,554)	2,851 (3,064)	227.3 (240.0)	3,079 (3,076)
Wave of Retirement	-23,508*** (1,864)	307.5 (198.6)	-23,200*** (1,865)	-24,766*** (2,250)	414.3* (236.2)	-24,352*** (2,250)
1 Wave Post-Retirement	-44,821*** (2,316)	-65.55 (243.0)	-44,887*** (2,323)	-48,613*** (2,692)	172.1 (285.0)	-48,441*** (2,698)
2 Waves Post-Retirement	-43,354*** (3,011)	-301.1 (291.4)	-43,655*** (3,020)	-46,909*** (3,642)	-163.6 (336.1)	-47,073*** (3,649)
3 Waves Post-Retirement	-42,597*** (3,742)	-593.9* (358.5)	-43,191*** (3,755)	-47,065*** (4,606)	-450.4 (420.9)	-47,516*** (4,618)
4 Waves Post-Retirement	-42,958*** (4,485)	-819.5* (426.3)	-43,777*** (4,501)	-46,754*** (5,635)	-613.2 (498.9)	-47,367*** (5,649)
5 Waves Post-Retirement	-43,457*** (5,293)	-839.2* (491.2)	-44,296*** (5,316)	-47,236*** (6,678)	-587.5 (576.3)	-47,824*** (6,699)
6 Waves Post-Retirement	-41,639*** (6,859)	-1,502** (634.4)	-43,141*** (6,888)	-46,166*** (8,689)	-1,234* (747.2)	-47,400*** (8,714)
Pre-Retirement Mean	56,362.94	1,536.075	57,899.02	60,972.2	1,312.081	62,284.28
% Change	-79.5%	-4.3%	-77.5%	-79.7%	13.1%	-77.8%
Observations	12,132	12,132	12,132	9,437	9,437	9,437
R-squared	0.269	0.011	0.267	0.285	0.013	0.282
Number of Individuals	1,300	1,300	1,300	1,001	1,001	1,001

Standard errors clustered by individual in parentheses.

All regressions also include wave and age dummies, and individual fixed effects.

*** p<0.01, ** p<0.05, * p<0.1

5. Quantitative results with timing risk only

To focus attention on the main feature of our model (timing risk), throughout this section we abstract from disability risk and from the disability insurance aspect of the Social Security program. In the next section we will add these features back into the model.

We begin by presenting quantitative results from a version of the model in which there is no Social Security taxation and no Social Security retirement benefits. Then we assess whether various social insurance arrangements (including Social Security) can mitigate the welfare cost of retirement timing risk.

5.1. Consumption, savings, and welfare without insurance

Figure 5 plots consumption over the life cycle for the case in which there is no Social Security taxation and no Social Security retirement benefits. The consumption function c_1^* is the optimal consumption path conditional on the individual still working. The domain of this function stretches from zero up to the maximum working age $t' = 50/75$ (age 75). As soon as the individual draws a retirement shock, he jumps onto the new optimal consumption path c_2^* . Although the retirement date is a continuous random variable in the model, for expositional purposes in the figure we show just four hypothetical shock dates (age 60, 65, 70, and 75). The figure helps to illustrate the magnitude of the distortions to consumption, relative to a safe world in which the individual would simply consume c^{NR} .

Pre-retirement consumption c_1^* starts out below no-risk consumption c^{NR} . The individual must be conservative during the earlier years because the timing of retirement is unknown. However, if he continues to stay working, then eventually the risk of very early retirement begins to dissipate and he responds by spending more aggressively and c_1^* rises above no-risk consumption c^{NR} .

Notice that the retirement shock is accompanied by a downward correction in consumption, with the earliest dates generating the largest corrections. Only those who draw the shock at the last possible moment will smooth their consumption across the retirement threshold. For example, if the shock hits at the average age of 65, then consumption will drop by about 13%.

Why does consumption always drop, even for those who experience a late shock? Because a shock at age t is always earlier than expected (in a mathematical sense) from the perspective of age $t - \epsilon$. In other words, at $t - \epsilon$ the individual expects the shock to occur later than it actually occurs, and therefore he turns out to be poorer at t than he anticipated at $t - \epsilon$. Hence, the consumption drop is the result of rational expectations over retirement timing risk.

The drop in consumption at retirement in our model is consistent with a large literature that documents a drop in consumption roughly in the range of 10%-30%.²⁶ There have been a variety of explanations for the drop, including the cessation of work-related expenses, consumption-leisure substitutability, home production, and various behavioral explanations such as the sudden realization that one's private assets are insufficient to keep spending at pre-retirement levels. Our paper clarifies the role that uncertainty about the timing of retirement could play in helping to explain the drop.

Our predictions are also consistent with the conjecture that the drop in consumption is anticipated (Hurd and Rohwedder (2006), Ameriks, Caplin and Leahy (2007)). While the precise date of retirement is a random variable that takes individuals in our model by surprise, the drop in consumption upon retirement is all part of a rational, forward-looking plan. Individuals in our model at time zero cannot say for sure how big the drop will be, but they can say how big the drop will be conditional on the date of retirement.

In addition, retirement timing uncertainty is a powerful channel that may help to explain precautionary savings balances that otherwise seem large. For instance, Scholz, Seshadri and Khitatrakun (2006) estimate that as much as 80% of Americans in the HRS have asset balances that *exceed* the optimal amount of savings from a life-cycle optimization perspective. In their model households face longevity risk, earnings risk, and medical expense risk but the date of retirement is known with certainty. In our baseline calibration with timing uncertainty only (no disability risk), individuals in their 50's who live with retirement timing uncertainty would accumulate between 18% to 37% more savings by that age than otherwise identical individuals who know that they will retire at the expected age of 65. In other words, a significant portion of observed savings for retirement may actually be due to uncertainty about the date of retirement.²⁷

²⁶For instance, see Hamermesh (1984), Mariger (1987), Bernheim, Skinner and Weinberg (2001), Hurd and Rohwedder (2006), Hurst (2006), Haider and Stephens (2007), and Ameriks, Caplin and Leahy (2007) among others.

²⁷We obtain these estimates as follows. We compare asset holdings for two individuals, one who *knows* he will retire at age 65 (which is model time $t = 0.533$), and one who *expects* to retire at age 65 but faces uncertainty about the retirement date. In both cases, we assume the individual knows that he will not be disabled when he retires, $d = 0$. If the individual knows the retirement date $t = 0.533$ and the disability status $d = 0$, then he consumes $c(z|t, d) = c(z|0.533, 0)$, which is based on an initial wealth endowment $k(0|t, d) = k(0|0.533, 0)$. For comparison with the risky world, we use this consumption path to compute an asset path $a(z)$ with initial condition $a(0) = 0$ and law of motion

$$\frac{da(z)}{dz} = ra(z) + (1 - \tau)w(z) - c(z|0.533, 0) \text{ for } z \leq 0.533.$$

Then, the amount of additional savings that can be attributed to the precautionary motive to hedge retirement timing risk

Finally, the full welfare cost Δ to individuals who live with retirement timing uncertainty and no insurance is 4.26%. That is, the individual would be willing to give up 4.26% of his total lifetime consumption in order to fully insure the timing uncertainty and thereby live in a safe world with comparable expected wealth. Moreover, the timing premium alone is $\Delta_0 = 2.95\%$, which is the fraction of total lifetime consumption that he would give up just for early information about the timing of the shock.²⁸

Given the size of the welfare cost of timing uncertainty, it is natural to consider whether the predominant social insurance arrangement presently in place (Social Security) succeeds or fails to mitigate this cost, and to consider alternative arrangements that could potentially do better. This is the subject of the next subsection of the paper.

5.2. Policy experiments

In this subsection we consider four insurance arrangements: (1) U.S. Social Security retirement insurance, (2) first-best insurance that perfectly protects the individual from timing risk, (3) a simple policy in which benefits are completely independent of the individual's earnings history, and (4) a hybrid system as in Japan, the UK, Spain, and other countries in Europe with a benefit component that is unrelated to earnings and a component that is earnings based. Our goal is to evaluate whether Social Security can mitigate the cost of timing uncertainty, and to compare Social Security to the simple policy and the hybrid system.

Our first policy experiment is to add Social Security taxes and retirement benefits to the model. When we do this, the baseline welfare cost Δ falls from 4.26% without Social Security to 4.05% with Social Security, and the timing premium drops from $\Delta_0 = 2.95\%$ without Social Security to $\Delta_0 = 2.80\%$ with Social Security. Thus, Social Security reduces the welfare cost of timing uncertainty by a small amount. However, this small reduction is not because Social Security is providing timing *insurance*. Instead, Social Security boosts the individual's expected wealth in our baseline calibration, making him less sensitive to

is $k_1^*(z)/a(z) - 1$ for $z \leq 0.533$.

²⁸The welfare cost of retirement timing uncertainty is much larger if people do not accumulate precautionary savings balances. To see this, consider an individual who incorrectly assumes that he will retire with certainty at the mean age of 65 (model time $t = 0.533$). He therefore follows the optimal consumption path conditional on this retirement date, $c(z|t, d) = c(z|0.533, 0)$, where we continue to assume temporarily that there is no risk of disability. The individual follows this path, rather than the optimal path $c_1^*(z)$, for all z before shock date t , at which point he depletes his available wealth in an optimal, deterministic way over the remainder of the life cycle. To compute the welfare cost of retirement timing uncertainty, we compute the timing premium Δ_0 as usual but with $c(z|0.533, 0)$ replacing $c_1^*(z)$ in the calculation of expected utility. We find $\Delta_0 = 4.17\%$, as opposed to $\Delta_0 = 2.95\%$ when the individual self insures.

retirement risk.

Recall that we have modeled Social Security from the perspective of a single individual. This means that while Social Security taxes and benefits may need to balance at the aggregate level, at the individual level Social Security can have a non-zero effect on expected wealth. For the individual, the expected net present value of participating in Social Security (i.e., Social Security’s contribution to expected wealth) is

$$\mathbb{E}(NPV_{SS}) = - \int_0^{t'} \phi(t) \int_0^t e^{-rv} \tau w(v) dv dt + \int_0^{t'} \phi(t) SS(t|0) e^{-rt} dt.$$

At our baseline calibration this quantity is positive, which in turn means that a given loss in wage income is relatively small compared to when there is no Social Security program in place. This wealth effect explains why the welfare cost of retirement timing uncertainty is a little lower when Social Security exists than when it does not. But, as we will explain below, Social Security does not really help to insure the individual against retirement timing risk in a substantive way.

The U.S. Social Security system is among the smallest in the OECD. Only Switzerland, Canada, and Korea have slightly lower public pension tax rates. The average OECD rate is about twice the U.S. rate. Countries such as Austria, Finland, Greece, Turkey, and Germany are close to the mean, while Poland, Italy, Czech Republic, and the Netherlands all have rates that exceeds 30%. We run the experiment of doubling the size of the Social Security program in our model by doubling the tax rate τ and doubling benefits $SS(t|0)$. Doing this causes our baseline, full insurance welfare cost to drop a little further to $\Delta = 3.87\%$ and it causes the timing premium to drop to $\Delta_0 = 2.66\%$. Hence, even a very large Social Security system would not necessarily provide ample insurance against retirement timing uncertainty. The size of the system is not really the issue, it is the structure that prevents it from providing much insurance.

To make this point, suppose the individual participates in a first-best social insurance arrangement rather than Social Security. By “first-best” we mean that the individual is perfectly insured against retirement timing uncertainty by collecting a lump-sum payment $FB(t)$ upon retirement at t . We continue to assume wages are taxed at rate $\tau = 10.6\%$. The magnitude of this lump-sum payment is selected to make the individual indifferent about when the retirement shock is realized; and, to make a fair comparison with Social Security, we assume $FB(t)$ is wealth-neutral relative to Social Security in an expectation sense

(see Appendix F for full details). This gives

$$FB(t) = FB(0)e^{rt} + \int_0^t \left[rY(v) - \frac{dY(v)}{dv} - (1 - \tau)w(v) \right] e^{r(t-v)} dv$$

where

$$FB(0) = \int_0^{t'} \phi(t)SS(t|0)e^{-rt} dt - \int_0^{t'} \phi(t) \int_0^t \left[rY(v) - \frac{dY(v)}{dv} - (1 - \tau)w(v) \right] e^{-rv} dv dt.$$

Figure 6 plots $FB(t)$ versus $SS(t|0)$. Recall that both quantities represent the present value of retirement benefits as of the retirement date t . Notice that the first-best social insurance arrangement provides the individual with a big payment if he draws an early retirement shock, and a small payment if he draws a late shock. On the other hand, Social Security does just the reverse because of the positive relationship between benefits and wage earnings: individuals who suffer early retirement shocks must potentially include some zeros in the calculation of their average earnings, while individuals who draw late shocks have the highest possible average earnings. In this sense, Social Security is anti-insurance because it pays good in good states and it pays bad in bad states.

The obvious drawback, however, is that the first-best insurance arrangement creates a disincentive to work. A compromise between the first-best and the current system would be to make benefits independent of earnings. This would eliminate distortions to labor choices and also eliminate the implicit penalty on early retirement shocks. We will show that making retirement benefits completely independent of earnings can mitigate about one-third of the welfare costs of retirement timing uncertainty. We continue to hold taxes fixed at rate $\tau = 10.6\%$ on wage income, but with the twist that the individual collects the same benefits no matter when he draws the retirement shock. As with the other arrangements, we utilize the assumption that capital markets are complete by endowing the individual with a lump sum $SP(t)$ at retirement age t that reflects the value at t of a flow of benefits that start at age 65 (see Appendix G for a full explanation)

$$SP(t) = \frac{\int_0^{t'} \phi(t)SS(t|0)e^{-rt} dt \times \int_{40/75}^1 e^{r(t-v)} dv}{\int_0^{t'} \phi(t) \left(\int_{40/75}^1 e^{-rv} dv \right) dt}.$$

As with first-best insurance, we have parameterized the simple policy to be wealth-neutral relative to Social Security in order to make a fair comparison.

The baseline welfare cost of retirement timing uncertainty drops from 4.26% without insurance to 2.74% with the simple policy, and the timing premium drops from 2.95% without Social Security to 1.84% with the simple policy. In other words, simply breaking the link between benefits and earnings

would significantly increase the insurance value of Social Security.

If breaking the link is not politically feasible or desirable, it still is possible to provide partial coverage against retirement timing uncertainty while also encouraging labor force participation. To see this, consider a hybrid system that requires the same taxes during the working period but whose benefits are a convex combination of the U.S. Social Security retirement system and our simple policy. We assume a 50-50 split,

$$HY(t) = \frac{1}{2}SS(t|0) + \frac{1}{2}SP(t).$$

With this hybrid system in place, the baseline welfare cost of retirement timing uncertainty is 3.34%, and the timing premium is 2.27%. The hybrid system isn't able to match the cost reduction associated with the simple policy, but it clearly provides better insurance against retirement timing risk than the current Social Security system.

6. Disability

To provide a more comprehensive evaluation of the Social Security program's overall role in mitigating retirement uncertainty, we extend the model to include disability risk and a disability component within the Social Security program. In the extended model, individuals not only face uncertainty about the timing of retirement, they also face uncertainty about their disability status upon retirement. If the individual draws a disability shock along with the retirement shock, then he is unable to earn any part-time income during retirement. If the individual draws a retirement shock only (for instance, because of a plant closing), then he is able to collect part-time income after retirement. The former individual collects disability benefits as a bridge until he begins collecting retirement benefits. The latter individual collects retirement benefits only.

Figure 7 plots life-cycle consumption when the individual faces retirement timing risk and disability risk, and he participates in a Social Security program that includes a disability component in addition to a retirement component. Again, as with Figure 5, although retirement timing is a continuous random variable, we show just a few of the potential realizations in order to keep the picture informative. For each retirement shock date, we plot two c_2^* profiles. One profile corresponds to an individual who also draws a disability shock in addition to a retirement shock, and the other corresponds to an individual who does not draw a disability shock. The first individual collects disability benefits but no income from part-time work, while the second individual collects income from part-time work after retirement and no disability benefits.

For relatively late retirement shock dates (for example, beyond age 65), drawing the disability shock causes a loss in part-time income after retirement and does not lead to the payment of any disability benefits because the individual is already at the age in which he can collect Social Security retirement benefits. For these individuals, disability has a strictly negative effect on lifetime wealth. It is therefore intuitive that a retirement shock that is coupled with a disability shock causes a much bigger downward correction in consumption than a retirement shock alone would cause.

For early retirement shock dates, drawing the disability shock causes competing effects on lifetime wealth. On the one hand it reduces wealth because of lost earnings capacity after retirement, but on the other hand the individual collects disability benefits. If the shock date is early enough (age 45, for example), then the second effect can dominate and therefore disability benefits are generous enough that they more than replace lost part-time income in retirement.

Under our calibration, recall that the probability of becoming disabled upon retirement is much higher for those who draw an early retirement shock than for those who draw a late retirement shock. Because of this, disability insurance almost perfectly offsets the added disability risk that the individual faces, but it does *not* offset the timing risk. When we compute the joint welfare cost of timing risk and disability risk, while including both Social Security retirement and disability insurance, we get $\Delta = 3.94\%$. This is almost the same as when there is only timing risk and Social Security retirement benefits in the model ($\Delta = 4.05\%$). In other words, adding a second layer of risk and a second insurance component leaves the welfare cost almost unchanged, which suggests that the second insurance component is insuring the second risk but not the first risk. Finally, the portion of the full insurance welfare cost ($\Delta = 3.94\%$) that can be attributed to the timing premium is $\Delta_0 = 2.72\%$.

In sum, disability insurance seems to do a very good job of solving the disability risk problem but not the timing risk problem. That is, it does a nice job of replaing lost post-retirement (part-time) income due to the inability to work, but it does not solve the problem that the individual doesn't know when such a shock might strike. All of the welfare costs that we have discussed throughout the paper are summarized in Table 5. Finally, as a robustness check we redo all of our welfare calculations in Appendix H, under the assumption that the standard deviation of retirement timing uncertainty is significantly smaller than our baseline estimate. All of our essential conclusions continue to hold.

Table 5. Summary of Welfare Costs of Retirement Timing Risk & Disability Risk

Panel A: Timing Risk Only

	Full Insurance (Δ)	Timing Premium (Δ_0)
No Social Security	4.26%	2.95%
SS retirement only	4.05%	2.80%
Simple policy rule	2.74%	1.84%
50-50 hybrid policy	3.34%	2.27%

Panel B: Timing Risk and Disability Risk

	Full Insurance (Δ)	Timing Premium (Δ_0)
SS retirement & DI	3.94%	2.72%

7. Conclusion

There is a large literature that measures and assesses the economic impact of various life-cycle risks such as mortality risk, asset return risk, idiosyncratic earnings risk, and temporary unemployment risk, but less attention has been paid to retirement uncertainty. We document that many individuals end up retiring earlier or later than planned, by at least a few years, which can have dramatic consequences for lifetime budgeting. For instance, an individual who draws a one-standard deviation retirement shock of 6 years and retires unexpectedly at age 59 instead of 65 loses about one-sixth of his total (non-discounted) wage earnings. Moreover, the smaller amount of earnings must be spread over a longer retirement period. Not knowing when such a shock might strike makes planning for retirement a difficult task.

We build a detailed microeconomic model that involves dynamic decision making under uncertainty about the timing of retirement and uncertainty about one's potential for earning part-time income after retirement. We calibrate the following model features to our own estimates from a variety of data sources: survival probabilities are estimated from the Social Security cohort mortality tables; wage earnings are estimated from the 2014 CPS; the retirement timing p.d.f. is calibrated to match our estimate of the standard deviation between planned and actual retirement ages in the HRS; post-retirement earnings

are estimated from the HRS; the Social Security retirement and disability programs are calibrated to match the relevant aspects of the U.S. system; and, the probability of becoming disabled conditional on retirement is estimated from the HRS.

We use the calibrated model to compute the welfare cost of retirement timing risk. We find that the cost is quite large. Individuals would be willing to pay 4% of their total lifetime consumption to fully insure themselves against retirement timing risk. In fact, individuals would pay 3% just to know their date of retirement.

Finally, we consider the role of the Social Security retirement program in mitigating timing uncertainty. We find that Social Security retirement benefits provide almost no protection against timing risk. We also consider the role of the Social Security disability program in mitigating timing uncertainty. We find that disability insurance almost completely protects against the risk of lost part-time income during retirement, but it doesn't provide much protection against timing risk. In short, retirement timing risk is a large and costly risk that has not received very much attention in the literature, and existing social insurance arrangements do not already deal adequately with this risk.

References

- Alonso-Ortiz, Jorge.** 2014. “Social Security and Retirement across the OECD.” *Journal of Economic Dynamics and Control*, 47: 300–316.
- Ameriks, John, Andrew Caplin, and John Leahy.** 2003. “Wealth Accumulation and the Propensity to Plan.” *Quarterly Journal of Economics*, 118: 1007–1047.
- Ameriks, John, Andrew Caplin, and John Leahy.** 2007. “Retirement Consumption: Insights from a Survey.” *Review of Economics and Statistics*, 82: 265–274.
- Autor, David.** 2011. “The Unsustainable Rise of the Disability Roles in the United States: Causes, Consequences, and Policy Options.” NBER Working Paper.
- Banks, James, Richard Blundell, and Sarah Tanner.** 1998. “Is There a Retirement-Savings Puzzle?” *American Economic Review*, 88: 769–788.
- Benítez-Silva, Hugo, and Debra S. Dwyer.** 2005. “The Rationality of Retirement Expectations and the Role of New Information.” *The Review of Economics and Statistics*, 87: 587–592.
- Benítez-Silva, Hugo, Debra S. Dwyer, Wayne-Roy Gayle, and Thomas J. Muench.** 2008. “Expectations in Micro Data: Rationality Revisited.” *Empirical Economics*, 34: 381–416.
- Bernheim, B. Douglas.** 1989. “The Timing of Retirement: A Comparison of Expectations and Realizations.” In *David Wise (Ed.), The Economics of Aging*, University of Chicago Press.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg.** 2001. “What Accounts for the Variation in Retirement Wealth among U.S. Households.” *American Economic Review*, 91: 832–857.
- Caliendo, Frank N., Aspen Gorry, and Sita Slavov.** 2015. “The Cost of Uncertainty about the Timing of Social Security Reform.” Utah State University Working Paper.
- Campbell, John Y.** 2006. “Household Finance.” *Journal of Finance*, 61: 1553–1603.
- Casanova, Maria.** 2013. “Revisiting the Hump-Shaped Wage Profile.” UCLA Working Paper.
- Dwyer, Debra S., and Jianting Hu.** 1999. “Retirement Expectations and Realizations: The Role of Health Shocks and Economic Factors.” In *Olivia Mitchell, P. Brett Hammond, and Anna M. Rappaport (Eds.), Forecasting Retirement Needs and Retirement Wealth*, University of Pennsylvania Press.

- Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki.** 2014. “How Much Would You Pay to Resolve Long-Run Risk.” *American Economic Review*, forthcoming.
- French, Eric.** 2005. “The Effects of Health, Wealth, and Wages on Labour Supply and Retirement Behaviour.” *Review of Economic Studies*, 72: 395–427.
- Grochulski, Borys, and Yuzhe Zhang.** 2013. “Saving for Retirement with Job Loss Risk.” *Economic Quarterly*, 99: 45–81.
- Haider, Steven J., and Melvin Stephens.** 2007. “Is There a Retirement-Consumption Puzzle? Evidence Using Subjective Retirement Expectations.” *Review of Economics and Statistics*, 89: 247–264.
- Hamermesh, Daniel.** 1984. “Consumption During Retirement: The Missing Link in the Life Cycle.” *Review of Economics and Statistics*, 66: 1–7.
- Hurd, Michael, and Susann Rohwedder.** 2008. “The Retirement Consumption Puzzle: Actual Spending Change in Panel Data.” NBER Working Paper 13929.
- Hurd, Michael D., and Susann Rohwedder.** 2006. “Some Answers to the Retirement-Consumption Puzzle.” NBER Working Paper.
- Hurst, Erik.** 2006. “Grasshoppers, Ants and Pre-Retirement Wealth: A Test of Permanent Income Consumers.” University of Chicago Working Paper.
- Lucas, Robert E.** 2003. “Macroeconomic Priorities.” *American Economic Review*, 93: 1–14.
- Lusardi, Annamaria, and Olivia Mitchell.** 2007. “Baby Boomer Retirement Security: The Roles of Planning, Financial Literacy, and Housing Wealth.” *Journal of Monetary Economics*, 54: 205–224.
- Lusardi, Annamaria, and Olivia Mitchell.** 2008. “Planning and Financial Literacy: How Do Women Fare?” *American Economic Review*, 98: 413–417.
- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia Mitchell.** 2011. “Optimal Financial Literacy and Saving for Retirement.” Pension Research Council Working Paper.
- Mariger, Randall P.** 1987. “A Life-Cycle Consumption Model with Liquidity Constraints: Theory and Empirical Results.” *Econometrica*, 55: 533–557.
- McGarry, Kathleen.** 2004. “Health and Retirement: Do Changes in Health Affect Retirement Expectations?” *Journal of Human Resources*, 39: 624–648.

- Rogerson, Richard, and Johanna Wallenius.** 2009. "Micro and Macro Elasticities in a Life Cycle Model with Taxes." *Journal of Economic Theory*, 144: 2277–2292.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun.** 2006. "Are Americans Saving 'Optimally' for Retirement?" *Journal of Political Economy*, 114: 607–643.
- Smith, Sarah.** 2006. "The Retirement-Consumption Puzzle and Involuntary Early Retirement: Evidence from the British Household Panel Survey." *Economic Journal*, 116: C130–C148.
- Stokey, Nancy L.** 2014. "Wait-and-See: Investment Options under Policy Uncertainty." University of Chicago Working Paper.
- Szinovacz, Maximiliane, and Adam Davey.** 2005. "Predictors of Perceptions of Involuntary Retirement." *The Gerontologist*, 45: 36–47.
- van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie.** 2012. "Financial Literacy, Retirement Planning and Household Wealth." *Economic Journal*, 122: 449–478.
- Vidangos, Ivan.** 2009. "Household Welfare, Precautionary Saving, and Social Insurance under Multiple Sources of Risk." Federal Reserve Board Working Paper.
- Zayatz, Tim.** 2011. "Social Security Disability Insurance Program Worker Experience." Social Security Administration Actuarial Study 122.

Technical appendices

Appendix A: Measuring retirement uncertainty

This appendix describes the construction of the variables measuring an individual’s expected retirement age ($Eret$) and actual age at retirement (Ret), together with the computation of the standard deviation of $X = (Eret - Ret)$.

As described in Section 2, we use a sample of male respondents aged 51 to 61 in the first wave of the Health and Retirement Study (HRS). There are 4,541 male respondents in this age group in wave 1. Out of these, we drop 864 individuals whose retirement expectations were not elicited because they were already retired, disabled, or out of the labor force; 255 individuals for whom the retirement expectation is missing; and 175 individuals who are unemployed, and hence would be considered retired according to our definition below. This leaves us with 3,251 observations of the variable $Eret$. The details of sample selection are summarized in Table 6.

Table 6. Sample Selection for Variable $Eret$

Males Aged 51 to 61 in wave 1		4,545
Work status missing	4	
Unemployed	175	
Retired	613	
Disabled	189	
Not in the labor force	47	
Total dropped because not employed	1,028	
Males Aged 51 to 61 and Employed in wave 1		3,517
Proxy interview ($Eret$ not asked)	244	
Already retired	15	
Other missing	7	
Total dropped because of missing $Eret$	266	
Males Aged 51 to 61, Employed, and $Eret$ observed in wave 1 (Final Sample)		3,251

To be consistent with the wording of the questions used by the HRS to elicit retirement expectations, we define retirement as working zero hours. We follow individuals over time, and construct the variable

Ret using information on the month and year when they left their last job prior to retirement. There is a small number of observations (102, or 3% of the total sample) for which we do not observe the actual retirement year, but for which it is possible to obtain both an upper and a lower bound of their retirement date. We make the conservative assumption that they retired on the date within that interval that is closest to $Eret$.

If either the variable $Eret$ or Ret are measured with error, this will increase the standard deviation of X , and in turn overstate our measure of retirement uncertainty. We are particularly concerned about measurement error in the variable $Eret$. HRS respondents are allowed to report their expected retirement time as both an age or a specific year. All responses are then transformed into a retirement year, and this process is bound to generate some rounding error. We deal with this issue by allowing for plus/minus one year of error in $Eret$. We compute the variable X as

$$\min\{|(Eret - 1) - Ret|, |Eret - Ret|, |(Eret + 1) - Ret|\}.$$

Table 7 describes retirement outcomes as a function of retirement expectations in wave 1. There are 2,449 individuals in the sample, shown in column 1, who expect to retire before the end of the HRS panel. 1,893 (77%) of those actually retire within that period; 244 (10%) are still employed by the time they reach their expected retirement age, but their actual retirement age cannot be established because of attrition, truncation of retirement date, or death; 102 (4%) die and 210 (9%) are lost to attrition before their expected retirement date. The second column shows 17 individuals who expect to retire after the last wave in the HRS panel. 10 (59%) of those retire during the sample period, 2 (12%) die before the end of the panel, and the remaining 5 (29%) remain employed by the time they leave the sample. Column 3 shows retirement outcomes for 475 individuals who state on the first wave that they will never retire. 324 (65%) eventually retire before the end of the panel, while the remaining 35% are still employed when they exit the sample due to death, attrition, or truncation. Finally, the last column shows retirement outcomes for 310 individuals who state that they do not know when they will retire. 212 (68%) of those retire during the sample period, and the remaining 22% remain employed when last observed in the sample.

Table 7. Retirement Outcomes by *Eret* Category

	<i>Eret</i>			
	Expect to retire by wave 11	Expect to retire after wave 11	Will never retire	DK if they will retire
Retire during sample period	1,893	10	324	212
Work past <i>Eret</i> , retirement age not observed	244			
Die before <i>Eret</i>	102	2		
Exit sample before <i>Eret</i>	210			
Employed by last wave observed in the sample		5	151	98
Total	2,449	17	475	310

The value of the variable X can be computed directly from the data for individuals for whom both $Eret$ and Ret are observed. In cases when one of those two variables is missing, we can sometimes make a conservative assumption that allows us to assign a value to the variable X . Table 8 describes these assumptions in detail. Row 1 shows that X is computed as the difference between the expected and actual retirement age for the 1,903 (58% of the sample) individuals for whom both $Eret$ and Ret observed. The 244 (8%) individuals in row 2 are still employed by the time they reach their expected retirement age, so we know that they have made a mistake in their predictions. However, because of truncation or attrition they leave the sample before their retirement age can be observed, and the exact size of the difference between $Eret$ and Ret cannot be established. To be as conservative as possible, we assume that those individuals retire the first year after exiting the sample. The 5 (0%) individuals in row 3 expect to retire after the sample period and are still employed by the time they exit the panel. Because we have no evidence that they have made a mistake in their predictions, we assign a value of 0 to the variable X for this group. Row 4 shows 104 (3%) individuals who die before reaching their expected retirement age. We do not use these individuals in the computation of retirement timing uncertainty, as mortality risk is modeled separately. Row 5 shows 210 (6%) individuals who exit the sample because of truncation or attrition before their expected retirement age. Because we cannot establish whether they have made a mistake in their prediction, and any assumption to that regard would be *ad hoc*, we do not use these

individuals in the computation of uncertainty either.

The next two rows represent individuals who say they will never retire. For those in row 6 (324, or 10%) retirement is observed. We compute the size of the difference between their expected and actual retirement ages by subtracting the latter from the average life expectancy for this cohort, which is 76.5 years of age. Those in row 7 (151 or 5%) die or leave the sample before retirement is observed, and we assume the size of their mistake is 0.

Finally, individuals in the last two rows (310 or 10%) say they do not know when they will retire. It is particularly difficult to assign a value to the variable X without making *ad-hoc* assumptions, as we have no way of telling what their expected retirement age is. However, their eventual retirement behavior closely mirrors that of those who say they will never retire. The proportion retiring in every wave of the panel, as well as the proportion whose retirement is not observed during the sample period, are essentially the same for the two groups. Therefore, we compute X in the same way for the two groups.

Table 8. Computation of $X = Eret - Ret$

	X computed as	N
<i>Eret</i> observed		
1. <i>Ret</i> observed	$(Eret - Ret)$	1,903
2. Work past <i>Eret</i> , <i>Ret</i> not observed	$Eret - (\text{Age in last wave in sample} + 1)$	244
3. <i>Eret</i> is after sample period, <i>Ret</i> not observed	0	5
4. Dies or leaves sample before <i>Eret</i>	Not used	104
5. Leaves sample before <i>Eret</i>	Not used	210
Will never retire		
6. <i>Ret</i> observed	$(\text{Average life expectancy} - Ret)$	324
7. <i>Ret</i> not observed	0	151
DK when they will retire		
8. <i>Ret</i> observed	$(\text{Average life expectancy} - Ret)$	213
9. <i>Ret</i> not observed	0	97
Total		3,251

Table 9 shows the value of the standard deviation of X for different subsamples. The first column considers the baseline subsample of individuals aged 51 to 61 in wave 1. Within this age group, using

only individuals for whom both expected and actual retirement are observed (row 1) yields a standard deviation of 4.28. Adding individuals who work past their expected retirement age and for whom X is computed as discussed in Table 8, the standard deviation increases to 5.05 (row 2). Row 3 adds individuals who do not expect to retire before the end of the sample period and whose retirement is indeed not observed before that date. Because we are assuming that they make no mistakes in their predictions, the standard deviation decreases slightly, to 5.04. Column 4 adds individuals who say they will never retire, but whose retirement is observed. Assuming they expected to work until death, and using the average life expectancy for the cohort, increases the standard deviation to 6.54. Finally, adding individuals who do not expect to retire and who are still employed by the time they exit the sample reduces the standard deviation to 6.35.

The second and third columns of Table 9 compute the standard deviation for a younger (51 to 55) and an older (56 to 61) age group within the baseline sample. This computation is carried out to illustrate that retirement uncertainty declines slowly as retirement approaches, even for age groups very close to retirement age. The two age groups considered here are 5 years apart, on average, but the standard deviation of the variable X declines only between half a year and one year for the older group.

Table 9. Standard Deviation of X for Different Subsamples

	Sample	Baseline Age 51 to 61	Age 51 to 55	Age 56 to 61
1	<i>Ret</i> observed	4.28	4.59	3.88
2	1 + Work past <i>Eret</i> , <i>Ret</i> not observed	5.05	5.26	4.78
3	2 + <i>Eret</i> after sample period, <i>Ret</i> not observed	5.04	5.25	4.77
4	3 + Will never retire, <i>Ret</i> observed	6.54	6.93	6.05
5	4 + Will never retire, <i>Ret</i> not observed	6.35	6.73	5.88
6	5 + DK when they will retire, <i>Ret</i> observed	6.92	7.37	6.37
7	6 + DK when they will retire, <i>Ret</i> not observed	6.82	7.24	6.29

Appendix B: Solution to individual optimization problem

The individual's problem is solved recursively as in Caliendo, Gorry and Slavov (2015) and Stokey (2014) but modified extensively to fit the current setting.²⁹

Step 1. The deterministic retirement problem

The optimal consumption path $c(z)$ for $z \in [t, T]$ after the retirement shock has hit at date t solves

$$\max_{c(z)_{z \in [t, T]}} : \int_t^T e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz,$$

subject to

$$\frac{dK(z)}{dz} = rK(z) - c(z), \text{ for } z \in [t, T],$$

$$t \text{ and } d \text{ given, } K(t) = k(t) + B(t, d) \text{ given, } K(T) = 0.$$

It is straightforward to show that the solution to this deterministic control problem is

$$c_2^*(z|t, k(t), d) = \frac{(k(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}, \text{ for } z \in [t, T].$$

This solution, for an arbitrary $k(t)$ and for given realizations of t and d , will be nested in the continuation function in the next step.

Step 2. The time zero stochastic problem

Facing random variables t and d , at time zero the individual seeks to maximize expected utility

$$\max_{c(z)_{z \in [0, t']}} : \mathbb{E}_{t, d} \left[\int_0^t e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz + \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k(t), d)^{1-\sigma}}{1-\sigma} dz \right]$$

which can be rewritten as

$$\max_{c(z)_{z \in [0, t']}} : \int_0^{t'} \int_0^t \phi(t) e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz dt + \int_0^{t'} \left(\sum_d \theta(d|t) \phi(t) \mathcal{S}(t, k(t), d) \right) dt$$

²⁹Relative to Caliendo, Gorry and Slavov (2015) and Stokey (2014), the current paper has the added complication that the timing density is truncated, which in turn renders the usual Pontryagin first-order conditions insufficient to identify a unique optimum. We will elaborate more below.

where

$$\mathcal{S}(t, k(t), d) = \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k(t), d)^{1-\sigma}}{1-\sigma} dz.$$

Using a change in the order of integration, i.e., $\int_0^{t'} \int_0^t (\cdot) dz dt = \int_0^{t'} \int_z^{t'} (\cdot) dt dz$, we can write

$$\begin{aligned} \int_0^{t'} \int_0^t \phi(t) e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz dt &= \int_0^{t'} \int_z^{t'} \phi(t) e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dt dz \\ &= \int_0^{t'} [1 - \Phi(z)] e^{-\rho z} \Psi(z) \frac{c(z)^{1-\sigma}}{1-\sigma} dz \\ &= \int_0^{t'} [1 - \Phi(t)] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} dt. \end{aligned}$$

Using this result we can state the stochastic problem as a standard Pontryagin problem

$$\max_{c(t)_{t \in [0, t']}] : \int_0^{t'} \left\{ [1 - \Phi(t)] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \sum_d \theta(d|t) \phi(t) \mathcal{S}(t, k(t), d) \right\} dt$$

subject to

$$\mathcal{S}(t, k(t), d) = \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k(t), d)^{1-\sigma}}{1-\sigma} dz,$$

$$\frac{dk(t)}{dt} = rk(t) + (1 - \tau)w(t) - c(t),$$

$$k(0) = 0, k(t') \text{ free},$$

$$c_2^*(z|t, k(t), d) = \frac{(k(t) + B(t, d))e^{-\tau t}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}, \text{ for } z \in [t, T].$$

To solve, form the Hamiltonian \mathcal{H} with multiplier $\lambda(t)$

$$\mathcal{H} = [1 - \Phi(t)] e^{-\rho t} \Psi(t) \frac{c(t)^{1-\sigma}}{1-\sigma} + \sum_d \theta(d|t) \phi(t) \mathcal{S}(t, k(t), d) + \lambda(t) [rk(t) + (1 - \tau)w(t) - c(t)].$$

The necessary conditions include

$$\frac{\partial \mathcal{H}}{\partial c(t)} = [1 - \Phi(t)] e^{-\rho t} \Psi(t) c(t)^{-\sigma} - \lambda(t) = 0$$

$$\frac{d\lambda(t)}{dt} = -\frac{\partial \mathcal{H}}{\partial k(t)} = -\sum_d \theta(d|t) \phi(t) \frac{\partial \mathcal{S}(t, k(t), d)}{\partial k(t)} - \lambda(t)r,$$

where the usual transversality condition $\lambda(t') = 0$ is automatically satisfied by the Maximum Condition

(since $\Phi(t') = 1$ by definition). Note that

$$\begin{aligned}
\frac{\partial \mathcal{S}(t, k(t), d)}{\partial k(t)} &= \int_t^T e^{-\rho z} \Psi(z) [c_2^*(z|t, k(t), d)]^{-\sigma} \frac{\partial c_2^*(z|t, k(t), d)}{\partial k(t)} dz \\
&= \int_t^T e^{-\rho z} \Psi(z) \left[\frac{(k(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma} \right]^{-\sigma} \frac{e^{-rt} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} dz \\
&= \left[\frac{(k(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt}.
\end{aligned}$$

Using this result, together with the Maximum Condition, we can rewrite the multiplier equation as

$$\frac{d\lambda(t)}{dt} = - \sum_d \theta(d|t) \phi(t) \left[\frac{(k(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt} - [1 - \Phi(t)] e^{-\rho t} \Psi(t) c(t)^{-\sigma} r.$$

Now differentiate the Maximum Condition with respect to t

$$-\phi(t) \left\{ [e^{-\rho t} \Psi(t)] c(t)^{-\sigma} \right\} + [1 - \Phi(t)] \left\{ [-\rho e^{-\rho t} \Psi(t) + e^{-\rho t} \Psi'(t)] c(t)^{-\sigma} - \sigma [e^{-\rho t} \Psi(t)] c(t)^{-\sigma-1} \frac{dc(t)}{dt} \right\} = \frac{d\lambda(t)}{dt}$$

and combine the previous two equations and solve for $dc(t)/dt$

$$\frac{dc(t)}{dt} = \left(\frac{c(t)^\sigma e^{(\rho-r)t}}{\sigma \Psi(t)} \sum_d \theta(d|t) \left[\frac{(k(t) + B(t, d))e^{-rt}}{\int_t^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} - \frac{1}{\sigma} \right) \left(\frac{c(t)\phi(t)}{1 - \Phi(t)} \right) + \left[\frac{\Psi'(t)}{\Psi(t)} + r - \rho \right] \frac{c(t)}{\sigma},$$

which matches the Euler equation stated in the body of the paper.

The Euler equation, together with the law of motion for savings dk/dt and the initial condition $k(0) = 0$ are used to pin down solution consumption and savings *conditional* on $c(0)$, which has yet to be identified.

In general, in stochastic stopping time problems where there is no restriction on the state variable at the maximum stopping date—a setting that arises naturally if the timing p.d.f. is truncated—the usual Pontryagin first-order conditions for optimality are not sufficient to identify a unique solution. The transversality condition is redundant and the first-order conditions therefore produce a family of potential solutions rather than a unique solution. We provide a “work-around” that works in general and is easy to use. The answer is to use the limiting case of the transversality condition, together with the other first-order conditions, to derive what we refer to as a “stochastic continuity” condition to provide the needed endpoint restriction. This extra condition allows us to identify the unique solution.

We can identify $c(0)$ as follows. Rewrite the Maximum Condition as

$$e^{-\rho t} \Psi(t) c(t)^{-\sigma} = \frac{\lambda(t)}{1 - \Phi(t)}.$$

Noting the transversality condition and properties of the c.d.f.

$$\frac{\lambda(t')}{1 - \Phi(t')} = \frac{0}{0},$$

we can use L'Hôpital's Rule on this indeterminate expression

$$\lim_{t \rightarrow t'} e^{-\rho t} \Psi(t) c(t)^{-\sigma} = \lim_{t \rightarrow t'} \frac{\lambda(t)}{1 - \Phi(t)} = \lim_{t \rightarrow t'} \frac{d\lambda(t)/dt}{-\phi(t)} = \frac{d\lambda(t')/dt}{-\phi(t')}$$

and hence we can use the following as a boundary condition in lieu of the redundant transversality condition:

$$e^{-\rho t'} \Psi(t') c(t')^{-\sigma} = \frac{d\lambda(t')/dt}{-\phi(t')}.$$

Note that

$$\frac{d\lambda(t')}{dt} = - \sum_d \theta(d|t') \phi(t') \left[\frac{(k(t') + B(t', d)) e^{-rt'}}{\int_{t'}^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt'}$$

so the new boundary condition becomes

$$e^{-\rho t'} \Psi(t') c(t')^{-\sigma} = \sum_d \theta(d|t') \left[\frac{(k(t') + B(t', d)) e^{-rt'}}{\int_{t'}^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} e^{-rt'}.$$

Simplify

$$\begin{aligned} c(t') &= \left(\sum_d \theta(d|t') \left[\frac{(k(t') + B(t', d)) e^{-rt'}}{\int_{t'}^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right]^{-\sigma} \times e^{-(r-\rho)t'} \Psi(t')^{-1} \right)^{-1/\sigma} \\ &= \left(\sum_d \theta(d|t') \left[\frac{(k(t') + B(t', d)) e^{-rt'}}{\int_{t'}^T e^{-rv + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t'/\sigma} \Psi(t')^{1/\sigma} \right]^{-\sigma} \right)^{-1/\sigma} \\ &= \left(\sum_d \theta(d|t') [c_2^*(t'|t', k(t'), d)]^{-\sigma} \right)^{-1/\sigma}. \end{aligned}$$

In sum, we choose $c(0)$ so that the Euler equation dc/dt , together with dk/dt and the initial condition $k(0) = 0$ all imply “stochastic continuity” at time t' : $c(t') = \left(\sum_d \theta(d|t') [c_2^*(t'|t', k(t'), d)]^{-\sigma} \right)^{-1/\sigma}$. Note that we literally have continuity if d is deterministic, $c(t') = c_2^*(t'|t', k(t'), d)$. For the more general case

where d is stochastic, there is continuity between marginal utility and expected marginal utility.

Appendix C: Welfare decomposition with Jensen's inequality

Here we prove using Jensen's inequality that the timing premium is smaller than the value of full insurance.

Making use of the following equations

$$\begin{aligned}
c^{NR}(t) &= k(0)G(t) \\
k(0) &= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t)k(0|t, d) \right) dt \\
G(t) &\equiv \frac{e^{(r-\rho)t/\sigma}\Psi(t)^{1/\sigma}}{\int_0^T e^{-rv+(r-\rho)v/\sigma}\Psi(v)^{1/\sigma} dv} \\
c(z|t, d) &= k(0|t, d)G(z) \\
U^{NR} &= \int_0^T e^{-\rho t}\Psi(t) \frac{c^{NR}(t)^{1-\sigma}}{1-\sigma} dt \\
EU(t, d) &= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^T e^{-\rho z}\Psi(z) \frac{c(z|t, d)^{1-\sigma}}{1-\sigma} dz \right) \right) dt,
\end{aligned}$$

we note that

$$\begin{aligned}
U^{NR} &= \frac{k(0)^{1-\sigma}}{1-\sigma} \int_0^T e^{-\rho t}\Psi(t)G(t)^{1-\sigma} dt \\
&= \frac{\left[\int_0^{t'} (\sum_d \theta(d|t)\phi(t)k(0|t, d)) dt \right]^{1-\sigma}}{1-\sigma} \int_0^T e^{-\rho t}\Psi(t)G(t)^{1-\sigma} dt \\
EU(t, d) &= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^T e^{-\rho z}\Psi(z)G(z)^{1-\sigma} \frac{k(0|t, d)^{1-\sigma}}{1-\sigma} dz \right) \right) dt \\
&= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \frac{k(0|t, d)^{1-\sigma}}{1-\sigma} \right) dt \int_0^T e^{-\rho z}\Psi(z)G(z)^{1-\sigma} dz.
\end{aligned}$$

By Jensen's inequality,

$$\frac{\left[\int_0^{t'} (\sum_d \theta(d|t)\phi(t)k(0|t, d)) dt \right]^{1-\sigma}}{1-\sigma} > \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \frac{k(0|t, d)^{1-\sigma}}{1-\sigma} \right) dt$$

which implies $U^{NR} > EU(t, d)$ and hence $\Delta > \Delta_0$. In other words, the individual would always pay more to have his expected wealth with certainty than he would pay for retirement information, because simply knowing his wealth is not as good as insuring his wealth.

Appendix D: Leisure

Suppose period utility is additively separable in consumption c and leisure l . In keeping with our main assumption that retirement is an uncertain event, utility from leisure is now an uncertain quantity as well. Early retirement brings extra utility from leisure while late retirement erodes utility from leisure.

Without loss of generality, we normalize instantaneous leisure time to $l = 0$ before retirement and $l = 1$ after retirement. We also normalize the instantaneous utility of leisure during the working period to $u(0) = 0$. The utility of leisure during retirement is $u(1)$. We assume $u' > 0$ and $u'' < 0$. For a given retirement realization t , the total lifetime utility from leisure is $\int_t^T e^{-\rho z} \Psi(z) u(1) dz$. The additive separability of consumption and leisure implies that consumption decisions are not influenced by the presence of leisure in the utility function. Hence, the individual will continue to follow $c_1^*(z)$ for all z before the retirement date t is realized and $c_2^*(z|t, k_1^*(t), d)$ for all z after the retirement date t and disability status d are realized.

Full insurance

For the case in which the individual is fully insured against retirement uncertainty, he collects with certainty his expected wealth as before and makes optimal consumption decisions over the life cycle as before, $c^{NR}(t)$. Concerning leisure, he receives at each moment t his expected leisure at that moment

$$l^{NR}(t) = \Phi(t) \times 1 + [1 - \Phi(t)] \times 0$$

which confers period leisure utility $u(\Phi(t))$ and total leisure utility $\int_0^T e^{-\rho t} \Psi(t) u(\Phi(t)) dt$, where $\Phi(t) = 1$ for all $t \geq t'$.

Equating utility from expected wealth and expected leisure to expected utility, and then solving for

Δ (willingness to pay to avoid uncertainty), gives the full insurance value of timing uncertainty

$$\begin{aligned}
& \int_0^T e^{-\rho t} \Psi(t) \frac{[c^{NR}(t)(1-\Delta)]^{1-\sigma}}{1-\sigma} dt + \int_0^T e^{-\rho t} \Psi(t) u(\Phi(t)) dt \\
= & \int_0^{t'} \left(\sum_d \theta(d|t) \phi(t) \left(\int_0^t e^{-\rho z} \Psi(z) \frac{c_1^*(z)^{1-\sigma}}{1-\sigma} dz + \int_t^T e^{-\rho z} \Psi(z) \frac{c_2^*(z|t, k_1^*(t), d)^{1-\sigma}}{1-\sigma} dz \right) \right) dt \\
& + \int_0^{t'} \phi(t) \left(\int_t^T e^{-\rho z} \Psi(z) u(1) dz \right) dt.
\end{aligned}$$

Now performing some algebra on the last term on both the left and right sides, including a change in the order of integration on the term on the right, we have

$$\begin{aligned}
\text{II} & \equiv \int_0^T e^{-\rho t} \Psi(t) u(\Phi(t)) dt \\
& = \int_0^{t'} e^{-\rho t} \Psi(t) u(\Phi(t)) dt + \int_{t'}^T e^{-\rho t} \Psi(t) u(1) dt \\
\text{III} & \equiv \int_0^{t'} \int_t^T \phi(t) e^{-\rho z} \Psi(z) u(1) dz dt \\
& = \int_0^{t'} \int_0^z \phi(t) e^{-\rho z} \Psi(z) u(1) dt dz + \int_{t'}^T \int_0^{t'} \phi(t) e^{-\rho z} \Psi(z) u(1) dt dz \\
& = \int_0^{t'} \int_0^z \phi(t) e^{-\rho z} \Psi(z) u(1) dt dz + \int_{t'}^T e^{-\rho z} \Psi(z) u(1) dz \\
& = \int_0^{t'} e^{-\rho z} \Psi(z) u(1) \Phi(z) dz + \int_{t'}^T e^{-\rho z} \Psi(z) u(1) dz \\
& = \int_0^{t'} e^{-\rho t} \Psi(t) u(1) \Phi(t) dt + \int_{t'}^T e^{-\rho t} \Psi(t) u(1) dt.
\end{aligned}$$

Using the concavity of u and the fact that $\Phi(t) < 1$ for all $t < t'$, it must be that

$$u(\Phi(t)) > u(1)\Phi(t) \text{ for all } t < t' \implies \text{II} > \text{III}.$$

Finally, this implies that Δ must be strictly larger when we include leisure in the utility function than when we do not. Hence, we are safe to ignore leisure and treat our calculations of the welfare cost of retirement uncertainty as a lower bound. While including leisure may at first glance seem to mitigate the welfare loss of timing uncertainty because early retirement shocks are accompanied by more leisure, the additive separability of utility prevents this from happening. Instead, retirement timing uncertainty simply implies that the individual faces risk over two (unrelated) margins, consumption as well as leisure,

and the presence of the second margin only amplifies his willingness to pay to avoid uncertainty.

Timing premium

Similar arguments can be made for the timing premium. With leisure in the period utility function, the timing premium Δ_0 is the solution to the following equation

$$\begin{aligned} & \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^T e^{-\rho z} \Psi(z) \frac{[c(z|t, d)(1 - \Delta_0)]^{1-\sigma}}{1 - \sigma} dz + \int_t^T e^{-\rho z} \Psi(z) u(1) dz \right) \right) dt \\ &= \int_0^{t'} \left(\sum_d \theta(d|t)\phi(t) \left(\int_0^t e^{-\rho z} \Psi(z) \frac{c_1^*(z)^{1-\sigma}}{1 - \sigma} dz + \int_t^T e^{-\rho z} \Psi(z) \left(\frac{c_2^*(z|t, k_1^*(t), d)^{1-\sigma}}{1 - \sigma} + u(1) \right) dz \right) \right) dt. \end{aligned}$$

The leisure terms cancel out and we are left with the same timing premium Δ_0 as when we ignore leisure. This is an immediate implication of the assumption that leisure is fixed before and after retirement. Early resolution of retirement uncertainty does not change leisure allocations over the life cycle, which means the individual isn't willing to pay any more for retirement information in this case than in the case without leisure.

Appendix E: Social Security

Because the individual faces uncertainty about becoming disabled, we must model Social Security in both states.

Without disability

Suppose the individual never becomes disabled but instead retires for other reasons (such as a health shock to a spouse or parent).

Let $\bar{w}(t)$ be the individual's average wage income corresponding to the last 35 years of earnings before retirement (which is virtually equivalent to the top 35 years of earnings given the wage profile that we are using), where t is the stochastic retirement age. If the individual draws a bad enough shock, some of these years will be zeros. If the individual draws a very good shock, then the average of his last 35 years can increase because wages are lowest at age 25 in our calibration.

Let $b(\bar{w}(t))$ be the constant, flow value of Social Security benefits if claimed at age 65. The individual receives this constant flow until death. Benefits are a piecewise linear function of an individual's average wage, where the kinks (bend points) are multiples of the economy-wide average wage \bar{e} . Social Security replaces 90% of $\bar{w}(t)$ up to the first bend point, 32% of $\bar{w}(t)$ between the first and second bend points,

15% of $\bar{w}(t)$ between the second and third bend points, and 0% of $\bar{w}(t)$ beyond the third bend point. The nominal values of the bend points change each year, but Alonso-Ortiz (2014) and others assume the bend points are the following multiples of the average economy-wide wage: $0.2\bar{e}$, $1.24\bar{e}$, and $2.47\bar{e}$.

To simplify, we assume the economy-wide average wage equals the average wage of an individual who draws a retirement shock at the average age (65)

$$\bar{e} = \bar{w}(40/75),$$

which means that the flow value of benefits claimed at 65 is

$$b(\bar{w}(t)) = \begin{cases} 90\% \times \bar{w}(t) & \text{for } \bar{w}(t) \leq 0.2\bar{e} \\ 90\% \times 0.2\bar{e} + 32\% \times (\bar{w}(t) - 0.2\bar{e}) & \text{for } 0.2\bar{e} \leq \bar{w}(t) \leq 1.24\bar{e} \\ 90\% \times 0.2\bar{e} + 32\% \times (1.24\bar{e} - 0.2\bar{e}) + 15\% \times (\bar{w}(t) - 1.24\bar{e}) & \text{for } 1.24\bar{e} \leq \bar{w}(t) \leq 2.47\bar{e} \\ 90\% \times 0.2\bar{e} + 32\% \times (1.24\bar{e} - 0.2\bar{e}) + 15\% \times (2.47\bar{e} - 1.24\bar{e}) & \text{for } 2.47\bar{e} \leq \bar{w}(t). \end{cases}$$

Finally, $SS(t|d)$ is the present discounted value (as of retirement date t) of Social Security benefits, conditional on disability status. Taking advantage of our assumption that capital markets are complete, and assuming $d = 0$, we endow the individual with the following lump sum at t ,

$$SS(t|d) = SS(t|0) = \left(b(\bar{w}(t)) \times \int_{40/75}^1 e^{-r(v-40/75)} dv \right) e^{r(t-40/75)}.$$

With disability

If the individual becomes disabled, we re-use notation and assume $\bar{w}(t)$ is his average wage income corresponding to the last 35 years of earnings, where t is the stochastic retirement age, and *no* zeros are included in the average if the individual draws a timing shock that leaves him with fewer than 35 years of work experience. Moreover, he begins collecting full benefits at the moment he retires (rather than waiting until age 65). Hence

$$SS(t|d) = SS(t|1) = \max \left\{ SS(t|0), b(\bar{w}(t)) \times \int_t^1 e^{-r(v-t)} dv \right\}.$$

The max operator is to recognize that a disability shock after $t = 40/75$ (age 65) can't lead to lower benefits than a system without disability. In other words, disability leads to higher total benefits if the shock is early and has no effect on total benefits if the shock happens late.

Appendix F: First-best insurance against timing risk

Let's assume the individual participates in a first-best arrangement that perfectly insures against retirement timing uncertainty by providing a lump-sum payment $FB(t)$ upon retirement at t . We continue to assume wages are taxed at rate τ .

Suppose there is no disability risk in the model. If so, then the present value (as of time zero) of total lifetime income, as a function of the retirement date t , is

$$PV_0(t) = \int_0^t e^{-rv}(1-\tau)w(v)dv + e^{-rt}Y(t) + e^{-rt}FB(t) \text{ for all } t \in [0, t'].$$

By definition, the first-best arrangement would make the individual indifferent about when the retirement shock is realized, hence it must satisfy

$$\frac{d}{dt}PV_0(t) = 0,$$

or

$$\frac{d}{dt}PV_0(t) = e^{-rt}(1-\tau)w(t) - re^{-rt}Y(t) + e^{-rt}\frac{dY(t)}{dt} - re^{-rt}FB(t) + e^{-rt}\frac{dFB(t)}{dt} = 0.$$

Simplify

$$\frac{dFB(t)}{dt} = rFB(t) + rY(t) - \frac{dY(t)}{dt} - (1-\tau)w(t).$$

The general solution to this differential equation is

$$FB(t) = \left(C + \int^t \left[rY(v) - \frac{dY(v)}{dv} - (1-\tau)w(v) \right] e^{-rv} dv \right) e^{rt}$$

where C is a constant of integration. Evaluate at $t = 0$ and solve for C

$$C = FB(0) - \int^0 \left[rY(v) - \frac{dY(v)}{dv} - (1-\tau)w(v) \right] e^{-rv} dv$$

which gives the particular solution

$$FB(t) = FB(0)e^{rt} + \int_0^t \left[rY(v) - \frac{dY(v)}{dv} - (1-\tau)w(v) \right] e^{r(t-v)} dv.$$

Notice that the level is not pinned down; the overall generosity of the first-best arrangement is indeterminate. To make a fair comparison with Social Security, we assume the first-best arrangement is

wealth-neutral relative to Social Security in an expectation sense

$$\int_0^{t'} \phi(t)FB(t)e^{-rt}dt = \int_0^{t'} \phi(t)SS(t|0)e^{-rt}dt,$$

which pins down $FB(0)$

$$FB(0) = \int_0^{t'} \phi(t)SS(t|0)e^{-rt}dt - \int_0^{t'} \phi(t) \int_0^t \left[rY(v) - \frac{dY(v)}{dv} - (1 - \tau)w(v) \right] e^{-rv} dv dt.$$

Appendix G. Simple policy

Independent of work history, suppose the government makes a fixed payment p from 65 forward that is not a function of past earnings. Utilizing the assumption that capital markets are complete, we endow the individual with the following lump sum at retirement age t ,

$$SP(t) = \left(p \times \int_{40/75}^1 e^{-r(v-40/75)} dv \right) e^{r(t-40/75)}.$$

To make a fair comparison with Social Security, we assume the simple policy is wealth-neutral relative to Social Security in an expectation sense

$$\int_0^{t'} \phi(t)SP(t)e^{-rt}dt = \int_0^{t'} \phi(t)SS(t|0)e^{-rt}dt,$$

which implies

$$p = \frac{\int_0^{t'} \phi(t)SS(t|0)e^{-rt}dt}{\int_0^{t'} \phi(t) \left(\int_{40/75}^1 e^{-r(v-40/75)} dv \right) e^{-r40/75} dt}.$$

Appendix H. Robustness

While our baseline estimate of the standard deviation of retirement timing uncertainty (6 years) is already conservative, here we entertain the implications of our most conservative estimate from our main sample (about 4.25 years). Our existing assumption $\mathbb{E}(t) = 40/75$ (age 65), combined with our new assumption $\sqrt{\text{var}(t)} = 0.05667$ (4.25 years), together imply the following parameterization of the beta density over retirement timing risk: $\gamma = 16.91626$ and $\beta = 4.22907$.

Table 10 reports the welfare costs for all of the scenarios that we consider. Notice that the magnitude of the cost of retirement uncertainty is still significant, and the essential policy lessons continue to go through. Social Security continues to provide very little insurance against timing risk, while breaking

part or all of the benefit-earning link creates some additional protection.

Table 10. Robustness Check with Most Conservative Estimate of Timing Risk


Panel A: Timing Risk Only

standard deviation	Full Insurance (Δ)		Timing Premium (Δ_0)	
	4.25 years	6 years	4.25 years	6 years
No Social Security	1.93%	4.26%	1.38%	2.95%
SS retirement only	1.80%	4.05%	1.28%	2.80%
Simple policy rule	1.30%	2.74%	0.91%	1.84%
50-50 hybrid policy	1.53%	3.34%	1.09%	2.27%

Panel B: Timing Risk and Disability Risk


standard deviation	Full Insurance (Δ)		Timing Premium (Δ_0)	
	4.25 years	6 years	4.25 years	6 years
SS retirement & DI	1.78%	3.94%	1.27%	2.72%

Figure 1. Simulated and Fitted Survival Probabilities



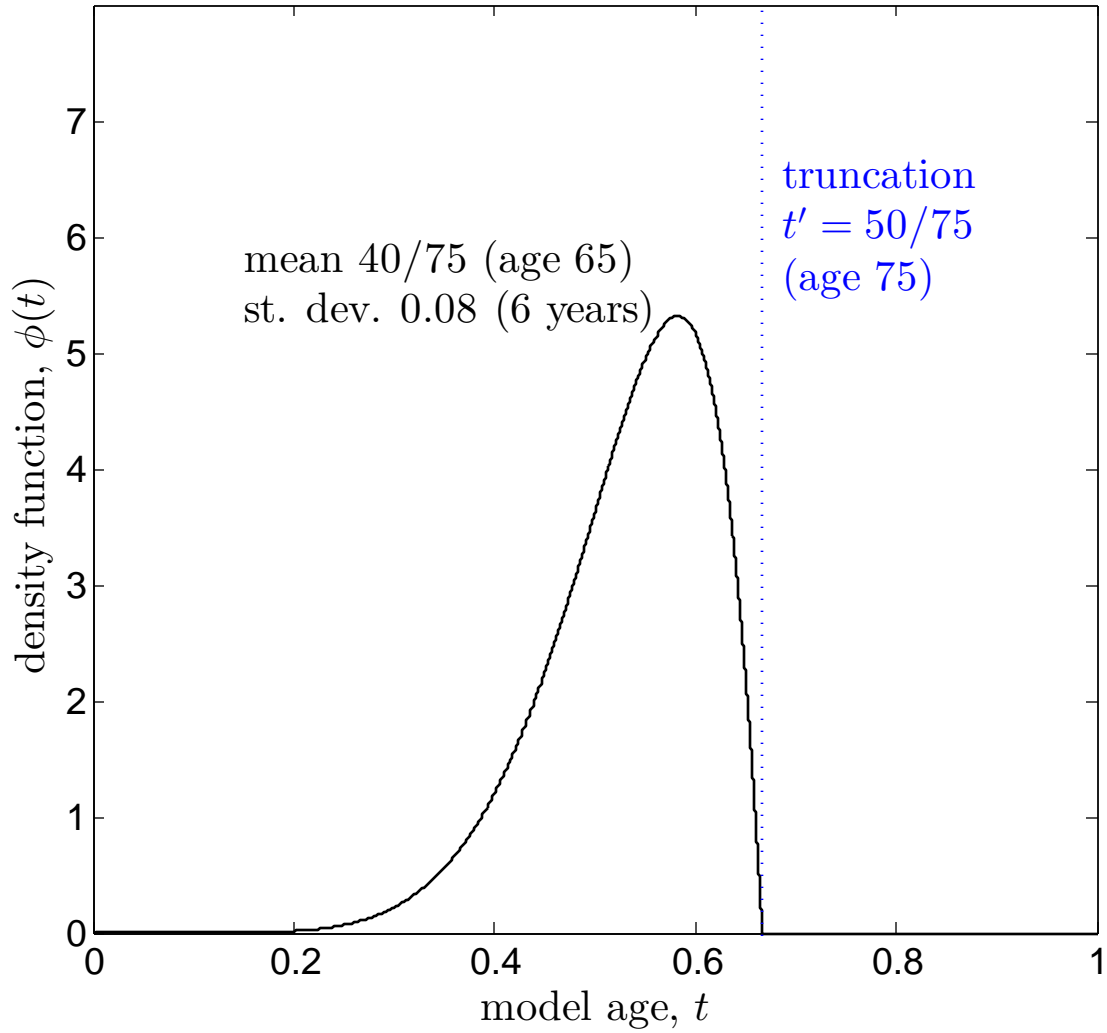
Fit $\Psi(t) = 1 - t^x$ to Social Security Administration cohort mortality tables.

Figure 2. Simulated and Fitted Wages, Age 16-75



Fifth-order polynomial fit to simulated male CPS data.

Figure 3. Calibrated p.d.f. over Retirement Timing Uncertainty



Truncated beta: $\phi(t) = t^{\gamma-1}(t' - t)^{\beta-1} \left\{ \int_0^{t'} t^{\gamma-1}(t' - t)^{\beta-1} dt \right\}^{-1}$.

Figure 4. Probability of Disability, Conditional on Retirement Age




Figure 5. Consumption over the life cycle with retirement timing uncertainty






Figure 6. U.S. Social Security vs. First-Best Insurance



$FB(t)$ and $SS(t|0)$ are lump-sum payments at the date of retirement, t .

Figure 7. Consumption over the life cycle with timing risk and disability risk



Dashed lines are c_2^* with $d = 0$, dotted lines are c_2^* with $d = 1$.