Adverse Selection in the Annuity Market and the Role for Social Security
Author(s): Roozbeh Hosseini
Source: *Journal of Political Economy*, Vol. 123, No. 4 (August 2015), pp. 941-984
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/10.1086/681593
Accessed: 30/08/2015 16:09

Adverse Selection in the Annuity Market and the Role for Social Security

Roozbeh Hosseini

Arizona State University

I study the role of social security in providing insurance when there is adverse selection in the annuity market. I calculate welfare gain from mandatory annuitization in the social security system relative to a laissez-faire benchmark, using a model in which individuals have private information about their mortality. I estimate large heterogeneity in mortality using the Health and Retirement Study. Despite that, I find small welfare gain from mandatory annuitization. Social security has a large effect on annuity prices because it crowds out demand by high-mortality individuals. Welfare gain would have been significantly larger in the absence of this effect.

I. Introduction

Mandatory annuitization is a key feature of the current US social security system. The value of this feature is derived from its ability to overcome potential inefficiencies due to adverse selection in the annuity market.[1]

I am grateful to Larry Jones and V. V. Chari for their gracious guidance throughout this project. I also thank Laurence Ales, Marco Bassetto, Neil Doherty, Mike Golosov, Amy Finkelstein, Narayana Kocherlakota, Pricila Maziero, Ellen McGrattan, Olivia Mitchell, Chris Phelan, Jim Poterba, José-Víctor Ríos-Rull, Richard Rogerson, Todd Schoellman, Gustavo Ventura, Pierre Yared, Steve Zeldes, and participants at several workshops, seminars, and conferences for helpful comments and discussion. The editor and two anonymous referees provided many helpful suggestions that substantially improved the paper. All remaining errors are mine. Data are provided as supplementary material online.

[1] Existence of adverse selection is well documented by Friedman and Warshawsky (1990), Finkelstein and Poterba (2002, 2004, 2006), and McCarthy and Mitchell (2003), among others.

Electronically published June 30, 2015
[*Journal of Political Economy,* 2015, vol. 123, no. 4]

The purpose of this paper is to quantify the value of mandatory annuitization in the current US social security system using a framework in which informational frictions in the annuity market are explicitly modeled.

To this end, I develop a dynamic life cycle model in which individuals have private information about their mortality. Uncertainty about the time of death generates demand for longevity insurance. In this environment individuals can purchase annuity contracts at linear prices. Contracts are nonexclusive, and insurers cannot observe individuals' trades. The lack of observability in my model implies that insurers cannot classify individuals by their risk types. As a result, the unit price of insurance coverage is identical for all agents. Individuals with higher mortality (who, on average, die earlier) demand little insurance (or nothing at all). This makes lower mortality types (types with higher risks of survival) more represented in the market. This, in turn, leads the equilibrium price of annuities to be higher than the overall actuarially fair value of their payment.

In this environment, I define and characterize a set of ex ante efficient (first-best) allocations. I show that these allocations are independent of individuals' mortality risk type and contingent only on survival, which is publicly observed. This feature implies that ex ante efficient allocations can be implemented by a system of mandatory annuitization in which every individual is taxed, lump-sum, before retirement and receives a benefit contingent on survival after retirement. The ex ante efficient allocation will be the benchmark for the best outcome that any social security system can achieve.

The environment I study has three important features. First, individuals know all relevant information about their mortality risk types at the beginning of life. This assumes away any possibility of insuring against the realization of risk type in the market and biases the results in favor of mandatory annuitization. Second, there is no heterogeneity other than mortality types. This implies that optimal policies are uniform across individuals. Finally, there are no distortionary effects of policy on labor supply and retirement decisions. Therefore, the focus will be only on inefficiencies caused by adverse selection and the beneficial role of mandatory annuitization.

A key object in the model is the distribution of mortality risk types. This distribution determines the extent of private information in the economy. Following the demography literature, heterogeneity in mortality risk is modeled as a frailty parameter that shifts the force of mortality (see, e.g., Vaupel, Manton, and Stallard 1979; Manton, Stallard, and Vaupel 1981; Butt and Haberman 2004). This parameter, once realized at birth, stays constant throughout one's lifetime. Individuals with a higher frailty parameter are more likely to die at any given age. I parameterize the initial distribution of mortality types and use data on subjective survival probabilities in the Health and Retirement Study (HRS) to estimate those parameters.[2]

The model is calibrated to match two key features in the data: (1) the average replacement ratio in the current US social security system (this determines the extent of annuitization through social security) and (2) the average fraction of retirement wealth that is annuitized outside social security (this determines the demand for annuities in the current US system).

The quantitative exercise of this paper consists of welfare comparisons between three economies: (1) an economy with no social security in which individuals share their longevity risks only through an annuity market, (2) the same economy with the addition of a social security system calibrated to the current US system, and (3) an economy in which ex ante efficient allocations are implemented. To highlight the importance of market response to policy, I report all welfare calculations under three assumptions: no annuity market, annuity markets with full information, and annuity markets with private information.

The three main findings of the paper are as follows: (1) The overall welfare gain from having mandatory annuitization through the current US social security system relative to a benchmark without social security is 0.07 percent of consumption. (2) Social security has a large effect on the annuity price. This effect comes as a result of crowding out of demand for annuities by low-survival individuals. This price effect has a negative welfare impact of 0.35 percent of consumption. In other words, in the absence of this price effect, the welfare gain from social security would have been as large as 0.42 percent. (3) In the absence of an annuity market, the welfare gain from annuitization through the current US social security system is 2.68 percent of consumption. This is a significant welfare gain and indicates the extent of uninsured survival risk in the absence of any insurance mechanism.

The upshot of these findings is that assessing how useful social security is in providing annuity insurance depends on assumptions about imperfect annuity insurance markets. If an annuity market is missing for exogenous reasons, then mandatory annuitization through social security can have large welfare gains. If annuity markets exist but are imperfect because of adverse selection, then mandatory annuitization can

[2] Hurd and McGarry (1995, 2002) and Smith, Taylor, and Sloan (2001) document that these probabilities are consistent with life tables and ex post mortality experience. They argue that they are good predictors of individuals' mortality.

be welfare improving. But they may also drive good risk types out of the annuity market and exacerbate adverse selection. This crowding-out effect can significantly reduce welfare gains from the policy. This is in line with findings by Golosov and Tsyvinski (2007) and Krueger and Perri (2011), who study endogenous insurance markets and responses to public provision of insurance.

Related literature.—The role of mandatory annuitization in the annuity market with adverse selection was first studied by Eckstein, Eichenbaum, and Peled (1985) and Eichenbaum and Peled (1987). The contribution of this paper is the quantitative assessment of the welfare gains due to mandatory annuitization in the current US system.

There is a broad literature on measuring the insurance value of annuitization for representative life cycle consumers (e.g., Kotlikoff and Spivak 1981; Mitchell et al. 1999; Brown 2001; Davidoff, Brown, and Diamond 2005). The exercise in this literature is to determine how much incremental, nonannuitized wealth would be equivalent to providing access to actuarially fair annuity markets.[3] A key feature of all these studies is the static comparison between full insurance and no insurance. In contrast, in the current paper I allow for annuitization through private annuity markets at retirement. This allows me to distinguish between risk sharing that is provided by the market and self-insurance and to study how risk sharing changes in response to changes in publicly provided insurance.

Einav, Finkelstein, and Shrimpf (2010) study the welfare cost of private information in the United Kingdom's mandatory annuity market. They report welfare gains from imposing further mandates that can implement the first-best. In contrast, in this paper the comparison is made between an economy with mandatory annuitization from social security and a laissez-faire economy. In both economies, participation in an annuity market is voluntary. However, outcome in a laissez-faire economy is inefficient because of adverse selection. The goal is to evaluate how successful the current social security system is in improving outcomes over laissez-faire.

The rest of the paper is structured as follows: Section II describes the environment, defines and characterizes efficient allocations, and defines the equilibrium. Section III describes the parametric specifications of the mortality model. Section IV describes the data and the calibration procedure. Section V reports results of welfare comparisons. Section VI explores sensitivity and robustness. Finally, Section VII presents conclusions.

[3] The study by Lockwood (2012) is an exception in that he considers the comparison between no annuity and an annuity available at actuarially unfair market rates.

II. Model

A. Environment

The economy starts at date 0 and ends at T $(1 \leq T < \infty)$. Individuals are born at the beginning of period 0 and face an uncertain life span. An individual who survives to age t faces the uncertainty of surviving to age $t+1$ or dying at the end of age t. Anyone who survives to age T will die at the end of that age. There is a set of individual *frailty* types, $\Theta = [\underline{\theta}, \bar{\theta}] \subseteq \mathbb{R}_+$. Frailty type $\theta \in \Theta$ determines the probability of survival to each age t. Individuals with lower θ have a higher probability of survival (and longer expected lifetimes). Individual type, θ, is private information.

Suppose there is a well-defined distribution $G_0 \in \Delta(\Theta)$ with full support. Let $P_t(\theta)$ be the probability of survival to age t at age 0. Therefore, the joint probability that an individual's type is in the set $Z \subseteq \Theta$ and survives to age t is $\mu_t(Z) = \int_{\theta \in Z} P_t(\theta)\, dG_0(\theta)$.

Individuals who die exit the economy. Therefore, in each age the distribution of types (conditional on survival) becomes more skewed toward the higher survival types (i.e., lower θ types). Let G_t be the distribution of types conditional on survival to age t; then the fraction of people with type in any set $Z \subseteq \Theta$ is

$$G_t(Z) = \frac{\displaystyle\int_{z \in Z} P_t(z)\, dG_0(z)}{\displaystyle\int_{\theta \in \Theta} P_t(\theta)\, dG_0(\theta)} \quad \forall Z \subseteq \Theta. \tag{1}$$

Individuals have time-separable utility over consumption, $u(\cdot)$, as long as they live. They also enjoy utility from leaving a bequest at the time of death, $v(\cdot)$. These functions are assumed to be twice continuously differentiable with u', $v' > 0$ and u'', $v'' < 0$ and satisfy the usual Inada conditions. Let $x_t(\theta) = P_{t+1}(\theta)/P_t(\theta)$ be the one-period survival rate for type θ (probability of surviving to age $t+1$ conditioned on being alive at t). Then type θ's utility out of a given sequence of consumption, c_t, and bequest, b_t, is

$$\sum_{t=0}^{T} P_t(\theta)\beta^t \{u(c_t) + [1 - x_{t+1}(\theta)]\beta v(b_t)\}, \quad 0 < \beta \leq 1.$$

Each individual is endowed with a unit of labor endowment that is inelastically supplied for age-dependent wage w_t in every period $t \leq J < T$. All individuals work until age J and then retire. There is also a saving technology with gross rate $R \geq 1/\beta$.

An *allocation* is a map from agents' type to a positive real line, that is,

$$c_t : \Theta \to \mathbb{R}_+, \quad 0 \le t \le T,$$
$$b_t : \Theta \to \mathbb{R}_+, \quad 0 \le t \le T.$$

Here, $c_t(\theta)$ is the consumption of all θ type individuals conditioned on their survival at age t (and, similarly, $b_t(\theta)$ is the bequest that type θ leaves if he dies at the end of age t).

An allocation is *feasible* if

$$\int \sum_{t=0}^{T} \frac{P_t(\theta)}{R^t} \left[c_t(\theta) + \frac{1 - x_{t+1}(\theta)}{R} b_t(\theta) \right] dG_0(\theta) = \int \sum_{t=0}^{J} \frac{P_t(\theta)}{R^t} w_t dG_0(\theta). \tag{2}$$

Individuals face two types of risks here. From the ex ante point of view (before birth), they face the risk of their type realization. Individuals whose type θ implies a higher survival probability need more resources to finance consumption through their lifetimes relative to those types who have lower survival. Also, upon realization of frailty types θ, individuals face the risk of outliving their assets.

B. *Ex Ante Efficient Allocation*

A benchmark for perfect risk sharing against both realization of types and time of death is the ex ante efficient (or first-best) allocation. That is the solution to the problem of a social planner who maximizes the expected discounted utility of individuals behind the veil of ignorance, that is, before agents are born:

$$\max_{c_t(\theta), b_t(\theta) \ge 0} \int \sum_{t=0}^{T} P_t(\theta) \beta^t \{ u(c_t(\theta)) + [1 - x_{t+1}(\theta)] \beta v(b_t(\theta)) \} dgG_0(\theta)$$

subject to (2).

It is straightforward to verify that the allocations that solve the problem above must satisfy

$$c_t(\theta) = c_t(\theta') = c_t \quad \text{for all } \theta,\ \theta' \in \Theta,\ \ \forall t,$$
$$b_t(\theta) = b_t(\theta') = b_t \quad \text{for all } \theta,\ \theta' \in \Theta,\ \ \forall t,$$

and

$$u'(c_t) = \beta R u'(c_{t+1}) = \beta R v'(b_t).$$

As is evident from the above equations, the allocations do not depend on individuals' type θ. The intuition for this result is the following. In this environment, individuals are heterogeneous ex ante (i.e., they differ in the risk of survival) but identical ex post. There is no difference among dead individuals. There is also no difference among people who survive. Therefore, there is no reason that the planner should discriminate between them ex post.

The fact that allocations are independent of heterogeneous risk type means that a "one-size-fits-all" identical allocation not only is ex ante efficient under full information but also is incentive compatible and hence implementable even if risk type θ is private information. This means that the efficient allocation can be implemented by a lump-sum tax and transfer. An example of implementation is discussed in Section V.B.2.

Two key assumptions drive this result. One is that the planner (as well as individuals) is an expected utility maximizer. Removing this assumption leads to efficient allocations that are type specific. The other assumption is that mortality risk is the only heterogeneity in this environment. If individuals are heterogeneous in other characteristics (such as ability or taste), then the efficient allocations are type specific, and therefore, incentive compatibility constraints are not trivially satisfied.

C. *Competitive Equilibrium with Asymmetric Information*

An alternative insurance arrangement is the competitive equilibrium. Here, risk sharing is not perfect because of informational frictions in the annuity market.

1. Survival-Contingent Contracts

Individuals can purchase annuity contracts during the last period of work (model age J). One unit of annuity contract pays one unit of consumption good contingent on survival for as long as the agent survives starting at age $J + 1$. Contracts are assumed to be nonexclusive and cannot be contingent on an agent's past trades or the volume of the transaction. Contracts are linear in the sense that to purchase a unit of annuity coverage, the individual pays qa.[4]

2. Consumer Problem

Let k_t be the amount of noncontingent saving by the individual and b_t be the bequest left if the individual dies at the end of age t. The optimization problem faced by this individual is

[4] In this model, individuals choose to purchase an annuity at only one age even when they are allowed to trade at other ages (see Pashchenko [2013] for the proof). However, freedom to choose the age of purchase gives rise to a multiplicity problem. To avoid this problem I restrict the trade to the time of retirement.

$$\max_{c_t, b_t, k_{t+1}, a \geq 0} \sum_{t=0}^{T} P_t(\theta)\beta^t\{u(c_t) + [1 - x_{t+1}(\theta)]\beta v(b_t)\}$$

subject to

$$c_t + k_{t+1} = Rk_t + (1 - \tau)w_t \quad \text{for } t < J, \tag{3}$$

$$c_J + k_{J+1} + qa = Rk_J + (1 - \tau)w_t, \tag{4}$$

$$c_t + k_{t+1} = Rk_t + a + z \quad \text{for } t > J, \tag{5}$$

$$b_t = Rk_{t+1}, \tag{6}$$

$$k_0 \text{ is given,} \tag{7}$$

in which a denotes annuity coverage purchased, τ is the social security tax rate, and z is the social security benefit. Individuals cannot borrow and cannot sell an annuity. Note also that $x_{T+1}(\theta) = 0$ for all θ. Given price q, the type θ individual's demand for an annuity is $a(\theta; q)$ and aggregate demand for an annuity is $y(q) = \int a(\theta; q)\, dG_J$.

3. Social Security

There is a fully funded social security system that taxes individuals at ages 0 to J at rate τ (since labor is inelastically supplied, this is in fact a lump-sum tax) and transfers constant social security benefit z to everyone at ages $t > J$ for as long as they are alive. Social security, therefore, is in fact a mandatory annuity. Benefits and taxes satisfy the following budget constraint:[5]

$$z \int \sum_{t=J+1}^{T} \frac{P_t(\theta)}{R^t}\, dG_0(\theta) = \tau \int \sum_{t=0}^{J} \frac{P_t(\theta)}{R^t}\, w_t\, dG_0(\theta). \tag{8}$$

4. Annuity Insurers

There are a large number of insurers who sell life annuity contracts to individuals of age J. Faced with the aggregate demand for an annuity $y(\cdot)$ and the anticipated distribution of payouts $F(\cdot; q)$, they choose annuity price q to maximize

$$\max_{q \geq 0} qy(q) - \int \sum_{t=J+1}^{T} \frac{y(q)}{R^{t-J}} \frac{P_t(\theta)}{P_J(\theta)}\, dF(\theta; q). \tag{9}$$

[5] In reality the social security system in the United States is a much more complicated arrangement and has many other features embedded in it (progressivity, survival benefits, etc.). It is also set up as a pay-as-you-go system and is not fully funded. I abstract from all these aspects and focus on only one feature of the system: mandatory annuitization.

The distribution $F(\theta; q)$ determines what fraction of each unit of total annuity obligations by the insurer is to be paid to type θ. In the equilibrium—which is defined below—$F(\theta; q)$ is required to be consistent with individuals' demand for an annuity. Annuity insurers engage in Bertrand competition, and therefore, they make nonpositive profits.

5. Competitive Equilibrium

The equilibrium notion is similar to that in Bisin and Gottardi (1999, 2003) and Dubey and Geanakoplos (2001).

Definition 1. A competitive equilibrium with asymmetric information is the sequence of consumers' allocations, $(c_t^*(\theta),\ b_t^*(\theta),\ a^*(\theta),\ k_{t+1}^*(\theta))_{\theta\in\Theta}$, annuity insurer decisions, annuity price (q^*), anticipated distribution of payouts by insurers, (F^*), and social security policy $(\tau,\ z)$ such that

1. $(c_t^*(\theta),\ a^*(\theta),\ k_{t+1}^*(\theta))_{\theta\in\Theta}$ solves the consumer's problem for all $\theta \in \Theta$ given annuity price q^*;
2. q^* is the lowest price such that

$$q^* = \int \sum_{t=J+1}^{T} \frac{1}{R^{t-J}} \frac{P_t(\theta)}{P_J(\theta)}\, dF(\theta; q^*)$$

   if $\int a(\theta; q^*)\, dG_J > 0$; otherwise

$$q^* = \sup_{\theta} \sum_{t=J+1}^{T} \frac{1}{R^{t-J}} \frac{P_t(\theta)}{P_J(\theta)};$$

3. allocations are feasible:

$$\begin{aligned}\int \sum_{t=0}^{T} \frac{P_t(\theta)}{R^t} \left[c_t^*(\theta) + \frac{1 - x_{t+1}(\theta)}{R} b_t^*(\theta) \right] dG_0(\theta) \\ = \int \sum_{t=0}^{J} \frac{P_t(\theta)}{R^t}\, w_t\, dG_0(\theta);\end{aligned} \tag{10}$$

4. F^* is consistent with consumers' choices; that is, for any price q, the fraction of total annuity coverage bought by individuals with type in $Z \subseteq \Theta$ is

$$F^*(Z; q) = \frac{\displaystyle\int_{\theta\in Z} a^*(\theta; q)\, dG_J(\theta)}{\displaystyle\int_{\theta\in\Theta} a^*(\theta; q)\, dG_J(\theta)}$$

   and with positive mass only on $\underline{\theta}$ if $a^*(\theta) = 0$ for all θ in which $G_J(\cdot)$ is defined in equation (1); and
5. social security budget balances (eq. [8]).

Using the zero-profit condition and consistency conditions (condition 4 in the equilibrium definition), we get the equation for equilibrium price:

$$q^* \int a(\theta; q^*) dG_J(\theta) = \int a(\theta; q^*) \sum_{t=J+1}^{T} \left[\frac{P_t(\theta)}{P_J(\theta)} \frac{1}{R^{t-J}} \right] dG_J(\theta). \tag{11}$$

In this environment, individuals with a higher probability of survival demand more annuity insurance at any price. Since they survive with higher probability, they are more likely to claim the insurance they have purchased. Any unit of coverage that is sold to these individuals is more risky from the point of view of insurers. On the other hand, individuals with a lower probability of survival are less risky for insurers since they are less likely to survive and claim insurance coverage. However, since they are less likely to survive, they purchase less insurance (relative to high survival types). As a result, the insurers are left with a pool of claims more likely to be materialized than the average probability of survival in the population. The risk in each insurer's pool is higher than what is implied by the average risk of survival. Therefore, the equilibrium price of an annuity is higher than the actuarially fair value of its payout. This is the essence of adverse selection in this environment.

Social security leads to more severe adverse selection and higher annuity prices in the equilibrium. An increase in the social security tax causes everyone to reduce demand for an annuity in the market. Such an increase has a larger effect on demand for an annuity by high mortality types. The reason is that an increasing tax (and benefits) of social security has two effects. On the one hand, it substitutes for annuities and therefore reduces demand in the market. This effect is the same for all types. On the other hand, it provides annuities at cheaper rates (than are available in the market). This generates an income effect that increases demand. But the magnitude of this income effect depends on the probability of survival. Therefore, this effect is larger for low mortality types. These individuals survive with a higher probability and are more likely to collect social security benefits. Therefore, the overall reduction in annuity demand is larger for high mortality types than for low mortality types. As a result, increasing social security increases risk in the private annuity pool in the market. This leads to higher equilibrium prices.[6]

III. Model of Mortality

In what follows, aging is modeled as a continuous-time process. Later, I derive discrete-time age-specific probabilities.

Individuals are indexed by their frailty types, $\theta \in \mathbb{R}_+$. Let $h_t(\theta)$ be the force of mortality of an individual at age t with a frailty of θ. Frailty can

[6] See the online supplemental appendix for formal arguments in a two-period model.

be modeled in many ways. I follow Vaupel et al. (1979) and Manton et al. (1981) and assume the following:

$$\frac{h_t(\theta)}{h_t(\theta')} = \frac{\theta}{\theta'} \tag{12}$$

or, alternatively,

$$h_t(\theta) = \theta h_t.$$

An individual with a frailty of 1 might be called the *standard* individual. Let h_t be the force of mortality for the standard individual (note that this is, in general, different from the average population force of mortality). The frailty index shifts the force of mortality. Furthermore, an individual's frailty does not depend on age. Therefore, $\theta > \theta'$ means that an individual with frailty θ has a higher likelihood of death at any age t than an individual with frailty θ', on the condition that they are both alive at age t. Let $H_t(\theta)$ be the cumulative mortality hazard. Then

$$H_t(\theta) = \int_0^t h_s(\theta)\,ds = \theta \int_0^t h_s\,ds = \theta H_t, \tag{13}$$

in which H_t is the cumulative mortality hazard for the standard individual. Finally, the probability that an individual of type θ survives to age t is

$$P_t(\theta) = \exp\left[-H_t(\theta)\right] = \exp(-\theta H_t). \tag{14}$$

Therefore, if an individual of type θ has a 50 percent chance of survival to age t, an individual of type 2θ has a 25 percent chance of survival to the same age.

Let $g_0(\theta)$ be the density of frailty at age $t = 0$. Also let $\bar{P}_t$ be the overall survival probability in the population. In other words, it is the fraction of all individuals (across all θ types) who survive to age t. Therefore, the relationship between $\bar{P}_t$ and $P_t(\theta)$ is the following:

$$\bar{P}_t = \int_0^\infty P_t(\theta) g_0(\theta)\,d\theta. \tag{15}$$

Note that individuals with higher values of frailty θ will have a higher probability of dying and are more likely to die early. This changes the distribution of frailty types who are alive at each age t. The conditional density of all types θ who survive to age t can be found by applying Bayes's rule:

$$g_t(\theta) = \frac{P_t(\theta) g_0(\theta)}{\int_0^\infty P_t(\theta) g_0(\theta)\,d\theta} = \frac{P_t(\theta) g_0(\theta)}{\bar{P}_t}. \tag{16}$$

As the population ages, the distribution of frailty types who survive tilts toward the lower values of θ. This implies that the overall average mortality hazard in the population does not correspond to an individual's mortality hazard. The relationship between the average population mortality hazard, $\bar{h}_t$, and the individual mortality hazard, $h_t(\theta)$, is represented by the following equation:

$$\bar{h}_t = \int_0^\infty \theta h_t g_t(\theta)\, d\theta = h_t \int_0^\infty \theta g_t(\theta)\, d\theta = h_t \mathrm{E}[\theta \mid t], \tag{17}$$

in which $\mathrm{E}[\theta|t]$ is the mean frailty among survivors to age t. Note that since individuals with higher frailty die earlier and the distribution of types becomes skewed toward lower values of θ as the population ages, mean frailty in the population decreases; that is, $\mathrm{E}[\theta|t]$ is a decreasing function of t. This implies that, overall, the population at each age t dies at a slower rate than individuals (unless g_0 is degenerate). Consequently, knowing the overall mortality rate, $\bar{h}_t$, which can be computed from life tables, is not enough to find individuals' mortality hazard rates. To uncover individuals' mortality hazard rates, further assumptions on the shape of the distribution g_0 are needed.

I assume that the initial distribution of individual frailty, θ, is lognormal:

$$\log(\theta) \sim \mathcal{N}(0,\ \sigma_\theta).$$

A zero mean is assumed without loss of generality. Using equation (17), one can always scale h_t up or down to be consistent with population mortality hazard rates in data for any nonzero mean. For any given σ_θ, let $g_0(\theta; \sigma_\theta)$ be the probability density function of lognormal distribution $\ln\mathcal{N}(0,\ \sigma_\theta)$. Then, equation (15) can be used to find to find the baseline cumulative mortality hazard, H_t, for any age t:

$$\bar{P}_t = \int_0^\infty \exp(-\theta H_t) g_0(\theta; \sigma_\theta)\, d\theta. \tag{18}$$

The values for $\bar{P}_t$ at each age can be calculated from cohort life tables. In the model I assume that a period is 5 years, that individuals enter the economy at the age of 30, and that everyone dies at or before age 110. Given the variance of the initial distribution of frailty at birth, σ_θ^2, equation (18), together with my assumption about frailty (eqq. [12] and [14]), uncovers individuals' survival probabilities at each age t. These survival probabilities are, by construction, consistent with life table data at every age. That means, for any variance of initial distribution, σ_θ^2, that overall population survival in the model is exactly equal to survival probabilities calculated from the life table. However, to estimate the variance of the initial distribution, more information is needed. I use data on subjective survival probabilities in the HRS to estimate σ_θ^2. The estimation procedure is described in the next section and the Appendix.

IV. Data and Calibration

In order to perform the quantitative exercise of the paper, three sets of parameters are needed: (1) the initial distribution of frailty and the time path of morality hazard, (2) preference parameters, and (3) policy parameters (social security taxes and transfers).

I first describe the data and the procedure used to estimate the initial distribution of frailty and the time path of morality hazard. I then describe the data that I use to calibrate preference parameters. Finally, I describe the calibration procedure for preference parameters and policy parameters.

A. Individual Survival Probabilities

In order to estimate the parameters of the initial distribution of frailty, I use individual subjective survival probabilities from the HRS. The HRS is a biennial panel survey of individuals born in the years 1931–41, along with their spouses. In 1992, when the first round of data were collected, the sample was representative of the community-based US population aged 51–61. The baseline sample contains 12,652 observations. The survey has been conducted every 2 years since. The HRS collects extensive information about health, cognition, economic status, work, and family relationships, as well as data on wealth and income. The particular observation on survival probabilities that I use comes from the following survey question: “Using any number from 0 to 10 where 0 equals *absolutely no chance* and 10 equals *absolutely certain*, what do you think are the chances you will live to be 75 and more?” Hurd and McGarry (1995, 2002) analyzed HRS data on subjective survival probabilities and found that responses aggregated quite closely to predictions of life tables and varied appropriately with known risk factors and determinants of mortality. Also, Smith et al. (2001) found that subjective survival probabilities are good predictors of actual survival and death.

Although the above-mentioned studies point to the potential usefulness of these responses as probabilities, there is a drawback. Gan, Hurd, and McFadden (2005) noticed the existence of focal points (0 or 1) in responses.[7] They propose a Bayesian updating procedure for recovering subjective survival probabilities.

[7] They report that 30 percent of responses in wave 1 and 19 percent of responses in wave 2 are 0s or 1s.

I follow Gan et al.'s (2005) approach and assume that individuals' true beliefs regarding their survival probability are unknown to the econometrician. However, the distribution of beliefs is known (which is taken as a Bayesian prior). The goal is to estimate the standard deviation of this distribution. I assume that subjective survival probabilities are reported with error. The difference between reported probabilities and true probabilities is modeled as a reporting error. The distribution of reporting errors is also parameterized and estimated. Starting from a prior distribution on survival probabilities and observing the report, I obtain a posterior distribution over types. This posterior distribution can be used to form a survival likelihood function. With ex post mortality and survival data, the likelihood function can be used to estimate the parameters of the model. The details of the estimation procedure are laid out in the Appendix.

This estimation procedure identifies the standard deviation of the initial distribution of frailty types. Once this is known, equation (18) can be used to back out the baseline cumulative mortality hazard, H_t (this is the mortality hazard of type $\theta = 1$). In equation (18), $\bar{H}_t = -\log(\bar{P}_t)$ and $\bar{P}_t$ is the average survival probability from Cohort Life Tables for the Social Security Area by Year of Birth and Sex for males of the 1930 birth cohort (table 7 in Bell and Miller [2005]).

Once the baseline cumulative mortality hazard, H_t, is known, equation (14) can be used to compute individuals' survival probabilities $P(\theta)$. Computed survival probabilities and their implied life expectancies are plotted in figure 1. The top panel shows the probability of survival to each age. The bottom panel shows the implied life expectancy at each age. Since frailty is not observable, interpreting the degree of heterogeneity from the variance of the initial distribution of frailty is not straightforward. However, heterogeneity in frailty implies heterogeneity in life expectancy at each age. The estimation implies a standard deviation of 6 years for life expectancy at age 30, which indicates a large degree of heterogeneity. As a comparison, note that the gap in life expectancy between males and females at age 30 for the birth cohort of 1930 is 5 years (Bell and Miller 2005, table 11). Another comparison can be the gap in life expectancy between college-educated and less than high school–educated males at age 25, which is about 6 years (Richards and Barry 1998). In what follows, I assume that these subjective probabilities are true probabilities and represent the true risk of survival for each frailty type.

B. *Annuitized Wealth at Retirement*

I use the data on individual and household wealth and income in wave 7 of the HRS (year 2004) to document the amount of annuitized wealth

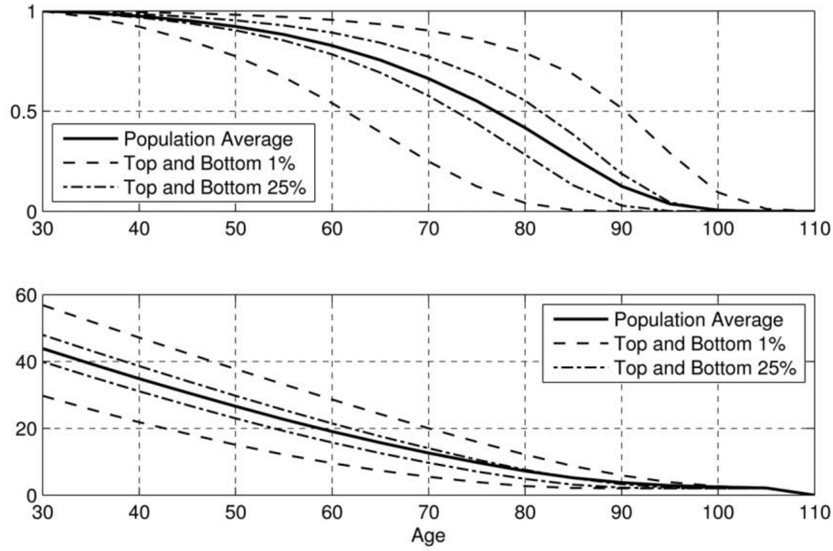

FIG. 1.—The top panel shows the probability of survival to each age. The bottom panel shows life expectancy at each age. Solid lines represent values for all the population. Dashed lines represent values for different quantiles of frailty, θ.

held by respondents.[8] I then calculate, on average, what fraction of individuals' *retirement wealth* (defined below) is in the form of private annuities and defined-benefit pensions. This is used as a calibration target in Section IV.C.

1. Annuitized Wealth in the HRS

I restrict my sample to male respondents who are older than 60 and collecting social security retirement benefits in wave 7. I use self-reported data on pension and annuity income as well as social security retirement benefits to compute the expected present discounted value of annuities/pension and social security benefits. To discount future payments, I use population survival probabilities for males of the 1930 birth cohort (Bell and Miller 2005, table 7) and a real interest rate of 3 percent. From now on, I refer to these calculated present values as *annuity and pension wealth* and *social security wealth*, respectively.

The HRS also collects data on household wealth, which includes financial assets, housing equity, and other assets. Financial assets include individual retirement account (IRA) balances; stock and mutual fund

[8] The reason for choosing wave 7 is that this is the period in which most of the respondents in wave 1 become early retirees and start collecting social security and pension/annuity income.

values; bond funds; checking, savings, money market, and certificates of deposit account balances; and trusts, less unsecured debt. Housing equity is the value of the home less mortgages and home loans. Other assets include the net value of other estates, vehicles, and businesses. I construct total retirement wealth for each individual as the sum of social security wealth, annuity and pension, and wealth data reported in the HRS.[9] I refer to this sum as *retirement wealth*. For each observation, I calculate what fraction of retirement wealth is annuitized through social security, private annuities, or employer-provided defined-benefit pensions. I report the average of this fraction over all individuals in table 1 (I calculate this fraction for each observation and report the average over all observations). The figures are similar for both married and single individuals.[10]

These numbers indicate that, on average, 10 percent of retirement wealth takes the form of private annuities or employer-provided defined-benefit pensions for males 60 years and older in wave 7 of the HRS. However, only a small fraction of annuitized wealth takes the form of private annuity contracts that people actively purchase using their accumulated retirement assets.

Table 2 shows the fraction of individuals in the sample who report positive annuity or defined-benefit pension income. As seen in columns 3 and 6 in table 2, there is a significant fraction of individuals in the sample (about half) who receive annuity income either through private annuities that they have purchased using their accumulated retirement assets or through employer-sponsored defined-benefit pension plans.

In summary, tables 1 and 2 show that individuals, on average, have a nontrivial amount of annuitized wealth, other than social security, at retirement. Also, a significant fraction receive annuity income from sources other than social security. However, table 2 also highlights a well-known fact that only a small fraction of people actively buy private annuity contracts. Most persons who receive annuity income other than social security get the annuity through employers' defined-benefit pension plans, which are in fact group annuity insurance arrangements purchased by employers on their behalf. Also, table 1 shows that most annuitized wealth other than social security is held in the form of defined-benefit pension entitlements.

2. Lump-Sum Withdrawals in Defined-Benefit Pension Plans

Workers who have a defined-benefit pension plan through their employers have limited control over the amount of benefits to which they are

[9] For HRS wealth data I use total wealth including secondary residence as reported in RAND HRS data, version M.

[10] See the online supplemental appendix for the complete table by age and marital status.

TABLE 1
AVERAGE FRACTION OF RETIREMENT WEALTH THAT IS ANNUITIZED

	Employer-Sponsored Pensions (%) (1)	Private Annuities (%) (2)	Social Security (%) (3)
All	9.5	.5	34.5
60–64	11.8	.3	41.2
65–69	11.0	.3	37.0
70–74	10.4	.6	37.5
75–79	8.8	.7	33.2
80+	6.0	.7	26.4

SOURCE.—Author's calculations using wave 7 of RAND HRS data, version M.
NOTE.—All calculations are for male respondents who are older than 60 and collecting social security retirement benefits in wave 7.

TABLE 2
FRACTION OF POPULATION WITH A PRIVATE ANNUITY OR A DEFINED-BENEFIT PENSION

	In Their Own Name (%)			In Own or Spouse's Name (%)		
	Employer-Sponsored Pension (1)	Private Annuity (2)	Pension or Annuity (3)	Employer-Sponsored Pension (4)	Private Annuity (5)	Pension or Annuity (6)
All	45.6	4.8	48.0	51.5	6.2	53.9
60–64	40.0	1.8	40.5	51.5	6.2	53.9
65–69	41.8	2.5	42.9	47.5	3.3	48.7
70–74	46.0	4.5	48.5	52.6	6.1	54.9
75–79	49.7	5.6	52.7	56.7	7.0	59.7
80+	48.9	8.8	52.9	54.4	10.8	58.7

SOURCE.—Author's calculations using wave 7 of RAND HRS data, version L.
NOTE.—All calculations are for male respondents who are older than 60 and collecting socials security retirement benefits in wave 7.

entitled.[11] However, a large fraction of these workers now have an option to claim benefits in the form of an annuity (default option) or lump-sum withdrawals (of the present discounted value of annuity benefits). There is overwhelming evidence from a variety of sources that a significant fraction of defined-benefit pension plans allow for lump-sum withdrawal of benefits, and their number is growing. According to the US Department of Labor (1995), the fraction of defined-benefit participants with access to any type of lump-sum option grew from 14 percent to 23 percent between 1991 and 1997. In 2005, 52 percent of all private industry workers with defined-benefit pension plans had lump-sum withdrawal options available to them at retirement (US Department of Labor 2007). Burman,

[11] The annuity benefit usually is based on an employee's average salary and length of service with the employer. With each year of service, a worker accrues a benefit equal to either a fixed dollar amount per month or year of service or a percentage of his or her final average pay.

Coe, and Gale (1999) report that on the basis of a 1993 employee benefit survey of the Current Population Survey, 58 percent of workers with defined-benefit pensions were eligible for a lump-sum distribution. Finally, Hurd and Panis (2006) study the 1992–2000 waves of the HRS and find that, on the basis of self-reported data, about 48 percent of full-time workers in the private sector had the option of lump-sum withdrawals upon job separation.[12] Perhaps more importantly, Hurd and Panis show that a large fraction of individuals choose to receive benefits as an annuity when they are presented with the lump-sum withdrawal option (24 percent expect to receive future annuity benefits, 56.2 percent draw current benefits, and about 15 percent cash out or roll their accrued benefits into an IRA). Benartzi, Previtero, and Thaler (2011) analyze more than 103,000 payout decisions from 112 different defined-benefit plans provided by a large plan administrator between 2002 and 2008. They find that 49 percent of participants who retire between ages 50 and 75 with at least 5 years of job tenure and account balances of $5,000 chose to collect benefits as annuities when they were given the option of lump-sum withdrawals.

This evidence suggests that a significant fraction of workers who received annuity income through defined-benefit pension plans are faced with the choice of collecting benefits as an annuity or a lump sum. There is little information on whether annuities offered in defined-benefit plans are more or less attractive than those offered on the market. However, the Internal Revenue Code regulates conversion between lifetime income benefits and lump sums in defined-benefit pension plans by prescribing mortality tables and discount rates to use in the calculations. Benartzi et al. (2011) show that the amount of (minimum) lump-sum withdrawals offered through these plans is comparable to the cost of the purchase of annuities with similar payments. This means that financial calculations made by these individuals are very similar to those of a purchase of an annuity contract. An important difference, however, is that in almost all plans, receiving benefits as an annuity is the default option. This may be an important factor in deciding whether to receive income as an annuity or a lump sum (however, as argued above, a nontrivial fraction of workers do choose to take lump-sum withdrawals). Also, although the number of plans that offer lump-sum withdrawal options is growing, there are still plans that offer no such options at retirement. (However, almost all plans offer lump-sum withdrawals when workers quit their jobs.)

[12] Not all of these job separations are due to retirement. A portion of them are due to job changes. However, depending on the size of benefits accrued and whether the worker is vested in the plan or not, she or he does have the option of making lump-sum withdrawals or leaving the accrued benefit as is and receiving annuity income at retirement.

The model described in Section II is very simple and stylized in many respects. The only source of annuity income other than social security in that model is the purchase of private annuity contracts. Motivated by the discussion above, I compare annuity insurance that is purchased in the model to the sum of private annuity and defined-benefit pension income in the HRS data. This is a simple way to account for all annuity income that individuals have without significantly complicating the model.[13]

C. *Benchmark Calibration*

The model period is 5 years. Individuals enter the model at age 30, and no one survives past 110. All individuals are endowed with the same hump-shaped earnings profile between ages 30 and 65, when they all retire. I use the US cross-sectional labor endowment efficiency profile estimated by Hansen (1993).

The utility function is constant relative risk aversion with coefficient of risk aversion γ over consumption and bequests:

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \quad \text{and} \quad v(b) = \xi \frac{b^{1-\gamma}}{1-\gamma}.$$

The term $\xi > 0$ is the weight on the bequest in the utility function and is identical for every individual. The higher ξ is, the higher the value of the bequest for individuals and the lower the demand for annuities.

For the curvature parameter of the utility function, γ, I use a benchmark value of $\gamma = 2$ and explore alternative values of $\gamma = 1$ and $\gamma = 4$ in Section VI.A. The annual discount factor is $\beta = 0.97$ and the annual real interest rate is 3 percent. Equation (8) is used to find social security taxes and benefits that match an average replacement rate of 45 percent.[14] Finally, to find the weight on the bequest, I solve the model using survival profiles calculated in Section IV.A and choose ξ such that, on average, the annuity wealth in the model is equal to 10 percent of total retirement wealth. To maintain consistency with calculations in the data, the present discounted values of annuity and social security income in the model are calculated using population survivals (rather than individual survivals).

Table 3 shows the calibrated parameters. All parameters other than γ and ξ will remain unchanged across various experiments and robustness

[13] I have also performed calibration and welfare calculations under three alternative assumptions. The results are reported in the online supplemental appendix.

[14] I assume that returns on social security are the same as market returns (3 percent annually) in the fully funded system I consider. This implies that a lower tax rate is required to balance the budget, relative to what is usually found under a pay-as-you-go system (in which returns are tied to demographic parameters).

TABLE 3
Benchmark Calibration

Parameter	Description	Value	Notes
T	Maximum age	17	Real life age of 110–114
J	Retirement age	7	Real life age of 60–65
w_t	Earnings profile		Taken from Hansen (1993)
β	Discount factor	.97	Annual
R	Real returns	1.03	Annual
γ	Risk aversion parameter	2	Benchmark value*
σ_θ	Standard deviation of log(9)	.56	Estimated using HRS response of subjective survival probabilities (see the Appendix)
$P_t(\theta)$	Individual survivals		Constructed to match life table for the 1930 birth cohort (eq. [18])
τ	Social security tax	.08	Finance social security replacement rate of .45
ξ	Weight on bequest	.9	Average annuitized wealth at retirement = 10%

Note.—The fraction with an annuity (not targeted) is 0.53.
* See Sec. VI.A for sensitivity.

exercises. Slightly more than half of the individuals hold an annuity at retirement in the model. This is in line with evidence reported in table 2. It shows that the model does a reasonable job of matching the actual rate of annuitization in the HRS data.

V. Findings

The main goal of the quantitative exercise is to find how much welfare is gained, ex ante, from implementation of mandatory annuitization in the current US social security system. I also report what would be welfare gains from implementing the ex ante efficient allocation described in Section II.B. The first calculation serves as a benchmark for how successful the current US social security system is in mitigating adverse selection. The second calculation is a benchmark for how much could possibly be gained.

To highlight the role of market structure in my calculations, I perform the quantitative exercise under three different assumptions as described below.

Economy with no survival-contingent assets.—This is an economy without any annuity contracts. Individuals can accumulate noncontingent assets only at rate R, which they leave as bequests if they die. I refer to this arrangement as *autarky*. If there is no social security, there will be no survival risk sharing in this economy. Note that in this economy the supply of annuity contracts is exogenously fixed at zero. So there will be no market response to a policy. Only individual consumption and saving allocations are different across scenarios with and without social security.

Economy with annuity markets and full information.—This is similar to the model economy described in Section II with the exception that the price of the annuity contract is assumed to be actuarially fair for each type. In other words, type θ pays the price $q^{AF}(\theta)$, which is the actuarially fair value of the unit annuity coverage that he purchases. I refer to this arrangement as an *annuity market with full information.* Note that in this economy the price of an annuity for each person is determined by his mortality risk type θ. Therefore, only the demand for annuity coverage (and not its price) will be affected by social security policy.

Economy with annuity markets and private information.—This is the full-blown model economy described in Section II. Individuals can purchase an annuity at the period before retirement. However, because of private information, there is adverse selection in the annuity market and the price is not actuarially fair. I refer to this arrangement as an *annuity market with private information.* In this economy the participation in the annuity market and the price of private annuity contracts depend on the level of social security tax and benefits.

I solve the model under all three assumptions above and calculate the ex ante welfare difference between the economy without social security and the economy with the current US replacement ratio. I calculate welfare as the percentage increase in lifetime income an individual requires, ex ante, without social security in order to be as well off as with social security.

A. *Welfare Gains from the Current US System*

Table 4 shows ex ante welfare gains from the current US replacement ratio relative to an economy without social security. The first row shows the welfare gain under the autarky assumption. As is expected in this case, the welfare gain is large since social security is the only source for annuity income at retirement.

The second row shows the welfare gain under the assumption of full information. As is expected, the welfare gain from social security is small in this case since every individual has access to an actuarially fair annuity.

TABLE 4
BENCHMARK WELFARE CALCULATIONS: EX ANTE WELFARE GAINS FROM ANNUITIZATION IN THE CURRENT US SYSTEM

	Welfare Gain (%)
Autarky	2.68
Annuity market with full information	.01
Annuity market with private information	.07
Annuity market with private information (with annuity prices fixed)	.42

NOTE.—Welfare gains are reported relative to an economy without social security.

The third row shows the main welfare calculation. That is the welfare gain from the current US replacement ratio under the assumption that mortality types are private information and there is adverse selection in the annuity market. Note that this number is much smaller than that in the first row. There are three reasons for this. First, in contrast to the autarky case, in the absence of social security, individuals can purchase an annuity in the private annuity market. Therefore, the outcome without social security is not as bad as in the autarky case. Second, as is the case with any forced insurance, individuals with low risk suffer losses. This will reduce the ex ante welfare gain from the policy. Third, social security crowds out good risk types (higher mortality types) in the annuity market. This leads to more severe adverse selection and increases in the market price of private annuities in the presence of social security. In fact, the market price of private annuities is 14 percent higher in the economy with the current US replacement ratio relative to an economy without social security.[15] To highlight the importance of this price effect, the fourth row reports welfare calculations under the private information assumption; but if the market price of annuities is held at that level, it would be in an economy without social security. If prices are held fixed, welfare gains from the current US replacement ratio would be higher by about 0.35 percent.

This finding highlights the key message of the paper. In the absence of social security, all individuals join the annuity market. In particular, the high-mortality individuals who are the good risk types will buy annuities. This leads to a better insurance pool, lower prices, and better risk sharing, ex ante. Public provision of annuities through social security crowds out private annuity markets. In particular, it runs good risk types (high mortality types) out of the markets. This in turn leads to high risks in the insurance pool and high prices. The calculation above shows that these price effects can have sizable welfare implications. Comparison of welfare gains in the first and third rows shows that ignoring endogenous responses of insurance markets to the policy can lead to very different conclusions about the usefulness of the policy. This is also in line with results by Golosov and Tsyvinski (2007) and Krueger and Perri (2011), who study endogenous insurance markets and the welfare implications of crowding out of private insurance by public insurance.

Figure 2 shows ex post welfare gains/losses by frailty type θ. The thick black line is the welfare gain from the current US replacement ratio under the assumption of annuity markets with private information. Welfare gains are positive and small for 80 percent of the population. The

[15] This finding is not new. Walliser (2000) also finds that in the absence of social security, annuity prices would be lower, although he reports much smaller numbers (2–3 percent). Part of the reason is that the annuity contracts he considers are not life annuities. Instead they are one-period survival-contingent bonds.

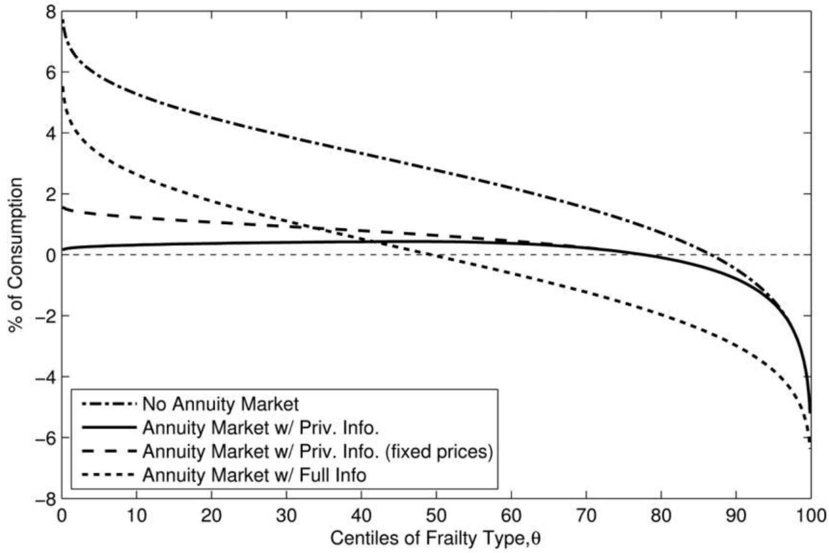

FIG. 2.—Ex post welfare gains/losses from the current US system relative to an economy with no social security for various assumptions regarding annuity markets.

20 percent with the highest mortality do not gain. In fact they suffer big welfare losses. These are individuals who have little value for annuitization, would rather not pay social security taxes, and instead increase their consumption at younger ages. The gray dashed line shows the same calculation but with the annuity price fixed at the level it would be when there is no social security. In this case welfare gains are higher for lower frailty types. This highlights the fact that the effect of social security on annuity prices is most harmful for the lowest mortality types. These are the risk types who value annuities most.

The dashed and dotted line and the dotted line show welfare gains under autarky and full information, respectively. Gains under autarky are significantly higher for lower mortality types and vanish as we move toward higher mortality types. Finally, gains and losses under the full-information assumption almost offset each other. But even under full information, some individuals make huge gains from social security and almost equal numbers suffer huge losses (these gains and losses are due to cross-subsidies across mortality types).

The results presented so far report only how useful current US social security is in mitigating adverse selection in the annuity market. The upshot from these findings is that social security has little welfare gain ex ante (and ex post), mainly because it lowers the welfare of high mortality types and worsens the adverse selection problem in the market for low mortality types. These results, however, do not mean that the welfare cost of private information in annuity markets is small. Nor do they necessarily mean that annuity markets provide large benefits (over the autarky alternative).

To assess the cost of private information in the annuity market, table 5 reports the ex ante welfare difference between an economy with full information and an economy with private information. This is the percentage increase in lifetime income that individuals require in the private information economy in order to be as well off as in the full information economy. The calculation is done both with and without the current US social security replacement ratio. The ex ante cost of private information is about 0.38 percent.

Table 6 shows the ex ante welfare benefit of having access to an annuity market with private information relative to autarky. This is the percentage increase in lifetime income individuals require in the autarky economy in order to be as well off as in the economy with annuity markets and private information. This is a measure of how useful the annuity market is in providing longevity risk sharing. As is expected, the annuity market is more valuable when there is no social security (and hence there is no alternative source for annuitized income). The value of the annuity market is 2.79 percent in the economy without social security and 0.17 percent in the economy with the current US social security replacement ratio. This shows that in this model, in the absence of social security, a substantial amount of survival risk can be shared through the private annuity market, even though the market suffers from inefficiencies due to adverse selection.

B. *Optimal Policy*

I discuss two benchmarks for optimal policy. In the first part I keep the policy instrument as before, that is, an age-independent tax rate during working ages and a constant social security benefit after retirement. I then find the best combination of tax and benefits that maximizes ex ante welfare. In the second part I remove restrictions on policy and discuss implementation of ex ante efficient allocations discussed in Section II.B.

TABLE 5
Cost of Private Information: Welfare Losses Due to Private Information

	Welfare Loss (%)
Without social security	.38
With social security	.32

Note.—This table reports the welfare difference between an economy with an annuity market and private information, and an economy with an annuity market and full information.

TABLE 6
VALUE OF AN ANNUITY MARKET: WELFARE GAINS FROM HAVING ACCESS TO AN ANNUITY MARKET

	Welfare Gain (%)
Without social security	2.79
With social security	.17

NOTE.—This table reports the welfare difference between an economy with an annuity market and private information, and an autarky economy.

1. Optimal Social Security Tax and Retirement Benefit

Table 7 shows the optimal replacement ratio and tax and ex ante welfare gains associated with them under the three assumptions regarding markets (autarky, full information, and private information). In all cases welfare gains are reported relative to an economy without social security. Recall that in the calibration the social security replacement ratio is 45 percent. Column 1 in table 7 shows optimal replacement ratios. The optimal replacement ratio under the autarky assumption is much higher than in the two other cases. In fact, when annuity markets are present (both with full information and with private information), optimal replacement ratios are smaller than 45 percent (the benchmark calibrated value for the current US system). The reason is that in this model there is a considerable demand for life insurance (survival benefits) at younger ages when individuals have few assets. A lower replacement ratio means lower taxes. Lower taxes allow individuals to accumulate assets faster, which they can leave as bequests if they die young. This is not a big concern in older ages since after retirement individuals can purchase an annuity (or receive social security), which insures one side of their mortality/survival risk.

2. Implementing Ex Ante Efficient Allocation

In this section I relax the restriction on policy instruments and describe a simple policy to implement ex ante efficient allocations in this model

TABLE 7
OPTIMAL SOCIAL SECURITY TAX AND RETIREMENT BENEFIT

	Optimal Replacement Ratio (1)	Optimal Tax (2)	Welfare Gain (%) (3)
Autarky	.58	.10	2.83
Annuity market with full information	.24	.04	.10
Annuity market with private information	.29	.05	.08

NOTE.—Welfare gains are reported relative to an economy without social security.

and report welfare gains from implementing these allocations. I refer to these allocations and the policy that implements them as *first-best*.

Recall from Section II.B that an ex ante efficient allocation is independent of mortality type. Everyone receives the same allocation independent of type. Also, the allocation must satisfy

$$u'(c_t) = \beta R u'(c_{t+1}) = \beta R v'(b_t),$$

in which b_t is the bequest left if the individual dies at the end of age t. I will maintain the assumption that $\beta R = 1$; hence allocations are constant over age. Let (c^*, b^*) be the ex ante efficient level of consumption (contingent on survival) and bequest (contingent on death). I propose a system of a social security tax rate, τ^*, a social security retirement benefit, z^*, and a sequence of survival benefits $(l_0^*, l_1^*, \ldots, l_T^*)$ that pays l_t^* if the individual dies at the end of age t. Consider the consumer problem of Section II.C.2 with the proposed social security policy

$$\max_{c_t, b_t, k_{t+1}, a \geq 0} \sum_{t=0}^{T} P_t(\theta)\beta^t \{u(c_t(\theta)) + [1 - x_{t+1}(\theta)]\beta v(Rk_{t+1}(\theta) + l_t^*)\}$$

subject to

$$\begin{aligned} c_t + k_{t+1} &= Rk_t + (1-\tau^*)w_t \quad \text{for } t < J, \\ c_J + k_{J+1} + qa &= Rk_J + (1-\tau^*)w_J, \\ c_t + k_{t+1} &= Rk_t + a + z^* \quad \text{for } t > J, \\ k_0 &= 0. \end{aligned}$$

PROPOSITION 1. Suppose that $\beta R = 1$ and let (c^*, b^*) be the ex ante efficient level of consumption and bequest. A social security policy (τ^*, z^*, l_t^*) such that

$$z^* = c^* + \left(\frac{1}{R} - 1\right) b^*,$$

$$l_t^* = \begin{cases} 0 & \text{for } t \geq J \\ (1-\tau^*)w_{t+1} - c^* + \left(1 - \dfrac{1}{R}\right) b^* + \dfrac{l_{t+1}^*}{R} & \text{for } t < J, \end{cases}$$

and

$$\int \left[\sum_{t=0}^{J} \frac{P_t(\theta)}{R^t} \frac{1 - x_{t+1}(\theta)}{R} l_t^* + \sum_{t=J+1}^{T} \frac{P_t(\theta)}{R^t} z^*\right] dG_0(\theta) = \tau^* \int \sum_{t=0}^{J} \frac{P_t(\theta)}{R^t} w_t \, dG_0(\theta)$$

implements (c^*, b^*).

Proof. The goal is to show that taking the policy (τ^*, z^*, l_t^*) as given, an individual will choose (c^*, b^*). I first show that for any type θ, if an individual purchases zero annuity, he will choose allocation (c^*, b^*). Then I show that given these choices, purchasing zero annuity is optimal. Consider the individual's first-order condition

$$P_t(\theta)u'(c_t) = P_{t+1}(\theta)\beta R u'(c_{t+1}) + P_t(\theta)[1 - x_{t+1}(\theta)]\beta R v'(Rk_{t+1}(\theta) + l_t^*).$$

Notice that if the individual chooses $c_t(\theta) = c^*$ and $k_{t+1}(\theta) = (b^* - l_t^*)/R$, the first-order condition is satisfied since a property of allocation (c^*, b^*) is that $u'(c^*) = v'(b^*)$. Also, it is straightforward to check that these choices satisfy consumers' budget constraints and the government's budget constraints by construction.

Now consider the following annuity price:

$$q = \sup_{\theta} \sum_{t=J+1}^{T} \frac{P_t(\theta)}{P_J(\theta)R^{t-J}}.$$

With $c_t(\theta) = c^*$, we have

$$qu'(c_J(\theta)) \geq \sum_{t=J+1}^{T} \frac{P_t(\theta)}{P_J(\theta)} \beta^{t-J} u'(c_t(\theta)),$$

and hence no one chooses to purchase an annuity. QED

Table 8 shows welfare gains from implementing the first-best relative to an economy with no social security. The calculation is done for all three assumptions regarding annuity markets (autarky, full information, and private information). Welfare gains are large in all cases. The first row shows the ex ante welfare gain from implementing the first-best relative to an autarky alternative. This is the ex ante welfare difference between a no-insurance arrangement and the best insurance arrangement. As is expected this number is large.

The second row is the ex ante welfare gain from implementing ex ante efficient allocation relative to a full-information alternative. Although

TABLE 8
Ex Ante Welfare Gains from Implementing the First-Best

	Welfare Gains (%)
Autarky	4.32
Annuity market with full information	1.11
Annuity market with private information	1.50

Note.—Welfare gains are reported relative to an economy without social security.

this number is much smaller than in the first row, it is still quite large. It appears that even the annuity market with full information is far from the first-best (but similarly far from autarky). This is, for the most part, the result of ignoring the life insurance market in the model.[16]

The third row shows the welfare gain from implementing the first-best in the economy with an annuity market and private information. These numbers are also much smaller than those in the first row, which is an indication that even an annuity market with private information goes a long way in providing survival risk sharing (this is also evident from table 6). However, the fact that these welfare gains are large is also an indication that there is substantial uninsured survival risk (and survival heterogeneity) in the environment. Comparing these numbers with the third row in table 4, we see that in the current US system the combination of social security and the annuity market is able to provide insurance that is worth a tiny fraction of the gain achievable under first-best (0.06 vs. 1.50).

The left panel in figure 3 shows the policies that implement the first-best. We can see that implementation requires large survival benefits at very young ages. This highlights, once again, that part of these large welfare gains are due to life insurance aspects of first-best policies (and not the annuity aspect of it).[17]

The right panel in figure 3 shows ex post welfare gains for different mortality types. About 95 percent of individuals gain from implementing ex ante efficient allocations (in a private information economy), and these gains are significant for a large fraction of mortality types.

VI. Robustness

I explore robustness of the results presented in Section V. First, I examine how choosing a different risk aversion parameter affects calibration and welfare calculations. Calibration of bequest parameters and welfare numbers is somewhat sensitive to the choice of risk aversion parameters. Second, I extend the model to include preference heterogeneity over bequests. The main findings of Section V are quite robust to this extension. Finally, I introduce heterogeneity in earnings profiles. Although income heterogeneity affects welfare numbers, it does not alter the main conclusions.[18]

[16] An active life insurance market together with the annuity market considered here can get very close to the ex ante efficient allocation under full information.

[17] Although social security provides survival benefits, these benefits are paid primarily to survivors of older/retired workers (unless the survivor cares for a young child). The survival benefits that come out of the first-best policy here are very different from those in the current US social security system. They are age dependent and are paid only to young workers.

[18] The online supplemental appendix contains more robustness exercises with respect to various parameters and assumptions.

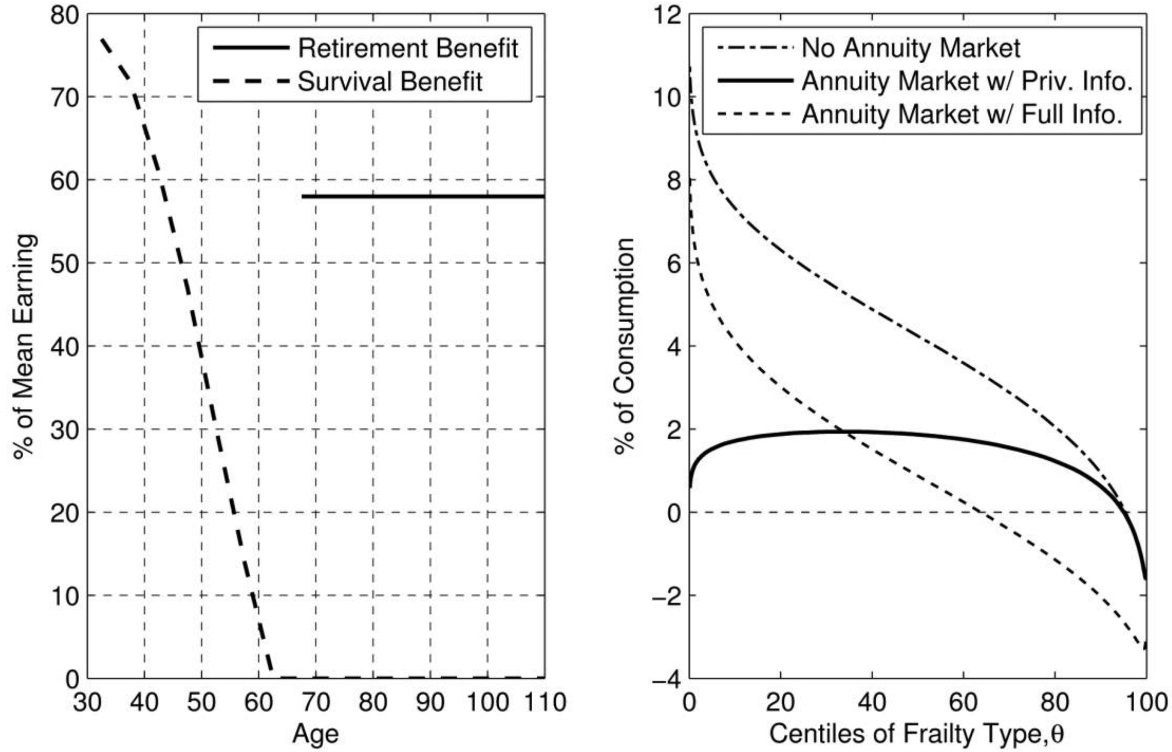

FIG. 3.—The left panel shows the first-best survival benefits and retirement benefits relative to mean earnings. The right panel shows ex post welfare gains/losses from implementing the first-best relative to an economy with no social security for three assumptions regarding annuity markets.

A. *Sensitivity to Risk Aversion Parameters*

Before I present calibration and welfare calculations under different values for risk aversion parameters, it is important to discuss how risk aversion affects equilibrium outcomes. The annuity price is higher in an economy with a lower risk aversion parameter.[19] The intuition for this result is as follows. In an economy with a lower value of a risk aversion parameter, individuals have less desire for a smooth consumption profile and are more sensitive to intertemporal prices. In this model there are always low-mortality individuals who find annuity prices lower than the actuarially fair price of their risk types. These individuals buy more annuity if risk aversion is lower. On the other hand, there are always individuals with high mortality who find annuities to be more expensive than the actuarially fair price of their risk types. These individuals buy less annuity if risk aversion is lower. This implies that at lower risk aversion the difference between the demand for annuity for high and low mortality is higher; that is, the profile of annuity purchase is steeper. This leads to a more severe adverse selection problem and higher annuity prices. Therefore, for lower values of risk aversion, the annuity price is higher and annuities are less attractive. On the other hand, for higher values of risk aversion, annuity prices are lower and annuities are more attractive.

The discussion above implies that to match the target of annuitized wealth at retirement, the calibrated value of ξ (the bequest parameter) must be lower (higher) for lower (higher) risk aversion. Otherwise, there will be too little or too much annuitization in the model. Table 9 shows the calibration for three different values of risk aversion parameters, $\gamma = 1$, 2, and 4. The bottom row shows the fraction of the population who buy an annuity in each case. For $\gamma = 1$, the annuity market is the least attractive of all. Therefore, only 41 percent purchase an annuity. This is still in line with data reported in table 2. For $\gamma = 4$, the annuity market is very attractive and 69 percent purchase an annuity. For this case the model predicts too many individuals who purchase an annuity.

Welfare calculations are presented in table 10. As is expected, welfare gains from social security are higher for lower values of risk aversion. For reasons discussed above, for $\gamma = 1$, the adverse selection is more severe, and therefore, individuals will gain more, ex ante, from public provision of annuities. Also, the calibrated value for ξ is lower, which means that individuals have more value for annuity income in general. For the same reasons, the cost of private information is also higher for $\gamma = 1$ and the benefit from access to annuity markets is lower.

On the other hand, welfare gains from the current US system are negative for $\gamma = 4$. Part of the reason is that annuity markets perform better if risk aversion is higher, and therefore, the adverse selection problem is

[19] The online supplemental appendix contains a formal proof for a two-period model.

TABLE 9
CALIBRATED BEQUEST PARAMETERS FOR VARIOUS RISK AVERSION VALUES

	Calibrated Weight on Bequest		
	$\gamma = 1$ (1)	$\gamma = 2$ (2)	$\gamma = 4$ (3)
ξ	.4	.9	6.0
Fraction with annuity (not targeted)	.41	.53	.69

NOTE.—For each value of the risk aversion parameter, ξ is chosen such that the average fraction of annuitized wealth at retirement is 10 percent at ages 65–70. All other parameters are the same as the benchmark (table 3). The benchmark calibration is $\gamma = 2$.

TABLE 10
WELFARE CALCULATIONS FOR VARIOUS RISK AVERSION VALUES

	$\gamma = 1$ (1)	$\gamma = 2$ (2)	$\gamma = 4$ (3)
	Ex Ante Welfare Gains from Annuitization in the Current US System (%)		
Autarky	2.91	2.68	−.54
Annuity market with full information	.51	.01	−2.82
Annuity market with private information	.69	.07	−2.87
Annuity market with private information (with annuity prices fixed)	1.09	.42	−2.70
	Welfare Losses Due to Private Information (%)		
Without social security	.58	.38	.08
With social security	.40	.32	.13
	Welfare Gains from Having Access to an Annuity Market (%)		
Without social security	2.3495	2.79	2.53
With social security	.15	.17	.13
	Ex Ante Welfare Gains from Implementing the First-Best (%)		
Autarky	3.75	4.32	9.05
Annuity market with full information	.78	1.11	6.26
Annuity market with private information	1.37	1.50	6.35

NOTE.—The benchmark calibration is $\gamma = 2$.

less severe. But negative welfare gains are mostly due to the fact that the bequest parameter required to match the calibration target in this case is very large ($\xi = 6.0$). Thus annuity income is valued less in this case. Note that even under full information the welfare gain is negative, which is another indication that annuity income is not valued. Therefore, ex post welfare losses for high mortality types are very large. At the same time, gains for low mortality types are small. These lead to overall welfare losses.

The gains from implementing the first-best, however, are large in this case. As discussed in the previous section, the first-best has an important life insurance component that is highly valued if risk aversion is $\gamma = 4$ and the bequest parameter is $\xi = 6.0$.

In all cases welfare gains from annuitization in the current US system are much smaller in the economy with annuity markets and private information relative to autarky. Also, the crowding-out effect of social security on prices is significant in all cases. For example, without the effect on prices, the welfare gain would have been more than 1 percent in the $\gamma = 1$ case.

B. *Heterogeneity in Preferences*

I extend the model to include heterogeneity in the bequest parameter, ξ, as well as heterogeneity in mortality. The model that I solve is very similar to the model in Einav et al. (2010). To capture possible correlation between the bequest parameter ξ and the mortality parameter θ, I assume that they are joint lognormal with correlation coefficient ρ:

$$\begin{bmatrix} \log(\theta) \\ \log(\xi) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ \mu_\xi \end{bmatrix}, \begin{bmatrix} \sigma_\theta^2 & \rho\sigma_\theta\sigma_\xi \\ \rho\sigma_\theta\sigma_\xi & \sigma_\xi^2 \end{bmatrix} \right).$$

For marginal distribution of θ, I assume the same parameters as estimated in the Appendix and reported in table 3. I use the estimation of Einav et al. for σ_ξ and ρ and calibrate the mean parameter μ_ξ to match the average fraction of wealth annuitized at retirement (similar to the calibration of the strategy in Sec. IV.C). The risk aversion parameter is equal to the benchmark value of $\gamma = 2$. Table 11 shows the calibrated

TABLE 11
CALIBRATION SUMMARY FOR PREFERENCE HETEROGENEITY

	Calibrated Average (Log of) Weight on Bequest			
	$\sigma_\xi = .099$, $\rho = .88$ (1)	$\sigma_\xi = .198$, $\rho = .88$ (2)	$\sigma_\xi = .099$, $\rho = .088$ (3)	$\sigma_\xi = 0$, $\rho = 0$ (4)
μ_ξ (mean of [log] bequest parameter)	−.07	−.03	−.1	−.1
Fraction with annuity (not targeted)	.52	.51	.53	.53

NOTE.—The risk aversion parameter in all cases is $\gamma = 2$. The weight is chosen to match average annuitized wealth of 10 percent at ages 65–70. In col. 1, $\sigma_\xi = .099$ is the estimation from Einav et al. (2010). The benchmark calibration is $\sigma_\xi = 0$.

mean parameter μ_ξ. All other parameters are fixed at the values reported in table 3.

Column 1 of table 11 uses the benchmark estimation in Einav et al. (2010) for σ_ξ and ρ. These are estimated as σ_β and ρ in their table 3. Column 2 uses the same value for ρ as column 1, but with a σ_ξ that is twice as large. Column 3 uses the same value for σ_ξ as column 1, but with ρ 10 times smaller. Finally, column 4 is the benchmark calibration discussed in Section IV.C. In all specifications the average annuitized wealth of 10 percent is matched. And the average values for the parameter ξ are close to 0.9 in all cases. Also, in all specifications, close to half of the population purchase annuities.

Table 12 presents a summary of all welfare calculations discussed in the previous section. All welfare numbers are very close to the numbers found in the benchmark model.

Note that with preference heterogeneity, the ex ante efficient (the first-best) allocations are no longer uniform across types. Therefore, they are not incentive compatible and cannot be implemented with type-independent policies. On the other hand, characterizing and implementing the ex ante efficient allocation in an economy with heterogeneity and multidimensional private information is a difficult problem and is outside the scope of this paper. Therefore, the ex ante efficient (or first-best)

TABLE 12
WELFARE CALCULATIONS WITH PREFERENCE HETEROGENEITY

	$\sigma_\xi = .099$, $\rho = .88$ (1)	$\sigma_\xi = .198$, $\rho = .88$ (2)	$\sigma_\xi = .099$, $\rho = .088$ (3)	$\sigma_\xi = 0$, $\rho = 0$ (4)
	Ex Ante Welfare Gains from Annuitization in the Current US System (%)			
Autarky	2.62	2.53	2.67	2.68
Annuity market with full information	−.03	−.06	.00	.00
Annuity market with private information	.06	.04	.06	.07
Annuity market with private information (with annuity prices fixed)	.40	.37	.41	.42
	Welfare Losses Due to Private Information (%)			
With social security	.39	.40	.38	.39
Without social security	.30	.29	.32	.32
	Welfare Gains from Having Access to an Annuity Market (%)			
With social security	2.73	2.67	2.78	2.76
Without social security	.17	.17	.17	.17

NOTE.—The risk aversion parameter in all cases is $\gamma = 2$. In col. 1, $\sigma_\xi = .099$ is the estimation from Einav et al. (2010). The benchmark calibration is $\sigma_\xi = 0$.

comparisons that are performed in previous sections are no longer appropriate here and are not reported in table 12.

Figure 4 shows the ex post welfare gains from the current US system for the private information economy (similar to what is presented in fig. 2) for the estimated parameters in Einav et al. (2010) (col. 1 in table 12). The thick solid lines represent all the combinations of type (θ, ξ) that experience the same welfare gain. The thin straight line shows the cutoff for annuity purchases. All types to the right of that line do not purchase an annuity under the current US system. Finally, the gray dashed rings demonstrate the area in the distribution of (θ, ξ) that contains 50 percent and 99 percent of the mass. It is plotted to indicate where the bulk of welfare gains/losses lie.

Individuals with high mortality (high θ) and a high value of bequest (high ξ) suffer the most from social security (up to 5 percent). Similarly to the case in figure 2, the largest gains are in the middle of the distribution. Once again the reason is that a significant part of the gain from annuitization in the current US system is lost as a result of the effect the system has on annuity prices (up to 0.42 percent). That is the explanation for low welfare gains for individuals with low mortality (low θ) and a low value of bequest (low ξ).

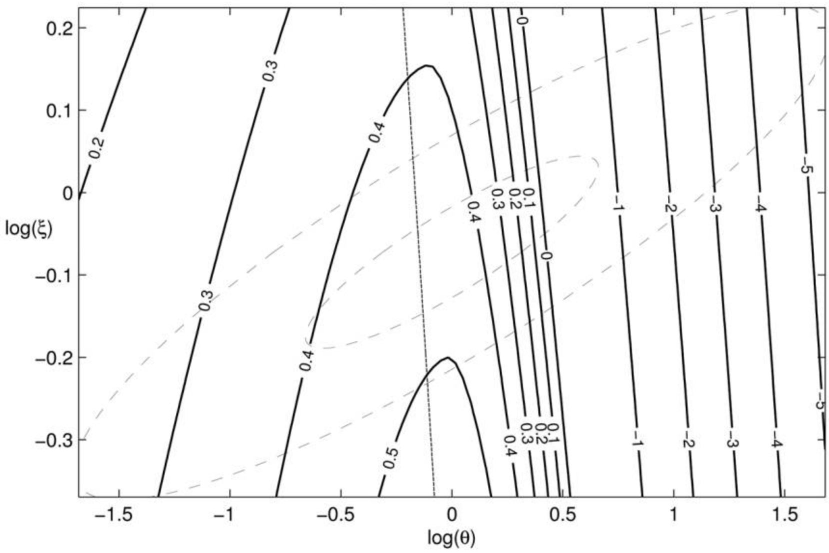

FIG. 4.—Ex post welfare gains/losses form the current US replacement ratio relative to an economy with no social security. The calculations are presented for $\sigma_\epsilon = 0.099$ and $\rho = 0.88$. The thin straight line shows the cutoff for annuity purchase. All types to the right of that line do not purchase an annuity under the current US system. Finally, the gray dashed rings demonstrate the area in the distribution of (θ, ξ) that contains 50 percent and 99 percent of the mass.

C. Heterogeneous Earnings Profiles

So far I have maintained the assumption that all individuals have the same age-varying profile of earnings. Although individuals will choose different levels of savings and therefore have different assets when they retire, there is no heterogeneity in their earnings ability in the model. This assumption was chosen to simplify the environment and focus on the annuitization and cross-insurance across mortality types rather than redistributional effects across income types. In this section I introduce income heterogeneity in a very simple and tractable way that also allows me to incorporate the correlation between mortality and income.

To maintain comparability with previous sections, I continue to assume the same distribution for mortality type θ. I also assume no heterogeneity in preferences (as in the benchmark calibration) and assume risk aversion of $\gamma = 2$. Heterogeneity in earnings is introduced by assuming that individuals have earnings $\epsilon \cdot w_t$, where w_t is the same common hump-shaped earning profile used in the previous sections of the paper (taken from Hansen [1993]) and ϵ is an age-independent scale factor. The scale factor ϵ has a lognormal distribution with mean zero and standard deviation σ_ϵ. I assume that θ and ϵ have the following joint distribution:

$$\begin{bmatrix} \log(\theta) \\ \log(\epsilon) \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma_\theta^2 & \rho_\epsilon \sigma_\theta \sigma_\epsilon \\ \rho_\epsilon \sigma_\theta \sigma_\epsilon & \sigma_\epsilon^2 \end{bmatrix} \right).$$

To calibrate parameters of income heterogeneity, I choose σ_ϵ to match a Gini coefficient of earnings equal to 0.5, which is the average value for earnings of prime-age workers (Budría et al. 2011). I then choose the correlation parameter ρ_ϵ to match mortality ratios by income quintile as reported in Cristia (2009).[20] These mortality ratios measure the likelihood of death for a person in a particular income group relative to the average population of the same age. The mortality ratios calculated by Cristia for men aged 50–64 are presented in table 13 (col. 2). For example, column 2 indicates that a male between ages 50 and 64 who is in the bottom income quintile is 1.63 times more likely to die than a random person of the same age in the population.

Given the standard deviation on earnings σ_ϵ and mortality σ_θ, I choose the correlation ρ_ϵ to minimize the distance between the mortality ratio in the model and data. The fit of the model is reported in table 13. Once the parameters of the joint distribution of (θ, ϵ) are known, I use the model to calibrate ξ by matching the average fraction of wealth that is annuitized at retirement. The target value is 10 percent as in previous

[20] He uses data from the 1984, 1993, 1996, and 2001 panels of the Survey of Income and Program Participation, matched to several files administered by the Social Security Administration. See Cristia (2009) for details on data and calculations of mortality ratios.

TABLE 13
Calibration of Income Mortality Correlation: Mortality Ratios for Men Aged 50–64

	Model (1)	Data (2)
Bottom income quintile	1.62	1.63
2nd income quintile	1.17	1.10
3rd income quintile	.95	.99
4th income quintile	.77	.68
Top income quintile	.55	.61

Note.—Data in col. 2 are taken from table 2 in Cristia (2009).

TABLE 14
Model Calibration with Income Heterogeneity: Parameters for Earnings Profiles and Weight on Bequest

Parameter	Value	Notes
σ_ϵ	.95	Earnings Gini of .45
ρ_ϵ	−.79	Match mortality ratio by income quintile*
ξ	.37	Average fraction of wealth annuitized = 10%

Note.—The risk aversion parameter in all cases is $\gamma = 2$. The fraction with an annuity (not targeted) is .53.
* Taken from table 2 in Cristia (2009).

sections. Table 14 shows the calibrated parameters as well as the fraction of the population who hold annuities. As was the case in previous sections, this fraction is in line with data in table 2.

Table 15 shows the results of the main welfare calculations, that is, ex ante welfare gains from the current US social security system.[21] Column 1 shows welfare numbers in the model with the heterogeneous earnings profile. Column 2 shows results in the benchmark calibration. Note that in the economy with annuity markets (both with full information and with private information), welfare gains from the current US system are negative. Low mortality types in this economy are more likely to be poor. Therefore, they suffer more from paying taxes while young, especially since the likelihood of receiving benefits is low (because of their high mortality). In this economy social security will redistribute from poor and high-mortality individuals to rich and low-mortality individuals.

These numbers indicate that the main finding of the paper is not sensitive to adding earning heterogeneity to the model. Still there is a large crowding-out effect from social security, which significantly reduces its potential gains. On top of that, when there is heterogeneity in earnings,

[21] Other calculations in this case are not very informative. For example, it is obvious that welfare gains from implementing the first-best are huge. But most of the gain comes from redistributing heterogeneous income. It is, however, interesting to investigate an environment when both income (or ability) and mortality are private information. This is left for future research.

TABLE 15
WELFARE CALCULATIONS WITH HETEROGENEOUS EARNINGS PROFILE: EX ANTE WELFARE GAINS FROM ANNUITIZATION IN THE CURRENT US SYSTEM

	WELFARE GAINS (%)	
	Model with Income Heterogeneity (1)	Benchmark (2)
Autarky	1.20	2.68
Annuity market with full information	−1.90	.00
Annuity market with private information	−.35	.07
Annuity market with private information (with annuity prices fixed)	−.22	.42

NOTE.—The risk aversion parameter in all cases is $\gamma = 2$.

social security has negative redistributive effects on poor individuals who are more likely to have high mortality.

VII. Conclusion

This paper investigates the welfare-improving role of social security as the provider of mandatory annuities in a model in which there is adverse selection in the annuity market. I found that social security can improve welfare ex ante. However, these ex ante welfare gains are small. There are two reasons for these findings. First, many high-mortality individuals prefer less annuitization than they receive from social security. These individuals incur a loss in the presence of social security. Second, gains to low-mortality individuals are also small because of high annuity prices in the presence of social security. Social security leads to a crowding effect that reduces the annuity demand by good risk types (high mortality types). This leads to more severe adverse selection in the private annuity market and high premiums. These higher premiums hurt low mortality types.

A key to this mechanism is the assumption that annuity contracts are nonexclusive and premiums are a linear function of coverage. Only under this assumption does social security have a crowding effect on annuity demand by high mortality types, which leads to higher annuity prices.[22]

This assumption is motivated by the available evidence that suggests that a model with nonexclusive contracts and linear pricing is a good approximation of how annuity markets work. For example, Cannon and Tonks (2008) study the voluntary annuity market in the United Kingdom and provide evidence on the actual pricing formula and premiums

[22] When contracts are exclusive and insurers use nonlinear pricing to screen annuitants, the main mechanisms described above are not operative. In such an environment, high mortality types are rationed and low mortality types pay a premium that more closely reflects their risk characteristics. Therefore, social security improves insurance coverage for everyone without any effect on the price of annuities. Therefore, it can have large welfare gains. See the online supplemental appendix for more discussion and the welfare calculations under alternative assumptions.

charged by annuity insurers for various levels of annuity coverage. They find that, except for very low coverage, the unit price of annuities does not vary with purchased coverage. Finkelstein and Poterba (2004) also estimate a hedonic pricing equation using the UK annuity market data and find no evidence of convexity in prices (which is an implication of the exclusivity in certain settings) or a major quantity discount.

Although the available evidence from annuity markets suggests that linear pricing is a good approximation, this need not be the case under an alternative social security policy. It is possible that in the absence of social security, annuity markets expand not only through lower unit prices of annuity coverage but also by introducing new sets of contracts. The response of the equilibrium set of contracts to policy and quantitative implications of that is another topic for future research.

The results also depend on the type of annuity contracts available. I allow individuals to purchase only single-premium life annuities. In reality, annuity insurers offer a variety of products and attempt to screen annuitants through specific features of contracts (other than price or coverage). For example, they attempt to lure higher mortality types to purchase a contract featuring a payment to survivors. I abstract from these features and focus on one single margin, namely, purchasing or not purchasing a given annuity contract. Future research should explore how the results are affected if individuals are allowed to choose from a variety of products.

Appendix

Estimation of Mortality Heterogeneity

In this appendix I describe the procedure for estimating the distribution of mortality heterogeneity using the HRS data on subjective survival probabilities. The main data that are used are the responses to the following question: “Using any number from 0 to 10 where 0 equals *absolutely no chance* and 10 equals *absolutely certain*, what do you think are the chances you will live to be 75 and more?” Here I follow Gan et al.’s (2005) procedure to estimate heterogeneity in mortality using data for male respondents who answered the question about survival probability in wave 1 (year 1992). I restrict the sample to male respondents between ages 50 and 72 who have answered the subjective survival probability question in wage 1. This leaves me with 5,083 individuals. I follow the status of each individual until wave 10 (year 2010) or until he exits the survey or dies.

A. *Notation and Assumptions on the Mortality Model*

Before I describe the estimation procedure, recall the notation and assumptions of the mortality model:

- θ: frailty index.
- $H_t(\theta)$: cumulative mortality hazard for an individual of type θ. I assume that

$$H_t(\theta) = \theta H_t,$$

in which H_t is independent of θ.
- $P_t(\theta)$: probability that an individual with type θ survives to age t from birth; note that by definition we have

$$P_t(\theta) = \exp\left[-H_t(\theta)\right] = \exp(-\theta H_t).$$

- $g_0(\theta;\sigma_\theta)$: density of the initial type distribution; $g_0(\theta;\sigma_\theta)$ is a lognormal distribution with mean of log equal to zero and standard deviation of log equal to σ_θ.
- $\bar{P}_t$: average (life table) probability of survival to age t from birth:

$$\bar{P}_t = \int_0^\infty \exp(-\theta H_t) g_0(\theta;\sigma_\theta)\, d\theta. \tag{A1}$$

Also let $\bar{H}_t = -\log(\bar{P}_t)$ denote the life table cumulative mortality hazard.
- $g_t(\theta;\sigma_\theta)$: density of types who survive to age t:

$$g_t(\theta;\sigma_\theta) = \frac{g_0(\theta;\sigma_\theta)\exp(-\theta H_t)}{\bar{P}_t}.$$

Note that the baseline mortality hazard, H_t, can be computed from equation (A1) if we know the standard deviation of the initial distribution.

Suppose that a respondent i has frailty type θ^i. Then the true probability that this individual survives to age 75 conditioned on being alive at age t is

$$\frac{P_{75}(\theta^i)}{P_t(\theta^i)} = \exp[-\theta^i(H_{75} - H_t)].$$

Let r_t^i be the report that this person makes about the probability of survival to age 75. Following Gan et al. (2005), suppose that the report is random and let $f(r_t^i \mid P_{75}(\theta^i)/P_t(\theta^i))$ be the density of the distribution of reports conditioned on the true probability of survival being $P_{75}(\theta^i)/P_t(\theta^i)$. I follow Gan et al. and assume that, conditional on $P_{75}(\theta^i)/P_t(\theta^i)$, r_t^i has a censored normal distribution; that is, there is a μ_t^i and σ_f such that

$$f\left(r \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right) = \phi\left(\frac{r-\mu_t^i}{\sigma_f}\right) \quad \text{for } 0 < r < 1,$$

and

$$\Pr\left[r = 0 \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right] = 1 - \Phi\left(\frac{\mu_t^i}{\sigma_f}\right),$$
$$\Pr\left[r = 1 \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right] = 1 - \Phi\left(\frac{1-\mu_t^i}{\sigma_f}\right),$$

in which $\phi(\cdot)$ is the standard normal probability density function and $\Phi(\cdot)$ is the standard normal cumulative distribution function. Furthermore, I assume that each individual makes no error on average:

$$\mathrm{E}\left[r \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right] = \frac{P_{75}(\theta^i)}{P_t(\theta^i)}.$$

Therefore, for each $P_{75}(\theta^i)/P_t(\theta^i)$, the following restriction must hold:

$$\frac{P_{75}(\theta^i)}{P_t(\theta^i)} = \Pr\left[r = 1 \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right] + \int_0^1 rf\left(r \mid \frac{P_{75}(\theta^i)}{P_t(\theta^i)}\right) dr. \tag{A2}$$

Note that this implies that the (uncensored) normal distribution has the same variance for all θ types and all ages. However, the mean depends on type and also on the age at which the report is being made (hence μ_t^i is indexed by both i and t). Given the true probability of survival $P_{75}(\theta^i)/P_t(\theta^i)$ and the standard deviation σ_f, μ_t^i can be computed by solving the following equation:

$$\begin{aligned}\frac{P_{75}(\theta^i)}{P_t(\theta^i)} = &\left[\Phi\left(\frac{1-\mu}{\sigma_r}\right) + \Phi\left(\frac{\mu}{\sigma_r}\right) - 1\right]\left[\mu_{\theta a t} - \sigma_r \frac{\phi\left(\frac{1-\mu}{\sigma_r}\right) - \phi\left(\frac{\mu}{\sigma_r}\right)}{\Phi\left(\frac{1-\mu}{\sigma_r}\right) + \Phi\left(\frac{\mu}{\sigma_r}\right) - 1}\right] \\ &+ \left[1 - \Phi\left(\frac{1-\mu}{\sigma_r}\right)\right].\end{aligned} \tag{A3}$$

B. *Estimation Procedure*

The prior on the density of types alive at age t is given by $g_t(\theta; \sigma_\theta)$. Once a report r_t^i is observed, we can form a posterior about the respondent i's type, given his report. I denote this posterior density by $\hat{g}_t(\theta \mid r_t^i; \sigma_\theta)$ and

$$\hat{g}_t(\theta \mid r_t^i; \sigma_\theta) = \frac{g_t(\theta; \sigma_\theta) f\left(r_t^i \mid \frac{P_{75}(\theta)}{P_t(\theta)}\right)}{\int_0^\infty g_t(\eta; \sigma_\theta) f\left(r_t^i \mid \frac{P_{75}(\eta)}{P_t(\eta)}\right) d\eta}.$$

We can use this posterior to form expectations about respondent i's frailty, given the report r_t^i. Let

$$\hat{\theta}(r_t^i) = \int_0^\infty \theta \hat{g}_t(\theta \mid r_t^i; \sigma_\theta) d\theta$$

be the conditional expectation of frailty type. Using $\hat{\theta}(r_t^i)$ and H_t, we can estimate the true probability of survival to any age T for individual i (conditional on being alive at t):

$$\hat{P}^i(t, T) = \exp[-\hat{\theta}(r_t^i)(H_T - H_t)],$$

or

$$\hat{P}^i(t,\ T) = \exp\left\{-\frac{\hat{\theta}(r_t^i)}{\sigma_\theta^2}\left[\exp\left(\frac{\bar{H}_T}{\sigma_\theta^2}\right) - \exp\left(\frac{\bar{H}_t}{\sigma_\theta^2}\right)\right]\right\}. \tag{A4}$$

The HRS has observations on

1. self-reported probability of survival to age 75 (I use the reports in only the first wave of the interview, which was 1992); this is used to construct a posterior mean of frailty types;
2. age at the time of the report (t^i);
3. exit information and age at exit (T^i) (a) for a respondent who died whose age we know at the last interview before death and (b) for a respondent who exited the survey before dying whose age we have at the last interview.

Therefore, we have right-censored observations, and censoring points are different for each observation. The HRS survey is biannual. If a person has died, we know that the person has died sometime between the previous wave and the current wave. So the likelihood function is

$$\mathcal{L} = \prod_{i \text{ is dead}} [\hat{P}^i(t^i,\ T^i) - \hat{P}^i(t^i,\ T^i + 2)] \prod_{i \text{ is not dead}} \hat{P}^i(t^i,\ T^i).$$

Then the log likelihood is

$$\log \mathcal{L} = \sum_{i \text{ is dead}} \log[\hat{P}^i(t^i,\ T^i) - \hat{P}^i(t^i,\ T^i + 2)] + \sum_{i \text{ is not dead}} \log \hat{P}^i(t^i,\ T^i),$$

in which $\hat{P}^i(t^i,\ T^i)$ is defined by equation (A4).

The parameters of $g_0(\cdot)$ and $f(\cdot|\cdot)$ are estimated by maximizing the above log likelihood function. There are two parameters that we need to estimate. One is the standard deviation of the logarithm of initial type distribution, σ_θ. The other is the standard deviation of the (uncensored) normal distribution that is used to construct $f(\cdot|\cdot)$. Therefore, the likelihood function is a function of two variables, σ_θ and σ_f. To evaluate the likelihood function, we use the following procedure:

- Given a guess of σ_θ and σ_f and using $\bar{H}_t$ (computed from the life table), we can compute H_t.
- For each respondent i we should find posterior density $\hat{g}_t(\theta \mid r_t^i; \sigma_\theta)$. This is then used to evaluate $\int_0^\infty \theta \hat{g}_t(\theta \mid r_t^i; \sigma_\theta)\, d\theta$ numerically to find an estimate of frailty for respondent i ($\hat{\theta}(r_t^i)$). Therefore, we need to know the value of $\hat{g}_t(\theta \mid r_t^i; \sigma_\theta)$ on a finite number of points (on a grid of θ). For each of these points, we can find μ_t^i from equation (A3). Once that is known, $f(\cdot \mid P_{75}(\theta)/P_t(\theta))$ can be evaluated (for each point on the grid). We can now compute $\hat{\theta}(r_t^i)$.

TABLE A1
Estimation Results

σ_θ	σ_f
.5607	.3481
(.0485)	(.0449)

Note.—Standard errors are in parentheses.

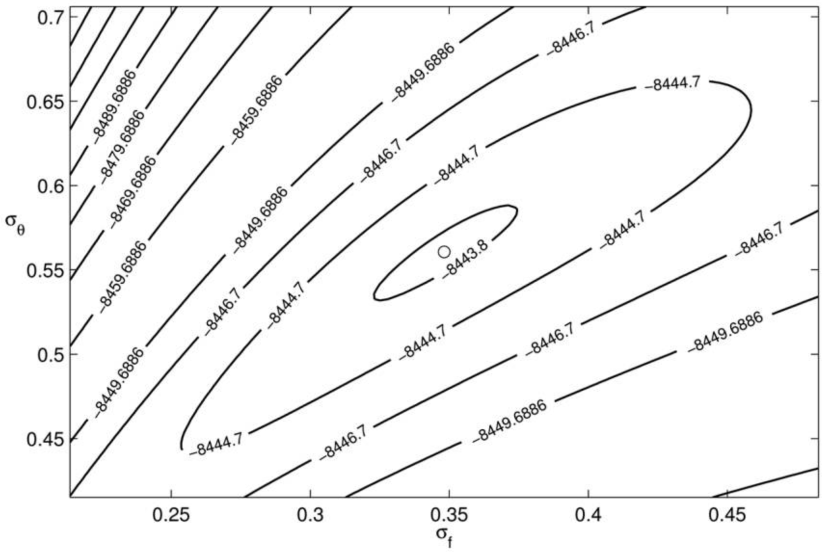

Fig. A1.—Likelihood contours. The estimated vector $(\sigma_f^*,\ \sigma_\theta^*) = (0.5607,\ 0.3481)$ is marked by an "o."

- Given $\hat{\theta}(r_t^i)$ and H_t, the probability of survival to any age $(\hat{P}^i(t,\ T))$ can be computed for respondent i.
- The likelihood of respondent i's survival to age T^i, conditioned on being alive at age t, is $\hat{P}^i(t,\ T)$.

The likelihood that respondent i dies between ages T^i and $T^i + 2$ is $\hat{P}^i(t^i,\ T^i) - \hat{P}^i(t^i,\ T^i + 2)$. Once the log likelihood is evaluated, we can use standard procedures for numerical optimization to find its maximum. Table A1 shows the results of the estimation. The numbers in parentheses are standard errors. Figure A1 shows contours for the likelihood function.

References

Bell, Felicitie C., and Michael L. Miller. 2005. "Life Tables for the United States Social Security Area 1900–2100." Actuarial Study 120, Social Security Admin., Washington, DC.

Benartzi, Shlomo, Alessandro Previtero, and Richard H. Thaler. 2011. "Annuitization Puzzles." *J. Econ. Perspectives* 25 (4): 143–64.
Bisin, Alberto, and Piero Gottardi. 1999. "Competitive Equilibria with Asymmetric Information." *J. Econ. Theory* 87 (1): 1–48.
———. 2003. "Competitive Markets for Non-exclusive Contracts with Adverse Selection: The Role of Entry Fees." *Rev. Econ. Dynamics* 6 (2): 313–38.
Brown, Jeffrey R. 2001. "Private Pensions, Mortality Risk, and the Decision to Annuitize." *J. Public Econ.* 82:29–62.
Budría, Santiago, Javier Díaz-Giménez, Vincenzo Quadrini, and José-Víctor Ríos-Rull. 2011. "Facts on the Distributions of Earnings, Income and Wealth in the United States: 2007 Update." *Fed. Reserve Bank Minneapolis Q. Rev.* 34 (1): 2–31.
Burman, Leonard E., Norma B. Coe, and William G. Gale. 1999. "Lump Sum Distributions from Pension Plans: Recent Evidence and Issues for Policy and Research." *Nat. Tax J.* 52 (3): 553–62.
Butt, Zoltan, and Steven Haberman. 2004. "Application of Frailty-Based Mortality Models Using Generalized Linear Models." *Astin Bull.* 34 (1): 175–97.
Cannon, Edmund, and Ian Tonks. 2008. *Annuity Markets.* Oxford: Oxford Univ. Press.
Cristia, Julian P. 2009. "Rising Mortality and Life Expectancy Differentials by Lifetime Earnings in the United States." *J. Health Econ.* 28 (5): 984–95.
Davidoff, Thomas, Jeffrey R. Brown, and Peter Diamond. 2005. "Annuities and Individual Welfare." *A.E.R.* 95:1573–90.
Dubey, P., and J. Geanakoplos. 2001. "Insurance Contracts Designed by Competitive Pooling." Cowles Foundation Discussion Paper no. 1315, Yale Univ.
Eckstein, Zvi, Martin Eichenbaum, and Dan Peled. 1985. "Uncertain Lifetimes and the Welfare Enhancing Properties of Annuity Markets and Social Security." *J. Public Econ.* 26 (3): 303–26.
Eichenbaum, Martin S., and Dan Peled. 1987. "Capital Accumulation and Annuities in an Adverse Selection Economy." *J.P.E.* 95 (2): 334–54.
Einav, Liran, Amy Finkelstein, and Paul Shrimpf. 2010. "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market." *Econometrica* 78 (3): 1031–92.
Finkelstein, Amy, and James Poterba. 2002. "Selection Effects in the United Kingdom Individual Annuities Market." *Econ. J.* 112 (476): 28–50.
———. 2004. "Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market." *J.P.E.* 112 (1): 183–208.
———. 2006. "Testing for Adverse Selection with 'Unused Observables.'" Working Paper no. 12112, NBER, Cambridge, MA.
Friedman, Benjamin, and Mark Warshawsky. 1990. "The Cost of Annuities: Implications for Saving Behavior and Bequests." *Q.J.E.* 105:135–54.
Gan, Li, Michael D. Hurd, and Daniel McFadden. 2005. "Individual Subjective Survival Curves." In *Analyses in the Economics of Aging*, edited by David A. Wise, 377–411. Chicago: Univ. Chicago Press (for NBER).
Golosov, Mikhail, and Aleh Tsyvinski. 2007. "Optimal Taxation with Endogenous Insurance Markets." *Q.J.E.* 122 (2): 487–534.
Hansen, Gary D. 1993. "The Cyclical and Secular Behaviour of the Labour Input: Comparing Efficiency Units and Hours Worked." *J. Appl. Econometrics* 8 (1): 71–80.
Hurd, Michael D., and Kathleen McGarry. 1995. "Evaluation of the Subjective Probabilities of Survival in the Health and Retirement Study." In "The Health and Retirement Study: Data Quality and Early Results," special issue, *J. Human Resources* 30:S268–S292.

———. 2002. "The Predictive Validity of Subjective Probabilities of Survival." *Econ. J.* 112 (482): 966–85.

Hurd, Michael D., and Constantijn Panis. 2006. "The Choice to Cash Out Pension Rights at Job Change or Retirement." *J. Public Econ.* 90 (12): 2213–27.

Kotlikoff, Laurence J., and Avia Spivak. 1981. "The Family as an Incomplete Annuities Market." *J.P.E.* 89 (2): 372–91.

Krueger, Dirk, and Fabrizio Perri. 2011. "Public versus Private Risk Sharing." *J. Econ. Theory* 146 (3): 920–56.

Lockwood, Lee M. 2012. "Bequest Motives and the Annuity Puzzle." *Rev. Econ. Dynamics* 15 (2): 226–43.

Manton, Kenneth G., Eric Stallard, and James W. Vaupel. 1981. "Methods for Comparing the Mortality Experience of Heterogeneous Populations." *Demography* 18 (3): 389–410.

McCarthy, David, and Olivia S. Mitchell. 2003. "International Adverse Selection in Life Insurance and Annuities." Working Paper no. 9975, NBER, Cambridge, MA.

Mitchell, Olivia S., James Poterba, Mark Warshawsky, and Jeffrey Brown. 1999. "New Evidence on the Money's Worth of Individual Annuities." *A.E.R.* 89 (5): 1299–1318.

Pashchenko, S. 2013. "Accounting for Non-annuitization." *J. Public Econ.* 98:53–67.

Richards, Hugh, and Ronald Barry. 1998. "US Life Tables for 1990 by Sex, Race, and Education." *J. Forensic Econ.* 11 (1): 9–26.

Smith, V. Kerry, Donald H. Taylor, and Frank A. Sloan. 2001. "Longevity Expectations and Death: Can People Predict Their Own Demise?" *A.E.R.* 91 (4): 1126–34.

US Department of Labor. 1995. *Retirement Benefits of American Workers: New Findings from the September 1994 Current Population Survey.* Washington, DC: Pension and Welfare Benefits Admin., Office Res. and Econ. Analysis.

———. 2007. "National Compensation Survey: Employee Benefits in Private Industry in the United States, 2005." Bulletin no. 2589, US Dept. Labor, Washington, DC.

Vaupel, James W., Kenneth G. Manton, and Eric Stallard. 1979. "The Impact of Heterogeneity in Individual Frailty on the Dynamics of Mortality." *Demography* 16 (3): 439–54.

Walliser, Jan. 2000. "Adverse Selection in the Annuities Market and the Impact of Privatizing Social Security." *Scandinavian J. Econ.* 102 (3): 373–93.