

SOCIAL SECURITY AND LONGEVITY RISK: THE CASE OF RISKY BEQUEST INCOME

ERIN N. COTTLE HUNT

Lafayette College

FRANK N. CALIENDO

Utah State University

This paper quantifies the welfare gains from Social Security when individuals face uninsurable longevity risk. While past researchers have studied this basic question, we do so from a unique perspective. In contrast to traditional macroeconomic models that abstract from specific linkages between parents and children, in our model, children are born to specific parents whose longevity is uncertain. And because parental asset holdings evolve over the life of the parent, children face uninsurable bequest income risk in addition to their own longevity risk. We find that Social Security improves ex ante expected utility by 3.4% of lifetime consumption (for the second generation). Because our baseline analysis assumes full information and optimal hedging of longevity risk, we treat these welfare gains as a conservative estimate, and we show that the gains are significantly larger when individuals fail to hedge their longevity risk.

Keywords: Social Security, Longevity Risk, Bequest Income Risk

1. INTRODUCTION

When individuals are young and accumulating assets for retirement, they do not know precisely how long they will live. This longevity risk presents a major challenge in determining how much to save because the amount needed to finance a comfortable living standard during retirement depends crucially on the total length of the retirement period. Living up to the age of 70 requires a far different level of savings than living up to the age of 100. Typical models in economic theory assume that individuals make consumption/saving plans that optimally hedge this risk, but even in that case such hedging involves large amounts of precautionary saving that go unspent when death occurs early. Absent a bequest motive, such unspent assets represent an inefficiency that opens the door for a second-best public program like Social Security.

Through collective risk sharing, Social Security helps individuals to hedge longevity risk. It does so by paying benefits to retirees as a life annuity that lasts

Address correspondence to: Frank N. Caliendo, Department of Economics and Finance, 3565 Old Main Hill, Utah State University, Logan, UT 84322-3565, USA. e-mail: frank.caliendo@usu.edu. Phone: 435 797 2963.

as long as they live, which is especially beneficial when longevity risk is uninsurable or expensive to insure in the private annuity market (due to adverse selection). This is the conventional logic in support of Social Security's longevity insurance role. The purpose of this paper is to test the conventional logic and quantify the individual welfare gains from Social Security when individuals face uninsurable longevity risk. We document large welfare gains from Social Security's provision of longevity insurance, on the order of 3.4% of lifetime consumption.

While a number of researchers have studied this issue and have performed similar calculations, this paper does so from the unique perspective of *sequential* welfare analysis of Social Security. We use a model in which bequest income is transmitted to the next generation through *explicit linkages* between parents and children. This contrasts with standard models that do not build explicit linkages between generations but instead assume that the aggregate assets of the deceased are either rebated to the surviving population in a deterministic way as in typical macroeconomic models [Feigenbaum (2008)] or simply do not factor into the model at all as in typical microeconomic models [Bütler (2001)].¹

We begin our sequential welfare analysis of Social Security with generation 1. They are the first generation to ever exist, and they participate in a Social Security retirement program. Social Security taxes are paid on wage earnings and benefits are collected as a life annuity during retirement. Social Security is efficiently financed, which allows us to abstract from the well-known effects of an inefficiently financed system. Individuals have full information about the longevity risk that they face and are fully rational. They have no private options for insuring their longevity risk, and they do not have a bequest motive so there is no downside to annuitization (for this generation). Using *ex ante* expected utility as the measure of welfare, the welfare gains to generation 1 from participating in Social Security are enormous and amount to 11.5% of lifetime consumption at our preferred level of risk aversion. This is consistent with conventional wisdom as well as with previous research measuring the welfare gains of Social Security with incomplete insurance markets [Hubbard and Judd (1987), and İmrohoroğlu et al. (1995)] and the vast literature that computes the welfare gains of annuitization in general [Kotlikoff and Spivak (1981), Mitchell et al. (1999), Brown (2001), Davidoff et al. (2005), and Lockwood (2012)].

Next, we move beyond conventional wisdom and into the heart of our analysis, generation 2. We maintain all of the assumptions concerning the economic setting, but with one twist: individuals in generation 2 are born to specific parents from generation 1 whose longevity is uncertain. And because asset holdings evolve over the life of generation 1, individuals in generation 2 face uninsurable risk over the magnitude of bequest income that they receive in addition to their own longevity risk. The effect of Social Security on welfare is now more interesting because Social Security carries both welfare benefits and welfare costs. On the benefit side, Social Security helps individuals in generation 2 to insure their own longevity risk, as was the case with generation 1. However, on the cost side, Social Security causes generation 1 to accumulate less private wealth for

retirement, and hence Social Security crowds out bequest income received by generation 2. Social Security's ability to provide net welfare gains through collective risk sharing requires that the benefits of longevity insurance dominate the cost of lost bequest income.

Ideally, to weigh these benefits and costs, one needs a model such as ours with explicit linkages between generations. We know from past studies that the manner in which bequest income is modeled has first-order effects on the welfare gains from longevity insurance through Social Security: the welfare gains can be very large in partial equilibrium models that ignore the impact of Social Security on the crowding out of bequest income to the next generation (as is the case in our generation 1 analysis) and can be *zero* in general equilibrium models that treat bequest income in an anonymous, deterministic fashion [Caliendo et al. (2014)]. The key to our model, however, is that the bequest income received by generation 2 is a *risky* income source. Therefore, a full understanding of Social Security's longevity insurance role must take into account that while Social Security crowds out bequest income, that income is risky. The riskiness of bequests reduces (but does not eliminate) the cost associated with Social Security.

The combined longevity risk and bequest income risk facing individuals in generation 2 is very costly. The Ex Ante Efficient consumption allocation—which is the solution to a Pareto planner problem that utilizes the Law of Large Numbers to fully insure the individual against both risks—is generally worth about 20% of the individual's total lifetime consumption. Social Security is nowhere near as effective as perfect insurance, but it does provide remarkably large welfare gains to generation 2—on the order of 3.4% of lifetime consumption at our preferred level of risk aversion.

We pay particular attention to the ability of the individual to hedge their own longevity risk. We attempt to place bookends on the welfare gains from Social Security by considering two examples. As a lower bound, we assume that the individual optimally hedges their longevity risk. Alternatively, as an upper bound, we consider an individual who does *not* hedge their longevity risk at all. Instead, the individual follows a consumption/saving plan that is based on deterministic survival to the age of life expectancy. That is, the individual correctly forecasts their own life expectancy, but they treat this as a certain event and follow a plan that is optimal based on this naïve assumption. The individual, therefore, makes no effort to hedge their longevity risk (they base their decision on the mean while ignoring the risk). Thus, they follow a plan that exhausts their assets at their life expectancy, and if they happen to live beyond their life expectancy (and therefore truly outlive their private assets), then they survive by consuming their Social Security benefits. Of course, this is perhaps only one of the multiple ways to study suboptimal hedging, but this description fits with the standard, casual motivation of why we need a Social Security program. We show that the welfare gains become significantly larger under this type of suboptimal hedging behavior, leading us to conclude that the welfare gains reported above should be treated as a conservative estimate of the usefulness of Social Security as a risk sharing device.

The issue of suboptimal hedging of longevity risk is important and has not received much attention even though it is key to understanding Social Security's potential role in sharing longevity risk. Recall that in a standard Yaari (1965) life-cycle model with rational behavior and full information about the distribution of longevity risk, the individual optimally self insures their longevity risk and does not actually face the risk of outliving their assets. But Social Security's life annuity feature is often motivated on the grounds that it provides protection against this very risk [Brown (2000)], and our analysis is careful to account for this possibility.

We also consider the long-run welfare implications of Social Security by extending our analysis to a model with three generations. We show the magnitude of bequest income compounds from one generation to the next, as well as the riskiness of bequest income across generations due to expansion in the number of permutations of mortality combinations among one's ancestral line. These longer term intergenerational wealth dynamics have important and nuanced effects on Social Security's role as a risk sharing device. In this setting we find that the welfare gains of Social Security are still significant.

Our paper is an extension of Caliendo et al. (2014). Specifically, we extend the two-period model in their appendix—which provides qualitative results on the direction of the welfare effect of Social Security when bequest income is risky—to a continuous-time dynamic setting in order to quantify the welfare gains from Social Security when bequest income is risky.

2. SETTING

In our initial setup, we focus exclusively on Social Security's longevity insurance role. All individuals across and within both generations share the same wage income, and they behave rationally in the face of longevity risk. Later in the paper, we systematically relax these assumptions.

Age is continuous and is indexed by t . The consumption spending of an individual at age t is $c(t)$ and period utility is denoted $u(c(t))$ with $u' > 0$ and $u'' < 0$. There are two generations, which we call generation 1 and generation 2. Individuals from both generations face longevity risk and make optimal consumption and saving decisions in the face of this risk. A Social Security program is introduced for the first time in generation 1. We do not model any generations before generation 1. For both generations, private annuity markets are closed, which means that—beyond self insurance—longevity risk is insurable only through collective risk sharing in Social Security.²

At age $t = 0$, the first generation enters the workforce. This generation is continuously divisible with unit mass at $t = 0$. The probability an individual from this generation survives to age t is $S(t)$, and death occurs no later than $t = T$, hence $S(0) = 1$ and $S(T) = 0$. They collect wage income w and pay Social Security taxes at rate τ until retirement occurs at $t = t_R$. After retirement, they collect a constant Social Security benefit b until death. Social Security benefits are paid as a life

annuity. We abstract from any type of political or fiscal risk and instead assume Social Security benefits are certain [Bütler (1999), Gomes et al. (2012), Nelson (2020), Luttmer and Samwick (2018), Kitao (2018), Caliendo et al. (2019) and Cottle Hunt (2020)]. Any income not consumed flows into a zero-interest storage technology $B(t)$, with the constraints $B(0) = B(T) = 0$. The Social Security system has a balanced budget and is financed efficiently, hence

$$b = \tau w \frac{\int_0^{t_R} S(t)dt}{\int_{t_R}^T S(t)dt}.$$

The second generation is exactly like the first generation in terms of their survival probabilities and wage endowment. They differ only in that they collect an inheritance from the first generation. To distinguish the asset holdings of the two different generations, any income not consumed by generation 2 flows into a zero-interest storage technology $k(t)$ (rather than $B(t)$), with the constraints $k(0) = k(T) = 0$.

In standard macroeconomic models, there is no explicit linkage between generations. The total assets of the deceased at a moment in time are anonymously rebated back to all of the living in equal, deterministic amounts. In contrast, in reality bequest income flows from parents to children, and the inheritance of the children is stochastic because the date of death of the parent is stochastic and parental asset holdings vary across time. We explicitly model the intergenerational linkage by assuming that each child from generation 2 is assigned to a specific parent from generation 1. The asset holdings of the parent upon death are transmitted one-for-one to the child. Because the parent’s asset holdings naturally vary over the life cycle and the parent faces longevity risk, the child faces uncertainty about bequest income in addition to their own longevity risk.

To simplify our analysis somewhat, we assume risky bequest income is transmitted to generation 2 at the beginning of their economic life span and is, therefore, resolved upfront. Of course, in reality, both the magnitude and timing of bequest income are uncertain. Abstracting from timing risk allows us to focus cleanly on a number of concepts in a tractable way. However, at the end of our paper, we solve and simulate the full blown model with timing risk, where both the timing and magnitude of bequest income are uncertain. This extension serves to enhance Social Security’s longevity insurance role beyond our baseline model that features early resolution of bequest uncertainty.

3. GENERATION 1

Generation 1 solves a standard Yaari problem of maximizing expected utility in the face of longevity risk

$$\max : \int_0^T S(t)u(c(t))dt,$$

subject to

$$\begin{aligned} \dot{B}(t) &= (1 - \tau)w - c(t), \text{ for } t \in [0, t_R], \\ \dot{B}(t) &= b - c(t), \text{ for } t \in [t_R, T], \\ B(0) &= B(T) = 0. \end{aligned}$$

For CRRA utility $u(c) = c^{1-\sigma} / (1 - \sigma)$, the solution is

$$c^*(t) = \frac{(1 - \tau)wt_R + b(T - t_R)}{\int_0^T S(t)^{1/\sigma} dt} S(t)^{1/\sigma}.$$

And the optimal asset holdings of generation 1 conditional on death at $t = x$ are

$$\begin{aligned} B^*(x) &= \int_0^x [(1 - \tau)w - c^*(t)]dt, \text{ for } x \in [0, t_R], \\ B^*(x) &= \int_0^{t_R} [(1 - \tau)w - c^*(t)]dt + \int_{t_R}^x [b - c^*(t)]dt, \text{ for } x \in [t_R, T]. \end{aligned}$$

The welfare gain from Social Security, Δ_{SS} , is the solution to

$$\int_0^T S(t)u(c_{SS}^*(t)(1 - \Delta_{SS}))dt = \int_0^T S(t)u(c_{LF}^*(t))dt,$$

where $c_{LF}^*(t)$ is Laissez-Faire optimal consumption with $\tau = b = 0$ and $c_{SS}^*(t)$ is optimal consumption with a Social Security system in place.

To conduct our numerical analysis, we assume an individual begins work at the age of 18 ($t = 0$), retires at the age of 67 ($t = t_R = 49$), and lives up to a maximum age of 100 ($t = T = 82$). We normalize the wage to unity ($w = 1$). We set $\tau = 0.106$, which corresponds to the employer and employee payroll tax for the Old Age Survivor Insurance portion of Social Security. This generates a replacement rate of 35% ($b = 0.35$). For our baseline analysis, we set the CRRA parameter $\sigma = 3$, which standard in the literature. We explore a range of values for this risk aversion parameter in our sensitivity analysis.

The survival function $S(t)$ is calibrated to ensure that the ratio of workers to retirees is 3.3, which is approximately the average value in the USA during the period 2000–2010.³ The survival function $S(t) = 1 - (t/T)^{2.58}$ produces a ratio of workers to retirees $\int_0^{t_R} S(t)dt / \int_{t_R}^T S(t)dt = 3.3$. The life expectancy at $t = 0$ implied by this survival function is age 77.1, which corresponds closely to the life expectancy from birth for the USA reported by the World Bank of 78.5 years. Our parameters are summarized in Table 1.⁴

The welfare gain from Social Security for generation 1 is $\Delta_{SS} = 11.5\%$. An individual in generation 1 would be willing to pay 11.5% of their lifetime consumption in order to have access to the annuitization provided through Social Security. This is consistent with conventional wisdom. The size of this welfare gain is independent of the CRRA parameter σ .⁵ The paths of consumption and saving over the life cycle are depicted in Figures 1 and 2. Consumption is higher

TABLE 1. Summary of Baseline Calibration of Parameters

Parameters	
Social Security:	
$\tau = 10.6\%$	Social Security payroll tax
$b = 0.35$	Balances the Social Security budget
Demographics and misc:	
$T = 82$	Normalized maximum life span (age 18–100)
$t_R = 49$	Length of career (age 18–67)
$S(t) = 1 - (t/T)^{2.58}$	Survival function yielding 3.3 workers:retirees
$w = 1$	Normalized wage
$\sigma = 3$	Midpoint CRRA value from literature

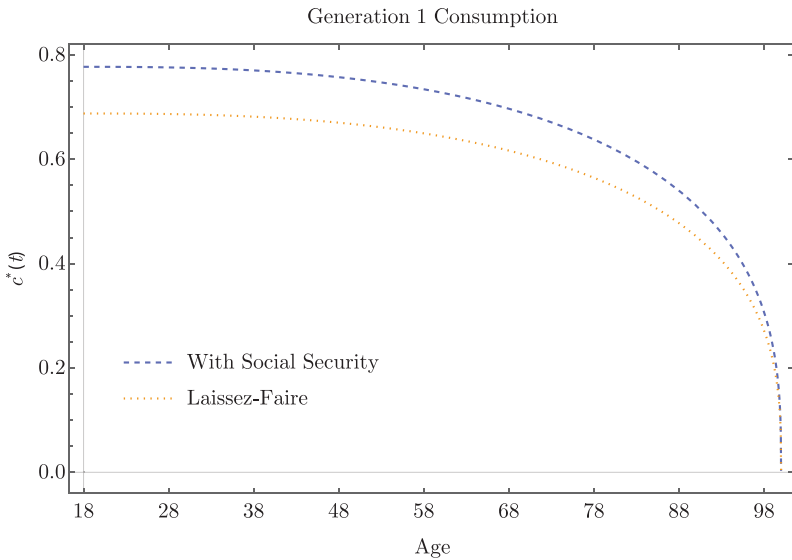


FIGURE 1. Generation 1 consumption $c^*(t)$ over the life cycle with Social Security (dashed blue line) and without (dotted yellow line) for $\sigma = 3$.

with Social Security than in the Laissez-Faire economy. Social Security crowds out saving, and so $B^*(t)$ is lower at every age with Social Security than without. This means the accidental bequests passed to generation 2 are smaller with Social Security than in the Laissez-Faire world.

4. GENERATION 2

Generation 2 faces a different optimization problem, because it receives an uncertain amount of bequest income. With probability $-\dot{S}(x)$ an individual in

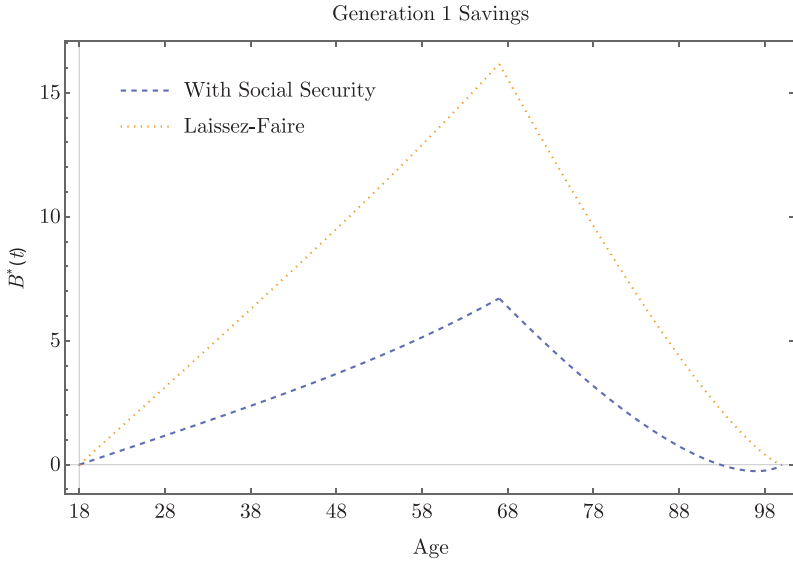


FIGURE 2. Generation 1 saving $B^*(t)$ over the life cycle with Social Security (dashed blue line) and without (dotted yellow line) for $\sigma = 3$.

generation 1 will die at age x holding assets $B^*(x)$. Hence, for the individual in generation 2, bequest income is a random variable with moments

$$\mathbb{E}(B^*) = \int_0^T (-\dot{S}(x)) B^*(x) dx,$$

$$\text{var}(B^*) = \int_0^T (-\dot{S}(x)) (B^*(x) - \mathbb{E}(B^*))^2 dx.$$

Throughout our paper, the key to our analysis is to carefully account for how Social Security affects the distribution of bequest income that passes from generation 1 to generation 2. Social Security crowds out the asset holdings of generation 1 which, in turn, reduces the bequests that generation 2 can expect to receive. But Social Security also reduces the variance of risky bequest income, and both effects have important welfare implications.

But to set the stage, we first derive the Ex Ante Efficient allocation for generation 2, which is the consumption path for an individual in generation 2 that solves a Pareto planner problem that perfectly hedges survival risk and bequest risk through the Law of Large Numbers.

4.1. Ex Ante Efficient Allocation

How big is the problem we are trying to solve? In other words, how costly is the individual’s exposure to the multiple risks they face? To answer this question,

we compute the Ex Ante Efficient allocation as the solution to a Pareto planner’s problem to create a benchmark for welfare comparisons.

We focus on the welfare of generation 2. The planner maximizes aggregate utility

$$\max \int_0^T S(t)u(c(t))dt,$$

subject to an aggregate resource constraint

$$\int_0^T S(t)c(t)dt = \int_0^{tR} S(t)w dt + \mathbb{E}(B^*),$$

where $\mathbb{E}(B^*) = \int_0^T (-\dot{S}(x)) B^*(x)dx$ is the expected bequest from generation 1 to generation 2.

The Lagrangian

$$\mathcal{L} = \int_0^T S(t)u(c(t))dt + \lambda \left(\int_0^{tR} S(t)w dt + \mathbb{E}(B^*) - \int_0^T S(t)c(t)dt \right),$$

has the first-order condition

$$S(t)u_c(c(t)) = S(t)\lambda \text{ for all } t,$$

and hence consumption is constant

$$c(t) = u_c^{-1}(\lambda).$$

We denote First Best consumption as c_{FB} , and using the aggregate resource constraint

$$c_{FB} = \frac{\int_0^{tR} S(t)w dt + \mathbb{E}(B^*)}{\int_0^T S(t)dt}.$$

This consumption path provides a perfect ex ante hedge against longevity and bequest risk. Expected utility is

$$\mathbb{E}(U) = \int_0^T S(t)u(c_{FB})dt.$$

4.2. The Second Best: Social Security with Optimal Hedging

We begin with the very best case for generation 2: the individual is fully rational and optimally hedges their longevity risk. The purpose of this section is to evaluate the welfare effects of Social Security and to compare it to the welfare gains associated with the Ex Ante Efficient allocation (or First Best allocation).

Generation 2 faces the following optimization problem, conditional on realization of x

$$\max \int_0^T S(t)u(c(t))dt,$$

subject to

$$\begin{aligned} \dot{k}(t) &= (1 - \tau)w - c(t), \text{ for } t \in [0, t_R], \\ \dot{k}(t) &= b - c(t), \text{ for } t \in [t_R, T], \\ k(0) &= B^*(x), \\ k(T) &= 0, \end{aligned}$$

where $B^*(x)$ is the bequest received from a specific individual from generation 1 who died at age x , which is the asset holdings of that individual at their age of death. The solution for CRRA is

$$c(t|x) = \frac{(1 - \tau)wt_R + b(T - t_R) + B^*(x)}{\int_0^T S(t)^{1/\sigma} dt} S(t)^{1/\sigma},$$

and ex ante expected utility (welfare) before the realization of $B^*(x)$ is

$$\mathbb{E}(U) = \int_0^T \int_0^T (-\dot{S}(x)) S(t)u(c(t|x))dt dx.$$

A **Sequential Equilibrium** with optimal hedging of longevity risk is comprised of generation 2 allocations $(c(t|x), k(t|x))_{t \in [0, T]}$ for each x that solve each individual’s optimization problem conditional on bequest income that is realized at time $t = 0$ and is distributed according to the PDF $(-\dot{S}(x))$ and optimal asset holdings of generation 1, $B^*(x)$, and Social Security tax rate τ and benefit b that jointly balance the Social Security budget.

The welfare gain from Social Security Δ_{SS} is

$$\begin{aligned} &\int_0^T \int_0^T (-\dot{S}(x)) S(t)u(c_{SS}(t|x)(1 - \Delta_{SS}))dt dx \\ &= \int_0^T \int_0^T (-\dot{S}(x)) S(t)u(c_{LF}(t|x))dt dx. \end{aligned}$$

And for comparison, the welfare gain associated with the Ex Ante Efficient allocation Δ_{FB} is

$$\int_0^T S(t)u(c_{FB}(1 - \Delta_{FB}))dt = \int_0^T \int_0^T (-\dot{S}(x)) S(t)u(c_{LF}(t|x))dt dx.$$

The welfare gains of Social Security and of the Ex Ante Efficient allocation for generation 2 are presented in Table 2 under the heading “Optimal Hedging.” With our baseline CRRA utility parameter $\sigma = 3$, the welfare gain associated with the First Best consumption is 20.2%. Ex ante, before bequest income is known, an individual facing longevity and bequest income risk would be willing to pay 20.2% of lifetime consumption for perfect insurance. In the absence of insurance, the only way for the individual to hedge their longevity risk is to engage in precautionary saving. This decreases their consumption over the life cycle and leaves resources on the table.

TABLE 2. Welfare gains of Social Security Δ_{SS} and of the Ex Ante Efficient allocation Δ_{FB} for generation 2 given optimal hedging, for different CRRA parameters σ . The right half of the table shows the expected value and the standard deviation of bequests inherited by generation 2 with Social Security (B_{SS}^*) and in Laissez-Faire (B_{LF}^*)

σ	Optimal Hedging		Bequests from Gen 1 to Gen 2			
	Δ_{SS}	Δ_{FB}	$\mathbb{E}(B_{SS}^*)$	$SD(B_{SS}^*)$	$\mathbb{E}(B_{LF}^*)$	$SD(B_{LF}^*)$
0.5	1.4%	9.9%	-3.66	1.49	1.98	2.09
1	2.1%	14.5%	-1.01	0.99	4.33	3.55
1.5	2.5%	16.8%	0.45	1.51	5.62	4.22
2	2.9%	18.3%	1.37	1.87	6.44	4.58
3	3.4%	20.2%	2.47	2.26	7.41	4.97
4	3.8%	21.4%	3.11	2.47	7.97	5.17
5	4.2%	22.3%	3.52	2.60	8.34	5.30

The large welfare gain associated with the First Best allocation suggests there is room for Social Security to improve lifetime wellbeing through a second-best allocation. We find the welfare gain of Social Security is 3.4%. Before their bequest is known, an individual would be willing to give up 3.4% of their lifetime consumption in order to live in an economy with Social Security. Clearly, Social Security is not as valuable as the First Best allocation, but it is an improvement over the Laissez-Faire economy. Social Security provides about one sixth of the welfare gain associated with the First Best allocation.

The welfare gains of Social Security are attributable to the annuity benefit payment which provides partial longevity insurance. Social Security crowds out private saving and thus reduces the bequests $B^*(x)$ that generation 2 receives. If bequest income were certain, this crowding out would harm the individual so much that it would completely unwind the gains from Social Security [Caliendo et al. (2014)]. However, in our model, bequest income is risky. So, the reduction in the size of the bequest is partially offset by the reduction in risk that the individual faces. This is evident by comparing the expected value and standard deviation of bequests received by generation 2 in the far right columns of Table 2. Expected bequests are larger and the variance is higher in the Laissez-Faire economy compared to the economy with Social Security. From the perspective of generation 2, Social Security crowds out bequest income, but it also reduces the variance of that income. For example, with the baseline value of $\sigma = 3$, Social Security reduces the size of the expected bequest $\mathbb{E}(B^*)$ from 7.41 to 2.47 compared to the Laissez-Faire economy. With wages normalized to $w = 1$, the reduction in expected bequests from 7.41 to 2.47 is equivalent to a loss of 4.94 years of wage earnings. Although Social Security reduces the size of bequests, it also reduces the riskiness of bequest income by reducing the variance. The standard deviation of bequest income is 4.97 without Social Security and falls to 2.26 with Social

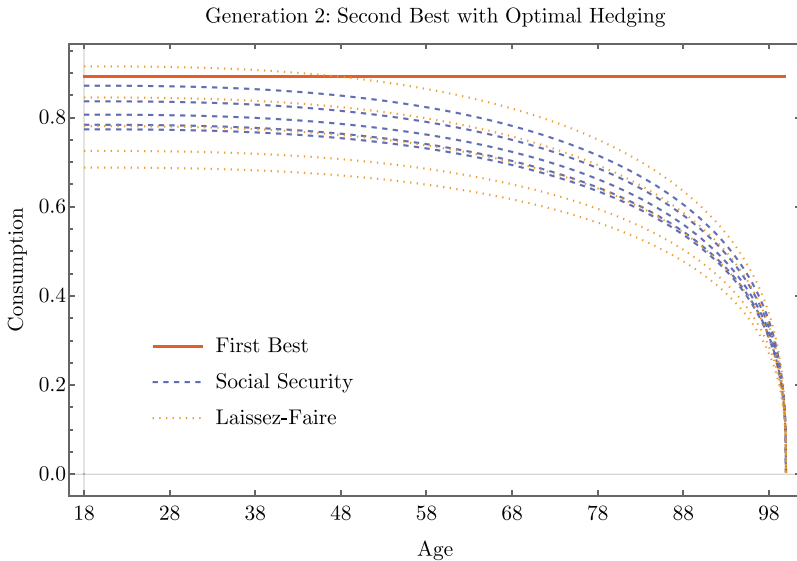


FIGURE 3. Generation 2 welfare effects. First Best consumption (solid red line) and consumption with Social Security (dashed blue lines) and without Social Security (dotted yellow lines) given different possible bequests $B^*(x)$ with $\sigma = 3$.

Security (when $\sigma = 3$). From an ex ante perspective, this reduction in risk partially offsets the reduction in expected bequest income. Social Security crowds out a risky income stream and replaces it with a life annuity.

The welfare gain of Social Security is also visible in the path of lifetime consumption. Figure 3 plots the path of consumption for generation 2 given different hypothetical realizations of bequests $B^*(x)$ along with the First Best consumption. While bequest income is a continuous random variable, the figure plots consumption for the maximum and minimum possible bequest as well as a few intermediate bequests for illustrative purposes. Ex ante, before bequest income is known, and individual faces the possibility of any of the consumption paths. First Best consumption is a perfectly smooth consumption profile, depicted in the graph as a solid red line. The Laissez-Faire consumption paths are yellow dotted lines. If an individual receives the maximum possible bequest in the Laissez-Faire world, their consumption starts out higher than the First Best amount, but falls over the life cycle. Lower bequests result in similar downward sloping consumption profiles that begin with lower consumption. Social Security crowds out bequest income, so consumption with the maximum possible bequest is lower. However, Social Security also reduces the riskiness of bequests, and so the possible consumption paths are closer together. They are shown as dashed blue lines in the graph. In expectation, an individual is better off with Social Security.

As another illustration, Figure 4 plots the path of consumption for an individual in generation 2 who receives the expected bequest $\mathbb{E}(B_{LF}^*)$ in the Laissez-Faire

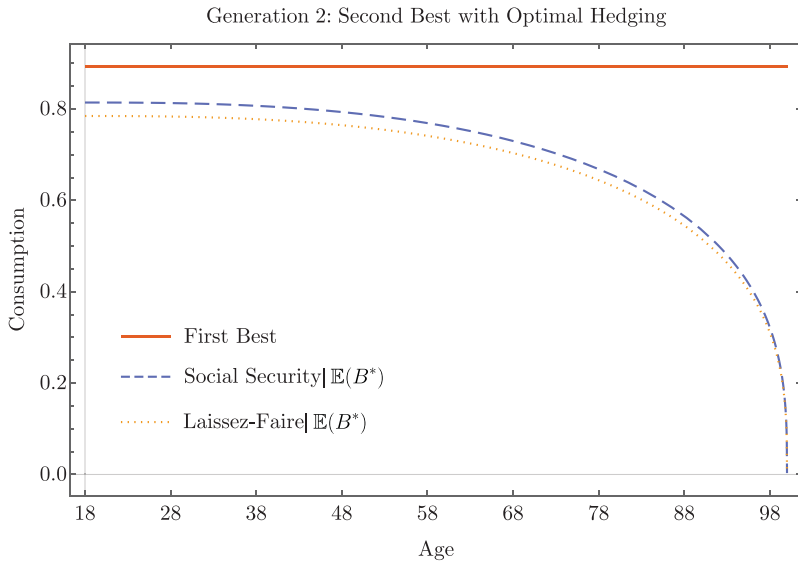


FIGURE 4. Generation 2 welfare effects. First Best consumption (solid red line) and consumption with Social Security (dashed blue line) and without Social Security (dotted yellow line) given expected bequests $\mathbb{E}(B_{SS}^*)$ and $\mathbb{E}(B_{LF}^*)$ with $\sigma = 3$.

economy (dotted yellow line) and the expected bequest $\mathbb{E}(B_{SS}^*)$ with Social Security (the dashed blue line). Consumption with Social Security and the corresponding expected bequest is everywhere higher than in the Laissez-Faire world. Social Security brings the individual closer to the First Best consumption because its positive longevity insurance effect overpowers the lost bequest income.

5. ROBUSTNESS

We explore the sensitivity of our results along several different margins in this section.

5.1. Wage Heterogeneity

In our analysis so far, we have intentionally abstracted from wage inequality and from the redistributive effects of Social Security’s benefit-earning formula. In this section, we set taxes and benefits according to their statutory values, and we consider the welfare gains from Social Security under a variety of assumptions about one’s own lifetime earnings as well as the lifetime earnings of one’s parent. Differences in earnings across generations interact with the dynamics of risky bequest income to create an important set of results about Social Security’s risk sharing role (explained below). In addition to longevity insurance, now Social Security has a redistributive role, as in Conesa and Krueger (1999), Huggett and

Ventura (1999), Cremer et al. (2008), Fehr and Habermann (2008), and Cremer and Pestieau (2011), to name a few.

Our baseline model assumes that all workers in both generations receive the same (average) wage. We relax that assumption in this section and consider wage heterogeneity. Social Security taxes wages up to a tax cap (currently \$137,700 in 2020), and pays benefits according to a progressive formula. The benefit rule is a piecewise linear function of earnings, with bend points that correspond to 0.2, 1.24, and 2.47 times the average wage, as in Alonso-Ortiz (2014). We consider a low-wage worker who earns just enough to receive benefits at the first bend point of the earning rule (earnings of \$11,300; benefit replacement rate of 90%). We also consider a high-wage worker who earns the maximum taxable wage (\$137,700), and thus receives the maximum Social Security benefit (replacement rate of 28%), as well as a worker with an average wage (\$55,700) who receives a replacement rate of 43%. We normalize the average wage to 1.

We begin with two scenarios: an average worker in generation 2 who had a low-wage parent in generation 1 and an average worker in generation 2 who had a high-wage parent in generation 1. We initially restrict our analysis to an average worker in generation 2 to avoid conflating welfare gains from redistribution provided by a progressive benefit-earning rule with the longevity insurance provided through the annuity benefit. Our results are presented in Table 3. The top row of the table shows results for an average worker in generation 2 who had a low-wage parent, the bottom row shows an average worker in generation 2 who had a high-wage parent, and the middle row features an average-wage worker in generation 2 with an average-wage parent.

The size of the bequest a worker in generation 2 receives depends positively on their parent's wage. If their parent had a low-wage, they will receive a smaller bequest, and if their parent had a high-wage, they receive a larger bequest (on average). And because bequest income is risky, the welfare gain associated with the Ex Ante Efficient allocation for generation 2 is increasing in parent wages (generation 1 wages). The ex ante welfare gains of Social Security for generation 2 is decreasing in the parent's wage. As the average size of bequests increases, the relative harm of crowding out bequest income increases. Thus, for an average-wage individual in generation 2 who had a high-wage parent, the welfare gain of longevity insurance provided through Social Security is 2.1% of consumption, compared to 12.2% for an average-wage individual in generation 2 who had a low-wage parent.

We continue our heterogeneity analysis in Table 4, by holding the wage of generation 1 constant (at the average), and varying the wage of generation 2. The top row of the table shows a low-wage generation 2 individual who had an average-wage parent, and the bottom row shows a high-wage individual in generation 2 who had an average-wage parent. The welfare gains of perfect insurance, Δ_{FB} are decreasing in generation 2's wage. This is because bequest income is a relatively smaller share of lifetime income for a high-wage individual in generation 2. The welfare effect of Social Security Δ_{SS} is non-monotonic in generation 2

TABLE 3. Welfare gains of Social Security Δ_{SS} and of the Ex Ante Efficient allocation Δ_{FB} for generation 2 given optimal hedging, for different parent wages (generation 1) and holding generation 2 wages fixed at the average wage. The right half of the table shows the expected value and the standard deviation of bequests inherited by generation 2 with Social Security (B_{SS}^*) and in Laissez-Faire (B_{LF}^*)

Parent wage	Generation 2		Bequests from Gen 1 to Gen 2			
	Δ_{SS}	Δ_{FB}	$\mathbb{E}(B_{SS}^*)$	$SD(B_{SS}^*)$	$\mathbb{E}(B_{LF}^*)$	$SD(B_{LF}^*)$
Low	12.2%	18.7%	-0.81	0.29	1.48	0.99
Average	6.4%	20.2%	1.46	1.75	7.41	4.97
High	2.1%	24.1%	8.10	6.62	18.30	12.29

TABLE 4. Welfare gains of Social Security Δ_{SS} and of the Ex Ante Efficient allocation Δ_{FB} for generation 2 given optimal hedging, for different generation 2 wages, holding parent wages (generation 1) constant at the average. The right half of the table shows the expected value and the standard deviation of bequests inherited by generation 2 with Social Security (B_{SS}^*) and in Laissez-Faire (B_{LF}^*)

Gen 2 wage	Generation 2		Bequests from Gen 1 to Gen 2			
	Δ_{SS}	Δ_{FB}	$\mathbb{E}(B_{SS}^*)$	$SD(B_{SS}^*)$	$\mathbb{E}(B_{LF}^*)$	$SD(B_{LF}^*)$
Low	5.3%	31.0%	1.46	1.75	7.41	4.97
Average	6.4%	20.2%	1.46	1.75	7.41	4.97
High	3.4%	19.0%	1.46	1.75	7.41	4.97

wages. This is due to two competing effects of Social Security. As generation 2 wages increase, Social Security becomes less valuable due to the progressivity of the benefit-earning rule. All else equal, this would lead the welfare gain of Social Security to be decreasing in generation 2 wages. However, Social Security also crowds out bequest income. Holding the wages of generation 1 constant at the average, reduced bequest income is worse for low-wage individuals in generation 2. All else equal, this would lead the welfare gain of Social Security to be increasing in generation 2 wages. Combining the two effects the welfare gain of Social Security is larger for an average individual in generation 2 than for either a low- or high-wage individual in generation 2. These competing effects can be seen more clearly in Figure 5 described below.

The welfare effects of Social Security for generation 2 are presented in Figure 5. This figure plots the generation 2 welfare effect of Social Security Δ_{SS} holding their wage constant at the average, varying the generation 1 wages as a dotted blue line. Consistent with the result in Table 3, the welfare gain of Social Security is decreasing in generation 1 wages. The average wage is normalized to 1 in the figure. The figure also plots the generation 2 welfare effect of Social

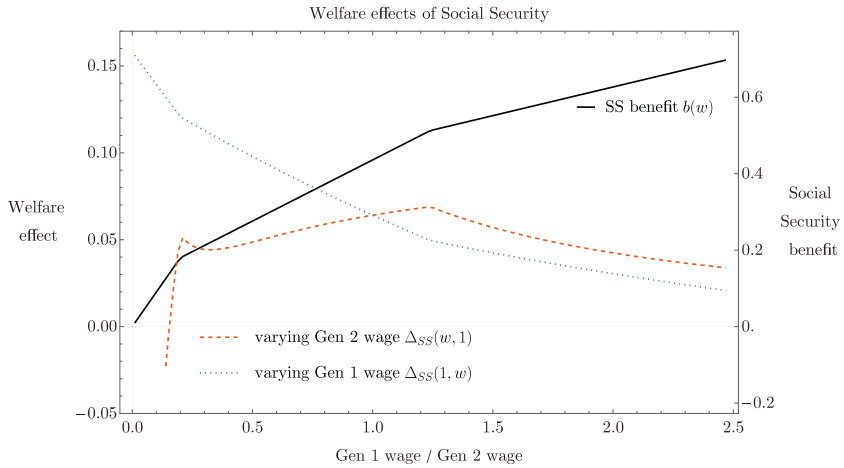


FIGURE 5. Generation 2 welfare effect of Social Security with wage heterogeneity. The notation $\Delta_{SS}(w_2, w_1)$ indicates generation 2 welfare as a function of generation 2 wages (w_2) and generation 1 wages (w_1). The dashed red line shows the welfare effect of Social Security $\Delta_{SS}(w, 1)$ for different generation 2 wages, holding generation 1 wages fixed at the average wage $w = 1$. The dotted blue line shows the welfare effect of Social Security $\Delta_{SS}(1, w)$ for generation 2, holding their wage constant at 1, varying the generation 1 wages. The solid black line plots the Social Security benefit-earning rule.

Security Δ_{SS} for different generation 2 wage levels, holding generation 1 wages constant at the average as a dashed red line. Consistent with Table 4, the welfare effect of Social Security is non-monotonic in generation 2’s wage. The welfare effect is highest for a generation 2 individual with a slightly above-average wage which corresponds to the second bend point on the Social Security benefit-earning rule. Individuals earning higher wages have a smaller welfare gain because of the progressivity of the benefit-earning rule. Individuals earning lower wages also have a smaller welfare gain. This is due to the crowding out of bequests. For very low-wage generation 2 workers, the welfare effect of Social Security is negative, holding generation 1 wages constant at the average. For these generation 2 individuals, a decrease in bequest income (because of Social Security) constitutes a large reduction in total lifetime income, which reduces their expected utility. Notice that by including bequest income in the analysis, Social Security ceases to be an entirely progressive program with low-income individuals actually better off without Social Security.

5.2. Life-Cycle Wage Dynamics, Interest Rate, and Discount Rate Sensitivity

In our baseline parameterization, we assume that wages are constant over the life cycle and that both the interest rate and discount rate are zero. In this section, we

relax these assumptions by adding life-cycle wage dynamics, a positive interest rate, and a discount rate. These modifications change the optimal paths of consumption and saving for both generations. The equations will not be provided here, as they are straightforward to derive.

We assume that wages follow a hump-shaped pattern according to the profile in Gourinchas and Parker (2002). While life-cycle wages depend on a variety of factors, the main point here is simply to capture a hump-shaped wage profile in contrast to the baseline flat wage profile. After extrapolating to obtain wage estimates for age groups that fall outside their original sample, a fifth-order polynomial provides a close fit to their data

$$w(t) = \lambda(1 + 0.0244t - 0.0001t^2 + 4(10)^{-5}t^3 - 2(10)^{-6}t^4 + 2(10)^{-8}t^5), \text{ for } t \in [0, t_R],$$

where λ is chosen such that the average wage equals 1. Introducing a hump-shaped wage profile changes the savings choices of individuals slightly. Individuals continue to smooth consumption over the life cycle, so this results in relatively less saving while young and relatively more saving in middle age. From the perspective of generation 2, this impacts their own saving/consumption choices directly, as well as impacting the distribution of possible bequests from generation 1. However, in practice, these differences are quantitatively small. The welfare gain of Social Security for generation 2 is 3.38% when both generations earn wages according to the equation above, with the interest rate and discount rate both set to zero. This compares to a welfare gain of 3.39% in the baseline (constant wage) parameterization.

Adding a positive interest rate increases the incentives for both generations to save. This increases the average size of bequests passed from generation 1 to generation 2. Social Security continues to crowd out bequest income. From the perspective of generation 2, this is more harmful than in the baseline, because bequests are larger on average. Additionally, a positive interest rate makes the economy (more) dynamically efficient and makes pay-as-you-go Social Security relatively less attractive than private savings. This drives the welfare gain of Social Security negative. In contrast, adding a positive discount rate to the model reduces the incentive to save, as individuals value future consumption relatively less than current consumption.

When the discount rate and interest rate are both set to 3.5% and wages follow the hump-shaped profile above, the welfare gain of Social Security is -20.0% . This large negative number is driven almost entirely by the positive interest rate and the dynamic efficiency of the economy. For example, if we turn off the interest rate, by setting it back to zero, and leave a positive discount rate and hump-shaped wage profile, the welfare gain of Social security is 2.5%, which is much closer to the baseline welfare gain of 3.4%.

Hence, while the particular assumptions that we make about the discount rate and the wage rate certainly affect the shape of the savings profiles over the life

cycle, these parameters do not play a major role in our welfare conclusions. Our baseline assumption of zero interest, however, is nontrivial. Larger rates of interest weaken Social Security’s welfare role, as is commonly understood, though the effect on welfare is especially strong in our model with bequest income.

5.3. Long Run

In our sequential equilibrium setting, expected bequest wealth builds upon itself with each successive generation as individuals die with unspent assets. Additionally, the dispersion of bequest wealth also grows with each successive generation because the permutations of mortality combinations among one’s ancestral line become multiplied. To explore the implications of such intergenerational wealth dynamics, we introduce a third generation into our analysis and we assess the welfare effects of Social Security on generation 3.

Like generation 2, generation 3 receives an uncertain bequest, allocated at age 0. The optimal asset holdings of generation 2 conditional on death at age $t = y$ and conditional on death of generation 1 at age $t = x$ are

$$k^*(y|x) = B^*(x) + \int_0^y [(1 - \tau)w - c(t|x)]dt, \text{ for } y \in [0, t_R],$$

$$k^*(y|x) = B^*(x) + \int_0^{t_R} [(1 - \tau)w - c(t|x)]dt + \int_{t_R}^y [b - c(t|x)]dt, \text{ for } y \in [t_R, T].$$

Note that the expected bequest to generation 3 is

$$\mathbb{E}(k^*) = \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) k^*(y|x) dx dy,$$

and the variance is

$$\text{var}(k^*) = \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) (k^*(y|x) - \mathbb{E}(k^*))^2 dx dy.$$

The Ex Ante Efficient allocation (i.e. First Best) for an individual in generation 3 is

$$c_{FB} = \frac{\int_0^{t_R} S(t)w dt + \mathbb{E}(k^*)}{\int_0^T S(t)dt}.$$

However, an individual living in a decentralized economy instead faces the following optimization problem, conditional on realization of the mortality of their ancestors (x and y). For notational clarity, we will name the asset holdings of generation 3 at age t as $a(t)$ to distinguish from the assets of generation 2 ($k(t)$) and generation 1 ($B(t)$):

$$\max \int_0^T S(t)u(c(t))dt,$$

subject to

$$\begin{aligned} \dot{a}(t) &= (1 - \tau)w - c(t), \text{ for } t \in [0, t_R], \\ \dot{a}(t) &= b - c(t), \text{ for } t \in [t_R, T], \\ a(0) &= k^*(y|x), \\ a(T) &= 0. \end{aligned}$$

The solution for CRRA utility is

$$c(t|x, y) = \frac{(1 - \tau)wt_R + b(T - t_R) + k^*(y|x)}{\int_0^T S(t)^{1/\sigma} dt} S(t)^{1/\sigma},$$

and ex ante expected utility before the realization of $k^*(y|x)$ is

$$\mathbb{E}(U) = \int_0^T \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) S(t)u(c(t|x, y))dtdxdy.$$

A **Sequential Equilibrium** with optimal hedging of longevity risk is comprised of generation 3 allocations $(c(t|x, y), a(t|x, y))_{t \in [0, T]}$ for each (x, y) pair that solve each individual’s optimization problem conditional bequest income that is realized at time $t = 0$ and is distributed according to the joint PDF $(-\dot{S}(y))(-\dot{S}(x))$ and optimal asset holdings of generation 2, $k^*(y|x)$, and Social Security tax rate τ and benefit b that jointly balance the Social Security budget.

The welfare gain from Social Security Δ_{SS} is

$$\begin{aligned} &\int_0^T \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) S(t)u(c_{SS}(t|x, y))(1 - \Delta_{SS})dtdxdy \\ &= \int_0^T \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) S(t)u(c_{LF}(t|x, y))dtdxdy. \end{aligned}$$

And for comparison, the welfare gain associated with the Ex Ante Efficient allocation Δ_{FB} is

$$\int_0^T S(t)u(c_{FB}(1 - \Delta_{FB}))dt = \int_0^T \int_0^T \int_0^T (-\dot{S}(y)) (-\dot{S}(x)) S(t)u(c_{LF}(t|x, y))dtdxdy.$$

The bequests passed from generation 2 to generation 3, as well as the welfare gains for generation 3 are presented in Table 5. The magnitude of bequest income compounds from one generation to the next. The average bequest received by generation 3 (in Table 5) is larger in magnitude than the average bequest received by generation 2 (in Table 2). This is because lifetime income (wages and bequests) is higher in generation 2 than in generation 1. Thus, over the life cycle, generation 2 holds more assets and leaves larger bequests to generation 3 on average.

The riskiness of bequest income also compounds across generations due to expansion in the number of permutations of mortality combinations among one’s ancestral line. For example, an individual in generation 3 inherits a very small bequest if their ancestors in generation 1 and 2 both died very young, before

TABLE 5. Welfare gains of Social Security Δ_{SS} and of the Ex Ante Efficient allocation Δ_{FB} for generation 3. The right half of the table shows the expected value and the standard deviation of bequests inherited by generation 3 with Social Security (k_{SS}^*) and in Laissez-Faire (k_{LF}^*)

σ	Generation 3		Bequests from Gen 2 to Gen 3			
	Δ_{SS}	Δ_{FB}	$\mathbb{E}(k_{SS}^*)$	$SD(k_{SS}^*)$	$\mathbb{E}(k_{LF}^*)$	$SD(k_{LF}^*)$
1	0.6%	14.5%	-1.17	0.95	5.03	4.11
1.5	1.0%	17.0%	0.53	1.62	6.68	5.05
2	1.3%	18.6%	1.65	2.12	7.76	5.59
3	1.9%	20.8%	3.03	2.70	9.08	6.19
4	2.5%	22.2%	3.84	3.03	9.86	6.51
5	3.1%	23.3%	4.38	3.23	10.37	6.71

they had a chance to accumulate assets. Similarly, an individual in generation 3 inherits a very large bequest if their ancestors in generation 1 and 2 both died near retirement age when assets reach their peak. The potential dispersion of bequests for generation 3 is much larger than for generation 2.

The increase in magnitude and riskiness of bequests passed to generation 3 have important and nuanced effects on risk sharing. The welfare gain of perfect insurance, Δ_{FB} , is larger for generation 3 than for the previous generation. For example, for our baseline risk-aversion parameter $\sigma = 3$, the welfare gain of perfect insurance for generation 2 is 20.2% while for generation 3 is 20.8%. The welfare effect of Social Security Δ_{SS} is slightly smaller for generation 3 than generation 2, however, the welfare gain still remains relatively large and positive. For example, with $\sigma = 3$, the welfare gain for generation 3 is 1.9% of consumption.

In sum, as we move sequentially through the generations in our model, expected bequest income grows and so does the riskiness of bequest income. Such long-term wealth dynamics have important and nuanced effects on Social Security's role as a risk sharing device. Generation 3 is wealthier (in expectation) than generation 2, because bequest income received by generation 3 is greater than that received by generation 2. And this makes Social Security's crowding out of bequest income especially painful to generation 3, because bequest income represents a large share of generation 3's budget. However, the bequest received by generation 3 is also riskier than the bequest received by generation 2, and this partially mitigates the cost associated with the crowding out of bequest income. The first force dominates the second and generations 3's welfare gains from Social Security are smaller than the gains experienced by generation 2, but the welfare gains are still significant.

6. NO HEDGING

To complete our analysis, we consider the possibility that Social Security protects the individual from outliving their assets. Recall that in a standard

Yaari life-cycle model the individual optimally self insures their longevity risk and does not actually face the risk of outliving their assets. However, Social Security’s life annuity feature is often motivated on the grounds that it provides protection against this very risk. Indeed, Brown (2000) states, “If all individuals have near-perfect information and behaved rationally, they would all adequately provide for the contingency of living to advanced ages. Unfortunately, due to factors such as imperfect information about mortality risk, ... it is quite likely that some individuals will fail to adequately insure themselves against outliving their resources.” In this section, we wish to study Social Security’s longevity insurance role when we relax the optimal hedging assumption.⁶

We consider generation 2, but now we assume the individual does not hedge their longevity risk and instead solves a deterministic problem of consumption and saving that is based on their life expectancy. That is, the individual assumes they will die with certainty at $\mathbb{E}(t) \equiv \int_0^T S(t)dt$, and formulates a consumption/saving plan that is optimal based on this naïve assumption. In general, when a decision maker fails to hedge a given risk, it means they are basing their decisions on the mean of a random variable while ignoring the variance. This is exactly what our individual is doing when they base their consumption/saving decision on their expected length of life $\mathbb{E}(t)$ —they are ignoring all of the risk surrounding this expectation. But notice that the individual is not totally irrational either. They correctly forecast their life expectancy, and they solve an optimal retirement saving problem based on this forecast. Their only error lies in the failure to account for risk.⁷

The deterministic saving plan of the naïve individual involves exhausting their assets completely at their life expectancy. If the individual happens to die before their life expectancy, then they die with nonzero assets. If the individual survives beyond their life expectancy (and therefore outlives their private assets), then they survive by consuming the Social Security benefits only, and if there is no Social Security program then their consumption is zero until death. CRRA utility is defined when consumption is zero only if $\sigma < 1$, so we limit our attention to this case.

The generation 2 optimization problem with no hedging, conditional on realization of x , is

$$\max \int_0^{\mathbb{E}(t)} u(c(t))dt,$$

subject to

$$\dot{k}(t) = (1 - \tau)w - c(t), \text{ for } t \in [0, t_R],$$

$$\dot{k}(t) = b - c(t), \text{ for } t \in [t_R, \mathbb{E}(t)],$$

$$k(0) = B^*(x),$$

$$k(\mathbb{E}(t)) = 0,$$

where $B^*(x)$ is the bequest received from a specific individual from generation 1 who died at age x , which is the asset holdings of that individual at their age of death. We assume generation 1 does in fact solve an optimal hedging problem, so that the distribution of bequest income $B^*(x)$ continues to be the same as before in order to facilitate welfare comparisons.

The solution for CRR is

$$\hat{c}(t|x) = \begin{cases} \frac{(1 - \tau)wt_R + b(\mathbb{E}(t) - t_R) + B^*(x)}{\mathbb{E}(t)} & \text{for } t \in [0, \mathbb{E}(t)]. \\ b & \text{for } t \in [\mathbb{E}(t), T]. \end{cases}$$

A **Sequential Equilibrium** with no hedging of longevity risk is comprised of generation 2 allocations that solve the naïve problem above for each x conditional on bequest income that is realized at time $t = 0$ and is distributed according to the PDF $(-\dot{S}(x))$ and optimal asset holdings of generation 1, $B^*(x)$, and Social Security tax rate τ and benefit b that jointly balance the Social Security budget.

The ex ante welfare gain from Social Security Δ_{SS} is

$$\int_0^T \int_0^T (-\dot{S}(x)) S(t)u(\hat{c}_{SS}(t|x)(1 - \Delta_{SS}))dt dx = \int_0^T \int_0^T (-\dot{S}(x)) S(t)u(\hat{c}_{LF}(t|x))dt dx.$$

And for comparison, the welfare gain associated with the Ex Ante Efficient allocation Δ_{FB} is

$$\int_0^T S(t)u(c_{FB}(1 - \Delta_{FB}))dt = \int_0^T \int_0^T (-\dot{S}(x)) S(t)u(\hat{c}_{LF}(t|x))dt dx.$$

Setting $\sigma = 0.5$, we find that the welfare gain from Social Security Δ_{SS} is 2.5% and the welfare gain associated with the Ex Ante Efficient allocation Δ_{FB} is 16.7%. Both welfare gains are about 1.7 times larger than in the baseline case where the individual optimally hedges their longevity risk (as in Section 4.2). When the individual hedges their own longevity risk, the welfare gain of Social Security is only 1.4% and the welfare associated with the First Best allocation is 9.9%. The welfare gains are larger in this no hedging exercise because the individual faces the risk of outliving their assets and having zero consumption in old age. This is relatively more costly than the risk of dying with unconsumed assets in the fully rational model with longevity risk. Thus, the traditional, rational model significantly understates the welfare gains from collective sharing of longevity risk through Social Security.

As a final comparison, we consider the case of an individual who “over-hedges” and saves as if they were going to live to the maximum age T with certainty. The individual perfectly smooths their consumption over the maximum possible life span by saving too much compared to the fully rational model. The welfare gain of the First Best consumption compared to the over-hedged consumption is 22.4% for $\sigma = 0.5$. This is larger than in the fully rational model (welfare gain of 9.9%) because the individual over-saves in this example. In this setting, Social Security is still welfare enhancing. Social Security improves the ex ante expected lifetime utility of the individual by crowding out some of their saving in two ways. Social

Security taxes reduce take-home pay, which reduces the individual's ability to save. Since they were saving too much, this is an improvement. Social Security benefits also reduce the incentive to save, since they provide an annuity payment in retirement. The ex ante welfare gain from Social Security is 1.4% for $\sigma = 0.5$. The welfare gain of Social Security is the same as in the baseline example (this is true for any σ). This is because the welfare gain works through the same channel. In both cases, Social Security improves welfare by providing an annuity benefit that crowds out private savings.

7. LATE RESOLUTION OF RISK: TIMING UNCERTAINTY WITH DYNAMIC HEDGING

In the analysis above, we took a step toward a more realistic bequest process by allowing for uncertainty in the magnitude of bequest income received by generation 2, but we simplified the analysis by assuming that this risk is resolved immediately for generation 2. Of course, in reality generation 2 does not know the magnitude of the bequest income they will receive nor do they know the *timing*. In this section, we solve the full blown model with uncertainty over the timing and magnitude of bequest income.

To do this, we construct a continuous-time overlapping generations model. *Calendar time* is continuous and is indexed by t . At $t = 0$ the first generation enters the workforce. Hence, for generation 1, the index t is also equal to the individual's age. The probability an individual from this generation survives to age/time t is $S(t)$, and death occurs no later than $t = T$, hence $S(0) = 1$ and $S(T) = 0$ as usual.

At calendar date $t = t_0$, the second generation enters the workforce and lives to a maximum date of $t = t_0 + T$. The probability that an individual in generation 2 survives to date t is $P(t)$. As usual, $P(t_0) = 1$ and $P(t_0 + T) = 0$. Because we assume stationarity in longevity risk across the two generations,

$$P(t) = S(t - t_0),$$

since an individual in generation 2 is $(t - t_0)$ years old at calendar date t .

We focus on the dynamic stochastic problem facing an individual in generation 2. This is not a standard problem because optimal decisions must reflect an *additional precautionary saving motive* resulting from bequest risk: the child knows that they will receive a one-time injection of wealth, but the magnitude and timing of this injection are uncertain. Hence, the child must solve a dynamic hedging problem in order to properly save in the face of this risk. To solve this problem, we utilize timing uncertainty methods from the stochastic control literature [see Stokey (2016), Caliendo et al. (2019), and Caliendo et al. (2020)].

We break into two cases. The first case is when the parent dies before the child enters the workforce. In this case, the parent's assets are held in a trust and are made available to the child at $t = t_0$ when the child enters the workforce. The second case is when the parent is still alive when the child enters the workforce. While the first case is a standard consumption/saving problem, the second case is much more involved. Of course, the survival probabilities of the parent determine

the likelihood that the child operates in the first or second case, and the second case is much more likely.

To compress notation a little bit, we define disposable income as follows:

$$y(t) \equiv \begin{cases} (1 - \tau)w, & \text{for } t \in [t_0, t_0 + t_R], \\ b, & \text{for } t \in [t_0 + t_R, t_0 + T]. \end{cases}$$

7.1. Case 1. The Deterministic Problem of a Child Whose Parent Died Early

First consider an individual from generation 2 who enters the workforce at t_0 as an orphan. This happens with probability $(1 - S(t_0))$. The child is endowed at t_0 with $B^*(t)$ where $t \in [0, t_0]$ is the date of death of the parent, and the child solves the following deterministic problem

$$\max_{c(z)_{z \in [t_0, t_0+T]}} : \int_{t_0}^{t_0+T} P(z)u(c(z))dz,$$

subject to

$$\begin{aligned} \dot{k}(z) &= y(z) - c(z), \\ k(t_0) &= B^*(t), \quad k(t_0 + T) = 0, \end{aligned}$$

where $B^*(t)$ is the optimal asset holdings of the parent at date t . The solution to the above problem is

$$c_d^*(z|B^*(t)) = \frac{B^*(t) + \int_{t_0}^{t_0+T} y(v)dv}{\int_{t_0}^{t_0+T} P(v)^{1/\sigma} dv} P(z)^{1/\sigma}, \text{ for } z \in [t_0, t_0 + T].$$

The subscript “ d ” stands for deterministic.

7.2. Case 2. The Stochastic Problem of a Child Whose Parent is Living

Next, we model the problem of a child whose parent is still alive at time t_0 . In this case, the child must solve a dynamic stochastic problem with a variable inheritance occurring at an unknown date. The child must form an optimal plan in the face of not only their own longevity risk but also that of the parent. The child has full information about the distribution of the risks that they face (the survival probabilities of the parent and child are known to the child). The recursive solution to the stochastic problem of the child whose parent is alive at time t_0 goes as follows.

Step 1. The post-inheritance problem

Suppose the parent passes away at time $t \geq t_0$. Upon learning the timing and magnitude of the inheritance at date t , the child solves⁸

$$\max_{c(z)_{z \in [t, t_0+T]}} : \int_t^{t_0+T} P(z)u(c(z))dz,$$

subject to

$$\begin{aligned} \dot{K}(z) &= y(z) - c(z), \\ t \text{ given, } K(t) &= k(t) + B^*(t), K(t_0 + T) = 0, \end{aligned}$$

where $K(t)$ is the total financial wealth of the child, which is the sum of his private savings $k(t)$ plus his inheritance $B^*(t)$, which is the optimal asset holdings of the parent at date t . The solution is

$$c_2^*(z|t, k(t), B^*(t)) = F(t, k(t), B^*(t))P(z)^{1/\sigma}, \text{ for } z \in [t, t_0 + T],$$

where we have compressed notation for simplicity

$$F(t, k(t), B^*(t)) \equiv \frac{k(t) + B^*(t) + \int_t^{t_0+T} y(v)dv}{\int_t^{t_0+T} P(v)^{1/\sigma} dv}.$$

The “2” subscript denotes that this is the solution to the second-stage (post-inheritance) portion of the life cycle.

Step 2. The pre-inheritance stochastic problem

Working backwards to time t_0 , we solve a dynamic stochastic problem of an individual who hedges timing risk by accounting for the continuation value of accumulated assets and the fact that the shock can happen at any time along the continuum $[t_0, T]$. Because the timing of the death of the parent t is a random variable, the child’s expected utility from the perspective of entering the workforce (at time t_0), conditional on the parent being alive at time t_0 , is

$$\begin{aligned} \mathbb{E}(U)_{t \geq t_0} &= \int_{t_0}^T \left(-\dot{S}(t) \frac{1}{S(t_0)} \right) \\ &\quad \times \left(\int_{t_0}^t P(z)u(c(z))dz + \int_t^{t_0+T} P(z)u(c_2^*(z|t, k(t), B^*(t)))dz \right) dt. \end{aligned}$$

Notice with the outer integral, we are integrating up to the last possible date of death of the parent, $t = T$. To reduce clutter, we can drop the scalar constant $1/S(t_0)$ from expected utility without consequence. Then, with a change in the order of integration on the first double integral, $\int_{t_0}^T \int_{t_0}^t (\cdot) dz dt = \int_{t_0}^T \int_z^T (\cdot) dt dz$, we have

$$\begin{aligned} \mathbb{E}(U)_{t \geq t_0} &= \int_{t_0}^T \int_z^T (-\dot{S}(t)) P(z)u(c(z)) dt dz + \int_{t_0}^T \int_t^{t_0+T} (-\dot{S}(t)) P(z)u(c_2^*(\cdot)) dz dt \\ &= \int_{t_0}^T S(z)P(z)u(c(z)) dz + \int_{t_0}^T \int_t^{t_0+T} (-\dot{S}(t)) P(z)u(c_2^*(\cdot)) dz dt \\ &= \int_{t_0}^T S(t)P(t)u(c(t)) dt + \int_{t_0}^T \int_t^{t_0+T} (-\dot{S}(t)) P(z)u(c_2^*(\cdot)) dz dt. \end{aligned}$$

With this algebra behind us, we can now express the stochastic optimization problem as a Pontryagin problem

$$\max_{c(t)_{t \in [t_0, T]}} : \int_{t_0}^T \left(S(t)P(t)u(c(t)) + (-\dot{S}(t)) \int_t^{t_0+T} P(z)u(c_2^*(z|t, k(t), B^*(t)))dz \right) dt,$$

subject to

$$\begin{aligned} \dot{k}(t) &= y(t) - c(t), \text{ for } t \in [t_0, T], \\ k(t_0) &= 0, k(T) \text{ free,} \\ c_2^*(z|t, k(t), B^*(t)) &= F(t, k(t), B^*(t))P(z)^{1/\sigma}, \text{ for } z \in [t, t_0 + T]. \end{aligned}$$

The Hamiltonian \mathcal{H} with multiplier $\lambda(t)$ is

$$\mathcal{H} = S(t)P(t)u(c(t)) + (-\dot{S}(t)) \int_t^{t_0+T} P(z)u(c_2^*(\cdot))dz + \lambda(t)[y(t) - c(t)],$$

with necessary conditions

$$\begin{aligned} \frac{\partial \mathcal{H}}{\partial c(t)} &= S(t)P(t)c(t)^{-\sigma} - \lambda(t) = 0, \\ \dot{\lambda}(t) &= -\frac{\partial \mathcal{H}}{\partial k(t)} = \dot{S}(t) \int_t^{t_0+T} P(z)c_2^*(\cdot)^{-\sigma} \frac{\partial c_2^*(\cdot)}{\partial k(t)} dz, \\ \lambda(T) &= 0. \end{aligned}$$

Upon taking the partial derivative $\partial c_2^*(\cdot)/\partial k(t)$, the multiplier equation can be simplified (with some algebra) to

$$\dot{\lambda}(t) = \dot{S}(t) [F(t, k(t), B^*(t))]^{-\sigma}.$$

Differentiate the Maximum Condition with respect to t and combine with the multiplier equation

$$\left(\frac{d}{dt} (S(t)P(t)) \right) c(t)^{-\sigma} - \sigma S(t)P(t)c(t)^{-\sigma-1} \dot{c}(t) = \dot{S}(t) [F(t, k(t), B^*(t))]^{-\sigma},$$

and solve for $\dot{c}(t)$

$$\dot{c}(t) = \left(\frac{\dot{S}(t)}{S(t)} + \frac{\dot{P}(t)}{P(t)} \right) \frac{c(t)}{\sigma} - \frac{\dot{S}(t)}{S(t)P(t)} [F(t, k(t), B^*(t))]^{-\sigma} \frac{c(t)^{\sigma+1}}{\sigma}.$$

Finally, we need to make use of the Transversality Condition, and we do so as follows. Rewrite the Maximum Condition

$$P(t)c(t)^{-\sigma} = \frac{\lambda(t)}{S(t)},$$

and note from the Transversality Condition that we have the indeterminate form 0/0 when $t = T$ so

$$P(T)c(T)^{-\sigma} = \lim_{t \rightarrow T} P(t)c(t)^{-\sigma} = \lim_{t \rightarrow T} \frac{\lambda(t)}{S(t)} = \lim_{t \rightarrow T} \frac{\dot{\lambda}(t)}{\dot{S}(t)} = \frac{\dot{\lambda}(T)}{\dot{S}(T)}$$

or

$$P(T)c(T)^{-\sigma} = \frac{\dot{S}(T) [F(T, k(T), B^*(T))]^{-\sigma}}{\dot{S}(T)},$$

which further simplifies to

$$c(T) = P(T)^{1/\sigma} F(T, k(T), B^*(T)).$$

We refer to this as the limiting case of the Transversality Condition, and it provides the needed endpoint condition to fully identify the solution.

Having now solved the model analytically, we summarize how the analytical solution is used to numerically simulate Case 2 on the computer:

- Step 1. Guess initial consumption $c(t_0)$.
- Step 2. “Shoot” forward to obtain the timepath of $c(t)$ and $k(t)$ over the interval $t \in [t_0, T]$ using the following system of equations and boundary condition

$$\dot{c}(t) = \left(\frac{\dot{S}(t)}{S(t)} + \frac{\dot{P}(t)}{P(t)} \right) \frac{c(t)}{\sigma} - \frac{\dot{S}(t)}{S(t)P(t)} [F(t, k(t), B^*(t))]^{-\sigma} \frac{c(t)^{\sigma+1}}{\sigma},$$

for $t \in [t_0, T]$,

$$\dot{k}(t) = y(t) - c(t), \text{ for } t \in [t_0, T],$$

$$k(t_0) = 0.$$

- Step 3. Check to see if the limiting case of the Transversality Condition is satisfied,

$$c(T) = P(T)^{1/\sigma} F(T, k(T), B^*(T)).$$

If so, then the guessed value $c(t_0)$ is correct and we have identified the optimal pre-inheritance contingency path $\{c_1^*(t), k_1^*(t)\}$ and the algorithm stops. If not, go back to Step 1 and repeat.

7.3. Welfare Gain from Social Security

To summarize, if the parent dies at some date t before the child enters the workforce (i.e. $t \leq t_0$), the parent’s assets at the time of death $B^*(t)$ are held in a trust and passed along to the child at time t_0 and the child’s consumption follows the deterministic path c_d^* for the duration of the child’s life cycle. If instead the parent dies at some date t after the child has entered the workforce, (i.e. $t > t_0$), then the child follows the optimal hedging, contingent consumption path c_1^* up to the death of the parent. After that, the child follows a new path c_2^* that depends on the timing of the shock t , the level of own assets accumulated $k_1^*(t)$, and the magnitude of the inheritance $B^*(t)$. Of course, since the timing of the inheritance is uncertain, the child’s expected utility must take into account all possible dates of death of the parent.

Integrating over all potential dates of death of the parent t , while accounting for the probability of death $(-\dot{S}(t))$ and optimal parental asset holdings at death $B^*(t)$, the ex ante expected utility of a child who employs an optimal, dynamic hedging strategy is

$$\begin{aligned} \mathbb{E}(U) = & \int_0^{t_0} (-\dot{S}(t)) \left(\int_{t_0}^{t_0+T} P(z)u(c_d^*(z|B^*(t)))dz \right) dt \\ & + \int_{t_0}^T (-\dot{S}(t)) \left(\int_{t_0}^t P(z)u(c_1^*(z))dz \right. \\ & \left. + \int_t^{t_0+T} P(z)u(c_2^*(z|t, k_1^*(t), B^*(t)))dz \right) dt. \end{aligned}$$

The welfare gain from Social Security is the fraction of lifetime consumption the child would give up, ex ante before risks are realized, to live in a world with Social Security,

$$\Delta_{SS} = 1 - \left(\frac{\mathbb{E}(U)_{LF}}{\mathbb{E}(U)_{SS}} \right)^{\frac{1}{1-\sigma}}.$$

To compute Δ_{SS} , we use all of our baseline parameters from Table 1. The only additional parameter that needs to be specified is t_0 (the number of years between generation 1 and generation 2 entering the workforce). We set $t_0 = 30$. Our computational method allows us to match the limiting case of the Transversality Condition with great precision, and we calculate that the welfare effect of Social Security is 4.46%, which is somewhat larger than our baseline value (3.4%).

Hence, in addition to being very tractable, our baseline analysis with early resolution of bequest risk is also conservative in its estimation of the welfare effect of Social Security. Here, with later resolution of bequest risk, Social Security’s welfare effect is bigger. By adding a second layer of risk to the model (i.e. bequest *timing* risk in addition to bequest *magnitude* risk), Social Security’s crowding out of bequest income becomes less costly than in our baseline model, which in turn implies that Social Security is even more effective at hedging longevity risk.

8. CONCLUSION

Not knowing the timing of death creates a major challenge for retirement planning. The length of the retirement period is a key factor in formulating a saving plan that can provide financial security during the retirement years of life. Rational individuals in textbook economic theory optimally hedge longevity risk by accumulating precautionary saving balances to avoid running out of assets before death. Absent a bequest motive, such hedging is inefficient because individuals tend to die with unspent asset. This inefficiency creates a wedge between

the Ex Ante Efficient allocation and an equilibrium allocation, thereby opening the door for a second-best program such as Social Security to potentially improve welfare.

In this paper, we revisit a basic question: How well does Social Security insure longevity risk? We approach this question from a unique angle of sequential welfare analysis. We study an initial generation (generation 1) that participates in an efficiently financed Social Security program. Consistent with conventional wisdom about the welfare gains from annuitization, this generation experiences massive welfare gains from Social Security.

Then, we consider a second generation (generation 2) that is the same as generation 1 in almost every respect, with one crucial exception: this generation inherits bequest income from generation 1, and this bequest income is risky because the timing of death of generation 1 is uncertain and the asset holdings of generation 1 vary with age. Hence, unlike traditional models, our model accounts for specific linkages between parents and children in order to model bequest income risk. Now, Social Security carries both benefits and costs: on the benefit side, it provides collective risk sharing of longevity risk because it pays benefits as a life annuity; but on the cost side, it crowds out bequest income from the previous generations. The key insight in our analysis is that the net welfare effect depends crucially on the riskiness of bequest income: yes it is costly to crowd out bequest income, but the magnitude of this cost is partially mitigated by virtue of the fact that Social Security is crowding out something that is inherently risky. Essentially, our model is a continuous-time, computational adaptation of the two-period model developed in the appendix of Caliendo et al. (2014).

In the face of both bequest income risk and longevity risk, we find that the Ex Ante Efficient allocation confers vastly higher utility for individuals in generation 2 than ex ante expected utility. A Social Security program calibrated to the US economy does not come close to providing the level of insurance that the Ex Ante Efficient allocation can provide, but it still provides a remarkably large welfare gain (on the order of 3.4% at our preferred level of risk aversion). In our view, Social Security is quite successful in its longevity insurance role.

We pay particular attention to the degree to which individuals hedge their longevity risk. We consider two bookend examples in an attempt to bound the welfare gains from Social Security. In one scenario, the individual behaves in the standard, rational manner and fully accounts for longevity risk when saving for the future. In the other case, the individual makes a saving plan that is based on the correct life expectancy, but the individual ignores the risk associated with this expectation. In the former case, there is no risk of outliving one's assets, whereas in the latter case the individual solves a suboptimal problem and will indeed run out of assets if they live longer than expected. We show that the welfare gains from Social Security can be substantially larger when individuals fail to hedge their longevity risk. In other words, standard economic models tend to underestimate the potential gains associated with Social Security's annuitization feature.

NOTES

1. De Nardi (2004) is perhaps the most notable exception of this assumption in the macroeconomic literature; in her model, children inherit earning ability from their parent as well as receiving a bequest “at a random date which depends on their parent’s death.”

2. For a discussion the potential interplay between Social Security and the problem of adverse selection in private annuity markets, see Hosseini (2015).

3. Of course, the ratio of workers to retirees is projected to decline in the coming decades as life expectancies continue to improve. This strains the Social Security budget and creates the need for reform. Optimal Social Security reform is not our focus but is the subject of a large literature including Huang et al. (1997), Conesa and Garriga (2008), Kitao (2014), and McGrattan and Prescott (2017) among many other examples.

4. Individuals who enjoy a longer life span will reap greater welfare gains from Social Security’s annuitization feature, and the wealthy tend to live longer than the poor. This fact tends to unwind some of the progressivity that is otherwise implied by the benefit-earning formula [see Liebman (2002), Coronado et al. (1999, 2002, 2011), Goda et al. (2011), and Bishnu et al. (2019)].

5. The parameter σ does not change the welfare gain of generation 1 because we are working with CRRA utility and because our welfare metric is a percentage reduction in consumption that enters the period utility function multiplicatively.

6. While a number of papers consider the role of Social Security when individuals fall short of saving optimally for retirement [Feldstein (1985), İmrohoroğlu et al. (2003), Andersen and Bhattacharya (2011), and Bucciol (2011) to name a few], we do so in a unique way by assuming people ignore longevity risk.

7. Brown (2000) conjectures that behavior such as we assume here, may be another explanation for the lack of private annuitization in general. If people don’t appreciate the risk associated with their life expectancy, then they wouldn’t really care to buy insurance that hedges this risk.

8. Standing at time $t \geq t_0$, the probability the child survives to date z is the conditional probability $P(z)/P(t)$, but we can use the unconditional probability $P(z)$ in the maximand without consequence because $1/P(t)$ is a scalar constant.

REFERENCES

- Alonso-Ortiz, J. (2014) Social security and retirement across the OECD. *Journal of Economic Dynamics and Control* 47, 300–316.
- Andersen, T. and J. Bhattacharya (2011) On myopia as rationale for social security. *Economic Theory* 47(1), 135–158.
- Bishnu, M., N. L. Guo and C. S. Kumru (2019) Social security with differential mortality. *Journal of Macroeconomics* 62, 103077.
- Brown, J. R. (2000) How Should We Insure Longevity Risk In Pensions And Social Security? Issues in Brief ib-4, Center for Retirement Research.
- Brown, J. R. (2001) Private pensions, mortality risk, and the decision to annuitize. *Journal of Public Economics* 82(1), 29–62.
- Bucciol, A. (2011) A note on social security welfare with self-control problems. *Macroeconomic Dynamics* 15(4), 579–594.
- Bütler, M. (1999) Anticipation effects of looming public-pension reforms. *Carnegie-Rochester Conference Series on Public Policy* 50, 119–159.
- Bütler, M. (2001) Neoclassical life-cycle consumption: A textbook example. *Economic Theory* 17(1), 209–221.
- Caliendo, F. N., M. Casanova, A. Gorry and S. Slavov (2020) Retirement timing uncertainty: Empirical evidence and quantitative evaluation. *working paper*.
- Caliendo, F. N., Aspen G. and S. Slavov (2019) The cost of uncertainty about the timing of social security reform. *European Economic Review* 118, 101–125.

- Caliendo, F. N., N. L. Guo and R. Hosseini (2014) Social security is not a substitute for annuity markets. *Review of Economic Dynamics* 17(4), 739–755.
- Conesa, J. C. and C. Garriga (2008) Optimal fiscal policy in the design of social security reforms*. *International Economic Review* 49(1), 291–318.
- Conesa, J. C. and D. Krueger (1999) Social security reform with heterogeneous agents. *Review of Economic Dynamics* 2(4), 757–795.
- Coronado, J. L., D. Fullerton and T. Glass (1999) *Distributional Impacts of Proposed Changes to the Social Security System*, pp. 149–186. MIT Press.
- Coronado, J. L., D. Fullerton and T. Glass (2002) *Long-Run Effects of Social Security Reform Proposals on Lifetime Progressivity*. In Feldstein and Liebman (2002), pp. 149–206.
- Coronado, J. L., D. Fullerton and T. Glass (2011) The progressivity of social security. *The B.E. Journal of Economic Analysis & Policy* 11(1).
- Cottle Hunt, E. (2020) Adaptive learning, social security reform, and policy uncertainty. *Journal of Money, Credit, and Banking*, forthcoming.
- Cremer, H., P. Donder, D. Maldonado and P. Pestieau (2008) Designing a linear pension scheme with forced savings and wage heterogeneity. *International Tax and Public Finance* 15(5), 547–562.
- Cremer, H. and P. Pestieau (2011) Myopia, redistribution and pensions. *European Economic Review* 55(2), 165–175.
- Davidoff, T., J. R. Brown and P. A. Diamond (2005) Annuities and individual welfare. *American Economic Review* 95(5), 1573–1590.
- De Nardi, M. (2004) Wealth inequality and intergenerational links. *The Review of Economic Studies* 71(3), 743–768.
- Fehr, H. and C. Habermann (2008) Risk sharing and efficiency implications of progressive pension arrangements*. *The Scandinavian Journal of Economics* 110(2), 419–443.
- Feigenbaum, J. (2008) Can mortality risk explain the consumption hump? *Journal of Macroeconomics* 30(3), 844–872.
- Feldstein, M. (1985) The optimal level of social-security benefits. *Quarterly Journal of Economics* 100(2), 303–320.
- Feldstein, M. and J. B. Liebman (2002) *The Distributional Aspects of Social Security and Social Security Reform*. University of Chicago Press.
- Goda, G. S., J. B. Shoven and S. N. Slavov (2011) *Differential Mortality by Income and Social Security Progressivity*, pp. 189–204. University of Chicago Press.
- Gomes, F. J., L. J. Kotlikoff and L. M. Viceira (2012) The excess burden of government indecision. *Tax Policy and the Economy* 26(1), 125–164.
- Gourinchas, P.-O. and J. A. Parker (2002) Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Hosseini, R. (2015) Adverse selection in the annuity market and the role for social security. *Journal of Political Economy* 123(4), 941.
- Huang, H., S. İmrohoroğlu and T. J. Sargent (1997) Two computations to fund social security. *Macroeconomic Dynamics*, 1(1):7–44.
- Hubbard, R. G. and K. L. Judd (1987) Social security and individual welfare: Precautionary saving, borrowing constraints, and the payroll tax. *The American Economic Review* 77(4), 630.
- Huggett, M. and G. Ventura (1999) On the distributional effects of social security reform. *Review of Economic Dynamics* 2(3), 498–531.
- İmrohoroğlu, A., S. İmrohoroğlu and D. H. Joines (1995) A life-cycle analysis of social-security. *Economic Theory* 6(1), 83–114.
- İmrohoroğlu, A., S. İmrohoroğlu and D. H. Joines (2003) Time-inconsistent preferences and social security*. *The Quarterly Journal of Economics* 118(2), 745–784.
- Kitao, S. (2014) Sustainable social security: Four options. *Review of Economic Dynamics* 17(4), 756–779.
- Kitao, S. (2018) Policy uncertainty and cost of delaying reform: The case of aging Japan. *Review of Economic Dynamics* 27, 81–100.

- Kotlikoff, L. J. and A. Spivak (1981) The family as an incomplete annuities market. *Journal of Political Economy* 89(2), 372–391.
- Liebman, J. B. (2002) *Redistribution in the Current U.S. Social Security System*. In Feldstein and Liebman (2002), pp. 11–48.
- Lockwood, L. M. (2012) Bequest motives and the annuity puzzle. *Review of Economic Dynamics* 15(2), 226–243.
- Luttmer, E. F. P. and A. A. Samwick (2018) The welfare cost of perceived policy uncertainty: Evidence from social security. *American Economic Review* 108(2), 275–307.
- McGrattan, E. R. and E. C. Prescott (2017) On financing retirement with an aging population. *Quantitative Economics* 8(1), 75–115.
- Mitchell, O. S., J. M. Poterba, M. J. Warshawsky and J. R. Brown (1999) New evidence on the money's worth of individual annuities. *American Economic Review* 89(5), 1299–1318.
- Nelson, J. (2020) Welfare implications of uncertain social security reform. *Public Finance Review*, 48(4), 425–466.
- Stokey, N. L. (2016) Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics* 21, 246–265.
- Yaari, M. E. (1965) Uncertain lifetime, life insurance, and the theory of the consumer. *Review of Economic Studies* 32(2), 137–150.