

ARTICLE

Social Security safety net with rare event risk

Erin Cottle Hunt^{1,*}  and Frank N. Caliendo²

¹Department of Economics, Lafayette College, Easton, PA 18042, USA

²Department of Economics and Finance, Utah State University, Logan, UT, USA

*Corresponding author. Email: cottle@lafayette.edu. Phone: 610 330 3373

Abstract

The US Social Security program was created in part as an explicit response to the Great Depression. We evaluate the performance of Social Security as a protective safety net against a rare episode of sudden and significant loss of private wealth such as the Great Depression. We construct a model in which a rare event causes a shock to wealth at an unknown time. How well does Social Security function as a safety net against such risk? The answer depends critically on whether households optimize in the face of this risk. If the household has full information on the distribution of rare event risk and solves a dynamic stochastic problem to hedge this risk, then Social Security is unnecessary along this dimension. Alternatively, if the household does not account for rare event risk in its financial planning, then Social Security can provide very large welfare gains as a safety net.

Keywords: Social Security; rare event risk; asset risk

1. Introduction

As all investors who hold risky assets such as US equities can attest, there are essentially two types of risk to worry about. First, there is high-frequency risk. This is the day-to-day, or minute-to-minute etc., volatility that causes constant, idiosyncratic fluctuations in one's wealth. This risk is most problematic for investors with very short investment horizons, such as those making daily allocation decisions, and is somewhat less relevant for investors with longer holding periods. For example, while it can be stressful to watch one's wealth fluctuate with each high-frequency shock, those who are investing for the long run understand that a longer holding period can tend to smooth out much of the day-to-day noise. In fact, for relatively young individuals saving for retirement, the traditional advice is to take a long-run perspective that capitalizes on high equity returns while essentially ignoring the day-to-day volatility.

However, the second risk—low-frequency risk or rare event risk—remains problematic even for those investing for the long run. This risk refers to rare events that unpredictably lead to the sudden destruction of significant wealth.¹ The COVID-19 outbreak in the Spring of 2020 triggered the S&P 500 to suddenly lose approximately one-third of its value within just a few weeks. (Of course, within the same year, equity prices recovered rapidly alongside record earnings among technology companies, record low interest rates from rapid expansion in the money supply, and anticipation of a vaccine). But not all rare event shocks have been followed by rapid recovery. The S&P 500 reached a peak in October of 2007, lost more than half of its value in the Great Recession, and did not fully recover until more than 5 years had passed. This episode was problematic for workers on the verge of retirement who then faced the grim prospect of retiring with far less financial security than they had supposed. The same can be said of other rare events such as the dot.com recession in 2000 which featured a crash that erased almost half of the value of the S&P

500 and a slow recovery of asset prices lasting 7 years, as well as the Great Depression which quickly erased 86% of the value of the S&P 500 followed by a painfully slow recovery lasting 25 years.

Social Security was created in 1935 in response to the Great Depression with the goal to ensure the financial security of the elderly. The purpose of this paper is to examine the welfare role of Social Security as a safety net against rare event risk. To study this issue, we blend traditional Social Security welfare analysis—which compares *ex ante* lifetime utility with and without Social Security—with the dynamic stochastic control literature. We study a life-cycle model in which individuals make optimal consumption and saving decisions in the face of a large, negative shock to their private financial wealth occurring at an unknown time. Losing a large fraction of one's financial wealth is relatively painless if the individual has not yet accumulated very much wealth, but such a shock can have a major impact on the individual's standard of living if it occurs close to retirement when wealth is at a maximum. Not knowing when such a shock might occur complicates financial planning, and we consider a model in which Social Security is, essentially, a safe asset within the individual's retirement portfolio.

In our baseline calibration, we consider the worst-known financial shock in modern history, the Great Depression. The shock wipes out 86% of the individual's accumulated assets, matching the lost value in the S&P 500 from peak to trough, and this shock occurs at an unknown time. Both before and after the shock hits, we assume the individual's assets grow at 8% per year to reflect long run real average equity returns in the USA. The individual has full information about the distribution of the rare event risk that they face, and they hedge this risk by optimizing their consumption and saving decisions as the solution to a dynamic stochastic life-cycle saving problem. In this setting (with full information and perfect optimization), we do not find a safety net rationale for Social Security. Instead, we find that Social Security reduces lifetime welfare by 7.64%.

However, perfect hedging and full information are strong assumptions that drive these negative results. Moving beyond our baseline analysis, we attempt to carefully study Social Security's safety net role when there are limits to the individual's optimization and information. We consider three extensions in which optimization or information are imperfect. First, since the dynamic hedging problem is hard to solve, we study the case in which the individual fails to hedge rare event risk but otherwise saves rationally in the sense of consumption smoothing over the life cycle. Second, we consider the case in which the individual behaves according to a sophisticated dynamic hedging problem but does so under a misspecification of risk. Misspecification seems particularly plausible given the difficulty with inferring the timing distribution of an event that is rarely observed. And third, we study the case in which a maximin policy maker bases their evaluation of Social Security on whether it improves individual welfare under the worst-case timing scenario of the shock hitting at the moment of retirement, which creates maximum destruction of wealth. In all of these cases, the welfare effect of Social Security can be far larger than in the baseline model, and in multiple cases, the effect is a large welfare gain rather than a welfare loss.

In sum, if individuals have full information about rare asset shocks and optimally hedge this risk by following the solution to a sophisticated model, then they do not need Social Security as a safety net. They do just fine on their own. However, if individuals fall short of optimal hedging or if there are limits to the information about rare event risk, there are plausible scenarios in which Social Security is quite valuable as a safety net.

Given Social Security's status as the largest government program in operation, it's not surprising that an enormous amount of literature considers its welfare role. One strand of this literature focuses on evaluating Social Security's effectiveness as a solution to the problem of inadequate private saving for retirement.² Another strand focuses on Social Security's role as a hedge against longevity risk.³ And another strand considers its role in hedging lifetime earnings risk or reducing inequality.⁴ Still other risks such as disability risk and survivor risk to the dependent spouse and children also have been thoroughly examined.⁵ Interestingly, however, less work has been done to understand Social Security's role in insuring individuals against rare event asset shocks, despite the fact that Social Security was created in response to the worst financial shock in modern American

history. This paper contributes to the open question of how well Social Security protects people against such risks.

A few studies consider the role of Social Security when private asset returns are risky. For example, Feldstein and Rangelova (1998) assess the size of Social Security benefits relative to what could be financed privately with a portfolio of stocks and bonds whose annual returns are uncertain. Like our baseline model, they do not find much scope for Social Security as a safety net against risky financial markets. We add to their work by considering decision-making under uncertainty about the timing of rare event shocks like the Great Depression.

Peterman and Sommer (2019a) find that Social Security improved the welfare of most people who lived through the Great Depression, thereby explaining the program's creation if not explaining its persistence. Interestingly, the welfare gains in their model are not because Social Security provided a safety net during a period of crises—indeed, the welfare gains would be even larger in their model had the program been implemented during a period of economic tranquility. Instead, the welfare gains to the initial generation are because they paid very little in taxes relative to the large, windfall benefits that they received. Also, the Great Depression shock takes agents in their setting by surprise, making their analysis a study of *ex post* welfare. We ask a different question. We consider *ex ante* welfare analysis. The main feature of our analysis is decision-making in the face of uncertainty about the timing of a large shock to asset values.⁶

2. A model of optimal hedging under full information

2.1. Life-cycle consumption with rare event risk

Time is continuous and indexed by t . An individual is born at $t = 0$, retires at date t_R , and dies no later than $t = T$. The probability of surviving to age t is $\Psi(t)$. The individual receives income $y(t)$ over the life cycle which consists of after-tax labor income before retirement and Social Security benefits after retirement. Consumption is denoted $c(t)$ and saving is denoted $k(t)$. The interest rate is r ; this is the rate of return at all times other than when the shock hits. The individual discounts future utility at the rate ρ .

The instantaneous utility function is the CRRA function:

$$u(c) = \frac{c(t)^{1-\sigma}}{1-\sigma}. \quad (1)$$

The household faces a risk that they will experience a negative wealth shock at some point during their working life that makes the share S of their assets disappear. The shock date t_1 is a random variable with probability density $\phi(t_1)$ for $t_1 \in [0, \infty]$. The one-time negative wealth shock reduces the household's private savings by fraction S . That is, the household only has $(1 - S)$ of their assets remaining after the shock. The household knows the size of S , but does not know when S will occur.

The rare event financial shock that we model has an asymmetric effect on positive and negative asset holdings. That is, we allow individuals to hold negative assets (debt), and in that case there is no change in the magnitude of their debt when a financial shock hits. The shock only reduces positive asset valuations.⁷ Therefore, notationally, we denote the state dependency of the shock

$$S(k(t_1)) = \mathbf{1}\{k(t_1) > 0\}S, \quad (2)$$

where $\mathbf{1}\{k(t_1) > 0\}$ is an indicator function that is equal to one when assets are positive and zero when assets are zero or negative.

The household optimization problem is solved recursively, using the timing uncertainty methods from the stochastic control literature as in Cottle Hunt and Caliendo (2021), Caliendo et al. (2020), Caliendo et al. (2019), and Stokey (2016).

2.1.1. Postshock problem

The negative wealth shock S occurs at time $t_1 \leq \infty$. If the shock occurs before time T , the household solves the following deterministic, fixed-end point problem from the vantage point of the shock date:

$$\max_{c(t)} \int_{t_1}^T e^{-\rho t} \Psi(t) u(c(t)) dt, \tag{3}$$

subject to

$$\dot{K}(t) = y(t) - c(t) + rK(t) \quad \text{for } t \in [t_1, T] \tag{4}$$

$$K(t_1) = k(t_1)(1 - S(k(t_1))) \tag{5}$$

$$K(T) = 0 \tag{6}$$

$$k(t_1) \text{ given, } S(k(t_1)) \text{ given, } t_1 \text{ given.} \tag{7}$$

Here, $K(t)$ is the total financial wealth of the household after the shock takes place. $K(t_1)$ is the share of private savings $k(t_1)$ that remain immediately after the shock $S(k(t_1))$ occurs.

The solution to this problem is the consumption path $c_2^*(t|t_1, k(t_1), S(k(t_1)))$ which is the optimal path of consumption after the wealth shock has occurred. The subscript 2 indicates this is the path of post-shock consumption for $t \in [t_1, T]$,

$$c_2^*(t|t_1, k(t_1), S(k(t_1))) = \frac{k(t_1)(1 - S(k(t_1))) + \int_{t_1}^T e^{-r(v-t_1)} y(v) dv}{\int_{t_1}^T e^{-r(v-t_1) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}. \tag{8}$$

Note changing variable notation for convenience, we replace t with z and t_1 with t . Hence, consumption on $z \in [t, T]$ for shock date t is

$$c_2^*(z|t, k(t), S(k(t))) = \frac{k(t)(1 - S(k(t))) + \int_t^T e^{-r(v-t)} y(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)z/\sigma} \Psi(z)^{1/\sigma}. \tag{9}$$

2.1.2. Preshock problem

Next we solve the dynamic stochastic problem of a household that hedges the timing risk of shock $S(k(t))$, from the perspective of time 0.

The household's expected utility is

$$\begin{aligned} \mathbb{E}(U) &= \int_0^T \phi(t) \left(\int_0^t e^{-\rho z} \Psi(z) u(c(z)) dz + \int_t^T e^{-\rho z} \Psi(z) u(c_2^*(z|t, k(t), S(k(t)))) dz \right) dt \\ &+ \int_T^\infty \phi(t) \left(\int_0^T e^{-\rho z} \Psi(z) u(c(z)) dz \right) dt, \end{aligned} \tag{10}$$

which can be rewritten more compactly using the properties of iterated integrals

$$\mathbb{E}(U) = \int_0^T \left[\left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) u(c(t)) + \phi(t) U_2(t, k(t), S(k(t))) \right] dt \tag{11}$$

where

$$U_2(t, k(t), S(k(t))) = \int_t^T e^{-\rho z} \Psi(z) u(c_2^*(z|t, k(t), S(k(t)))) dz \tag{12}$$

and $c_2^*(z|t, k(t), S(k(t)))$ is given by (9).

Thus, the preshock household’s stochastic optimization problem can be written neatly as a Pontryagin problem

$$\max_{c(t)_{t \in [0, T]}} \int_0^T \left[\left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) u(c(t)) + \phi(t) U_2(t, k(t), S(k(t))) \right] dt \tag{13}$$

subject to

$$\dot{k}(t) = rk(t) + y(t) - c(t) \quad \text{for } t \in [0, T] \tag{14}$$

$$k(0) = 0, k(T) = 0. \tag{15}$$

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

$$\begin{aligned} \mathcal{H} = & \left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) u(c(t)) \\ & + \phi(t) U_2(t, k(t), S(k(t))) + \lambda(t) (rk(t) + y(t) - c(t)). \end{aligned} \tag{16}$$

The necessary conditions are

$$\frac{\partial \mathcal{H}}{\partial c(t)} = \left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) c(t)^{-\sigma} - \lambda(t) = 0 \tag{17}$$

$$\dot{\lambda}(t) = - \frac{\partial \mathcal{H}}{\partial k(t)} = -\phi(t) \frac{\partial U_2(t, k(t), S)}{\partial k(t)} - \lambda(t)r. \tag{18}$$

Note, substituting $c_2^*(z|t, k(t), S(k(t)))$ into $U_2(t, k(t), S(k(t)))$ the partial derivative

$$\frac{\partial U_2(t, k(t), S(k(t)))}{\partial k(t)} = \left(\frac{k(t)(1 - S(k(t))) + \int_t^T e^{-r(v-t)} y(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right)^{-\sigma} (1 - S(k(t))) e^{-rt}. \tag{19}$$

Insert (19) and (17) into (18)

$$\begin{aligned} \dot{\lambda}(t) = & -\phi(t) \left(\frac{k(t)(1 - S(k(t))) + \int_t^T e^{-r(v-t)} y(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right)^{-\sigma} (1 - S(k(t))) e^{-rt} \\ & - r \left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) c(t)^{-\sigma}. \end{aligned} \tag{20}$$

Differentiate (17) with respect to t

$$\begin{aligned} 0 = & -\phi(t) e^{-\rho t} \Psi(t) c(t)^{-\sigma} + \left(\int_t^\infty \phi(t_1) dt_1 \right) (\dot{\Psi}(t) e^{-\rho t} - \rho \Psi(t) e^{-\rho t}) c(t)^{-\sigma} \\ & - \left(\int_t^\infty \phi(t_1) dt_1 \right) \sigma e^{-\rho t} \Psi(t) c(t)^{-\sigma-1} \dot{c}(t) - \dot{\lambda}(t). \end{aligned} \tag{21}$$

Combine (20) and (21) to obtain the Euler equation in closed form

$$\begin{aligned} \dot{c}(t) = & \left(\frac{c(t)^{\sigma+1}}{\Psi(t)} \left(\frac{k(t)(1 - S(k(t))) + \int_t^T e^{-r(v-t)} y(v) dv}{\int_t^T e^{-r(v-t) + (r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} \right)^{-\sigma} (1 - S(k(t))) e^{(\rho-r)t} - c(t) \right) \\ & \times \left[\frac{\sigma}{\phi(t)} \int_t^\infty \phi(t_1) dt_1 \right]^{-1} + \left(\dot{\Psi}(t) \frac{1}{\Psi(t)} + r - \rho \right) \frac{c(t)}{\sigma}. \end{aligned} \tag{22}$$

The optimal path of consumption and assets prior to the shock $(c_1^*(t), k_1^*(t))_{t \in [0, T]}$ obey the Euler equation and the constraints.

2.2. Computational method

Although our model is solved analytically for both the preshock and postshock timepaths of consumption and savings, we must employ a recursive computational method to identify the initial consumption level. We utilize our analytical solutions to identify initial consumption as follows:

Step 1. Compute $c_2^*(z|t, k(t), S(k(t)))$ for given shock date t , given asset holdings at that date $k(t)$, and given shock size $S(k(t))$.

Step 2. Guess $c(0)$.

Step 3. Given the guessed value of $c(0)$ together with $k(0) = 0$, use the Euler equation $\dot{c}(t)$ and law of motion for capital $\dot{k}(t)$, which themselves depend on the post-shock consumption path $c_2^*(z|t, k(t), S(k(t)))$, to shoot forward in time to construct the pre-shock timepaths $c(t)$ and $k(t)$.

Step 4. Check the terminal value of pre-shock capital to see if $k(T) = 0$. If yes, then we have found our solution pre-shock consumption and saving paths $c_1^*(t)$ and $k_1^*(t)$. If not, go back to step 2 and repeat.

2.3. Calibration

Our baseline calibration strategy is to maximize the potential safety net feature of Social Security by assuming that a modern individual faces risk of a financial shock that is proportional to the stock market crash of the Great Depression. This means that we will calibrate things like income, career length, longevity, and Social Security to modern specifications, and we will assume the individual faces financial risk like that experienced by past generations who endured the Great Depression.⁸

We assume the individual enters the model at age 18 ($t = 0$), retires at age 67 ($t_R = 49$), and lives to a maximum age of 100 ($T = 82$). We calibrate the survival function and wage income following Cottle Hunt and Caliendo (2021). We calibrate the survival function as

$$\Psi(t) = 1 - (t/T)^{2.58} \tag{23}$$

to ensure that the ratio of workers to retirees is 3.3, which is approximately the average value in the USA during the period 2000–2010.⁹ 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. We parameterize $y(t)$ to represent after-tax income before retirement and Social Security benefits after retirement. We use a 5th order polynomial to fit life-cycle wages to the data in Gourinchas and Parker (2002)

$$w(t) = \alpha(1 + 0.0244t - 0.0001t^2 + 4(10)^{-5}t^3 - 2(10)^{-6}t^4 + 2(10)^{-8}t^5) \tag{24}$$

where α is chosen such that the average wage is equal to one.

We set the Social Security tax rate to $\tau = 0.106$, which corresponds to the employer and employee Old Age Survivors portion of Social Security. We assume the Social Security system has a balanced PAYGO budget, hence

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

Thus, $y(t)$ is given by

$$y(t) = \begin{cases} (1 - \tau)w(t), & t \leq t_R \\ b, & t > t_R. \end{cases} \tag{25}$$

Table 1. Summary of baseline calibration of parameters

Parameters	
<i>Social Security</i>	
$\tau = 10.6\%$	Social Security payroll tax
$b = 0.348$	Balances the Social Security budget
<i>Rare event shock</i>	
$\phi(t)$	Equation (26) in text, uniform distribution
$S = 0.86$	S&P loss during the Great Depression
<i>Demographics and misc</i>	
$T = 82$	Normalized maximum lifespan (age 18–100)
$t_R = 49$	Length of career (age 18–67)
$\Psi(t) = 1 - (t/T)^{2.58}$	Survival function yielding 3.3 workers:retirees
$w(t)$	Polynomial fitted to Gourinchas and Parker (2002)
$\sigma = 3$	Midpoint CRRA value from literature
$r = 0.08$	Interest rate set to real return on equities
$\rho = 0$	Discount rate

We set the interest rate $r = 0.08$ which corresponds to the average long-term real return on equities. We set $\sigma = 3$, which is a mid-point of the standard values used in the literature. Finally, we set the discount rate $\rho = 0$.

As our baseline parameterization for the rare event shock, we assume the risk of experiencing the shock is uniform and that the individual cannot escape the shock—it is sure to hit during their life cycle. Hence

$$\phi(t) = \begin{cases} 1/T, & t \leq T \\ 0, & t > T. \end{cases} \quad (26)$$

As a robustness check later in the paper, we explore a distribution that has positive mass beyond the maximum age T . We calibrate the size of the shock S to correspond to the Great Depression. The S&P lost 86% of its value during the Great Depression, and so we set $S = 0.86$.¹⁰ Our baseline parameters are summarized in Table 1.

2.4. Welfare

Our goal is to assess Social Security's role as a safety net against rare event financial shocks to wealth. Our stylized model leaves out a number of Social Security's features like disability insurance, spousal benefits, survivor insurance, and a benefit-earning rule that redistributes wealth from the rich to the poor. While these are important features, we view the simplicity of our baseline analysis as an advantage because it allows us to focus cleanly on the safety net rationale, which is our purpose. As a result, we are not in a position to make statements about Social Security's overall welfare effects. Instead, we are drawing conclusions about a single mechanism in particular.

We measure the welfare effect of Social Security as the percentage of lifetime consumption the individual would be willing to give up to live in a world with Social Security relative to a Laissez Faire world with no Social Security.

Notationally, let $\mathbb{E}(U_{SS}^*)$ stand for expected utility in a world with Social Security and following optimal consumption rules under uncertainty. Likewise let $\mathbb{E}(U_{LF}^*)$ stand for expected utility in a Laissez Faire world without Social Security and following optimal consumption rules in that

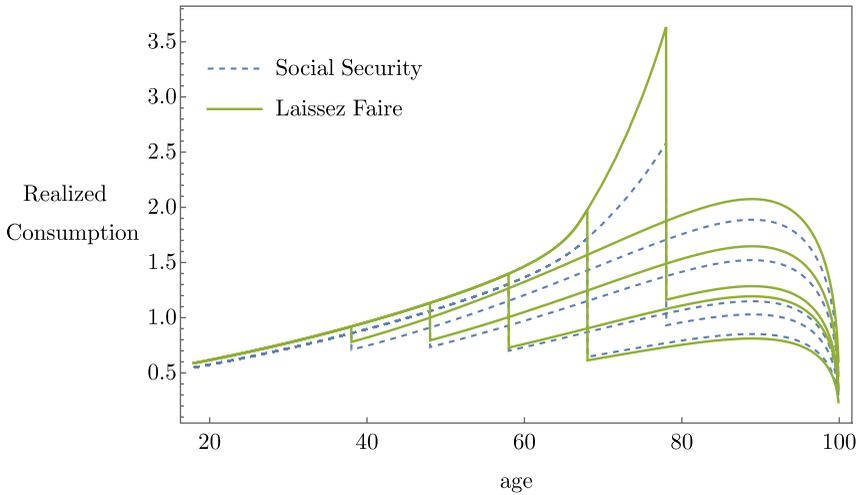


Figure 1. Realized consumption given $S = 0.86$, with and without Social Security.

environment. With CRRA utility, the fraction of lifetime consumption the individual would give up to participate in Social Security is

$$\Delta = 1 - \left(\frac{\mathbb{E}(U_{LF}^*)}{\mathbb{E}(U_{SS}^*)} \right)^{\frac{1}{1-\sigma}} \tag{27}$$

We calibrate the size of the shock to correspond to the Great Depression: $S = 0.86$. Even in this extreme case, we find the welfare effect of Social Security to be negative, $\Delta = -7.64\%$. That is, an individual would need to receive an additional 7.64% of lifetime consumption in order to compensate for participating in Social Security.

What is the intuition for this result? The basic answer is that US equities are still a very good deal, even when individuals face catastrophic, rare event risk. An 8% return compounded annually, together with the risk of a one-time shock equivalent to the Great Depression (86% instantaneous loss in wealth), still compares favorably to the implicit rate of return in a Pay-As-You-Go Social Security program. And this is true even after accounting for the life annuity feature of Social Security.

The intuition is also visible in the possible timepaths of consumption over the life cycle with and without Social Security. Figure 1 plots a few possible paths of realized life cycle consumption with and without Social Security, given different possible dates for the shock (specifically ages 38, 48, 58, 68 and 78). The solid green line plots the Laissez Faire consumption paths and the dashed blue lines are with Social Security. Consumption is generally higher without Social Security.

This negative welfare result is important because in some ways we have stacked our analysis in favor of Social Security in our choice of modeling assumptions. We assume the financial shock is guaranteed to happen; we assume the shock will be the biggest possible shock in modern history; and, we assume the individual’s only choice for private saving is a risky asset. By assuming the individual faces an inescapable, large shock with no access to a safe asset, the individual is left completely exposed to risk and Social Security’s safety net role should be maximized. And yet we do not find a safety net role for Social Security.

However, we will show in Section 3 that our baseline assumptions about optimal hedging under full information are driving the negative welfare result. Relaxing these assumptions will create space for large welfare gains.

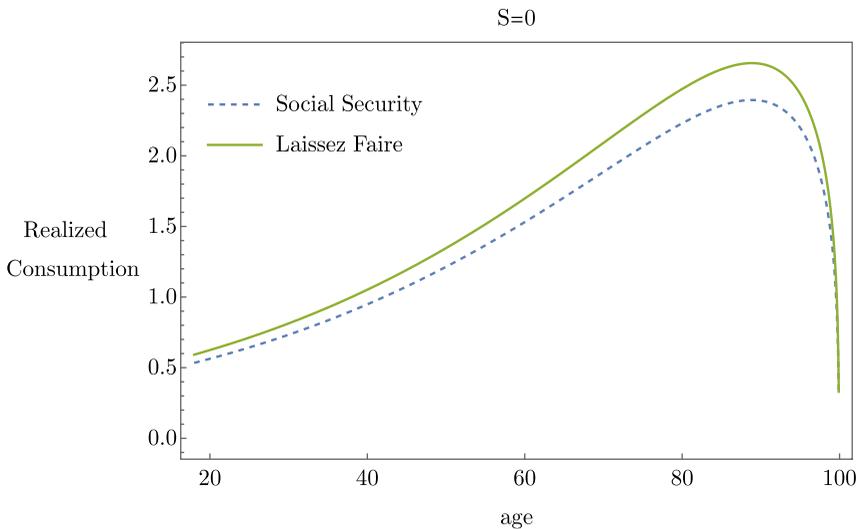


Figure 2. Realized consumption given $S = 0$, with and without Social Security.

2.5. Decomposition

We decompose the welfare effect of Social Security into the portion that comes from the provision of longevity insurance and the portion that comes from the provision of a financial safety net. To do this, we turn off financial risk by setting the asset shock to $S = 0$ and we find that $\Delta = -10.91\%$. Consumption with and without Social Security when $S = 0$ is depicted in Figure 2. Since the shock is set to zero, there is only one realized consumption path and consumption is everywhere higher without Social Security. This welfare effect is worse (more negative) than the combined welfare effect of insurance against both longevity and asset risk which was -7.64% in our baseline model. This suggests that Social Security is less harmful in the presence of rare event risk.

As an alternative decomposition, we turn off both the asset risk and the asset reward by setting $S = 0$ and $r = 0$. We find that $\Delta = 11.57\%$ in this case. This large positive welfare effect is due to the well-known annuitization feature of Social Security. An individual would rather pay Social Security taxes and receive a Social Security annuity during retirement rather than invest in a riskless asset with zero return. Consumption with and without Social Security when $S = 0$ and $r = 0$ is depicted in Figure 3. In this case, consumption is everywhere higher with Social Security.

2.6. Robustness checks

Some of the parameters that we use in our baseline analysis are unobservable and could have direct bearing on our welfare calculations. In this section, we briefly discuss these issues before addressing more fundamental extensions.

Our baseline value of 3 for the coefficient of relative risk aversion (σ) falls within a standard range of values used in the life-cycle consumption literature, though values as low as 0.5 and as high as 5 are also used. The welfare effect of Social Security is -8.64% and -7.08% with these values, compared to our baseline value -7.64% . Hence, while different levels of risk aversion change the specific welfare effect of Social Security, our overall conclusion about its role as a safety net is unchanged.

Similarly, our use of a 0% discount rate is not uncommon, but it is also common to see values as high as 5% used in the life-cycle literature. The welfare effect of Social Security is -7.89% under this assumption. Again, no material change to our conclusions thus far.

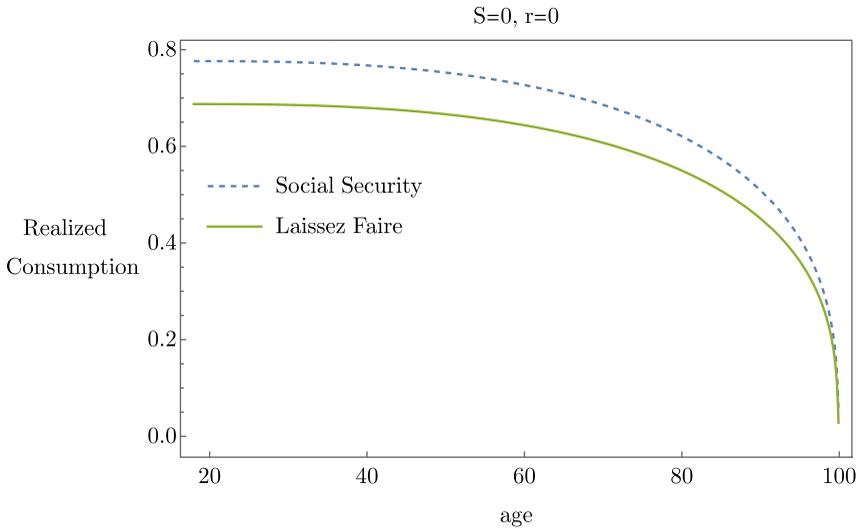


Figure 3. Realized consumption given $S = 0$ and $r = 0$, with and without Social Security.

We assumed the financial shock was distributed uniformly over the life cycle of the individual. This is a reasonable starting assumption and it helps to stack the analysis in favor of Social Security: if the individual cannot escape experiencing the shock at some point (because the upper support of the shock coincides with the maximum age of the individual), then Social Security would be more attractive as a safety net than if there was a chance the individual may never experience the negative shock. However, our uniform assumption has a subtle side effect: if the shock has not yet happened, then as the individual ages they become increasingly sure that it will happen to them. Indeed, as the individual approaches the maximum model age, the probability that the shock strikes will spike to 100% conditional on it not happening yet. This is a mathematical definition of any truncated distribution, and yet there is no reason that a rare financial shock would increase in relative likelihood as an individual ages.

As an alternative, we consider the exponential density

$$\phi(t) = \gamma e^{-\gamma t}, \text{ for } t \in [0, \infty] \text{ and } \gamma > 0. \tag{28}$$

This density is memoryless: if the shock has not occurred by date x , the likelihood of the shock occurring between x and $x + \delta$ does not depend on x . As an example, we calibrate γ by assuming there is a 50% chance that a financial crash like the Great Depression does not happen within the individual's maximum lifespan; hence, $\gamma = -\ln(0.5)/82$. The welfare effect of Social Security is -7.42% under this assumption. Note that while the need for a safety net is intuitively weakened if there is a chance that such a net is never needed, the exponential distribution considered here also rearranges the shock probabilities over the lifespan of the individual, with the most weight during the earliest years of life. This secondary effect virtually washes out the first effect and leaves the overall welfare effect of Social Security quite similar to our baseline estimate.

Lastly, we assumed a real interest rate of $r = 0.08$ which corresponds to the average annual real return on equities in the USA over the last 50–100 years. The welfare effect of Social Security is decreasing in the interest rate. That is, Social Security reduces welfare less (or increases welfare more) if the interest rate is lower. Holding the other parameters fixed at their baseline values, the welfare effect of Social Security is -2.47% for an interest rate of $r = 0.06$, zero for $r = 0.054$, 1.82% for $r = 0.05$, and 7.9% for $r = 0.04$.

3. Departures from optimal hedging under full information

There may be limits on the degree to which people are able to hedge rare event risk in actuality. In the analysis above, the individual has perfect information about the distribution of timing risk, and then is able to solve a complex dynamic stochastic hedging problem to protect themselves against this risk. To study the importance of these strong assumptions, in this section, we consider three possibilities.

First, we assume the individual does not hedge the risk at all but otherwise behaves optimally in saving for retirement in the face of survival uncertainty. The policy maker is aware of the true distribution of rare event risk and evaluates the performance of Social Security through the lens of whether it improves the ex ante expected utility of the individual. We find that the large welfare losses from Social Security in the baseline analysis with optimal hedging now become large welfare gains when the individual fails to hedge.

Second, we consider the case in which the true distribution of risk is unknown to the individual. Consequently, the individual solves the dynamic stochastic hedging problem featured in the baseline analysis above, but they do so under a misspecification of rare event risk. We find that the direction of the misspecification is important: relative to the baseline model with correct risk specification, Social Security is less negative for individuals who underestimate the risk they face. However, Social Security is even more harmful for individuals who overestimate the risk they face.

Third, we study the case in which the policy maker evaluates Social Security based on its performance under the worst possible shock timing, which is the date of retirement when the individual’s wealth is at a maximum. We find this dampens the welfare losses associated with Social Security in our baseline parameterization and leads to welfare gains if the risk is distributed exponentially. In other words, simply switching the policy maker’s welfare criteria from expected utility to maximin can be enough to justify Social Security as a safety net against rare event risk.

3.1. Failing to hedge

Up to this point in the paper, we have assumed the individual has full information about the risks that they face and that they behave optimally in the face of this risk. In this case, we found above that Social Security is unable to confer any welfare gains (as a safety net against rare event risk), at least at the parameterizations that we have highlighted. Our baseline model highlights that an individual who saves optimally in the face of this risk does not need a Social Security program to help manage this risk.

In this section, we consider an alternative model. In actuality, for a variety of reasons, individuals may not solve a sophisticated dynamic hedging problem. For instance, individuals may lack information on the nature of rare event risk, and without any information on the distribution of the risk, they may choose to ignore it and instead base their financial plans on expected (average) returns. Similarly, in financial planning, people often say things like, “when saving over long time horizons, equities pay better than bonds on average” and then use this fact as a justification for basing financial projections on a rate of return to saving that equals the average long-run return to equities, without explicitly modeling rare event risk. Moreover, the complexity of optimal hedging may limit its attractiveness as a saving plan, with individuals instead opting for simple rules. Whatever the cause for failing to hedge, in this section we study the welfare effect of Social Security when rare event risk exists but is not incorporated into the household’s saving plan. Unlike the previous section which struggles to find a rationale for Social Security as a safety net, in this section, the welfare gains of Social Security are large.

We consider a household that plans consumption and saving to maximize lifetime utility, ignoring the possibility of rare event risk. The household solves the following deterministic problem:

$$\max_{c(t)} \int_0^T e^{-\rho t} \Psi(t) u(c(t)) dt, \tag{29}$$

subject to

$$\dot{k}(t) = y(t) - c(t) + rk(t) \quad \text{for } t \in [0, T] \tag{30}$$

$$k(0) = 0, k(T) = 0. \tag{31}$$

The solution to this problem is the *planned* consumption path

$$c_1^{nh}(t) = \frac{\int_0^T e^{-rv} y(v) dv}{\int_0^T e^{-rv+(r-\rho)v/\sigma} \Psi(v)^{1/\sigma} dv} e^{(r-\rho)t/\sigma} \Psi(t)^{1/\sigma}. \tag{32}$$

The superscript nh on $c_1^{nh}(t)$ indicates this is planned consumption with no hedging. The subscript on $c_1^{nh}(t)$ indicates that this is the consumption path the household follows until the rare event shock occurs (if it occurs at all).

If the rare event shock occurs during the household’s lifetime (at date t), we assume the household reoptimizes and follows the postshock consumption path $c_2^*(z|t, k(t), S(k(t)))$ for $z \in [t, T]$, as given in equation (9). The consumption path $c_2^*(z|t, k(t), S(k(t)))$ is optimal conditional on shock date t , assets $k(t)$, and shock $S(k(t))$. The assets the household owns at the time of the shock are determined by the law of motion equation (30) and the planned consumption path $c_1^{nh}(t)$.

Since this household does not hedge the rare event risk, they are surprised if the shock occurs in their lifetime. The realized life-cycle consumption path of the household given shock date t is $c_1^{nh}(z)$ prior to the shock for $z < t$, and $c_2^*(z|t, k(t), S(k(t)))$ for $t \leq z \leq T$ following the shock.

We measure the welfare effect of Social Security from the perspective of a social planner who knows that the household does not hedge against the rare event risk. The social planner knows that the household will follow consumption path $c_1^{nh}(z)$ prior to the shock and the consumption path $c_2^*(z|t, k(t), S(k(t)))$ after the shock. The planner calculates ex ante expected utility by computing the life-cycle utility for the realized path of consumption the household would experience for every possible shock date and weighing each path by how likely the shock is to occur at that date. Ex ante expected utility is given by

$$\mathbb{E}(U^{nh}) = \int_0^T \left[\left(\int_t^\infty \phi(t_1) dt_1 \right) e^{-\rho t} \Psi(t) u(c_1^{nh}(t)) + \phi(t) U_2(t, k(t), S(k(t))) \right] dt \tag{33}$$

where

$$U_2(t, k(t), S(k(t))) = \int_t^T e^{-\rho z} \Psi(z) u(c_2^*(z|t, k(t), S(k(t)))) dz \tag{34}$$

and $c_2^*(z|t, k(t), S(k(t)))$ is given by (9).

Notationally, let $\mathbb{E}(U_{SS}^{nh})$ stand for expected utility in a world with Social Security and no hedging. Likewise let $\mathbb{E}(U_{LF}^{nh})$ stand for expected utility in a Laissez Faire world without Social Security and no hedging. With CRRA utility, the consumption equivalent variation is

$$\Delta^{nh} = 1 - \left(\frac{\mathbb{E}(U_{LF}^{nh})}{\mathbb{E}(U_{SS}^{nh})} \right)^{\frac{1}{1-\sigma}}. \tag{35}$$

The consumption equivalent variation is the percentage of lifetime consumption the individual who does not hedge rare event risk would be willing to give up to live in a world with Social Security relative to a Laissez Faire world with no Social Security. Note that our welfare comparison takes into consideration the distribution of rare event risk, even though the household itself does not incorporate this information into their consumption and saving plan. As in the previous section, a positive number indicates that Social Security makes the individual better off.

Using our baseline parameters, we find $\Delta^{nh} = 16.39\%$. That is, a household who does not hedge rare event risk is made better off by participating in Social Security by an amount equivalent to about 16% of lifetime consumption. This large positive welfare effect stands in stark contrast to the negative welfare effect we found for households who optimally hedge rare event risk. Social Security is able to confer large welfare gains to households who do not hedge rare event risk by crowding out the risky assets the household owns and offering a safe asset (albeit, an asset with a low rate of return.). Households who do not hedge their rare event risk are made much better off in this scenario.

The positive welfare effect of Social Security is robust across several different parameter combinations, including when the shock distribution is calibrated to be an exponential distribution as in equation (28), in which case $\Delta^{nh} = 6.38\%$.

We also conduct an alternative welfare comparison to see how much better off the household who does not hedge rare event risk would be if they were able to optimally hedge. Specifically, we compute the percentage of lifetime consumption the individual who does not hedge rare event risk would be willing to give up to optimally hedge rare event risk:

$$\Gamma = 1 - \left(\frac{\mathbb{E}(U_{LF}^{nh})}{\mathbb{E}(U_{LF}^*)} \right)^{\frac{1}{1-\sigma}} \tag{36}$$

Here, $\mathbb{E}(U_{LF}^{nh})$ stands for expected utility following the no-hedging consumption path $c_1^{nh}(z)$ prior to the shock and the consumption path $c_2^*(z|t, k(t), S(k(t)))$ following the shock in the Laissez Faire economy. Similarly, $\mathbb{E}(U_{LF}^*)$ indicates expected lifetime utility with optimal hedging.

This comparison allows us to see how much a household is harmed by not hedging rare event risk. Using our baseline parameters, we find the welfare gain of optimal hedging $\Gamma = 27.44\%$. The welfare gain of optimal hedging in Laissez Faire essentially shows the size of the problem of ignoring rare event risk. Participating in Social Security is not as good as optimally hedging, but does a reasonably good job helping households hedge rare event risk (by providing a safe asset). Recall, the welfare gain of Social Security is 16.39% for households that do not hedge rare event risk. Hence, the welfare gain of Social Security is about 60% of the welfare gain of optimally hedging, for households that ignore rare event risk.

Using the exponential shock distribution as in (28), we find the welfare gain of optimal hedging $\Gamma = 14.14\%$ in Laissez Faire. With this parameterization, Social Security is able to generate about 45% of the welfare gain of optimally hedging rare event risk.

3.2. Misspecification of risk

Rare event shocks are, by definition, seldom observed in reality. Therefore, it may be difficult for decision makers to know the distribution of the risk. To infer the distribution, one would need to observe many historical instances. Such inference is possible with high frequency, stationary risk but not with truly rare events.

To further connect our model to reality, in this section, we explore the possibility that the true distribution of timing risk is unknown to the individual decision maker. Instead, the individual makes decisions to hedge risk that are based upon an incorrect distribution of the risk. This exercise tells us how the welfare gains from Social Security are affected by risk misspecification, which seems especially relevant given the difficulty of drawing accurate inferences about the distribution of rare event risk.

We assume the risk of the rare event is distributed exponentially as in equation (28)

$$\phi(t) = \gamma e^{-\gamma t}, \text{ for } t \in [0, \infty] \text{ and } \gamma > 0.$$

Table 2. Welfare effect of Social Security, given different levels of misspecification of risk

Welfare effect of Social Security		
ω	Perceived prob. shock within life	Welfare
4γ	0.94	-7.91%
2γ	0.75	-7.85%
γ	0.5 (true probability)	-7.42%
$1/2\gamma$	0.29	-5.98%
$1/4\gamma$	0.16	-3.51%

Suppose the true distribution has a parameter $\gamma = -\ln(0.5)/82$ which corresponds to a 50% chance of the shock occurring within the individual’s maximum lifespan. However, the individual decision-maker does not know γ and assumes the shock is instead distributed

$$\phi(t) = \omega e^{-\omega t}, \text{ for } t \in [0, \infty], \tag{37}$$

where ω is a multiple of γ . We consider four cases: $\omega = 4\gamma$, $\omega = 2\gamma$, $\omega = 1/2\gamma$, and $\omega = 1/4\gamma$, which correspond to the individual believing the shock will occur within the maximum lifespan with probability 0.94, 0.75, 0.29, and 0.16, respectively.

We solve for the individual’s optimal consumption-saving rule just like in Section 2, but now we assume the individual applies the wrong distribution of risk [equation (37)] and therefore optimizes the wrong expected utility function. Hence, the individual is sophisticated in their optimization behavior but they lack knowledge of the true distribution of risk. Then, in the welfare analysis, we apply the true distribution of risk [equation (28)] in our assessment of the welfare effect of Social Security.

Our results are presented in Table 2. We find that the welfare effect of Social Security is decreasing in the perceived likelihood of the rare event shock occurring. That is, individuals who believe the shock is more likely to occur than the true distribution are harmed more by Social Security than if they knew the risk. Similarly, individuals who do not believe the shock is as likely to occur are less harmed by Social Security. This is because Social Security is more helpful (less negative welfare effect) for households who are not as prepared for a rare event to strike. And for those who overestimate the risk they face—and therefore “over hedge”—adding an additional safety net is problematic.

3.3. Maximin policy making

In this section, we consider a final permutation on the information structure of the problem. We imagine a policy maker who does not have reliable information on the distribution of rare event risk. They know rare event risk exists, but do not have enough information to construct a distribution for policy evaluation. Hence, as an example, we study the case in which the policy maker evaluates the success of the Social Security safety net based upon its ability to insure against the worst case scenario, which, in our model would be the realization of the timing shock at the moment of retirement when the individual’s wealth is at a peak. Such unfortunate timing would create the largest possible loss in wealth to the individual, and we assume for the moment that this is the very scenario that the policy is designed to address. Hence, the policy maker evaluates the success of Social Security based upon its performance under this specific realization of risk.¹¹

The policy maker calculates maximin expected utility as

$$U^{mm} = \int_0^T e^{-\rho t} \Psi(t) u(c^{mm}(t)) dt \tag{38}$$

where c^{mm} indicates the consumption an individual would experience if the shock hits at the date of retirement:

$$c^{mm} = \begin{cases} c_1^*(t) & t < t_R \\ c_2^*(t|t_R, k(t_R), S(k(t_R))) & t \geq t_R. \end{cases} \tag{39}$$

The welfare effect of Social Security from the perspective of the policy maker is the fraction of lifetime consumption the individual would give up to participate in Social Security (given the individual hedged optimally, but the shock happened exactly at the date of retirement),

$$\Delta^{mm} = 1 - \left(\frac{(U_{LF}^{mm})}{(U_{SS}^{mm})} \right)^{\frac{1}{1-\sigma}}. \tag{40}$$

We find the welfare effect of Social Security to be -2.28% using our baseline parameters and the uniform shock distribution. The welfare effect is less negative compared to the baseline, because the policy maker is only considering the worst possible shock date, rather than forming expected utility based on all possible shock dates. Social Security still reduces welfare because the internal rate of return is low compared to the risky asset—even when the shock hits at the worst possible moment.

However, the welfare effect is 3.55% using the exponential shock distribution. Social Security is welfare enhancing in this scenario because the individual optimally hedged risk under the assumption that the shock would only occur within their lifetime with probability 0.5, whereas the policy maker evaluates Social Security assuming the shock hits at the worst possible moment with probability 1.

Recall that the welfare effect of Social Security is -7.42% when the welfare criteria is calculated using expected utility [equation (27)] and the rare event risk is distributed exponentially. However, we have just shown, that if the welfare criteria is instead a maximin criteria—comparing utility of the worst possible realization of the shock rather than taking an expectation over all possible realizations of the shock—Social Security is welfare improving by the equivalent of 3.55% . In both cases, behavior is fully rational and the individual optimally hedges rare event risk. The only difference is which possible lifetime utility paths the policy maker includes in their welfare criteria. If all possible paths are included (and weighted according to the distribution of risk), Social Security is welfare reducing; however, if only the worst path is included, Social Security is welfare improving.

4. Inequality and redistribution

Our analysis thus far has focused on a representative (average) wage earner. In this section, we expand our analysis to consider Social Security’s safety net role in the presence of wage inequality.

We continue to hold the Social Security tax fixed at its statutory rate (10.6%), but now we consider different replacement rates. In actuality, the government computes the average of an individual’s highest 35 years of earnings and then calculates Social Security benefits as a piecewise linear function of average earnings, with kinks or bend points at 0.2, 1.24, and 2.47 times the economy-wide mean wage. Social Security replaces 90% of an individual’s average wage up to 0.2 times the economy-wide mean wage, then replaces 32% of an individual’s average wage between 0.2 and 1.24 times the economy-wide mean wage, and finally replaces 15% of an individual’s average wage between 1.24 and 2.47 times the economy-wide wage. Wages beyond 2.47 times the economy-wide mean are not taxed and do not affect benefits. Hence, the lowest wage earners experience the highest replacement rate (90%), and the replacement rate is progressively less as wages increase beyond the first bend point.

Table 3. Welfare effect of Social Security for the median earner of each income quintile, with optimal hedging and no hedging, for two different shock distributions

Welfare effect of Social Security					
Income quintile	Replacement rate	Optimal hedging		No hedging	
		Uniform	Exponential	Uniform	Exponential
Lowest	0.900	-0.96%	-0.83%	26.86%	13.33%
2nd	0.577	-4.76%	-4.03%	22.72%	10.48%
3rd	0.463	-6.04%	-5.51%	20.38%	8.95%
4th	0.402	-6.74%	-6.35%	18.81%	7.93%
Highest	0.303	-7.93%	-7.76%	15.46%	5.79%

We consider five different Social Security replacement rates that correspond to income quintiles. Our welfare results are presented in Table 3. Specifically, we assume that wages are distributed and Social Security benefits are calculated as in Cottle Hunt and Caliendo (2022a) and calculate the replacement rate for the median earner for each earnings quintile. Then, we calculate the welfare effect of Social Security with our baseline uniform shock distribution (26) and with the alternative exponential shock distribution (28).

We calculate the welfare effect of Social Security assuming optimal hedging and find that Social Security makes individuals in all five earning quintiles worse off. Higher earners are harmed the most because Social Security benefits are progressive, and thus, high earners receive the lowest implicit rate of return.

Consider an individual with wages at or below the first bend point and hence a 90% replacement rate. Holding all other parameters fixed at the baseline calibration, the welfare effect of Social Security still is slightly negative. Intuitively, these individual's experience a much higher internal rate of return on Social Security (4.4%), but this internal rate still is lower than most of the possible lifetime ex post rates of return on private saving. An individual who draws the very worst possible timing shock (at the moment of retirement) and losses 86% of their total accumulated wealth, effectively earns an ex post rate of return on their private saving of 3.75%. All other shock dates confer a higher rate of return on private saving than this. There is a specific window of time around the date of retirement for which the shock would need to hit in order for Social Security to outperform private saving, and this window is narrow enough that even a risk averse, low-income individual would prefer the risk-reward tradeoff found in private assets over participating in Social Security. Hence, even when making Social Security as attractive as possible by considering the highest possible replacement rate, while also assuming away any private options for insuring longevity risk, the performance of Social Security does not compare favorably to that of (risky) private assets.

We also calculate the welfare effect of Social Security assuming the individual does not hedge the rare event risk as in Section 3, and find that Social Security makes all groups better off with the largest welfare gains accruing to the lowest earners. These results are consistent with our earlier analysis: Social Security improves the ex ante expected utility of households who are unable to optimally hedge rare event risk by providing a safe asset.

In our model, Social Security provides direct insurance against longevity risk by paying retirement benefits as an annuity. In an ex ante sense, Social Security also provides insurance against having low lifetime earnings by using a progressive benefit earning rule. The welfare effects of Social Security are largest for individual's with the lowest lifetime earnings.¹² Although our model does not include idiosyncratic wage shocks (like unemployment or periods of low earnings), Social Security would provide insurance against those shocks through the progressive benefit earning rule as well.¹³

5. Discussion and limitations

Our analysis focuses on differences between households who optimally hedge rare event risk and those who fail to hedge. We find the ability of Social Security to function as a safety net (and improve ex ante expected utility) depends critically on whether or not the household hedges rare event risk. Our analysis includes two simplifications that likely have level effects on our welfare calculations. First, we assume the rare event shock happens at most once during a household's life cycle. Second, we assume fixed factor prices.

5.1. One shock

The household in our model faces a risk of a rare event shock happening in their lifetime once (if the shock is distributed uniformly) or possibly not at all (if the shock is distributed exponentially). Of course, in reality, a household could experience a rare event shock that depresses the value of their assets multiple times. Modeling persistent or recurring shocks could amplify the insurance benefits of Social Security. In that sense, the welfare effects we find could be viewed as a lower bound.

From a conceptual standpoint, our single shock assumption has the disadvantage that the individual essentially enters a risk-free world after the shock hits, whereas in reality the individual still faces the risk that another rare shock might strike. This disadvantage is partially mitigated by two considerations. First, we focus on ex ante welfare and the individual faces the possibility of the shock hitting at any moment (or never) during the life cycle, which in turn creates a setting in which even a single shock has a powerful effect on welfare and decision making. Second, we have to allow the individual to get hit with multiple rare event shocks of the size of our calibrated shock (86% destruction of wealth with a 25-year recovery time), then we would be subjecting the individual to shocks that exceed the US experience over the last 100 years. While there have been multiple shocks over the last century (dot.com, Great Recession, etc), none of these shocks come close to the Great Depression in terms of relative loss in equity prices. In other words, while the individual only gets hit with one shock, it is a massive one.

Finally, our paper is an attempt to consider asset risk from a novel angle. Typically, returns on risky assets are modeled as random disturbances from a stationary distribution, without explicitly considering timing risk over very large rare shocks. We reverse this tradition by focusing on the latter and abstracting from the former. Of course, both risks are realistic but our focus on rare event risk is particularly relevant to our research question about safety nets.

5.2. Partial equilibrium

Second, we assumed exogenous factor prices. If we were to model a general equilibrium overlapping generations economy, the introduction of Social Security would crowd out private savings which would reduce the wage and increase the interest rate. These changes to factor prices would dampen the welfare effects of Social Security (if the economy was dynamically efficient). This channel has been well documented in Slavov et al. (2019), Blau (2016), Conesa and Garriga (2008), and Kitao (2014) among many others.

A full analysis of the welfare effect of Social Security would need a general equilibrium treatment of factor prices and bequest income, as Social Security affects the level of aggregate capital accumulation as well as the level of wealth that is transmitted across generations. In this paper, we do not attempt a full analysis of Social Security but instead attempt to clarify a single aspect of Social Security, and, in particular, to show the role of full optimization and full information. By showing that Social Security's welfare effect can swing dramatically, from very negative to very positive as we move from full information/optimization to imperfections along these margins, we have identified some of the most significant modeling features that are of first-order importance to the research question.

Because we are specifically interested in the role of strong neoclassical assumptions like optimal hedging and full information in shaping our welfare conclusions, our partial equilibrium setting allows us to conveniently experiment with a number of different assumptions without needing to recalibrate to the larger macroeconomy and without needing to specify how the degree of optimizing behavior and information processing are distributed across households throughout the economy.

6. Conclusion

Social Security is the largest government program in existence and it affects Americans of all ages. Part of the initial rationale for creating Social Security in 1935 was to ensure that Americans are protected from financial shocks like the catastrophic stock market crash during the Great Depression. That particular shock was the worst shock to financial assets in the last 100 years, though more recent shocks such as the Great Recession and the COVID-19 pandemic have brought intense financial losses, especially to certain sectors of the economy, and have been particularly problematic for households who are close to retirement. While a large literature focuses on the welfare effects of Social Security, its role as a safety net against rare event risk to asset values has received less attention and is the purpose of our paper.

We build a model in which a rare but very large negative shock to financial wealth hits the household at an unknown time. The shock is especially painful if it occurs near retirement when the individual's accumulated wealth is at a maximum. But since the timing of the shock is uncertain, individual's must solve a dynamic stochastic life-cycle consumption and saving problem, and we utilize a recursive technique from the continuous-time control literature to obtain a closed-form solution. We calibrate the size of the shock to the Great Depression, which brought an 86% decline in equity values from peak to trough.

We consider a number of assumptions about the individual's ability to hedge rare event risk and about the policy makers information. We find the answer to whether or not Social Security provides a safety net against rare event risk depends critically on these information and optimization assumption.

We show that if a household has full information on the distribution of rare event risk and solves a dynamic stochastic problem to hedge this risk by building up optimal precautionary savings balances, then Social Security is unnecessary as a safety net. Alternatively, if the household does not account for rare event risk in its financial planning and instead leaves itself exposed to rare shocks without a proper buffer, then Social Security can provide very large welfare gains as a safety net.

Notes

1 We focus on negative rare events, even though in reality positive rare events also exist.

2 See Feldstein (1985), Docquier (2002), İmrohoroğlu et al. (2003), Gul and Pesendorfer (2004), Cremer et al. (2007, 2008, 2009), Fehr et al. (2008), Findley and Caliendo (2008), Kumru and Thanopoulos (2008), Pestieau and Posse (2008), Buccioli (2011), Cremer and Pestieau (2011), Caliendo (2011), Andersen and Bhattacharya (2011), Caliendo and Findley (2013), and Guo and Caliendo (2014), among many others.

3 See Diamond (1977), Diamond (2004), Feldstein (2005), and Chetty and Finkelstein (2013), among many others.

4 See Diamond (1977), Huggett and Ventura (1999), Nishiyama and Smetters (2007), İmrohoroğlu and Kitao (2012), Bagchi (2016), Huggett and Parra (2010), Cottle Hunt and Caliendo (2022b), and others.

5 See Cottle Hunt and Caliendo (2021), Li (2018), Kaygusuz (2015), and others.

6 In Peterman and Sommer (2019a), the Great Depression destroys 24% of the economy's capital stock, which matches estimates of the loss in value of the nation's aggregate fixed assets. In contrast, in our paper, the agent is planning/bracing for a valuation shock that destroys 86% of their privately accumulated financial wealth, which matches the loss to a specific asset during the Great Depression (equities). We do not consider labor supply or shocks to labor productivity. In related work, Peterman and Sommer (2019b) find that Social Security helped to mitigate the welfare loss associated with the Great

Recession, which caused sudden shocks to wealth and employment. Agents in their model are not aware that a shock of this magnitude can happen and instead make consumption and saving plans based on a fixed rate of return on asset holdings. Hence, they conduct ex post welfare analysis after a surprise shock has just hit. Alternatively, we consider ex ante welfare analysis for agents who are aware of catastrophic asset risk and solve a dynamic stochastic problem with uncertainty over the timing of this risk.

7 Financial shocks that reduce the likelihood of repayment are beyond the scope of this paper. Individuals in our model follow consumption and saving plans whereby any debts are fully repaid by the end of the life cycle.

8 We are not attempting to provide a historical account of the welfare effects of the Great Depression, and hence we are not calibrating all parameters to the 1930s.

9 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.

10 Recall that it took 25 years for the S&P 500 index to fully recover to its peak level: it was not until September 1954 that the index finally returned to its previous high from August 1929. As a validation of our calibration, we note that our assumptions embody the same slow recovery time. With a sudden loss of wealth share S , it would take n years for wealth to return to its previous level if the post-shock growth rate in wealth is r ,

$$n = \frac{1}{r} \ln \left(\frac{1}{1-S} \right).$$

With $S = 0.86$ and $r = 0.08$, we obtain $n = 24.6$, which is a close fit to the actual recovery speed following the Great Depression.

11 Equivalently, we could motivate this section by imagining that the policy maker has full information about the timing of the risk but cares only about the worst-case scenario in policy evaluation.

12 In our model, all household's face the same mortality risk. However, in reality, life-time earnings and mortality are related. The redistributive role of Social Security in light of differential mortality is explored in several papers such as Garrett (1995), Goda et al. (2011), Bishnu et al. (2019), Sheshinski and Caliendo (2021), Coronado et al. (1999, 2000, 2002, 2011), and Liebman (2002).

13 In overlapping generations models, Social Security can also provide insurance against aggregate shocks through inter-generational redistribution. See Olovsson (2010) and Harenberg and Ludwig (2019) for the interactions of aggregate and idiosyncratic risk on the welfare effects of Social Security.

References

- Andersen, T. and J. Bhattacharya. (2011) On myopia as rationale for social security. *Economic Theory* 47(1), 135–158.
- Bagchi, S. (2016) Is the social security crisis really as bad as we think? *Macroeconomic Dynamics* 20(3), 737–776.
- Bishnu, M., N. L. Guo and C. S. Kumru. (2019) Social security with differential mortality. *Journal of Macroeconomics* 62, 103077.
- Blau, D. M. (2016) Pensions, household saving, and welfare: A dynamic analysis of crowd out. *Quantitative Economics* 7(1), 193–224.
- Buccioli, A. (2011) A note on social security welfare with self-control problems. *Macroeconomic Dynamics* 15(4), 579–594.
- Caliendo, F. N. (2011) Time-inconsistent preferences and social security: Revisited in continuous time. *Journal of Economic Dynamics and Control* 35(5), 668–675.
- Caliendo, F. N., M. Casanova, A. Gorry and S. Slavov. (2023) Retirement timing uncertainty: Empirical evidence and quantitative evaluation. *Review of Economic Dynamics*. <https://doi.org/10.1016/j.red.2023.01.002>
- Caliendo, F. N. and T. S. Findley. (2013) Limited computational ability and social security. *International Tax and Public Finance* 20(3), 414–433.
- Caliendo, F. N., A. Gorry and S. Slavov. (2019) The cost of uncertainty about the timing of social security reform. *European Economic Review* 118, 101–125.
- Chetty, R. and A. Finkelstein. (2013) Chapter 3 - Social insurance: Connecting theory to data. In: A. J. Auerbach, R. Chetty, M. Feldstein and E. Saez. (eds.), *Handbook of Public Economics*, Vol. 5, pp. 111–193. Amsterdam: Elsevier.
- Conesa, J. C. and C. Garriga. (2008) Optimal fiscal policy in the design of social security reforms. *International Economic Review* 49(1), 291–318.
- Coronado, J. L., D. Fullerton and T. Glass. (1999) *Distributional Impacts of Proposed Changes to the Social Security System*, pp. 149–186. Cambridge, MA: MIT Press.
- Coronado, J. L., D. Fullerton and T. Glass. (2000) Long run effects of social security reform proposals on lifetime progressivity. National Bureau of Economic Research, Working Paper 7568.
- Coronado, J. L., D. Fullerton and T. Glass. (2002) Long-run effects of social security reform proposals on lifetime progressivity. In: M. Feldstein and J. B. Liebman (eds.), *The Distributional Aspects of Social Security and Social Security Reform*, pp. 149–206. Chicago, IL: University of Chicago Press.

- Coronado, J. L., D. Fullerton and T. Glass. (2011) The progressivity of social security. *The B.E. Journal of Economic Analysis & Policy* 11(1), 70.
- Cottle Hunt, E. and F. Caliendo. (2022a) Social security and risk sharing: The role of economic mobility across generations. *International Tax and Public Finance*. <https://doi.org/10.1007/s10797-022-09761-x>
- Cottle Hunt, E. and F. N. Caliendo. (2022b) Social security and longevity risk: An analysis of couples. *Journal of Public Economic Theory* 24(3), 547–579.
- Cottle Hunt, E. N. and F. N. Caliendo. (2021) Social security and longevity risk: The case of risky bequest income. *Macroeconomic Dynamics* 26(6), 1647–1678.
- Cremer, H., P. De Donder, D. Maldonado and P. Pestieau. (2007) Voting over type and generosity of a pension system when some individuals are myopic. *Journal of Public Economics* 91(10), 2041–2061.
- 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., P. D. Donder, D. Maldonado and P. Pestieau. (2009) Forced saving, redistribution, and nonlinear social security schemes. *Southern Economic Journal* 76(1), 86–98.
- Cremer, H. and P. Pestieau. (2011) Myopia, redistribution and pensions. *European Economic Review* 55(2), 165–175.
- Diamond, P. (1977) A framework for social security analysis. *Journal of Public Economics* 8(3), 275–298.
- Diamond, P. (2004) Social security. *American Economic Review* 94(1), 1–24.
- Docquier, F. (2002) On the optimality of public pensions in an economy with life-cyclers and myopes. *Journal of Economic Behavior & Organization* 47(1), 121–140.
- Fehr, H., C. Habermann and F. Kindermann. (2008) Social security with rational and hyperbolic consumers. *Review of Economic Dynamics* 11(4), 884–903.
- Feldstein, M. (1985) The optimal level of social-security benefits. *Quarterly Journal of Economics* 100(2), 303–320.
- Feldstein, M. (2005) Rethinking social insurance. *American Economic Review* 95(1), 1–24.
- Feldstein, M. and E. Rangelova. (1998) Individual risk and intergenerational risk sharing in an investment-based social security program. National Bureau of Economic Research, Working Paper 6839.
- Findley, T. S. and F. N. Caliendo. (2008) The behavioral justification for public pensions: A survey. *Journal of Economics & Finance* 32(4), 409–425.
- Garrett, D. M. (1995) The effects of differential mortality rates on the progressivity of social security. *Economic Inquiry* 33(3), 457–475.
- Goda, G. S., J. B. Shoven and S. N. Slavov. (2011) *Differential Mortality by Income and Social Security Progressivity*, pp. 189–204. Chicago, IL: University of Chicago Press.
- Gourinchas, P.-O. and J. A. Parker. (2002) Consumption over the life cycle. *Econometrica* 70(1), 47–89.
- Gul, F. and W. Pesendorfer. (2004) Self-control, revealed preference and consumption choice. *Review of Economic Dynamics* 7(2), 243–264.
- Guo, N. L. and F. N. Caliendo. (2014) Time-inconsistent preferences and time-inconsistent policies. *Journal of Mathematical Economics* 51, 102–108.
- Harenberg, D. and A. Ludwig. (2019) Idiosyncratic risk, aggregate risk, and the welfare effects of social security. *International Economic Review* 60(2), 661–692.
- Huang, H., S. İmrohoroğlu and T. J. Sargent. (1997) Two computations to fund social security. *Macroeconomic Dynamics* 1(1), 7–44.
- Huggett, M. and J. Parra. (2010) How well does the U.S. social insurance system provide social insurance? *Journal of Political Economy* 118(1), 76–112.
- 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. (2003) Time-inconsistent preferences and social security. *The Quarterly Journal of Economics* 118(2), 745–784.
- İmrohoroğlu, S. and S. Kitao. (2012) Social security reforms: Benefit claiming, labor force participation, and long-run sustainability. *American Economic Journal: Macroeconomics* 4(3), 96–127.
- Kaygusuz, R. (2015) Social security and two-earner households. *Journal of Economic Dynamics and Control* 59, 163–178.
- Kitao, S. (2014) Sustainable social security: Four options. *Review of Economic Dynamics* 17(4), 756–779.
- Kumru, C. S. and A. C. Thanopoulos. (2008) Social security and self control preferences. *Journal of Economic Dynamics and Control* 32(3), 757–778.
- Li, Y. (2018) Economic analysis of social security survivors insurance. *International Economic Review* 59(4), 2043–2073.
- Liebman, J. B. (2002) Redistribution in the current U.S. Social Security system. In: M. Feldstein and J. B. Liebman (eds.), *The Distributional Aspects of Social Security and Social Security Reform*, pp. 11–48. Chicago, IL: University of Chicago Press.
- McGrattan, E. R. and E. C. Prescott. (2017) On financing retirement with an aging population. *Quantitative Economics* 8(1), 75–115.
- Nishiyama, S. and K. Smetters. (2007) Does social security privatization produce efficiency gains? *Quarterly Journal of Economics* 122(4), 1677–1719.

- Olovsson, C. (2010) Quantifying the risk-sharing welfare gains of social security. *Journal of Monetary Economics* 57(3), 364–375.
- Pestieau, P. and U. Possen. (2008) Prodigality and myopia—Two rationales for social security. *The Manchester School* 76(6), 629–652.
- Peterman, W. B. and K. Sommer. (2019a) A historical welfare analysis of social security: Whom did the program benefit? *Quantitative Economics* 10(4), 1357–1399.
- Peterman, W. B. and K. Sommer. (2019b) How well did social security mitigate the effects of the great recession? *International Economic Review* 60(3), 1433–1466.
- Sheshinski, E. and F. N. Caliendo. (2021) Social security and the increasing longevity gap. *Journal of Public Economic Theory* 23(1), 29–52.
- Slavov, S., D. Gorry, A. Gorry and F. N. Caliendo. (2019) Social security and saving: An update. *Public Finance Review* 47(2), 312–348.
- Stokey, N. L. (2016) Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics* 21, 246–265.