Social Security Safety Net with Rare Event Risk

Erin Cottle Hunt* Frank N. Caliendo†

Lafayette College Utah State University

November 19, 2021

Abstract

The S&P 500 lost 86% of its value during the Great Depression, representing the worst aggregate decline in equity valuations in modern history. More recent examples of sudden loss of investor wealth include the dot.com bubble, Great Recession, and COVID-19 corrections. Because the US Social Security program was created in 1935 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 destruction of private wealth. We construct a model in which a rare event causes devastating shock to wealth occurring at an unknown time during life cycle. Despite the intuitive appeal of Social Security as a safety net, we do not find an economic basis for such a justification, even when the magnitude of the asset shock is calibrated to the Great Depression.

*cottlee@lafayette.edu, corresponding author
†frank.caliendo@usu.edu

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

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

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.[2]

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. We do this by constructing a life-cycle model in which the individual faces a one-time risk of losing a substantial portion of their private wealth. The timing of the shock is unknown. It may happen at any point during the life cycle (or never, if the individual is lucky).

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

[2] Of course, investors who are less diversified have experienced many other deep corrections in valuations, such as those from being heavily invested in a single asset (as in the case of a company stock buying program, etc.).

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. To further emphasize a worst case calibration scenario, the individual cannot escape the shock. Instead, the shock is guaranteed to occur within the individual's lifetime, although the timing of the shock is uncertain.[3] Both before and after the shock hits, we assume the individual's assets grow at 8% per year to reflect long run average equity returns in the US. 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.

We do not find a safety net rationale for Social Security in this setting. Instead, we find that a Social Security program reduces lifetime welfare by 7.64%. That is, the individual would need to receive an additional 7.64% of lifetime consumption in order to compensate them for the welfare losses imposed by Social Security.

Importantly, in our baseline calibration individuals face longevity risk in addition to rare event asset risk, and hence the mandatory annuitization feature of Social Security is present. Social Security pays benefits as a life annuity that lasts as long as the individual lives, and this feature by itself provides rather large welfare gains assuming that the individual lacks

[3] Of course, the individual would prefer that the shock happen very early in the life cycle when they haven't accumulated much to lose, but they don't control the timing of the shock and must therefore plan for the possibility of a more painful shock deeper in the life cycle when accumulated wealth is substantial.

access to private insurance markets. But adding a risky asset whose rewards reflect that of US equities and whose risk is calibrated to the Great Depression, not only reduces the welfare effect of Social Security but reduces it by so much that the overall effect becomes negative.

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, low frequency 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 well 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. This is an interesting finding because Social Security's role as a safety net against shocks to private wealth is a common argument for the program's existence.

Our negative welfare result is important because we have stacked our analysis in favor of Social Security in our choice of a number of modeling assumptions. First, we assume the financial shock is guaranteed to happen. In reality some people may never experience a terrible shock, either because of luck or through diversification. Second, we assume not only that the shock is going to hit during the lifetime of the individual, but it will be the biggest possible shock in modern history. Third, the individual's only choice for private saving is a risky asset. Fourth, insurance markets are closed and there are no private options for insuring longevity risk and asset risk. Hence, by assuming the individual faces an inescapable, large shock with no access to a safe asset or private insurance, 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.[4]

Given Social Security's status as the largest government program in operation, it's not surprising that an enormous 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 (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 Possen (2008), Bucciol (2011), Cremer and Pestieau (2011), Caliendo (2011), Andersen and Bhattacharya (2011), Caliendo and Findley (2013), and Guo and Caliendo (2014), among many others). Another strand focuses on Social Security's role as a hedge against longevity risk (see Diamond (1977), Diamond (2004), Feldstein (2005), and Chetty and Finkelstein (2013), among many others). And another strand considers its role in hedging lifetime earnings risk or reducing inequality (Diamond (1977), Huggett and Ventura (1999), Nishiyama and Smetters (2007), İmrohoroğlu and Kitao (2012), Bagchi (2016), Huggett and Parra (2010), Cottle Hunt and Caliendo (2020), and others). Still other risks such as disability risk and survivor risk to the dependent spouse and children also have been thoroughly examined. 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.

[4]Our result does depend on our assumption that individuals save rationally in the face of risk, and that they have full information about the risks that they face. If we leave the Neoclassical paradigm and assume individuals do not behave rationally, then Social Security can improve welfare. But there are already many research papers that illustrate this point and we would not even need asset risk to rationalize Social Security if we dive into behavioral models.

A few studies consider the role of Social Security when private asset returns are risky. For example, Feldstein and Ranguelova (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 paper, 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, large 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. Agents in our model are aware of the rare event risk that they face. They know that a Great Depression style asset shock could hit, but they don't know when. The main feature of our analysis is decision making in the face of uncertainty about the timing of a large shock to asset values. Like their conclusion, we do not find a safety net role either.[5]

In related work, Peterman and Sommer (2019b) find that Social Security helped to mit-

[5] 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.

igate the welfare loss associated with the Great Recession, which caused sudden shocks to wealth and employment. The Great Recession takes households completely by surprise: 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.

2 Model

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:

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

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.[6] Therefore, notationally, we denote the state dependency of the shock

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

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

Post-shock 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:

[6] 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.

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

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

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

Pre-shock problem. Next we solve the dynamic stochastic problem of a household that hedges the timing risk of shock $S(k(t))$.

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 pre-shock 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

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

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

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

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

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

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

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

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

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 pre-shock and post-shock 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 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.[7]

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 US during the period 2000-2010.[8] The life expectancy at $t = 0$ implied by this

[7]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.

[8]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.

survival function is age 77.1, which corresponds closely to the life expectancy from birth for the US 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)

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

where λ 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

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

We set the interest rate $r = 0.08$ which corresponds to the average long-term real return of 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

Parameters	
Social Security:	
$\tau = 10.6\%$	Social Security payroll tax
$b = 0.348$	balances the Social Security budget
Rare Event Shock:	
$\phi(t)$	equation 26 in text, uniform distribution
$S = 0.86$	S&P loss during the Great Depression
Demographics and misc:	
$T = 82$	normalized maximum lifespan (age 18 to 100)
$t_R = 49$	length of career (age 18 to 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

Table 1: Summary of Baseline Calibration of Parameters

hit during their life cycle. Hence:

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

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 $S = 0.86$.[9] Our baseline parameters are summarized in Table 1.

[9]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.

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 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.

In our model, the individual has full information about the risks that they face (longevity risk and financial risk) and they behave optimally in the face of these risks. Hence, our calculation of the welfare gains from Social Security should be treated as a lower bound. It is possible for the program to confer larger gains than what we report here if individuals make suboptimal decisions in the face of uncertainty.

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

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

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 be need to receive an additional 7.89% of lifetime consumption in order to compensate for participating in Social Security.

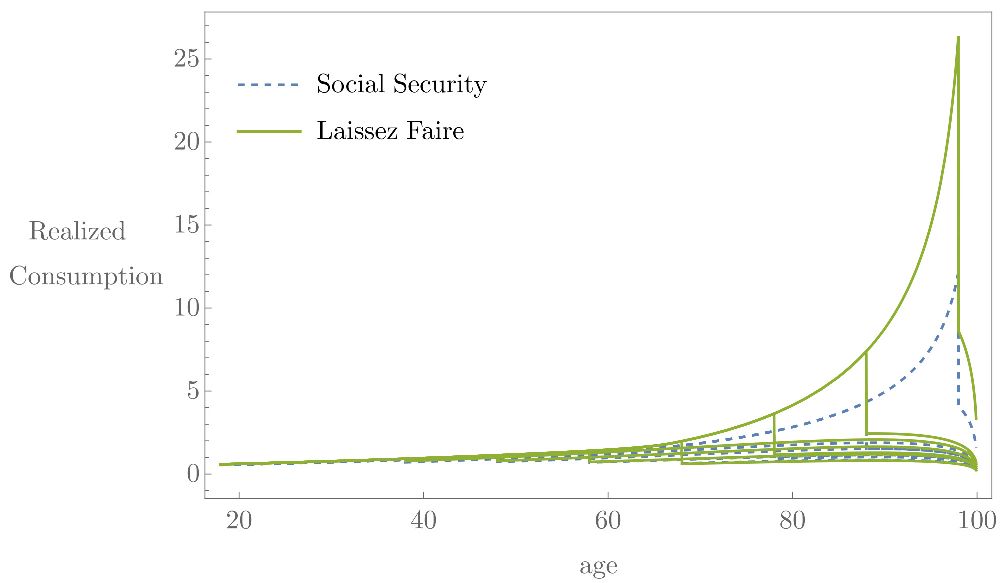

Figure 1: Realized consumption given $S = 0.86$, with and without Social Security.

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 well 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 48, 58, 68, 78, and 88). 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.

Consider a simple example to illustrate. Suppose the individual saves \$1 per year across the 49-year working period, compounded continuously at $r = 8\%$, and upon retirement annuitizes the balance across the maximum lifespan. Further assume that the one-time shock is the same size as the Great Depression ($S = 0.86$), and that it hits at the worst possible time (the moment the individual retires when wealth is at a maximum). With these assumptions, the retirement annuity a is the solution to the following equation

$$\int_{t_R}^{T} ae^{-rt}dt - \int_{0}^{t_R} e^{-rt}dt \times (1 - S) = 0.$$

Under these assumptions, the retirement annuity $a = 7.45$. Alternatively, an asset with rate of return r^* and no risk ($S = 0$) would confer the same retirement annuity if $r^* = 3.75\%$. In this worst case scenario where the worst possible shock (Great Depression) happens at the worst possible time (retirement), the individual experiences a "lifetime" ex post rate of return that is equivalent to 3.75%, provided the individual is invested in equities that pay 8% both before and after the financial shock. A shock that hits at an earlier (or later) time leaves the individual in a better position with an ex post rate of return that exceeds 3.75% and approaches 8% as the timing of the shock approaches the boundaries of the life cycle.

How do these rates of return compare to Social Security? The longitudinal internal rate of return on Social Security that accounts for its life annuity feature (i.e., the payment of benefits for as long as the individual lives), is the value of i that solves the following equation:

$$\int_{t_R}^{T} e^{-it} b dt - \int_{0}^{t_R} e^{-it} \tau w(t) dt = 0.$$

Given our parameterization of the length of the working and retirement periods, the wage profile over the life cycle, and Social Security taxes and benefits that balance the aggregate budget, the internal rate of return on Social Security is 2.0%. Clearly, in a simple rate of return comparison such as this, even the worst outcome for equities in our model (3.75% lifetime rate of return) beats the longitudinal internal rate of return on Social Security.

These calculations illustrate why there are no welfare gains associated with a safety net for rare event risk in our model. No amount of risk aversion would cause the individual to prefer to pay taxes that ultimately net them a 2.0% return in Social Security over a risky asset whose worst possible lifetime return is 3.75%.

Our negative welfare result is important because we have stacked our analysis in favor of Social Security in our choice of a number 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.

To overturn our negative welfare result, we would need to either decrease the rate of return on equities r or increase the magnitude of the risk S beyond the historical values that we are using in our calibration. In fairness to Social Security as a safety net against financial shocks, a Great Depression shock is not actually the worst possible outcome. Some individuals could experience a total financial loss ($S = 1$) at some point over the course of their lifetimes. For example, an individual who is heavily invested in a single enterprise such as a family business or company stock could experience a total financial loss if that enterprise fails. And even the 86% loss in financial wealth during the Great Depression was the loss to the S&P 500 *index*, implying that some investors who were less diversified than the index did better and worse than the index.

In the case of a total financial loss occurring at the worst possible date (retirement), the individual's lifetime rate of return on equities is -100%. The lifetime rate of return would be higher if the total loss happens at any other time and approaches 8% as the timing of the shock approaches the boundaries of the life cycle. But now there is at least the possibility that Social Security's internal rate of return beats equities, and with a high enough level of risk aversion, there will be a positive welfare role for Social Security as a safety net against rare event risk. In fact, as long as the size of the shock is greater than about 94%, it will be theoretically possible to rationalize Social Security on the grounds that it pays better than the worst outcomes associated with equities.

2.5 Decomposition

We decompose the welfare gain from 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 -7.89% in our baseline model. This suggests that Social Security is less harmful in the presence of rare event risk.

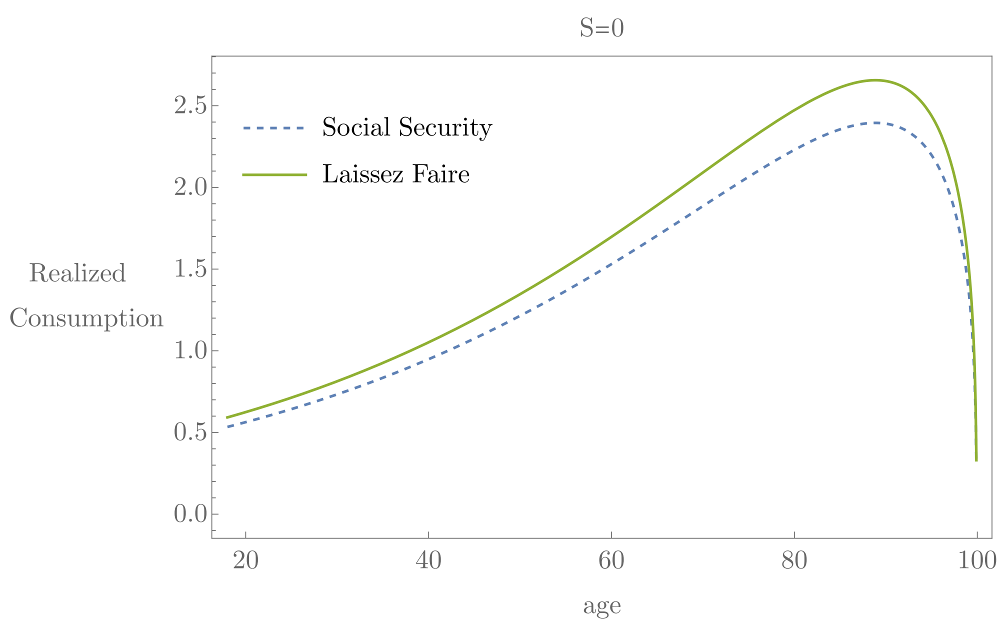

Figure 2: Realized consumption given $S = 0$, with and without Social Security.

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 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 that has a zero percent 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.

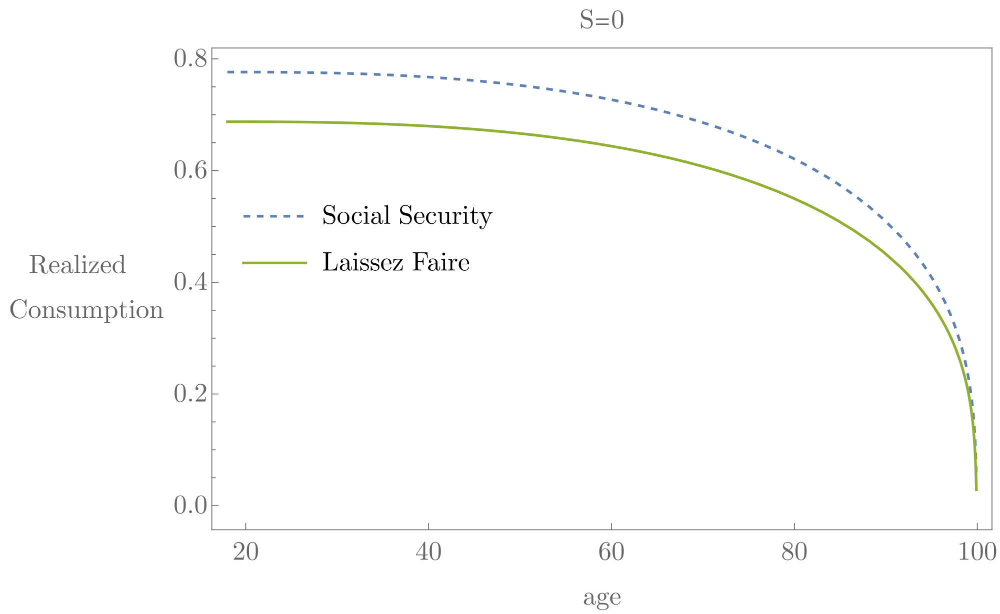

Figure 3: Realized consumption given $S = 0$ and $r = 0$, with and without Social Security.

3 Robustness Check

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.

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.

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

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

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.

4 Inequality

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 piece-wise 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.

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. Recall that 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 trade-off 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.

5 Conclusion

Social Security is the biggest 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 little 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. Importantly, there are no safe assets in our model which means that there is no way for the individual to escape financial risk in saving for retirement. On the other hand, Social Security is a safe asset with guaranteed benefits that are paid out as a life annuity. These modeling assumptions tilt our welfare analysis in favor of finding a positive welfare role for Social Security.

However, despite the fact that the individual in our model faces the risk of massive losses in financial wealth and cannot escape this risk, we do not find a role for Social Security as a safety net against such risk. Instead, the individual in our model is better off without Social Security. The basic intuition for our result is this: even when accounting for the possibly of a large loss in wealth on the scale of the Great Depression, US equities still represent a more attractive asset than Social Security. In other words, mandating that individuals hold a safe asset (Social Security) in addition to a risky asset (equities) does not improve welfare in our model. This result is robust to the degree of risk aversion of the individual.

Of course, this does not necessarily imply that Social Security is welfare reducing overall. Social Security is a complex program with a number of aspects that might improve welfare (longevity insurance, redistribution, forced saving for those who do not save rationally, disability insurance, survivor insurance, and spousal benefits). We do not take away from any of these potential roles. However, it is important that a commonly held rationale for Social Security—a safety net against large and devastating shocks to private savings—may not be a strong argument after all.

References

Andersen, T. and Bhattacharya, J. (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.

Bucciol, 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., Casanova, M., Gorry, A., and Slavov, S. (2020). Retirement timing uncertainty: Empirical evidence and quantitative evaluation. *working paper*.

Caliendo, F. N. and Findley, T. S. (2013). Limited computational ability and social security. *International Tax and Public Finance*, 20(3):414 – 433.

Caliendo, F. N., Gorry, A., and Slavov, S. (2019). The cost of uncertainty about the timing of social security reform. *European Economic Review*, 118:101 – 125.

Chetty, R. and Finkelstein, A. (2013). Chapter 3 - Social Insurance: Connecting Theory to Data. In Auerbach, A. J., Chetty, R., Feldstein, M., and Saez, E., editors, *Handbook of Public Economics, Vol. 5*, pages 111 – 193. Elsevier.

Conesa, J. C. and Garriga, C. (2008). Optimal fiscal policy in the design of social security reforms*. *International Economic Review*, 49(1):291–318.

Cottle Hunt, E. and Caliendo, F. (2020). Social Security and Risk Sharing: The Role of Economic Mobility across Generations. working paper.

Cottle Hunt, E. N. and Caliendo, F. N. (2021). Social security and longevity risk: The case of risky bequest income. *Macroeconomic Dynamics*, page 1–32.

Cremer, H., De Donder, P., Maldonado, D., and Pestieau, P. (2007). Voting Over Type and Generosity of A Pension System When Some Individuals Are Myopic. *Journal of Public Economics*, 91:2041–2061.

Cremer, H., Donder, P., Maldonado, D., and Pestieau, P. (2008). Designing a linear pension scheme with forced savings and wage heterogeneity. *International Tax and Public Finance*, 15(5):547–562.

Cremer, H., Donder, P. D., Maldonado, D., and Pestieau, P. (2009). Forced Saving, Redistribution, and Nonlinear Social Security Schemes. *Southern Economic Journal*, 76(1):86–98.

Cremer, H. and Pestieau, P. (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., Habermann, C., and Kindermann, F. (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 Ranguelova, E. (1998). Individual risk and intergenerational risk sharing in an investment-based social security program. Working Paper 6839, National Bureau of Economic Research.

Findley, T. S. and Caliendo, F. N. (2008). The behavioral justification for public pensions: a survey. *Journal of Economics & Finance*, 32(4):409 – 425.

Gourinchas, P.-O. and Parker, J. A. (2002). Consumption over the life cycle. *Econometrica*, 70(1):47–89.

Gul, F. and Pesendorfer, W. (2004). Self-Control, Revealed Preference and Consumption Choice. *Review of Economic Dynamics*, 7(2):243 – 264.

Guo, N. L. and Caliendo, F. N. (2014). Time-inconsistent preferences and time-inconsistent policies. *Journal of Mathematical Economics*, 51:102 – 108.

Huang, H., İmrohoroğlu, S., and Sargent, T. J. (1997). Two computations to fund social security. *Macroeconomic Dynamics*, 1(1):7–44.

Huggett, M. and Parra, J. (2010). How well does the u.s. social insurance system provide social insurance? *Journal of Political Economy*, 118(1):76–112.

Huggett, M. and Ventura, G. (1999). On the distributional effects of social security reform. *Review of Economic Dynamics*, 2(3):498 – 531.

İmrohoroğlu, A., İmrohoroğlu, S., and Joines, D. H. (2003). Time-Inconsistent Preferences and Social Security*. *The Quarterly Journal of Economics*, 118(2):745–784.

İmrohoroğlu, S. and Kitao, S. (2012). Social Security Reforms: Benefit Claiming, Labor Force Participation, and Long-Run Sustainability. *American Economic Journal: Macroeconomics*, 4(3):96–127.

Kitao, S. (2014). Sustainable social security: Four options. *Review of Economic Dynamics*, 17(4):756 – 779.

Kumru, c. S. and Thanopoulos, A. C. (2008). Social Security and Self Control Preferences. *Journal of Economic Dynamics and Control*, 32(3):757 – 778.

McGrattan, E. R. and Prescott, E. C. (2017). On Financing Retirement with an Aging Population. *Quantitative Economics*, 8(1):75 – 115.

Nishiyama, S. and Smetters, K. (2007). Does social security privatization produce efficiency gains? *Quarterly Journal of Economics*, 122(4):1677–1719.

Pestieau, P. and Possen, U. (2008). Prodigality and Myopia–Two Rationales for Social Security. *The Manchester School*, 76(6):629–652.

Peterman, W. B. and Sommer, K. (2019a). A historical welfare analysis of social security: Whom did the program benefit? *Quantitative Economics*, 10(4):1357–1399.

Peterman, W. B. and Sommer, K. (2019b). How well did social security mitigate the effects of the great recession? *International Economic Review*, 60(3):1433–1466.

Stokey, N. L. (2016). Wait-and-see: Investment options under policy uncertainty. *Review of Economic Dynamics*, 21:246 – 265.