Relaxing Household Liquidity Constraints through Social Security

Sylvain Catherine
Wharton

Max Miller
Wharton

Natasha Sarin*
PennLaw & Wharton

May 8, 2020

Abstract

More than a quarter of working-age households in the United States do not have sufficient savings to cover their expenditures after a month of unemployment. We explore proposals to alleviate financial distress arising from the COVID-19 pandemic. We show that giving workers early access to just 1% of their future Social Security benefits allows most households to maintain their current consumption for at least two months. Unlike other approaches (like early access to retirement accounts, stimulus relief checks, and expanded unemployment insurance), access to Social Security serves the needs of workers made vulnerable by the crisis, but does not increase the overall liabilities of the federal government or have distortionary effects on the labor market.

Keywords: Covid-19, Social Security, Household Finance
JEL codes: E21, G51, H55, H12

*Preliminary and in progress. All remaining errors are our own.

1 Introduction

The COVID-19 pandemic has pushed US unemployment to its highest level since the Great Depression. Most households do not have sufficient savings to weather this unexpected calamity (Bhutta and Dettling, 2018, Board of Governors of the Federal Reserve System, 2017). This paper discusses the potential of allowing workers to tap their Social Security wealth to finance consumption today. We show that distributing just 1% of the present value of scheduled benefits provides significant liquidity to households. This approach may well be superior to alternatives already legislated because it delivers relief to households that need it most, allowing them to borrow at near-zero rather than sky-high private borrowing rates.

This conclusion is not surprising. As Figure 1 illustrates, Social Security benefits are relatively evenly distributed across the wealth distribution, whereas the value of retirement accounts and liquid savings is concentrated in the top decile of the wealth distribution. Social Security is hugely significant to most Americans: it represents more than 30 trillion dollars and nearly 60% of the wealth of the bottom 90% (Catherine et al., 2020). This means that, in exchange for relatively small cuts in scheduled benefits, the government can send sizable checks to most households in a fiscally neutral way. Specifically, we show that just a 1% cut in scheduled benefits is able to provide sufficient liquidity to allow most households to continue to meet their expenditures for at least two months in case of unemployment.

From households' point of view, this policy allows them to borrow against their retirement benefits at historically low interest rates. Households who do not need this loan can choose to invest the money in government bonds and should be indifferent. From the point of view of the government, this policy transforms implicit Social Security liabilities into public debt but leaves its overall long-run obligations unchanged.

Figure 1: Distribution of various forms of wealth

This figure shows the distribution of Social Security wealth, retirement wealth, and liquid wealth across deciles of the marketable wealth distribution. The red bar denotes the per household average present value of scheduled Social Security benefits, the green bar shows the per household average amount of retirement savings, and the blue bar displays the per household average liquid wealth. We calculate the present value of Social Security benefits by simulating workers' earnings trajectories and matching this data with the SCF based on current earnings. Retirement accounts are defined as IRA accounts, thrift accounts, or any current or future defined contribution pension obligations and come from the SCF. Liquid wealth is defined as all wealth held in transactions accounts, certificates of deposit, mutual funds, stocks, bonds, and also come from the SCF.

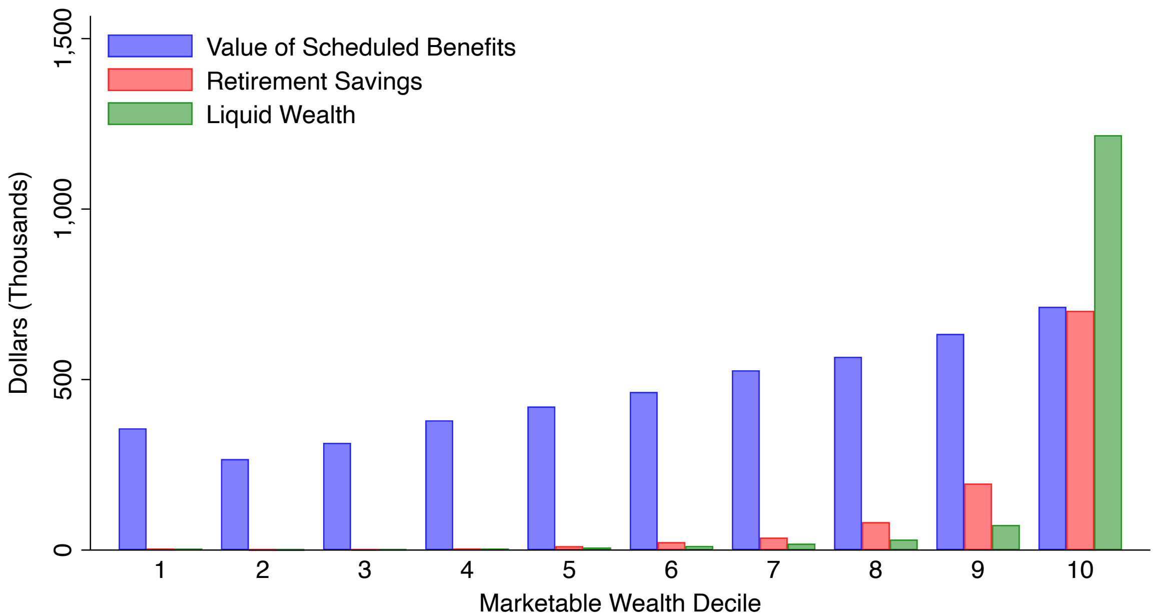

We compare this approach to already enacted alternatives: allowing workers to tap retirement accounts without penalty, $1,200 stimulus checks, and the extension of unemployment insurance by $600 per week. To do so, we use the 2016 Survey of Consumer Finances (SCF) to measure how long it takes for households to run out of liquid savings in case of unemployment. Then, we study how the distribution of this measure of liquidity constraints varies in response to each policy. Distributing 1% of the value of scheduled benefits allows 75% of households to go through 3 months of unemployment without cutting their consumption, which is longer than most alternative approaches.

This paper adds to several strands of literature. First, we contribute to the growing literature on the economic impact of COVID-19 and evaluation of policies aimed at stemming it (Baker et al., 2020, Bartik et al., 2020, Gormsen and Koijen, 2020). In recent work, Biggs and Rauh (2020) contemplate a closely related policy: allowing workers to access Social Security wealth today and delay retirement to repay these benefits.[1] Under current law, we estimate that a 1% cut in benefits can be offset by claiming benefits six weeks later. More generally, our estimates of the market value of Social Security benefits and the implications of policy for household liquidity can easily be extended to evaluate alternative approaches.

We also add to the literature on the optimal design of public savings programs. In the U.S. and many other countries, public savings are designed to be illiquid to supplement the private market for longevity insurance, plagued by adverse selection problems (Abel, 1986, Hosseini, 2015); and to deter overconsumption by behavioral households suffering from present bias (Beshears et al., 2019). We extend this literature by pointing out that the optimal mandatory savings rate cannot be static and should be revisited in moments of crisis when households' liquidity needs are pronounced. Much work has advocated the provision of lump-sum benefits of Social Security wealth to discourage early retirement, noting households' preferences for one-time payouts that enable them to pay down mortgages or other debt (Maurer and Mitchell, 2018, Maurer et al., 2016). Our proposal builds on this insight, suggesting that lump sum payments in this crisis would provide households a way to finance expenditure at record-low rates of interest, with the alternative being sky-high private borrowing rates.

The remainder of our paper proceeds as follows. Section 2 describes the Social Security program and our approach for valuing Social Security wealth, and estimates the consequences of early access to Social Security wealth across the age distribution. Section 3 compares this approach to other alternatives to increasing households' liquidity, including tapping retirement accounts, stimulus checks, and extended unemployment benefits. Section 4 concludes.

[1] Importantly, unlike Biggs and Rauh (2020) we contemplate a universal program; rather than an opt-in approach. The latter raises concerns about adverse selection stemming from individuals choosing to withdraw today because of hidden information about the future value of their Social Security wealth and expected longevity (Abel, 1986, Eckstein et al., 1985).

2 Valuing scheduled benefits

In this section, we estimate how much can be paid immediately to American households in exchange for a small cut in future Social Security benefits. Because benefits are determined based on individuals' historical earnings, the present value of benefits depends on age and workers' earnings trajectories.

We estimate the market value of a benefit cut in two steps. First, we compute expected benefits by simulating earnings trajectories and apply the Social Security benefit formula assuming all workers retire at full retirement age. Second, we discount expected benefits using the real yield curve[2] implied by Treasury inflation-protected securities (TIPS) and taking into account the long run correlation between Social Security and stock market returns, following the approach of (Catherine et al., 2020).

2.1 Expected benefits

Simulating earnings To forecast benefits, we simulate earnings using the income process estimated in Guvenen et al. (2019). Specifically, we assume that a worker i earnings at age t are:

$$L_{it} = L_{1,t} \cdot L_{2,it}. \tag{2.1}$$

where $L_{1,t}$ is the average wage in the economy and $L_{2,it}$ represents the idiosyncratic component of earnings. The latter evolves as follows:

[2]The construction of this series is detailed in Section A.1.

$$\text{Level of idiosyncratic earnings:} \quad L_{2,it} = (1-\nu_t^i)e^{\left(g(t)+\alpha^i+\beta^i t+z_t^i+\varepsilon_t^i\right)} \tag{2.1.1}$$

$$\text{Persistent component:} \quad z_t^i = \rho z_{t-1}^i + \eta_t^i \tag{2.1.2}$$

$$\text{Innovations to AR(1):} \quad \eta_t^i \sim \begin{cases} \mathcal{N}(\mu_{\eta,1}, \sigma_{\eta,1}^2) & \text{with prob. } p_z \\ \mathcal{N}(\mu_{\eta,2}, \sigma_{\eta,2}^2) & \text{with prob. } 1-p_z \end{cases} \tag{2.1.3}$$

$$\text{Initial condition of } z_t^i\text{:} \quad z_0^i \sim \mathcal{N}(0, \sigma_{z,0}^2) \tag{2.1.4}$$

$$\text{Transitory shock:} \quad \varepsilon_t^i \sim \begin{cases} \mathcal{N}(\mu_{\varepsilon,1}, \sigma_{\varepsilon,1}^2) & \text{with prob. } p_\varepsilon \\ \mathcal{N}(\mu_{\varepsilon,2}, \sigma_{\varepsilon,2}^2) & \text{with prob. } 1-p_\varepsilon \end{cases} \tag{2.1.5}$$

$$\text{Nonemployment duration:} \quad \nu_t^i \sim \begin{cases} 0 & \text{with prob. } 1-p_\nu(t, z_t^i) \\ \min\{1, \exp\{\lambda\}\} & \text{with prob. } p_\nu(t, z_t^i) \end{cases} \tag{2.1.6}$$

$$\text{Prob. of Nonemp. shock:} \quad p_\nu^i(t, z_t) = \frac{e^{\xi_t^i}}{1+e^{\xi_t^i}}, \text{ where } \xi_t^i = a + bt + cz_t^i + dz_t^i i \tag{2.1.7}$$

where z_i is a component of earnings with persistence ρ and innovations drawn from a mixture of normal distributions. Transitory shocks ε_i also have a normal mixture distribution. Finally, workers can experience a period of unemployment with probability p which depends on age, earnings and gender, and whose length follows an exponential distribution. We refer readers to Guvenen et al. (2019)'s study for more details.

Benefit formula Social Security benefits depend on individuals historical earnings and are computed in three steps. First, past taxable earnings are wage-indexed, which means that they are adjusted to reflect the growth in nominal wages up to the year a worker reaches age 60. In a second step, the average indexed yearly earnings ("AIYE") is determined by taking the mean of the best 35 years of indexed earnings. Finally, benefits are computed as a concave function of the AIYE. Specifically, benefits equal the sum of 90% of the share of the AIYE below the first Social Security "bend point" ($11,112 in 2019), 32% of the share AIYE between the first and second bend point ($66,996) and 15% of the remaining part of the AIYE. Since the 1980's, these bend points have tracked the evolution of earnings, representing 0.21 and 1.25 times the national wage index L_1. We assume that they will keep evolving that way. Hence, the value of benefits is a piece-wise linear function of the AIYE:

$$\text{Benefits}_i = \begin{cases} 0.9 \times AIYE_i & \text{if } AIYE/L_{1,60} < 0.21 \\ 0.1218 \times L_{1,60} + 0.32 \times AIYE & \text{if } 0.21 \leq AIYE/L_{1,60} < 1.25 \\ 0.3343 \times L_{1,60} + 0.15 \times AIYE & \text{if } 1.25 \leq AIYE/L_{1,60}, \end{cases} \tag{2.2}$$

where $L_{1,60}$ is the level of wage index when a worker turns 60.

2.2 Market value

We need to determine the present value of a stream of benefits protected against inflation, backed by the Federal government and indexed on the national wage index. We define the present value of expected benefits as:

$$\text{Value of Benefits}_{it} = \sum_{s=66}^{T} \left(\prod_{k=t}^{s-1} (1 - m_{ik}) \right) \Psi_s \mathbb{E} \left[\text{Benefits}_{it} \right] \tag{2.3}$$

where T is the maximum age, m_{ik} is mortality at age k, and Ψ_s is the appropriate discount factor for benefits paid at age s.

When discounting benefits, we take into account that wage indexation exposes the government to systematic risk because of the correlation between market returns and the wage index. This contemporaneous correlation is small but Benzoni et al. (2007) argue that the labor and stock markets are cointegrated, which reduces the present value of benefits substantially (Geanokoplos and Zeldes (2010), Catherine (2019)). To take this into account, we model the evolution of the log national index l_1 and the log cumulative market returns s_t as in Benzoni et al. (2007):

$$\begin{cases} dl_{1,t} = & \left((\phi - \kappa) y_t + \mu - \delta - \frac{\sigma_l^2}{2} \right) dt + v_1 dz_{1,t} \\ ds_t = & \left(\mu + \phi y_t - \frac{\sigma_s^2}{2} \right) dt + \sigma_s dz_{2,t} \\ dy_t = & -\kappa y_t + \sigma_l dz_{1,t} - \sigma_s dz_{2,t} \end{cases} \tag{2.4}$$

In these equations, $\mu - \delta$ determines the unconditional log aggregate growth rate of earnings and v_1 its volatility. μ and σ_s represent expected stock market log returns and their volatility. The state variable y_t keeps track of whether the labor market performed better or worse than the stock market relative to expectations. Finally, κ determines the strength of the cointegration between the labor and stock markets.

In Catherine et al. (2020), we show that the market beta of a "wage bond" paying a single cash flow indexed to the value of $L_{1,n}$ in n years is:

$$\beta_{L_{1,n}} = \left(1 - \frac{\phi}{\kappa}\right)\left(1 - e^{-\kappa n}\right) \tag{2.5}$$

and we demonstrate that, under the no-arbitrage condition, the expected return on such a bond is:

$$\mathrm{E}\left[r_{L_{1,n}}\right] = \beta_{L_{1,n}}\left(\mu - r\right) + r \tag{2.6}$$

where r is the risk-free rate. Therefore, for workers below age 60, the appropriate discount factor for a benefit expected at age s is:

$$\Psi_s \approx \left[\prod_{k=t+1}^{60}\left(1 + \beta_{L_{1,60-k}}\left(\mu - r\right) + f_k\right)\prod_{s=n+1}^{k}\left(1 + f_k\right)\right]^{-1}, \tag{2.7}$$

where f_k is the forward real interest rate between years $k-1$ and k.

2.3 Calibration and validity

We calibrate the dynamics of idiosyncratic earnings using the benchmark estimation of Guvenen et al. (2019). In Catherine et al. (2020), we use the same simulation strategy to estimate the value of future benefits, net of future payroll taxes, from 1989 to 2016. We validate this approach by showing that, when using the same macroeconomic assumptions, we can track very well the evolution of aggregate Social Security obligations reported by the Office of the Chief Actuary of the SSA. Moreover, we also show that our simulation produces full-retirement benefits that match those we observe in the SCF for different gender and cohorts. Finally, the income process estimated in Guvenen et al. (2019) matches a very large numbers of moments of the cross-section and dynamics of earnings.

We calibrate the model in Section 2.2 as in Benzoni et al. (2007). These authors estimate $\kappa = .16$ and $\phi = .08$ using US macroeconomic data from 1929 to 2004. This calibration implies a market beta of 0.5 for very distant Social Security benefits. We assume an equity premium of $\mu - r = 0.06$. We use the TIPS yield curve of April 2020 to compute forward interest rates. Finally, we assume that a 1% growth rate for the national wage index.

2.4 Results

We simulate past and future earnings for 800,000 workers per cohort, producing a cross-section of 36 million observations for the year 2020. The simulated dataset includes age, average past taxable earnings and the present value of expected benefits. We use this simulated data to estimate how much can be paid to workers today in exchange for a small cut in old-age benefits. Our focus is on working age (20 to 61 year-old) individuals. The answer to this question is a function of workers age and earnings histories (those who have contributed more to Social Security have greater benefits).

Panel A of Figure 2 illustrates this fact. The present value of a 1% cut is highest for workers who are approaching retirement because they borrow against more imminent cash flows. In contrast, it is less significant for workers who have just entered the labor force and who will start repaying this loan in forty years. But importantly, across the age and earnings distribution, just a 1% decrease in benefits significantly boosts liquidity by providing more than $2,000 to the large majority of workers. In dual-earner households, the provision of liquidity would be twice as large.

To illustrate this point another way, we consider what cut in benefits would be required to deliver workers $2,500 today (Figure 2, Panel B), enough to finance roughly one month of (median) household consumption. For all but the lowest earners, the decrease in future benefits is minor: for 40 year old individuals earning the median income of around $34,000, a $2,500 check represents between 0.7% (75th percentile) and 0.9% (25th percentile) of future benefits. It is possible to imagine supporting household consumption for several months through an approach like this one, with relatively minor implications for retirement wealth.

It is worth noting that is not the case for workers close to retirement with limited past earnings, for whom a $2,500 check today could represent between 5-10 percent of future benefits, which will scale quickly should this approach be adopted for several months. This is a population who has not accrued much Social Security wealth (e.g., because of little time spent in the workforce).

Figure 2: Price of early Social Security check

This figure shows the correspondence between benefit cuts and check size as a function of workers' age and the average past taxable earnings. Panel A shows how much can be paid in exchange for a 1% benefit cut. Panel B shows the benefit cuts corresponding to a $2,500 check. The graphs are constructing by simulating data following the procedure outlined in Section 2.

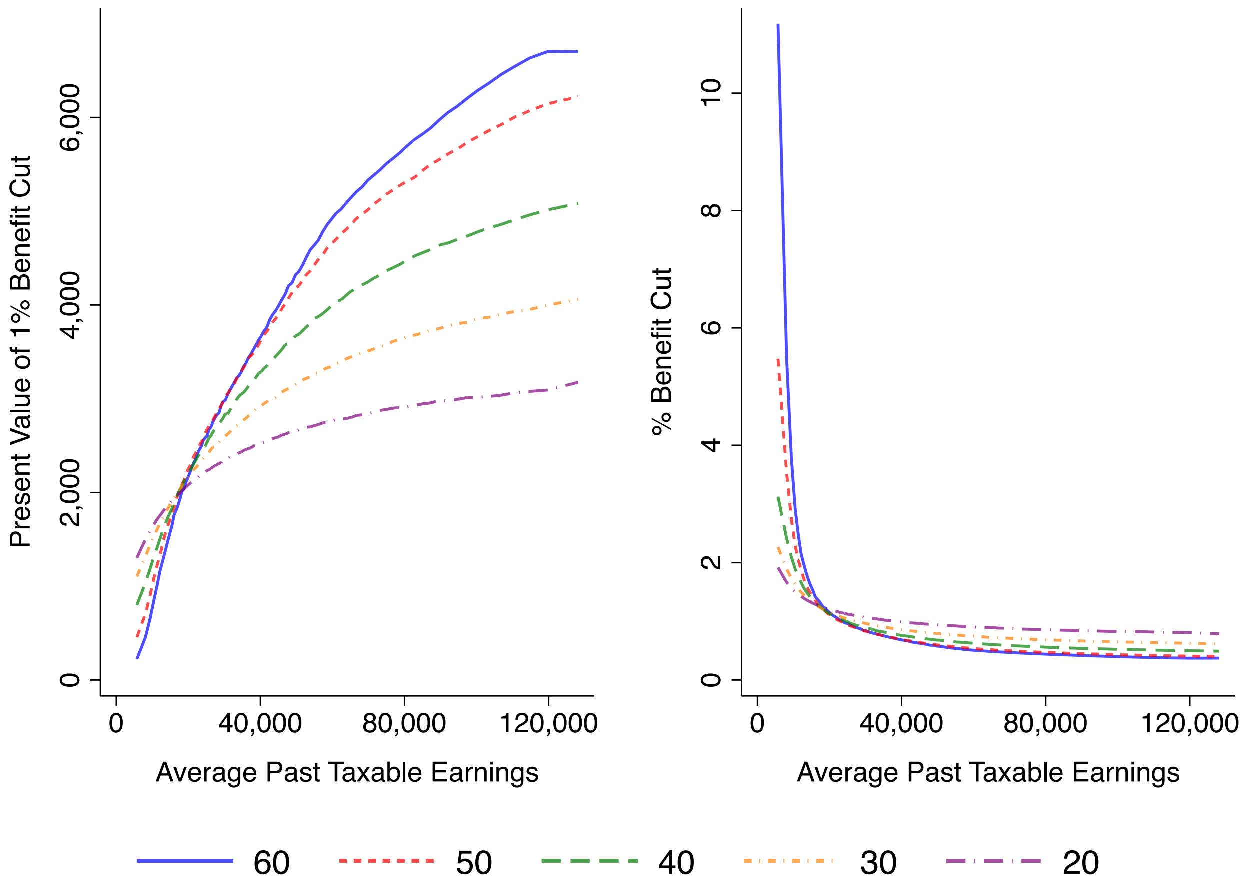

3 Relaxing housing liquidity constraints

We next quantify the magnitude of households' liquidity constraints and consider how they are exacerbated by COVID-19. We then document the extent to which a small cut in future Social Security benefits redresses them. We compare this approach to already legislated household support: penalty-free access to retirement accounts, stimulus checks, and a significant expansion of unemployment benefits.

3.1 Time to cash shortfall

We start by estimating how long it takes for households to run out of cash when they are on unemployment benefits. This depends on their liquid wealth, the generosity of unemployment benefits and their consumption level. We define the variable "Days to shortfall" as:

$$\text{Days to Shortfall} = \frac{\text{Liquid Wealth}}{\text{Consumption} + \text{Housing and Fixed Expenses} - \text{Unemployment Insurance}} \tag{3.1}$$

where the denominator represents daily expenditures minus insurance benefits. Two categories of households are more likely to run out of cash faster: *(i)* those with rent and mortgage payments and *(ii)* those with low liquid wealth-to-earnings ratios. We build these variables using the 2016 SCF, which provides detailed information on wealth, income, and expenditures by household.

First, we assume that unemployment insurance covers 50% of after-tax income. In reality, the benefit formula varies by state and takes into account workers' earnings and employment histories. However, our assumption is broadly consistent with the 45% average replacement rate reported by the Department of Labor for 2019. After-tax income is computed using the federal tax code and taking into account income, family composition and deductions (see Appendix A.2).

Housing and fixed expenditures include rent, mortgage payments, property taxes, co-op, and mobile home fees, car lease payments, as well as other loan payments. The details of these expenses is reported in the SCF (see Appendix A.3). We assume that consumption of other goods and services represent 60% of after-tax income. Our calibration implies an average saving rate of 6%, which matches the aggregate personal savings rate over the last 20 years (See FRED series PSAVERT).

Finally, liquid wealth is constructed as in Bhutta and Dettling (2018) and includes transactions accounts, certificates of deposit, mutual funds, stocks, and bonds. Using these estimates yields a proxy of Equation 3.1 that can be observed in the data, which is the measure we use for the remainder of the paper.

Figure 3 shows the share of households who can maintain their consumption up to 30, 60 or 90 days when unemployed, for each decile of marketable wealth. Unsurprisingly, wealthy households can afford to remain unemployed for longer. But the differences are stark: for those in the bottom three deciles of the marketable wealth distribution, more than 80 percent cannot cover three months of expenditures should they become unemployed. In the top decile, less than 5% face the same issue. Age largely explain this finding: workers who have just entered the labor force have yet to accumulate significant precautionary savings.

Figure 3: Time to shortfall by decile of wealth

This figure shows the fraction of households who can maintain their consumption on standard unemployment benefits for three, two, and one month before running out of cash, by decile of marketable decile.

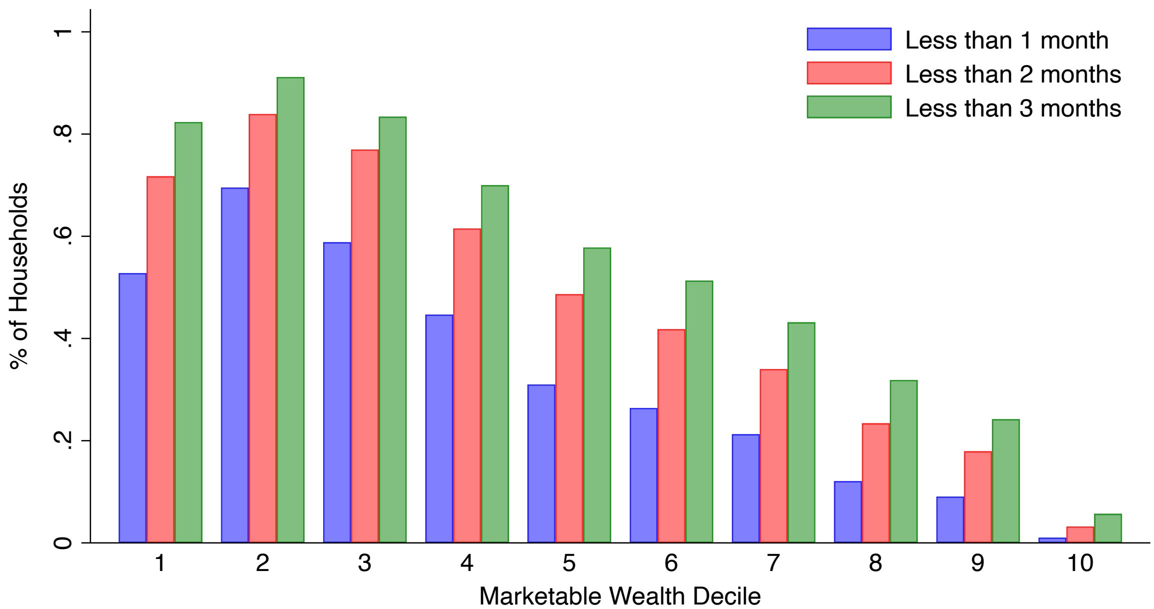

3.2 Impact of COVID-19 without intervention

In Figure 5, Panel A, we consider the implication of the counterfactual world in which aggressive stimulus efforts had not been undertaken to provide liquidity to households in need. We illustrate how our measure of days until cash shortfall is distributed throughout the population.

Importantly, we adjust the SCF sample weights such that our sample is representative of workers who have lost their jobs as a consequence of the novel coronavirus crisis (as of March 2020). This crisis disproportionately impacts particular industries (e.g. food services and entertainment) and is more likely to cause unemployment among young and less educated workers. Using data from the Current Population Survey (CPS) and the Bureau of Labor Statistics (BLS), we estimate the probability of becoming employed in that last six weeks as a function of industry, education, and age. We then adjust the SCF weights by multiplying them by the model implied probability of unemployment and dividing by the mean of this variable, a procedure detailed in Section A.5.

Figure 4, we show the fraction of households who would run out of cash after a given number of days if their income was reduced to standard unemployment benefits. Panel A shows this distribution for the entire population of working-age households whereas Panel B uses our adjusted weights to be more representative of households having claimed unemployment benefits in March 2020. Panel C shows the differences in between the two panels and illustrates the fact that unemployment induced by the pandemic disproportionately impacts households with low liquidity.

Overall, American households do not have sufficient liquid saving to weather the COVID-19 crisis. If displaced workers were only receiving unemployment benefits to supplement on average 50% of lost wages (as in normal times), more than 25% of working age households would not be able to meet their current expenditures after a month of unemployment, and 50% cannot last more than 75 days.

Figure 4: Effects of different weights on days to shortfall

This figure shows the number of days until the exhaustion of savings for households with at least one person aged 20 to 61 in the household in the event of unemployment when there is no intervention and under different weights. Days to shortfall is defined as liquid wealth divided by expenditures less income under employment insurance, as described in Section 3. Panel A shows the fraction of individuals (in percent) of the SCF that fall in each five-day days to shortfall bucket under the normal SCF population weights. Panel B shows the same thing, except using the weights that emphasize households that are more likely to become unemployed. Panel C shows the difference between these two. The probability of unemployment weights are derived using a logistic regression on the CPS data, where an indicator variable for new employed is the dependent variable and indicator variables for employment sector, race, education, and age. We take the expected probability of unemployment from this regression model and divide by its mean to obtain the unemployment multiplier, which we multiply by the SCF weights. This process is described in detail in Section A.5.

A. SCF Weights

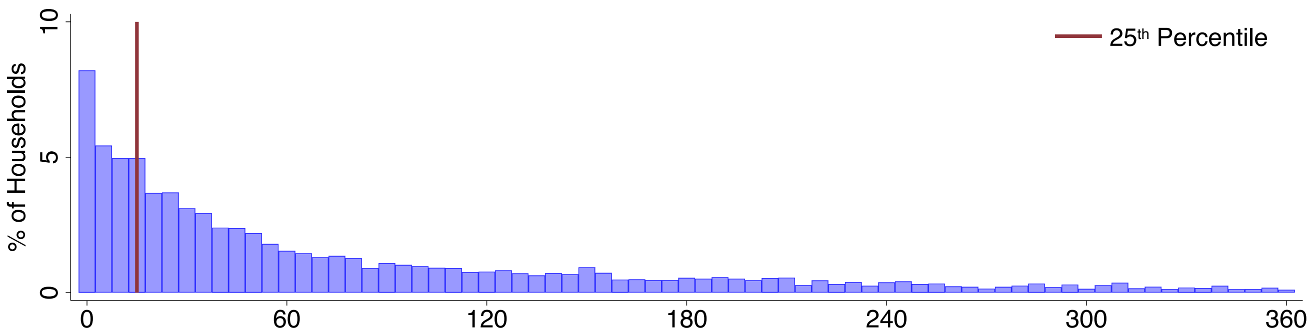

B.Probability of Unemployment Weighted

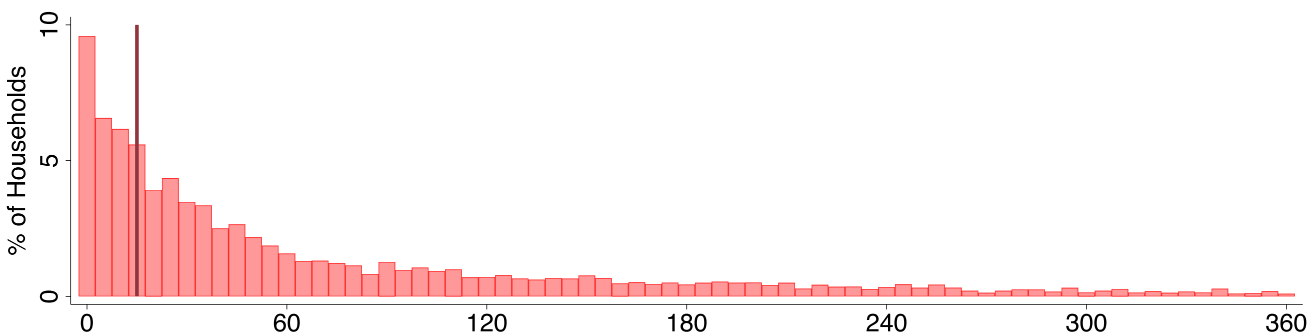

C. Relative Likelihood of Unemployment

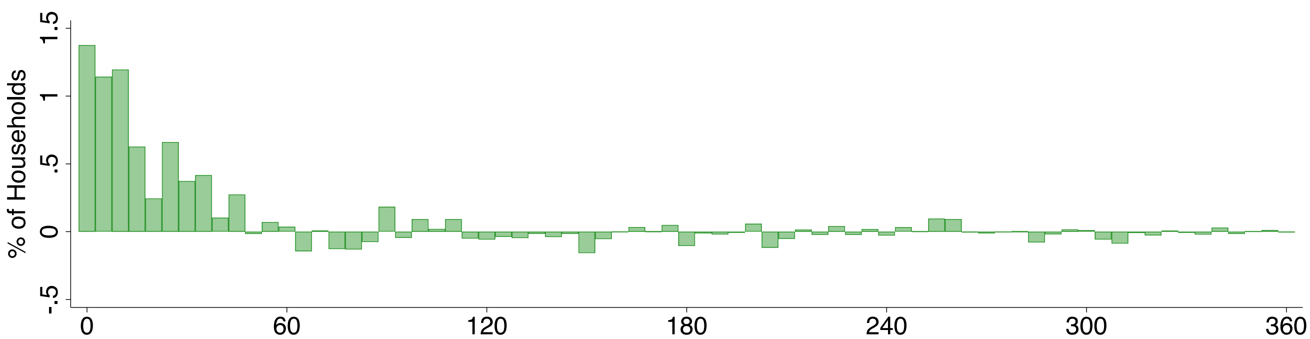

3.3 Tapping Social Security benefits

What would be the effect of allowing households to borrow against 1% of scheduled Social Security benefits? To analyze the quantitative effects of this policy, we must estimate the present value of Social Security benefits for each household in the SCF. To do this, we simulate a data set of 36 million individuals using the procedure described in Section 2.1, which contains age, sex, the present value of future benefits, average past taxable wage earnings, and current wage earnings. We then match the SCF to the simulated data by randomly assigning each individual in the SCF to a simulated outcome with the same age, sex, and wage income, a procedure which is detailed in Section A.4.

Early access to 1% of Social Security benefits are a boon to the liquidity of the most vulnerable households, as illustrated in Panel A of Figure 5. Under this policy, the bottom 25% of the marketable wealth distribution have an additional 88 days on average until they are no longer able to cover their current consumption, and the 25th percentile in terms of liquidity shortfall now has an additional two-and-a-half months of support, and the median is nearly six months. Even this small cut in benefits supports finances more consumption than most of the alternatives already legislated, as discussed below below.

3.4 Retirement accounts withdrawals

Penalty-free access to retirement accounts, as provided for by Congress' COVID-19 stimulus package, has a much more muted effect on household liquidity (Panel B). Under this policy, the bottom 25% of the marketable wealth distribution have an additional 29 days on average before they are no longer able to finance their consumption, and the 25th percentile has only 9 days of support, and the median is only 3 months.

This is because, unlike Social Security (which accrues evenly across the wealth distribution), the vast majority of workers made most vulnerable by the crisis do not have the funds in their retirement accounts to finance consumption today. Only half of workers have a retirement account, and in the bottom decile of marketable wealth, only 31% have non-zero retirement savings. Second, even for those who could gain liquidity by accessing retirement accounts, this would require liquidation of investment assets in the midst of a dramatic downturn (the S&P dropped by 16% in

March alone). While allowing penalty-free retirement account withdrawals does help the median family before our reweighting procedure[3], it does little for the poorest households most likely to be displaced by the pandemic.

3.5 Stimulus checks

Congress legislated a one-time issuance of $1,200 COVID-19 relief for all individuals earning less than $75,000[4]. The objective was to provide a quick bridge to families in dire financial straits because of nationwide lockdowns. This one-time stimulus was expensive: costing over $290 billion (Fink, 2020). It also boosts household liquidity by less than providing households just 1% of their accrued Social Security benefits early: the median household receives $2,200 from the stimulus, but $4,300 from a 1% cut in future benefits.

To be sure, policymakers would do well to sustain stimulus measures to support household consumption by continued fiscal expenditure. But the consequences for the budget deficit are significant: already legislated responses to the COVID-19 crisis have caused debt to balloon to over 100% of GDP (Swagel, 2020). There is widespread disagreement on the effect of large government debts and deficits in the economics literature (Blanchard, 2019, Rogoff, 2016). Given the low interest rate environment and lack of inflationary concerns, substantial focus on deficits at present is misplaced. But it is worth noting that funding household consumption through Social Security is budget neutral and allows for liquidity constraints to be relaxed for a few months at least without increasing the government debt burden, and with minute consequences for retirement savings.

3.6 Supplemental unemployment benefits

Since the onset of this crisis, unemployment benefits were increased by an extra $600 weekly through the end of July. As with relief checks, this measure is expensive: it costs around $260 billion. While the median household receives $600 per week, there are still a large plurality

[3]The median time to shortfall is nearly 8 months under this policy using the normal SCF weights.

[4]For heads of households, this number is increased to $112,500, and for couples filing jointly, the amount is $150,000. For people making over this the benefits are gradually phased out. Further, joint filers receive $2,400 in stimulus plus an additional $500 for each qualifying dependent. For more information on how this is constructed, see Section A.6.

of households that need more than this to avoid a shortfall. For those in the 25th percentile, this proposal provides an additional 110 days of liquidity, 20 days more than what 1% of Social Security benefits delivers.[5] But policy can be designed differently so that the Social Security approach delivers more to households, e.g. 2% of benefits today lengthens the time to cash shortfall by more than supplemental UI. Additionally, allowing workers access to Social Security wealth introduces fewer labor market distortions. For more than 50 percent of those displaced, unemployment benefits now exceed normal wages. There is a long literature the labor market impact of generous unemployment benefits, which can discourages workers from re-entering the workforce (Fredriksson and Holmlund, 2006, Lalive et al., 2006, Lentz, 2009).[6].

3.7 Other considerations and concerns

Despite policies already enacted to support households, many will find themselves unable to meet their financial obligations in the coming months. For those without access to credit, the result will be delinquency on obligations like rent and mortgage payments, that could result in eviction or bankruptcy. For those with access to credit, borrowers (many subprime) will take out loans at sky-high rates. An advantage of allowing workers to access a small portion of future Social Security benefits today is that this allows all households to benefit from the low interest rate environment. Workers will essentially be financing consumption needs today by taking a loan from themselves, at a near-zero interest rate. Those who need funds to tide them over will have them; and those who do not, can save.

One issue for policymakers to weigh is that the lump-sum payment of Social Security benefits will hasten the depletion of the Social Security trust fund by a few years. Thus policymakers will be forced to weigh entitlement reform, like increases in taxes or cuts for beneficiaries, sooner.

[5]This is an overstatement, since we estimate the impact on household liquidity of an additional $600 in UI for each month going forward. We also assume that all households that are unemployed receive these benefits, which appears to not be the case, as state UI programs are overloaded.

[6]In the current climate with nationwide lockdowns, concerns about distortions are overstated, and there is a host of evidence that generous UI during recessions is optimal (Crépon et al., 2013). Yet as the economy reopens, these issues may become more relevant. Already, those who have been provided loans from the government that will be forgiven only if they maintain payroll report struggling to re-employ workers (Morath, 2020)

Another potential concern with providing access to future Social Security benefits is that this decreases the funds they will have to finance consumption in retirement. Indeed, Social Security was introduced in the aftermath of the Depression to ensure the elderly did not die in poverty, and today for the vast majority of Americans these savings are their largest source of income after leaving the workforce. But our analysis demonstrates how a minuscule benefit cut can allow them to stay afloat for several months: even to provide workers $2,500 today, future benefits will be cut by on average 0.5% percent (Figure 2). And conceptually, the objective of any program of forced savings is to provide liquidity when households are in need. That is precisely why retirement accounts allow (with a penalty) for early withdrawal: so those in need can finance consumption today.

Figure 5: Days to cash shortfall under different policies

This figure shows the number of days until working-age households run out of cash in case of unemployment under different policies. Time to Shortfall is defined as liquid wealth divided by daily expenditures minus daily unemployment benefits, which we assume covers 50% of after-tax income. Each bin represents a 5-day increment and the graphs report the percentage of households who would run out of cash within these 5 days. The light blue bars in each graph show the no intervention case. Panel A refers to our policy proposal, in which everyone receives a check equal to 1% of the present value of expected benefits. Panel B shows the scenario in which can withdraw from their retirement accounts without penalty. Panel C shows the effect of giving \$1,200 checks to households using the policy outlined in the CARES Act. Panel D shows the results with extended unemployment insurance, also as specified in the CARES Act. The red, vertical lines represent the 25th percentile of the each time to shortfall variable.

A. Early access to 1% of Social Security Benefits

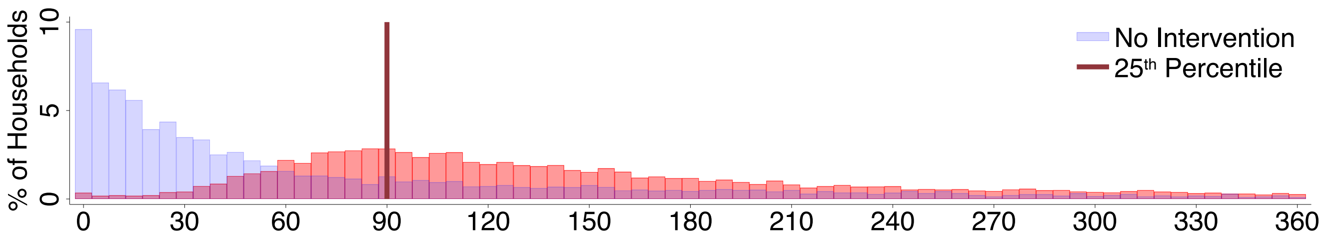

B. Early access to 401K accounts

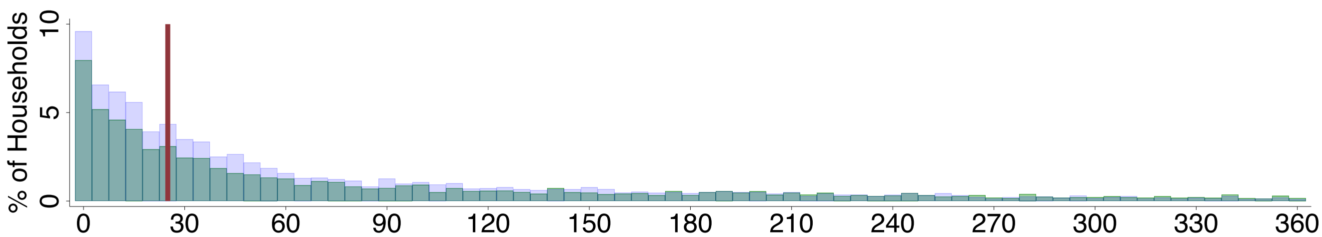

C. \$1,200 stimulus checks

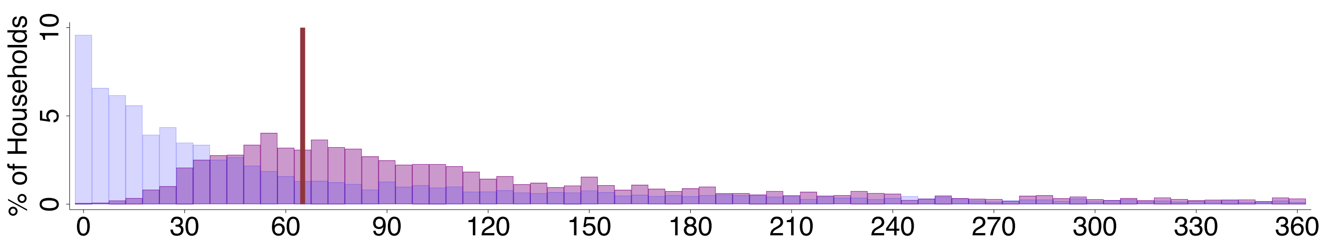

D. Supplemental unemployment insurance of \$600 per week

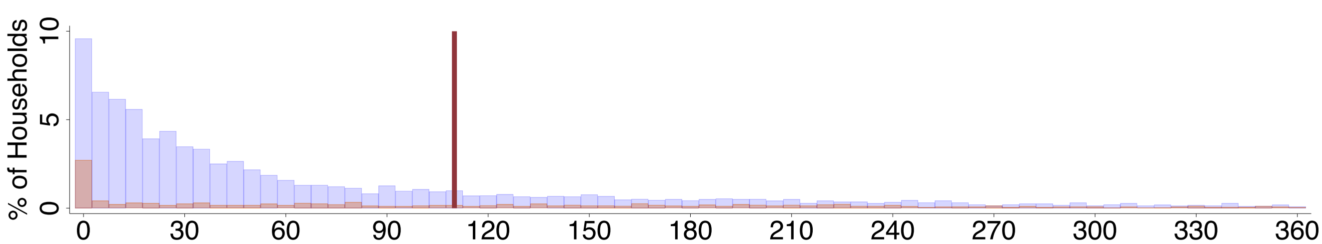

4 Conclusion

In the United States, Social Security wealth is designed to be illiquid to provide longevity insurance that safeguards retirees in old age. The result is that for most American workers, illiquid forced savings exceed the liquid wealth they have on hand to finance consumption shocks. But optimal illiquidity is time-varying, and in downturns like this current crisis, there is a case to be made for allowing workers to access their illiquid Social Security wealth.

We illustrate the potential of this approach by carefully computing the market value of workers Social Security benefits based on their age, earnings history, and estimated future earnings trajectories, adapting the approach of Catherine et al. (2020). We show that a minimal cut in scheduled Social Security benefits of just 1% is sufficient to finance household expenditure for two months. This provides more liquidity to households most vulnerable than alternative approaches already enacted, like penalty-free withdrawals from retirement savings accounts, and stimulus checks. It is also fiscally neutral and unlikely to introduce labor market distortions.

To be sure, Social Security benefits are the main source of income for the retired, and there is a case to be made for their expansion to better provide for the elderly. This proposal by no means pushes against that view, but rather suggests that the optimal savings rate for households is time-varying and should be allowed to adjust in times of crisis.

References

Abel, Andrew B., "Capital Accumulation and Uncertain Lifetimes with Adverse Selection," *Econometrica*, 1986, *54*, 1079–1097.

Baker, Scott R., Nicholas Bloom, Steven J. Davis, and Stephen J. Terry, "COVID-Induced Economic Uncertainty," *NBER Working Paper No. 26983*, 2020.

Bartik, Alexander W., Marianne Bertrand, Zoë B. Cullen, Edward L. Glaeser, Michael Luca, and Christopher T. Stanton, "How Are Small Businesses Adjusting to COVID-19? Early Evidence from a Survey," *NBER Working Paper No. 26989*, 2020.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, "Portfolio Choice over the Life-Cycle when the Stock and Labor Markets Are Cointegrated," *Journal of Finance*, October 2007, *62* (5), 2123–2167.

Beshears, John, James Choi, David Laibson, Brigitte C. Madrian, and William Skimmyhorn, "The effect of automatic enrolment on debt," 2019.

Bhutta, Neil and Lisa Dettling, "Money in the Bank? Assessing Families' Liquid Savings using the Survey of Consumer Finances," *FEDS Note*, 2018.

Biggs, Andrew G. and Joshua D. Rauh, "Funding Direct Payments to Americans Through Social Security Deferral," *Stanford University Graduate School of Business Research Paper No. 3580533*, 2020.

Blanchard, Olivier, "Public Debt and Low Interest Rates," *American Economic Review*, 2019, *109*, 1197–1229.

Board of Governors of the Federal Reserve System, "Report on the Economic Well-Being of U.S. Households in 2016," Technical Report 2017.

Catherine, Sylvain, "Labor Market Risk and the Private Value of Social Security," *Working paper*, 2019.

_ , Max Miller, and Natasha Sarin, "Social Security and Trends in Inequality," 2020.

Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora, "Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment," *The Quarterly Journal of Economics*, 2013, *128* (2), 531–580.

Eckstein, Zvi, Martin Eichenbaum, and Dan Peled, "Uncertain lifetimes and the welfare enhancing properties of annuity markets and social security," *Journal of Public Economics*, 1985, *26*, 303–326.

Fink, Jenni, "Stimulus Checks Cost $290 Billion. A Fraction of That Could have Changed Response to Coronavirus Outbreak, Experts Say," 2020.

Fredriksson, Peter and Bertil Holmlund, "Improving Incentives in Unemployment Insurance: A Review of Recent Research," *Journal of Economic Surveys*, 2006, *20*, 357–386.

Geanokoplos, John and Stephen P. Zeldes, "Market Valuation of Accrued Social Security Benefits," 2010, pp. 213–233.

Gormsen, Niels Joachim and Ralph S. J. Koijen, "Coronavirus: Impact on Stock Prices and Growth Expectations," *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-22*, 2020.

Gürkaynak, Refet, Brian Sack, and Jonathan Wright, "The TIPS Yield Curve and Inflation Compensation," *Finance and Economics Discussion Series*, 2008, *2008* (5).

Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song, "What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?," *SSRN Electronic Journal*, 2019.

Hosseini, Roozbeh, "Adverse Selection in the Annuity Market and the Role for Social Security," *Journal of Political Economy*, 2015, *123*, 941–984.

Lalive, Rafael, Jan van Ours, and Josef Zweimüller, "How Changes in Financial Incentives Affect the Duration of Unemployment," *The Review of Economic Studies*, 2006, *73* (4), 1009–1038.

Lentz, Rasmus, "Optimal Unemployment Insurance in an Estimated Job Search Model with Savings," *Review of Economic Dynamics*, 2009, *12*, 37–57.

Maurer, Raimond and Olivia S. Mitchell, "Evaluating Lump Sum Incentives for Delayed Social Security Claiming," *Public Policy and Aging Report*, 2018, *28*, S15–21.

_ , _ , **Ralph Rogalla, and Ivonne Siegelin**, "Accounting and Actuarial Smoothing of Retirement Payouts in Participating Life Annuities," *Insurance: Mathematics and Economics*, 2016, *71*, 268–283.

Morath, Eric, "Coronavirus Relief Often Pays Workers More Than Work," 2020.

Rogoff, Kenneth, *Progress and Confusion: The State of Macroeconomic Policy*, Cambridge: MIT Press,

Swagel, Phill, "CBOs Current Projections of Output, Employment, and Interest Rates and a Preliminary Look at Federal Deficits for 2020 and 2021," 2020.

INTERNET APPENDIX

A Data appendix

In this section, we give a detailed account of the data methodology employed in the main text.

A.1 Constructing the real yield curve

To obtain the real yield curve, we use estimates from Treasury Inflation-Protected Securities (TIPS) the Federal Reserve Board of Governors, based on the methodology in Gürkaynak et al. (2008).[7] These data provide real, annualized zero coupon yields for government securities from 2 to 20 years. To obtain the one yield yield, we use the annualized two year rate. To obtain longer horizon estimates of the yield curve, we take the 19-to-20 year forward rate, given by $f_{t+19\to t+20} = \frac{1+r_{t,t+20}}{1+r_{t,t+19}}$ to be the long-run real interest rate, and iterate this rate on the 20-year yield to extend the yield curve. This is given mathematically by $1+r_{t,t+20+h} = \left((1+r_{t,t+20})^{20}(f_{t+19\to t+20})^h\right)^{\frac{1}{20+h}}$, which we use to extend the real yield curve out to 100 years. We use the most recent data available for this calculation, which as of this writing is from April 9th, 2020, which is shown in Figure A.1.

[7] The data can be found at https://www.federalreserve.gov/data/tips-yield-curve-and-inflation-compensation.htm.

Figure A.1: Extended Real Yield Curve

This figure shows the real yield curve used for the calculation of the present value of future Social Security benefits. Values are in annualized spot rates. The one year ahead is set equal to the annualized two year rate. Rates beyond 20 years are obtained by iteratively applying the 19-to-20 year forward rates to the spot rate, as described in Section A.1.

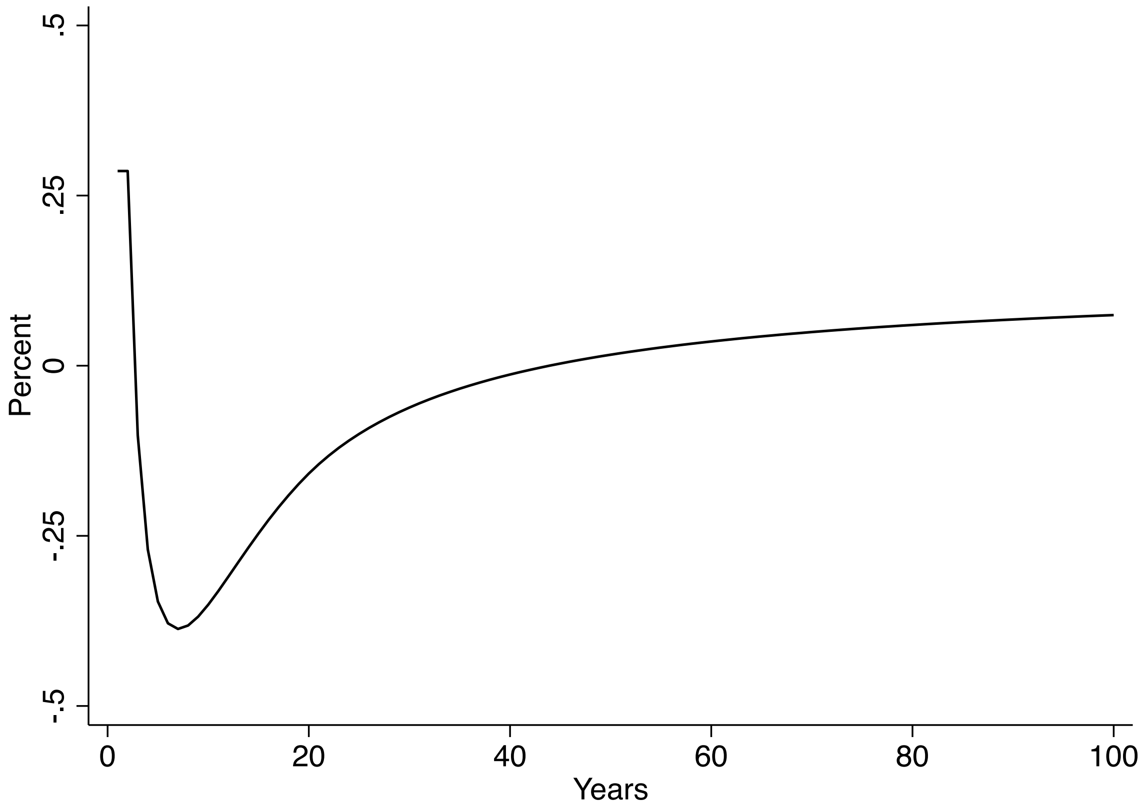

A.2 Estimating taxes in the SCF

While there is insufficient data to arrive at the exact tax payment a household makes using the SCF data, a reasonable estimate can be achieved. To do this, we apply the tax code in a straightforward way to arrive at after-tax income. To do this we start by deducting personal exemptions, the standard deduction, and interest payments on student loans.

In calculating these, we make the simplifying assumption based on variable X5746, which asks about the filing behavior of each household. Namely, for married couples that file jointly or separately, we use the tax brackets for married, joint filers, and, for everyone else we use the head-of-household tax bracket. The first assumption is made because we do not observe income for each member of the household, so treating them as joint filers is a requirement and likely only overstates taxes for those who are relatively wealthy, as wealthier couples have more to gain from filing separately. The second assumption is made because we do not observe how single households file. This is assumption will likely understate the tax burden for individuals.

Personal exemptions in 2016 are equal to $4,050 for each qualifying dependent. To arrive at this number we multiply $4,050 by the number of children in the household. Further, we phase out these exemptions using IRS rules, namely for each $2,050 above $285,350 and $311,300 for heads of households and joint filers, respectively, we subtract 2% of the exemption until the full amount is exhausted. Similarly, we apply the Standard Deduction for all households, which is $9,300 and $12,600 for heads of households and joint filers, respectively, using the same phase out procedure via IRS policy. Finally, we allow for up to $2,500 of student loan payments to be deducted annually.

We subtract these deductions from total income in the SCF and then apply the appropriate tax rate based on the progressive tax brackets used by the IRS in 2016, taken from the Tax Foundation. These give us annual estimates for taxes paid during the year. We then subtract this from income to arrive at after-tax income.

A.3 Calculation of fixed expenses in the SCF

The Survey of Consumer Finances (SCF) contains information about loan expenditures that we use to measure a household's proximity to a cash shortfall, which we define as the number of days until the household is unable to meet current obligations based on current liquid savings. Our definition of fixed expenses are those which cannot be changed or renegotiated easily, and are therefore unlikely to change in a crisis setting. These include regular living expenses like rent and fees for apartments, houses, condominiums, and mobile homes, mortgage payments, property taxes, car lease payments, and non-mortgage loan payments.

Some of the variables we use are included in the SCF raw files and are not present in the cleaned extracts produced by the Survey of Consumer Finance division at the Federal Reserve Board of Governors. Payments for rent are in the SCF raw data file under variable X708, which can be adjusted into a monthly variable by using the frequency of payment variable X709. In fact, for each variable we discuss, there is an associated frequency variable which is always one plus the original variable number. Mobile home payments come from three different variables. The first is X602, which is the cost of renting the mobile home when the respondent owns the site, but not the home. Variable X612 corresponds to respondents who own the mobile home and rent the site, and variable X619 corresponds to people who rent both.[8] Co-op fees are also given in the SCF by variable X703, and property taxes by X721. Finally, we add in car lease payments which are given by variables X2105 (first car lease, if applicable) and X2112 (second car lease, if applicable).

These variables are combined with the TPAY variable from the cleaned SCF extract. The TPAY variable represents all monthly loan payments the household makes which is equal to the sum of MORTPAY, which is total

[8] There are no non-zero observations for X602 in the 2016 survey. There are 415 non-zero observations for variable X612, and 360 non-zero observations for variable X619.

mortgage debt payments, CONSPAY, which is total non-mortgage non-revolving consumer debt, and REVPAY, which is total revolving debt excluding home equity lines of credit (HELOCs). MORTPAY includes all mortgage payments for home mortgages, mortgages for other residential properties, payments on land contracts, payments on certain types of lines of credit. CONSPAY includes payments on auto loans, student loans, installment loans, margin loans, loans against insurance policies, other loans, and loans against pension plans. The REVPAY variable includes credit card payments and other lines of credit not included in MORTPAY. For the median person in the SCF, these payments make up roughly 9.9% of before tax income.

A.4 Merging the simulated data to the SCF

To calculate the net present value of future benefits in the SCF, we must merge the simulated data to the actual data. To do this, we generate a sample of 36 million individuals using the simulation where the sample consists of age, sex, current wage income, AIYE, and the present value of future Social Security benefits. We then round the wage income variable to the nearest $2,500 and then generate an identifier for each observation within each age, sex, and current wage bucket.

To merge this data with the SCF, we must split household wage earnings between people in multi-earner households. To do this, we use data from the SCF on self reported wages for each earner in the household on wage income. However, these self reported wages will often differ from the Internal Revenue Service, Form 1040, Box 7 income reported by the SCF in the cleaned extracts. Therefore, we use these information from the self reported wage data to ascertain how wages are split within the household. More detail on this procedure is given bellow in Section A.4.1. We then round these split wage data to the nearest 2,500 and randomly generate an identifier to be merged with the simulated data. This is in essence treating the present value of Social Security for each SCF respondent as a random draw from the simulation, conditioning on current wage income. From there, we take 1% of the combined present value of future benefits as the check that the household will receive in our policy.

A.4.1 Splitting household wage income in the SCF

To split the WAGEINC variable from the cleaned SCF extract between household earners we rely on data from the raw SCF files. In particular, we use variables on self reported wages, which are X4112, X4509, X4712, and X5109 which are the wage earnings on the first and second (if applicable) jobs of the first and second members of the household, respectively. These are then adjusted to annual frequencies using variables X4113, X4510, X4713, and X5110. For single earner households, splitting the wage is easy; we assign 100% of the wage to the single earner. For dual earning households, we some together the total wages for each member and assign to each person the corresponding fraction of WAGEINC. For example, if if self-reported earnings of $75,000 and $25,000 for the first and second persons in the household, and the IRS Form 1040, Box 7 income is $80,000, then the first person will be assigned $60,000 and the second person $20,000.

A.5 Re-weighting the SCF by likelihood of unemployment

In our days to shortfall calculations in Section 3, we alter the nationally representative weights in the SCF to overweight respondents who work in sectors that are most likely to be unemployed due to the COVID-19 crisis, and young and less educated workers. To do this we rely on data from the SCF raw data files and the CPS from the BLS.

The SCF contains data on the industry of employment of each respondent. However, for privacy purposes, in the public data, the detailed industry information is aggregated into 7 sectors which broadly correspond to the overarching sectors in the Census Bureau's industry classification system. For dual earner households, this information is available for each person, and is given by variables X7402 and X7412 for the first and second members of the household, respectively. Also, for these households, there will be two re-weighting variables, one for each person. To aggregate this to the household level, we income weight the re-weighting variables to come up with an aggregate household weight multiplier.

To calculate these reweighing multipliers, we use the CPS data. The CPS data allows us to observe characteristics of the recently unemployed such as their age, sex, level of education, industry of employment, and occupation. Using this data, we match the detailed industry classifications in the CPS to the more aggregated classification available in the public SCF files. We then identify all respondents have become unemployed in the last 6 weeks, excluding new entrants, and run a logistic regression of this indicator variable for new unemployment on dummy variables for level of education[9], five-year age cohort, race, and SCF industry. The model we estimate is of the following form

$$p_i = \frac{1}{1 + \exp\{-(\beta_{i,\text{Race}} + \beta_{i,\text{Education}} + \beta_{i,\text{Industry}} + \beta_{i,\text{Age}})\}} \tag{A.1}$$

where each i is a distinct race, education, industry, and age combination.

We then calculate the reweighting multiplier by taking the expected probability of new unemployment for each age, industry, race, and education category, dividing by the mean. For example, a 20 year old working in the Wholesale and Retail Trade, Bars, and Restaurants sector with some college has an expected probability of becoming newly unemployed of approximately 5%. The mean expected probability of becoming newly unemployed in the sample is around 2% for the March 2020 CPS. This means that the reweighting multiplier is 2.5.

We then merge these reweighting multipliers by industry, age, and educational attainment. In dual earning households, this gives us two multipliers to apply. To determine the household multiplier, we weight the multipliers of each person by their relative contribution to household income, using the same approach as in Section A.4.1.

[9]For this, we map the CPS educational attainment variable into the EDCL variable from the cleaned SCF extract.

A.6 Calculating the value of other policies

In the main text, we examine three policies: 1) early access to retirement savings, 2) $1,200 stimulus checks as in the CARES Act, and 3) supplemental unemployment benefits of $600 per week. To calculate the change in days to shortfall under early access to retirement savings, we add the RETQLIQ from the cleaned SCF extract to the variable for liquid wealth (which is the sum of the variables LIQ, CDS, NMMF, STOCKS, and BOND).

To calculate the effect of $1,200 stimulus checks, we apply the formula from the CARES act to our estimate for taxable income. This means that every head of household making under $112,500 (for simplicity, we assume all single households file as heads of households) or every joint filer making less than $150,000 gets the full amount of the stimulus, equal to $1,200 for single households, $2,400 for two person households, with an additional $500 for each qualifying dependent under 17 years of age. These checks are phased out by $5 for every $100 a couple makes beyond this amount until they set to zero for single, head of household earners making more that $136,500, and joint households making more than $198,000.

Finally, we incorporate the supplemental unemployment insurance by adding $7,200 dollars to annual taxable income to come up with additional taxes under this proposal, as the supplemental benefits are taxable, but only last until the end of July. Next, we add $2,400 to monthly after-tax income in the event of unemployment. This overstates the value of this program, as the additional benefits will expire at the end of July.