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).[13] These data provide real, annualized zero coupon yields for government securities from 2 to 20 years. To obtain the one year 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.

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

[13] This data is provided by the Federal Reserve, and available here.

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.

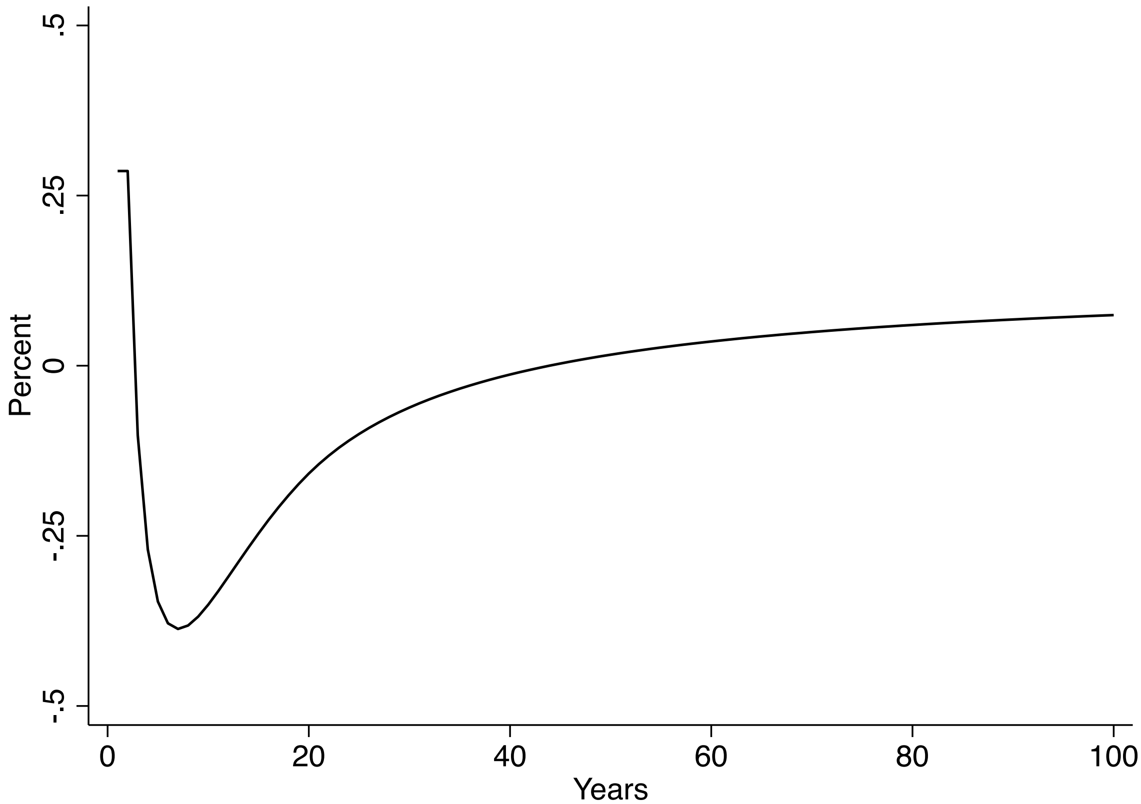

rules, namely for each \$2,500 above \$285,350 and \$311,300 for heads of households and joint filers, respectively, we subtract 2% of the exemption until the full amount is exhausted. Similarly, we apply the Standard Deduction for all households, which is \$9,300 and \$12,600 for heads of households 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.[14] 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 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 household with at least one member aged 20 and 61 in the SCF, these payments make up roughly 26.9% of before tax income and on average 34.4%.

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

[14] 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.

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[15], 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. Figure A.2 shows the days to shortfall results in the no intervention case under the SCF weights compared to our reweighting procedure. Under the reweighted data, a greater weight is placed on people least able to weather the COVID-19 crisis.

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

Figure A.2: 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

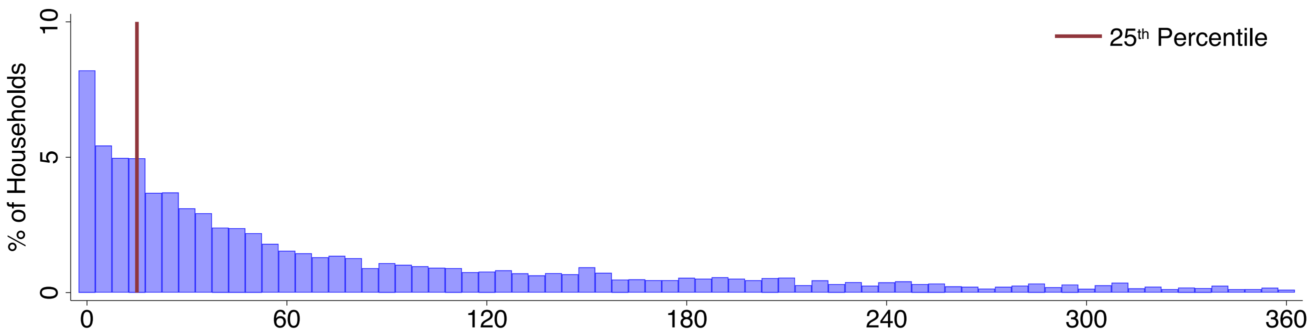

B.Probability of Unemployment Weighted

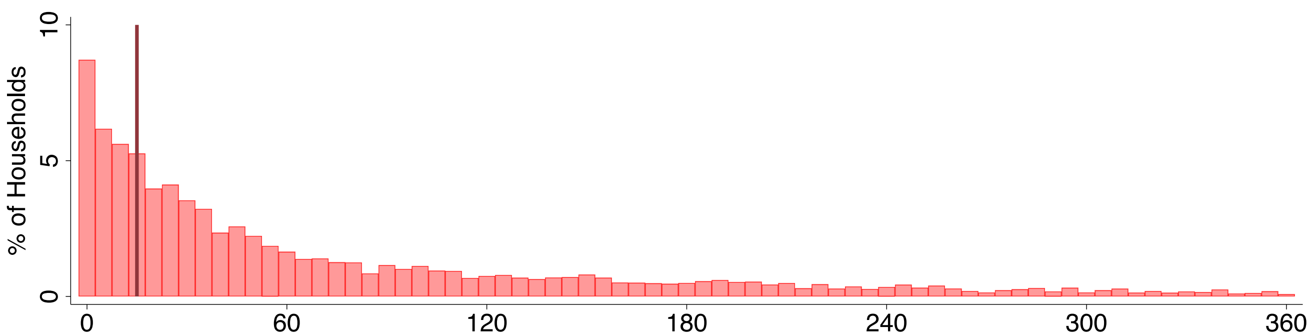

C. Relative Likelihood of Unemployment

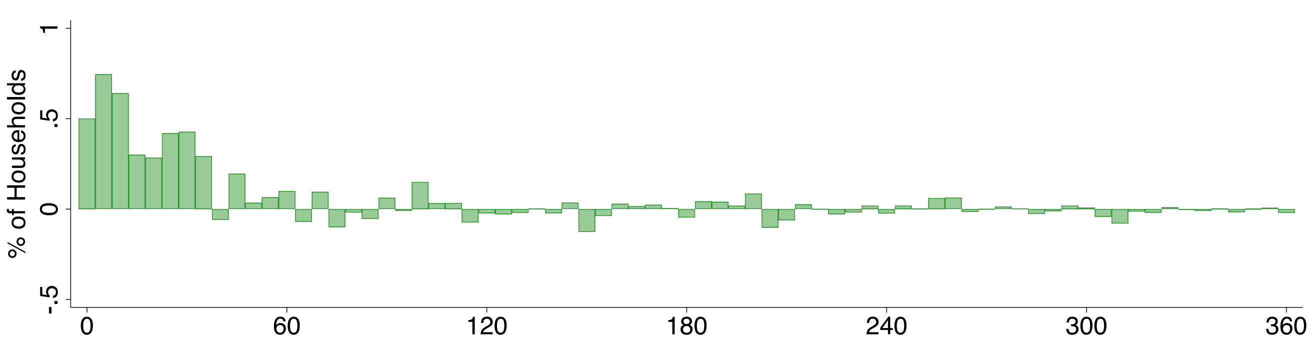

A.6 Calculating the value of other policies

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). Further, we provide the number of households who do not have access to these types of funds by marketable wealth decile. This is shown in Figure A.3. Note that households in the bottom of the wealth distribution who are most likely to be affected by COVID-19 also have the highest likelihood of having no retirement savings.

Figure A.3: No retirement savings across the wealth distribution

This figure shows the fraction of households in the SCF with no retirement savings across the marketable wealth distribution. Retirement accounts are taken from the SCF extract quasi-liquid retirement accounts variable (RETQLIQ).

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. The program is set to only run for three months though, so we calculate the time to shortfall as the minimum of the time to shortfall under a permanent addition of $2,400 to 50% of after-tax income or the time to shortfall under 50% of after-tax income plus 91 days.

In additional to different value to households, these programs also have different budgetary implications. Ta-ble A.1 shows the total outlays, number of eligible recipients, per person benefit, budgetary impact, and liquidity provided to the bottom 25% in terms of time to shortfall.

Table A.1: Costs of different proposals

This table shows the current outlays, per capita benefits, changes in government liabilities, and amount of liquidity provided for four programs discussed in the main text. Outlays for the expanded UI program and stimulus payments are taken from the Center for a Responsible Federal Budget, and the outlays for the Social Security proposals are estimated using the SCF by multiplying the mean benefit times the number of recipients. The number of people eligible for expanded UI is the April 2020 level of unemployment, the number eligible for stimulus checks is taken from a press release from the United States Treasury Department, the number eligible for the Social Security program is the sum of weights for respondents aged $20-61$, where the weights are doubled for two person households to arrive at the per capita figure. Per person benefits are calculated by dividing the first column by the second column. For the expanded UI and stimulus programs, the change in government liabilities is equal to the outlays; for the Social Security programs, this is zero as the current outlays are offset by reductions in future benefit payments. The liquidity provided is calculated by the authors, the procedure for which is outlined in Section 3.

Proposal	Outlays Today	People Eligible	Per Person	Government Liability	Liquidity Provided (25th Percentile)
Expanded UI	$260B	23M	$11,304	$260B	105 days
Stimulus payments	$290B	195M	$1,487	$290B	65 days
1% of SS benefits	$416B	144M	$2,884	$0	94 days
$2,500 advance on SS	$323B	144M	$2,243	$0	92 days

A.7 Differences in mortality rates

Our approach values future Social Security benefits based on average life expectancy across the population. This ignores the covariance between wealth and life expectancy: wealthier individuals life longer, and this has implications for the valuation of Social Security (Catherine et al., 2020). Simply stated: the value of Social Security benefits for those with below-average life expectancy will fall, and for the wealthy with above-average life expectancy, it will rise.

Depending on the policy chosen, differential mortality will have varied implications. Should a 1% cut in benefits be pursued, those at the bottom in the distribution will receive more than 1% of expected benefits unless mortality differences are factored into benefits calculations, which seems politically untenable. If instead the more adminstrable lump-sum payment is pursued, then the liquidity infusion households receive today will be unchanged, but the implications for future benefits will be more pronounced for the bottom of the wealth distribution.

B Alternative assumptions

B.1 Retirement age

To arrive at a value of Social Security, we assume that individuals choose to retire and begin claiming benefits once they reach full retirement age. In practice, claimants can retire early and receive benefits below the full retirement level; or retire late and receive benefits above the full retirement level. In practice, though, early or delayed retirement does not impact Social Security's value, because changes to benefits are priced in an actuarially fair way. A simple example is illustrative: individuals who retire one year before full retirement receive 6.7% lower benefits, but these benefits are provided for an extra year (or 5.1% more checks on average). These two effects

offset each other. It is certainly possible that individuals will make retirement choices that result in their benefits being less, or more, than the full retirement age baseline: There may be adverse selection problems, so people who know their life expectancy is less (greater) than average will be want to retire early (later). From the point of view of the government – and our valuation exercise – since cuts for early benefits/benefits for late claiming are fairly priced on average, the issue is moot (Munnell and Chen, 2019).

B.2 Macroeconomic variables

Here, we provide our results under various macro-economic assumptions. In Figure B.1, we show the percentage of benefits that will be eroded (added) under more (less) severe macroeconomic assumptions. Here, the benefits adjustment for alternative wage growth rates is calculated relative to the baseline case used in the main text. This is done by applying

$$\%\ Adjustment(age) = \frac{(1+g_{alt})^{60-age}}{\Pi_{t=1}^{60-age}\Big(1+g(t)\Big)} - 1$$

where g_{alt} is the alternative constant growth rate and $g(t)$ is the SSA wage growth rate for horizon t. These results are located in Panel A. Similarly, we repeat this exercise for the equity premium, where the adjustment becomes

$$\%\ Adjustment(age) = \left(\frac{r+\beta(age)(\mu-r)_{alt}}{r+\beta(age)(\mu-r)_{base}}\right)^{60-age} - 1$$

where $\beta(age) = \left(1 - \frac{\phi}{\kappa}\right)\left(1 - e^{-\kappa(60-age)}\right)$. These results are located in Panel B.

Figure B.1: Benefit adjustment under different assumptions

This figure shows shows the percent benefit adjustment for each age under different macroeconomic assumptions. Panel A reports results for different assumptions about wage growth. Panel B reports different assumptions about equity premia. For wage growth, we assess alternative paths of 0%, 0.5%, 1.5%, and 2% constant wage growth relative to the baseline wage growth projections in the SSA report from the Office of the Chief Actuary. For equity premia, we assess alternative paths of 4%, 5%, 7%, and 8% constant equity premia relative to a 6% equity premium in the baseline case.

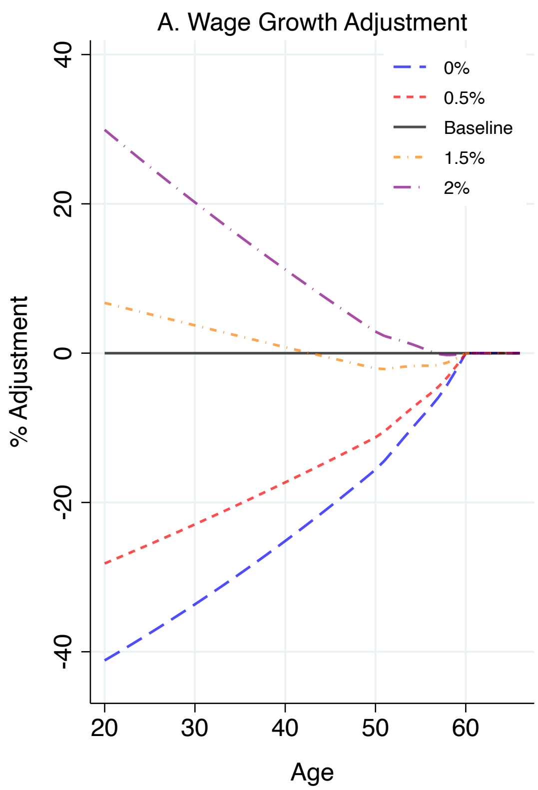

We also perform the simulation exercise from Section 2 under two sets of extreme assumptions. The first set of assumptions is the low growth, high equity premium scenario shown in Figure B.2. Here, report the size of 1% of future simulated benefits (Panel A) and the percent of benefits that would need to be cut to obtain \$2,500 (Panel B). For this exercise, we assume the equity premium will be a large 8% (as opposed to it's historical value of approximately 6%) and that real wage growth is zero. Even under this extreme scenario, many people will be provided substantial liquidity, with individuals making \$40,000 per year on average receiving an average check of at least \$1,000, regardless of age.

Figure B.2: Price of early Social Security check – Low growth, high equity premium

This figure shows the relationship 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. For this exercise, we assume an equity premium of 8% and wage growth of 0% for all simulated data.

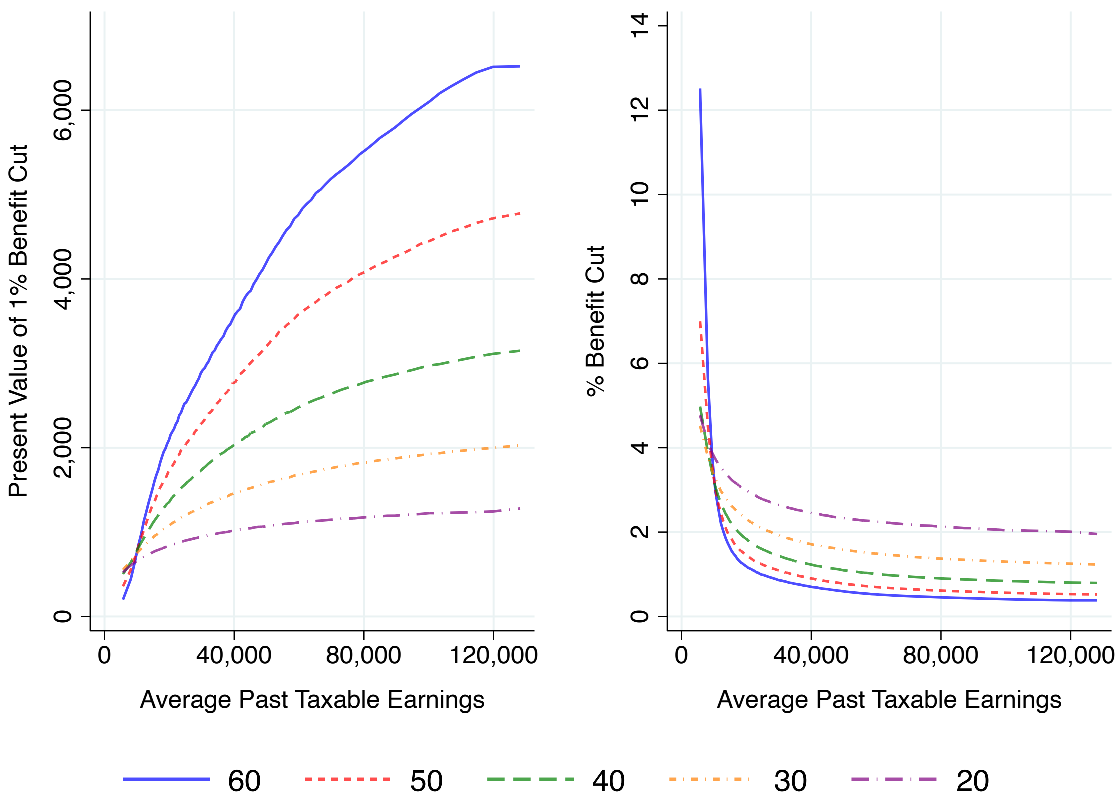

The second set of assumptions is the high growth, low equity premium scenario shown in Figure B.3. Here, report the size of 1% of future simulated benefits (Panel A) and the percent of benefits that would need to be cut to obtain $2,500 (Panel B). For this exercise, we assume the equity premium will be a small 4% (as opposed to it's historical value of approximately 6%) and that real wage growth to be 2%. Under this more generous scenario, people will be provided with even more liquidity, with individuals making $40,000 per year on average receiving an average check of at least $1,000, regardless of age.

Figure B.3: Price of early Social Security check – High growth, low equity premium

This figure shows the relationship 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. For this exercise, we assume an equity premium of 4% and wage growth of 2% for all simulated data.

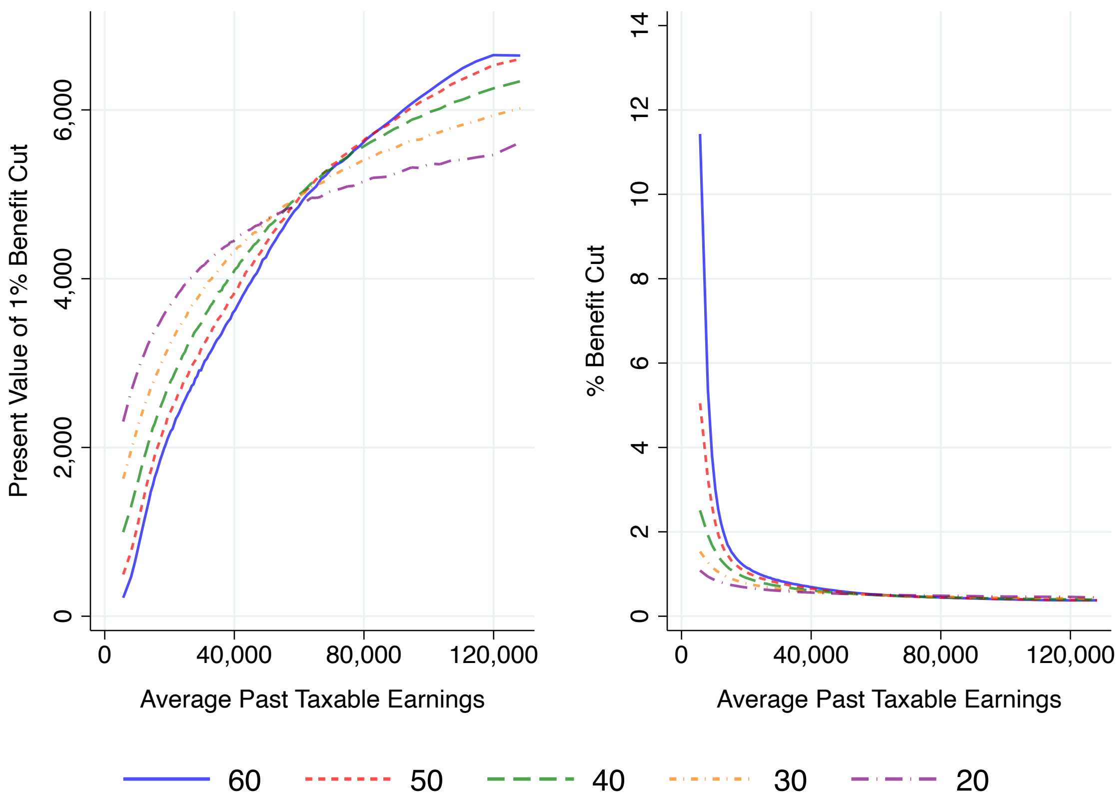

B.3 Policy risk

Related to the robustness with respect to the macroeconomic assumptions is the robustness of the results to the funding shortfall that Social Security will experience in the next 10-15 years. To see how a funding shortfall would impact our results. We apply the shortfall estimates from the SSA Office of the Chief Actuary 2019 report using a uniform benefits cut, such that incoming tax revenue exactly covers outgoing benefits. The SSA makes a range of projections along this dimension, so for the purpose we will use the most extreme set of assumptions that they provide, which correspond to roughly a 40% cut in benefits after 10-years. The results for this exercise are provided in Figure B.4. Event under this extreme scenario, many people will be provided substantial liquidity, with individuals making $40,000 per year on average receiving an average check of at least $1,500, regardless of age.

Figure B.4: Price of early Social Security check – Funding gap

This figure shows the relationship between benefit cuts and check size as a function of workers' age and the average past taxable earnings under the Social Security Administration's most in the most extreme funding gap case. 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. Data on the funding gap are taken from the SSA annual reports.

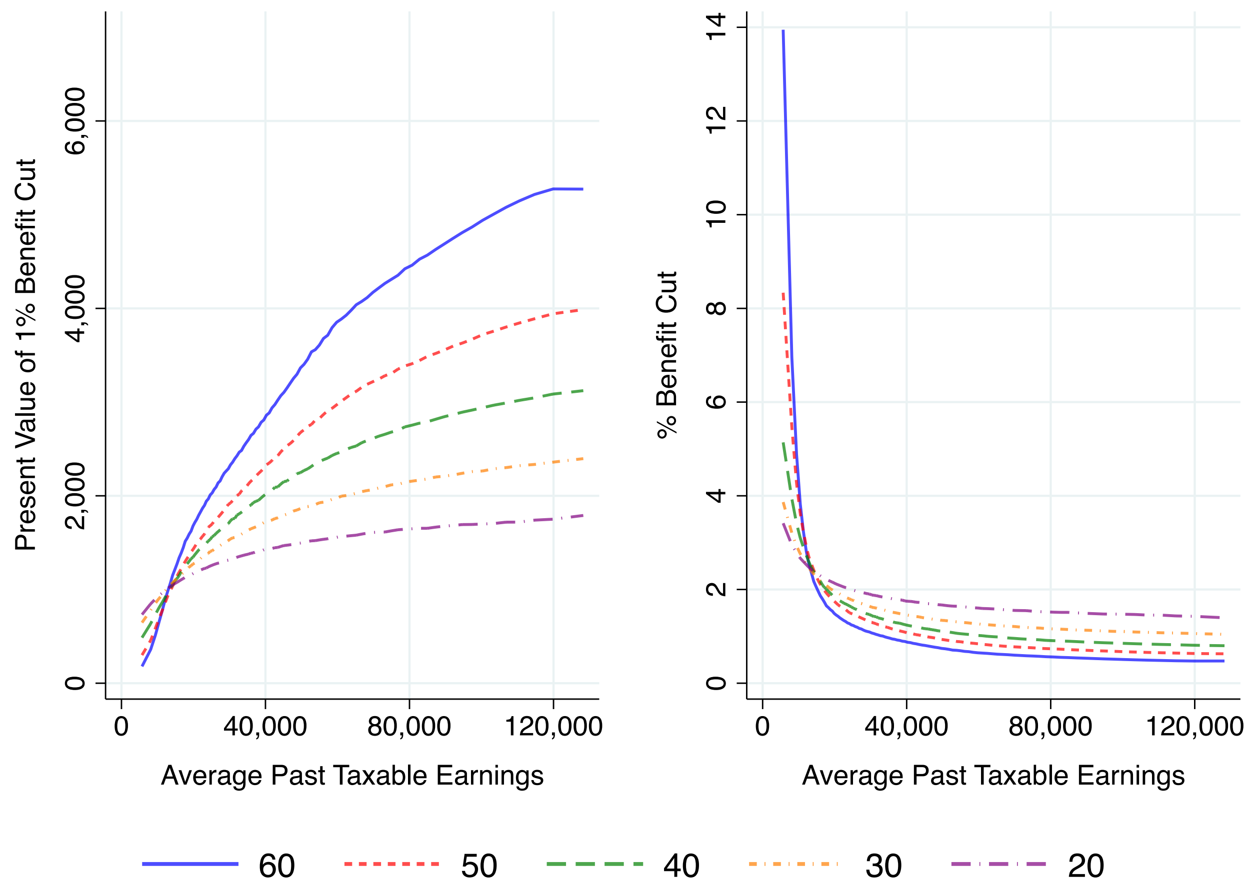

B.4 Savings rate by income quintiles

In the main text, we assume that discretionary consumption is 60% of after-tax income. Here, we discuss the validity of this assumption, and see how conclusions would be changed by relaxing it. Figure B.5 shows the savings rates implied by our measure by income quintile relative to the third income quintile. Note, that savings rates are rising sharply, consistent with the results of Dynan et al. (2004). Further, the slope of this line is also broadly consistent with Dynan et al. (2004). Note that the variation in savings rates here are driven by differential amounts of fixed expenses, the construction of which is outlined in Section A.3.

Figure B.5: Savings rates by income quintiles

This figure shows the implied average savings rate for each income quintile in the SCF under our methodology. Each estimate is relative to the 3rd quintile savings rate. Standard errors are bootstrapped.

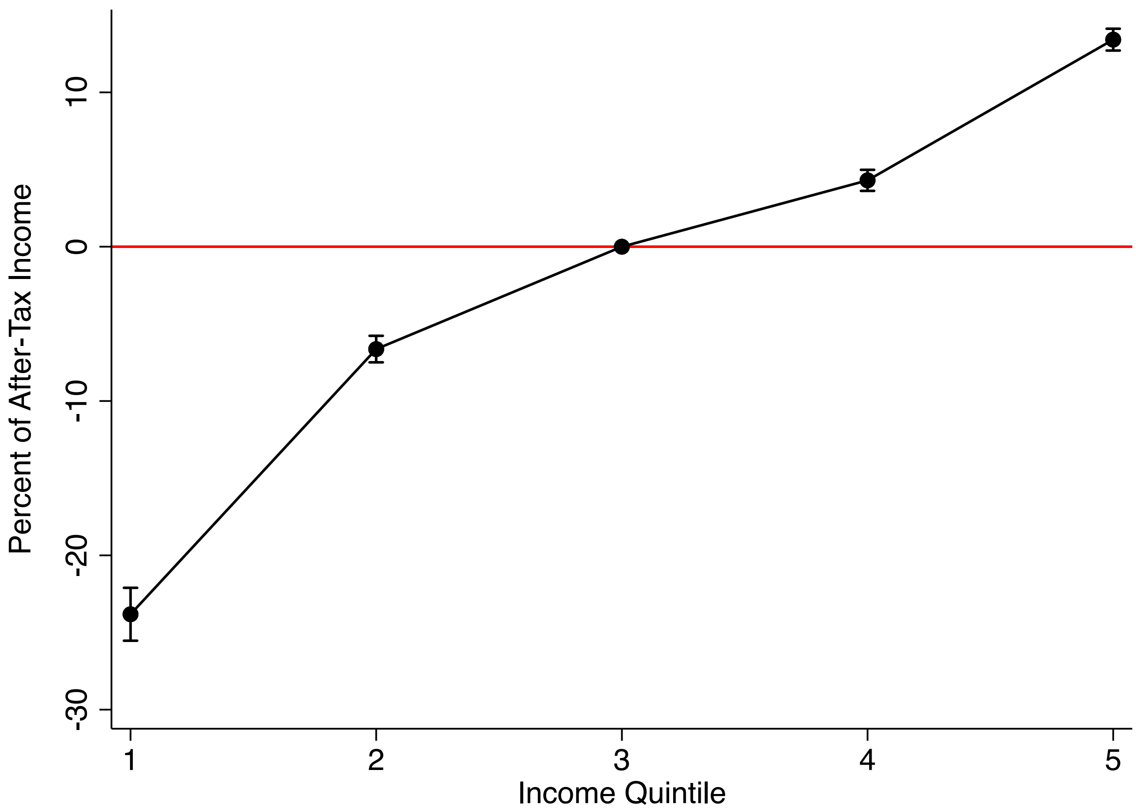

However, relaxing this assumption such that discretionary consumption is 50% of after-tax income does not qualitatively change the results. Figure B.6 shows the time to shortfall measure under this assumption. Now, the policy of early access to one percent of Social Security benefits provides 120 days of liquidity to the bottom 25th percentile, as opposed to 35 days for early access to retirement funds, 85 days under the $1,200 stimulus checks, and 110 days under extended unemployment benefits. If anything, these results are more strongly in favor of the early access to Social Security benefits, than the exercise in the main paper.

Figure B.6: Days to cash shortfall under different policies – Lower consumption

This figure shows the number of days until working-age households run out of cash in case of unemployment under different policies when consumption is cut by 10% in response to the pandemic. 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 workers 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 of $600 in extra unemployment insurance, as provided for by 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

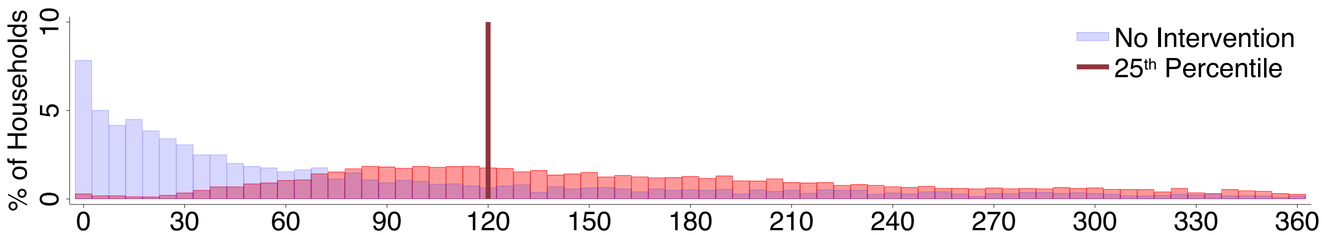

B. Early access to 401K accounts

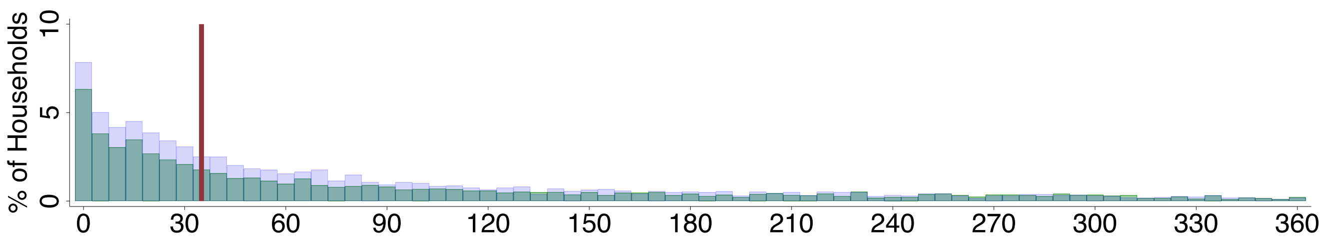

C. $1,200 stimulus checks

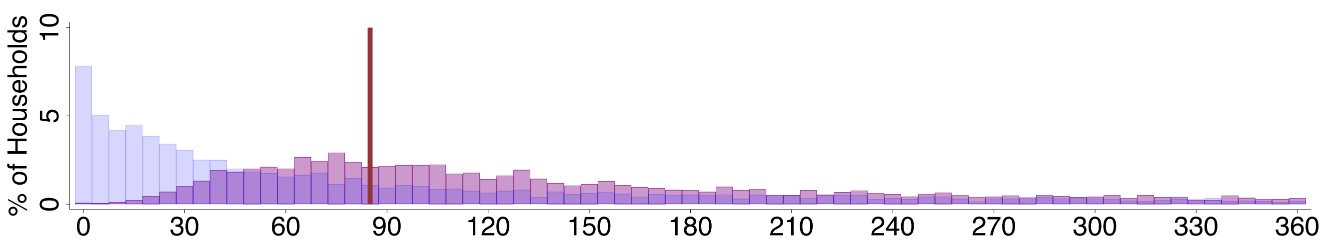

D. Supplemental unemployment insurance of $600 per week

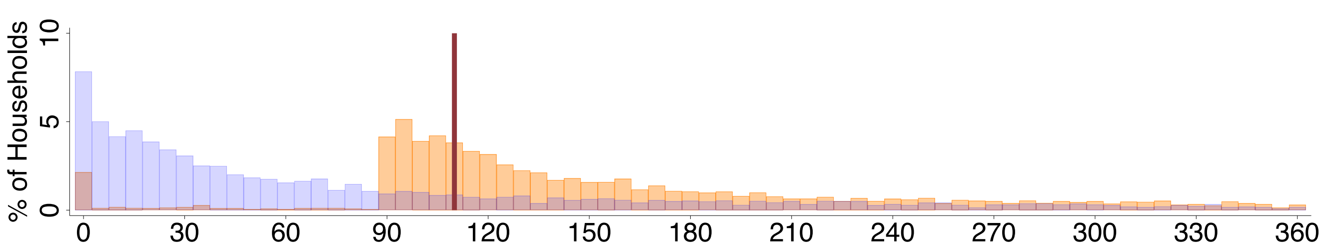

C Calibration for the simulation

Table C.1: Calibration of labor income process

Parameter	Value	Parameter	Value
ρ	0.958	$\sigma_{\varepsilon,2}$	0.063
p_z	21.9%	σ_α	0.298
$\mu_{\eta,1}$	-0.147	$\sigma_\beta \times 10$	0.185
$\sigma_{\eta,1}$	0.457	$\text{corr}_{\alpha\beta}$	0.976
$\sigma_{\eta,2}$	0.139	$a_\nu \times 1$	-3.2740
$\sigma_{z_1,0}$	0.667	$b_\nu \times t$	-0.8935
λ	0.001	$c_\nu \times z_t$	-4.5692
p_ε	12.6%	$d_\nu \times t \times z_t$	-2.9203
$\mu_{\varepsilon,1}$	0.236	$a_{z_1} \times 1$	0.2191
$\sigma_{\varepsilon,1}$	0.343		