THE FULL RECESSION: PRIVATE VERSUS SOCIAL COSTS OF COVID-19

Juan Carlos Córdoba*, Marla Ripoll† and Siqiang Yang‡

December, 2020

Abstract

Official recession figures ignore the costs associated with the loss of human life due to COVID-19. This paper constructs "full recession" measures that take into account the death toll. Key for the estimates are the number of dead, the individual's willingness to accept mortality risk, and society's willingness to accept inequality. Our model features tractable heterogeneity, constant relative risk aversion to mortality risk, and age-specific survival rates. Using an estimated death toll of 300 thousand people for the US during a year, and a 5% recession, we find that the corresponding full recession is 19% on average across individuals, 11% for a median voter, 7% for a planner with mild aversion to inequality, and 12% for a planner with larger aversion. Regarding the overall cost of the pandemic, we find that individuals would be willing to pay, on average, 41% of one-year consumption to fully avoid the pre-lockdown 1.9 million deaths from COVID-19. A median voter would be willing to pay 23%, a social planner with mild aversion to inequality only 11%, while a log-planner would pay 73%.

Key words: social welfare, value of statistical life, mortality risk aversion, consumption risk aversion, Epstein-Zin-Weil preferences, random walk, pandemics

JEL Codes: I14, I31, J17

1 Introduction

As the COVID-19 pandemic spreads around the world, the death toll keeps increasing and economic activity stagnates. Individuals are confronted with difficult choices between risking their lives and trying to carry out normal economic and non-economic activities. The mortality risks involved are very heterogeneous, increasing with age and certain medical conditions, and affecting disadvantaged demographic groups at significantly higher rates. In the aggregate, societies face a trade-off between the loss of human life and a non-trivial drop in living standards for an unforeseeable future. In a situations like these, traditional measures of economic activity, such as GDP, may be misleading indicators of the society's well-being since they do not typically incorporate any adjustment for the loss of human life. A more reliable indicator of well-being during pandemics, one that could

*Department of Economics, Iowa State University. E-mail: cordoba@iastate.edu

†Department of Economics, University of Pittsburgh. E-mail: ripoll@pitt.edu

‡Corresponding author. School of Economics, Nankai University. E-mail: siy10@nankai.edu.cn

be used to guide policy, would need to include a sensible adjustment for the increased risk at the individual level, and the death toll at the aggregate level.

This paper constructs robust microfounded measures of well-being useful in the presence of mortality shocks. Our indicators seek to answer two interrelated questions: *(i)* what is the imputed total cost of the pandemic? and *(ii)* what is the full, or effective, yearly recession that takes into account the death toll? As we show, the answers to these questions depends crucially on five elements: the death toll, the magnitude of the mortality risk involved, the willingness of individuals to accept mortality risk, the age-distribution of the population, and the willingness of societies to accept inequality.

There is a literature that has estimated the aggregate consequences of changes in mortality risk, due for example to public health interventions, the AIDS pandemic, or the end of wars. Papers include Becker *et al.* (2005), Jones and Klenow (2016), Cordoba and Ripoll (2017) and Murtin *et al.* (2017) among others. A common assumption in this literature, including the recent papers on COVID-19, is that individuals are expected utility maximizers. In Cordoba and Ripoll (2017, CR henceforth), two of us demonstrated key limitations of the expected utility model when applied to issues of mortality risk. A fundamental issue is that while the expected utility model could guarantee a constant coefficient of relative risk aversion for non-mortality risk, it may at the same time imply strongly increasing relative aversion to *mortality* risk. For example, consider Hall *et al.* (2020) who have recently provided welfare costs estimates of the pandemic using an expected utility model. In their model, the value of a year of life for an average individual is five times average consumption. But this value is only two times consumption for an individual with half the average consumption. Individuals with 1/6 of the average consumption, or below, would not value life at all. Such a strongly decreasing valuation of life is not supported by the existing evidence, as argued by CR.[1]

This paper provides a flexible framework to evaluate individual, social, and distributional welfare effects of COVID-19 and other maladies. The paper builds on CR by introducing incomplete markets, incomplete annuities, age-dependant survival rates, and tractable heterogeneity. A distinguishing feature of this new framework is the disentangling of three key preference parameters that are otherwise embodied in a single parameter in the expected utility model. The three parameters are: *(i)* the elasticity of intertemporal substitution; *(ii)* the coefficient of relative risk aversion to non-mortality risks; and *(iii)* the coefficient of relative risk aversion to mortality risk. This disentangling, an extension of the original Epstein-Zin-Weil (EZW) preferences, allows to preserve a constant relative aversion to mortality risk for any value of the elasticity of intertemporal substitution. As a result, the valuation of life is proportional to consumption for both low and high-income individuals.

In addition to generalizing the framework in CR, a key contribution of the paper is to characterize the social planner's problem for the case in which the EZW preferences feature two distinct risk aversion parameters, namely mortality and non-mortality risk. The tractability of the model allows us to derive closed-form solutions for some parameters and clearly show how the private and

[1] See Section 2.3 in Cordoba and Ripoll (2017).

social costs of the pandemic are related. We also derive expressions to show the role of the social planner's inequality aversion on the social welfare costs. Furthermore, the framework developed in this paper preserves the same advantages of that in CR. First, the valuation of changes in mortality rates depends on expected survival, which could rationalize why the elderly, who have less expected survival, would be willing to pay non-trivial amounts to extend their shorter remaining life span. Second, the coefficient of relative risk aversion to mortality risk becomes the natural parameter to match the valuation of life. In the expected utility model, the parameter utilized to match this target is typically a level parameter which artificially creates income effects not seen in the data. These considerations are important in evaluating the welfare costs of a pandemic when there is inequality in consumption. In sum, the EZW utility we use has distinct quantitative implications for the value of life for young adults, the elderly, and individuals with low consumption levels, all of which are important in analyzing the distributional welfare implications of COVID-19.

We calibrate the model to recent US data. The age-varying parameters of the consumption process are estimated from 1999-2017 PSID data. Survival probabilities by age are obtained from the National Vital Statistics, and the demographic structure of the population from the US Census, both for the most recent available year, 2017. The key parameter, the coefficient of relative aversion to mortality risk, is calibrated to match plausible Values of Statistical Life commonly used in the health literature. The age-specific decreases in survival probabilities due to COVID-19, before interventions, are calculated from the pre-lockdown data in Menachemi *et al.* (2020) and Ferguson *et al.* (2020). We then use the model to calculate various willingness to pay to avoid the pandemic, which provides estimates for the overall burden of the pandemic. In addition, we use the model to compute "full recessions", that is, consumption equivalent percentage drops in welfare due to the combination of an economic recession, the death toll, and the increase in mortality risk. Finally, we also characterize the trade-off between lives and consumption by tracing a frontier for the average and median voter, as well as the planner.

Regarding the total burden of the pandemic, our analysis provides the following insights. First, in our preferred calibration, the welfare costs of a one-year COVID-19 pandemic are high: on average (population-weighted), individuals would be willing to reduce consumption by 41.2% during one year to completely avoid the deaths of 1.9 million people due to the pandemic. This death toll is computed from the pre-lockdown age-specific fatality rates of Ferguson *et al.* (2020) and the aggregate death rate from Menachemi *et al.* (2020). Our preferred calibration features an elasticity of intertemporal substitution equal to one, a widely used value in quantitative macro, as well as the absence of annuity markets, a reasonable benchmark given the limited use of annuities in the data. The average welfare costs are reduced to 35.9% if perfect annuities are introduced, since more insurance is available in this case. Regardless of the scenario, the welfare costs of COVID-19 are high.

We find a large dispersion of the welfare costs across ages. In our preferred calibration, while the average 25-year old would only pay 5% of one-year consumption to avoid COVID-19, the average 45-year old would pay 23%, and the average 60-year old would pay 80%. These welfare costs parallel the dramatic differences in the estimated survival reductions from COVID-19 by age group in the

absence of a lockdown: while the fatality rate for a 25-year old is 0.015%, it is 0.08% for a 45 year old, and 1.16% for a 60-year old (Menachemi *et al.*, 2020 and Ferguson *et al.* 2020). In fact, the standard deviation of the age-specific welfare costs is 41.2%, capturing the substantial disagreement across individuals in different age groups.

Third, we also find a disagreement between the average welfare cost and the median voter: the latter is always willing to accept a lower reduction in consumption than the average. This result is robust to different scenarios. Taking into account the US distribution of the population above age 18, the median voter is age 46. In our preferred benchmark calibration, the median voter is willing to accept 22.7% less consumption during the year of the pandemic, while the average would reduce consumption by 41.2% during one year to avoid the pandemic. That the average is above the median reflects the fact that the welfare costs rapidly increase with age under COVID-19. In fact, we find a sharp increase in welfare costs after age 40 and costs above 90% for those older than 65.

Another main finding is that the extent to which the social planner cares about inequality has sizable effects on the aggregate willingness to pay to avoid the pandemic. In our preferred calibration, a social planner with lower aversion to inequality would be willing to pay 10.8% but this number increases up to 72.7% for a log planner. When the planner cares less about inequality, the utility of the young weights relatively more since they have higher continuation value. In this case the planner is willing to pay less to avoid the pandemic because younger people face substantially lower probability of dying from COVID-19. As the planner's inequality aversion increases, the relative weight of the elderly increases, resulting in a higher willingness to pay.

Fifth, in a scenario where COVID-19 lasts one year but the consumption costs extend beyond one period, the welfare costs are still sizable. We find that if individuals pay costs for about 8 years, which simulates lingering effects from a crisis similar to the 2007-2009 Great Recession, then the average welfare costs on the first year are 30.9% in our preferred calibration. In this case the median voter is willing to cut consumption by 12.1% on that first year. In addition, if consumption costs were closer to permanent, and individuals had to accept reductions in consumption for at least 30 years, then they would on average accept a reduction of 12.2% on the first year, while the median voter would accept a 2.9% reduction.

Our model allows us to compute measures of "full recessions." We use our model to simulate various combinations of economic recessions, resulting from lockdown policies, and the loss of lives due to the pandemic. We find that a 5% contraction in consumption during the year of the pandemic combined with a lost of 300,000 lives corresponds to a 18.8% full recession for the average individual, a 10.6% full recession for the median voter, a 7.1% full recession for a planner with a relatively low aversion to inequality, and a 12.4% for a planner with higher aversion to inequality. When the lives lost go up to 400,000, the corresponding full recessions are 21.9% for the average, 12.4% for the median voter, 7.7% for the planner with low inequality aversion, and 14.8% for the planner with higher aversion.

To interpret these figures, take first the case of the planner with lower aversion to inequality and 400,000 lives lost. Since for the planner the full recession is 7.7%, the additional 2.7% corresponds

to the planner's valuation of the lives lost. Assuming a US GDP of $22 trillion, this roughly corresponds to a planner's valuation of $1.5 million per life lost. But for the planner with higher aversion to inequality, this valuation becomes $5.4 million per life lost. For the median voter, who is age 46, the implied marginal valuation of life is roughly $4.1 million, while the implied marginal valuation of life for the average individual is $9 million. The latter figure reflects the fact that those above age 60 would be willing to pay more to reduce their risk exposure as they not only face higher mortality rates already, regardless of the pandemic, but also larger death rates from the pandemic.

Last, we find that the consumption-lives frontier varies substantially across the average, the median voter and the planner. While the average (population-weighted) frontier is more concave and reflects a higher willingness to sacrifice consumption to save lives, the frontier for the median voter tends to be less concave and reflects lower willingness to give up consumption. Starting at 1.9 million deaths, which correspond to projections consistent with Menachemi *et al.*'s (2020) (aggregate death rate of 0.58%), the average individual is initially willing to reduce consumption sharply to save lives: the first 1.1 million lives are traded-off by a 32.9% consumption cut in the year of the pandemic. But the trade-off is smaller for the next 800 thousand lives, adding a extra 8% consumption cut. The concavity of the average frontier reflects the large variation in willingness to pay across individuals of different ages. For the median voter, the frontier is less concave, sacrificing 13.9% of consumption to save the first 1.1 million lives, an additional 8.8% for the next 800 thousand people.

Our paper complements a number of other recent papers on COVID-19, but it is most closely related to Hall *et al.* (2020). Hall *et al.* (2020) also compute the welfare costs of COVID-19. Relative to them our model includes incomplete markets, inequality, and a non-expected utility representation. In this respect, we develop a rich framework that merges the traditional analysis of the welfare costs of economic shocks, consumption inequality, and the literature on the value of life over the life cycle. Although rich, our framework is tractable. In addition to our welfare costs calculations, we also explore the dispersion of the welfare cost across ages and the disagreement between the median voter and the average individual. We also explore the potential scenario in which the consumption costs of COVID-19 extend beyond one year. Finally, we compute measures of full recession and a consumption-lives frontier.

The remainder of the paper is organized as follows. Section 2 presents the model for the individual, derives the solution and its properties, and discusses the notion of value of statistical life. It also includes all the aggregation results and the theoretical derivation of the welfare costs of COVID-19. Although the model exhibits substantial heterogeneity, in Section 2 we derive closed-form solutions for the aggregate economy. The calibration procedure in explained in Section 3 and all numerical results are presented in Section 4. The relationship to the literature is discussed in Section 5. Section 6 concludes.

2 Model

The exercise of this paper is similar in spirit to Lucas (1987, 2003), Tallarini (2000), Cordoba and Verdier (2008), Jones and Klenow (2016) and Cordoba and Ripoll (2017) among others. For a given path of individual consumption and mortality rates over the life cycle, and for given individual and social welfare functions, individual and social welfare is calculated for a benchmark in which a pandemic is occurring. The burden of the pandemic and the full recession estimates for individuals and for a planner are calculated from their willingness to pay to avoid either the risk and death toll of the full pandemic, or the risk and death toll associated with only the first year of the pandemic.

2.1 Consumption of the alive

Time is discrete and denoted by t. There is an initial distribution of population of ages $a \geq 0$ denoted by $M_{a,0}$. Let $\pi_{a,t}$ be the probability of surviving to age a at time t conditional on having survived to age $a-1$ at $t-1$. The distribution of population at t satisfies $M_{a,t} = \pi_{a,t} M_{a-1,t-1}$ for $a > 0$ and $M_{0,t} = n^t$, where n is the birth rate.

Let $I(a,t)$ be the set individuals of age a alive at time t and $c_{i,a,t}$ be the consumption of individual i at age a and time t, where $i \in I(a,t)$. Assume that (log) consumption follows a random walk

$$\ln c_{i,a,t} = \ln c_{i,a-1,t-1} + \varepsilon_a, \tag{1}$$

where $\varepsilon_a \sim N\left(\mu_a - \eta_a^2/2, \eta_a^2\right)$. The age-dependent drift and variance would allow us to replicate a realistic life-cycle pattern of consumption. Assume that age-0 consumption is drawn from a log-normal distribution,

$$\ln c_{i,0,t} \sim N\left(\ln c_{0,t} - \Omega^2/2, \Omega^2\right) \text{ for } i \in I(0,t), \tag{2}$$

where Ω^2 represents the initial (age-0) consumption inequality. The average consumption of cohort t at age 0 is then $E\left[c_{i,0,t}|_t\right] = e^{\ln c_{0,t} - \Omega_t^2/2 + \Omega_t^2/2} = c_{0,t}$. A constant rate of aggregate economic growth is introduced by assuming that $c_{0,t} = c_{0,0} e^{gt}$, where g is the growth rate. We normalize $c_{0,0} = 1$. Under these assumptions it follows that the unconditional distribution of consumption at any point in time is log-normal,

$$\ln c_{i,a,t} \sim N\left(g\left(t-a\right) + \upsilon_a - \Omega^2/2 - \Phi_a^2/2, \Omega^2 + \Phi_a^2\right) \text{ for } i \in I(a,t), \tag{3}$$

where $\upsilon_a = \sum_{j=1}^{a} \mu_j$ and $\Phi_a^2 = \sum_{j=1}^{a} \eta_j^2$ for $a \geq 1$.Intuitively, $g\left(t-a\right)$ reflects economic growth, υ_a life-cycle growth, Ω^2 initial inequality and Φ_a^2 reflects life-cycle inequality. Given that (log) consumption is a random walk, inequality increases with age. Using the law of large numbers, (3) also describes the distribution of consumption, or inequality, at any point in time.

Let effective consumption be defined as $\bar{c}_{i,a,t} = \lambda_{i,a,t} c_{i,a,t}$ where $\lambda_{i,a,t}$ is a proportional shift parameter that will be used to calculate welfare under various scenarios. In the baseline $\lambda_{i,a,t} = 1$

for all i, a and t.

2.2 Preferences

We model individual utility by adopting an extension of the EZW framework due to CR.[2] This extension disentangles intertemporal substitution, mortality risk aversion, and aversion to consumption risk. The utility of individual i at age a and time t is described by

$$\overline{V}_{i,a,t} = \left[\pi_{a,t} V_{i,a,t}^{1-\gamma} + (1-\pi_{a,t})\, D_{i,a,t}^{1-\gamma}\right]^{\frac{1}{1-\gamma}}, \text{ with } \gamma \geq 0,\ \gamma \neq 1, \tag{4}$$

where $V_{i,a,t} \geq 0$ is the utility of being alive, $D_{i,a,t} \geq 0$ the perceived utility upon dead, and γ is *the coefficient of mortality aversion*, a parameter describing attitudes toward mortality risk. In addition,

$$V_{i,a,t} = \begin{cases} \left[(1-\beta_a)\, \overline{c}_{i,a,t}^{1-\sigma} + \beta_a \left(E_a\left[\overline{V}_{i,a+1,t+1}^{1-\theta}\right]\right)^{\frac{1-\sigma}{1-\theta}}\right]^{\frac{1}{1-\sigma}} & \text{for } \sigma \geq 0, \sigma \neq 1, \\ \overline{c}_{i,a,t}^{1-\beta_a} \left(E_a\left[\overline{V}_{i,a+1,t+1}^{1-\theta}\right]\right)^{\frac{\beta_a}{1-\theta}} & \text{for } \sigma = 1, \end{cases}$$

where β_a is an age-specific discount factor, $1/\sigma$ is the standard intertemporal substitution, and θ captures overall aversion to risk.[3] Parameter θ is the traditional coefficient of risk aversion in EZW preferences as it will become clear shortly.

In what follows we focus on the case in which $D_{i,a,t} = 0$, which ensures that being alive is always preferred to death.[4] Such normalization requires the restriction $\gamma \in (0,1)$ as otherwise the strong complementarity between V and D would render $\overline{V} = 0$ too. Preferences are thus described by

$$\overline{V}_{i,a,t} = \pi_{a,t}^{\frac{1}{1-\gamma}} V_{i,a,t}, \text{ with } \gamma \in (0,1), \tag{5}$$

and

$$V_{i,a,t} = \begin{cases} \left[(1-\beta_a)\, \overline{c}_{i,a,t}^{1-\sigma} + \beta_a \pi_{a+1,t+1}^{\frac{1-\sigma}{1-\gamma}} \left(\left(E_a\left[V_{i,a+1,t+1}^{1-\theta}\right]\right)^{\frac{1-\sigma}{1-\theta}}\right)\right]^{\frac{1}{1-\sigma}} & \text{for } \sigma \geq 0, \sigma \neq 1, \\ \overline{c}_{i,a,t}^{1-\beta_a} \left(\pi_{a+1,t+1}^{\frac{1}{1-\gamma}} E_a\left[V_{i,a+1,t+1}^{1-\theta}\right]^{\frac{1}{1-\theta}}\right)^{\beta_a} & \text{for } \sigma = 1. \end{cases} \tag{6}$$

Equation (6) is a Bellman equation for $V_{i,a,t}$ alone with an "effective" discount factor on future utility of $\beta_a \pi_{a+1,t+1}^{(1-\sigma)/(1-\gamma)}$ rather than the standard $\beta_a \pi_{a+1,t+1}$. Restriction $\gamma \in [0,1)$ guarantees that

[2] See Section 5.3 in Cordoba and Ripoll (2017).

[3] Regarding the age-specific discount factors, in the calibration we will follow the traditional approach of matching a target for the risk-free rate. Given that consumption growth and survival rates may vary with age, hitting the target may require age-specific discount factors as explained in Section 3.3.

[4] Section 4.2 considers a case with $D > 0$.

higher survival results in higher utility. When $\pi = 1$ for all a, these preferences correspond to the standard EZW preferences with θ as the only coefficient of relative risk aversion.

Utility function (6) accommodates preferences for early or late resolution of both mortality and consumption uncertainty. In particular, when $\sigma > \gamma$ individuals exhibit a preference for late resolution of mortality uncertainty, while $\theta > \sigma$ implies a preference for early resolution of consumption uncertainty. In this case individuals prefer not to know in advance when they will die but would like to know their consumption realizations early on. CR document evidence on a preference for late resolution of uncertainty regarding death, while the finance literature provides evidence for early resolution of consumption uncertainty. Finally, the case $\sigma = 1$, popular in quantitative macro, corresponds to a version of the risk-sensitive preferences in Tallarini (2000), but generalized to include aversion to mortality risk.

The preferences above include as a special case an expected utility representation with constant relative risk aversion. To see this more clearly define the monotonic transformation $X = V^{1-\sigma}/\left[(1-\beta)(1-\sigma)\right]$ for $\sigma \geq 0$, $\sigma \neq 1$ and $\beta \in (0,1)$, to rewrite (6) as

$$X_{i,a,t} = \frac{\overline{c}_{i,a,t}^{1-\sigma}}{1-\sigma} + \beta_a \pi_{a+1,t+1}^{\frac{1-\sigma}{1-\gamma}} \left(\left(E_a \left[\left(X_{i,a+1,t+1} \right)^{\frac{1-\theta}{1-\sigma}} \right] \right)^{\frac{1-\sigma}{1-\theta}} \right).$$

The case $\theta = \sigma$ is a standard expected utility model for consumption risk, but not for mortality risk. The case $\sigma = \gamma$ correspond to expected utility for mortality risk, but not for consumption risk. The case $\theta = \sigma = \gamma$ is the full expected utility model for both consumption and mortality risk. In this respect our utility representation is more flexible, allowing us to capture additional features on individuals' attitudes toward mortality and consumption risk.

2.3 Closed-form solution of the value function

Closed-form solutions are possible given the set up of the problem, as shown in the following proposition.

Proposition 1. The utility of individual $i \in I(a,t)$ is given by

$$\overline{V}_{i,a,t} = \pi_{a,t}^{\frac{1}{1-\gamma}} V_{i,a,t} = \pi_{a,t}^{\frac{1}{1-\gamma}} A_{i,a,t} \overline{c}_{i,a,t} \text{ for all } i,\ t \text{ and all } a, \tag{7}$$

where $A_{i,a,t}$ is the saddle path solution of the following first order difference equation,

$$A_{i,a,t} = \begin{cases} \left[1 - \beta_a + \beta_a \left[\pi_{a+1,t+1}^{\frac{1}{1-\gamma}} A_{i,a+1,t+1} \frac{\lambda_{i,a+1,t+1}}{\lambda_{i,a,t}} e^{\mu_{a+1} - \theta \eta_{a+1}^2/2} \right]^{1-\sigma} \right]^{\frac{1}{1-\sigma}} & \text{for } \sigma \geq 0, \sigma \neq 1, \\ \left[\pi_{a+1,t+1}^{\frac{1}{1-\gamma}} A_{i,a+1,t+1} \frac{\lambda_{i,a+1,t+1}}{\lambda_{i,a,t}} e^{\mu_{a+1} - \theta \eta_{a+1}^2/2} \right]^{\beta_a} & \text{for } \sigma = 1. \end{cases} \tag{8}$$

Proof. Guess that value function of the alive has the form $V_{i,a,t} = A_{i,a,t}\bar{c}_{i,a,t}$ where $A_{i,a,t}$ is a coefficient to be determined. Under this guess, and using the fact that $\bar{c}_{i,a,t} = \lambda_{i,a,t}c_{i,a,t}$, it follows that:

$$\begin{aligned} E\left[V_{i,a+1,t+1}^{1-\theta}|_{a,t}\right] &= \left(A_{i,a+1,t+1}\lambda_{a+1,t+1}\right)^{1-\theta} E\left[c_{i,a+1,t+1}^{1-\theta}|_{a,t}\right] \\ &= \left[A_{i,a+1,t+1}\lambda_{a+1,t+1}c_{i,a,t}e^{\mu_{a+1}-\theta\eta_{a+1}^2/2}\right]^{1-\theta}. \end{aligned}$$

Substituting this expression into (6) we can write

$$V_{i,a,t} = \begin{cases} \lambda_{i,a,t}c_{i,a,t}\left[1-\beta_a+\beta_a\pi_{a+1,t+1}^{\frac{1-\sigma}{1-\gamma}}\left(A_{i,a+1,t+1}\frac{\lambda_{i,a+1,t+1}}{\lambda_{i,a,t}}e^{\mu_{a+1}-\theta\eta_{a+1}^2/2}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}} & \text{for } \sigma\geq 0, \sigma\neq 1, \\ \lambda_{i,a,t}c_{i,a,t}\left[\pi_{a+1,t+1}^{\frac{1}{1-\gamma}}A_{i,a+1,t+1}\frac{\lambda_{i,a+1,t+1}}{\lambda_{i,a,t}}e^{\mu_{a+1}-\theta\eta_{a+1}^2/2}\right]^{\beta_a} & \text{for } \sigma=1, \end{cases} \tag{9}$$

which satisfies the guess if $A_{i,a,t}$ satisfies (8).

In the baseline, $\lambda_{i,a,t} = 1$ and $A_{i,a,t} = A_{a,t}$, so that $A_{i,a,t}$ is independent of i. The linearity of the utility in consumption is worth stressing and results from the fact that the EZW aggregator is homogeneous of degree one. An implication is that a utilitarian planner will not be averse to inequality, as we discuss below.

2.4 Some properties of the value function

The sequence $A_{i,a,t}$, utility per unit of consumption, is a central object because it contains all the welfare information (rate of time preference, mortality and consumption risk aversion, and intertemporal substitution) as well as survival rates. Intuitively, one would expect that older individuals would derive less utility from consumption as their life span diminishes. This is in fact the case in the calibrated exercises. To gain some intuition, consider the case in which all parameters $(\mu_a, \eta_a, \pi_a, \beta_a)$ become constant after a certain age $a \geq a^*$. In particular, survival is still positive but becomes constant, arbitrarily small and equal to π. Sequence $A_{i,a,t}$ then converges to a terminal value $\bar{A}$ which can be written as

$$\bar{A} = \begin{cases} (1-\beta_a)^{\frac{1}{1-\sigma}}\left[1-\beta_a\left(\pi^{1/(1-\gamma)}e^{\mu-\theta\eta^2/2}\right)^{1-\sigma}\right]^{\frac{-1}{1-\sigma}} & \text{for } \sigma\geq 0, \sigma\neq 1, \\ \left[\pi^{\frac{1}{1-\gamma}}e^{\mu-\theta\eta^2/2}\right]^{\frac{\beta_a}{1-\beta_a}} & \text{if } \sigma=1. \end{cases} \tag{10}$$

It is easy to show that starting from $\bar{A}$ and iterating backwards results in a sequence $A_{i,a,t}$ that decreases with age given that survival probabilities decrease with age. Consider first the case when $\sigma = 1$. Starting from $\bar{A}$ and going one period backwards, $A_{i,a^*-1,t-1} > \bar{A}$ as long as $[\pi_{a^*-1,t-1}/\pi]^{\beta_a/(1-\gamma)} > 1$, which always holds. Similarly, for the case $\sigma < 1$, it can be shown

that $A_{i,a^*-1,t-1} > \bar{A}$ as long as $\pi_{a^*-1,t-1}^{1-\sigma/(1-\gamma)} > \pi^{1-\sigma/(1-\gamma)}$ which always holds. Finally, when $\sigma > 1$ we have that $A_{i,a^*-1,t-1} > \bar{A}$ as long as $\pi_{a^*-1,t-1}^{1-\sigma/(1-\gamma)} < \pi^{1-\sigma/(1-\gamma)}$, which always holds when $\sigma > 1$. We verify in our calibrated model that even for age-varying μ_{a+1} and η_{a+1}, $A_{i,a,t}$ decreases with age.

2.5 Value of statistical life

Although the elderly derive lower continuation utility from consumption than the young, the willingness to pay to reduce the risk of dying is not necessarily lower for the elderly. To see this we now compute the value of statistical life (VSL) in the model, which captures the trade-offs between survival and consumption. The VSL corresponds to the marginal rate of substitution between survival and consumption defined as:

$$VSL_{i,a,t} = \frac{\partial \overline{V}_{i,a,t}/\partial \pi_{a,t}}{\partial \overline{V}_{i,a,t}/\partial c_{i,a,t}} = \frac{1}{1-\gamma}\frac{1}{1-\beta_a}\frac{1}{\pi_{a,t}}\frac{V_{i,a,t}^{1-\sigma}}{c_{i,a,t}^{-\sigma}} \text{ for any } \sigma. \tag{11}$$

Using (7), this expression simplifies to

$$VSL_{i,a,t} = \frac{1}{1-\gamma}\frac{1}{1-\beta_a}\frac{A_{a,t}^{1-\sigma}}{\pi_{a,t}}c_{i,a,t} \text{ for any } \sigma. \tag{12}$$

The equation above shows the special role of γ, the coefficient of relative mortality aversion, in determining the VSL. The individual willingness to pay for survival increases with γ and $\lim_{\gamma \to 1} VSL_{i,a,t} = \infty$. This expression will be used to calibrate γ below. The time preference parameter β_a plays an analogous role to that of γ but the mechanism is different. While a larger γ represents a larger aversion to mortality risk, a larger β_a represents an increase in the weight that individuals place on the future, which increases the value of surviving.

The effect of $A_{a,t}$ on the VSL depends on the value of σ, and it is only active if $\sigma \neq 1$. Since $A_{a,t}$ decreases with age, then the VSL would tend to increase with age if $\sigma > 1$, but it would decrease it if $\sigma < 1$. The intuition for this result is that when $\sigma > 1$ there is complementary in consumption over time, so everything else equal, those with fewer periods left to live will be willing to pay more to extend life.

Higher survival, $\pi_{a,t}$, reduces the willingness to pay for survival in two ways: a first-order direct effect through term $1/\pi_{a,t}$, and an indirect effect through $A_{i,a,t}$. The direct effect reflects an increase in the expected marginal utility of consumption as current consumption only takes places if the individual survives. Finally, the VSL is linear in consumption.

2.6 Annuities and asset pricing

Given an exogenous process for consumption, an approach in the finance literature is to use first-order conditions, or Euler equations, to price assets such as riskless bonds or equity. In this section,

we use this approach to price a one-period riskless bond. This price is used later to calibrate preference parameters.

In the absence of mortality risk, a riskless bond pays one unit of consumption in any state of the world. In the presence of mortality risk, an annuitized riskless bond pays one unit of consumption only if the individual survives, while a non-annuitized riskless bond pays even if the individual dies. Let $p^A_{a,t} = \pi_{a,t}/(1+r)$ and $p^{NA} = 1/(1+r)$ be the prices of these two bonds respectively and let $p^\delta_{a,t} = \delta p^A_{a,t} + (1-\delta)\, p^{NA}$ the price of a bond that is partly annuitized to an extent determined by $\delta \in [0,1]$. If consumers' only savings vehicle is the semi-annuitized bond, $p^\delta_{a,t}$ is the relevant marginal rate of transformation. Focusing on the case $\sigma \neq 1$, it can be shown that the Euler equation corresponding to this asset market structure and (6) is given by

$$p^\delta_{a,t}\left(1-\beta_a\right) c^{-\sigma}_{i,a,t} = \beta_a \pi^{\frac{1-\sigma}{1-\gamma}}_{a+1,t+1} E_{a,t}\left[\left(\frac{V_{i,a+1,t+1}}{\widetilde{V}_{i,a+1,t+1}}\right)^{\sigma-\theta}\left(1-\beta_{a+1}\right)\left(c_{i,a+1,t+1}\right)^{-\sigma}\right], \tag{13}$$

where $\widetilde{V}_{i,a+1,t+1} = \left(E_a\left(V^{1-\theta}_{i,a+1,t+1}\right)\right)^{\frac{1}{1-\theta}}$ is the certainty equivalent. The left-hand side is the marginal cost of buying the semi-annuitized bond, while the right-hand side is the marginal benefit. This equation reduces to a traditional Euler equation for a riskless bond when $\pi = 1$, $\sigma = \theta = \gamma$ and $\beta_a = \beta$. In the more general case, the effective discount factor is given by $\widetilde{\beta}_a \equiv \beta_a \pi^{(1-\sigma)/(1-\gamma)}_{a+1,t+1}$ and the stochastic discount factor is adjusted by term $\left(V_{i,a+1,t+1}/\widetilde{V}_{i,a+1,t+1}\right)^{\sigma-\theta}$, reflecting risk-taking behavior.

Using (7), equation (13) simplifies to

$$p^\delta_{a,t} = \frac{1-\beta_{a+1}}{1-\beta_a}\beta_a \pi^{\frac{1-\sigma}{1-\gamma}}_{a+1} e^{-\sigma\mu_{a+1}+\theta(1+\sigma)\eta^2_{a+1}/2}. \tag{14}$$

Below in the quantitative exercise we consider the cases of perfect annuity markets, $\delta = 1$, and no annuity markets, $\delta = 0$.

2.7 Social welfare

Consider now social welfare according to a fictional benevolent social planner. For concreteness, suppose the planner only cares about the welfare of the population alive at time 0.[5] Define social welfare as

$$W_0 = \left[\sum_a \int_{i\in I(a,0)} \overline{V}^{1-\psi}_{i,a,0} di\right]^{\frac{1}{1-\psi}} = \left[\sum_a \int_{i\in I(a,0)} \left(\pi^{\frac{1}{1-\gamma}}_{a,0} A_{i,a,0}\lambda_{i,a,0} c_{i,a,0}\right)^{1-\psi} di\right]^{\frac{1}{1-\psi}}, \tag{15}$$

[5] A more general social welfare function will include the welfare of all potential individuals, present and future. Implicitly, we are assuming a social discount factor of zero.

where curvature parameter ψ describes planner's aversion to inequality. A classical utilitarian planner is represented by the case $\psi = 0$. As mentioned above, the EZW aggregator is homogenous of degree one, which results in the linearity of $\overline{V}$ with respect to c. Absent any curvature in the individual value function, a classical utilitarian planner would have no aversion to inequality.[6] Although it might be natural to assume $\psi = \sigma$, since σ reflects the curvature of the utility flow, this is not obvious in the case of EZW utility which features two additional curvature parameters: γ and θ. We leave ψ as a separate curvature parameter for the planner and report results for several values of ψ in the quantitative section.

The expression above can be simplified whenever $\lambda_{i,a,0} = \lambda_{a,0}$, which happens to hold in this model due to the homotheticity of utility. In this case $A_{i,a,0} = A_{a,0}$ as well. Thus,

$$W_0 = \left[\sum_a \left(\pi_{a,0}^{\frac{1}{1-\gamma}} A_{a,0} \lambda_{a,0} \right)^{1-\psi} \int_{i \in I(a,0)} c_{i,a,0}^{1-\psi} di \right]^{\frac{1}{1-\psi}}.$$

Furthermore, since $c_{i,a,0}$ is distributed log-normal according to (3), then by the law of large numbers

$$\begin{aligned} \int_{i \in I(a,0)} c_{i,a,0}^{1-\psi} di &= M_{a,0} E\left[c_{i,a,0}^{1-\psi} \right] = M_{a,0} e^{(1-\psi)\left(\upsilon_a - g a - \overline{\Omega}_a^2/2 \right) + (1-\psi)^2 \overline{\Omega}_a^2/2} \\ &= M_{a,0} e^{(1-\psi)\left(\upsilon_a - g a - \psi \overline{\Omega}_a^2/2 \right)} \end{aligned}$$

where $\overline{\Omega}_a^2 = \Omega^2 + \Phi_a^2$. Therefore, social welfare can be written as

$$W_0 = \left[\sum_a M_{a,0} \left(\pi_{a,0}^{\frac{1}{1-\gamma}} A_{a,0} \lambda_{a,0} e^{\upsilon_a - g a - \psi \overline{\Omega}_a^2/2} \right)^{1-\psi} \right]^{\frac{1}{1-\psi}}. \tag{16}$$

Social welfare can be seen as a weighted composite of cohort-specific average utilities defined as $\widehat{V}_a \equiv \pi_{a,0}^{1/(1-\gamma)} A_{a,0} \lambda_{a,0} e^{\upsilon_a - g a - \psi \overline{\Omega}_a^2/2}$. $\widehat{V}_a$ tends to increase with age due to life-cycle consumption growth (υ_a), but tends to decrease with age due to lower survival, lower $A_{a,0}$, lower initial consumption (if $g > 0$), and larger inequality. If both $M_{a,0}$ and $\widehat{V}_a$ decrease with age, then social welfare would reflect to a larger extent the welfare of the young rather than the old. A larger ψ, however, could counteract a declining role of the old on social welfare unless life-cycle inequality is relatively large.

2.8 Willingness to pay to avoid COVID-19

In this section we derive consumption equivalent measures or the willingness to pay (WTP) to avoid COVID-19. We calculate measures for individuals and for the social planner. It is convenient to make explicit two sequences of parameters determining individual welfare at time 0: survival rates,

[6] The EZW preferences could be modified to introduced an extra parameter of risk aversion to atemporal time-0 risk, or aversion to wealth gambles. However, such parameter is not an inequality aversion parameter and it would be unclear how to calibrate it.

$\boldsymbol{\pi}_{a,0} = \{\pi_{a+s,s}\}_{s\geq 0}$, and consumption shifter parameters $\boldsymbol{\lambda}_{i,a,0} = \{\lambda_{i,a+s,s}\}_{s\geq 0}$. While $\boldsymbol{\lambda}_{i,a,0} = \mathbf{1}$ in the baseline, in this section we use $\boldsymbol{\lambda}_{i,a,0}$ to calculate the WTP. Using this notation, express (7) as:

$$\overline{V}_{i,a,0} \equiv \overline{V}\left(\boldsymbol{\pi}_{a,0}, \boldsymbol{\lambda}_{i,a,0}, c_{i,a,0}\right) = \pi_{a,0}^{\frac{1}{1-\gamma}} \cdot A\left(\boldsymbol{\pi}_{a+1,1}, \boldsymbol{\lambda}_{i,a,0}\right) \cdot \lambda_{i,a,0} \cdot c_{i,a,0}. \tag{17}$$

where $A\left(\boldsymbol{\pi}_{a+1,1}, \boldsymbol{\lambda}_{i,a,0}\right)$ is defined by the recursion in (8). Let $\boldsymbol{\pi}^*_{a,0}$ denote baseline survival rates in the absence of the pandemic. Baseline time-0 utility is given by $\overline{V}^*_{i,a,0} = \overline{V}\left(\boldsymbol{\pi}^*_{a,0}, \mathbf{1}, c_{i,a,0}\right)$.

A pandemic could be described as a sharp drop in survival rates, from $\boldsymbol{\pi}^*_{a,0}$ to $\boldsymbol{\pi}_{a,0}$. We will model the COVID-19 pandemic as a drop in survival probabilities for only one period, $t = 0$. Survival rates remain at their baseline level for other periods. In our calibration, a period will be a year. This description of the pandemic assumes that a treatment or a vaccine is discovered during the initial year. Let $\Delta\pi_{a,0} = \pi_{a,0} - \pi^*_{a,0} \leq 0$ for all a while $\Delta\pi_{a,t} = 0$ for all $a \geq 0$ and all $t > 0$. Initial utility, $\overline{V}_{i,a,0} = \pi_{a,0}^{1/(1-\gamma)} V_{i,a,0}$, is thus affected by the pandemic while the utility of surviving individuals, $V_{i,a,0}$, is not affected because it depends only on future survival rates.

To calculate equivalent variation measures, we compare two economies: one in which individuals face the pandemic but keep the baseline consumption levels, and one in which individuals enjoy baseline survival probabilities but lower consumption. The welfare cost calculation computes the consumption shift, $\boldsymbol{\lambda}_{i,a,0} \leq \mathbf{1}$, that would make individuals indifferent between these two economies. The WTP can then be defined as $WTP_{i,a,0} = 1 - \boldsymbol{\lambda}_{i,a,0}$. We calculate the required sequence $\boldsymbol{\lambda}_{i,a,0}$ under two scenarios. In the first scenario, the cost is fully paid at $t = 0$ so that $\boldsymbol{\lambda}_{i,a,0} = \{\lambda_{i,a,0}, 1, 1, ..\}$. In a second scenario, the cost is paid gradually over time according to $\boldsymbol{\lambda}_{i,a,0} = \left\{\lambda_{i,a,0}, \lambda^{\rho}_{i,a,0}, \lambda^{\rho^2}_{i,a,0}, ..\right\}$, where $\rho \in [0, 1)$ is a persistence parameter. The first scenario is a special case of the second one when $\rho = 0$. The case $\rho > 0$ seeks to describe to a situation in which avoiding the pandemic could trigger a long-lasting recession. Sequence $\boldsymbol{\lambda}_{i,a,0}$ is fully characterized its first element $\lambda_{i,a,0}$.

2.8.1 Individual WTP with $\rho = 0$

Let $\lambda_{i,a,0}$ be implicitly defined by the equation:

$$\overline{V}\left(\boldsymbol{\pi}^*_{a,0} + \Delta\boldsymbol{\pi}_{a,0}, \mathbf{1}, c_{i,a,0}\right) = \overline{V}\left(\boldsymbol{\pi}^*_{a,0}, \left[\lambda_{i,a,0}, \mathbf{1}\right], c_{i,a,0}\right). \tag{18}$$

The left-hand side is the expected utility of individual i under the pandemic. The right hand side is the utility without the pandemic but with a level of effective initial consumption just low enough to deliver the same utility as the one obtained under the pandemic. Using equations (17) and (8), this expression simplifies to

$$\left(\pi^*_{a,0} + \Delta\pi_{a,0}\right)^{\frac{1}{1-\gamma}} A^*_{a,0} = \left(\pi^*_{a,0}\right)^{\frac{1}{1-\gamma}} A_{a,0}\lambda_{i,a,0},$$

where

$$A_{a,0}\lambda_{i,a,0} = \begin{cases} \left[\left(A^*_{a,0}\right)^{1-\sigma} + (1-\beta_a)\left(\lambda_{i,a,0}^{1-\sigma} - 1\right)\right]^{\frac{1}{1-\sigma}} \text{ for } \sigma \geq 0, \sigma \neq 1, \\ \lambda_{a,0}^{1-\beta_a} A^*_{a,0} \text{ for } \sigma = 1, \end{cases} \tag{19}$$

and $A^*_{a,0} = A\left(\boldsymbol{\pi}^*_{a+1,1}, \mathbf{1}\right)$ is the baseline level of $A_{a,0}$. A closed-form solution for $\lambda_{i,a,0}$ is then given by

$$\lambda_{i,a,0} = \lambda_{a,0} = \begin{cases} \left[1 + \left[\left(\frac{\pi^*_{a,0}+\Delta\pi_{a,0}}{\pi^*_{a,0}}\right)^{\frac{1-\sigma}{1-\gamma}} - 1\right] \frac{\left(A^*_{a,0}\right)^{1-\sigma}}{1-\beta_a}\right]^{\frac{1}{1-\sigma}} \text{ for } \sigma \geq 0, \sigma \neq 1, \\ \left(\frac{\pi^*_{a,0}+\Delta\pi_{a,0}}{\pi^*_{a,0}}\right)^{\frac{1}{1-\gamma}\frac{1}{1-\beta_a}} \text{ for } \sigma = 1. \end{cases} \tag{20}$$

Notice that the WTP is independent of i but depends on age: individuals of the same age are willing to pay the same fraction of their initial consumption to avoid the pandemic. According to (20), for the case $\sigma = 1$, $\lambda_{a,0}$ is mostly determined by the extent to which survival drops due to the pandemic ($\Delta\pi_{a,0} < 0$), a magnitude affected by exponent $1/(1-\gamma) \times 1/(1-\beta_a)$. In the case of COVID-19, survival decreases by much more for older individuals and therefore they are more willing to pay. Notice also the link between the WTP in (20) and the VSL-to-consumption ratio in (12). For the case $\sigma = 1$, exponent $1/(1-\gamma) \times 1/(1-\beta_a)$ equals $VSL_{a,0}/c_{a,0} \times \pi^*_{a,0}$. Therefore for $\sigma = 1$,

$$\lambda_{a,0} = \left(\frac{\pi^*_{a,0} + \Delta\pi_{a,0}}{\pi^*_{a,0}}\right)^{\frac{VSL_{a,0}}{c_{a,0}}\pi^*_{a,0}},$$

which provides a link between the WTP and the VSL-to-consumption ratio. As we explain below, we calibrate γ to match $VSL_{a,0}/c_{a,0}$ at age $a = 40$. The larger this ratio, the larger the calibrated γ and the WTP.

When $\sigma \neq 1$ there is an additional effect that comes from $A^*_{a,0}$, and a first-order effect from β_a. Focusing on the case $\sigma > 1$, which is more relevant in quantitative macro, recall that since $A^*_{a,0}$ decreases with age, this effect tends to make the WTP higher for older individuals. The intuition here is again related to the complementarity in consumption across ages for the case $\sigma > 1$, which makes life relatively more value for those who have less periods left. An additional effect occurs if β_a decreases with age, as it would be the case in the calibrated model with $\sigma > 1$. This effect tends to decrease the WTP for older individuals.

The average WTP is calculated as the population-weighted average

$$\overline{WTP}_0 = 1 - \sum\nolimits_a M_{a,0}\lambda_{a,0}.$$

2.8.2 Social WTP with $\rho = 0$

Consider now the social WTP measured by single proportional shift, $\lambda_{i,a,0} = \lambda_{a,0} = \lambda_0$, for everyone in the economy. Let $W_0^{pandemic}$ denote social welfare under the COVID-19 pandemic at time $t = 0$, and W_0^* the social welfare in the economy with the baseline survival but lower consumption. The

one equation in one unknown defining λ_0 is $W_0^{pandemic} = W_0^*$. Using (16), this expression becomes

$$\sum\nolimits_a M_{a,0}\left[\left(\pi_{a,0}^* + \Delta\pi_{a,0}\right)^{\frac{1}{1-\gamma}} A_{a,0}^* e^{\upsilon_a - ga - \psi\overline{\Omega}_a^2/2}\right]^{1-\psi} = \sum\nolimits_a M_{a,0}\left[\left(\pi_{a,0}^*\right)^{\frac{1}{1-\gamma}} A_{a,0}\lambda_0 e^{\upsilon_a - ga - \psi\overline{\Omega}_a^2/2}\right]^{1-\psi}$$

where

$$A_{a,0}\lambda_0 = \begin{cases} \left[\left(A_{a,0}^*\right)^{1-\sigma} + (1-\beta_a)\left(\lambda_0^{1-\sigma} - 1\right)\right]^{\frac{1}{1-\sigma}} \text{ for } \sigma \geq 0, \sigma \neq 1, \\ \lambda_0^{1-\beta_a} A_{a,0}^* \text{ for } \sigma = 1, \end{cases} \tag{21}$$

and $A_{a,0}^*$ is the baseline computed from the recursion in (8) with $\lambda_{a,t} = 1$ for all t. Substituting this expression into the previous equation, and cancelling the common component Ω^2 in $\overline{\Omega}_a^2$, gives an implicit solution for the case $\sigma \geq 0, \sigma \neq 1$,

$$\begin{aligned} & \sum\nolimits_a M_{a,0}\left(\pi_{a,0}^* + \Delta\pi_{a,0}\right)^{\frac{1-\psi}{1-\gamma}}\left(A_{a,0}^* e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi} \\ = & \sum\nolimits_a M_{a,0}\left(\pi_{a,0}^*\right)^{\frac{1-\psi}{1-\gamma}}\left[\left(A_{a,0}^*\right)^{1-\sigma} + (1-\beta_a)\left(\lambda_0^{1-\sigma} - 1\right)\right]^{\frac{1-\psi}{1-\sigma}}\left(e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi}. \end{aligned} \tag{22}$$

In order to better understand the role of the inequality aversion parameter ψ on the social WTP, and to also see the relationship between the social and age-specific WTP, notice that for the case $\sigma = 1$ and $\beta_a = \beta$ for all a, there is a closed-form solution for λ_0 given by

$$\lambda_0 = \left[\frac{\sum_a M_{a,0}\left(\pi_{a,0}^* + \Delta\pi_{a,0}\right)^{\frac{1-\psi}{1-\gamma}}\left(A_{a,0}^* e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi}}{\sum_a M_{a,0}\left(\pi_{a,0}^*\right)^{\frac{1-\psi}{1-\gamma}}\left(A_{a,0}^* e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi}}\right]^{\frac{1}{(1-\beta)(1-\psi)}}. \tag{23}$$

The following proposition derives some results for this special case.

Proposition 2. For the case $\sigma = 1$ and $\beta_a = \beta$ for all a, the λ_0 for a planner with aversion to inequality parameter ψ can be written as

$$\lambda_0 = \left[\sum\nolimits_a f_{a,0}(\psi)\left(\lambda_{a,0}\right)^{(1-\beta)(1-\psi)}\right]^{\frac{1}{(1-\beta)(1-\psi)}} \tag{24}$$

where $\lambda_{a,0}$ is the age-specific WTP as defined in equation (20), and $f_{a,0}(\psi)$ is a density function given by

$$f_{a,0}(\psi) = \frac{M_{a,0}\left(\left(\pi_{a,0}^*\right)^{\frac{1}{1-\gamma}} A_{a,0}^* e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi}}{\sum_a M_{a,0}\left(\left(\pi_{a,0}^*\right)^{\frac{1}{1-\gamma}} A_{a,0}^* e^{\upsilon_a - ga - \psi\Phi_a^2/2}\right)^{1-\psi}}. \tag{25}$$

In addition, the marginal rate of substitution between consumption and survival for the

planner is given by

$$MRS\left(\psi\right)=-\frac{1}{1-\gamma}\frac{1}{1-\beta}\sum\nolimits_{a}f_{a,0}(\psi)\frac{\left(\pi_{a,0}^{*}\right)^{\frac{1-\psi}{1-\gamma}}}{\pi_{a,0}^{*}}. \tag{26}$$

Proof. Density (25) and equation (24) follow from (23). In order to derive (26), notice that the WTP of the planner is given by

$$\alpha\left(\boldsymbol{\Delta\pi}\right)=1-\lambda_{0}=1-\left[\sum\nolimits_{a}f_{a,0}(\psi)\left(\frac{\pi_{a,0}^{*}+\Delta\pi_{a,0}}{\pi_{a,0}^{*}}\right)^{\frac{1-\psi}{1-\gamma}}\right]^{\frac{1}{(1-\beta)(1-\psi)}}$$

where we have used equation (20) for the case $\sigma=1$ to substitute for $\lambda_{a,0}$. The marginal rate of substitution consumption and survival can be calculated as $MRS\left(\psi\right)=\lim_{\boldsymbol{\Delta\pi}\to0}\frac{\alpha(\boldsymbol{\Delta\pi})}{\Delta\pi}$,and using L'Hopital's rule

$$MRS\left(\psi\right)=\lim_{\boldsymbol{\Delta\pi}\to0}\frac{1}{\Delta\pi}\left[1-\left[\sum\nolimits_{a}f_{a,0}(\psi)\left(\frac{\pi_{a,0}^{*}+\Delta\pi_{a,0}}{\pi_{a,0}^{*}}\right)^{\frac{1-\psi}{1-\gamma}}\right]^{\frac{1}{(1-\beta)(1-\psi)}}\right],$$

which results in (26).

Proposition 2 provides three main insights. First, equation (24) suggests that the social costs of the pandemic λ_0 can be expressed as a CES function of the age-specific social costs ($\lambda_{a,0}$). In fact, since β is close to 1, equation (24) is close to a Cobb-Douglas of the form

$$\lambda_{0}\approx\prod\nolimits_{a}\lambda_{a,0}^{f_{a,0}(\psi)},$$

so social welfare is approximately a geometric average of age-specific welfare, where the weights are given by densities $f_{a,0}\left(\psi\right)$. Consider the cases with $\psi=1$ (log planner) and $\psi=0$ (linear planner, no aversion to inequality), with

$$f_{a,0}\left(1\right)=\frac{M_{a,0}}{\sum_{a}M_{a,0}}\text{ and }f_{a,0}\left(0\right)=\frac{M_{a,0}\left(\pi_{a,0}^{*}\right)^{\frac{1}{1-\gamma}}A_{a,0}^{*}e^{\upsilon_{a}-ga}}{\sum_{a}M_{a,0}\left(\pi_{a,0}^{*}\right)^{\frac{1}{1-\gamma}}A_{a,0}^{*}e^{\upsilon_{a}-ga}}.$$

When $\psi=1$, the planner weights the age-specific welfare costs just using population weights. When $\psi=0$, weights depend not only on population, but also on survival probabilities ($\pi_{a,0}^{*}$) and utilities ($A_{a,0}^{*}$). Since both survival and continuation utility are larger for the young, then the young weight relatively more when $\psi=0$ than when $\psi=1$. Since younger individuals are relatively less affected by the pandemic, their WTP ($1-\lambda_{a,0}$) is lower, resulting in the planner exhibiting a lower WTP

$(1-\lambda_0)$ when $\psi = 0$ than when $\psi = 1$, as we verify in the calibrated exercise. Notice also that in the intermediate case in which $\psi = \gamma$ the weights $f_{a,0}(\psi)$ are given by

$$f_{a,0}(\gamma) = \frac{M_{a,0}\pi^*_{a,0}\left(A^*_{a,0}e^{\upsilon_a - ga - \gamma\Phi_a^2/2}\right)^{1-\gamma}}{\sum_a M_{a,0}\pi^*_{a,0}\left(A^*_{a,0}e^{\upsilon_a - ga - \gamma\Phi_a^2/2}\right)^{1-\gamma}},$$

which closely resemble an utilitarian planner in the expected utility case: utilities are weighted by population and the expression is linear in survival probabilities. We will consider this case below in the calibration, as it provides an intuitive benchmark for the planner's curvature parameter.

Second, in addition to providing insights on the role of parameter ψ, Proposition 2 also shows the role of cross-sectional inequality on the value of λ_0. Notice first from (25) that weights the weights $f_{a,0}(\psi)$ are only affected by life-cycle inequality, Φ_a^2. Since Φ_a^2 is increasing in age, the effective weight of each term in the summation gets smaller for higher a. Therefore for given $0 < \psi < 1$, the weight on the elderly, who are the ones most affected by the pandemic, gets downplayed in the computation of social welfare. In other words, the fact that inequality increases during the life-cycle tends to reduce the aggregate WTP to avoid the pandemic.

Third, the planner's $MRS(\psi)$ in (26) is computed using the notion that the WTP for a change in survival $\Delta\pi$ is equal to the marginal rate of substitution between consumption and survival times change in survival. Notice that equation (26) resembles a weighted average of the VSL-to-consumption ratio in equation (12). In this respect, $MRS(\psi)$ corresponds to a "social" VSL. Consider again the cases $\psi = 1$ and $\psi = 0$, for which the $MRS(\psi)$ is given by

$$MRS(1) = -\frac{1}{1-\gamma}\frac{1}{1-\beta}\frac{\sum_a M_{a,0}/\pi^*_{a,0}}{\sum_a M_{a,0}}$$

$$MRS(0) = -\frac{1}{1-\gamma}\frac{1}{1-\beta}\frac{\sum_a \left(M_{a,0}\left(\pi^*_{a,0}\right)^{\frac{1}{1-\gamma}}A^*_{a,0}e^{\upsilon_a - ga}\right)/\pi^*_{a,0}}{\sum_a M_{a,0}\left(\pi^*_{a,0}\right)^{\frac{1}{1-\gamma}}A^*_{a,0}e^{\upsilon_a - ga}}.$$

As mentioned above, again in this case when $\psi = 0$ the planner weights the young relatively more than when $\psi = 1$. Therefore for a given distribution of population shares, a linear planner ($\psi = 0$) is less willing to transfer resources from the elderly to the young than the log planner ($\psi = 1$).

2.8.3 The case of persistent consumption costs

In this section we consider the case in which the consumption costs are paid gradually over time, although the drop in survival from COVID-19 only happens in period $t = 0$. This case captures what in reality could be the prolonged costs of the pandemic in terms of income and consumption. Specifically, assume that

$$\ln \lambda_{a,t+1} = \rho \ln \lambda_{a,t} \tag{27}$$

for $t \geq 0$, where $\rho \in [0,1)$ characterizes the persistence of consumption costs. The scenario we considered so far, where welfare costs are computed for one period, corresponds to the case $\rho = 0$. When $\rho > 0$ the welfare costs are paid gradually, up to some period $t = t^*$ whose length is increasing in ρ. By construction, the WTP when $\rho > 0$ will be gradually declining over time, and the WTP at $t = 0$ will be smaller the larger the ρ is. In other words, the WTP at time $t = 0$ is smaller the larger t^* is.

The computation of the individual and the aggregate welfare costs from $t = 0$ to $t = t^*$ is more complicated than what was shown above for the case $\rho = 0$. The reason is that the $A_{a,0}\lambda_{a,0}$ in equation (19) and the $A_{a,0}\lambda_0$ in (21) cannot longer be expressed as functions of the baseline $A^*_{a,0}$ because consumption goes back to the baseline in $t = t^*$, not in $t = 1$. For the case $\rho > 0$ we have that

$$A_{a,0} = \begin{cases} \left[(1-\beta_a) + \beta_a \left[\pi_{a+1,1}^{\frac{1}{1-\gamma}} A_{a+1,1} \lambda_{a,0}^{\rho-1} e^{\mu_{a+1} - \theta\eta^2_{a+1}/2}\right]^{1-\sigma}\right]^{\frac{1}{1-\sigma}} & \text{for } \sigma \geq 0, \sigma \neq 1, \\ \left[\pi_{a+1,1}^{\frac{1}{1-\gamma}} A_{a+1,1} \lambda_{a,0}^{\rho-1} e^{\mu_{a+1} - \theta\eta^2_{a+1}/2}\right]^{\beta_a} & \text{for } \sigma = 1. \end{cases}$$

In this case one needs to compute separate sequences to obtain each $A_{a,0}$, as older individuals will reach terminal age a^* before time $t = t^*$, while younger individuals will reach this terminal age after period $t = t^*$. These different $A_{a,0}$ sequences reflect the fact that individuals experience the "transition" from $t = 0$ to $t = t^*$ at different ages over their life cycle.[7] Notice that by construction $\lambda_{a,t^*} = \lambda_{t^*} = 1$.

3 Calibration

We now calibrate the model to recent US data. The following parameters need to be calibrated: the inverse of intertemporal substitution (σ), mortality risk aversion (γ), consumption risk aversion (θ), the means and variances of the individual age-specific consumption shocks ($\mu_a - \eta^2_a/2$ and η^2_a), the mean and variance of the lognormal distribution for age-0 consumption ($\ln c_{0,t} - \Omega^2_t/2$ and Ω^2_t), the distribution of population by age (M_a), the conditional survival probabilities by age (π_a), and the age-specific discount factors (β_a).

The first step of our calibration involves computing and estimating exogenous parameters. First, we compute the conditional survival probabilities π_a using the CDC 2017 life tables for the US (National Vital Statistics, 2019). This is the most recent year available. Second, the distribution of population by age is obtained from the US Census for 2017. Both of these variables are displayed in Figure 1. Last, we use PSID data to estimate the means and variances of the individual age-specific consumption shocks ($\mu_a - \eta^2_a/2$ and η^2_a), as well as the mean and variance of the lognormal distribution for age-0 consumption ($\ln c_{0,t} - \Omega^2_t/2$ and Ω^2_t). Details for this step are provided below. We also set $g = 0$ to consider a stationary state.

[7] A technical appendix for this computation is available upon request.

Once we compute these exogenous parameters, our calibration strategy proceeds as follows. Since we would like to illustrate the role of σ in evaluating the distributional welfare effects of COVID-19, we choose different values of σ and report results for each. In the benchmark calibration we set $\sigma = 1$, but we also report a calibration with $\sigma = 2$, since these are commonly used values of σ in quantitative macro. In the benchmark calibration, we also set $\theta = 1$, so that $\sigma = \theta = 1$ and there is no separate role for aversion to consumption risk. We also report alternative calibrations with $\theta = 10$ and $\theta = 20$, which are commonly used values in the finance literature. We set $r = 2\%$ and back out the age-specific discount factors (β_a) so that the Euler equation holds for every age, as we explain in more detail below. Since β_a depends on all the exogenous parameters and σ, we recalibrate this age-specific discount factors for every σ we report. Last, we calibrate the mortality risk aversion parameter γ to match the ratio of the value of statistical life to consumption at age 40. We now provide additional details on the calibration.

3.1 Consumption process

We use 1999-2017 PSID biennial consumption data in order to compute cross-sectional means and variances by age. The PSID has measured consumption since 1999, accounting for about 70% of the consumption measured in the CEX, and has been found to provide cross-sectional life-cycle estimates of consumption that are very similar across surveys (Charles *et al.*, 2007). As it is standard practice, we omit the SEO component of the PSID. We measure non-durable consumption as the sum of the following categories: food (home, away and delivery), utilities, transportation (gasoline, auto insurance, vehicle repair, parking, bus, cab and other, and excluding vehicle loan and down payment), housing expenditures (rent and house insurance, excluding mortgages and property taxes), childcare, education and health. We impute homeowners rent as 6% of the home value reported by households.

We use the full sample of households observed during 1999-2017, whose heads of household's birth cohorts are between 1934 and 1991, and whose ages are between 23 and 77. Birth cohort and age restrictions are imposed so we have enough observations in each. We have a total of 45,290 observations over the ten PSID waves, or 4,529 households per wave. Although the PSID is longitudinal, the fact that it is biennial posses challenges in using individual-level consumption growth rates to directly estimate the mean (μ_a) and variance of the shock (η_a). Therefore we use cross-sectional consumption data in the estimation and back up the corresponding values for μ_a and η_a.

In order to construct cross-sectional means and variances of consumption by age, we first remove the effects of differences in family size across households by regressing log consumption on a full set of family size dummies. As it is standard in the literature, we construct 5-year overlapping age groups, and use the middle age in the bracket to name each age group. Since we restrict the sample to those ages 23 to 77, we have age groups between 25 and 75. Due to the biennial nature of the PSID data, we also construct non-overlapping 2-year cohort groups. In particular, every odd year cohort group is merged with the previous even cohort. Cohort groups are named with the even

cohort. We then have cohort groups from 1934 to 1990 (even years).

Using the residuals of the regression on family size dummies, we compute the mean and the variance (with appropriate PSID weights) for each age-group and cohort-group cell. Finally, we separately regress the mean and the variance of each cell on a full set of age-group dummies (age group 25 omitted). We call the coefficients of these dummies the "raw" age profile of means and variances of consumption. As it is standard in the literature, we also run two additional regressions to separately control for time (year 1999 omitted) and cohort effects (1934 cohort omitted).

Figure 2 displays the three estimated age profiles of the (cross-sectional) mean of log consumption for ages 25 to 75: raw, cohort effects removed, and time effects removed. All estimates show an increasing mean that starts flattening around age 50 and slightly decreases after age 70. This pattern is consistent with what Aguiar and Hurst (2013) document using 1980-2003 CEX data. In fact, what explains the flattening around age 50 is the inclusion of housing services in the measure of non-durable consumption (see their Figure 1).[8] Although the levels of the profile when time or cohort effects are removed are sensitive to the omitted year or cohort, this does not affect the calibration of the mean and variance of the individual consumption shock, which only depend on the changes age by age.

Figure 3 displays the estimated age profiles of the (cross-sectional) variance of log consumption. For the model's calibration we follow the literature and use the age profile when cohort effects are removed. For this case Figure 3 shows that the variance initially decreases, but then starts increasing rapidly between ages 40 and 55, after which it increases as a slower rate. Except for the initially decreasing portion, our estimates are consistent with those obtained by others using the CEX (see Heathcote *et al.* 2010 and Aguiar and Hurst, 2013). We find that the initially decreasing variance is driven by two factors: the inclusion of the very recent cohorts, particularly those born in the 1980s (Millennials); and the inclusion of childcare, education and health expenditures in the definition of non-durable consumption.[9]

In order to calibrate the model we use a Hodrick-Prescott filter to smooth out the estimates of the mean and variance of consumption by age as reported in Figures 2 and 3. We use the estimates when cohort effects are removed. We set age-0 in the model to be age 25 in the data. Figure 4 displays the calibrated mean and standard deviation of the individual consumption shock (ε_a). In particular it displays the age-varying mean (μ_a) and standard deviation (η_a), as well as constant values for these variables (constant μ and η). The constant $\mu_a = \mu$ and $\eta_a = \eta$ are used as robustness checks and are computed using a larger Hodrick-Prescott parameter.

For ages older than 75 we set $\mu_a = \mu_{75}$ and $\eta_a = \eta_{75}$. Therefore, we assume mean consumption continues to decrease slightly after age 75. Since we are interested in comparing the average WTP, $\overline{WTP}_0$, with the WTP of the median voter, and since those older than 18 years of age can vote, we need to extrapolate μ_a and η_a for ages between 18 and 24. Since we do not have enough observations

[8] When housing services are excluded mean log consumption starts declining at age 50 as in the standard consumption hump. Results available upon request.

[9] We are able to obtain the same increasing pattern of the variance of consumption starting at age 25 reported by Aguiar and Hurst (their Figure 1B) when we follow their definition of non-durable consumption and exclude childcare, education and health expenditures. In fact, we are also able to confirm with our data that food consumption is the main category driving the increasing variance of consumption by age.

to estimate μ_a and η_a before age 24, we assume that $\mu_a = \mu_{26}$ and $\eta_a = \eta_{26}$ for ages 19 through 25, and then use the estimated μ_{25} and η_{25} to construct μ_{18} and η_{18} so as to preserve the mean and variance of the lognormal distribution for age-0 consumption. We check that our results are not sensitive to this extrapolation.

3.2 Time horizon

The model so far has been specified with age-specific survival rates in infinite horizon. As discussed in Cordoba and Ripoll (2017), if horizon is finite so that dying becomes a sure state in finite time, and the utility upon dying is zero, then the utility of the alive is also zero whenever $\sigma \geq 1$. This result arises due to the implied strong complementary of consumption overtime. This implication of the model is not present when $0 < \sigma < 1$. A tractable way to close the model for any σ, while still maintaining homotheticity, is to assume that after age a^* the survival probability is constant but strictly positive, even if close to zero, a sort of "perpetual old" assumption. We use this approach in our benchmark calibration. We check the robustness of our results to an alternative formulation in which the utility upon dead is positive once death becomes a sure state: $D_{i,a,t} > 0$ for ages above a^*. Under this alternative, the horizon is strictly finite and non-homothetic.

The benchmark calibration implements the perpetual old assumption by assuming that after certain age a^* all parameters are age invariant. Suppose that for $a \geq a^* = 99$, $\mu_a = \mu$, $\eta_a^2 = \eta^2$, and $\pi_{a,t} = \pi$. In this case $A_{a,t}$ in equation (8) is constant and given by the terminal value of $\bar{A}$ specified in (10). Given $\bar{A}$ we compute the sequence of $A_{a,t}$ going backwards from age a^* and using equation (8).

3.3 Age-specific discount factors

We follow the standard practice of calibrating the discount factor to match a given target for the risk free interest of rate according to the Euler equation. In our life-cycle model Euler equations (14) are age specific due to the age-specific survival rates and a potentially age-specific consumption growth process. As a result, age-specific discount factors are needed to match the risk free interest target for every age group.[10] We set $r = 2\%$ and use Equation (14) to calibrate β_a given an assumption on annuity markets (δ). We present results for the case with no annuity markets ($\delta = 0$) and perfect annuities ($\delta = 1$). To back out β_a, solve for β_a as a function of β_{a+1},

$$\frac{1-\beta_a}{\beta_a} = \frac{1+r}{\delta\pi_a + 1 - \delta}\left(1-\beta_{a+1}\right)\pi_{a+1}^{\frac{1-\sigma}{1-\gamma}} e^{-\sigma\mu_{a+1}+\theta(1+\sigma)\eta_{a+1}^2/2}. \tag{28}$$

This equation suggests that β_a typically depends on age. A constant β would result only in specific cases such as: *(i)* $\delta = 0$ (no annuities), $\sigma = 1$ and $\sigma\mu_{a+1} = \theta\left(1+\sigma\right)\eta_{a+1}^2/2$; *(ii)* $\delta = 0$, $\sigma = 1$, $\mu_{a+1} = \mu$ and $\eta_{a+1} = \eta$; or *(iii)* $\delta = 1$ (perfect annuity markets), $\sigma = \gamma$ (expected utility case for

[10] Murphy and Topel (2006) and Cordoba and Ripoll (2017) use age-specific health indexes to match the target.

mortality risk), $\mu_{a+1} = \mu$ and $\eta_{a+1} = \eta$. But in general, for arbitrary δ and $\sigma \neq 1$, an age-dependent β_a has to be backed-out for the Euler equation of this model to hold at every age.

Like before, assume that for $a \geq a^* = 99$ parameters are constant so that $\mu_a = \mu$, $\eta_a^2 = \eta^2$, and $\pi_{a,t} = \pi$. In this case a terminal $\bar{\beta}$ is given by

$$\bar{\beta} = \frac{1}{1+r} \frac{\delta\pi + 1 - \delta}{\pi^{\frac{1-\sigma}{1-\gamma}} e^{-\sigma\mu + \theta(1+\sigma)\eta^2/2}}. \tag{29}$$

The sequence of β_a's can be then computed backwards from age a^* using equation (28) for any σ and δ. Equation (28) provides some insights into the profile of β_a. First, since the estimated μ_a and η_a in Figure 4 are small in magnitude, the quantitative role of term $e^{-\sigma\mu_{a+1} + \theta(1+\sigma)\eta_{a+1}^2/2}$ will be second order. If $\delta = 0$, then β_a mostly counteracts $\pi_{a+1}^{(1-\sigma)/(1-\gamma)}$, which is the only age-varying term, or $\beta_a \approx \pi_{a+1}^{(\sigma-1)/(1-\gamma)}$. For the most relevant case with $\sigma > 1$, since $\gamma < 1$ then we have that β_a is decreasing in age. If $\delta = 1$, then β_a counteracts $\pi_{a+1}^{(\gamma-\sigma)/(1-\gamma)}$ or $\beta_a \approx \pi_{a+1}^{(\sigma-\gamma)/(1-\gamma)}$. Here again, for $\sigma > 1$ we have that β_a is decreasing in age more strongly than in the no-annuity case, $\delta = 0$. Although β_a varies with age in our model , it is important to remember that the effective discount factor is not β_a but it is $\widetilde{\beta}_a \equiv \beta_a \pi_{a+1,t+1}^{(1-\sigma)/(1-\gamma)}$. From this perspective, allowing β_a to vary with age does not imply that individuals effectively become more patient or impatient as they age. We verify this property in our calibrated model.

3.4 Aversion to mortality risk

We follow Cordoba and Ripoll (2017) and calibrate the mortality risk aversion parameter γ to match the average ratio of the value of statistical life (VSL) to consumption at age 40.[11] From equation (12) above, this ratio is given by

$$\frac{VSL_{i,a,t}}{c_{i,a,t}} = \frac{1}{1-\beta_a} \frac{1}{1-\gamma} \frac{A_{a,t}^{1-\sigma}}{\pi_{a,t}}.$$

We calibrate γ so that the average ratio $VSL_{i,a,t}/c_{i,a,t}$ at age 40 is 150, a number similar to the one in Cordoba and Ripoll (2017) and Murphy and Topel (2006). With a level of consumption of \$45,000, this corresponds to a VSL of \$6.75 million. We also report results for the case of a VSL-to-consumption of 200, which corresponds to a VSL of \$9 million. This value is closer in magnitude to the one used by Hall *et al.* (2020).

Given our calibration strategy for β_a, it is instructive to discuss how it interacts with other determinants of the VSL in the equation above. Consider first the case in which $\sigma = 1$, there are no annuities ($\delta = 0$) and the parameters of the consumption shock process are constant across ages ($\mu_{a+1} = \mu$ and $\eta_{a+1} = \eta$). In this case, equation (14) implies that β_a is constant across ages, so that the VSL is only affected by $\pi_{a,t}$ and $c_{i,a,t}$. Therefore unless there is a strong drop in consumption

[11] We calibrate the model to match the VSL-to-consumption ratio rather than the VSL level because the PSID data we use in our calibration includes about 70% of the consumption categories of the CEX.

at older ages, the VSL tends to increase in age.

Next, continue to assume $\sigma = 1$ and constant μ and η, but now assume there are perfect annuity markets ($\delta = 1$). In this case, β_a is decreasing with age since β_a curbs the effect of $\pi_{a+1,t+1}^{(\gamma-\sigma)/(1-\gamma)}$ to satisfy the Euler equation (14). Therefore with perfect annuity markets the effect of β_a counteracts the effect of all $\pi_{a,t}$ and $c_{i,a,t}$ on the VSL. It turns out that in our calibrated model the effect of β_a dominates, making the VSL decreasing in age. This analysis illustrates how the presence of annuity markets influences the profile of the VSL through its effect on β_a.

Last, consider the case of $\sigma > 1$, a commonly used case in quantitative macro. In this case it is easy to see that, regardless of the annuity market structure, β_a decreases with age, which tends to make the VSL decrease with age. However, we find that in the calibrated model the effects from $A_{a,t}^{1-\sigma}$, $1/\pi_{a,t}$ and $c_{i,a,t}$ dominate, making the VSL increase with age.

3.5 Benchmark calibration

The benchmark with $\sigma = \theta = 1$ and $\delta = 0$ features two additional characteristics, both of which can be seen in equation (28). First, in the benchmark parameter γ does not affect the Euler equation and therefore does not affect the calibrated β_a. This implies that γ is completely identified from the targeted VSL/c ratio in equation (12). Second, with $\sigma = 1$ and $\delta = 0$, discount factor β_a only reflects the variations in μ_a and η_a. Since these variations are quantitatively small, β_a would be almost constant across ages. These multiple desirable features make the benchmark our preferred calibration.

For the benchmark calibration we obtain a mortality risk aversion parameter of $\gamma = 0.675$. Since $\sigma > \gamma$, the benchmark calibration implies a preference for late resolution of uncertainty regarding mortality. Similar results were found in Cordoba and Ripoll (2017). Figure 5 displays additional features of the benchmark calibration. The top two panels show the estimated cross-sectional mean and variance of log consumption from the PSID. Our age-varying μ_a allows us to estimate a reasonable consumption profile, with mean consumption falling slowly after age 70. As implied by the random walk process, the cross-section variance of consumption increases with age, a feature documented by others in the literature (Aguiar and Hurst, 2013).

The bottom-left panel of Figure 5 shows the implied Gini coefficients of consumption. As shown in Cowell (1995), the Gini coefficient for the lognormal distribution is given by

$$Gini_{a,t} = 2\Theta\left(\sqrt[2]{\left(\Omega_t^2 + \Phi_a^2\right)/2}\right) - 1$$

where $\Theta(.)$ is the cumulative density of a standard normal distribution. Recall from equation (3) that the variance of the lognormal cross-sectional distribution of consumption at age a is given by $\Omega_t^2 + \Phi_a^2$. As shown in Figure 5, while the Gini of consumption starts at around 0.26 at age 18, it goes up to 0.33 at age 60, suggesting our estimates from the PSID are a good representation of US consumption inequality.

Finally, the bottom-right panel of Figure 5 displays the calibrated age-varying β_a. As expected

for the benchmark calibration with $\sigma = 1$ and $\delta = 0$, β_a varies very little across ages, between 0.984 and 0.974.

4 Results

4.1 Distributional welfare costs of COVID-19

In order to simulate the effect of COVID-19 we follow Ferguson *et al.* (2020) and Menachemi *et al.* (2020). From Ferguson *et al.* (2020) we take the age-specific pre-lockdown fatality rates, but we adjust them to obtain the aggregate fatality rate from Menachemi *et al.* (2020). While Ferguson *et al.* (2020) predict an age-adjusted fatality rate of 1.1%, Menachemi *et al.* (2020) estimate it to be 0.58%, a more reasonable number. Since the log of age-specific fatality rates from Ferguson *et al.* (2020) are linear in age, we can easily adjust them to be consistent with the aggregate fatality rate of Menachemi *et al.* (2020). We use (log) interpolation to compute the decrease of survival probabilities for each age using the following fatality rates by age bracket from: 0.001% for ages 0-9; 0.003% for ages 10-19; 0.015% for ages 20-29; 0.04% for ages 30-39; 0.08% for 40-49; 0.3% for 50-59; 1.2% for 60-69; 2.7% for 70-79; and 4.9% for ages 80+. Both Ferguson *et al.* (2020) and Menachemi *et al.* (2020) provide estimates before the implementation of a lockdown, allowing us to compute the WTP to avoid COVID-19 absent any additional policy that would alter the death rates.

Table 1 summarizes the welfare costs of COVID-19 for different assumptions regarding annuity markets, σ, and the VSL-to-consumption ratio. The results on the table refer to the one-time WTP to avoid the pandemic ($\rho = 0$). The table reports the following three statistics: the average WTP ($\overline{WTP}_0$), which corresponds to the population-weighted average of age-specific WTP ($1 - \lambda_{a,0}$); the WTP of the median voter who is 46 years old (median age of population above 18); and the standard deviation of the age-specific WTP.

For our benchmark calibration with no annuity markets, $\sigma = 1$ and $VSL/c = 150$, we find that the costs of COVID-19 are high: the average WTP to avoid the pandemic is 41.2% of one-year consumption. More interestingly, we find that the median voter is willing to pay much less, 22.7%. In addition, we find substantial disagreement across individuals: the standard deviation of the distribution of age-specific WTP is 41.2% of one-year consumption.

Figure 6 displays some relevant variables for the benchmark calibration. The top-left panel reports VSL/c, which is calibrated to 150 at age 40. As shown, the calibrated VSL/c is U-shape in age. Recall from equation (12) that when $\sigma = 1$ the VSL/c is independent of $A_{a,t}$. Since π_a decreases with age, but β_a also decreases with age, there are two competing forces: for younger ages the effect of β_a dominates, making the VSL/c decrease with age, but at older ages the VSL/c increases with age because the effect of lower π_a is stronger. Notice in particular that VSL/c increases after age 75 because β_a becomes constant after that age: with $\delta = 0$ and $\sigma = 1$, the calibrated β_a only captures the changes in μ_a and η_a, but these are assumed to be constant after age 75. The top-right panel of Figure 6 displays the percentage change of continuation utility due

to the COVID-19 pandemic. Although a one-year event, the welfare change is large for the oldest who face relatively large survival losses, much larger than those due to influenza.

The bottom-left panel of Figure 6 displays the distribution of age-specific WTP $(1 - \lambda_{a,0})$, which sharply increases with age and is close to 100% after age 80. Recall from the formula of the one-time age-specific WTP (20), that when $\sigma = 1$ and $\delta = 0$, since β_a is almost constant across ages, the WTP is mostly driven by the age-specific change in survival $\Delta\pi_{a,0} < 0$. As the change in the survival probability is substantially different across ages, going from almost zero at age 20, to 0.003 at age 50, to 0.049 at ages 80+, the dispersion of the WTP across ages is also substantial.

Finally, the bottom-right panel of Figure 6 shows the contribution of the age-specific WTP to the average WTP $(\overline{WTP}_0)$: although the oldest are willing to pay a large percentage of their consumption to avoid COVID, they represent a lower percentage of the population (Figure 1), contributing less to the 41.2% average in Table 1. At the opposite end, although the youngest represent a larger share of the population, they are willing to pay the least, contributing less to the average WTP. It is those around age 60 the ones who contribute the most to the average WTP.

As shown in Table 1, for the case of no annuity markets, when $\sigma = 2$ the WTP are lower: the average WTP is 34.3% of one-year consumption, the median voter's WTP is 20.5% and the standard deviation of the WTP is 35.2%. Although lower, these numbers still indicate larger welfare costs of COVID-19. The top-left panel of Figure 7 illustrates this point by comparing the distribution of age-specific WTP when $\sigma = 1$ and $\sigma = 2$ for the case of no annuities. As seen in the figure, when $\sigma = 2$ the WTP is lower after age 40. The reason can be traced to asset pricing and the calibration of β_a: for the Euler equation to hold at older ages with the given consumption process, the elder would have to become less patient, or β_a has to be lower. This effect decreases the WTP of the elder.

Table 1 also reports the WTP for the case of perfect annuity markets. The main observation for this case is that relative to the case with no annuities, the average WTP, the WTP of the median voter and the standard deviation of the WTP are lower. These lower statistics can again be attributed to the WTP of the elder. This is shown on the top-right panel of Figure 7 for the case $\sigma = 1$: with perfect annuity markets, the WTP is lower starting at age 40, but particularly so after age 75 when the WTP declines with age. Recall from equation (14), that with $\sigma = 1$ and $\delta = 1$, the effective discount factor is $\widetilde{\beta}_a \equiv \beta_a/\pi_{a+1}$. Therefore, the β_a that satisfies the Euler equation almost exactly cancels the effect of decreasing π_{a+1}, or $\beta_a \approx \pi_{a+1}$. The intuition here is that with perfect annuities the return from saving, for the survivors, is higher than with no annuities which would induce higher consumption growth. For the assumed consumption process to continue satisfy Euler Equations, individuals would need to be more impatient, at every age, in the model with annuities. As seen in Figure 1, π_a is almost one at age 18, but it goes down to about 0.71 at age 100. This implies that β_a differs more across ages when $\sigma = 1$ and $\delta = 1$. This lower β_a for the elder reduces the WTP relative to the case with no annuities.

Finally, Table 1 also shows the corresponding WTP for the case $VSL/c = 200$. Naturally, the higher the VSL/c the larger the average WTP as well as the WTP for the median voter. This scenario also displays substantial dispersion of the distribution of age-specific WTP. The scenario

with $VSL/c = 200$ more closely resembles the calibration in Hall *et al.* (2020). The most relevant comparison corresponds to the case with perfect annuity markets and $\sigma = 2$. In this case we obtain an average WTP of 34.3%. They obtain a comparable value of 27.7%, although their is lower because they simulate a lower aggregate fatality rate (0.44%).

4.2 Robustness checks to alternative parameters

We check the robustness of our results to alternative parameter values. Table 2 summarizes the results when we change certain parameter values, one at a time. In particular, we report results for alternative values of θ, for a calibration of the model that assumes constant $\mu_a = \mu$ and $\eta_a = \eta$, and for a finite-horizon calibration that considers $D_{i,a,t} \neq 0$. Table 2 includes the benchmark calibration with $\sigma = \theta = 1$.

Table 2 shows that alternative values of θ have little effects on the WTP results. For both $\theta = 10$ and $\theta = 20$ the average and the median voter WTP are higher than in the calibrated benchmark ($\sigma = \theta = 1$) with or without annuity markets, but the differences are quantitatively small. Recall from equation (20) that when $\sigma = 1$ the age-specific WTP depends on $\Delta\pi_{a,0}$, the calibrated β_a, and the calibrated γ. As seen in equation (28), parameter θ affects the calibrated β_a only through term $e^{-\sigma\mu_{a+1}+\theta(1+\sigma)\eta^2_{a+1}/2}$: a larger θ results in a lower β_a. However, since η^2_a is very small in magnitude, the overall quantitative effect of θ on β_a is small. For the case of no annuity markets, with $\theta = 10$ we obtain a calibrated $\gamma = 0.877$, while for $\theta = 20$ we obtain $\gamma = 0.894$. In this respect, higher consumption risk aversion results in a higher mortality risk aversion in the calibration. The bottom-left panel of Figure 7 displays the full distribution of age-specific WTP for $\theta = 1$ and $\theta = 10$. As shown, there is little difference, although those between ages 50 and 60 have slightly higher WTP when $\theta = 10$. This is the case because the standard deviation of the consumption shock (η_a) is higher at those ages (Figure 4).

Table 2 also shows that for the case of constant $\mu_a = \mu$ and $\eta_a = \eta$ (but still maintaining $\sigma = \theta = 1$), the WTP results only change slightly for both the no annuities and the perfect annuity cases. The bottom-right panel of Figure 7 displays the full distribution of age-specific WTP for constant and age-varying μ_a and η_a for the case of no annuities. This is not surprising since the age-varying estimated values of μ_a and η_a only affect the results through term $e^{-\sigma\mu_{a+1}+\theta(1+\sigma)\eta^2_{a+1}/2}$, which tends to have a quantitatively second-order effect. The bottom-right panel of Figure 7 displays the full distribution of age-specific WTP for constant and age-varying μ_a and η_a. Having said this, it is important to notice that the decrease in survival rates from COVID-19 are so large, particularly for the elder, that the WTP results are mostly driven by $\Delta\pi_{a,0}$. This is specially true in the case of $\sigma = 1$. If we were to use our model to compute the WTP to avoid pandemics with smaller effects in survival probabilities, other features of our model like age-varying μ_a and η_a would play a larger role.

In the benchmark calibration we closed the model by assuming that the horizon is infinite, but that after certain age $a \geq a^* = 99$, the survival probability is constant and very low, but positive ($\pi_{a,t} = \pi$). This "perpetual old" assumption avoids that lifetime utility becomes zero when $\sigma \geq 1$.

Table 2 reports the robustness of the WTP results to an alternative way of closing the model, namely one in which the horizon is finite. In this case, rather than assuming that the utility upon death $D_{i,a,t} = 0$, we impute a positive value of consumption in the dead state for ages after which the survival probability is zero. In particular, a non-homotheticity in utility ($D_{i,a,t} > 0$) is assumed for ages over $a^* = 99$.

The main difficulty of solving this case is that with the non-homotheticity after age 99, utility is no longer linear in consumption. In this finite-horizon model, utility in period a^* is given by $\overline{V}_{a^*} = D_{a^*}$ while utility in period $a^* - 1$ is given by

$$V_{a^*-1} = (c_{a^*-1})^{1-\beta_{a^*-1}} (D_{a^*})^{\beta_{a^*-1}}.$$

It can be shown that in this case utility is of the form

$$V_{i,a,t} = A_{a,t} (\lambda_{a,t} c_{i,a,t})^{\phi_a}$$

where $\phi_a = 1 - \beta_a + \beta_a \phi_{a+1}$ and

$$A_{a,t} = \left(\pi_{a+1,t+1}^{\frac{1}{1-\gamma}} A_{a+1,t+1} \left(\frac{\lambda_{a+1,t+1}}{\lambda_{a,t}} e^{\mu_{a+1} + (\phi_{a+1}(1-\theta)-1)\eta_{a+1}^2/2} \right)^{\phi_{a+1}} \right)^{\beta_a}.$$

The last two recursions can be solved using the following terminal conditions $\phi_{a^*-1} = 1 - \beta_{a^*-1}$, and $A_{a^*-1,t^*-1} = (D_{a^*})^{\beta_{a^*-1}}$ where $D_{a^*} > 0$.

As shown in Table 2, our results are robust to this alternative way of closing the model. We set $D_{a^*} = 0.001$ and find that for the benchmark calibration with $\sigma = \theta = 1$, the average and the median voter's WTP are virtually identical to the benchmark for both the cases with no annuities and annuities. These results are robust to other values of D_{a^*}, but we set $D_{a^*} = 0.001$ to guarantee a decreasing profile for $A_{a,t}$ with age.

4.3 Persistent consumption costs

In this section we extend the computations to consider the case in which the consumption costs are persistent ($\rho > 0$). In this case, although the pandemic decreases the survival probabilities only in year $t = 0$, the consumption costs persist up to time $t = t^*$. We compute the WTP each year from 0 to t^*. Table 3 presents the results for this extension, in particular for the cases $\rho = 0.5$ and $\rho = 0.9$. Value $\rho = 0.5$ implies $t^* \approx 8$, while $\rho = 0.9$ implies $t^* \approx 30$. The former illustrates lingering effects of a crisis similar to the 2007-2009 Great Recession, while the latter illustrates the limiting case as $\rho \to 1$.

Table 3 reports the WTP statistics for $t = 0$, while Figure 8 includes plots for the full transition of WTP from $t = 0$ to t^*. As seen in Table 3, the WTP during period $t = 0$ falls as the persistence increases: individuals are willing to pay less to avoid the pandemic if they will be paying a percentage

of consumption for a longer period of time. For instance, for the case with no annuities, the average WTP at $t = 0$ falls from 41.2% when all costs are paid in one period ($\rho = 0$), to 30.9% when $\rho = 0.5$ and to 12.2% when $\rho = 0.9$. As seen in the table, a similar pattern is observed for the WTP of the median voter. Finally, the disagreement among individuals of different ages still persists as seen on the standard deviation of the WTP across ages.

The overall message of Table 3 is that even when the consumption costs are persistent, the average WTP to eliminate a one-year COVID-19 pandemic is still large. For the benchmark economy with $\sigma = \theta = 1$ and no annuity markets, simulating a $\rho = 0.5$ as a potentially realistic scenario for a lengthy recovery results in an average WTP of 30.9% of time-0 consumption. In the presence of annuity markets this number falls to 23.4%, still a large welfare cost.

The top-left panel in Figure 8 shows the distribution of the WTP by age at time $t = 0$. As ρ increases, the WTP falls for all ages. In the limit, with $\rho = 0.9$, those above age 80 are willing to pay much less than when $\rho = 0$, between 55 and 65% of time-0 consumption. The top-right panel in Figure 8 displays the average WTP from period $t = 0$ to $t = t^*$ for different values of ρ. Notice that by construction the WTP falls over time. The bottom-left panel shows how the WTP for the median voter evolves over time, while the bottom-right displays the case of the average 65-year old individual.

4.4 Social welfare costs

Table 4 presents the aggregate costs of COVID-19 from the perspective of the planner. All exercises on the table use the benchmark calibration with $\sigma = \theta = 1$ and $VSL/c = 150$. Recall from section 2.7 that the planner's curvature parameter is given by ψ. When $\psi = 0$ the planner has linear preferences and is indifferent to the inequality across individuals of different ages. We also consider other cases with more curvature, including one in which $\psi = \gamma$ and one in which $\psi = \sigma = \theta = 1$. Recall that the EZW utility representation we have three curvature parameters: σ, γ and θ, making it reasonable to explore how the aggregate costs of the pandemic change when ψ equal each of these values.

As seen in Table 4, the aggregate costs of the pandemic are quite different across various values of ψ. As ψ goes from $\psi = 0$ to $\psi = 1$, the aggregate WTP increases. As mentioned in the discussion of equation (23) and Proposition 2, when $\psi = 0$ the planner gives the highest relative weight to the young and is willing to pay less to avoid the pandemic. In contrast, as ψ increases towards one, the planner gives relatively less weight to the young and is willing to pay more to avoid COVID-19.

Table 4 also shows a difference between the cases with and without annuities. Overall, the aggregate WTP is higher when annuities are perfect. This result emerges because the calibrated profiles of A_a are different in these two cases. Specifically, in the absence of annuities A_a rapidly declines with age, while the decline is much slower when annuities are perfect.

Our preferred scenario is that with no annuities, since annuities are rarely used in practice. In this case, a social planner who does not care about inequality ($\psi = 0$) would be willing to pay only 4.1% to avoid COVID-19. However, this number changes rapidly as inequality aversion increases:

when $\psi = \gamma = 0.675$, the aggregate WTP more than doubles to 10.8%. The WTP of the median voter, 22.7%, roughly coincides with that of the planner when $\psi = 0.81$. In addition, the average WTP, 41.2%, coincides with that of the planner when $\psi = 0.9$. As we continue to increase ψ up to the log-planner with $\psi = \sigma = \theta = 1$, the aggregate WTP becomes 72.7%. In sum, in the case of no annuities the aggregate WTP is sensitive to the planner's inequality aversion, spanning a wide range of values, including the average WTP and the WTP of the median voter.

As mentioned before, the case $\psi = \gamma = 0.675$ is interesting because the planner's weight on age-specific WTP is linear in both the population share and the survival probability, getting closer to the case of a utilitarian planner with expected utility. Although the planner is willing to pay relatively little in this case, it is important to emphasize that higher levels of inequality aversion can also rationalize a planner who is willing to pay as much as the median voter and as much as the average.

4.5 Full recessions

All our computations so far have been made using the pre-lockdown fatality rates from Ferguson *et al.* (2020) and Menachemi *et al.* (2020). These are the appropriate fatality rates to compute the WTP to fully avoid the pandemic. In practice, pandemics cannot be fully avoided and in the case of COVID-19 lockdown policies have been implemented to avoid the loss of lives, which have resulted in contraction of economic activity. In this section we use our model to calculate "full recessions", which correspond to a combined measure of consumption loss and lives lost. Our notion of full recession is analogous to that of full income, which has been used in cross-country literature to compute a measure of welfare that augments income with life expectancy (Jones and Klenow, 2016; Cordoba and Ripoll, 2017).

We compute full recessions using the following version of individual utility in equation (18) for the case $\rho = 0$

$$\overline{V}\left(\boldsymbol{\pi}^*_{a,0} + \Delta\boldsymbol{\pi}_{a,0}, \left[\lambda^r_{i,a,0}, \mathbf{1}\right], c_{i,a,0}\right) = \overline{V}\left(\boldsymbol{\pi}^*_{a,0}, \left[\lambda^{fr}_{i,a,0}, \mathbf{1}\right], c_{i,a,0}\right),$$

where $\lambda^r_{i,a,0}$ corresponds to the time-0 change in consumption due to the associated lockdown recession. The left-hand side is the expected utility of individual i under the one-year lockdown recession and the change in survival probabilities from the pandemic. The right-hand side is the utility without the pandemic but with a level of effective initial consumption adjusted by $\lambda^{fr}_{i,a,0}$ and that delivers the same utility as the left-hand side. We assume $\lambda^r_{i,a,0} = \lambda^r_0$, so that the lockdown recession is the same for all individuals and all ages. Similar to what we found before, $\lambda^{fr}_{i,a,0}$ only depends on age. We refer to $\alpha^{fr}_{a,0} = (1 - \lambda^{fr}_{a,0})$ as the full recession in the sense that it is the drop in consumption that would reflect both the change in survival due to the pandemic and the economic contraction from the lockdown.

Table 5 presents the value of the full recessions in terms of the percentage of consumption lost. We consider three possible recession scenarios, with drops of 3%, 5% and 10% of consumption

during the year of the pandemic. We also consider different death rates, with 300 and 400 thousand lives lost. Table 5 includes the computation of full recessions from the perspective of the average individual (population weighted), the median voter, the planner with $\psi = \gamma$ and the planner with $\psi = 0.9$. The table uses our preferred calibration with no annuity markets, $\sigma = \theta = 1$ and age-varying discount factors.

As shown in Table 5, a 5% contraction in consumption during the year of the pandemic combined with a lost of 300,000 lives corresponds to a 18.8% full recession for the average individual, a 10.6% full recession for the median voter, a 7.1% full recession for a planner with a relatively low aversion to inequality, and 12.4% for the planner with higher inequality aversion. When the lives lost go up to 400,000, the corresponding full recessions are 21.9% for the average, 12.4% for the median voter, 7.7% for the lower-inequality aversion planner, and 14.8% for the planner with higher aversion. In order to interpret these findings, first notice that for the planner with lower inequality aversion, a 5% recession becomes a 7.7% full recession when 400,000 lives are lost: the additional 2.7% corresponds to the planner's valuation of 400,000 lives lost. Using a back-of-the-envelope computation where US GDP is $22 trillion, this roughly corresponds to a planner's valuation of about $1.5 million per life lost. Pindyck (2020) has suggested $1 million as a reasonable VSL in evaluating social preferences. But for the planner with higher aversion to inequality, this valuation becomes $5.4 million per life lost. On the other hand, the median voter, who is age 46, values each life at roughly $4.1 million, while on average each life lost is valued at $9 million. The reason why the average individual values life much more than the low inequality aversion planner is that those above age 75 would be willing to pay a sizeable amount to save their lives, since the pandemic reduces their already shorter life spans. The planner with low inequality aversion ($\psi = \gamma = 0.675$) places an effectively lower weight on older individuals, whose continuation utility is lower. This planner's valuation of life aligns with that of younger individuals whose survival is affected little by the pandemic. Alternatively, the planner with higher inequality aversion places a higher weight on older individuals and is willing to pay more.

4.6 Consumption-lives frontier

In this section we use our model to compute the trade-off between consumption and lives along a frontier. We compute this frontier by using the following version of equation (18)

$$\overline{V}\left(\boldsymbol{\pi}^*_{a,0} + \Delta\boldsymbol{\pi}_{a,0}, c_{i,a,0}\right) = \overline{V}\left(\boldsymbol{\pi}^*_{a,0} + \varphi\Delta\boldsymbol{\pi}_{a,0}, \left[\lambda_{i,a,0}, \mathbf{1}\right], c_{i,a,0}\right). \tag{30}$$

The left-hand side is the expected utility of individual i under the COVID-19 death rates at pre-lockdown levels from Ferguson *et al.* (2020). The right-hand side is the utility under a milder pandemic, with $\varphi \in [0, 1]$, but with a level of initial consumption adjusted by $\lambda_{i,a,0}$ that delivers the same utility as the left-hand side. At one extreme, when $\varphi = 0$, we obtain the WTP $(1 - \lambda_{i,a,0})$ to avoid all COVID deaths, which are the same ones reported in Table 1. At the other extreme, with $\varphi = 1$, by definition we obtain $\lambda_{i,a,0} = 1$ or a WTP of zero: if no deaths can be avoided, then

no consumption would be traded off.

It can easily be shown that the solution for $\lambda_{i,a,0}$ in (30) is then given by the following version of equation (20)

$$\lambda_{i,a,0} = \lambda_{a,0} = \begin{cases} \left[1 + \left[\left(\frac{\pi^*_{a,0}+\Delta\pi_{a,0}}{\pi^*_{a,0}+\varphi\Delta\pi_{a,0}}\right)^{\frac{1-\sigma}{1-\gamma}} - 1\right] \frac{\left(A^*_{a,0}\right)^{1-\sigma}}{1-\beta_a}\right]^{\frac{1}{1-\sigma}} & \text{for } \sigma \geq 0, \sigma \neq 1, \\ \left(\frac{\pi^*_{a,0}+\Delta\pi_{a,0}}{\pi^*_{a,0}+\varphi\Delta\pi_{a,0}}\right)^{\frac{1}{1-\gamma}\frac{1}{1-\beta_a}} & \text{for } \sigma = 1. \end{cases}$$

Table 6 and Figure 9 present our computation of the consumption-lives frontier. The table uses our preferred calibration with no annuity markets, $\sigma = \theta = 1$ and age-varying discount factors. As mentioned, the average individual is willing to trade off 41.2% of one-year consumption to avoid all COVID deaths, while the median voter would pay 22.7%, the low inequality aversion planner ($\psi = \gamma$) 10.8%, and the high inequality aversion planner 41.0%. These are the same values reported for our preferred calibration in Tables 1 and 4. At the other end, if no COVID deaths can possibly be avoided, 1.9 million people would die according to the fatality rates of Menachemi *et al.* (2020), a 0.58% death rate. In this case the WTP would be zero, as shown in the last row of Table 6.

As can seen in Figure 9, the frontier for the average individual is relatively more concave than that of the median voter or the planners. Starting at 1.9 million deaths, the average individual is initially willing to reduce consumption sharply to save lives: the first 1.1 million lives are traded-off by a 32.9% consumption cut in the year of the pandemic. But the trade-off is smaller for the next 800 thousand lives, adding a extra 8% consumption cut. The concavity of the average frontier reflects the large variation in willingness to pay across individuals of different ages. For the median voter, the frontier is less concave, sacrificing 13.9% of consumption to save the first 1.1 million lives, an additional 8.8% for the next 800 thousand people.

Figure 9 also shows how, although the consumption-lives frontier for the high inequality aversion planner starts and ends at the same point as the average, the planner's is more linear and falls faster. This reflects the fact that the planner is weighting individuals differently than with just population weights: the planner also takes into account the survival probability and the continuation utility of individuals.

5 Discussion and relationship to the literature

In this section we further discuss the properties of our model, in particular its predictions regarding the VSL, and we also relate our results to the literature on the costs of the COVID-19 pandemic. For this purpose, we consider a version of the model with only mortality risk but no heterogeneity or consumption risk. We can drop the i subscripts, and for simplicity we also drop subscript t in this section.

5.1 VSL and the value of a year-life

Focusing on the case $\sigma \neq 1$, utility is given by

$$\overline{V}_a = \pi_a^{\frac{1}{1-\gamma}} V_a \text{ where } V_a = \left[(1-\beta_a)\, c_a^{1-\sigma} + \beta_a \pi_{a+1}^{\frac{1-\sigma}{1-\gamma}} V_{a+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

Iterating forward on the last equation above the following extensive form for V_a can be obtained:

$$V_a = \left[\sum\nolimits_{s=a}^{\infty} S_{a,s}\, (1-\beta_s)\, c_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{31}$$

where $S_{a,a} = 1$ and $S_{a,s} = \prod_{j=a}^{s-1} \beta_j \pi_{j+1}^{(1-\sigma)/(1-\gamma)}$ for $s > a$. The budget constraint is given by

$$b_a = \sum\nolimits_{s=a}^{\infty} P_{a,s}\, [c_s - y_s] \tag{32}$$

where $P_{a,a} = 1$ and $P_{a,s} = \prod_{j=a+1}^{s} p_j = (1+r)^{a-s} \prod_{j=a+1}^{s} (\delta \pi_j + 1 - \delta)$ for $s > a$. Using the Euler equation from maximizing V_a subject to lifetime budget constraint (32) results in

$$V_a^{1-\sigma} = (1-\beta_a)\, c_a^{-\sigma} \sum\nolimits_{s=a}^{\infty} P_{a,s} c_s. \tag{33}$$

The expression above can be used to express the VSL in equation (11) as

$$VSL_a = \frac{1}{1-\gamma} \frac{1}{\pi_a} \left[\sum\nolimits_{s=a}^{\infty} P_{a,s} c_s \right], \tag{34}$$

and defining the value of a year-life as $v_a = c_a/(1-\gamma)$ we can express the VSL as the discounted sum of the values of year-life

$$VSL_a = \frac{1}{\pi_a} \left[\sum\nolimits_{s=a}^{\infty} P_{a,s} v_s \right],$$

which is a standard expression in the literature (see Murphy and Topel, 2006; and Cordoba and Ripoll, 2017). In order to see how the assumptions about annuities affect the VSL, assume momentarily that $v_s = v$. In this case, if there are no annuity markets ($\delta = 0$) we have that

$$VSL_a = \frac{v}{\pi_a} \frac{1+r}{r},$$

which implies that the VSL strictly increases with age, a novel insight. Most papers in the literature consider the case of perfect annuity markets, even though the use of annuities is quite limited in practice.

In the case of perfect annuity markets ($\delta = 1$), we have that

$$VSL_a = \frac{v}{\pi_a} \sum\nolimits_{s=a}^{\infty} \left(\prod\nolimits_{j=a}^{s} \pi_j \right) \equiv \frac{v}{\pi_a} LE_a, \tag{35}$$

where LE_a is the remaining life expectancy at age a. In this case, the VSL could fall with age if LE_a falls faster than π_a.

5.2 Relationship to the literature

Our welfare computation is most related to Hall *et al.* (2020) and complements their analysis. We now show how we can obtain their framework and their main formulas for the costs of the pandemic from ours. For this purpose, in addition to eliminating consumption risk and assuming $\theta = \sigma$, we add the assumption that $\gamma = \sigma$ to obtain expected utility. As in Hall *et al.* we also assume that $\beta_a = \beta$ and that consumption is constant so that $c_a = c$. In this case, equation (31) reads

$$V_a^{1-\sigma} = (1-\beta)\, c^{1-\sigma} \sum\nolimits_{s=a}^{\infty} \left(\prod\nolimits_{j=a}^{s-1} \beta_j \pi_{j+1} \right) \equiv (1-\beta)\, c^{1-\sigma} ELE_a \tag{36}$$

which allows us to write the VSL in equation (11) as

$$VSL_a = \frac{c}{1-\sigma} \frac{ELE_a}{\pi_a} \tag{37}$$

where ELE_a stands for "effective" life expectancy. Equation (36) states that continuation utility equals the utility flow $(1-\beta)\, c^{1-\sigma}$ times the effective life expectancy. This equation is almost the same as in Hall *et al.* (2020), except that theirs is simplified by assuming $\beta = 1$, so that instead of obtaining ELE_a they directly obtain remaining life expectancy LE_a as follows

$$VSL_a = \frac{c}{1-\sigma} \frac{LE_a}{\pi_a}. \tag{38}$$

Our EZW framework does not allow for the $\beta = 1$ simplification because due to the CES nature of our utility function, $(1-\beta)$ is the weight of the current utility flow. However, notice from equation (35) that if we make $v = c/(1-\sigma)$, then with perfect annuity markets we can obtain the same expression as in (38). Therefore, when there are perfect annuity markets, no consumption risk, and $\theta = \gamma = \sigma$, we obtain similar expressions for the VSL as in Hall *et al.*

We now compare the costs of a one-year COVID-19 pandemic with Hall *et al.* (2020). Recall that the pandemic takes place at time $t = 0$. Referring back to equation (20), which we rewrite here in logs for the case $\sigma \neq 1$ and $\gamma = \sigma$

$$\ln \lambda_{a,0} = \frac{1}{1-\sigma} \ln \left[1 + \left[\frac{\pi^*_{a,0} + \Delta\pi_{a,0}}{\pi^*_{a,0}} - 1 \right] \frac{\left(A^*_{a,0}\right)^{1-\sigma}}{1-\beta} \right],$$

define the WTP as $\alpha_{a,0} = 1 - \lambda_{a,0}$ and notice that for small $\alpha_{a,0}$, $\ln \lambda_{a,0} \simeq -\alpha_{a,0}$. Linearizing $\ln \lambda_{a,0}$ around $\Delta\pi_{a,0} = 0$ we have that

$$\ln \lambda_{a,0} \simeq \frac{1}{1-\beta} \frac{1}{1-\sigma} \frac{\left(A^*_{a,0}\right)^{1-\sigma}}{\pi^*_{a,0}} \Delta\pi_{a,0}$$

which using (12) can be written as

$$\alpha_{a,0} \simeq -\ln \lambda_{a,0} \simeq -\frac{VSL_{a,0}}{c_{a,0}} \Delta\pi_{a,0}$$

and using (37) we obtain

$$\alpha_{a,0} \simeq -\ln \lambda_{a,0} \simeq -\frac{\Delta\pi_{a,0}}{1-\sigma} \frac{1}{\pi_{a,0}} ELE_{a,0},$$

an expression similar to Hall *et al.* (2020) for the cost of the pandemic. Since they assume $\beta = 1$ they directly obtain $LE_{a,0}$ instead of $ELE_{a,0}$. The other difference is that term $(1/\pi_{a,0})$ is present in our formula but not in Hall *et al.* (2020) due to the timing of our model, where the pandemic hits at time $t = 0$ and directly affects the survival probability that period.

In terms of quantitative results, Hall *et al.* (2020) obtain welfare costs of 27.7% for their preferred calibration with $\sigma = 2$, although for a lower aggregate fatality rate (0.44%). Our model and theirs differ in many dimensions, but as shown in Table 1 for the case of $\sigma = 2$, we obtain average welfare costs between 29.7% and 34.3% depending on the target VSL/c. Although these results are overall similar, the main differences can be traced to the distributional welfare effects. Our analysis here has emphasized the WTP of the mean voter and the standard deviation of the WTP across ages. We have also modeled the planner's social welfare function to include aversion to inequality. Because of the nature of our preferences, in our model the average WTP across ages does not coincide with the aggregate WTP by the planner, even when the planner does not care about inequality.

A fundamental underlying difference between our framework and that of Hall *et al.* (2020) is that our utility is homothetic while theirs is not. An implication of this difference is that in our homothetic framework the young, who enjoy lower levels of consumption, do not value survival relatively less than individuals enjoying higher consumption. The income effects of non-homothetic utility are eliminated. Finally, our framework exhibits a preference for late resolution of mortality uncertainty and a preference for early resolution of consumption uncertainty, while in Hall *et al.* individuals are indifferent in both regards. An implication of this difference is that the elderly in our framework have a relatively higher WTP to avoid the pandemic because the marginal valuation of survival decreases with the level of survival, and theirs is particularly reduced by COVID-19. It is on these distributional effects of COVID-19 where our framework has distinct implications.

6 Concluding comments

The COVID-19 pandemic has brought back to light the difficult trade-offs between consumption and survival, as well as the inherent conflicts among individuals who bear very different costs from a rare but potentially fatal shock. Our analysis provides a framework to study these trade-offs as well as the distributional welfare effects of pandemics. We find that the welfare costs of COVID-19 are large; that the differences in the willingness to pay to avoid the pandemic across individuals of different ages is substantial; and that the willingness to pay of the median voter is less than the average. Our framework can also rationalize planners with very different perspectives about the pandemic.

We hope our framework is useful to analyze the trade-offs between consumption and survival from other mortality shocks. We think it provides a meaningful and rigorous integration of traditional models of economic shocks, with models of mortality shocks. In future work we plan to extend our model to include other aspects that may affect the value of life, such as leisure. We also plan to further explore the value of life from the perspective of the social planner.

References

[1] Aguiar, M. and Hurst, E. (2013), "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121, 437-492.

[2] Becker, G., Philipson, T. and Soares, R. (2005), "The Quantity and Quality of Life and the Evolution of World Inequality", *American Economic Review*, 95, 277-291.

[3] Charles, K., Danziger, S., Li, G. and Shoeni, R. (2007), "Studying Consumption with the Panel Study of Income Dynamics: Comparisons with the Consumer Expenditure Survey and an Application to the Intergenerational Transmission of Well-Being," *Finance and Economics Discussion Series*, 2007-16, Federal Reserve Board.

[4] Cordoba, J.C., and Ripoll, M. (2017), "Risk Aversion and the Value of Life," *Review of Economic Studies*, 84, 1472-1509.

[5] Cordoba, J.C., and Verdier, G. (2008), "Inequality and Growth: Some Welfare Calculations," *Journal of Economic Dynamics & Control*, 32, 1812-1829.

[6] Cowell, F. (1995), *Measuring Inequality*. Oxford University Press, 2nd edition.

[7] Epstein, L., and Zin, S.E. (1989), "Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns: A Theoretical Framework", *Econometrica*, 57, 937-969.

[8] Epstein, L. and Zin, S.E. (1991), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis", *Journal of Political Economy*, 99, 263-286.

[9] Ferguson, N., Laydon, D., Nedjati-Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba, Z., Cuomo-Dannenburg, G. *et al.* (2020), "Report 9: Impact of Non-Pharmaceutical interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand," Imperial College COVID-19 Response Team, March.

[10] Hall, R., Jones, C.I. and Klenow, P. (2020), "Trading Off Consumption and COVID-19 Deaths," *Federal Reserve Bank of Minneapolis Quarterly Review,* 42, 2-13.

[11] Heathcote, J., Perri, F. and Violante, G. (2010), "Unequal We Stand: An Empiricial Analysis of Economic Inequality in the United States, 1967-2006," *Review of Economic Dynamics*, 13, 15-51.

[12] Jones, C. and Klenow, P. (2016), "Beyond GDP? Welfare Across Countries and Time," *American Economic Review*, 106, 2426–2457.

[13] Lucas, R.E., Jr. (2003), "Macroeconomic Priorities," *American Economic Review*, 93, 1–14.

[14] Lucas, R.E., Jr. (1987), *Models of Business Cycles*. Basil Blackwell, New York.

[15] Menachemi, N., Yiannoutsos, C., Dixon, B., Duszynski, T., Fadel, W., Wools-Kaloustian, K., Unruh Needleman, N., Box, K., Caine, V., Norwood, C., Weaver, L., Halverson, P. (2020), "Population Poin Prevalence of SARS-CoV-2 Infection Based on a Statewide Random Sample - Indiana, April 25-29, 2020," *Morbidity and Mortality Weekly Report*, 69, 960-964.

[16] Murphy, K. and Topel, R. (2006), "The Value of Health and Longevity," *Journal of Political Economy*, 114, 871-904.

[17] Murtin, F., Boarini, R., Cordoba, J.C. and Ripoll, M. (2017), "Beyond GDP: Is there a law of one shadow price?" *European Economic Review*, 100, 390-411.

[18] Pindyck, R. (2020), "COVID-19 and the Welfare Effects of Reducing Contagion," *NBER Working Paper* No. 27121, May.

[19] Tallarini, T, (2000), "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45, 507-532.

[20] Viscusi, W.K. and Aldy, J. (2003), "The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World", *Journal of Risk and Uncertainty*, 27, 5-76.

[21] Weil, P. (1990), "Nonexpected Utility in Macroeconomics", *Quarterly Journal of Economics*, 105, 29–42.

TABLE 1
Welfare costs of COVID-19
Willingness to pay to avoid a one-year pandemic (% of consumption)

	VSL/c = 150			**VSL/c = 200**		
	Average	Median voter	Standard deviation	Average	Median voter	Standard deviation
No annuity markets						
σ = 1	41.2%	22.7%	41.2%	45.5%	29.0%	40.5%
σ = 2	34.3%	20.5%	35.2%	38.3%	25.6%	35.6%
Perfect annuity markets						
σ = 1	35.9%	21.1%	32.9%	41.6%	27.1%	35.2%
σ = 2	29.7%	19.2%	25.9%	34.3%	24.2%	28.0%

Notes: Table displays the willingness to pay (WTP) for different values of σ and different VSL-to-consumption ratios at age 40 (VSL/c). COVID-19 age-specific death rates are fitted from Ferguson *et al*. (2020) in combination with Menachemi *et al*. (2020). As in the benchmark calibration, discount values are age-dependent and $\theta = 1$. Average WTP corresponds to the weighted average of age-specific WTP using population weights. Median voter WTP is that of individuals age 46, which corresponds to the median of the population above age 18. Standard deviation is computed over the age-specific WTP.

TABLE 2

Robustness checks to alternative parameters

Willingness to pay to avoid a one-year pandemic (% of consumption)

Alternative parameters	Average age-specific	Median voter	Standard deviation
No annuity markets			
Benchmark calibration	41.2%	22.7%	41.2%
$\theta = 10$	42.9%	23.7%	41.4%
$\theta = 20$	43.9%	24.0%	41.5%
Constant μ and η	42.3%	23.9%	41.6%
Non-homothetic finite horizon	41.4%	22.8%	41.3%
Perfect annuity markets			
Benchmark calibration	35.9%	21.1%	32.9%
$\theta = 10$	40.2%	22.7%	38.1%
$\theta = 20$	42.5%	23.4%	40.2%
Constant μ and η	36.5%	21.7%	33.3%
Non-homothetic finite horizon	35.9%	21.1%	33.1%

Notes: Benchmark calibration refers to the case $\sigma = 1$, VSL/c = 150, $\theta = 1$ and age-varying discount factors. Each robustness check is introduced separately. Average WTP corresponds to the weighted average of age-specific WTP using population weights. Median voter WTP is that of individuals age 46, which corresponds to the median of the population above age 18. Standard deviation is computed over the age-specific WTP.

TABLE 3

Persistent consumption costs

Willingness to pay during year of pandemic (% of consumption)

	Average age-specific	Median voter	Standard deviation
No annuity markets			
Benchmark calibration ($\rho = 0$)	41.2%	22.7%	41.2%
$\rho = 0.5$	30.9%	12.1%	40.8%
$\rho = 0.9$	12.2%	2.9%	25.0%
Perfect annuity markets			
Benchmark calibration ($\rho = 0$)	35.9%	21.1%	32.9%
$\rho = 0.5$	23.4%	11.1%	25.3%
$\rho = 0.9$	8.2%	2.6%	13.7%

Notes: Same as in Table 2. In the benchmark calibration consumption costs are not persistent ($\rho = 0$) and are all paid in one period. For the cases with persistent consumption shocks ($\rho > 0$) the table displays the WTP on the first year (year of the COVID-19 pandemic).

TABLE 4

Aggregate willingness to pay according to social planner (% of consumption)
Alternative social welfare functions

	Aggregate WTP
No annuity markets	
$\Psi = 0$	4.1%
$\Psi = \gamma = 0.675$	10.8%
$\Psi = 0.81$	22.0%
$\Psi = 0.90$	41.0%
$\Psi = \sigma = \theta = 1$	72.7%
Perfect annuity markets	
$\Psi = 0$	40.1%
$\Psi = \gamma = 0.81$	58.2%
$\Psi = \sigma = \theta = 1$	63.1%

Notes: All exercises on the table use the benchmark calibration with $\sigma = 1$, VSL/c = 150, $\theta = 1$ and age-varying discount factors. $\Psi = 0$ represents the case in which the planner has linear preferences or is indifferent to inequality across ages. The rest of the scenarios introduce some curvature in the planner's utility ($\Psi > 0$).

TABLE 5

Full recessions: COVID-19 with concurrent economic contraction

(% of consumption)

Total number COVID deaths	Death rate	3% recession	5% recession	10% recession
Average				
300,000	0.09%	17.1%	18.8%	23.1%
400,000	0.12%	20.2%	21.9%	25.9%
Median voter				
300,000	0.09%	8.7%	10.6%	15.3%
400,000	0.12%	10.5%	12.4%	16.9%
Planner ($\psi = \gamma = 0.675$)				
300,000	0.09%	5.1%	7.1%	11.9%
400,000	0.12%	5.8%	7.7%	12.6%
Planner ($\psi = 0.9$)				
300,000	0.09%	10.6%	12.4%	17.1%
400,000	0.12%	13.0%	14.8%	19.3%

Notes: All exercises on the table use the benchmark calibration with no annuity markets, $\sigma = 1$, VSL/c = 150, $\theta = 1$ and age-varying discount factors. For the planner we consider two different scenarios of inequality aversion: $\psi = \gamma = 0.675$ and $\psi = 0.9$.

TABLE 6
The consumption - lives frontier
(% of consumption)

Total number of COVID deaths	Death rate	Average age-specific	Median voter	*Planner* ($\psi = \gamma = 0.675$)	*Planner* ($\psi = 0.9$)
0	0%	41.2%	22.7%	10.8%	41.0%
200,000	0.06%	39.5%	20.6%	9.7%	37.6%
400,000	0.12%	37.6%	18.4%	8.6%	34.1%
600,000	0.18%	35.5%	16.2%	7.5%	30.4%
800,000	0.24%	32.9%	13.9%	6.4%	26.4%
1,000,000	0.30%	30.0%	11.5%	5.3%	22.2%
1,200,000	0.36%	26.3%	9.1%	4.1%	17.8%
1,400,000	0.42%	21.6%	6.6%	2.9%	13.1%
1,600,000	0.48%	15.3%	4.0%	1.8%	8.1%
1,800,000	0.54%	6.3%	1.4%	0.6%	2.9%
1,900,000	0.58%	0.0%	0.0%	0.0%	0.0%

Notes: Same as Table 5. The largest number of deaths, 1.9 million people, correspond to the fatality rate estimated in Menachemi *et al*. (2020).

Figure 1. US demographic structure and survival probabilities - 2017

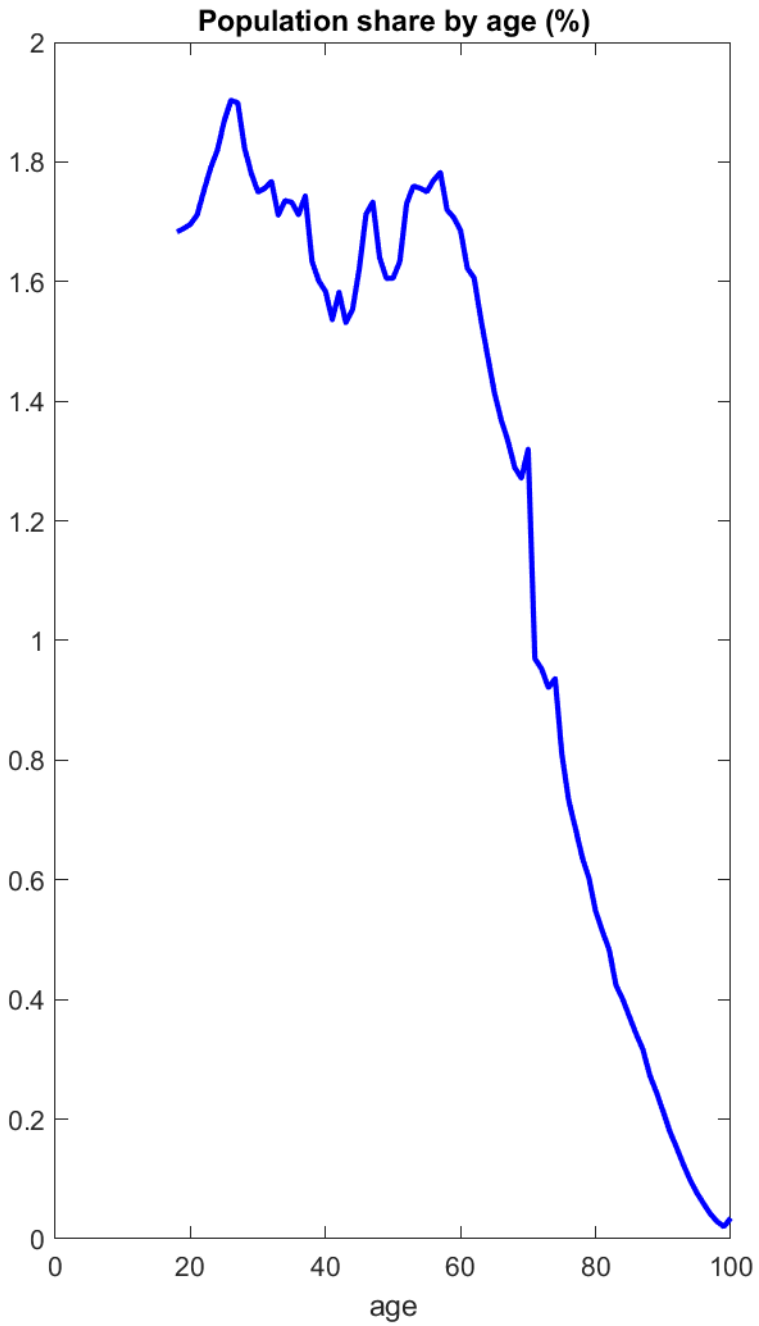

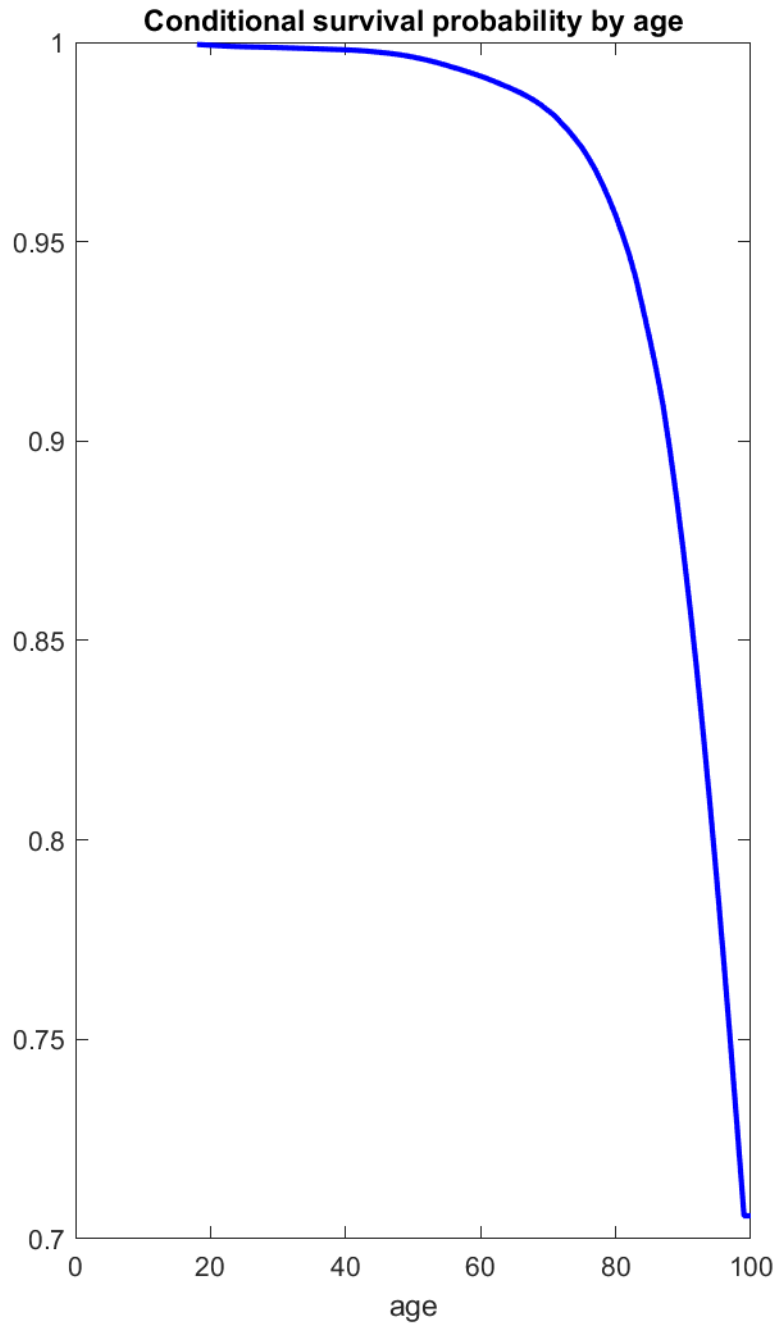

Figure 2. Age-profile of mean log consumption

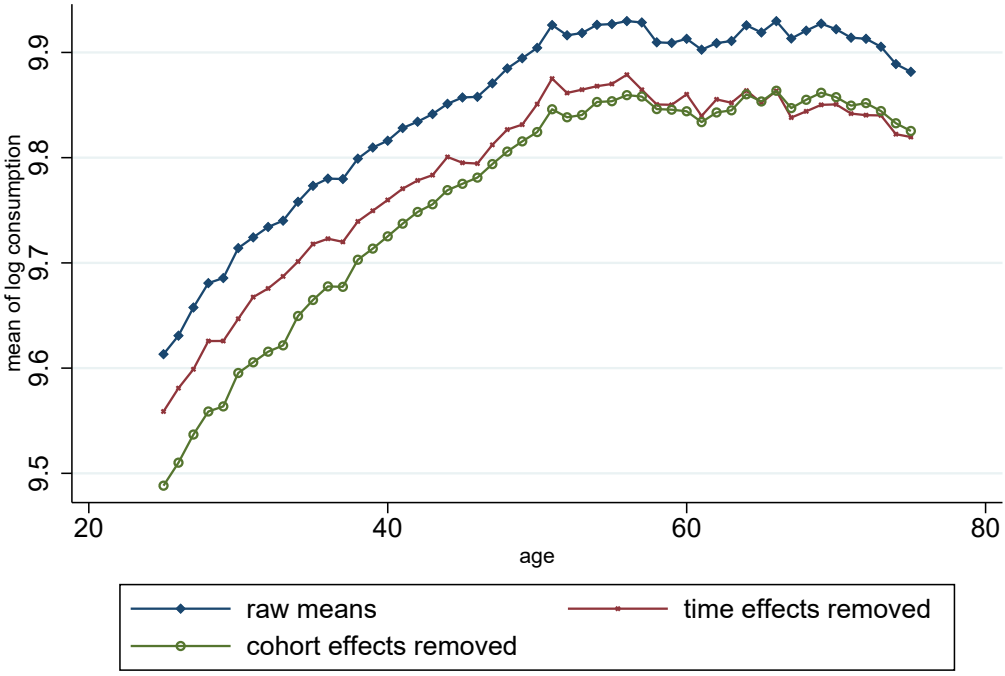

Figure 3. Age-profile of variance log consumption

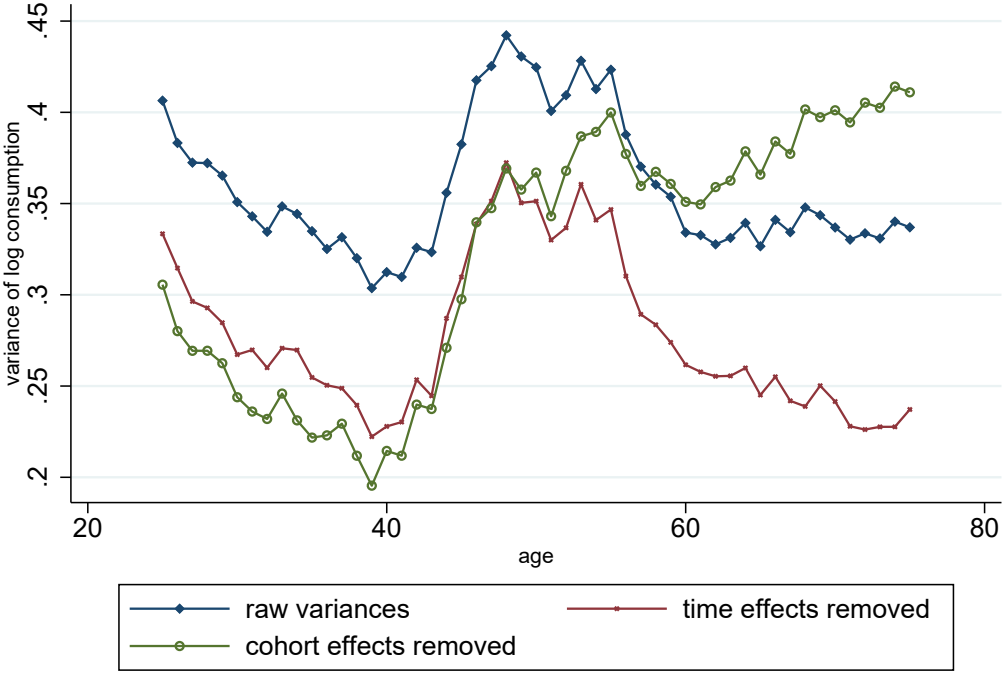

Figure 4. Individual consumption shock

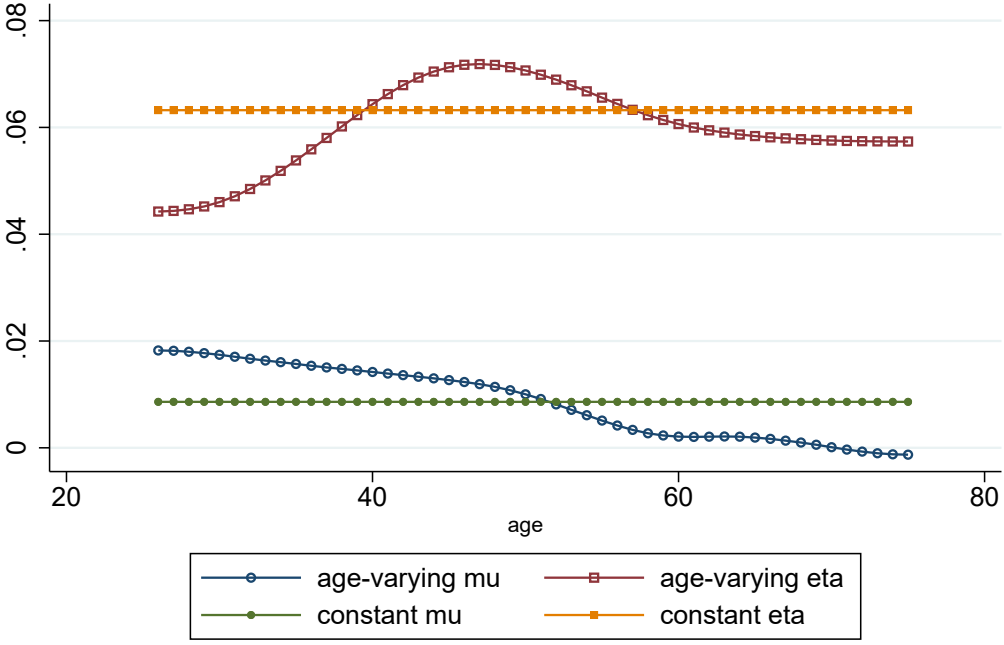

Figure 5. Parameters for benchmark calibration

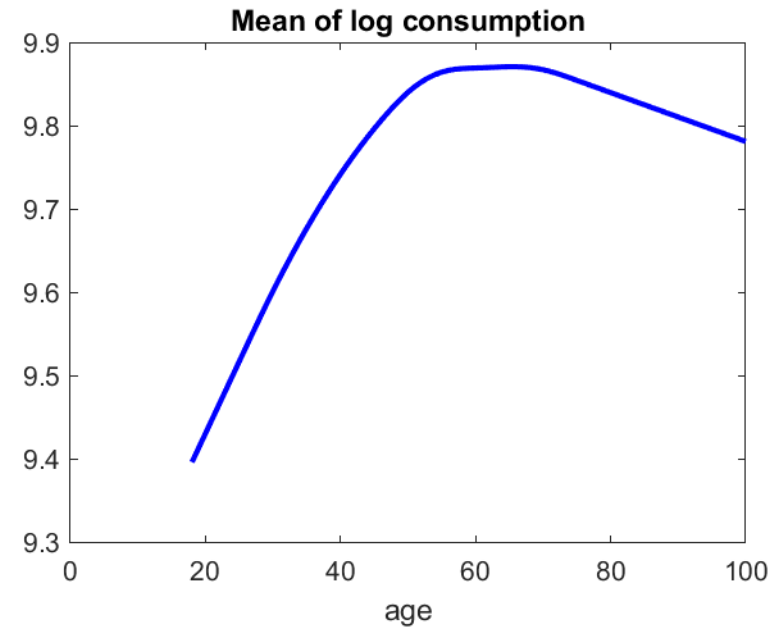

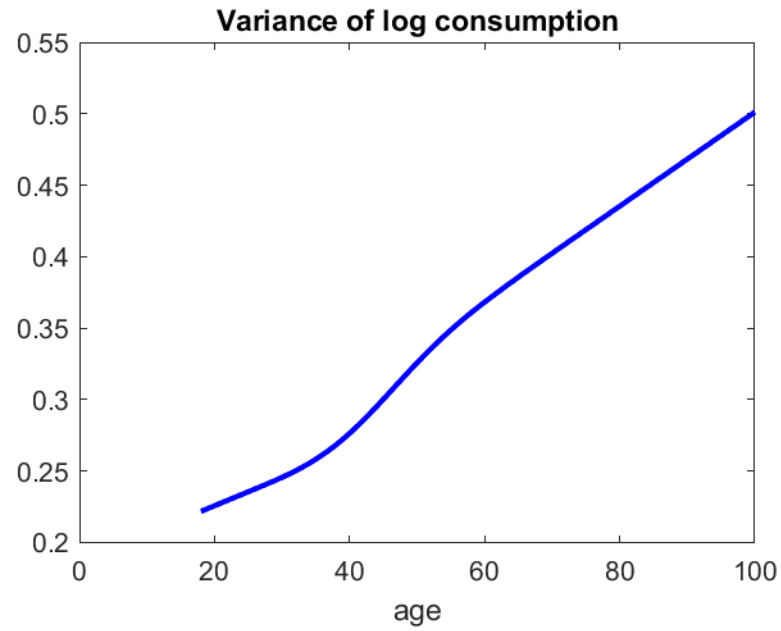

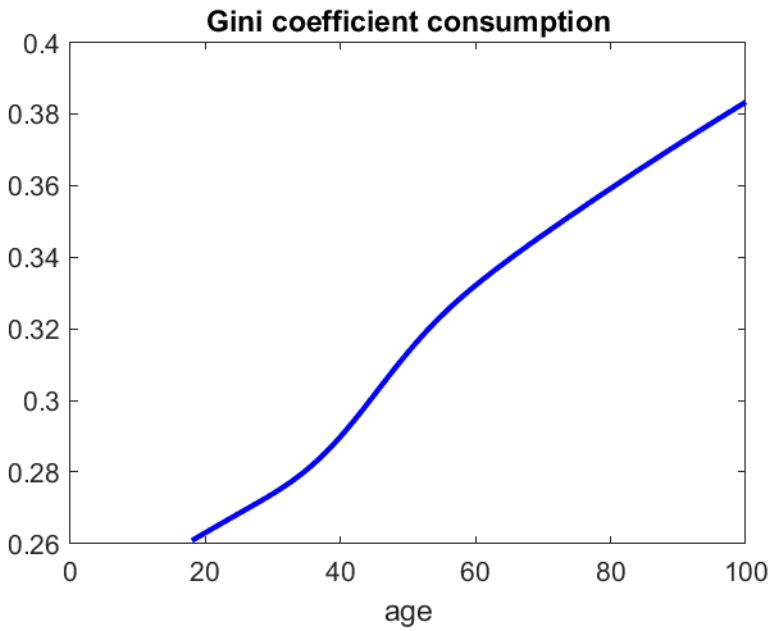

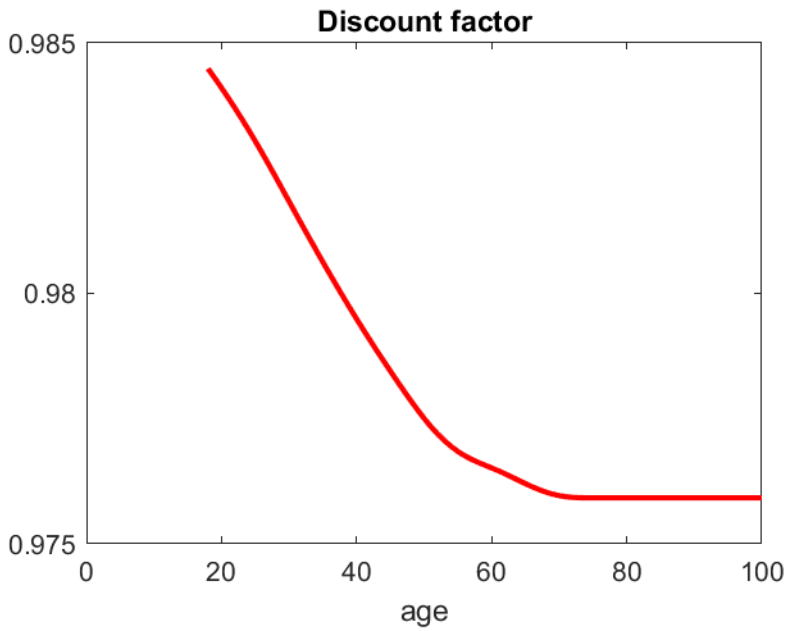

Figure 6. Distributional welfare effects of COVID-19 pandemic

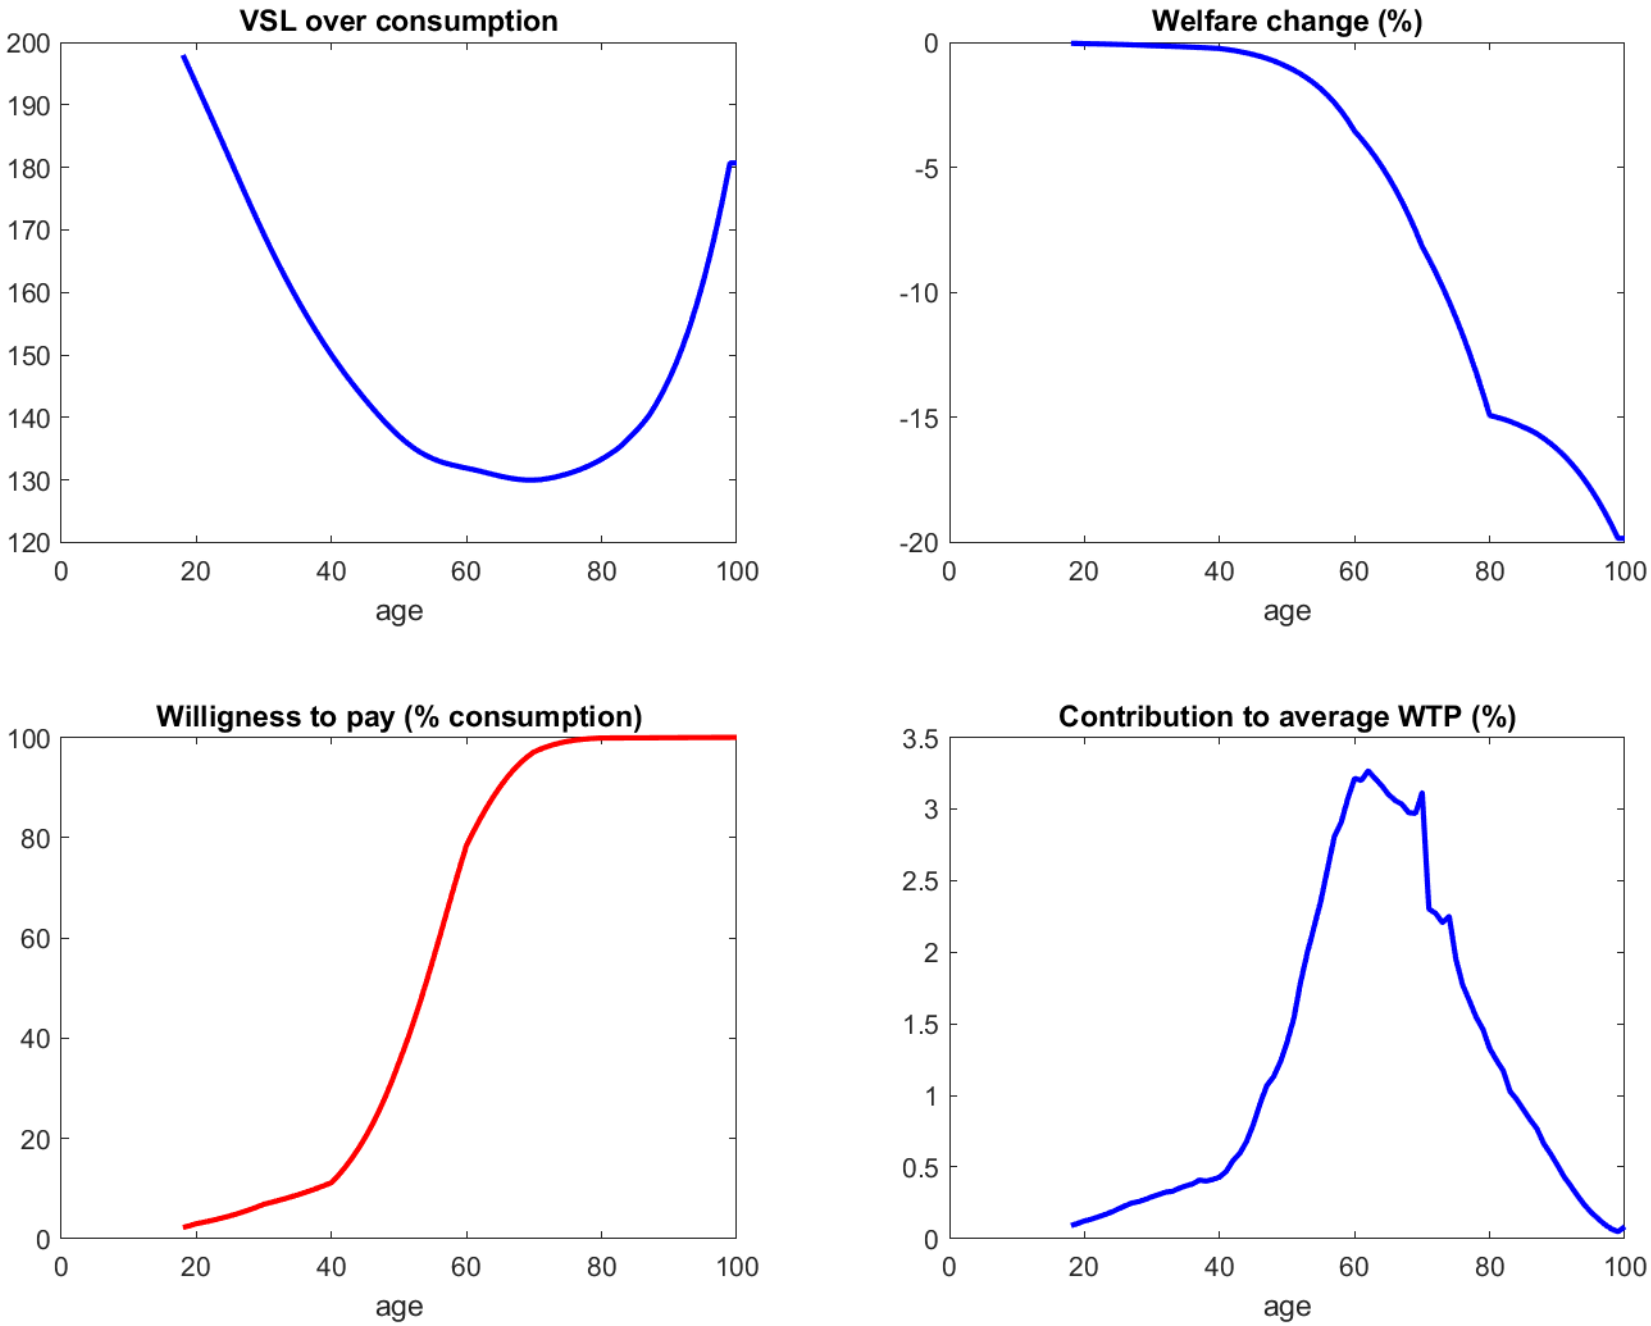

Figure 7. Robustness checks: Willingness to pay under alternative scenarios

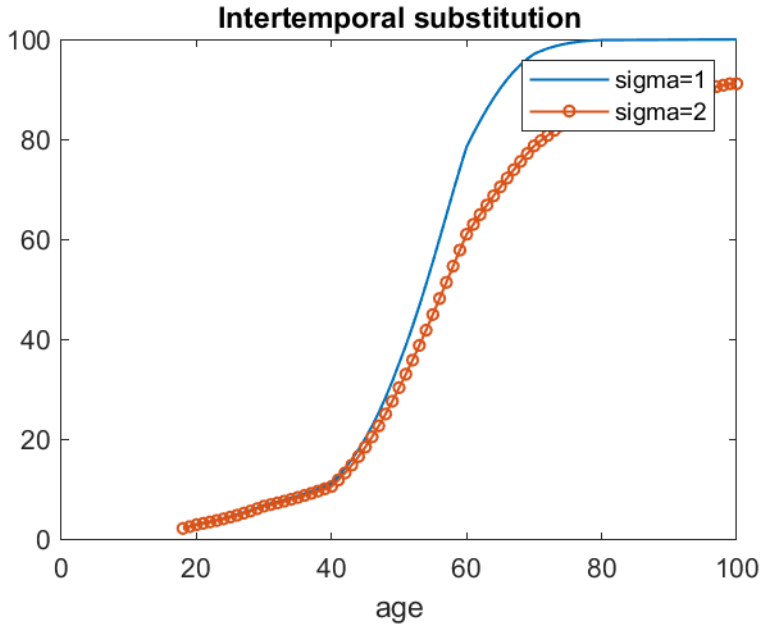

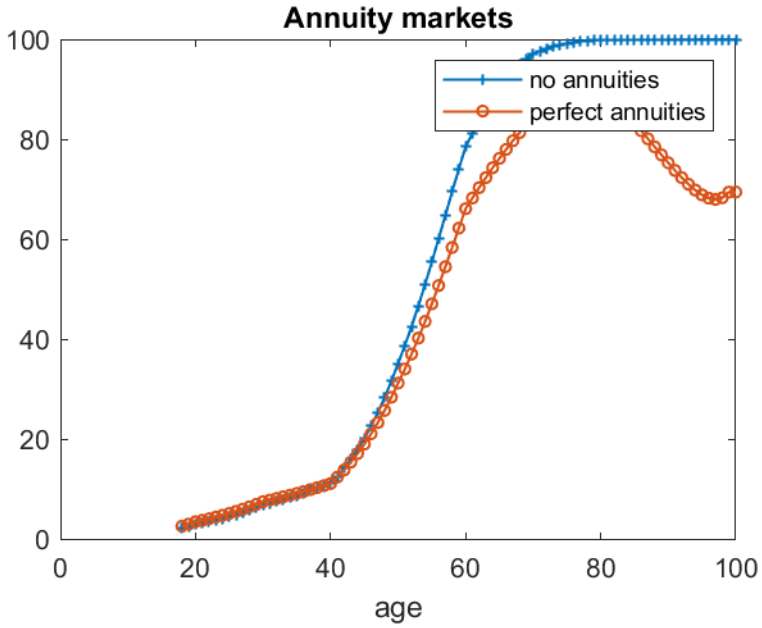

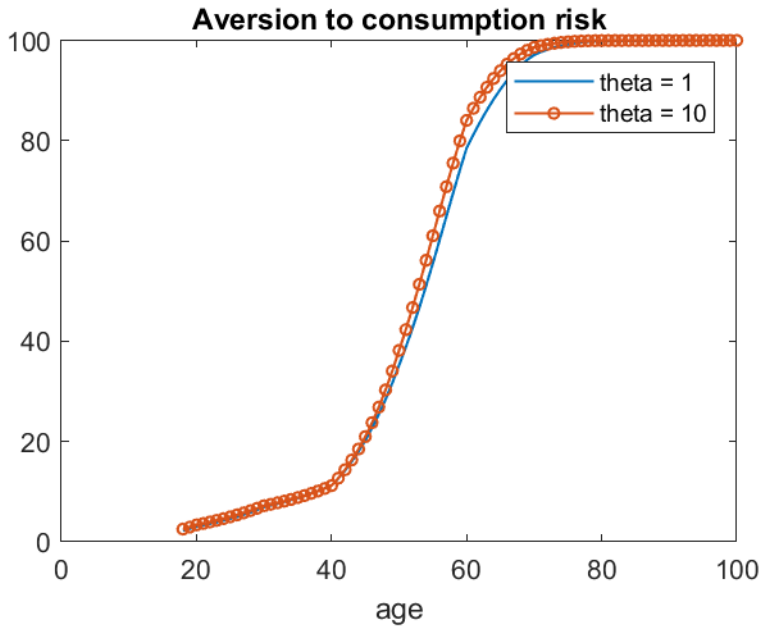

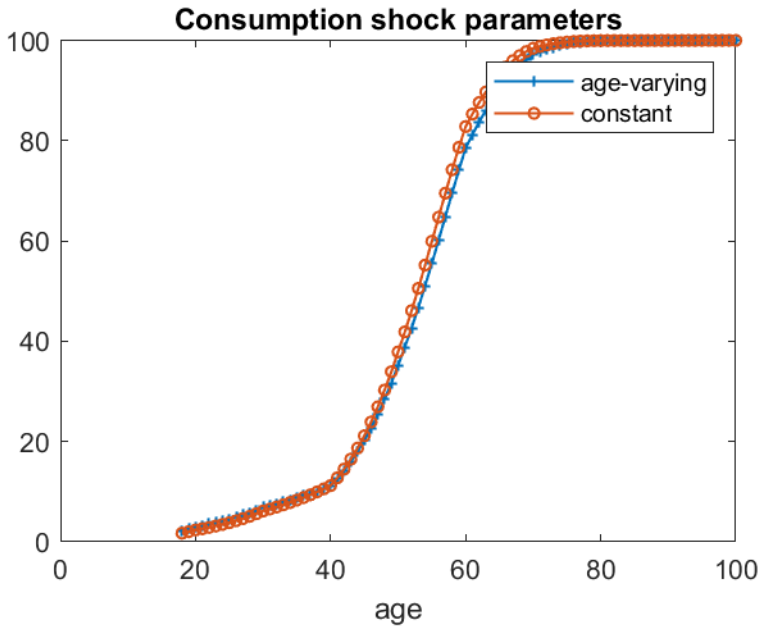

Figure 8. Willingness to pay under persistent consumption costs

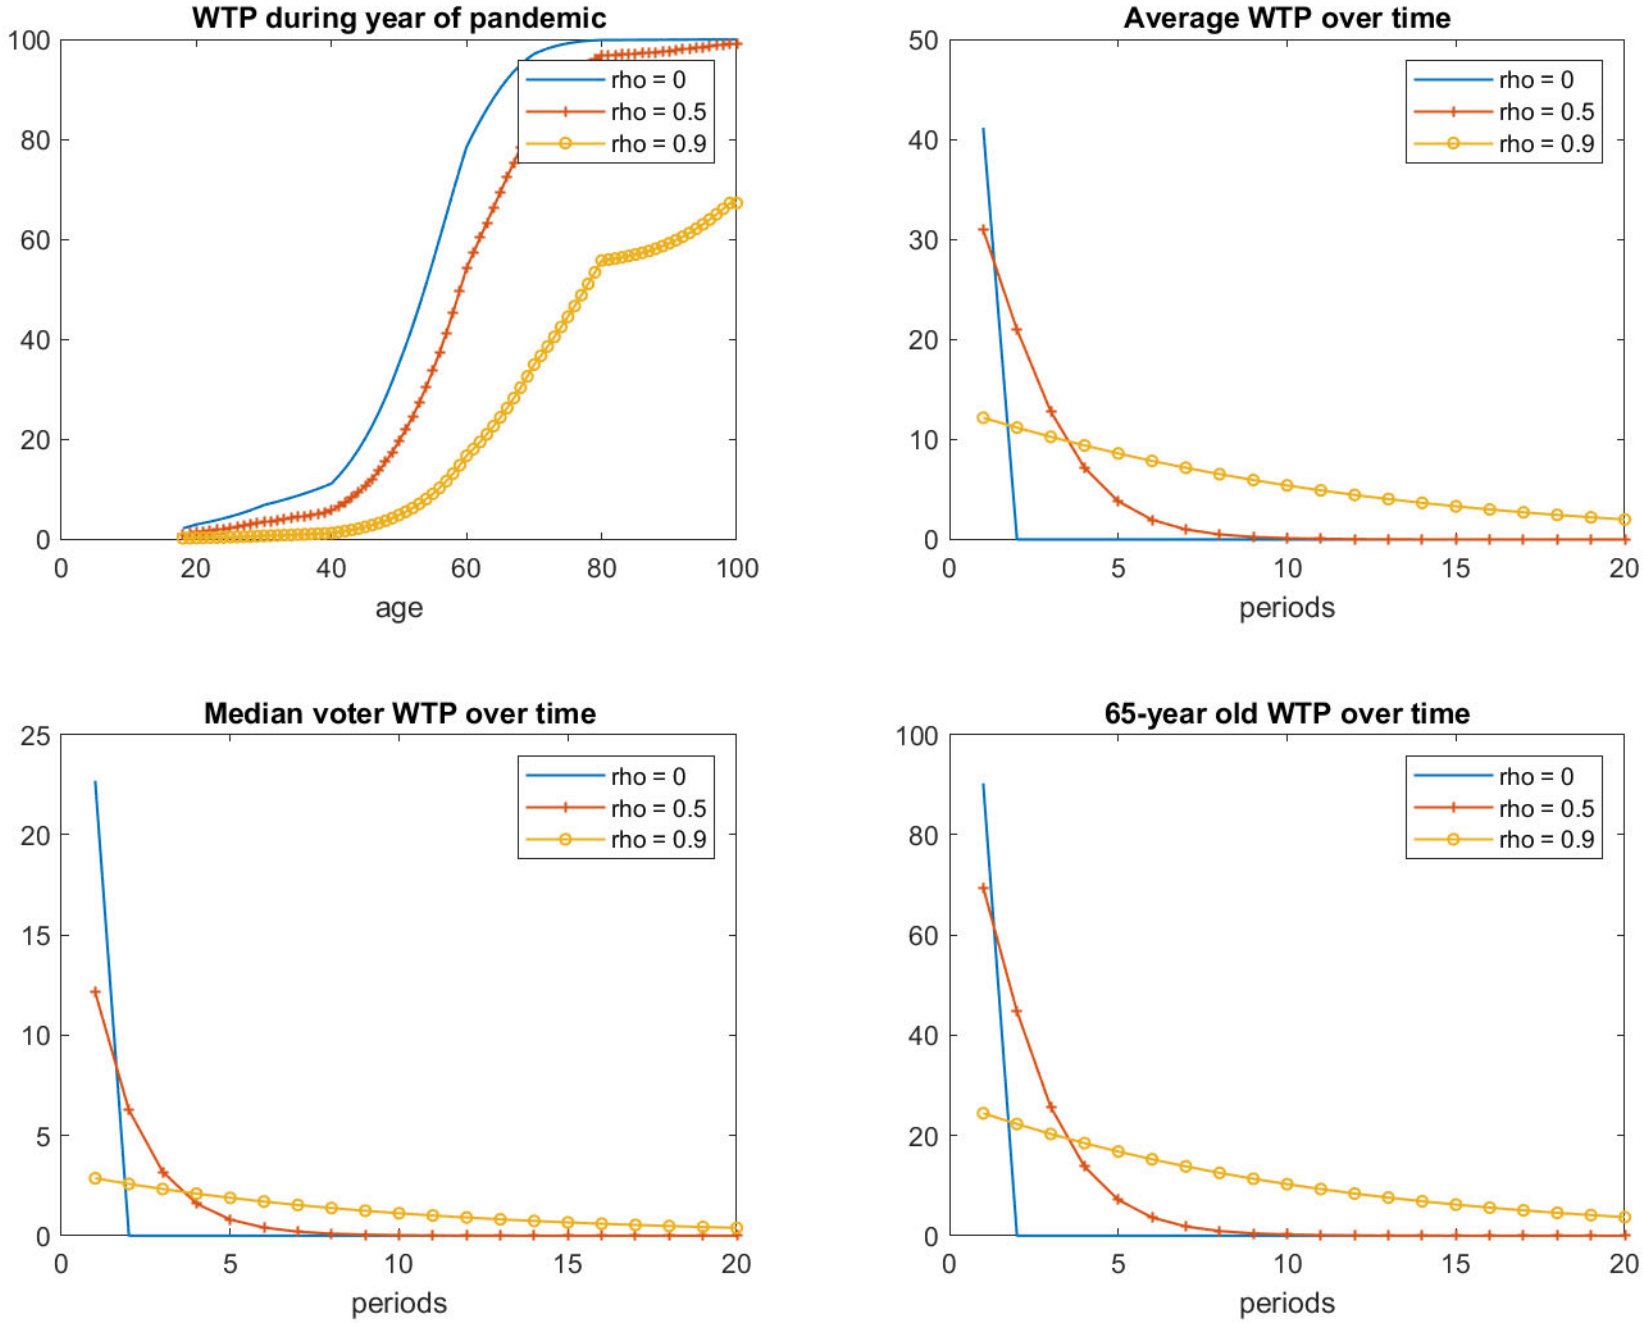

Figure 9. The consumption - lives frontier