Recursive Preferences and the Value of Life. Theory and Application to Epidemics

Antoine Bommier François Le Grand Daniel Harenberg*

June 19, 2021

Abstract

This article investigates how recursive preferences can be used in the context of lifecycle models featuring uncertain and endogenous lifespans. After pointing out difficulties arising with recursive models based on the frameworks of Epstein and Zin (1989) and Weil (1989), we explain the benefits of the risk-sensitive preferences of Hansen and Sargent (1995). These are well-behaved and can accommodate a positive value of life without constraining the risk aversion and intertemporal elasticity of substitution parameters. We apply the risk-sensitive framework to the consumption-mortality trade-off faced by a benevolent planner in a pandemic and show how risk aversion amplifies the role of the age distribution of deaths. Contrasting results for the Covid-19 pandemic with those obtained for the 1918 influenza highlights how risk aversion makes deaths at young age become a greater source of concern. As a consequence, the risk-sensitive model calls for weaker restrictions in the Covid-19 pandemic than the standard, additive model. The reverse finding is obtained when considering a pandemic that strongly affect young people, like the 1918 influenza.

Keywords: value of life, recursive utility, lifecycle models.

JEL codes: G11, J17.

*Bommier: ETH Zurich, abommier@ethz.ch; Le Grand: EMLyon Business School and ETH Zurich, legrand@em-lyon.com; Harenberg: ETH Zurich, Switzerland, and Oxford Economics Ltd., Germany, e-mail: dan.harenberg@gmail.com. The authors gratefully acknowledge the financial support of the Swiss-Re Foundation and the ETH Zurich Foundation and are thankful to Julien Hugonnier, Florian Pelgrin, and Pascal St-Amour for the feedback they provided on an earlier version of this paper.

1 Introduction

Major epidemics like the plague, the Spanish flu or the ongoing Covid-19 may cause the death of millions of people. A solution to reduce the number of fatalities involves limiting human interactions, which however may imply a huge economics cost. Societal decision-making therefore has to consider a trade-off between survival probabilities and wealth, which requires to set (implicitly or explicitly) an economic value for mortality risk reductions. The literature that addresses this trade-off typically uses the seminal theoretical framework introduced by Yaari (1965), where agents' lifetime utility is the sum of instantaneous utilities weighted by survival probabilities. As is well-known, such additive models lack flexibility to disentangle risk aversion from intertemporal elasticity of substitution (IES, henceforth). Using additive models may then prevent from properly accounting for risk aversion, which, however, is likely to be a major behavioral trait in the presence of mortality risk.

Following the works of Epstein and Zin (1989) and Weil (1989) (EZW, henceforth), recursive utility models have become the most popular tool to address the role of risk aversion in intertemporal contexts, with applications in numerous fields. While recursive preferences were initially developed to address long-standing puzzles in the macro-finance literature, they are now increasingly employed in other fields such as the economics of climate change, health economics, or household finance. In the current paper, we provide a theoretical investigation on how recursive models may help clarify the role of risk aversion when discussing the value of mortality risk reduction. We then highlight how accounting for risk aversion may provide new insights when considering wealth-survival trade-offs.

Our first contribution consists in pointing at major problems that emerge when applying the EZW framework to the context of an uncertain lifetime.[1] We explain that when applied to realistic mortality patterns and when the IES is smaller than

[1]See Hugonnier et al. (2013) for a continuous-time version and Córdoba and Ripoll (2017) for a discrete-time one. Córdoba et al. (2020) also use the EZW framework to price the mortality risk caused by the Covid-19 pandemic.

one, the recursion defining EZW preferences either features a negative value of life or admits a unique solution where utility is constant and equals zero everywhere. Non-constant solutions with a positive value of life only exist when mortality is constrained to remain small at all ages. Such a restriction involves assuming that agents are perpetually young, with a life expectancy that remains high at all ages. This makes such models unsuitable for studying value-of-life issues with realistic demographic data. In addition, these non-constant solutions imply that consumption at a given age and survival at that same age are substitutes (and not complements), leading to counterfactual predictions regarding lifecycle consumption profiles.

These serious shortcomings do not mean that all recursive preferences are inadequate for studying value-of-life questions. On the contrary, and this is the second contribution of our paper, we show that the risk-sensitive preferences of Hansen and Sargent (1995) allow the additive model to be consistently extended. In the risk-sensitive recursive framework, the value of life is mostly driven by a parameter that determines the utility gap between life and death. In contrast to homothetic EZW preferences, the IES and the risk aversion parameters (which can be varied independently from one another) can take any positive values without generating a discontinuity. Moreover, risk-sensitive preferences are monotone with respect to first-order stochastic dominance. This means that dominated strategies are never chosen, in contrary to what may occur when working with EZW preferences (see Bommier et al., 2017 or Bommier et al., 2020 for illustrations). Risk-sensitive preferences therefore provide an appealing theoretical framework for extending the literature on the relationship between the value of life and risk aversion, offering a way to complement the analyses of Eeckhoudt and Hammitt (2004), Kaplow (2005), Andersson and Treich (2011), and Bommier and Villeneuve (2012). We also explain that the risk-sensitive model can be easily generalized to account for ambiguity aversion. It then provides a pathway for extending to a multi-period setting the contributions of Treich (2010) and Bleichrodt et al. (2019) regarding the impact of ambiguity and ambiguity aversion.

In order to illustrate the relevance of using a recursive specification, we focus on two real-world cases, namely the Covid-19 and the 1918 influenza pandemics, and demonstrate how recursive preferences alter the conclusions regarding the optimal consumption-mortality trade-off. Our results highlight that the sign of corrections depends on whether the pandemic predominantly impacts older people (as with Covid-19) or younger people (as with the 1918 influenza outbreak). For Covid-19, accounting for risk aversion through recursive preferences would tend to reduce the amount of consumption that the social planner would be willing to sacrifice in order to limit mortality. Concretely, this would advocate for a quicker reopening of the economy, or a less severe lock-down, than what is advised when using usual additive preferences. For the 1918 influenza, the conclusions are the inverse. The reason for those findings is that models are usually calibrated upon observations of wage-risk trade-offs made by workers (thus of "middle-aged" people), while accounting for the mortality impact of pandemics requires to infer the value of mortality risk reduction at younger and older ages. A form of extrapolation is thus needed, which is typically achieved through the use of a specific model of individual preferences. The structure of the model adopted ends up playing a key role in this extrapolation. In particular, in comparison to additive models, recursive models end up giving more weight to particularly adverse consequences (i.e., death of younger people) and less weight to less dramatic consequences (i.e., death of older people). This of course reflects the very natural role of risk aversion which could not be properly investigated with the standard additive model.

The structure of our paper is as follows. In Section 2, we explain how recursive models can be used to discuss value-of-life issues. We start by addressing the shortcomings of the EZW homothetic specifications that account for the possibility of death and assume an IES below one. We then take a constructive approach and points to a well-behaved recursive framework that can be used to discuss value-of-life issues. Section 3 illustrates the application of risk-sensitive preferences to two real-world examples: the Covid-19 and 1918 influenza pandemics. Results are contrasted

with those of the standard additive model.[2] Section 4 concludes.

2 Studying the value of life with recursive models

2.1 Domain matters

Recursive models were initially developed in infinite horizon settings. To model choices under uncertain lifespans, we need to consider lives of unequal lengths, which is fundamentally different from an infinite or fixed horizon setting. To overcome this difficulty, instead of describing a life as a finite sequence of consumption periods, we can view it as an infinite sequence of per-period "realizations", where each realization is either "being dead" or "being alive, with some positive consumption". Mathematically speaking, the set of possible per-period realizations is $\mathbb{R}_+ \cup \{d\}$, where d is a symbol used to denote "being dead". Denoting by z_t the consumption at date t, a life is then represented by a sequence in the form:

$$(z_0, \ldots, z_T, d, d, \ldots) \in (\mathbb{R}_+ \cup \{d\})^\infty, \tag{1}$$

for some $T \geq 0$. Sequence (1) corresponds to a lifespan of $T+1$ periods, where $(z_0, \ldots, z_T) \in \mathbb{R}_+^{T+1}$ is the consumption profile when alive. Formally, if one excludes resurrection and immortality, the set of possible lives, denoted by L, is defined as:

$$L = \left\{(z_t)_{t\geq 0} \in (\mathbb{R}_+ \cup \{d\})^\infty : \ \exists T \in \mathbb{N}, \forall t > T, z_t = d \text{ and } \forall t \leq T, z_t \in \mathbb{R}_+\right\}. \tag{2}$$

As the set L is a subset of an infinite product space, recursive methods can be used in the standard way. We must, however, bear in mind that the underlying product space is not $\mathbb{R}_+^\infty$, as is usually the case, but $(\mathbb{R}_+ \cup \{d\})^\infty$. This means that period utility functions have to be defined on $\mathbb{R}_+ \cup \{d\}$, and not on $\mathbb{R}_+$. Moreover, in contrast to the usual case, the set $\mathbb{R}_+ \cup \{d\}$ is not always convex, which may raise

[2] Papers that analyze the Covid-19 pandemic while assuming the standard additive model include, among others, Hall et al. (2020), Hammitt (2020), Greenstone and Nigam (2020), and Robinson et al. (2020).

some technical issues.

2.2 Recursive models with mortality

Nearly all applied studies using recursive methods rely on the framework of Kreps and Porteus (1978) and assume a parametric form that can be expressed as:

$$U_t = f^{-1}(u(z_t) + \beta\phi^{-1}(E[\phi f(U_{t+1})]), \tag{3}$$

where $z_t \in \mathbb{R}_+$ is the consumption at time t, $\beta \in (0,1)$ is a time preference parameter, ϕ is an increasing function representing risk preferences, $E[\cdot]$ denotes the expectation operator, and $u : \mathbb{R}_+ \cup \{d\} \to \mathbb{R} \cup \{-\infty, \infty\}$ is the period utility function.[3] The function f is only a normalization device and can be any increasing function, with no impact on preferences. For example, one of the popular EZW specifications uses $f(x) = \frac{x^{1-\sigma}}{1-\sigma}$, but it is also possible to use $f(x) = \phi^{-1}(x)$, as in Kreps and Porteus (1978), or simply $f(x) = x$, which we will also use in the following.

We now discuss how to further parameterize the functions u and ϕ. It is important to note that, for the model to be well-defined, the function ϕ must have a domain that includes both $Im(u) = u(\mathbb{R}_+ \cup \{d\})$ and, because of the expectation operator, its convex hull, which may be strictly larger (in the sense of set inclusion) than $Im(u)$. Indeed, since $\mathbb{R}_+ \cup \{d\}$ is not convex, there is no reason for $Im(u)$ to be convex.

For our discussion, it is useful to consider conditional utilities, which are the utilities obtained conditionally on being dead or on being alive. For a dead agent, we simply get $U_t(d, d, \ldots) = f^{-1}\left(\frac{u(d)}{1-\beta}\right)$, which is the value of the period utility function, u, in the death state, d. Plugging this into (3), we find that the utilities of alive agents, denoted by V_t, are linked through the following recursion:

$$V_t = f^{-1}\left(u(z_t) + \beta\phi^{-1}\left(\pi_t E[\phi f(V_{t+1})] + (1-\pi_t)\phi\left(\frac{u(d)}{1-\beta}\right)\right)\right), \tag{4}$$

[3]Cases where the function u may take infinite values ($-\infty$ or ∞) are perfectly possible if $\phi(-\infty)$ or $\phi(\infty)$ are well-defined.

where π_t is the survival probability between dates t and $t+1$. Note that – with a slight abuse of notation – we still denote the expectation operator by $E[\cdot]$, although mortality risk is now treated separately.[4]

2.3 The period utility function u

The period utility function u is what determines, together with the time parameter β, preferences over deterministic prospects. A common specification is the case where preferences exhibit a constant IES. Formally, this means that there exist a positive scalar σ (the inverse of the IES) such that $\frac{-u''(z)}{zu'(z)} = \sigma$, for all $z \in (0, \infty)$. Leaving aside the case where $\sigma = 1$, one obtains by integration that $u(z) = K(\frac{z^{1-\sigma}}{1-\sigma} + u_l)$, where K and u_l are two constants.[5] The function u, defined over $\mathbb{R}_+ \cup \{d\}$, is then:

$$\begin{cases} u(z) = K\frac{z^{1-\sigma}}{1-\sigma} + Ku_l \text{ for } z \in \mathbb{R}_+, \\ u(d) = (1-\beta)u_d, \end{cases} \tag{5}$$

for some $K > 0$, $u_l \in \mathbb{R}$, and $u_d \in \mathbb{R} \cup \{-\infty, \infty\}$. Setting $K = 1$ corresponds to a normalization that is always possible. Another convenient choice is $K = 1 - \beta$, that implies with recursion (4) that the utility associated to a constant consumption flow $(z, z, \ldots)$ is $\frac{z^{1-\sigma}}{1-\sigma} + u_l$. We can further normalize the function u by setting either $u_l = 0$ or $u_d = 0$. But imposing an additional relation between u_l and u_d goes beyond a mere normalization. In particular, this would constrain the willingness-to-pay for mortality risk reduction. In other words, one can freely constrain u_l, or alternatively u_d, but constraining both simultaneously is not without loss of generality.

[4] Treating mortality separately is possible as long as mortality is independent of other risks.

[5] All the models we consider in the following can naturally be extended to account for the case where $\sigma = 1$. For the sake of simplicity, we will not consider the case $\sigma = 1$ in most of our discussion (Footnotes 6 and 18 will however explain how the case $\sigma = 1$ can be handled for EZW and risk-sensitive preferences).

2.4 The risk aversion function ϕ

The function ϕ does not play any role when considering deterministic prospects. Indeed, in absence of uncertainty, the recursion (3) simplifies into: $U_t = f^{-1}(u(z_t) + \beta f(U_{t+1}))$ where ϕ no longer appears. However, this function ϕ plays an important role as soon as there is some uncertainty. It then governs agents' risk aversion as well as their preference for the timing of resolution of uncertainty.

The function ϕ needs to be properly defined on the convex hull of $Im(u) = u(\mathbb{R}_+) \cup u_d$. In theory, this function does not need to be concave, but convexity issues in the agents' optimization problems could arise when considering nonconcave functions. Models in the economic and finance literature typically assume that the function ϕ is either affine (yielding an additive specification) or strictly concave. Increasing the concavity of ϕ amounts to increasing risk aversion (see Epstein and Zin, 1989, Section 4). The characterization of preference for the timing of resolution of uncertainty is slightly less straightforward as it is not directly related to the concavity of ϕ. Indeed, it follows from Kreps and Porteus (1978, Theorem 3) that preferences will feature preference for early (resp. late) resolution if the function $x \mapsto \phi(u(c) + \beta\phi^{-1}(x))$ is convex (resp. concave).

We discuss below three of the most common functional forms that can be found in the literature for ϕ: (i) affine, (ii) isoelastic, and (iii) exponential. We further assume specification (5) for the period utility function u. A summary of the properties of the preferences implied by these three functional forms can be found in Table 1 on page 23.

2.5 Affine ϕ (additive model)

Defining ϕ as an affine function is probably the most common choice in the literature on the value of life, and corresponds to the standard additive expected utility model of Yaari (1965). It is most often associated with the normalization $u_d = 0$, but renormalizing u by adding the same constant to both u_l and u_d, while keeping ϕ

unchanged, would have no impact on preferences. With $u_d = 0$ and $f(x) = x$, the recursion (4) defining the utility function of alive agents can be expressed as follows:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l + \beta \pi_t E[V_{t+1}]. \tag{6}$$

While the constant u_l can be ignored when considering choices under an exogenous mortality pattern, it plays a key role when discussing the welfare impacts of mortality risk. So, additionally requiring in recursion (6) that $u_l = 0$ might seem appealing for tractability reasons, but the value of mortality risk reduction would then be mostly pinned down by β and σ. In particular, if we further impose $\sigma > 1$ (along with $u_l = u_d = 0$), the value of mortality risk reduction will always be negative, for all consumption levels, which is a highly undesirable property when studying value of life matters. Independently of the value of u_l, a well-known limitation of the affine specification is that risk aversion and IES are intertwined.

It can be checked that, if consumption takes values in a compact interval $[z_{min}, z_{max}] \subset (0, \infty)$, the recursion (6) has for solution:

$$V_t^{add} = \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(\Pi_{j=t}^{\tau-1} \pi_j \right) \left(\frac{z_\tau^{1-\sigma}}{1-\sigma} + u_l \right). \tag{7}$$

The factor $\Pi_{j=t}^{\tau-1} \pi_j$ (with the usual convention that $\Pi_{j=t}^{t-1} \pi_j = 1$) is nothing other than the probability of being alive at age τ, conditional on being alive at age t. The utility V_t^{add} is thus a sum of instantaneous utilities weighted by the probability of being alive in the corresponding period and a discount factor.

2.6 Isoelastic ϕ (EZW specifications)

We focus here on the case of an isoelastic function ϕ, which corresponds to EZW preferences.[6] One problem with isoelastic functions relates to their definition sets.

[6] In this section, we keep assuming that $\sigma \neq 1$. The case $\sigma = 1$ could be obtained by setting $K = 1 - \beta$ in (5) and taking the limit $\sigma \to 1$ in the mathematical expressions of this section. We will not discuss this case any further here, as such a limit actually yields the risk-sensitive preferences of Section 2.7 (the fact that EZW preferences with $\sigma = 1$ actually corresponds to risk-sensitive preferences has already been noticed by Tallarini, 2000, among others).

They are never defined on the whole set $\mathbb{R}$ but either only on $\mathbb{R}_+$ (if ϕ is of the type x^α/α) or only on $\mathbb{R}_-$ (if ϕ is of the type $-(-x)^\alpha/\alpha$). Since ϕ needs to be defined on the convex hull of $u(\mathbb{R}_+) \cup u_d$, the model is well-defined if and only if the period utilities when alive, $\frac{z_t^{1-\sigma}}{1-\sigma} + u_l$, and when dead, u_d, always have the same sign. This implies that u_l and u_d must have the same sign as $1-\sigma$ (or be equal to zero).

All the studies using isoelastic functions ϕ that we are aware of, set $u_l = 0$. It is then convenient to take $\phi(x) = \frac{1}{1-\gamma}((1-\sigma)x)^{\frac{1-\gamma}{1-\sigma}}$ (where $\gamma \neq 1$ is a parameter driving risk aversion) and $f(z) = u(z) = \frac{z^{1-\sigma}}{1-\sigma}$ – which corresponds to the most popular EZW specification. With these elements, the recursion (4) defining the utility function V_t of alive agents can be expressed as follows:

$$V_t = \left(z_t^{1-\sigma} + \beta \left(\pi_t E[V_{t+1}^{1-\gamma}] + (1-\pi_t) u_d^{1-\gamma} \right)^{\frac{1-\sigma}{1-\gamma}} \right)^{\frac{1}{1-\sigma}}. \tag{8}$$

Although popular in applications, specification (8) raises a number of concerns. First, preferences represented by this specification are in general not monotone with respect to first-order stochastic dominance. This means that an agent endowed with such preferences can end up optimally choosing an allocation, while other allocations that are ex post always preferred to the chosen one are available. This is akin to opting for a dominated strategy in a game-theoretic context. Bommier et al. (2020) provide a two-period example illustrating that agents endowed with non-monotone preferences may end up saving *more* in presence of a mortality risk than what they would do if living two periods for sure. Second, one of the appealing features of EZW preferences is homotheticity that offers some tractability in applications. In specification (8), homotheticity requires to further impose $u_d^{1-\gamma} = 0$, such that recursion (8) becomes:

$$V_t = \left(z_t^{1-\sigma} + \beta \left(\pi_t E[V_{t+1}^{1-\gamma}] \right)^{\frac{1-\sigma}{1-\gamma}} \right)^{\frac{1}{1-\sigma}}. \tag{9}$$

Despite its appealing tractability features, specification (9) represents a severe restriction on preferences. More precisely, it does not enable to simultaneously have, for realistic mortality patterns, a well-defined utility, a positive willingness-to-pay for

mortality risk reduction, and an IES below 1. We investigate this statement further by distinguishing the cases $\gamma > 1$ and $\gamma < 1$.[7]

2.6.1 The homothetic case with $\gamma > 1$

When $\gamma > 1$, setting $u_d^{1-\gamma} = 0$ implies having $u_d = \infty$, which means that the utility of being dead is greater than the one of being alive and consuming any finite amount (since $u(z)$ is finite whenever z is finite). In other words, this means that the willingness-to-pay for mortality risk reduction is negative. This can be seen by computing the derivative of recursion (9) with respect to survival probability, which yields:

$$\frac{\partial V_t}{\partial \pi_t} = \frac{1}{1-\gamma} \pi_t^{\frac{\gamma-\sigma}{1-\gamma}} E[V_{t+1}^{1-\gamma}]^{\frac{1-\sigma}{1-\gamma}} V_t^{\sigma} < 0.$$

Though well-defined, the specification (9) with $\gamma > 1$ is not well-adapted to study questions related to mortality. For instance, this would imply the counterfactual statement that agents have a positive willingness-to-pay for goods that are likely to shorten their lives (they would also have preferences for hazardous cars, dangerous jobs, etc.).

2.6.2 The homothetic case with $\gamma < 1$

The case $\gamma < 1$ does not yield a negative value of mortality risk reduction, since in that case $u_d^{1-\gamma} = 0$ corresponds to $u_d = 0$. However, as we explain in the remainder of this section, this framework raises other concerns, being unable to cope with the case $\sigma > 1$, that is, with the empirically relevant case when IES is smaller than one (see Havránek, 2015). We explain in greater detail below why the case where $\gamma < 1$ and $\sigma > 1$ is problematic.[8]

[7] The case $\gamma = 1$ corresponds to $V_t = \left[z_t^{1-\sigma} + \beta \exp((1-\sigma)\pi_t E[\log(V_{t+1})])\right]^{\frac{1}{1-\sigma}}$ and raises the same concerns as the case $\gamma < 1$. For the sake of simplicity, we restrict our attention to $\gamma \neq 1$.

[8] Consistently with the rest of the paper, the analysis of this section is developed in a discrete-time framework, similar to the one in Córdoba and Ripoll (2017) and Córdoba et al. (2020). Although the mathematics are slightly more involved, the arguments developed here also apply to the continuous-time modeling of Hugonnier et al. (2013, 2020). See Bommier et al. (2021) for a detailed discussion of the continuous-time case.

Solutions to the homothetic EZW recursive model (9). First, consider the case where we assume that there exists a finite maximal lifespan $T_{\max}$, such that the probability that an agent living more than $T_{\max}$ years is null. If we consider that the population of concern is that of people currently alive and that we realize that the maximal lifespan $T_{\max}$ can be set arbitrarily large, this assumption seems pretty reasonable.[9] Since age $T_{\max}$ cannot be reached with positive probability, there must be $T < T_{\max}$ such that $\pi_T = 0$ and $V_{T+1} = 0$, where the last equality comes from the assumption of zero utility for death, $u_d = 0$. This implies that, for all $z_T \geq 0$, we have $V_T = (z_T^{1-\sigma} + \beta 0^{\frac{1-\sigma}{1-\gamma}} 0^{1-\sigma})^{\frac{1}{1-\sigma}} = 0$, when $\sigma > 1 > \gamma$. Similarly, by backward induction, we then obtain that $V_t = 0$ for all t, irrespective of consumption levels and mortality rates. Assuming a finite upper bound on lifespan therefore implies that the only solution to recursion (9) is the constant zero utility.

Furthermore, this result still holds if we relax the assumption of a maximal lifespan and approximate the distribution of observed lifespans with an "unlimited" survival pattern, where survival rates become low at greater ages. This statement is formalized in the following formal result.

Result 1 *Consider the utility function defined by the recursion (9) with $\gamma < 1 < \sigma$.*

1. *If there is a maximal lifespan (i.e., there exists T such that $\pi_T = 0$), the only solution to (9) is $V_t = 0$ for all t.*

2. *When death is never certain and survival probability decreases with age, there are two cases:*

- $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$*: the recursion admits a unique solution $V_t = 0$ for all t;*

- $\lim_{t\to\infty} \pi_t \geq \beta^{\frac{1-\gamma}{\sigma-1}}$*: the recursion admits multiple solutions: one being $V_t = 0$*

[9] This assumption is also consistent with demographic evidence. Jeanne Calment is reported to have the longest lifespan of 122 years and 164 days and is the only human to have lived beyond the age of 120 years. Maximal biological age is also supported by biological evidence (Weon and Je, 2009; Dong et al., 2016).

for all t, and the other solution being given by:

$$V_t = \left[\beta^{-t} \left(\Pi_{j=0}^{t-1} \pi_j \right)^{\frac{\sigma-1}{1-\gamma}} M + \sum_{s=t}^{\infty} \beta^{s-t} \left(\Pi_{j=t}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \qquad (10)$$

for some constant $M \geq 0$.

Result 1, which is proved in Appendix C, states that whenever survival rates become low at greater ages, the only solution to the recursive equation (9) is zero utility, that is $V_t = 0$ for all t. The recursive equation (9) admits a non-zero solution when the model is restricted to agents whose mortality rates are not greater than $1 - \beta^{\frac{1-\gamma}{\sigma-1}}$. For this solution to exist, agents should have a life expectancy that is never below $(1 - \beta^{\frac{1-\gamma}{\sigma-1}})^{-1}$, no matter their age.[10] Consequently, the utility function (10), being the only non-degenerate solution to (9), could be used in a model of perpetual youth, but not in a lifecycle model that accounts for actual mortality profiles.

From a technical point of view, one may notice that when $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$ the solution given in (10) corresponds to $V_t = 0$ because simple convergence considerations yield $\sum_{s=t}^{\infty} \beta^{s-t} \left(\Pi_{j=t}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} = \infty$. The lack of of non-degenerate solutions in the case where $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$ can therefore be seen as a convergence issue. Before discussing how such issues could be addressed (with no fully-satisfying outcome however), we first explain what occurs when these convergence issues are ignored.

Implications of ignoring convergence issues. It may be tempting to compute the first-order conditions implied by the recursion (9), without addressing convergence issues. However, this leads to counterfactual predictions that will now be explored. We start by computing the marginal rate of substitution (MRS) between consumption in period $t+1$ and consumption in period t. Ignoring that $V_t = 0$ due to convergence issues, one would derive from equation (9) that:

$$\frac{\partial V_t}{\partial z_{t+1}} \Big/ \frac{\partial V_t}{\partial z_t} = \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} z_{t+1}^{-\sigma} z_t^{\sigma}. \qquad (11)$$

[10] By way of illustration, taking $\beta = 0.97$, $\sigma = 2.0$, and $\gamma = 0.5$ implies that agents' life expectancy should remain above 65 years, independently of their age.

When $\gamma < 1 < \sigma$, this MRS is decreasing with the survival probability π_t. Formally:

$$\frac{\partial}{\partial \pi_t}\left(\frac{\partial V_t}{\partial z_{t+1}}\Big/\frac{\partial V_t}{\partial z_t}\right) < 0. \tag{12}$$

Equation (12) states that the less likely survival is at a given age, the more the agent wants to save resources for consumption at that age. The approach therefore assumes that consumption at a given age and survival at that same age are substitutes – while complementarity could be expected given the absence of bequest motive and thus of "utility of consumption" after death. In the limit case where $\pi_t \to 0$ (death is almost sure at the end of period t), we find that $\frac{\partial V_t}{\partial z_{t+1}}\Big/\frac{\partial V_t}{\partial z_t} \to \infty$. Thus, a marginal increase in consumption would be infinitely more valuable in period $t+1$ than in period t, in spite of the fact that survival in period $t+1$ is extremely unlikely. This seems counter-intuitive, and is at odds with what the additive expected utility model suggests.[11]

The MRS (11) finally yields counterfactual predictions regarding consumption-saving behavior. Consider a standard lifecycle consumption-saving problem. Agents, endowed with income y_t in period t, have to decide how much to consume and to save. Savings, denoted by w_t in period t, are assumed to pay a constant riskless return $1+r$. The agent's program can be written as:[12]

$$V_t(w_t, \pi_t) = \max_{z_t, w_{t+1}} \left(z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}(w_{t+1}, \pi_{t+1})^{1-\sigma}\right)^{\frac{1}{1-\sigma}}, \tag{13}$$

$$\text{s.t. } y_t + w_t = z_t + \frac{1}{1+r} w_{t+1}. \tag{14}$$

Ignoring convergence issues, the first-order conditions of (13)–(14) yield:

$$\frac{z_{t+1}}{z_t} = \left(\beta(1+r)\pi_t^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{\sigma}}, \tag{15}$$

which means that if $\sigma > 1 > \gamma$, mortality reduces agents' impatience instead of

[11] With the additive model (6), $\frac{\partial V_t^{add}}{\partial z_{t+1}} / \frac{\partial V_t^{add}}{\partial z_t} = \beta \pi_t z_{t+1}^{-\sigma} z_t^{\sigma}$. The MRS thus increases with survival probability and tends to zero when the survival probability π_t becomes infinitesimally small.

[12] This is the same problem as studied in Section 3 of Córdoba and Ripoll (2017), with health set to a constant.

contributing to it, as in the additive model for instance.[13] This is a consequence of the substitutability between consumption and survival exhibited in equation (12).[14]

We now examine the quantitative implications of equation (15) for consumption using benchmark parameters and realistic mortality profiles. The interest rate is assumed to be $r = 4\%$. Preference parameters are the same as in Footnote 10, $\beta = 0.97$, $\sigma = 2.0$, and $\gamma = 0.5$. Mortality rates are those of the total US population in 2018, as reported in the Human Mortality Database (HMD). In Figure 1, we

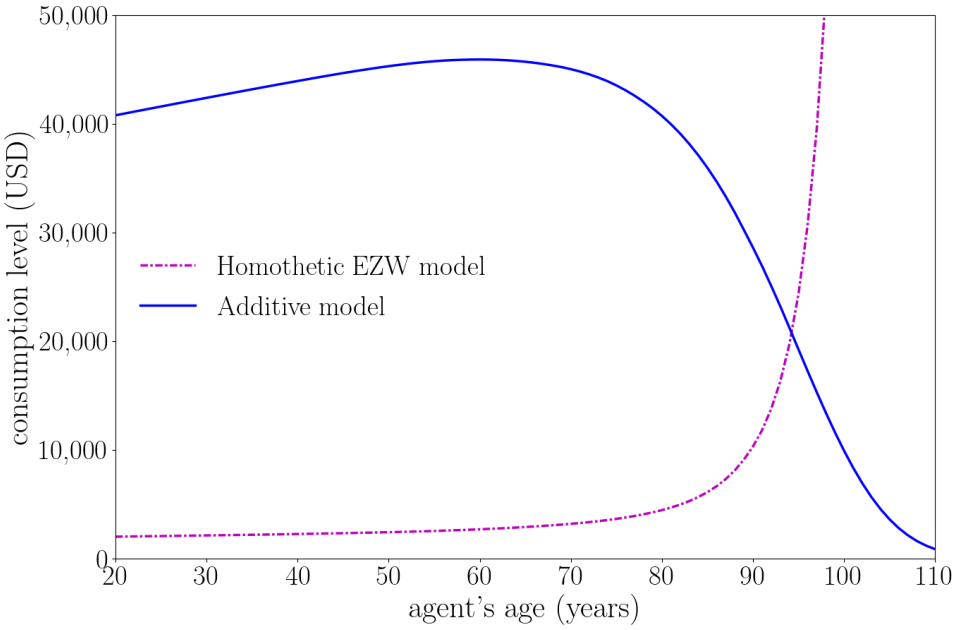

Figure 1: Consumption profiles implied by the additive and the homothetic EZW models.

plot the consumption path implied by the first-order equation (15). For the sake of comparison, we also plot the consumption path implied by the standard additive expected utility model for the same parameter values for β and σ. Lifetime wealth

[13] In the absence of annuities, survival probabilities have no impact on the budget constraint. The impact of survival probabilities on the optimal consumption profile is then a pure impatience effect.

[14] In a complementary analysis, Zhang et al. (2018) also remark that with the solution (15), the parameter σ, which is the inverse of the IES, may end up playing an unexpected role, qualitatively inconsistent with what would be obtained with an additive model. For example, when $\beta(1+r) = 1$, (15) implies a rate of consumption growth equal to $\frac{1-\sigma}{\sigma}\log(\pi_t)$. Since $\pi_t \in (0,1)$, this growth rate is always positive when $\sigma > 1$ and increases with σ. Thus, instead of moderating the variations of consumption over time, a low IES would in fact amplify such variations.

is normalized to $1,000,000, such that the yearly consumption at age 20 for the additive agent is close to $40,000. This normalization is of little importance, since preferences are homothetic. The homothetic EZW profile exhibits a consumption level that remains extremely low until age 100, but that then sky-rockets after that. This is very different from the level obtained with the additive specification. Note that the y-axis is truncated at $50,000 for greater legibility but under the homothetic EZW model, consumption in fact reaches $98,000 at age 100 and $14,000,000 at age 110.

We finally turn to the willingness-to-pay for mortality risk reduction (MRR), which we will define as the marginal rate of substitution between survival and consumption:[15]

$$MRR_t = \frac{\partial V_t}{\partial \pi_{t1}} \Big/ \frac{\partial V_t}{\partial z_t} \,. \tag{16}$$

This represents the amount of consumption an agent would be willing to pay for increasing their survival probability. From equation (9), we obtain:

$$MRR_t = \frac{1}{1-\gamma} z_t^{\sigma} \beta \pi_t^{\frac{\gamma-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}. \tag{17}$$

With $\gamma < 1 < \sigma$, the willingness-to-pay for mortality risk reduction is decreasing in the continuation utility V_{t+1} ($\partial MRR_t / \partial V_{t+1} < 0$). This means that as the possible future loss (measured by continuation utility V_{t+1}) increases, the agent will be less willing to avoid the loss. At the extreme, the agent has an infinite willingness-to-pay to marginally increase survival probability when they know that they will consume nothing and be miserable in the event of survival ($V_{t+1} = 0$). Conversely, their willingness-to-pay is zero when the agent knows that they will consume huge amounts

[15] The Environmental Protection Agency advises using this terminology, rather than the "value of a statistical life" (VSL). Note, moreover, that the literature generally defines VSL as the marginal rate of substitution between survival and wealth (and not consumption). In the absence of an annuity, this makes no difference and VSL, as usually defined, also equals the MRS shown in (16). When annuities are available, we have to account for the fact that lowering mortality rates may reduce the return of annuities. An adjustment must therefore be made. For example, with a perfect annuity market, VSL is the difference between the MRS (16) and the amount invested in annuities. The adjustment is independent of the agent's preferences, and would therefore have no impact on the subsequent discussion.

and have a superb life in the event of survival ($V_{t+1} = \infty$). Here again, the results are at odds with those obtained when using an additive expected utility specification, and with economic intuition.

Working with the homothetic EZW specification (9) therefore raises a number of concerns along several dimensions. It is worth noting that adding health (as in Córdoba and Ripoll, 2017) or age-dependent discount factors β_t (as in Córdoba et al., 2020) does not solve these issues, as we show in Appendix B.1.

Addressing convergence issues with the limit-model where $u_d \to 0$. One way to (partially) fix the convergence issues raised by homothetic EZW specification (9) involves using the non-homothetic EZW recursion (8) and taking the limit where $u_d \to 0$. Indeed, for any $u_d > 0$, if we assume that there is a finite maximal length of life, the recursion (8) provides a well-defined and non-zero sequence of utilities $(V_t)_{t\geq 0}$, and thereby a solid ground for the limit $u_d \to 0$.

Formally, starting from recursion (8), we renormalize the utility representation by setting $W_t = V_t/u_d$. The utilities $(W_t)_{t\geq 0}$ fulfill the recursion:

$$W_t = \left[\left(\frac{z_t}{u_d}\right)^{1-\sigma} + \beta\left(\pi_t W_{t+1}^{1-\gamma} + 1 - \pi_t\right)^{\frac{1-\sigma}{1-\gamma}}\right]^{\frac{1}{1-\sigma}}. \tag{18}$$

The crucial aspect is that when $u_d \to 0$, $W_t = V_t/u_d$ converges to a finite limit (as V_t also tends to zero), and not to infinity.[16]

Indeed since $\sigma > 1$, we have $\lim_{u_d\to 0}\left(\frac{z_t}{u_d}\right)^{1-\sigma} = 0$. Equation (18) therefore implies that $\lim_{u_d\to 0} W_t = \chi_t$ where $(\chi_t)_{t\geq 0}$ is defined by:

$$\chi_t = \beta^{\frac{1}{1-\sigma}}\left(\pi_t \chi_{t+1}^{1-\gamma} + 1 - \pi_t\right)^{\frac{1}{1-\gamma}}. \tag{19}$$

If $\beta^{\frac{1-\gamma}{1-\sigma}}\pi_t < 1$ for large t (which holds in the case of a finite life), the forward recursion admits a unique solution, which is finite for all t. Using (18) and the fact

[16]The limit V_t/u_d is also considered in Córdoba and Ripoll (2017, p. 1503), but assumed to be equal to ∞ as it was not noticed that V_t also tends to zero when $u_d \to 0$. Córdoba et al. (2020) consider the limit $u_d \to 0$ only when $\sigma = 1$. When $\sigma = 1$, the statement $V_t/u_d \to \infty$ is correct, but, as we show, it does not extend to the case where $\sigma > 1$.

that $\lim_{u_d \to 0} W_t = \chi_t$, one can compute the marginal rate of substitution between consumption in period $t+1$ and consumption in period t:

$$\lim_{u_d \to 0} \frac{\frac{\partial W_t}{\partial z_{t+1}}}{\frac{\partial W_t}{\partial z_t}} = \beta \pi_t \left(\pi_t + (1-\pi_t) \chi_{t+1}^{\gamma-1} \right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t} \right)^{-\sigma}, \tag{20}$$

This is different from the formula shown in (11). One can actually notice that the expression given in (11) would follow from (20) if one mistakenly assumed that $\chi_t = \infty$. As we explain in greater detail in Appendix B.2, the limit model where $u_d \to 0$ corresponds to a setting where consumption and survival are complementary (as with the standard additive model), but where the willingness-to-pay for mortality risk reduction is infinite. When employed in applications looking at consumption smoothing, while assuming exogenous mortality rates, the limit model yields predictions that are relatively close to those of the additive model, but very different from the predictions obtained from the homothetic specification (9). These aspects are illustrated in Figure 2 (with the same calibration as in Figure 1). The limit model is nevertheless unsuitable for dealing with endogenous mortality choices. Indeed, since it assumes that the value of life is infinite, it will systematically suggest that the optimal strategy would be to lower mortality as much as possible, independently of the related cost.

2.7 Risk-sensitive preferences (exponential ϕ)

The risk-sensitive specification (i.e., the case of an exponential function ϕ) was initially introduced by Hansen and Sargent (1995) in an infinite horizon setting and later adapted to the problem of intertemporal choice under uncertain lifespan in Bommier (2014) and Bommier et al. (2020).[17] Since the function $\phi : x \mapsto \frac{1-e^{-kx}}{k}$ is well-defined and increasing on the whole set $\mathbb{R}$, there is no domain problem. Indeed,

[17]The "multiplicative preferences" axiomatized in Bommier (2013) can also be viewed as a particular case of risk-sensitive preferences where β is set to 1. Such preferences can match empirical consumption profiles and have been used in Bommier and Villeneuve (2012) and Bommier and LeGrand (2014) to study the value of life and the demand for annuities, respectively.

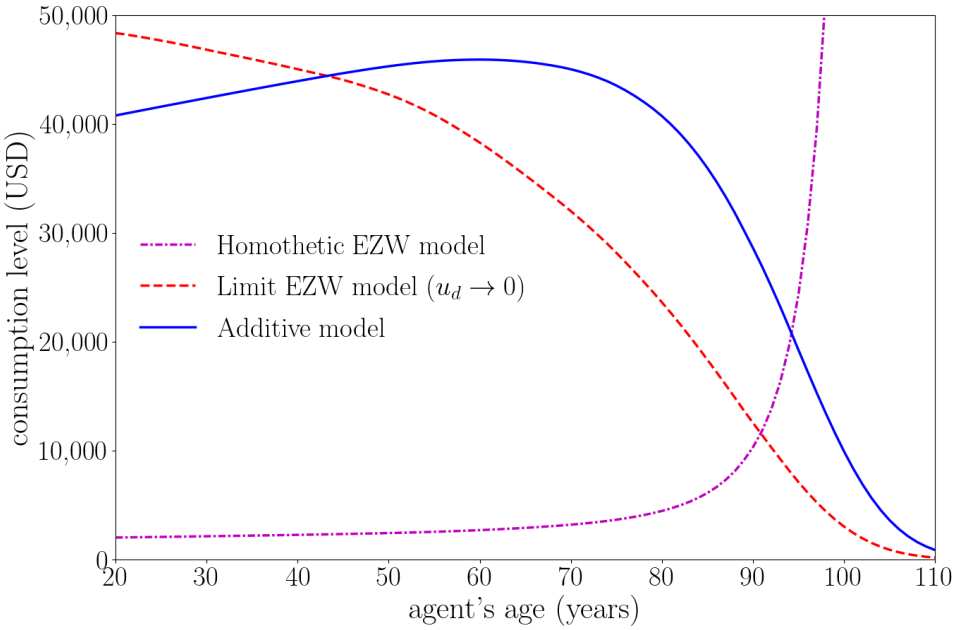

Figure 2: Consumption profiles implied by the additive, the homothetic EZW, and the $u_d \to 0$-limit models.

ϕ is defined on the convex hull of $u(\mathbb{R}_+) \cup u_d$, regardless of the choice of u and u_d. Moreover, we can easily check that, because ϕ is exponential, re-normalizing u by adding the same constant to both u_l and u_d, while keeping ϕ unchanged, has no impact on preferences. We can therefore set $u_d = 0$ without loss of generality. With $f(x) = x$, the recursion (4) defining the utilities $(V_t)_{t\geq 0}$ then becomes:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k}\log(\pi_t E[e^{-kV_{t+1}}] + 1 - \pi_t), \tag{21}$$

where the parameter k drives risk aversion. The case $k = 0$ yields by continuity the additive specification of recursion (6).[18]

The risk-sensitive specification offers a theoretically appealing framework, in which preferences are always well-defined and monotone. Indeed, as shown in Bommier et al. (2017), risk-sensitive preferences are monotone with respect to first-order

[18] In order to deal with the case of an IES equal to one (i.e., $\sigma = 1$), one would simply have to set:

$$V_t = \log(z_t) + u_l - \frac{\beta}{k}\log(\pi_t E[e^{-kV_{t+1}}] + 1 - \pi_t).$$

stochastic dominance. Monotonicity guarantees that dominated choices are ruled out and offers an intuitive interpretation for risk aversion. Indeed, risk aversion can then be understood in terms of how bad states are weighted compared to good states when making choice under uncertainty. Furthermore, flexibility is afforded by four degrees of freedom: σ, k, β, and u_l, which determine the IES, risk aversion, time preferences, and the utility gap between life and death (and thereby the value of mortality risk reduction), respectively.

As shown in Figure 3, risk-sensitive preferences generate plausible hump-shaped consumption profiles.[19] The predictions diverge from those of the additive model, because of the role of risk aversion which is extensively discussed in Bommier et al. (2020). The difference remains, however, quantitatively limited and would actually vanish for very small k.

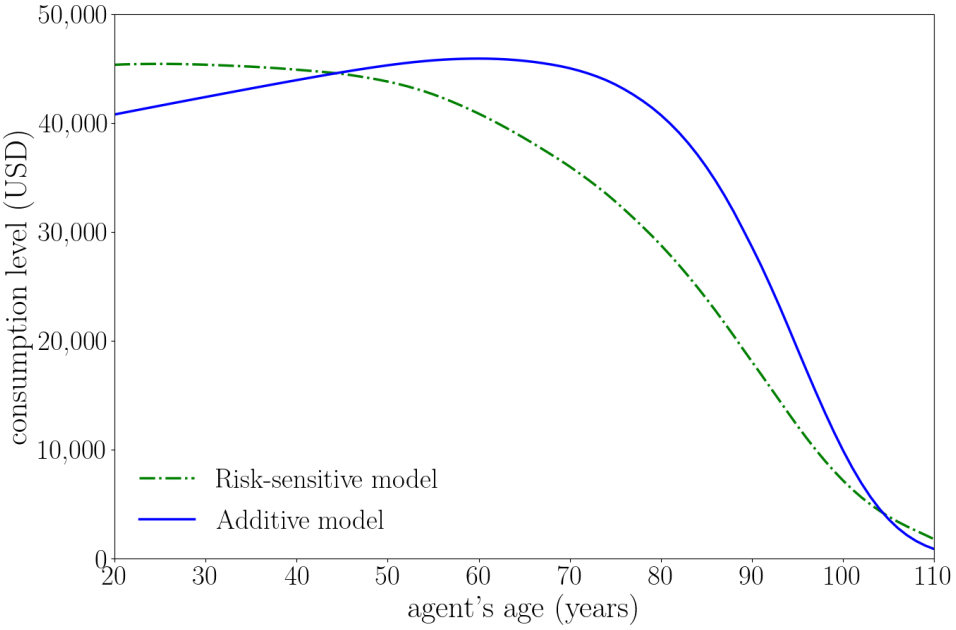

Figure 3: Consumption profiles implied by the additive and the RS models.

[19]For the calibration, we choose the value of k from Bommier et al. (2020) calibrated using annuity data and set u_l such that the value of life at 45 is 300 times the consumption at age 45. Note that as explained in Bommier et al. (2020), the results are not very sensitive to the precise target of the value-of-life calibration (as long as this target is large enough). The remaining of the calibration is the same as for Figures 1 and 2.

In such a framework the value of mortality risk reduction (16) has the following expression:

$$MRR_t = \frac{\beta}{k} z_t^{\sigma} \frac{1 - E[e^{-kV_{t+1}}]}{\pi_t E[e^{-kV_{t+1}}] + 1 - \pi_t}. \tag{22}$$

We observe in particular that the willingness-to-pay for mortality risk reduction is increasing in the continuation utility V_{t+1}. This means that people will make higher investment to reduce their mortality risk if they expect a pleasant future rather than a poor one. This is in line with what additive models suggest and with economic intuition but contrasts with the predictions of the homothetic EZW specification discussed in Section 2.6.

An extension of risk-sensitive preferences to account for ambiguity and ambiguity aversion. Mortality crises most often come with significant uncertainty regarding their magnitude. It may for example take months or years to quantify additional mortality related to a newly discovered disease (think of AIDS for example). Immediate action may nevertheless be needed. In such a case, one needs to account for ambiguity in survival probabilities and, possibly, for ambiguity aversion. An interesting aspect of the risk-sensitive framework is that it can easily be extended to account for ambiguity aversion while preserving monotonicity and differentiability (and hence tractability). A detailed presentation can be found in Bommier and LeGrand (2014).

When the survival probability is imperfectly unknown, and described by a random variable $\widetilde{\pi}_t$, the utility can be defined by the following recursion:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k_A} \log \left(E_{\widetilde{\pi}_t} \left[\exp \left(\frac{k_A}{k} \log \left(\widetilde{\pi}_t E[e^{-kV_{t+1}}] + 1 - \widetilde{\pi}_t \right) \right) \right] \right), \tag{23}$$

where k_A is the ambiguity aversion coefficient and k still denotes the risk aversion coefficient. The operator $E_{\widetilde{\pi}_t}[\cdot]$ is the expectation over $\widetilde{\pi}_t$. With $k_A = 0$ (ambiguity neutrality), the model reduces to the risk-sensitive specification (21) with $\pi_t = E[\widetilde{\pi}_t]$ (under ambiguity neutrality only the average survival probability matters). The

additive model (6) is obtained when $k_A = k = 0$. When $k = 0$ and $k_A > 0$ we get:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k_A} \log \left(E_{\widetilde{\pi}_t} \left[\exp(-k_A \widetilde{\pi}_t E[V_{t+1}]) \right] \right). \tag{24}$$

To connect this specification to the previous literature, note that (24) can be written as:

$$V_t = \phi^{-1} \left(E_{\widetilde{\pi}_t} \left[\phi \left(u(z_t) + \beta \widetilde{\pi}_t E[V_{t+1}] \right) \right] \right),$$

where $\phi(x) = \exp(-\frac{k_A}{\beta} x)$ and $u(z_t) = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l$. Written this way, the model appears to be a multi-period version of the Treich (2010) and Bleichrodt et al. (2019) models, while assuming constant absolute ambiguity aversion (which is necessary for preference monotonicity in multi-period settings), and additive separability of preferences under risk.[20] The more general monotone specification, given in (23), relaxes the assumption of additive separability of preferences under risk. As illustrated in André et al. (2020), this model can be used to jointly analyze the impacts of risk and ambiguity aversion.

We summarize the properties of the four models in Table 1. Overall, we view the risk-sensitive specifications (21) – or (23) in the presence of ambiguity – as the best way to extend the standard additive specification, affording flexibility to account for risk and ambiguity aversion, while avoiding the critical shortcomings of the isolelastic specification of Section 2.6. Importantly, the risk-sensitive specification maintains two fundamental features of the additive model: recursivity and monotonicity. The former, recursivity, is key for tractability. The latter, monotonicity, has long been considered as inherent to rationality (see e.g., Arrow, 1951) and is key to gain an intuitive understanding of the role of risk and ambiguity aversion (Bommier et al. 2017).

[20] The Treich (2010) and Bleichrodt et al. (2019) models also feature a bequest motive, which we have not considered so far. To account for bequest motives specification (23) should be replaced by:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k_A} \log \left(E_{\widetilde{\pi}_t} \left[\exp \left(\frac{k_A}{k} \log \left(\widetilde{\pi}_t E[e^{-kV_{t+1}}] + (1-\widetilde{\pi}_t) E[e^{-k_R W_{t+1}}] \right) \right) \right] \right),$$

where W_{t+1} would be the utility derived from bequests in period $t+1$.

Property	Additive	Risk-sensitive	EZW with $\gamma > 1$	EZW with $\gamma < 1$
Disentangling IES and risk aversion	No	Yes	Yes	Yes
Accommodating IES < 1 and realistic mortality	Yes	Yes	Yes	No
Accommodating IES > 1 and realistic mortality	Yes	Yes	No	Yes
Monotone (wrt FSD)	Yes	Yes	No	No

Note: γ is the risk aversion parameter in the EZW specification (see (8)).

Table 1: Summary of model properties

3 Application to epidemics

We now illustrate how risk-sensitive preferences provide new insights on the trade-off between consumption and mortality in the context of optimal epidemic mitigation. We build on the framework of Hall et al. (2020) and apply it to the pandemics of Covid-19 (Section 3.2 focusing on risk aversion and Section 3.4 for the inclusion of ambiguity and ambiguity aversion) and of the 1918 influenza (Section 3.3). One reason to focus on these two real-world cases is that they feature very contrasted age-specific mortality rates. We start with presenting the setup.

3.1 General case

We consider a population of size normalized to 1, initially containing agents of different ages a in proportions $(\omega_a)_a$, with $\omega_a \in [0,1]$ and $\sum_a \omega_a = 1$. Agents of age a are endowed with risk-sensitive preferences represented by utility function V_a defined in recursion (21). We assume that a pandemic (either Covid-19 or 1918 influenza in our applications) implies an age-specific impact on survival probabilities, which diminish from π_a to $\pi_a - \delta_a$ for one year. A benevolent planner seeks to determine which share α of current consumption agents are willing to give up in exchange for being rid of the excess mortality risk. We further simplify the framework by

assuming that consumption, z, is constant throughout ages and that agents' discount factor β is 1. Let $\lambda = 1 - \alpha$ and denote by $V_a(\lambda, \delta)$ the current utility of an agent of age a whose current consumption is λz (instead of z) and her next-period survival $\pi_a - \delta$. One has :

$$V_a(\lambda, \delta) = u(\lambda z) - \frac{1}{k} \log((\pi_a - \delta)e^{-kV_{a+1}(1,0)} + 1 - \pi_a + \delta), \tag{25}$$

where next-period utility is $V_{a+1}(1, 0)$ since the pandemic effects are assumed to last for one year only. The criterion of the benevolent planner is:

$$\begin{aligned} W(\lambda, (\delta_a)_a) &= \sum_a \omega_a V_a(\lambda, \delta_a) \\ &= u(\lambda z) - \frac{1}{k} \sum_a \omega_a \log((\pi_a - \delta_a)e^{-kV_{a+1}(1,0)} + 1 - \pi_a + \delta_a). \end{aligned} \tag{26}$$

The planner seeks to determine how much of the current consumption z can be reduced so as to offset in terms of welfare the extra mortality risk, which corresponds formally to the equivalence $W(1, (\delta_a)_a) = W(\lambda, 0)$, or using (25) and (26) to:

$$u(z) - u(\lambda z) = \frac{1}{k} \sum_a \omega_a \log \left(1 + \delta_a \frac{1 - e^{-kV_{a+1}(1,0)}}{\pi_a e^{-kV_{a+1}(1,0)} + 1 - \pi_a} \right). \tag{27}$$

If we assume that δ_a and λ are both small, we obtain:

$$\alpha = 1 - \lambda \approx \frac{1}{zu'(z)} \sum_a \delta_a \omega_a \frac{1}{k} \frac{1 - e^{-kV_{a+1}(1,0)}}{\pi_a e^{-kV_{a+1}(1,0)} + 1 - \pi_a}, \tag{28}$$

where the latter relationship can be shown to fall back on the expression in Hall et al. (2020, equation (4)) by taking the limit for $k \to 0$.

To interpret further equation (28), we can conduct a first-order Taylor expansion in k of expression (27) for α. We obtain the following approximation for small k:

$$\alpha \approx \underbrace{v \sum_a \delta_a \omega_a \mathbb{E}_{a+1}[\widetilde{T}]}_{\text{additive term}} + v \frac{ku(z)}{2} \left(\sum_a \delta_a \omega_a \left(\underbrace{(2\pi_a - 1)\, \mathbb{E}_{a+1}[\widetilde{T}]^2}_{\text{gain proportional to } \mathbb{E}_{a+1}[\widetilde{T}]^2} - \underbrace{\mathbb{V}_{a+1}[\widetilde{T}]}_{\text{loss due to risk}} \right) \right), \tag{29}$$

where $v = \frac{u(z)}{zu'(z)}$ is, as in Hall et al. (2020), the value of a year of life relative to consumption, $\mathbb{E}_{a+1}[\widetilde{T}]$ the life expectancy at age $a + 1$ and $\mathbb{V}_{a+1}[\widetilde{T}]$ the variance of

lifespans at age $a+1$. The expression (29) consists of three terms. The first term ("additive term") is the same as in the additive model, which is consistent with the fact that the risk-sensitive model reduces to the additive model when $k = 0$. The second term ("gain") is positive when $\pi_a > 0.5$, which is the case for all ages except for very old ages. For instance, in the HMD data we use, this is only the case for people of age 106 and older. This term is proportional to the square of life expectancy and reflects that agents with a long life expectancy are willing to pay more to be rid of the additional mortality risk that the epidemic generates (provided their survival probability is high enough).[21] The last term is proportional to the variance of lifespans at age $a+1$: the more uncertain the lifespan, the less the agent is willing to pay for avoiding the extra mortality risk. We expect the sum of the two last terms, scaled by the risk aversion parameter k, to be positive at younger ages, and to be negative at older ages. The overall impact of risk aversion is thus not clear-cut and may increase or decrease the value obtained with the additive model, depending on how the epidemic affects younger people as opposed to older people.

Introducing ambiguity and ambiguity aversion. We assume that there is some uncertainty regarding the mortality risk of the epidemic. Instead of being fully predictable the additional mortality risk at age a is described by a random variable, which we we denote $\tilde{\delta}_a$. We are still interested in the drop in consumption that would be acceptable in terms of welfare, from the planner's perspective, to be rid of the additional mortality. Note that this ambiguity has no impact in the case of the additive model, since the additive model features ambiguity neutrality. For the risk-sensitive model, we consider the utility $V_a(\lambda, \tilde{\delta}_a)$ of an agent of age a facing a current reduction in consumption equal to $1-\lambda$ and extra mortality implied by the epidemic. Due to the presence of ambiguity, we need to rely on recursion (23), which

[21] The fact that the effect is positive only if the survival probability is large enough is in line with well-known results in the literature on optimal prevention stating that risk aversion may fail to enhance optimal prevention when the probability of having an accident is not small. See in particular Dionne and Eeckhoudt (1985), or Jullien et al. (1999) among others.

yields:

$$V_a(\lambda, \tilde{\delta}_a) = u(\lambda z) \tag{30}$$
$$- \frac{1}{k_A} \log E_{\tilde{\delta}}^{\sim} \left[\exp\left(\frac{k_A}{k_R} \log\left((\pi_a - \tilde{\delta}_a) e^{-kV_{a+1}(1,0)} + 1 - \pi_a + \tilde{\delta}_a \right) \right) \right],$$

where $E_{\tilde{\delta}}^{\sim}[\cdot]$ is the expectation operator over the uncertainty in $\tilde{\delta}$. Note that agents' next-period utility is $V_{a+1}(1,0)$, as defined in equation (25), and does not feature any ambiguity since the pandemic effects, including ambiguity, are assumed to last for one year only.

The welfare criterion of the benevolent planner is similar to the one in the no-ambiguity case: $W(\lambda, (\tilde{\delta}_a)_a) = \sum_a \omega_a V_a(\lambda, \tilde{\delta}_a)$. The planner again aims to determine the reduction $\alpha = 1 - \lambda$ in current consumption z that, in terms of welfare, could offset the extra mortality uncertainty implied by the epidemic. This formally corresponds to: $W(1, (\tilde{\delta}_a)_a) = W(\lambda, 0)$, or to:

$$1 - \lambda^{1-\sigma} = \frac{1-\sigma}{z^{1-\sigma}} \frac{1}{k_A} \sum_a \omega_a \log E_{\tilde{\delta}} \left[\left(1 + \tilde{\delta}_a \frac{1 - e^{-kV_{a+1}(1,0)}}{\pi_a e^{-kV_{a+1}(1,0)} + 1 - \pi_a} \right)^{\frac{k_A}{k}} \right]. \tag{31}$$

When there is no ambiguity ($\tilde{\delta}_a = \delta_a$ for all a) or when there is ambiguity neutrality ($k_A = k$), expression (31) falls back on the formula (27) we had in the risk case.

3.2 The case of Covid-19

We now apply the computations of Section 3.1 to the Covid-19 pandemic. As in Hall et al. (2020), we assume that u features a constant IES (as in equation (5)) and its inverse is set to $\sigma = 2$. We also use their calibration for the consumption level z, set to $\$\,45,000$, and for the value of life at age 40, set to $\$\,10.4$m and used to determine u_l. The population shares $(\omega_a)_a$ are also those of the US total population in 2018, as reported by the US Census Bureau.[22]

We first consider the case where there is no ambiguity in survival probabilities

[22] `https://data.census.gov/cedsci/table?q=population&tid=ACSDP1Y2018.DP05`.

(the role ambiguity will be discussed in Section 3.4). For the Covid-19 age-specific mortality profile $(\delta_a)_a$, Hall et al. (2020) used the data of Ferguson et al. (2020) that was the only reliable source of data available when they wrote their paper. However, these data were estimated very early in the pandemic (March 2020), and we take advantage of more recent estimates that cover a broader set of observations. The data we use are taken from the meta-estimation of Levin et al. (2020), who report an exponential relationship between age and the infection fatality ratio (IFR), which is the proportion of people infected who die from the disease:

$$\log_{10} IFR_a = -3.27 + 0.0524 \times a, \tag{32}$$

where IFR_a is the IFR at age a and $\log_{10}$ is the log in base 10. This implies that the IFR increase by 12.82% every year of age – which is slightly higher than the 11.2% reported in Hall et al. (2020) based on Ferguson et al. (2020) estimates. Regarding the average mortality rate, we use the value of 0.69% which results from the age-specific IFR profile estimated by Levin et al. (2020) applied to the 2018 US population structure and a risk of contracting the Covid-19 of 65%, identical for all ages by assumption. This implies that we have: $\delta_a = 65\% \times IFR_a$. The infection probability of 65% corresponds to a reproductive number R_0 of 2.87, which is the mid-point estimate in the meta-review of Billah et al. (2020).[23]

Finally, the survival probabilities $(\pi_a)_a$ are chosen to be those of the US total population in 2018, as reported in the Human Mortality Database.[24] For the risk-sensitive model, we set the value of $k = 0.216$ based on Bommier and LeGrand (2014), who calibrated a risk-sensitive model with $\beta = 1$ to match annuity holdings.

We report in Table 2 the share α of consumption agents are willing to relinquish

[23] We report in Appendix A.1 the results using the data of Ferguson et al. (2020), so as to ease the comparison with Hall et al. (2020). Compared to the values reported here, the planner is willing to give up a *higher* share of consumption for avoiding Covid-19 mortality risk. This stems from the fact that Ferguson et al. (2020) data overestimated the IFR for young people (especially those of age below 40).

[24] `https://www.mortality.org/cgi-bin/hmd/country.php?cntr=USA&level=1`. Hall et al. (2020) use mortality data from the Social Security Administration, which differ very slightly from HMD data.

in order to be rid of the excess mortality risk of Covid-19. We do so for both the risk-sensitive and additive models and for each report the values α obtained from the linear approximation (equation (28)) and from the exact formula (equation (27)). We do not repeat here the expressions for the additive model as they are identical to those in Hall et al. (2020) and can be deduced by taking the limit $k \to 0$ in (27) and (28). When using the additive model, agents are willing to give up 50.4% of their current consumption for avoiding Covid-19 mortality risk according to the linear approximation, while this drop in consumption reduces to 33.5% according to the exact formula, which takes into account non-linearities.

When using the risk-sensitive model, the share α of consumption to relinquish is smaller than in the additive setup. With the exact formula, the acceptable drop in consumption equals 28.3% with risk-sensitive preferences, compared to 33.5% for additive ones. Taking risk aversion into account therefore diminishes the share of consumption agents are willing to give up, by 5 percentage points approximately (for the exact formula).

Computational method	Additive model	Risk-Sensitive model
Linear approximation (28)	50.4	39.6
Exact formula (27)	33.5	28.3

Table 2: The share α of consumption (in %) to relinquish to be rid of Covid-19 mortality risk. Computations are based on Levin et al. (2020) data and an average mortality rate of 0.69%

The differences between the results implied by the additive and risk-sensitive models are a direct implication of the role of risk aversion which leads to putting greater weight on more adverse consequences. To illustrate this, we plot in Figure 4 the parameters $(\alpha_a)_a$ of equation (28) as a function of age in both the additive and the risk-sensitive models. Each α_a represents the drop in consumption that a population only made of agents with age a is willing to accept to be rid of an extra mortality

risk of 0.1%.[25] It can be seen that the $(\alpha_a)_a$ are decreasing with age for both models, showing that, when assuming a flat consumption profile, younger agents are more willing to give up consumption than older ones for a given reduction in mortality risk. This reflects the fact that dying young is a more adverse event than dying old.

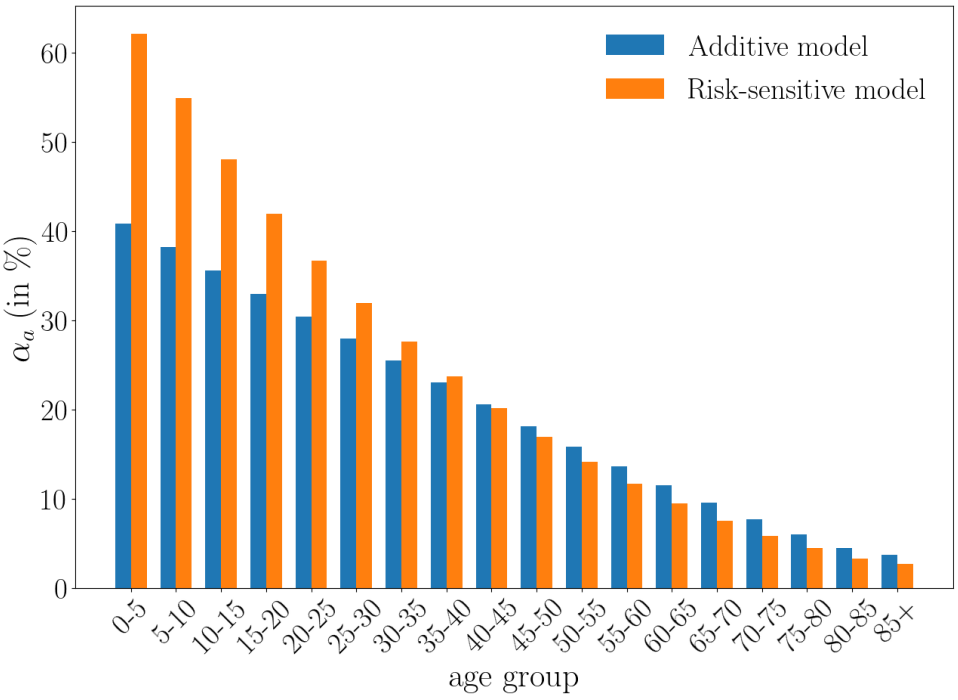

Figure 4: Profiles of the (α_a) parameters as a function of age.

However, the difference between additive and risk-sensitive models is that, due to greater risk aversion, short lives are comparatively a greater source of concern in the risk-sensitive model than in the additive model, which makes young agents willing to pay more with risk-sensitive preferences than with additive ones. The opposite holds for older agents who are more willing to pay in the additive model than in the risk-sensitive one. The calibration being performed on the willingness-to-pay for mortality risk reduction at age 40, the differences between models for age groups 35-40 and 40-45 are very modest. Since Covid-19 mainly affects older people, this

[25]Note that for a mortality risk higher than 0.1%, the value of α_a for some ages a could be higher than 100%. This simply reflects the limitation of the linear approximation of equation (28) that is only valid for low mortality risks. As can be seen in Table 2, the linear approximation is not very precise in the case of Covid-19 either for the additive or the risk-sensitive models.

explains why overall, the drop in consumption is smaller with the risk-sensitive model than with the additive one.

3.3 The case of 1918 influenza

We contrast the results of the Covid-19 with those obtained when considering the 1918 influenza, a disease that also affected young people strongly. We use the mortality data provided in Taubenberger and Morens (2006) based on Collins (1931). The age-profile of mortality risk is plotted in Figure 5.[26] As can be seen, the age mortality profile has a "W-shape", where young adults are also strongly affected by the disease. This shape is peculiar to the 1918 influenza, since regular influenza epidemics exhibit a U-shape mortality profile.

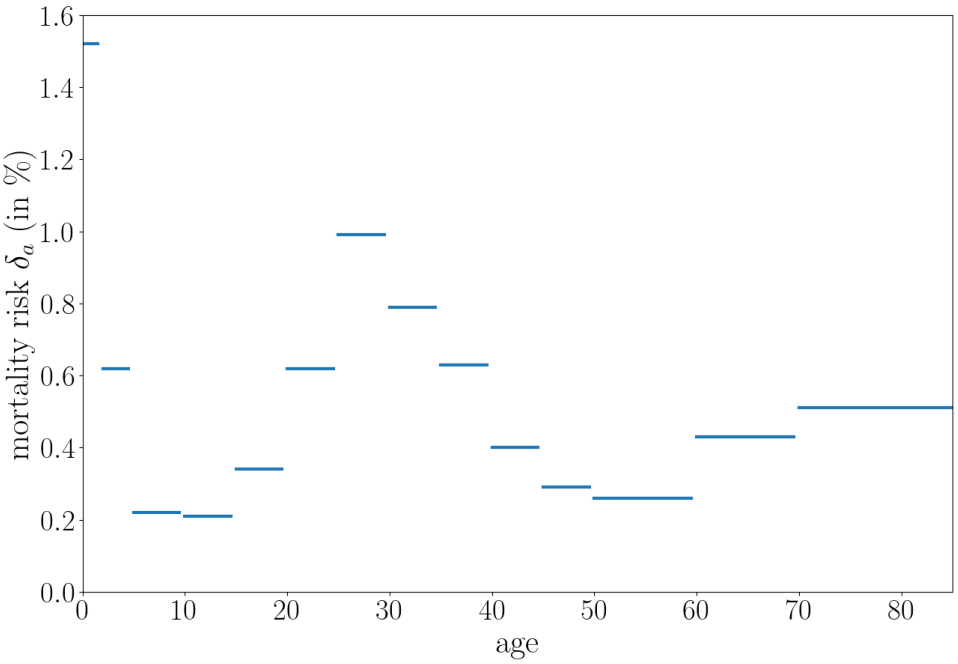

Figure 5: Age-profile of the mortality risk (δ_a) for 1918 influenza.

We report in Table 3 the share of consumption agents are willing to relinquish to be rid of the 1918 influenza mortality risk. This corresponds to the quantity α

[26]We directly use here the extra mortality risk implied by 1918 influenza. In the case of Covid-19, such data is not yet available and, as explained in Section 3.2, we estimate it based on the Infection Fatality Ratio (probability of dying once infected) and an infection probability of 65%.

computed using the exact formula (28). We do not report the values computed with the linear approximation here. Compared to Table 2, we only change the profile (δ_a), keeping the rest of the calibration unchanged. The value in Table 3 should therefore be interpreted as the willingness-to-pay to be rid of 1918 influenza mortality risk in the US of 2018 (and not of 1918). To ease the comparison with Covid-19, we consider two mortality patterns: the actual 1918 influenza one, corresponding to the mortality data of Figure 5 and a rescaled version that would yield the same average mortality rate as Covid-19. This "rescaled 1918 influenza scenario" can thus be seen as a fictive pandemic that would feature the same average mortality as Covid-19 but with the age-specific profile of the 1918-infleunza.[27] Unsurprisingly, the values of

Mortality pattern	Additive model	Risk-Sensitive model
1918 influenza	56.9	60.9
Rescaled 1918 influenza	64.8	68.4

Table 3: The share α of consumption (in %) to relinquish to be rid of 1918 influenza mortality risk, computed using the exact formula (28).

α for both models are higher for the 1918 influenza than for Covid-19, even when controlling for average mortality. This is due to the fact that younger people suffer comparatively much more from the 1918 influenza that from Covid-19. Furthermore, compared to Table 2, the relative outcomes of the additive and risk-sensitive models are reversed. With Covid-19 affecting older people disproportionately, the additive model tended to overestimate α compared to the risk-sensitive model. With the 1918 influenza also strongly affecting younger people, this is the opposite and the additive model tends to underestimate the value of α compared to the risk-sensitive model.

[27]The average mortality is actually slightly smaller for the 1918 influenza (about 0.5%) than for the Covid-19 (about 0.69%).

3.4 The case of Covid-19 with ambiguity

We now turn to a quantitative assessment of the impact of ambiguity for Covid-19. We only need to calibrate the uncertainty for Covid-19 mortality probability $(\tilde{\delta}_a)_a$, since the rest of the calibration is identical to the no-ambiguity case of Section 3.2. To do so, we take advantage of the regression results of Levin et al. (2020) and interpret the standard errors in their regression as uncertainty. More precisely, we assume that the probability of being infected remains equal to 65% and that $\tilde{\delta}_a = 65\% \times \widetilde{IFR}_a$, where $\widetilde{IFR}_a$ follows a log-normal distribution, with mean IFR_a and dispersion parameter σ_a. We thus have IID standard normal variables $(\tilde{\varepsilon}_a)_a$, $\tilde{\varepsilon}_a \sim_{IID} \mathcal{N}(0,1)$, such that:

$$\widetilde{IFR}_a = IFR_a \, e^{\sigma_a \tilde{\varepsilon}_a - \frac{\sigma_a^2}{2}}, \tag{33}$$

where the term $-\frac{\sigma_a^2}{2}$ guarantees that the mean of $\widetilde{IFR}_a$ is IFR_a. The specification for IFR_a is the point estimate in the regression of Levin et al. (2020) and is the same as in the no-ambiguity case seen in equation (32). We specify the dispersion σ_a to be equal to the standard error in the regression results of Levin et al. (2020), so that we have $\sigma_a = (0.07 + 0.0013 \times a)\log(10)$. To illustrate the implication of our calibration for uncertainty, we plot in Figure 8 of Appendix A.2 the mean value of $\widetilde{IFR}_a$ and its associated 95% confidence interval as a function of age. As can be seen, uncertainty about the IFR increases with age.

We use this calibration with equation (6) to compute the drop $\alpha = 1 - \lambda$ in current consumption z that would be acceptable for the planner to be rid of the extra Covid-19 mortality in the ambiguity case. To illustrate the role of ambiguity aversion, we compute the value of α as a function of the ambiguity aversion parameter, k_A. We report the values in Figure 6. The value of k_A is expressed in units of k (e.g., the value 1 on the x-axis corresponds to $k_A = k = 0.216$). As can be seen, the impact of ambiguity aversion is rather modest and a value of $k_A = 80k$ implies an increase in

α of approximately 1 percentage point compared to the no-ambiguity situation.[28]

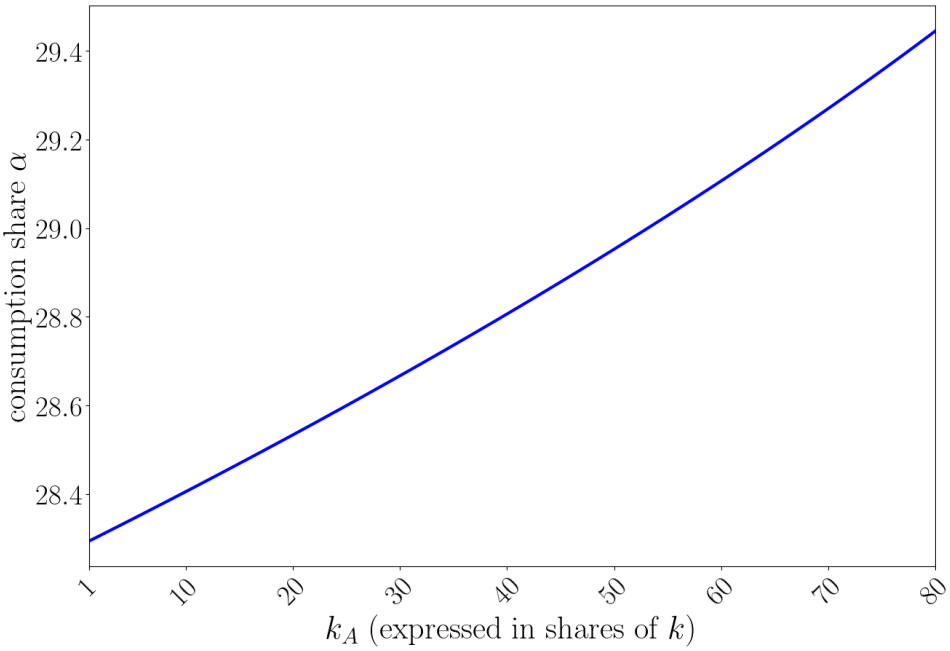

Figure 6: Share α as a function of k_A (expressed in k_A/k) in the risk-sensitive model with ambiguity.

In conclusion, ambiguity has a small quantitative effect in the context of Covid-19, even for large values of the ambiguity aversion parameter. This suggests that, when departing from the additive model, the main correction occurs when properly accounting for risk aversion. Accounting for the aversion one might have for uncertainties in risk estimates (that is for ambiguity aversion) is theoretically appealing, but seems to yield second-order corrections in our quantitative exercise.

4 Conclusion

Finding the right policy in the case of a sanitary crisis typically amounts to making trade-offs between mortality and consumption (or wealth). The recent Covid-19 crisis has shown how delicate such issues can be. Risk-reduction measures, such

[28] In Figure 9 of Appendix A.3, we show that further introducing uncertainty on the reproductive number R_0 (and thereby on the probability of becoming infected by Covid-19) adds another percent to α (still for the value $k_A = 80k$).

as lock-downs, have been strongly criticized both for being excessive and for not being severe enough. According to the micro-economic tradition, the choice of the social planner should be based on revealed preferences for mortality risk reduction. The difficulty, however, is that estimates on the willingness-to-pay for mortality risk reduction are most often based on samples of workers who, by nature, are not representative of the older and younger populations. Economists are then left with no other option than to make the best extrapolation they can. The standard approach involves using an additive model, which constrains risk aversion to be equal to the inverse of the IES. Such a property is anything but neutral as it drives how particularly bad consequences are weighted compared to not so adverse ones.

The most popular way to disentangle risk aversion and intertemporal substitutability consists of using recursive preferences as initially suggested by Kreps and Porteus (1978) and Epstein and Zin (1989). In the current paper, we have shown how this line of research can contribute to the value-of-life literature. Our message is twofold. First, we highlight that the homothetic EZW specification is unsuitable in the context of mortality risk. Second, we show that the risk-sensitive preferences initially introduced by Hansen and Sargent (1995), and shown to be the only class of monotone recursive preferences that disentangles risk aversion and IES, can provide new insights for lifecycle analysis. The main benefit of using such preferences is their ability to exhibit greater or lower levels of risk aversion which, unsurprisingly, is something that matters when considering mortality risks.

In practice, increasing risk aversion leads to exhibiting greater concern for deaths at younger ages, and relatively less for deaths at older ages, reflecting that the death of a young individual is considered as being a more dramatic consequence than the death of an elderly person. We illustrated the relevance of this point by contrasting Covid-19 with the 1918 influenza outbreak, but there are of course many other cases where accounting for risk aversion could yield interesting new insights: one could for example think of the opioid epidemic or of the gun violence public health epidemic.[29]

[29]See `https://www.whitehouse.gov/briefing-room/statements-releases/2021/04/07/`

Appendix

A Supplementary results

A.1 Considering the age-specific IFRs of Ferguson et al. (2020) for Covid-19

We report here the results implied when using the Covid-19 age-specific IFRs of Ferguson et al. (2020). This allows us to directly compare our results to those of Hall et al. (2020).

We report in Table 4 the share α that agents are willing to give up in order to get rid of the Covid-19 mortality risk. This is similar to the results of Table 2, but computed using the IFRs of Ferguson et al. (2020) instead of those of Levin et al. (2020). For the sake of comparison with Hall et al. (2020), we consider the same mortality rates, 0.81% and 0.44%, as they do. When using the additive model, with

Mortality rate δ	Additive model	Risk-Sensitive model
Using linear approximation (28)*:*		
0.81%	71.0	57.4
0.44%	38.6	31.2
Using exact formula (27)*:*		
0.81%	41.5	36.4
0.44%	27.8	23.7

Table 4: The share α of consumption (in %) to give up to get rid of Covid-19 mortality risk. Computations are based on Ferguson et al. (2020) IFR data and average mortality rates of 0.44% and 0.81%.

an average mortality rate of 0.81%, agents are willing to give up 71.0% of their current

fact-sheet-biden-harris-administration-announces-initial-actions-to-address-the-gun-violence-public-health-epidemic/ for a statement from the Biden-Harris administration regarding gun violence in the US (April 7, 2021).

consumption for avoiding Covid-19 mortality risk using the linear approximation, while this reduces to 41.5% when accounting for non-linearities and using the exact formula. These values are very close to the ones reported in Hall et al. (2020), which are 70.5% and 41.3%, respectively. The results with the average mortality rate of 0.44% are even closer to those of Hall et al. (2020), since the gap is 0.2% for the linear approximation and 0.1% for the exact formula.[30]

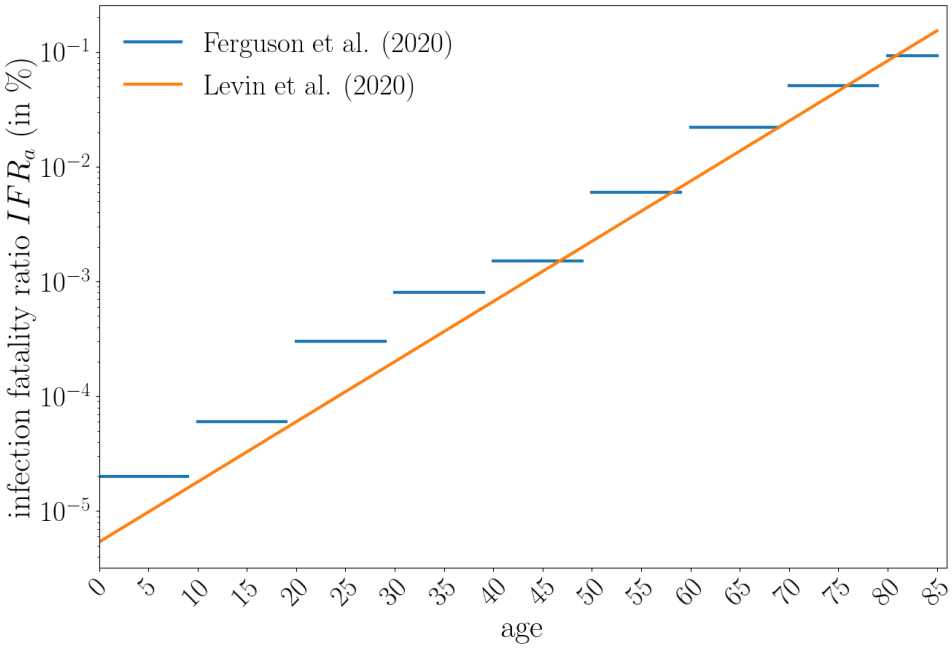

Figure 7: Comparison of age-infection-fatality-ratio profiles (in log-scale) for Covid-19 in Ferguson et al. (2020) and Levin et al. (2020).

Comparing to Table 2, we observe a drastic increase in α, when switching from the Levin et al. (2020) data and a mortality rate of 0.69% to the Ferguson et al. (2020) and a mortality of 0.81%. Indeed, the share raises from 50.4% to 71.0% with the linear approximation and from 33.5% to 41.5% with the exact formula. Approximately half of the reduction comes from the change in the average mortality

[30]There are several possible sources to explain these small differences. First, we calibrate u_l to obtain a VSL of \$10.8m, while they calibrate it to obtain a the value of a year of life (denoted v in their paper) equal to 6. The corresponding value of v in our paper is around 5.9. Second, we use mortality data for the total population from the human mortality database, while they use data from the Social Security (only provided for male and female).

rate (keeping an average mortality rate of 0.81% with the age-mortality profile of Levin et al., 2020 leads to $\alpha = 37.2\%$ for the exact formula in the additive model) and the other half in the change of the age-structure of the IFR profile. As seen in Figure 7, the data of Ferguson et al. (2020) tend to overestimate IFRs at young ages and to underestimate them at old ages, compared to the data of Levin et al. (2020). The gap is especially large for people of age smaller than 40 years.

A.2 Dispersion of the age mortality profile of Levin et al. (2020)

We plot in Figure 8 the 95% confidence interval of the age-specific IFR profile ($\widetilde{IFR}_a$) in presence of uncertainty. The expression of $\widetilde{IFR}_a$ can be found in equation (33) with the calibration given in the main text. We also plot the mean, corresponding to the quantity IFR_a in the main text (equation (32)).

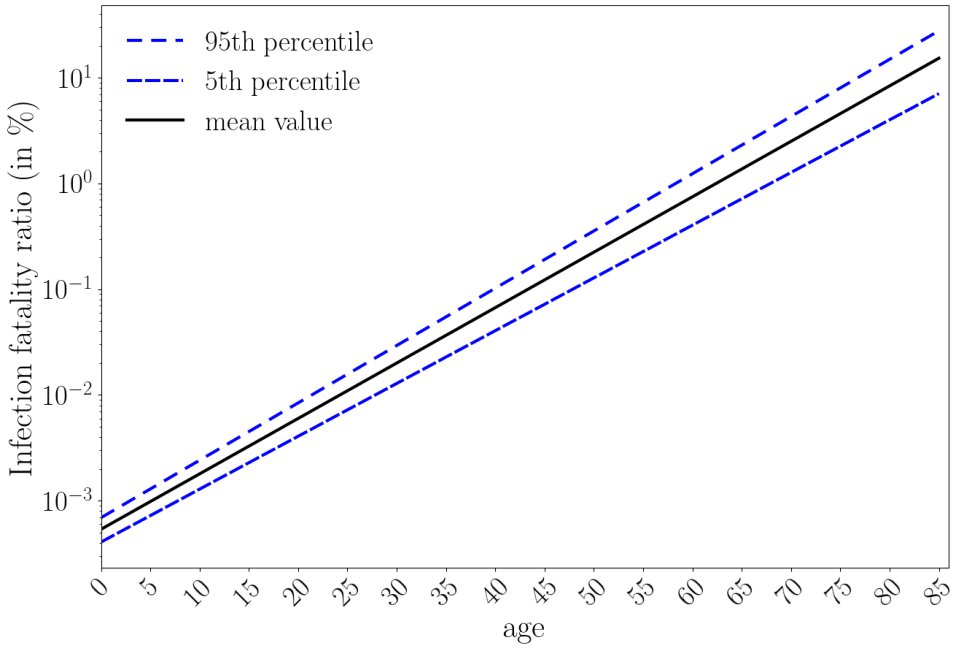

Figure 8: Mean and 95% confidence interval (in log-scale) of the age-specific IFR profile ($\widetilde{IFR}_a$) for Covid-19 in Levin et al. (2020).

A.3 Additional uncertainty on the reproductive number R_0

Billah et al. (2020) estimate the reproductive number R_0 to be equal to 2.87 with a 95% confidence interval equal to [2.39 ; 3.44]. We can observe that the logarithm of R_0 has an almost symmetric confidence interval, since $\log(3.44/2.87) \approx 0.181$ and $\log(2.87/2.39) \approx 0.183$. We will therefore make the assumption that R_0 follows a log-normal distribution with mean 2.87 and dispersion parameter σ_{R_0} equal to $\frac{1}{2}(\log(3.44/2.87) + \log(2.87/2.39)) \approx 0.1821$. We integrate it in our computation of α, similarly to equation (31) and obtain an ambiguous infection probability, denoted by $\tilde{p}_R$. Overall, the extra mortality combines two sources of ambiguity, one on the IFR and another one on the infection probability: $\tilde{\delta}_a = \tilde{p}_R \times \widetilde{IFR}_a$.

We plot in Figure 9 the value of the acceptable cut in consumption, α, as a function of k_A. When $k_A = 70k$, the value of α is approximately one percent higher than in the ambiguity case with a fixed R_0 (Figure 6) and two percents higher than in the no-ambiguity case. As can be seen, the quantitative impact of ambiguity remains of limited magnitude.

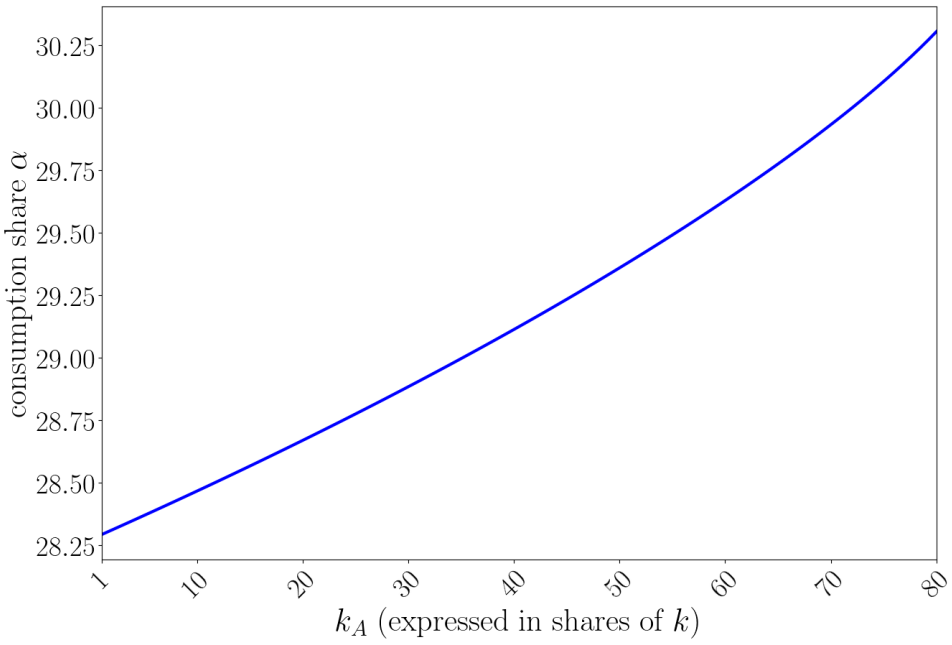

Figure 9: Share α as a function of k_A (expressed in k_A/k) in the risk-sensitive model with ambiguity, when R_0 is also ambiguous.

B About EZW preferences

B.1 Age-dependent health

In order to circumvent the difficulties raised by specification (9) and discussed in Section 2.6.2, Córdoba and Ripoll (2017) introduce an age-dependent variable H_t, which is interpreted as health. The recursive equation (9) becomes:

$$V_t = \left[H_t z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{34}$$

Since H_t depends on t, there are enough degrees of freedom to match any possible consumption profile. In principle, a rapidly declining profile H_t could potentially alter the counterfactual implications discussed in Section 2.6.2. However, with $\sigma > 1$, specification (34) assumes that utility *decreases* (rather than increases) with H_t, which hinders the interpretation of H_t as health.[31] In a similar vein, Córdoba et al. (2020) introduce age-dependent discount factors $(\beta_t)_{t\geq 0}$ which, as they explain, should decrease with age (when $\sigma > 1$) to counteract the term $\pi_t^{\frac{1-\sigma}{1-\gamma}}$. In other words, to avoid the pattern shown in Figure 1, one would have to assume that pure time preferences (which would govern impatience in absence of mortality risk) become rapidly stronger with age, so as to compensate for the implausible mortality effect discussed above.

The calibration of the "health parameters" (or of the time-varying discount factors β_t in Córdoba et al., 2020), and issues related to their potential endogeneity, may play a decisive role when discussing the impact of (exogenous or endogenous) mortality changes. Compare, for example, the utility functions in the model (34)

[31] Another option would be to consider $V_t = \left[h_t^{1-\sigma} z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$. In such a case, h_t would positively contribute to utility, but matching empirical consumption profiles would require this health profile h_t to increase and converge to ∞ at old ages. Again, this would be inconsistent with the interpretation of h as health.

and the standard additive specification of Murphy and Topel (2006):

$$V_t = \left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(\Pi_{j=t}^{\tau-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}} H_\tau z_\tau^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \text{ (from our equation (34))},$$

$$V_t^{MT} = \sum_{\tau=t}^{\infty} \beta^{\tau-t} \left(\Pi_{j=t}^{\tau-1}\pi_j\right) H_\tau^{MT} \left(\frac{z_\tau^{1-\sigma}}{1-\sigma} + u_l\right) \text{ (Murphy and Topel's model)}.$$

Córdoba and Ripoll (2017) calibrate the health profile so as to obtain the same consumption profile as Murphy and Topel (2006). In other words, they first set $H_t = \left(\Pi_{j=0}^{t-1}\pi_j\right)^{\frac{\sigma-\gamma}{1-\gamma}} H_t^{MT}$, where H_t^{MT} is the health profile chosen by Murphy and Topel (2006). A key point, however, is that they then assume the profile H_t to be exogenous and independent of survival rates when looking at the impact of mortality changes. This ultimately results in the two models forming radically different conclusions regarding the consequences of mortality decline. Murphy and Topel's model predicts that a decline in mortality would significantly increase the propensity to save (agents become more patient when survival probabilities increase), while the effect is much smaller or even opposite in the model (34) (agents become more impatient when survival probabilities increase).[32] Opting for one specification over the other will thus provide very different views regarding the impact of population aging. While both models match the same calibration targets, the implicit assumptions made about the role of health (considered as "good" in Murphy and Topel and as "bad" in (34)) lead to opposite conclusions.

B.2 The limit-model when $u_d \to 0$

We explain here why the limit-model (19) has different economic implications than the recursive representation (9). We compute marginal rates of substitutions implied by the recursive representation (18) when $u_d \to 0$. Since $\lim_{u_d \to 0} W_t = \chi_t$, we obtain

[32] Longevity extension also generates an income effect that adds to the impatience effect we emphasize, which explains why the overall impact can be ambiguous.

that when $u_d \to 0$:

$$\frac{\frac{\partial W_t}{\partial z_{t+1}}}{\frac{\partial W_t}{\partial z_t}} = \beta\pi_t \left(\pi_t + (1-\pi_t)W_{t+1}^{\gamma-1}\right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t}\right)^{-\sigma}, \tag{35}$$

$$\to_{u_d \to 0} \beta\pi_t \left(\pi_t + (1-\pi_t)\chi_{t+1}^{\gamma-1}\right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t}\right)^{-\sigma}, \tag{36}$$

and

$$\frac{\frac{\partial W_t}{\partial \pi_t}}{\frac{\partial W_t}{\partial z_t}} = u_d^{1-\sigma} \frac{\beta}{1-\gamma} z_t^{\sigma} \left(W_{t+1}^{1-\gamma} - 1\right) \left(\pi_t W_{t+1}^{1-\gamma} + 1 - \pi_t\right)^{\frac{\gamma-\sigma}{1-\gamma}} \to_{u_d \to 0} \infty.$$

Thus, the limit obtained when $u_d \to 0$ corresponds to a setting where the value of mortality risk reduction tends to infinity, while the marginal rate of substitution between consumption in period $t+1$ and consumption in period t converges to a well-defined finite limit that depends on the survival and consumption patterns.

Note that the MRS in (36) tends to zero when π_t tends to zero. This means that if survival from period t to period $t+1$ is very unlikely, the agent values an increase in consumption in period t much more than an increase in consumption in period $t+1$. While consistent with the additive model and economic intuition, this is contrary to the implications of recursion (9) (see Córdoba and Ripoll, 2017). Indeed, the recursion (9) implies, when $u_d = 0$, an MRS equal $\beta\pi_t^{\frac{1-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t}\right)^{-\sigma}$, which actually tends to ∞ (and not 0) when $\pi_t \to 0$. To fully clarify why our discussion sharply contrasts on these aspects with that of Córdoba and Ripoll (2017), it can be observed that the MRS derived from recursion (9) could seem to be obtained as the limit of equation (35) for $u_d \to 0$ if it is mistakenly assumed that $\lim_{u_d \to 0} W_{t+1} = \infty$. However, as we previously explained, a rigorous consideration of the limit (properly accounting for the indeterminate form $\frac{0}{0}$) shows that $\lim_{u_d \to 0} W_{t+1} = \chi_{t+1}$, where χ_{t+1} is finite.

C Proof of Result 1

Assume that there exists $t_0 \geq 0$ such that $V_{t_0} > 0$. It is then necessarily the case that $V_t > 0$ for all $t > t_0$.[33] Now, let us rewrite equation (9) as:

$$V_{t+1}^{1-\sigma} = \beta^{-1} \pi_t^{\frac{\sigma-1}{1-\gamma}} \left(V_t^{1-\sigma} - z_t^{1-\sigma} \right).$$

By iteration, we obtain that for all $t > t_0$:

$$V_t^{1-\sigma} = \beta^{-t} \left(\Pi_{j=t_0}^{t-1} \pi_j \right)^{\frac{\sigma-1}{1-\gamma}} \left(V_{t_0}^{1-\sigma} - \sum_{s=t_0}^{t-1} \beta^s \left(\Pi_{j=t_0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right). \tag{37}$$

Consider now the case where $t \to \infty$, while holding t_0 constant. Assuming that consumption is bounded from above, the right-hand side of (37) becomes negative, since $\sum_{s=0}^{t-1} \beta^s \left(\Pi_{j=0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma}$ diverges (as we are considering the case $\lim_{t\to\infty} \pi_t = 0$), while the left-hand side has to be positive. We thus obtain a contradiction, proving that there cannot exist a t_0 for which $V_{t_0} > 0$. This impossibility result holds when $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}} > 1$ for large t.[34]

[33] We previously established that if $V_t = 0$, then $V_\tau = 0$ for all $\tau \leq t$.

[34] A different way to prove this is by a fixed-point argument. When all V_t are positive, $V_t^{1-\sigma}$ is defined by the linear recursive equation $V_t^{1-\sigma} = z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}$, which is a contraction if and only if $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}} < 1$ for t sufficiently large. When $1 - \beta \pi_t^{\frac{1-\sigma}{1-\gamma}}$ is negative for large t, the linear recursive equation does not define a proper $V_t^{1-\sigma}$. It is worth noting that the term $1 - \beta \pi_t^{\frac{1-\sigma}{1-\gamma}}$ occurs in many instances in Córdoba and Ripoll (2017), e.g., in equations (27), (30), and (32), with, however, no mention that this term may be negative when $\gamma < 1 < \sigma$.

References

Andersson, H. and N. Treich (2011): "The Value of a Statistical Life," in *Handbook in Transport Economics*, ed. by A. de Palma, R. Lindsey, E. Quinet, and R. Vickerman, Cheltenham, UK: Edward Elgar, 396–424.

André, E., A. Bommier, and F. LeGrand (2020): "The Impact of Risk and Ambiguity Aversion on Annuity and Saving Choices," Working Paper, ETH Zurich.

Arrow, K. J. (1951): "Alternative Approaches to the Theory of Choice in Risk-Taking Situations," *Econometrica*, 19, 404–437.

Billah, M. A., M. M. Miah, and M. N. Khan (2020): "Reproductive Number of Coronavirus: A Systematic Review and Meta-Analysis Based on Global Level Evidence," *PLOS ONE*, 15, e0242128.

Bleichrodt, H., C. Courbage, and B. Rey (2019): "The Value of a Statistical Life under Changes in Ambiguity," *Journal of Risk and Uncertainty*, 58, 1–15.

Bommier, A. (2013): "Life Cycle Preferences Revisited," *Journal of European Economic Association*, 11, 1290–1319.

——— (2014): "Mortality Decline, Impatience and Aggregate Wealth Accumulation with Risk-Sensitive Preferences," CER-ETH Working Paper 14/194, ETH Zurich.

Bommier, A., D. Harenberg, and F. LeGrand (2021): "Recursive Preferences and the Value of Life: A Clarification," Working Paper, ETH Zurich.

Bommier, A., D. Harenberg, F. LeGrand, and C. O'Dea (2020): "Recursive Preferences, the Value of Life, and Household Finance," Discussion Paper 2231, Cowles Foundation.

Bommier, A., A. Kochov, and F. LeGrand (2017): "On Monotone Recursive Preferences," *Econometrica*, 85, 1433–1466.

Bommier, A. and F. LeGrand (2014): "Too Risk Averse to Purchase Insurance? A Theoretical Glance at the Annuity Puzzle," *Journal of Risk and Uncertainty*, 48, 135–166.

Bommier, A. and B. Villeneuve (2012): "Risk Aversion and the Value of Risk to Life," *Journal of Risk and Insurance*, 79, 77–103.

Collins, S. D. (1931): "Age and Sex Incidence of Influenza and Pneumonia Morbidity and Mortality in the Epidemic of 1928-29 with Comparative Data for the Epidemic of 1918-19: Based on Surveys of Families in Certain Localities in the United States following the Epidemics," *Public Health Reports (1896-1970)*, 46, 1909–1937.

Córdoba, J. C. and M. Ripoll (2017): "Risk Aversion and the Value of Life," *Review of Economic Studies*, 84, 1472–1509.

Córdoba, J. C., M. Ripoll, and S. Yang (2020): "The Full Recession: Private Versus Social Costs of Covid-19," Working Paper, Nankai University.

Dionne, G. and L. Eeckhoudt (1985): "Self-Insurance, Self-Protection and Increased Risk Aversion," *Economics Letters*, 17, 39–42.

Dong, X., B. Milholland, and J. Vijg (2016): "Evidence for a Limit to Human Lifespan," *Nature*, 538, 257–259.

Eeckhoudt, L. R. and J. K. Hammitt (2004): "Does risk aversion increase the value of mortality risk?" *Journal of Environmental Economics and Management*, 47, 13–29.

Epstein, L. G. and S. E. Zin (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Ferguson, N., D. Laydon, G. Nedjati Gilani, N. Imai, K. Ainslie, M. Baguelin, S. Bhatia, A. Boonyasiri, Z. Cucunuba Perez, G. Cuomo-Dannenburg, et al. (2020): "Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID19 Mortality and Mealthcare Demand," Report 9, London: Imperial College COVID-19 Response Team.

Greenstone, M. and V. Nigam (2020): "Does Social Distancing Matter?" *COVID Economics*, 7, 1–23.

Hall, R. E., C. I. Jones, and P. J. Klenow (2020): "Trading Off Consumption and COVID-19 Deaths," *Federal Reserve Bank of Minneapolis Quarterly Review*, 42.

Hammitt, J. K. (2020): "Valuing Mortality Risk in the Time of COVID-19," *Journal of Risk and Uncertainty*, 61, 129–154.

Hansen, L. P. and T. J. Sargent (1995): "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control*, 40, 968–971.

Havránek, T. (2015): "Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting," *Journal of the European Economic Association*, 13, 1180–1204.

Hugonnier, J., F. Pelgrin, and P. St-Amour (2013): "Health and (Other) Asset Holdings," *Review of Economic Studies*, 80, 663–710.

——— (2020): "Closing Down the Shop: Optimal Health and Wealth Dynamics near the End of Life," *Health Economics*, 29, 138–153.

Jullien, B., B. Salanié, and F. Salanié (1999): "Should More Risk-Averse Agents Exert More Effort?" *The Geneva Papers on Risk and Insurance - Theory*, 24, 19–28.

Kaplow, L. (2005): "The Value of a Statistical Life and the Coefficient of Relative Risk Aversion," *Journal of Risk and Uncertainty*, 31, 23–34.

Kreps, D. M. and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Levin, A. T., W. P. Hanage, N. Owusu-Boaitey, K. B. Cochran, S. P. Walsh, and G. Meyerowitz-Katz (2020): "Assessing the Age Specificity of Infection Fatality Rates for Covid-19: Systematic Review, Meta-Analysis, and Public Policy Implications," *European Journal of Epidemiology*, 35, 1123–1138.

Murphy, K. M. and R. H. Topel (2006): "The Value of Health and Longevity," *Journal of Political Economy*, 114, 871–904.

Robinson, L. A., R. Sullivan, and J. F. Shogren (2020): "Do the Benefits of COVID-19 Policies Exceed the Costs? Exploring Uncertainties in the Age-VSL Relationship," *Risk Analysis*, Early View.

Tallarini, T. D. J. (2000): "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics*, 45, 507–532.

Taubenberger, J. K. and D. M. Morens (2006): "1918 Influenza: the Mother of All Pandemics," *Emerging Infectious Diseases*, 12, 15–22.

Treich, N. (2010): "The Value of a Statistical Life under Ambiguity Aversion," *Journal of Environmental Economics and Management*, 59, 15 – 26.

Weil, P. (1989): "The Equity Premium Puzzle and the Risk-free Rate Puzzle," *Journal of Monetary Economics*, 24, 401–421.

Weon, B. M. and J. H. Je (2009): "Theoretical Estimation of Maximum Human Lifespan," *Biogerontology*, 10, 65–71.

Yaari, M. E. (1965): "Uncertain Lifetime, Life Insurance, and the Theory of the Consumer," *Review of Economic Studies*, 32, 137–150.

Zhang, S., M. Hardy, and D. Saunders (2018): "Epstein-Zin Preferences in Life-cycle Applications: Uncertain Lifetime and Consumption Smoothing," Working Paper, University of Waterloo.