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

Antoine Bommier François Le Grand Daniel Harenberg*

April 3, 2024

Abstract

This article investigates how recursive preferences can be used in the context of lifecycle models featuring uncertain and endogenous lifespans. We provide representation results showing how recursive preferences may be homothetic or fulfill a simple form of monotonicity with respect to first-order stochastic dominance, also known as ordinal dominance. While homotheticity appears to be very restrictive, constraining the intertemporal elasticity of substitution to be greater than one and risk aversion to be less than one, ordinal dominance points to the risk-sensitive preferences of Hansen and Sargent (1995), on which we focus in the second part of the paper. We then discuss the theoretical impact of risk aversion, and illustrate the relevance of our findings by looking at the consumption-mortality trade-offs faced by a benevolent planner during a pandemic.

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

JEL codes: G11, J17.

*Bommier: ETH Zurich, abommier@ethz.ch; Le Grand: Rennes School of Business, francois.le-grand@rennes-sb.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, Pascal St-Amour, Juan Carlos Córdoba, Marla Ripoll, anonymous referees and an editor of the journal for the feedback they provided on an earlier version of this paper. We also thank the participants at the 2023 SAET conference, at the workshops "Advances in the economics and operational control of Covid-19" organized by RSB and CNRS and "Frontiers of Modeling in Economic Epidemiology" organized by RSB and Rice University. This paper supersedes the working paper entitled *Recursive Preferences and the Value of Life: A Clarification* by the same authors (Bommier et al., 2021).

1 Introduction

Major epidemics such as the plague, 1918 influenza or the Covid-19 pandemic may cause the death of millions of people. One solution to reduce the number of fatalities is to limit human interactions, but this can have huge economic costs. Societal decision-making must therefore consider a trade-off between survival probabilities and wealth, which requires an economic value to be placed (implicitly or explicitly) on mortality risk reductions. A major difficulty is that epidemics, like most health threats, affect people very differently depending on their age. Covid-19, for example, was particularly deadly for the elderly, while the 1918 influenza impacted young adults more severely. Thus, epidemic management requires that due consideration be given to how the value of mortality risk reduction varies with the age of the people at risk. Intuitively, the mere fact that young people have (on average) more years to live than old people is enough to argue that the value of mortality risk reduction should depend on age. Indeed, it is well accepted that for a given number of expected deaths, a disease that would kill people in their 20s looks more dramatic than one that would only affect people after age 90. The challenging issue is about the quantitative adjustments to be made. By how much is a reduction in mortality at age 20 more valuable than a reduction in mortality at age 90? Lifecycle theory helps answer this question by emphasizing the role of the time horizon. The standard approach is to assume stationary preferences, so that the difference between young and old is not related to ad hoc assumptions about preferences, but to the fact that young and old people have different expectations – especially about how much time they have left to live. In practice, this approach is generally implemented using the seminal theoretical framework introduced by Yaari (1965), with agents' lifetime utility assumed to equal the sum of remaining (age-independent) instantaneous utilities weighted by survival probabilities and an exponentially decaying time-discounting factor. In such a framework, the agent's age has no direct effect on the agent's objective (preferences are stationary), but ends up playing a fundamental role since survival probabilities depend on age.

However, the widely used approach inherits major limitations from the model of Yaari (1965), which, due to its time-additive structure, cannot disentangle risk aversion from the intertemporal elasticity of substitution (IES, henceforth). It constrains intertemporal risk aversion to be equal to zero, thereby constraining risk aversion with respect to life duration in a very specific way (see Bommier, 2006). The impossibility of studying the role of risk aversion is not harmless. Indeed, since mortality is a risk, attitudes toward it are likely to depend strongly on risk aversion.

One way to avoid the lack of flexibility in Yaari's model is to retain the stationary assumption but relax time additivity. In infinite horizon settings, such a path was initiated by the seminal work of Epstein and Zin (1989) and Weil (1989) (EZW, henceforth), who build on Kreps and Porteus (1978) to propose models where utility is defined recursively (preferences are then stationary) with aggregators that are not necessarily linear. This provides sufficient flexibility to study the role of risk aversion in an intertemporal context. Such recursive models have become extremely popular. While they were initially developed to address long-standing puzzles in the macro-finance literature, they are now increasingly used in other fields such as the economics of climate change, health economics, or household finance. In the current paper, we provide a theoretical investigation of how recursive models can help clarify the role of risk aversion when discussing the value of mortality risk reduction. We then show how accounting for risk aversion can provide new insights when considering wealth-survival trade-offs, and in particular how the value of mortality risk reduction may depend on age.

Our paper consists of a theoretical part, with two contributions, followed by an application to epidemics. Our first theoretical contribution involves two representation results. We start from recursive preferences, that are monotonic and do not systematically predict a negative value of life. Monotonicity means that an agent's utility increases with consumption, while the assumption of a non-systematically negative value of life requires that an agent prefers being alive and consuming to being dead, at least for some consumption levels. Our first representation result is that if we further impose Homotheticity on these well-behaved preferences, we obtain either standard additive preferences or (homothetic) EZW preferences, but where the IES is restricted to be greater than one and the coefficient of risk aversion to be less than one. If instead of Homotheticity, we impose Ordinal Dominance (similar to monotonicity with respect to first-order stochastic dominance), we obtain risk-sensitive preferences à la Hansen and Sargent (1995), with no particular restriction on utility parameters. The only preferences that simultaneously satisfy Homotheticity and Ordinal Dominance are the standard additive preferences à la Yaari, which have the limitations mentioned above. One consequence of these representation results is that overcoming the limitations of the Yaari model requires abandoning either Ordinal Dominance or Homotheticity. Opting for Homotheticity and giving up Ordinal Dominance has three main drawbacks: (i) the agent may end up opting for dominated choices (see Bommier et al., 2017 or Bommier et al., 2020 for illustrations), (ii) the IES and risk aversion parameters are constrained in a way that is inconsistent with empirical evidence, and (iii) it involves assuming that rich and poor behave identically (up to a wealth scaling factor), which is also inconsistent with empirical evidence. Opting for Ordinal Dominance and discarding Homotheticity slightly reduces model tractability, but preserves model insights and does not constrain model parameters. Our takeaway from these representation results is that imposing Homotheticity for tractability reasons is costly. Meanwhile, imposing Ordinal Dominance leads to use the risk-sensitive framework, providing an appealing approach to discussing value of life matters, with sufficient flexibility to account for risk aversion.

Our second theoretical contribution is to show that working with risk-sensitive preferences allows one to derive results regarding the effective discount factor and the value of mortality risk reduction.[1] Regarding the effective discount factor, we

[1]We distinguish between the discount factor, which is the constant time preference parameter, denoted β in the paper, and the effective discount factor, which governs the agents' impatience and takes the impact of mortality into account. See Section 2.4.1 for a formal definition.

show that agents become more impatient (and have a lower effective discount factor) when: (i) they are more risk averse, (ii) they face a higher mortality probability, and (iii) their continuation utility is higher. The value of mortality risk reduction is also affected by risk aversion and continuation utility. From a purely theoretical perspective, the effect of risk aversion is ambiguous, reflecting general findings on the relationship between optimal prevention and risk aversion (see, e.g., Jullien et al., 1999). The effect of continuation utility is unambiguously positive: the willingness to reduce mortality risk is higher when an enjoyable life and longevity are expected in the event of survival. Numerical illustrations using a realistic mortality pattern complement these findings and provide a quantitatively clear picture: risk aversion tends to increase the value of mortality risk reduction, and also changes the relationship between age and the value of mortality risk reduction.

To illustrate the relevance of using a recursive specification, we focus on two real-world cases, the Covid-19 and 1918 influenza pandemics, and show how recursive preferences alter conclusions relating to the optimal consumption-mortality trade-off. Our results highlight the fact that the sign of the corrections depends on whether the pandemic predominantly affects older people (as in Covid-19) or younger people (as in 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. Specifically, this would argue in favor of a more rapid reopening of the economy or a less severe lock-down than is recommended when using additive preferences. For the 1918 influenza outbreak, the conclusions are the opposite. The reason for these findings is that models are typically calibrated using observations of wage-risk trade-offs made by workers (i.e., "middle-aged" people), while accounting for the mortality effects of pandemics requires us to infer the value of mortality risk reduction at younger and older ages. Thus, some form of extrapolation is required, which is typically accomplished by using a specific model of individual preferences. The structure of the model ultimately plays a key role in this extrapolation. Notably, compared to additive models, recursive models end up giving greater 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 explored with the standard additive model.

Some of our results seem to be in direct contradiction with other papers that have attempted to use recursive preferences to study value of life matters. In particular, this is the case of those by Córdoba and Ripoll (2017) and Hugonnier et al. (2013) and subsequent papers, where it is argued that homothetic EZW specifications could cope with an IES less than one.[2] This is in direct contradiction to the message of Proposition 2 of the current paper. We explain the origin of this difference in Appendix B, where we emphasize that using the models of Córdoba and Ripoll (2017) and Hugonnier et al. (2013) with an IES less than one requires an upper bound be set on mortality at older ages. In particular, these models cannot be applied to settings where life duration is bounded from above. When the assumption of a maximum life duration is relaxed, these models cannot cope with mortality patterns where life expectancy shrinks at old ages, unless one assumes an extremely low discount factor or extremely high consumption growth rates. Thus, they appear to be inappropriate for studying human behavior in old age. Moreover, we explain that when these models might be used (for example, when considering a perpetual youth model in which life expectancy remains high forever), they suggest counterintuitive predictions that are at odds with basic economic common sense.

The rest of the paper is organized as follows. In Section 2, we present our theoretical framework and derive our main theoretical contributions. In Section 3, we present our application of recursive models to epidemics. Section 4 concludes.

[2]Hugonnier et al. (2013) and Córdoba and Ripoll (2017), which were developed independently, provide very similar modeling approaches based on EZW preferences. The main difference is that Hugonnier et al. (2013) use a continuous-time setting, which makes the mathematics more complex, while Córdoba and Ripoll (2017) assume that time is discrete. Both papers have been used in several follow-up articles, including Hugonnier et al. (2020), Hugonnier et al. (2022), and Córdoba et al. (2023).

Appendix A sets out the proofs of our theoretical results. Appendix B clarifies why our results contrast with some of the previous literature.

2 Studying the value of life with recursive models

2.1 Temporal lotteries with uncertain lifetime

Recursive models were first suggested to model preferences over temporal lotteries in fixed or infinite horizon settings. To model choices under uncertain lifespans, we need to consider lives of unequal lengths. For the sake of simplicity we will assume that there is a finite upper bound T_{max} on how long a life can be.[3] We will use the dagger symbol † to describe the death state and define the set of temporal lotteries as follows:[4]

$$\begin{cases} D_t = \{\dagger\} & \text{for } t = T_{max}, \\ D_t = (\mathbb{R}_+ \times M(D_{t+1})) \cup \{\dagger\} & \text{for } t \in \{0, \dots, T_{max} - 1\}, \end{cases}$$

where for any set X the notation $M(X)$ denotes the set of simple lotteries with outcomes in X. As is usual for any element $x \in X$ we shall use the same notation x to denote the degenerate lottery in $M(X)$, which gives x with probability one.

An element of D_t different from † (thus reflecting a situation where the agent is alive in period t) will typically be denoted by a pair (c_t, m_t) where $c_t \in \mathbb{R}_+$ is the consumption in period t and $m_t \in M(D_{t+1})$ is a lottery over future states. For any $m_t \in M(D_{t+1})$, we define the survival probability $\pi(m_t)$, by $\pi(m_t) = 1 - Prob_{m_t}(\dagger)$, where $Prob_{m_t}(\dagger)$ is the occurrence probability of being dead in period $t+1$. When

[3]Since T_{max} can be arbitrarily large, there is no significant loss of generality. Moreover, the assumption of a finite lifespan is 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).

[4]The death state † is neither a consumption level nor a utility level. It has no other interpretation but describing a situation where the individual is dead.

$\pi(m_t) \neq 0$, we will also define $m_t^S \in M(D_{t+1} \setminus \{\dagger\})$ by:

$$m_t = \pi(m_t) m_t^S \oplus (1 - \pi_t(m_t))\dagger, \tag{1}$$

where $\oplus$ denotes the standard mixture operation over lotteries. The above equation thus simply states that m_t can be seen as a mixture, with weights $(1 - \pi(m_t))$ and $\pi(m_t)$, of a lottery that gives the death state for sure and a lottery m_t^S whose outcomes exclude immediate death. For an element $(c_t, m_t) \in D_t \setminus \{\dagger\}$, which describes the case of an agent alive in period t, the probability $\pi(m_t)$ is the probability of staying alive from period t to period $t+1$, and m_t^S is the lottery describing the distribution of outcomes in period $t+1$ conditional on being alive.

We (recursively) define a "multiplication by a scalar operation" over the spaces D_t as follows:

$$\begin{cases} \lambda\dagger = \dagger & \text{for all } \lambda \in \mathbb{R}_+, \\ \lambda(c_t, m) = (\lambda c_t, \lambda m) & \text{for all } \lambda \in \mathbb{R}_+ \text{ and } (c_t, m) \in D_t \setminus \{\dagger\}, \end{cases} \tag{2}$$

Thus, multiplying a temporal lottery by λ involves multiplying all (current and future) consumption levels by λ, with no impact on the survival probabilities.

As an example, the element

$$(c_1, m_1) = (c_1, \frac{1}{3}(c_2, \dagger) \oplus \frac{1}{3}(c_2', \dagger) \oplus \frac{1}{3}\dagger) \tag{3}$$

describes the case of an agent who is alive and consumes c_1 in period 1 and then dies with probability $\frac{1}{3}$ or survives with probability $\pi(m_1) = \frac{2}{3}$. If surviving in period 2, they consume either c_2 or c_2' with equal probabilities, and then die for sure at the end of period 2 (formally $m_1^S = \frac{1}{2}(c_2, \dagger) \oplus \frac{1}{2}(c_2', \dagger)$). In such a case, for any $\lambda \in \mathbb{R}_+$ one has:

$$\lambda(c_1, m_1) = (\lambda c_1, \frac{1}{3}(\lambda c_2, \dagger) \oplus \frac{1}{3}(\lambda c_2', \dagger) \oplus \frac{1}{3}\dagger).$$

A graphical representation of the temporal lottery (3) is provided in Figure 1.

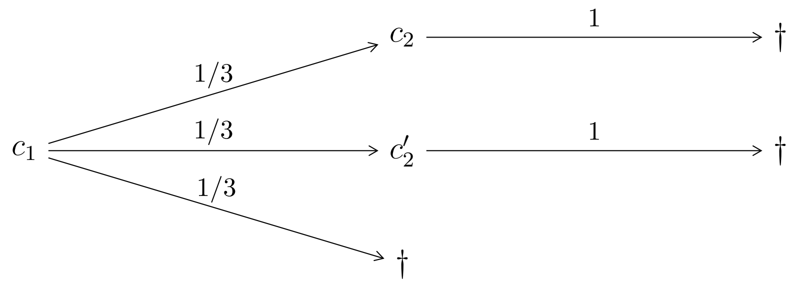

Figure 1: A graphical representation of the consumption lottery $(c_1, \frac{1}{3}(c_2, \dagger) \oplus \frac{1}{3}(c_2', \dagger) \oplus \frac{1}{3}\dagger)$.

2.2 Well-behaved recursive preferences

Most of the literature on the value of life follows the seminal paper of Yaari (1965) by assuming additively separable expected utility preferences, and combines this with an assumption of stationarity. The agent is then endowed with preferences over the sets D_t, for $t \in \{0, \ldots, T_{max}\}$, represented by utility functions $U_t : D_t \to Im(U_t) \subset \overline{\mathbb{R}}$ related by the following recursion:

$$\begin{cases} U_t(\dagger) = u_\dagger, \\ U_t(c_t, m_t) = u(c_t) + \beta E_m[U_{t+1}], \end{cases} \tag{4}$$

where $u_\dagger \in \overline{\mathbb{R}}$ is the utility associated with death, $u : \mathbb{R}_+ \to \overline{\mathbb{R}}$ is the period utility of consumption, $\beta > 0$ is the discount factor, and E_m the expectation operator with respect to the lottery m.[5] The model is sometimes introduced by providing an explicit solution to the recursion (4), where utility is obtained as the weighted sum of (future) expected per-period utilities, but for our purposes it is more convenient to base our discussion on the recursive formulation (4). Several features are worth emphasizing. First, recursivity precludes any ad hoc assumption of age-dependence in the formulation of agent preferences. In applications, age typically ends up playing a fundamental role since mortality (and hence the horizon) depends on age, but

[5] Note that there is no requirement that $u_\dagger$ belongs to the image of the instantaneous utility function u. In particular there does not need to exist a "death consumption equivalent" (that is a consumption level $c_\dagger$ such that $u(c_\dagger) = u_\dagger$).

this is different from assuming that tastes change with age. Second, uncertainty is resolved using a simple expectation operator: the model thus satisfies the basic axioms of the expected utility framework. Finally, although this property is not so well known, the model makes an assumption of correlation neutrality. Formally, for any consumption levels $c_t, c_{t+1}, c_{t+2}, c'_{t+1}, c'_{t+2} \in \mathbb{R}_+$:

$$U_t(c_t, \frac{1}{2}(c_{t+1}, (c_{t+2}, \dagger)) \oplus \frac{1}{2}(c'_{t+1}, (c'_{t+2}, \dagger)) = U_t(c_t, \frac{1}{2}(c_{t+1}, (c'_{t+2}, \dagger)) \oplus \frac{1}{2}(c'_{t+1}, (c_{t+2}, \dagger))$$

This equality (illustrated by the indifference shown in Figure 2) implies that the agent is indifferent to correlations between what happens in different periods. In particular, the agent is indifferent between a lottery paying off either two consecutive bad outcomes or two consecutive good outcomes and a lottery where good and bad outcomes are mixed.

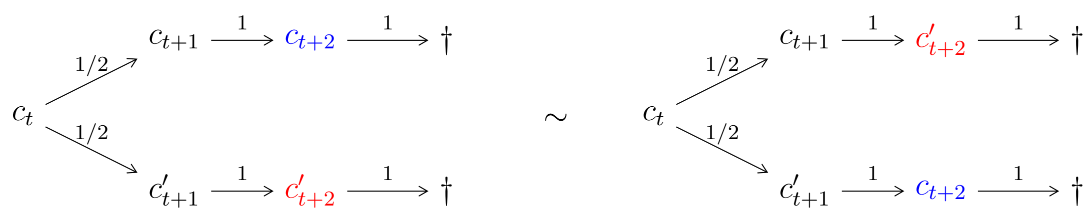

Figure 2: A graphical representation of correlation neutrality ($\sim$ is the indifference relationship).

Each of the assumptions embedded in (4) can be challenged. Among other things, it can be argued that stationarity or the axioms of expected utility are too restrictive. There is empirical and experimental evidence that contradicts these assumptions. Starmer (2000) reviews some of them, such as rejection of the betweenness axiom (a weak form of the independence axiom) found in Camerer and Ho (1994). Starmer (2000) also discusses the merits of other possible (static) models of choice under risk relaxing the linearity in probabilities, such as the rank-dependent model (Quiggin, 1982; Yaari, 1987) or the Kahneman and Tversky (1979) prospect theory. In a dynamic setting, another route is to relax stationarity, which allows for a wider class of intertemporal risk preferences. Nevertheless, when used in a setting where the question is how much effort a society should make to reduce mortality in the context of a pandemic, the assumptions of expected utility and of stationarity may seem normatively appealing. Stationarity avoids ad hoc age-dependent assumptions and the hurdles of defining a social choice objective for a population of time-inconsistent individuals (see Bernheim, 2009). The independent axiom endorsed by the expected utility theory can be justified by the wish to focus on consequentialist preferences (Hammond, 1988). More controversial, in our view, is the assumption of correlation neutrality, which constrains risk aversion (see Richard, 1975) and is contradicted by empirical evidence (Andersen et al., 2018 and Harrison et al., 2022). Since mortality is a risk, we consider it to be essential to study the role of risk aversion, but this cannot be done in the additive model.

Our goal in the current paper is to extend the framework (4), to gain flexibility while retaining key aspects of the additive model, such as preference stationarity and consequentialism. In practice, we use the dynamic choice theory of Kreps and Porteus (1978) and assume preferences over the sets D_t, for $t \in \{0, \ldots, T_{max}\}$, represented by utility functions $U_t : D_t \to Im(U_t) \subset \overline{\mathbb{R}}$ related through the following recursion:

$$\begin{cases} U_t(\dagger) = u_\dagger, \\ U_t(c_t, m_t) = u(c_t) + \beta\phi^{-1}(E_{m_t}[\phi(U_{t+1})]). \end{cases} \tag{5}$$

Compared to the recursion (4), the only new element is the function ϕ, which is assumed to be defined over a domain that includes the convex hull of $\cup_{t\in\{0,\ldots,T_{max}\}} Im(U_t)$, and to be strictly monotone (for ϕ^{-1} to be well-defined). In general, the recursion (5) admits no simple solution. However, since $D_{T_{max}} = \{\dagger\}$, for a given combination of $u_\dagger$, u, β and ϕ, there exists a unique sequence of functions U_t that fulfills (5). Starting from $t = T_{max}$ and applying the recursion (5) backwards gives a direct construction of the utilities U_t. This framework is similar to that used in Epstein and Zin (1989) and Weil (1989), except that it allows for the possibility of death and does not constrain the functions u and ϕ to take specific functional forms.

The functions u and ϕ will be assumed to be twice continuously differentiable.

The function u governs intertemporal substitutability of consumption and will be assumed to be concave. It will be said to be CRRA (for Constant Relative Risk Aversion) if $u(c) = \frac{c^{1-\sigma}}{1-\sigma} + u_l$, with $0 < \sigma \neq 1$, or $u(c) = \ln(c) + u_l$, for some constant $u_l \in \mathbb{R}$. The function ϕ governs risk aversion, with greater concavity reflecting greater risk aversion. As is well known, such recursive preferences may exhibit preference for the timing of resolution of uncertainty. Precisely, it follows from Kreps and Porteus (1978, Theorem 3) that preferences exhibit preference for early (resp. late) resolution if the function $x \mapsto \phi(u(c) + \beta\phi^{-1}(x))$ is convex (resp. concave). Recursive preferences also contrast with the standard additive framework for the possibility of exhibiting intertemporal correlation aversion. A formal definition of intertemporal correlation aversion can be found in Stanca (2021), who explains that a positive intertemporal correlation aversion is obtained whenever the function ϕ is concave.

Using survival probabilities and the mixture operation of equation (1), recursion (5) can be rewritten as $U_t(c_t, m_t) = u(c_t) + \beta\phi^{-1}(\pi_t(m_t)E_{m_t^S}[\phi(U_{t+1})] + (1 - \pi_t(m_t))\phi(u_\dagger))$. For greater legibility, we will simplify the notation in the remaining and simply write for any $(c_t, m_t) \in D_t \setminus \{\dagger\}$:

$$U_t(c_t, m_t) = u(c_t) + \beta\phi^{-1}\left(\pi_t E[\phi(U_{t+1})] + (1 - \pi_t)\phi(u_\dagger)\right), \tag{6}$$

where π_t implicitly stands for the probability $\pi(m_t)$ and $E[\cdot]$ for the expectation operator $E_{m_t^S}[\cdot]$.

Throughout the paper, we will assume that preferences fulfill two basic requirements, stipulating that greater consumption implies greater welfare and that at least some lives are considered to be worthwhile (i.e., better than death). Formally:

Definition 1 (Well-behaved preferences) *Recursive preferences with uncertain lifetime represented as in (5) are said to be well-behaved if:*

1. ***(Monotonicity)*** *The function* $(c_0, c_1) \in \mathbb{R}_+^2 \mapsto U_1(c_0, \frac{1}{2}(c_1, \dagger) \oplus \frac{1}{2}\dagger)$ *is strictly increasing.*

2. ***(Non-Systematically Suicidal)*** *For all* $t < T_{max}$*, there exists* $(c_t, m) \in$

$D_t \setminus \{\dagger\}$ *such that* $U_t(c_t, m) > u_\dagger$.

The assumptions "Monotonicity" and "Non-Systematically Suicidal", which we view as basic requirements, allow specific normalization to be made, simplifying the utility representation. Indeed:

Proposition 1 *Recursive preferences with uncertain lifetime are well-behaved if and only if they admit a utility representation* U_t *where:*

- *the period utility function* u *is strictly increasing, and not always negative (i.e., there exists* $c \in \mathbb{R}_+$ *such that* $u(c) > 0$*);*
- *we can set the normalization:* $u_\dagger = \phi(u_\dagger) = 0$.

The utility U_t *can then be defined by* $U_t(\dagger) = 0$ *and the following recursion:*[6]

$$U_t(c_t, m) = u(c_t) + \beta\phi^{-1}(\pi_t E[\phi(U_{t+1})]), \tag{7}$$
$$\text{with: } \phi(0) = 0 \text{ and } u' > 0.$$

For the remainder of the paper we will focus on well-behaved recursive preferences represented as in (7), and explore the additional restrictions that would be related to two additional assumptions, namely "Homotheticity" and "Ordinal Dominance".

2.3 Homotheticity and Ordinal Dominance

Preference homotheticity means that scaling all present and future consumption levels by the same (positive) factor does not change the ranking of lotteries. This is formalized in the following axiom:

Axiom 1 (Homotheticity) *For any* $t \geq 0$, $(c_t, m), (c'_t, m') \in D_t \setminus \{\dagger\}$ *and* $\lambda > 0$*:*

$$(U_t(c_t, m) \geq U_t(c'_t, m')) \Leftrightarrow (U_t(\lambda c_t, \lambda m) \geq U_t(\lambda c'_t, \lambda m')) .$$

[6]For the sake of conciseness, we do not specify definition domains in Proposition 1 (as well as in Propositions 2 and 3 below). The utility U_t is defined over D_t, while (c_t, m) is an element of $D_t \setminus \{\dagger\}$.

Homotheticity is a very popular assumption as it tends to simplify optimization problems. Indeed, under preference homotheticity, wealth has a basic scaling effect and can be easily removed from optimization problems. This is very convenient, especially in settings where wealth is impacted by some random factors (as asset returns, or random labor income) along the lifecycle.

Proposition 2 *Well-behaved recursive preferences with uncertain lifetime fulfill Homotheticity (Axiom 1) if and only if they admit a utility representation U_t fulfilling recursion (7) in which:*

- *either u is CRRA, not always negative ($u(c) > 0$ for some c), and ϕ is linear,*
- *or $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ for some $\sigma < 1$, and $\phi(x) = x^\rho$ for some $\rho > 0$. The parameter ρ governs risk aversion (with a larger ρ implying a lower degree of risk aversion), while $\frac{1}{\sigma}$ is the IES.*

Formally, utility U_t can then be defined by $U_t(\dagger) = 0$ and one of the following recursions:

- $U_t(c_t, m) = \frac{c_t^{1-\sigma}}{1-\sigma} + u_l + \beta \pi_t E[U_{t+1}]$ *with* $u_l > 0$*;*
- $U_t(c_t, m) = \frac{c_t^{1-\sigma}}{1-\sigma} + \beta \left(\pi_t E[U_{t+1}^\rho]\right)^{\frac{1}{\rho}}$ *with* $\sigma < 1$ *and* $\rho > 0$*.*

The first case corresponds to the standard additive specification, which lacks flexibility to study the role of risk aversion. The second case corresponds to Epstein-Zin-Weil preferences, with an IES, $\frac{1}{\sigma}$, above one and a coefficient of relative risk aversion, $1 - \rho(1 - \sigma)$, smaller than one. We explain in Appendix B why the above results contrast with the messages found in Hugonnier et al. (2013) and Córdoba and Ripoll (2017).

The other axiom we consider is Ordinal Dominance ("OD", hereafter).

Axiom 2 (Ordinal Dominance) *For all dates $0 \leq t < T_{max}$, consumption levels $c_t, c'_t \in \mathbb{R}_+$ and lotteries $m_1, m'_1, m_2, m'_2 \in M(D_{t+1})$, if:*

$$\begin{cases} U_t(c_t, m_1) \geq U_t(c'_t, m'_1), \\ U_t(c_t, m_2) \geq U_t(c'_t, m'_2), \end{cases}$$

then:

$$U_t(c_t, \frac{1}{2}m_1 \oplus \frac{1}{2}m_2) \geq U_t(c'_t, \frac{1}{2}m'_1 \oplus \frac{1}{2}m'_2).$$

OD is defined in similar fashion in Chew and Epstein (1990), or in Bommier et al. (2017) for temporal lotteries. It is adapted here to the context of uncertain lifetime. It states that if lottery (c_t, m_1) is preferred to (c'_t, m'_1) and (c_t, m_2) to (c'_t, m'_2), then the mixture of the two most preferred should also be preferred to the mixture of the two least preferred. OD is similar in spirit to a property of preference monotonicity with respect to first-order stochastic dominance. If one views taking decisions under uncertainty as playing a game against Nature, OD is akin to the elimination of dominated strategies. It thus appears to be a very natural assumption to model rational choice.

Proposition 3 *Well-behaved recursive preferences with uncertain lifetime fulfill OD (Axiom 2) if and only if they admit a utility representation U_t fulfilling recursion (7) in which $\phi(x) = \frac{1-e^{-kx}}{k}$ for some $k \neq 0$ or $\phi(x) = x$ and the function u is strictly increasing and not always negative ($u(c) > 0$ for some c). Utility U_t is then defined by $U_t(\dagger) = 0$ and the following recursion:*

$$U_t(c_t, m) = \begin{cases} u(c_t) - \frac{\beta}{k} \ln\left(\pi_t E\left[e^{-kU_{t+1}}\right] + 1 - \pi_t\right) & \text{for } k \neq 0, \\ u(c_t) + \beta\pi_t E[U_{t+1}] & \text{for } k = 0. \end{cases} \tag{8}$$

Specification (8) is nothing other than an adaptation of risk-sensitive preferences to a context of an uncertain lifetime. Indeed, Proposition 3 can be seen as an extension of the representation result of Bommier et al. (2017) to uncertain horizons but restricted to expected utility certainty equivalents, as in Kreps and Porteus (1978). From Bommier and LeGrand (2014), we also know that the preferences represented by recursion (8) verify a general notion of monotonicity with respect to first-order stochastic dominance that encompasses our definition of OD.

A consequence of Propositions 2 and 3 is that the only preferences fulfilling Homotheticity and OD are the standard additive preferences. This is formulated in the following corollary.

Corollary 1 *Well-behaved recursive preferences with uncertain lifetime fulfill Homotheticity (Axiom 1) and OD (Axiom 2) if and only if they can be represented as in (7) with a function u that is CRRA and not always negative, and $\phi(x) = x$.*

It follows from Corollary 1 that disentangling IES from risk aversion requires abandoning Homotheticity or OD. On the one hand, giving up Homotheticity and choosing OD leads to preferences that are overall well-behaved, even though they come with a slight disadvantage in terms of tractability. On the other hand, giving up OD and choosing Homotheticity leads to preferences that may end up opting for dominated choices. Moreover, IES is constrained to be above one, and risk aversion below one, which are restrictive assumptions that contradict empirical evidence. Finally, while the Homotheticity assumption is technically convenient, its empirical relevance has been regularly questioned. Indeed, it is well documented that rich and poor people do not behave identically (in relative terms) when making decisions relating to savings and financial investments, in direct contradiction to the Homotheticity assumption (see, for instance, Dynan et al., 2004 for saving behavior and Calvet and Sodini, 2014 for portfolio composition).

Overall, it seems to us that the overall balance is clearly in favor of OD and that the cost of maintaining Homotheticity is too high.

2.4 Risk-sensitive preferences

As shown in Proposition 3, the only well-behaved recursive preferences that satisfy Ordinal Dominance are risk-sensitive preferences, obtained when $\phi : x \mapsto \frac{1-e^{-kx}}{k}$, providing the recursion (8). Risk-sensitive preferences exhibit preference for early resolution of uncertainty when $k(1-\beta) > 0$ (either $k > 0$ and $\beta < 1$ or $k < 0$ and $\beta > 1$) and for late resolution of uncertainty when $k(1-\beta) < 0$. Indifference to the timing of resolution of uncertainty occurs when $k(1-\beta) = 0$ ($k = 0$ or $\beta = 1$). In such cases, risk-sensitive preferences actually yield expected utility preferences, the standard additive specification being obtained when $k = 0$ and the multiplicative specification of Bommier (2013) when $k \neq 0$ and $\beta = 1$. This multiplicative specification, which rules out pure time preferences ($\beta = 1$), was used in Bommier and Villeneuve (2012) and Bommier and LeGrand (2014) to study the value of life and the demand for annuities, respectively. Risk-sensitive preferences exhibit positive intertemporal correlation aversion when $k > 0$ and negative intertemporal correlation aversion when $k < 0$, regardless the value of β. The experimental evidence on intertemporal correlation aversion (see Andersen et al., 2018 and Harrison et al., 2022 for recent evidence using an online experiment in the US during the Covid-19 pandemic) would then support the assumption of a positive k.

Because they satisfy OD, risk-sensitive preferences allow for an intuitive interpretation of the impact of risk aversion. Intuitively, risk aversion can be understood in terms of how bad states are weighted relative to good states when making choices under uncertainty. As will be seen below, this is reflected in basic results regarding intertemporal consumption trade-offs and the value of mortality risk reduction. We discuss these aspects first through formal results, and then through a numerical example. In the remainder of the section, we assume that the only risk is mortality. As a result, the expectation symbol in (8) is no longer needed.[7]

2.4.1 Effective discount factor

The marginal rate of substitution between consumption in times $t+1$ and t is what governs the trade-off that an agent at time t makes between c_{t+1}, the consumption the agent will have in period $t+1$ if alive, and current consumption, c_t. With

[7] With the previous notation, an element of D_t can be written as $(c_t, \pi_t(c_{t+1}, m_{t+1}) \oplus (1-\pi_t)\dagger)$. For the sake of conciseness, we simply write U_t instead of $U_t(c_t, \pi_t(c_{t+1}, m_{t+1}) \oplus (1-\pi_t)\dagger)$.

risk-sensitive preferences this marginal rate of substitution is given by the following expression:

$$\frac{\frac{\partial U_t}{\partial c_{t+1}}}{\frac{\partial U_t}{\partial c_t}} = \frac{u'(c_{t+1})}{u'(c_t)} \beta \frac{\pi_t}{\pi_t + (1-\pi_t)e^{kU_{t+1}}},$$

As is standard, this marginal rate of substitution can be broken down into two factors: (i) a first one $\frac{u'(c_{t+1})}{u'(c_t)}$ that reflects the aversion for consumption fluctuation over time; and (ii) a second one, $\beta \frac{\pi_t}{\pi_t+(1-\pi_t)e^{kU_{t+1}}}$, called *effective* discount factor, obtained when controlling for time variation in consumption (that is when imposing $c_t = c_{t+1}$).[8] The effective discount factor simplifies to $\beta \pi_t$ in the standard additive model ($k = 0$) but is otherwise endogenous. The effective discount factor has the following characteristics.

Proposition 4 *The effective discount factor* $\beta \frac{\pi_t}{\pi_t+(1-\pi_t)e^{kU_{t+1}}}$ *is such that:*

1. *it increases with* π_t*;*
2. *it decreases with* k *(i.e. higher risk aversion makes people become more impatient) if* $u(c_t) > 0$ *for all* t*;*
3. *it decreases with continuation utility if* $k > 0$ *(and thus with future consumption, and future survival probabilities if* $u(c_t) > 0$ *for all* t*);*

The proof can be found in Appendix A.4. For an intuitive understanding of Proposition 4, it is useful to recall that giving up some consumption in period t for an increase in consumption in period $t+1$ involves risk, since consumption in period $t+1$ only occurs if the agent survives to period $t+1$. In the absence of bequest motives, consumption in period $t+1$ only makes sense if the agent survives to period $t+1$. This explains why the higher the survival probability the greater the marginal rate of substitution $\frac{\frac{\partial U_t}{\partial c_{t+1}}}{\frac{\partial U_t}{\partial c_t}}$, and the greater the effective discount factor. Points 2 and 3 in Proposition 4 simply reflect the general fact that the willingness to take a given

[8]We use the terminology "effective discount factor" to distinguish it from the time preference parameter β, which is also the "discount factor" that would be obtained in a context of certainty (that is in absence of mortality).

risk depends on risk aversion and the other risks at play. In our setting, the agent faces a mortality risk where death is a bad outcome if $u(c_t) > 0$ for all t (see the following Proposition 5 and its discussion). Reducing consumption in period t to increase consumption in $t+1$ involves placing a bet that pays nothing in the event of death but pays out in the event of survival. In other words, postponing consumption is a risk taking behavior that is positively correlated with the background risk generated by mortality (provided that $u(c_t) > 0$ for all t, implying that death is a bad outcome). Thus, the willingness to take this risk decreases with risk aversion (point 2 of Proposition 4). If the agent is "intertemporal correlation averse" ($k > 0$), then the willingness to postpone consumption decreases with the magnitude of the background mortality risk, and thus with the continuation utility which quantifies how much is lost in the event of death.

Note that, while the first point, i.e. the positive relationship between the effective discount factor and the survival probability is already present in the additive case, the second and third points only become apparent when considering specifications that are flexible enough to study the role of risk aversion and that allow for positive intertemporal correlation aversion.

2.4.2 Value of mortality risk reduction

The value of Mortality Risk Reduction (MRR) is defined as the marginal rate of substitution between survival and consumption.[9] It quantifies how much consumption in period t the agent is willing to relinquish to diminish their mortality risk from period t to period $t+1$, keeping everything else unchanged. Formally, we have:

$$MRR_t = \frac{\frac{\partial U_t}{\partial \pi_t}}{\frac{\partial U_t}{\partial c_t}}. \tag{9}$$

[9] In the absence of annuity and other assets whose returns depend on survival probability, the value of mortality risk reduction can identically be defined as the marginal rate of substitution between survival and wealth.

Another standard denomination for this marginal rate of substitution is the Value of a Statistical Life (VSL). Due to misunderstandings related to the "Value of Life" denomination, there is an ongoing debate around the appropriate terminology.[10] Independently of the denomination, this concept is extensively used in cost-benefit analysis for public policy design.

In the risk-sensitive framework, the value of mortality risk reduction (9) has the following expression if $k \neq 0$:

$$MRR_t = \frac{\beta}{k} \frac{1}{u'(c_t)} \frac{1 - e^{-kU_{t+1}}}{\pi_t e^{-kU_{t+1}} + 1 - \pi_t}. \tag{10}$$

This expression reduces to $MRR_t = \frac{\beta}{u'(c_t)} U_{t+1}$ in the additive model ($k = 0$), as is found in papers working with additively separable preferences (Hall et al., 2020 for a recent example). The following proposition summarizes the properties of the value of mortality risk reduction in the risk-sensitive model.

Proposition 5 *The value of mortality risk reduction in equation (10) is such that:*

1. *it is positive at all dates if $u(c_t) > 0$ for all t;*

2. *it increases with continuation utility (thus with future consumption, and with future survival probabilities if $u(c_t) > 0$ for all t);*

3. *it increases with π_t if $k > 0$;*

4. *the relationship with k is ambiguous in general (even when $k > 0$ and $u(c_t) > 0$ for all t).*

The proof of the proposition can be found in Appendix A.5. The first point of Proposition 5 states that having positive instantaneous utilities at all dates is a sufficient condition for having a positive value of mortality risk reduction. The positivity condition on instantaneous utilities should be interpreted as being alive

[10] See `https://www.epa.gov/environmental-economics/mortality-risk-valuation` for a discussion.

and consuming c_t is preferred to being dead, and results from our normalization $u_\dagger = 0$.[11] A positive value of mortality risk reduction means that the agent is willing to pay to reduce their mortality probability. A negative value would mean that the agent is willing to pay to increase their mortality probability.

The second point of the proposition states that the value of mortality risk reduction increases with the continuation utility (in the event of survival). The higher the "payoff" in case of survival, the more the agent is willing to pay for enjoying it. This effect, which is largely intuitive, is already present in the additive model. It implies that the value of mortality risk reduction increases with future consumption and future survival probabilities (if consumption is high enough to make life better than death).

In order to better understand the third and fourth points, which are absent with the additive specification, one has to realize that making investments today to lower tomorrow's mortality is not a plain risk reduction. Obviously, such investments lower the risk of dying (which is a risk reduction), but they also imply taking the risk of dying tomorrow after having made sacrifices today (which is an increase in risk). This latter aspect is more a concern when the risk of dying tomorrow is high (i.e., when π_t is low). When the agent is "intertemporal correlation averse" ($k > 0$), we then naturally find that the lower π_t, the lower the value of mortality risk reduction, which corresponds to the second point of Proposition 5. Also, the fact that making investments to lower mortality involves reducing some risks, while generating others explains why risk aversion has an ambiguous impact (last point of Proposition 5). This finding and the explanation we provide is fully 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 e.g., Dionne and Eeckhoudt, 1985, or Jullien et al., 1999).

[11] Without this normalization, the condition would have been that instantaneous utilities need to be greater than $(1-\beta)u_\dagger$.

2.4.3 A numerical illustration

We complement our theoretical results with simulations showing how risk aversion impacts lifecycle consumption profiles and the value of mortality risk reduction when using risk-sensitive preferences. For these simulations, we consider the problem of an agent who faces an exogenous mortality pattern, with an age-specific mortality rate corresponding to that of the total US population in 2018, as reported in the Human Mortality Database (HMD). The agent is endowed with a CRRA utility function of the form $u(c) = \frac{c^{1-\sigma}}{1-\sigma} + u_l$, where u_l is set such that the value of mortality risk reduction at age 40 is \$ 10 million in the additive model.[12] The interest rate is assumed to be $r = 4\%$. Preference parameters are set to standard values: $\beta = 0.98$, and $\sigma = 2.0$. Lifetime wealth is normalized such that the yearly consumption at age 40 for the additive agent is equal to \$ 45,000. For risk-sensitive preferences, we use the same calibration as in the additive model and furthermore set k to the value of Bommier et al. (2020) calibrated using annuity data. In order to highlight the role of risk aversion, we contrast the results obtained with a positive k (referred to as the "risk-sensitive model") to those obtained for $k = 0$ (referred to as the "additive model").

Figure 3 reports the consumption paths for both specifications. Both the additive and risk-sensitive models generate plausible hump-shaped consumption paths. The predictions diverge between the two models, due to the role of risk aversion, which makes the agent more impatient – as stated in the second point of Proposition 4. This is illustrated in Figure 3 by the consumption levels that are larger at younger ages (and lower at older ages) in the risk-sensitive model than in the additive model.

We plot in Figure 4 the age profiles for the value of mortality risk reduction in the additive and risk-sensitive models. Recall that the additive model is calibrated so that the value of mortality risk reduction is \$ 10 million at age 40 and that the difference between the two profiles is due to an increase in risk aversion. As noted in

[12]This is in the range of estimates for the US (in 2021). Viscusi (2021), for example, suggests a value of mortality risk reduction of about \$ 11 million.

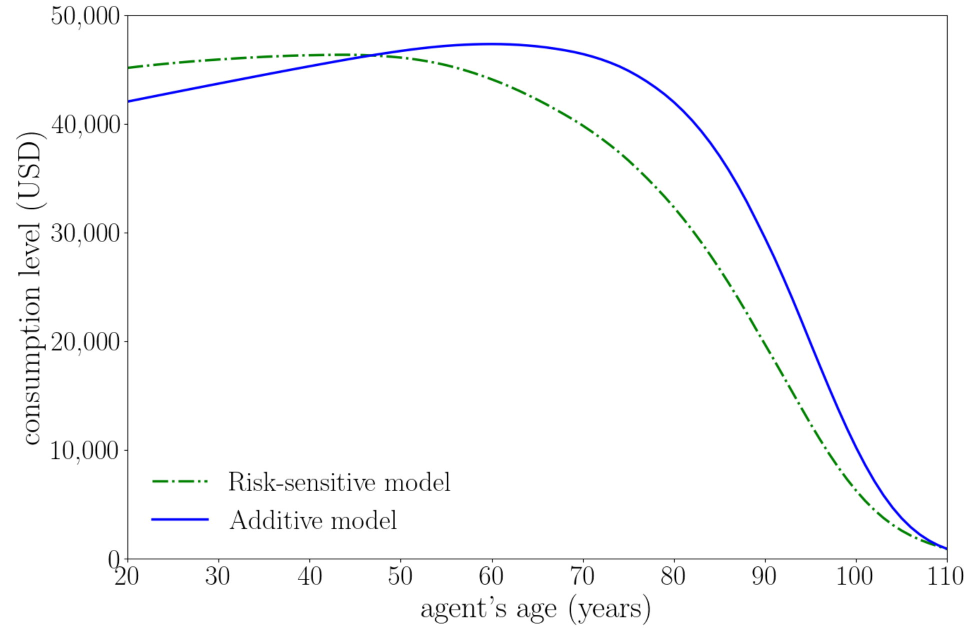

Figure 3: Consumption profiles implied by the additive and risk-sensitive models.

Proposition 5, the role of risk aversion is ambiguous.

In both models, the value of mortality risk reduction tends to decrease with age, reflecting the fact that older people have a shorter life expectancy, and thus a smaller stake in their mortality (the value of mortality risk reduction also depends on the marginal utility of current consumption, which explains why it increases slightly with age at the beginning of the life cycle with the additive model). It is important to note that the value of mortality risk reduction decreases more rapidly with age in the risk-sensitive specification than in the additive specification. The underlying reason is that risk aversion tends to magnify the willingness to avoid particularly dramatic outcomes (such as death at a young age), as compared to the willingness to avoid less adverse outcomes (such as death at an advanced age). As we shall see, this point is crucial to understanding how risk aversion can alter policy recommendations in the context of epidemic management.

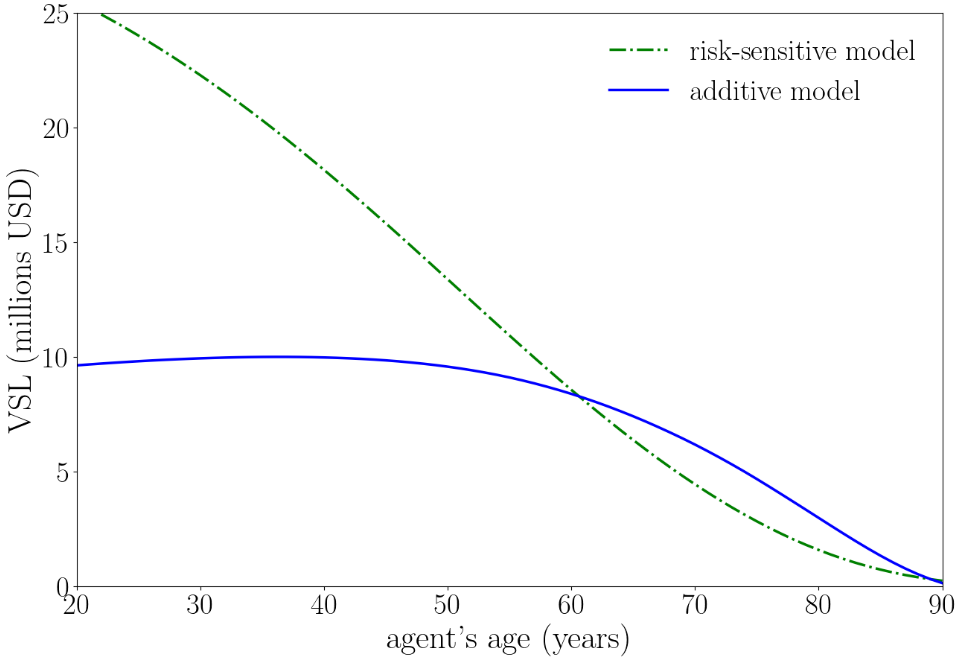

Figure 4: Profiles of value of mortality risk reduction implied by the additive and risk-sensitive models.

3 Application to epidemics

We now illustrate how risk-sensitive preferences provide new insights into 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 Covid-19 (Section 3.2) and 1918 influenza (Section 3.3) pandemics. One reason to focus on these two real-world cases is that they feature highly contrasted age-specific mortality rates. We start by presenting the setup.

3.1 General case

We consider a population of size normalized to 1, initially containing agents of different ages t in proportions $(\omega_t)_t$, with $\omega_t \in [0,1]$ and $\sum_t \omega_t = 1$. Agents of age t are endowed with risk-sensitive preferences represented by utility function U_t defined in recursion (8). 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_t to $\pi_t - \delta_t$ for one year. A benevolent planner seeks to determine which share α of current consumption agents are willing to relinquish in exchange for being rid of the excess mortality risk. We further simplify the framework by assuming that consumption, c, is constant throughout ages and that the agents' discount factor β is 1. Let $\lambda = 1 - \alpha$ and denote by $U_t(\lambda, \delta)$ the current utility of an agent of age t whose current consumption is λc (instead of c) and their next-period survival $\pi_t - \delta$. One has:

$$U_t(\lambda, \delta) = u(\lambda c) - \frac{1}{k} \log((\pi_t - \delta) e^{-kU_{t+1}(1,0)} + 1 - \pi_t + \delta), \tag{11}$$

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

$$\begin{aligned} W(\lambda, (\delta_t)_t) &= \sum_t \omega_t U_t(\lambda, \delta_t) \\ &= u(\lambda c) - \frac{1}{k} \sum_t \omega_t \log((\pi_t - \delta_t) e^{-kU_{t+1}(1,0)} + 1 - \pi_t + \delta_t). \end{aligned} \tag{12}$$

The planner seeks to determine how much of current consumption c can be reduced so as to offset in terms of welfare the extra mortality risk, which corresponds formally to the equivalence $W(1, (\delta_t)_t) = W(\lambda, 0)$, or using (11) and (12) to:

$$u(c) - u(\lambda c) = \frac{1}{k} \sum_t \omega_t \log\left(1 + \delta_t \frac{1 - e^{-kU_{t+1}(1,0)}}{\pi_t e^{-kU_{t+1}(1,0)} + 1 - \pi_t}\right). \tag{13}$$

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

$$\alpha = 1 - \lambda \approx \frac{1}{cu'(c)} \sum_t \delta_t \omega_t \frac{1}{k} \frac{1 - e^{-kU_{t+1}(1,0)}}{\pi_t e^{-kU_{t+1}(1,0)} + 1 - \pi_t}, \tag{14}$$

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

To interpret further equation (14), we can conduct a first-order Taylor expansion in k of expression (13) for α. We obtain the following approximation for small k:

$$\alpha \approx \underbrace{v\sum_t \delta_t\omega_t\mathbb{E}_{t+1}[\widetilde{T}]}_{\text{additive term}}+v\frac{ku(c)}{2}\left(\sum_t \delta_t\omega_t\left(\underbrace{(2\pi_t-1)\,\mathbb{E}_{t+1}[\widetilde{T}]^2}_{\text{gain proportional to }\mathbb{E}_{t+1}[\widetilde{T}]^2} - \underbrace{\mathbb{V}_{t+1}[\widetilde{T}]}_{\text{loss due to risk}}\right)\right), \quad (15)$$

where $v = \frac{u(c)}{cu'(c)}$ is, as in Hall et al. (2020), the value of a year of life relative to consumption. The factor $\mathbb{E}_{t+1}[\widetilde{T}]$ is life expectancy at age $t+1$ and $\mathbb{V}_{t+1}[\widetilde{T}]$ the variance of lifespans at age $t+1$. The expression (15) consists of three terms. The first term ("additive term") is the same as in Hall et al. (2020), 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 $\pi_t > 0.5$, which is the case for all ages except for very advanced ages. For instance, in the HMD data we use, it is only at ages greater than 105 that $\pi_t < 0.5$. 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). The last term is proportional to the variance of lifespans at age $t+1$: the more uncertain the lifespan, the less the agent is willing to pay to avoid 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 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.[13]

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 is CRRA, with a constant IES, $\frac{1}{\sigma}$ set to $\frac{1}{2}$ (thus with $\sigma = 2$). We also use their calibration for the consumption level c, set to $\$\,45,000$,

[13] The fact that the effect is positive only at young ages when the survival probability is large enough is in line with our result in Proposition 5 and its illustration in Section 2.4.3. As already mentioned, this is consistent with the literature on optimal prevention (Dionne and Eeckhoudt, 1985, or Jullien et al., 1999).

and for the value of life at age 40, set to \$ 10.4m and used to determine u_l. The population shares $(\omega_t)_t$ are also those of the US total population in 2018, as reported by the US Census Bureau.[14]

For the Covid-19 age-specific mortality profile $(\delta_t)_t$, 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:

$$\log_{10} IFR_t = -3.27 + 0.0524 \times t, \tag{16}$$

where IFR_t is the IFR at age t and $\log_{10}$ is the log in base 10. This implies that the IFR increases by 12.82% with 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: $\delta_t = 65\% \times IFR_t$. 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).

Finally, the survival probabilities $(\pi_t)_t$ are chosen to be those of the US total population in 2018, as reported in the Human Mortality Database.[15] 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.[16]

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

[15] 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.

[16] This calibration mostly focuses on replicating some average observed values at the population level – and also on making our work directly comparable to Hall et al. (2020). Other routes would

We report in Table 1 the share α of consumption agents are willing to relinquish in order to be rid of the excess mortality risk of Covid-19. We do so for both the risk-sensitive and additive models. For each, we report the values α obtained from the linear approximation (equation (14)) and from the exact formula (equation (13)).[17] When using the additive model, agents are willing to give up 50.4% of their current consumption to avoid 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 relinquish, by 5 percentage points approximately (for the exact formula), thus a relative reduction of about 16 percent.

Computational method	Additive model	Risk-sensitive model
Linear approximation (14)	50.4	39.6
Exact formula (13)	33.5	28.3

Table 1: 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 giving greater weight to more adverse consequences. To illustrate this, we plot in Figure 5 the parameters $(\alpha_t)_t$ of equation (14) as a function of age in both the additive

have been possible, and in particular using joint experimental estimates of risk aversion, discounting and IES (see Harrison et al., 2022 for a recent paper based on several waves during the Covid period).

[17]The corresponding expressions for the additive model are simply obtained by taking the limit $k \to 0$ in (13) and (14).

and the risk-sensitive models. Each α_t represents the drop in consumption that a population only made of agents with age t is willing to accept to be rid of an extra mortality risk of 0.1%.[18] It can be seen that the $(\alpha_t)_t$ 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

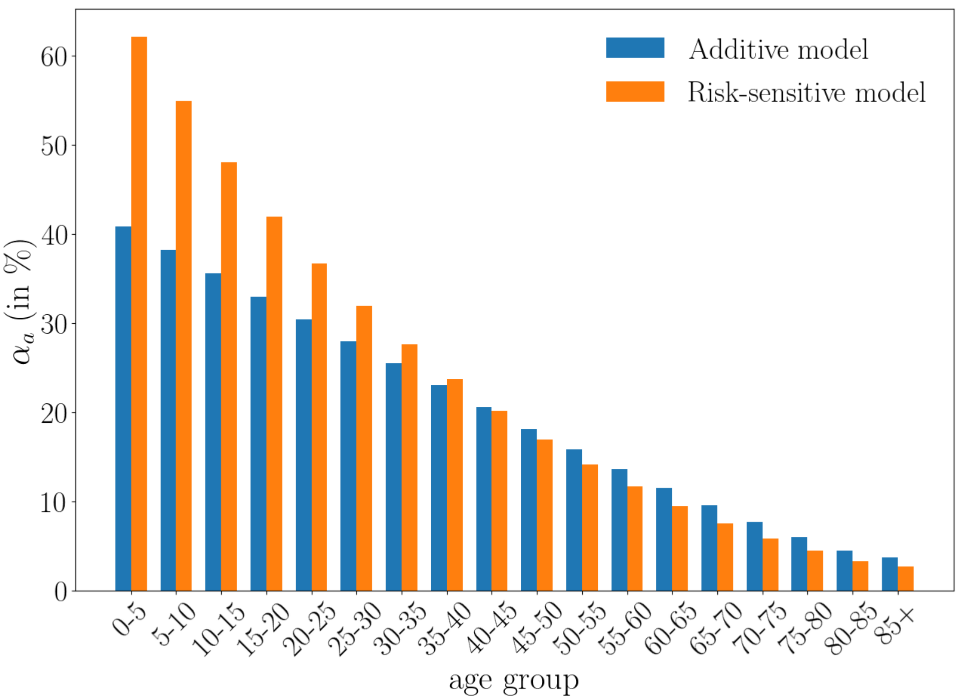

Figure 5: Profiles of the (α_t) parameters as a function of age.

than dying old. 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. This makes young agents willing to pay more for mortality risk reduction with risk-sensitive preferences than with additive ones. The opposite holds for older agents. As the calibration is performed on willingness-to-pay for mortality risk reduction at age 40, the differences

[18]Note that for a mortality risk higher than 0.1%, the value of α_t for some ages t could be higher than 100%. This simply reflects the limitation of the linear approximation of equation (14) that is only valid for low mortality risks. As can be seen in Table 1, the linear approximation is not very precise in the case of Covid-19 either for the additive or the risk-sensitive models.

between models for age groups 35-40 and 40-45 remain extremely modest. Since Covid-19 mainly affects older people, this explains why overall, the acceptable 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 for Covid-19 with those obtained when considering the 1918 influenza epidemic, a disease that heavily affected young people. 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 6.[19] 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 specific to the 1918 influenza, since regular influenza epidemics generally exhibit a U-shape mortality profile.

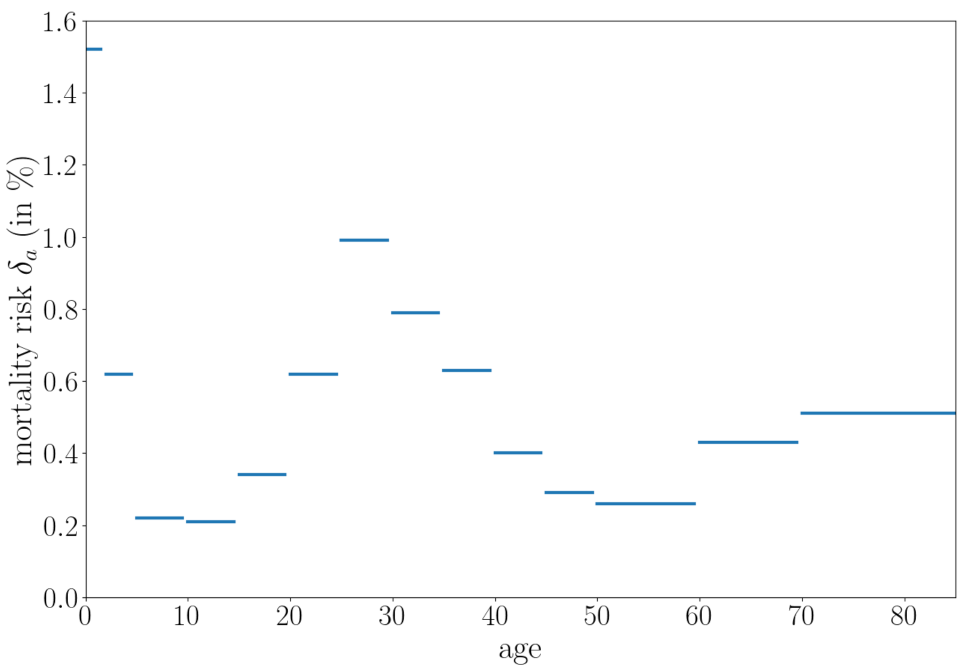

Figure 6: Age-profile of mortality risk (δ_t) for 1918 influenza.

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

[19] 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%.

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

Table 2: The share α of consumption (in %) to relinquish to be rid of 1918 influenza mortality risk, computed using the exact formula (14).

to be rid of the 1918 influenza mortality risk. This corresponds to the quantity α computed using the exact formula (14). We do not here report the values computed with the linear approximation. Compared to Table 1, we merely change the profile (δ_t), keeping the rest of the calibration unchanged. The value in Table 2 should therefore be interpreted as willingness-to-pay to be rid of 1918 influenza mortality risk in the US of 2018 (and not of 1918). To facilitate comparison with Covid-19, we consider two mortality patterns: the actual 1918 influenza one, corresponding to the mortality data of Figure 6, and a rescaled version that would yield the same number of deaths 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 influenza pandemic.[20]

Unsurprisingly, the values of α for both models are higher for 1918 influenza than for Covid-19, even when controlling for average mortality. This is due to the fact that younger people suffered comparatively much more from 1918 influenza that from Covid-19. Furthermore, compared to Table 1, 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.

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

4 Conclusion

Finding the appropriate policy in the case of a health crisis typically amounts to making trade-offs between mortality and consumption (or wealth). The recent Covid-19 crisis has shown how sensitive such issues can be. Risk-reduction measures, such as lock-downs, have been strongly criticized both for being excessive and for not being severe enough. According to the usual micro-economic approach, 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 negative 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 recursive homothetic specifications are unappealing in the context of mortality risk, constraining in particular the IES to be above one and risk aversion to be below one. Second, we point to the risk-sensitive preferences initially introduced by Hansen and Sargent (1995). They are shown to be the only recursive preferences that fulfill Ordinal Dominance and disentangle risk aversion and IES, providing 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. Moreover, the risk-sensitive framework does not constrain the IES (which does not even need to be constant).

In practice, we find that 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. However, the present research could open out new perspectives for other topical issues, such as the opioid or gun violence epidemics.

Appendix

A Proofs

We assume that u is twice continuously differentiable on $\mathbb{R}_+^* = \mathbb{R}_+ \setminus \{0\}$.

A.1 Proof of Proposition 1

Two preliminary lemmas. In the remainder, we call NSS the "Non-Systematically Suicidal" property of Definition 1. First, we prove the following lemma:

Lemma 1 *Recursive preferences represented with recursion (5) fulfill NSS iff they admit a utility representation U_t where the period utility function u verifies $u(c) > (1-\beta)u_\dagger$ for some c.*

Proof. If there exists c such that $u(c) > (1-\beta)u_\dagger$ then for all $t < T_{max}$ one has $U_t(c, \dagger) = u(c) + \beta u_\dagger > u_\dagger$ and preferences are NSS. Conversely, assume that $u(c) \leq (1-\beta)u_\dagger$ for all c. Then for all elements $(c, m) \in D_{T_{max}-1}$ one has $U_{T_{max}-1}(c, m) = u(c) + \beta u_\dagger \leq u_\dagger$. This in turn implies, by induction, that for all $t < T_{max}$ and all $(c_t, m) \in D_t$, one has $U(c_t, m) \leq u_\dagger$, in contradiction with NSS. ■

Second we explore the role of monotonicity.

Lemma 2 *Recursive preferences with recursion (5) are monotone (in the sense stated in Proposition 1) iff they admit a utility representation U_t where:*

- *the period utility function u is strictly increasing;*
- *$u_\dagger$ and $\phi(u_\dagger)$ are finite.*

Proof.

Consider $c_1' > c_1 \geq 0$ and $c_0 > 0$. Monotonicity implies:

$$U_1(c_0, \frac{1}{2}(c_1', \dagger) \oplus \frac{1}{2}\dagger) > U_1(c_0, \frac{1}{2}(c_1, \dagger) \oplus \frac{1}{2}\dagger), \tag{17}$$

or, equivalently:

$$\phi^{-1}\left(\frac{1}{2}\phi(u(c_1') + \beta u_\dagger) + \frac{1}{2}\phi(u_\dagger)\right) > \phi^{-1}\left(\frac{1}{2}\phi(u(c_1) + \beta u_\dagger) + \frac{1}{2}\phi(u_\dagger)\right). \tag{18}$$

Let us assume that $\phi(u_\dagger) = \pm\infty$. In that case, $\phi^{-1}\left(\frac{1}{2}\phi(u(c_1) + \beta u_\dagger) + \frac{1}{2}\phi(u_\dagger)\right)$ and $\phi^{-1}\left(\frac{1}{2}\phi(u(c_1') + \beta u_\dagger) + \frac{1}{2}\phi(u_\dagger)\right)$ are equal and inequality (18) does not hold. We deduce that $\phi(u_\dagger)$ is finite. Similarly, $u_\dagger$ is finite (otherwise $u(c_1)+\beta u_\dagger$ and $u(c_1')+\beta u_\dagger$ are equal and inequality (18) would not hold).

We further consider $c_0' > c_0$. Monotonicity implies $U_1(c_0', \frac{1}{2}(c_1, \dagger) \oplus \frac{1}{2}\dagger) > U_1(c_0, \frac{1}{2}(c_1, \dagger) \oplus \frac{1}{2}\dagger)$. Since $u_\dagger$ and $\phi(u_\dagger)$ are finite, it is then equivalent to: $u(c_0') > u(c_0)$, which implies that the function u is increasing on $\mathbb{R}_+$. By choosing $c_0 = 0$, we also obtain that $u(0) < \infty$ (one may however have $u(0) = -\infty$). ■

Proof of Proposition 1. We assume that NSS and Monotonicity hold. We have:

$$\begin{cases} U_t(\dagger) = u_\dagger \\ U_t(c_t, m) = u(c_t) + \beta\phi^{-1}(\pi_t(m)E_{m_S}[\phi(U_{t+1})] + (1 - \pi_t(m))\phi(u_\dagger)) \end{cases} \tag{19}$$

Since from Lemma 2, $u_\dagger$ is finite, we can posit for all t and for all $m \in D_t$, $\tilde{U}_t(m) = U_t(m) - u_\dagger$, as well as $\tilde{\phi}(x) = \phi(x + u_\dagger)$, which means that $\tilde{\phi}^{-1}(y) = \phi^{-1}(y) - u_\dagger$. Recursion (19) then becomes: $\tilde{U}_t(\dagger) = 0$ and $\tilde{U}_t(c_t, m) = u(c_t) + \beta\tilde{\phi}^{-1}(\pi_t(m)E_{m_S}[\tilde{\phi}(\tilde{U}_{t+1})] +$

$(1-\pi_t(m))\tilde{\phi}(0))$. Since $\tilde{U}_t$ represents the same preferences as U_t, we can set $u_\dagger = 0$ without loss of generality.

Furthermore, since $\phi(u_\dagger)$ is finite (Lemma 2), so is $\tilde{\phi}(0)$. We define: $\hat{\phi}(x) = \tilde{\phi}(x) - \tilde{\phi}(0)$ – such that $\hat{\phi}^{-1}(y) = \tilde{\phi}^{-1}(x + \tilde{\phi}(0))$. Thus, recursion (19) finally becomes recursion (7) (after dropping the tilde and hat decorations). Furthermore, observe that ϕ is defined up to a positive multiplicative factor: if preferences are represented by ϕ, they will also be represented by $\mu\phi$ for any $\mu > 0$.

A.2 Proof of Proposition 2

A.2.1 First step

Let $c_0, c_1 > 0$. Since $U_t(c_0, (c_1, \dagger)) = u(c_0) + \beta u(c_1)$, the marginal rate of substitution (MRS) between c_0 and c_1, denoted by MRS_{c_0,c_1}, is:

$$MRS_{c_0,c_1} = \beta \frac{u'(c_1)}{u'(c_0)}. \tag{20}$$

Homotheticity implies that $MRS_{c_0,c_1} = MRS_{\lambda c_0, \lambda c_1}$.[21] From (20), $\frac{u'(\lambda c_1)}{u'(\lambda c_0)}$ is independent of λ. Computing the log-derivative in λ for $\lambda = 1$ yields: $c_1 \frac{u''(c_1)}{u'(c_1)} = c_0 \frac{u''(c_0)}{u'(c_0)}$, which gives $c\frac{u''(c)}{u'(c)}$ constant. This implies that there exist $K, \sigma > 0$ and u_l, such that $u(c) = K\frac{c^{1-\sigma}}{1-\sigma} + Ku_l$ if $\sigma \neq 1$ or $u(c) = K\ln(c) + Ku_l$. In our applications, we will be able to normalize K to 1 without loss of generality (see Section A.2.2 below) and the function u is thus CRRA, in the sense provided in Section 2.2.

A.2.2 Second step

Let $c_0, c_1 > 0$ and $\pi \in (0,1)$. We have: $U_t(c_0, \pi(c_1, \dagger) \oplus (1-\pi)(c_2, \dagger)) = u(c_0) + \beta\phi^{-1}\left(\pi\phi(u(c_1)) + (1-\pi)\phi(u(c_2))\right)$. The MRS between c_0 and c_1, MRS_{c_0,c_1}, is then:

$$MRS_{c_0,c_1} = \beta\pi \frac{u'(c_1)}{u'(c_0)} \frac{\phi'(u(c_1))}{\phi'\left(\phi^{-1}\left(\pi\phi(u(c_1)) + (1-\pi)\phi(u(c_2))\right)\right)}.$$

[21] Let c_0' and c_1' such that $u(c_0) + \beta u(c_1) = u(c_0') + \beta u(c_1')$. We obtain $u'(c_0)\frac{\partial c_0}{\partial c_1} + \beta u'(c_1) = 0$. Using Homotheticity, we obtain that for any λ, $u(\lambda c_0) + \beta u(\lambda c_1) = u(\lambda c_0') + \beta u(\lambda c_1')$, which yields $u'(\lambda c_0)\frac{\partial(\lambda c_0)}{\partial(\lambda c_1)} + \beta u'(\lambda c_1) = 0 = u'(\lambda c_0)\frac{\partial c_0}{\partial c_1} + \beta u'(\lambda c_1)$.

Homotheticity implies $MRS_{c_0,c_1} = MRS_{\lambda c_0,\lambda c_1}$. Since we already proved that $\frac{u'(c_1)}{u'(c_0)} = \frac{u'(\lambda c_1)}{u'(\lambda c_0)}$, this implies that $\frac{\phi'(u(\lambda c_1))}{\phi'(\phi^{-1}(\pi\phi(u(\lambda c_1))+(1-\pi)\pi\phi(u(c_2))))}$ is independent of λ. Similarly, considering the MRS between c_0 and c_2, we obtain that $\frac{\phi'(u(\lambda c_2))}{\phi'(\phi^{-1}(\pi\phi(u(\lambda c_1))+(1-\pi)\pi\phi(u(c_2))))}$ is independent of λ. Taking the ratio, we deduce that $\frac{\phi'(u(\lambda c_1))}{\phi'(u(\lambda c_2))}$ is independent of λ. Taking the log derivative yields:

$$c_1 u'(\lambda c_1)\frac{\phi''(u(\lambda c_1))}{\phi'(u(\lambda c_1))} = c_2 u'(\lambda c_2)\frac{\phi''(u(\lambda c_2))}{\phi'(u(\lambda c_2))}, \tag{21}$$

or equivalently that $cu'(c)\frac{\phi''(u(c))}{\phi'(u(c))}$ is constant. There are two cases.

1. There exists c, such that $\phi''(u(c)) = 0$. In that case, from (21), we deduce that ϕ'' is null everywhere. Since $\phi(0) = 0$ and we impose $\phi'(0) = 1$, we obtain $\phi(x) = x$. Using the results of Section A.2.1, we deduce that (7) becomes: $U_t(c_t, m) = K\frac{c^{1-\sigma}}{1-\sigma} + Ku_l + \beta\pi_t U_{t+1}$ if $\sigma \neq 1$ or $U_t = K\ln(c) + Ku_l + \beta\pi_t U_{t+1}$. We can set $K = 1$ without loss of generality since the preferences represented by U_t and $K^{-1}U_t$ are the same. We obtain:

$$U_t(c_t, m) = \begin{cases} \frac{c^{1-\sigma}}{1-\sigma} + u_l + \beta\pi_t U_{t+1}, \\ \ln(c) + u_l + \beta\pi_t U_{t+1}, \end{cases}$$

   which is the standard Yaari (1965) model with a CRRA utility function.

2. For all c, $\phi''(u(c)) \neq 0$. There are two cases depending on whether the IES of u differs from 1 or not.

   (a) $u(c) = K\frac{c^{1-\sigma}}{1-\sigma} + u_l$ for some $K > 0$, $\sigma \neq 1$, and u_0. Using (21), we obtain that there exists ρ, such that for all c: $(y - u_l)\frac{\phi''(y)}{\phi'(y)} = \rho - 1$, where $y = K\frac{c^{1-\sigma}}{1-\sigma} + u_l$. This yields with $\phi(0) = 0$:

$$\phi(y) = \phi_0\left((y - u_l)^\rho - (-u_l)^\rho\right). \tag{22}$$

   We now consider $U_t(c_0, \pi(c_1, \dagger) \oplus (1-\pi)\dagger) = u(c_0) + \beta\phi^{-1}\left(\pi\phi(u(c_1))\right)$ and the MRS between c_0 and c_1. Using similar steps as above, we obtain that there exist a neighborhood B and a constant C, such that for all $y \in B$:

$$\frac{\phi(y)\phi''(y)}{\phi'(y)^2} = C. \tag{23}$$

Using (22), we deduce for all $y \in B$, $1 - \frac{(-u_l)^\rho}{(y-u_l)^\rho} = \frac{C\kappa}{\rho-1}$, which imposes $u_l = 0$. Using (22) and $\phi'(1) = \rho$ for the scaling normalization, we obtain $\phi(y) = y^\rho$. The condition $\phi(0) = 0$ implies $\rho > 0$. Recursion (7) becomes: $U_t(c_t, m) = K\frac{c_t^{1-\sigma}}{1-\sigma} + \beta(\pi_t E[U_{t+1}^\rho])^{\frac{1}{\rho}}$, where we can set $K = 1$ without loss of generality since the preferences represented by U_t and $K^{-1}U_t$ are the same. We finally obtain: $U_t(c_t, m) = \frac{c_t^{1-\sigma}}{1-\sigma} + \beta(\pi_t E[U_{t+1}^\rho])^{\frac{1}{\rho}}$ $(\rho > 0)$.

(b) $u(c) = K\ln(c) + u_l$ for some $K > 0$ and u_l. We use similar steps as in the previous case. Equation (21) implies that there exists k, such that $\frac{\phi''(y)}{\phi'(y)} = -k$ for all y, or after integration with $\phi(0) = 0$ and $\phi'(0) = 1$ (that we can impose without loss of generality): $\phi(y) = \frac{1-e^{-ky}}{k}$. Furthermore, equality (23) still holds and yields: $1 - e^{ky} = C$ for all y, which imposes $k = 0$. We fall back on the additive case.

A.3 Proof of Proposition 3

Let $c_0, c_1 > 0$ and $\pi \in (0, 1)$. Monotonicity implies that u is strictly increasing and that u^{-1} exists and is unique. For any $c > 0$, we define the function η_{c_0,c_1} as:

$$\eta_{c_0,c_1}(c) = u^{-1}\left(\phi^{-1}\left(\pi^{-1}\phi\left(\beta^{-1}(u(c_0) - u(c)) + \phi^{-1}\left(\pi\phi(u(c_1))\right)\right)\right)\right), \tag{24}$$

We can observe that by definition:

$$\eta_{c_0,c_1}(c_0) = u^{-1}\left(u(c_1)\right) = c_1, \tag{25}$$

$$u(c_0) + \beta\phi^{-1}\left(\pi\phi(u(c_1))\right) = u(c) + \beta\phi^{-1}\left(\pi\phi(u(\eta_{c_0,c_1}(c)))\right) \tag{26}$$

From (24) and (25), since u and ϕ are continuously differentiable, there exists a neighborhood $\widetilde{B}_{c_0}$ of c_0, where η_{c_0,c_1} exists and is continuously differentiable.

Let $c_1' > 0$ and $c_1' \neq c_1$. Similarly, we can find another neighborhood $\widehat{B}_{c_0}$ of c_0

such that the function $\eta_{c_0,c_1'}(c)$ defined as in (24) is continuously differentiable on $\widehat{B}_{c_0}$ and verifies for $c \in \widehat{B}_{c_0}$:

$$u(c_0) + \beta\phi^{-1}\left(\pi\phi(u(c_1'))\right) = u(c) + \beta\phi^{-1}\left(\pi\phi(u(\eta_{c_0,c_1'}(c)))\right). \tag{27}$$

We define $B_{c_0} = \widehat{B}_{c_0} \cap \widetilde{B}_{c_0}$, which is a non-empty open set. OD implies that equalities (26) and (27) lead to, for all $c \in B_{c_0}$:

$$\begin{aligned} u(c_0) + \beta\phi^{-1}\left(\frac{\pi}{2}\left(\phi(u(c_1)) + \phi(u(c_1'))\right)\right) = \\ u(c) + \beta\phi^{-1}\left(\frac{\pi}{2}\left(\phi(u(\eta_{c_0,c_1}(c))) + \phi(u(\eta_{c_0,c_1'}(c)))\right)\right). \end{aligned} \tag{28}$$

We compute the derivatives of equation (26) and its counterpart for $\eta_{c_0,c_1'}$, as well as of (28). After substitution, we obtain for $c \in B_0$, using $u'(c) > 0$ and $(\phi^{-1})'$:

$$\begin{aligned} \frac{1}{(\phi^{-1})'\left(\frac{\pi}{2}\left(\phi(u(\eta_{c_0,c_1}(c))) + \phi(u(\eta_{c_0,c_1'}(c)))\right)\right)} = \\ \frac{1}{2}\frac{1}{(\phi^{-1})'\left(\pi\phi(u(\eta_{c_0,c_1}(c)))\right)} + \frac{1}{2}\frac{1}{(\phi^{-1})'\left(\pi\phi(u(\eta_{c_0,c_1'}(c)))\right)}. \end{aligned}$$

We deduce that $\frac{1}{(\phi^{-1})'}$ is affine locally (see Aczél, 1966). Note that by changing π and c_0 we can cover the whole definition space of $\frac{1}{(\phi^{-1})'}$ and by continuity guarantee that $\frac{1}{(\phi^{-1})'}$ is actually affine globally. We deduce that there exist k and y_0, such that $\frac{1}{(\phi^{-1})'(y)} = -k(y - y_0)$. After integration, we find that there exists ϕ_0, such that: $\phi^{-1}(y) = -\frac{1}{k}\ln(\frac{y-y_0}{\phi_0})$, which implies $\phi(x) = \phi_0 e^{-kx} + y_0$. We further normalize ϕ using $\phi(0) = 0$ and $\phi'(0) = 1$ and obtain: $\phi(x) = \frac{1}{k}(1 - e^{-kx})$.

A.4 Proof of Proposition 4

As a preliminary, we prove the following Lemma.

Lemma 3 *If $u(c_t) > 0$ for all $t \in \{0, \ldots, T_{max} - 1\}$, we have $U_t > 0$ for all $t \in \{0, \ldots, T_{max} - 1\}$.*

The proof goes by backward induction. For $t = T_{max} - 1$, we have $U_{T_{max}-1} =$

$u(c_{T_{max}-1}) > 0$ by assumption. Let $t \in \{0, \ldots, T_{max}-2\}$ and assume that the result holds for $t+1$. We have by recursion: $U_t = u(c_t) - \frac{\beta}{k}\log(\pi_t e^{-kU_{t+1}} + 1 - \pi_t)$. Since $x \mapsto -\frac{\beta}{k}\log(\pi_t e^{-kx} + 1 - \pi_t)$ is increasing and $U_{t+1} \geq 0$ by induction hypothesis, we deduce that $U_t \geq u(c_t) > 0$. This concludes the induction and the proof of Lemma 3.

The proof of the first point of Proposition 4 is straightforward as U_{t+1} is independent of π_t. Denoting by $\Delta_t = \beta\frac{\pi_t}{\pi_t+(1-\pi_t)e^{kU_{t+1}}}$ the effective discount factor, we have $\frac{\partial \Delta_t}{\partial \pi_t} = \beta\frac{e^{kU_{t+1}}}{(\pi_t+(1-\pi_t)e^{kU_{t+1}})^2} > 0$.

For the second point, observe that the sign of $\frac{\partial \Delta_t}{\partial k}$ is the same as the one of $\frac{\partial(-kU_{t+1})}{\partial k}$.[22] Using the recursion (8) and defining $W_t = -kU_t$, we deduce:

$$W_{T_{max}-1} = -ku(c_{T_{max}-1}), \tag{29}$$

$$W_t = -ku(c_t) + \beta\log(\pi_t e^{W_{t+1}} + 1 - \pi_t) \text{ for } t < T_{max}-1. \tag{30}$$

We show by backward induction that $\frac{\partial W_t}{\partial k} < 0$ for all t. For $t = T_{max}-1$, we have from (29): $\frac{\partial W_{T_{max}-1}}{\partial k} = -u(c_{T_{max}-1}) < 0$, because we assume $u(c_t) > 0$ for all t. Let $t \in \{0, \ldots, T_{max}-2\}$ and assume that $\frac{\partial W_{t+1}}{\partial k} < 0$. Using (30), we get: $\frac{\partial W_t}{\partial k} = -u(c_t) + \frac{\beta\pi_t\frac{\partial W_{t+1}}{\partial k}}{\pi_t e^{W_{t+1}}+1-\pi_t}$, which implies $\frac{\partial W_t}{\partial k} < 0$ because of the induction hypothesis and the assumption that $u(c_t) > 0$ for all t. This concludes the induction and we have $\frac{\partial W_t}{\partial k} < 0$ for all t.

For the third point, we directly have $\frac{\partial \Delta_t}{\partial U_{t+1}} = \frac{-\beta k(1-\pi_t)e^{kU_{t+1}}}{(\pi_t+(1-\pi_t)e^{kU_{t+1}})^2} < 0$ if $k > 0$. Let us prove that U_{t+1} is increasing with future consumption and future survival probabilities. Let $t \in \{0, \ldots, T_{max}-1\}$ and $1 \leq s \leq T_{max}-t-1$. We first show that $\frac{\partial U_{t+1}}{\partial c_{t+s}} > 0$. Start with observing that because the utility defined in (8) is independent of the past, we have $\frac{\partial U_{t+s+1}}{\partial c_{t+s}} = 0$. We now prove that $\frac{\partial U_{t+s'}}{\partial c_{t+s}} > 0$ for all $s' \in \{1, \ldots, s\}$ by backward induction on s'. For $s' = s$, we have using $\frac{\partial U_{t+s+1}}{\partial c_{t+s}} = 0$: $\frac{\partial U_{t+s}}{\partial c_{t+s}} = u'(c_{t+s}) > 0$. Let $s' \in \{1, \ldots, s-1\}$ and assume that $\frac{\partial U_{t+s'+1}}{\partial c_{t+s}} > 0$. We have

[22] Though uncertain, the horizon in our setup is finite. So, the differentiability properties of intertemporal utility function (such as U_t) directly comes from the differentiability properties of the period utility function.

using (8):

$$\frac{\partial U_{t+s'}}{\partial c_{t+s}} = \frac{\beta e^{-kU_{t+s'+1}}}{\pi_{t+s} e^{-kU_{t+s'+1}} + 1 - \pi_{t+s}} \frac{\partial U_{t+s'+1}}{\partial c_{t+s}} > 0, \tag{31}$$

as a consequence of the induction hypothesis. This concludes the induction and we have $\frac{\partial U_{t+s'}}{\partial c_{t+s}} > 0$ for all $s' \in \{1, \ldots, s\}$. In particular, for $s' = 1$, $\frac{\partial U_{t+1}}{\partial c_{t+s}} > 0$ where $t \in \{0, \ldots, T_{max} - 1\}$ and $1 \leq s \leq T_{max} - t - 1$ are arbitrary.

The proof is similar for probabilities. We fix $t \in \{0, \ldots, T_{max} - 1\}$ and $1 \leq s \leq T_{max} - t - 1$. Since we have $\frac{\partial U_{t+s+1}}{\partial \pi_{t+s}} = 0$, we will prove that $\frac{\partial U_{t+s'}}{\partial \pi_{t+s}} > 0$ for all $s' \in \{1, \ldots, s\}$ by backward induction on s'. For $s' = s$: $\frac{\partial U_{t+s}}{\partial \pi_{t+s}} = \frac{\beta}{k} \frac{1 - e^{-kU_{t+s+1}}}{\pi_t e^{-kU_{t+s+1}} + 1 - \pi_t} > 0$. Let $s' \in \{1, \ldots, s-1\}$ and assume that $\frac{\partial U_{t+s'+1}}{\partial c_{t+s}} > 0$. We have, by induction hypothesis: $\frac{\partial U_{t+s'}}{\partial \pi_{t+s}} = \frac{\beta e^{-kU_{t+s'+1}}}{\pi_{t+s} e^{-kU_{t+s'+1}} + 1 - \pi_{t+S}} \frac{\partial U_{t+s'+1}}{\partial \pi_{t+s}} > 0$. This concludes the induction.

A.5 Proof of Proposition 5

The first point is implied by Lemma 3 and the fact that $U \mapsto \frac{\beta}{k} \frac{1}{u'(c_t)} \frac{1 - e^{-kU}}{\pi_t e^{-kU} + 1 - \pi_t}$ is increasing (independently of the sign of k). For the second point, we have from (10): $\frac{\partial MRR_t}{\partial U_{t+1}} = \beta \frac{1}{u'(c_t)} \frac{e^{-kU_{t+1}}}{(1 - \pi_t(1 - e^{-kU_{t+1}}))^2} > 0$, since $k > 0$ and $u'(c) > 0$ for all c. We conclude the third point using the same arguments as in Appendix A.4.

For the third point, using the expression of MRR_t in (10), we have: $\frac{\partial MRR_t}{\partial \pi_t} = \frac{\beta}{k} \frac{1}{u'(c_t)} \frac{(1 - e^{-kU_{t+1}})^2}{(1 - \pi_t(1 - e^{-kU_{t+1}}))^2}$, whose sign is the sign of k. This proves the second point.

We now turn to the last point. Let us assume $k > 0$ and $MRR_t > 0$. From (10):

$$\frac{1}{MRR_t} \frac{\partial MRR_t}{\partial k} = -\frac{1}{k} + \frac{e^{-kU_{t+1}}}{(1 - \pi_t(1 - e^{-kU_{t+1}}))(1 - e^{-kU_{t+1}})} \frac{\partial (kU_{t+1})}{\partial k}.$$

The first term $(-1/k)$ is negative while the second term is positive (from $\frac{\partial W_t}{\partial k} < 0$ proved in Appendix A.4). The overall sign is ambiguous.

B Issues when working with EZW preferences

In this section, we discuss issues related to working with EZW preferences in the context of an uncertain lifetime. To simplify the discussion, we assume that death is the only risk and that agents have no bequest motive.

Denoting by $v_\dagger$ the utility level associated with death, the utility V_t representing EZW preferences in the context of an uncertain lifetime is defined through the following recursion:

$$\begin{cases} V_t(\dagger) = v_\dagger, \\ V_t(c_t, m_t) = \left((1-\beta)c_t^{1-\sigma} + \beta\left(\pi_t V_{t+1}^{1-\gamma} + (1-\pi_t)v_\dagger^{1-\gamma}\right)^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{1-\sigma}}. \end{cases} \tag{32}$$

The parameters γ and σ are positive numbers reflecting risk aversion and the inverse of the IES, respectively.

This expression, which is also the one used in Córdoba and Ripoll (2017) (equation 2 in their paper), looks different from that of equation (6). However, this is actually only a question of utility normalization. For example, when $\sigma \neq 1$ and $\gamma \neq 1$, one can then set $U_t = \frac{V_t^{1-\sigma}}{(1-\beta)(1-\sigma)}$ and $u_\dagger = \frac{v_\dagger^{1-\sigma}}{(1-\beta)(1-\sigma)}$, representing the same preferences as V_t (remember that utility representation is defined up to an increasing transformation), and obtain a utility representation fulfilling the recursion (6). A similar transformation exists in the other cases where γ or σ is equal to 1.

Changing utility normalization is of course harmless and any representation may be used to discuss preference properties. In particular, it directly follows from Lemmas 1 and 2, in Appendix A.1, that when $v_\dagger > 0$, the recursion (32) represents well-behaved preferences in the sense of Definition 1 (i.e., preferences that are monotone and NSS). Such specifications are for instance used in Pashchenko and Porapakkarm (2022). They are not homothetic (unless $\sigma = \gamma$, providing the additive model) and do not fulfill OD (unless $\sigma = \gamma$ or $\sigma = 1$).[23]

B.1 The homothetic formulation

Most papers that have used the EZW specification in the value of life literature focus, at least for the main part of their analysis, on the homothetic specification

[23] When the IES is equal to 1 ($\sigma = 1$), the non-homothetic EZW specification is also a risk-sensitive one, as is well-known from Tallarini (2000).

obtained by choosing $v_\dagger$ such that $v_\dagger^{1-\gamma} = 0$. If $\gamma > 1$, this requires setting $v_\dagger = \infty$. Death then provides a higher utility than any consumption level, leading to models in which the value of life is systematically negative. This explains why the value of life literature focuses on the case $\gamma < 1$. This is consistent with the restriction we obtain in Proposition 2. However, there is some divergence in the ability to consider the case of an IES below 1 (i.e., $\sigma > 1$). While this is ruled out by Proposition 2, several papers, including Córdoba and Ripoll (2017), Córdoba et al. (2023), Hugonnier et al. (2013, 2020, 2022) claim that having $\sigma > 1$ is unproblematic. In fact, Córdoba and Ripoll (2017) argue that a key feature of the EZW specification is its ability to assume a homothetic specification while assuming a low IES. They devote Section 2.3 of their paper ("Homotheticity with low intertemporal substitution") to this point. This seems to be an essential feature since most of the applied literature tends to focus on the case where the IES is less than one.

In the following, we explain that working with a homothetic specification where $\gamma < 1$ (to have a positive value of life) and $\sigma > 1$ (to have a low IES) raises three main problems.

1. As shown in Section B.2, the existence of a non-zero solution requires an upper bound be set on mortality rates. This upper bound in never mentioned in Córdoba and Ripoll (2017). It is mentioned in Córdoba et al. (2023) in a footnote (footnote 28). It is introduced in Hugonnier et al. (2013, 2020, 2022) as a transversality condition.

2. As discussed in Section B.2, the required upper bound on mortality rates is very restrictive: for standard calibration it would imply that life expectancy remains unrealistically high during the lifecycle. The plausibility of such a bound is never discussed in any of the papers mentioned above.

3. We finally explain in Section B.3 that even in fictive settings where the upper bound on mortality rates is respected, the model yields predictions that are at odds with basic economic intuition. We also clarify the mechanisms behind

these counterintuitive predictions.

B.2 Homothetic EZW specifications with $\gamma < 1$ and $\sigma > 1$

The homothetic EZW specification is obtained when $v_\dagger^{1-\gamma} = 0$ in which case the recursion (32) reduces to:

$$V_t = \left[(1-\beta)c_t^{1-\sigma} + \beta\pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}. \tag{33}$$

Solution to the homothetic EZW recursion. A major issue with the specification (32) is that when the survival probabilities π_t are small (which is typically the case at old age), the effective discount factor $\beta\pi_t^{\frac{1-\sigma}{1-\gamma}}$ becomes larger than 1 (since $\frac{1-\sigma}{1-\gamma} < 0$). The recursion (33) then only admits constant solutions, which are of little use in practice. Formally, we can state the following result.

Proposition 6 *Consider the utility function defined by the recursion (33) with $\gamma < 1 < \sigma$ and assume that consumption is bounded from above ($c_t \leq \bar{c}$ for some $\bar{c} > 0$).*

1. *If there is a maximal lifespan (i.e., there exists T_{max} such that for all $t \geq T_{max}$, $\pi_t = 0$), the only solution to (33) is $V_t = 0$ for all t.*

2. *When death is never certain and survival probability decreases with age, there are three 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;*

- $\lim_{t\to\infty} \pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$*: the recursion admits multiple solutions: one being $V_t = 0$ for all t, and the other solution being given by:*

$$V_t = \left[\beta^{-t}\left(\Pi_{j=0}^{t-1}\pi_j\right)^{\frac{\sigma-1}{1-\gamma}} K + (1-\beta)\sum_{s=t}^{\infty}\beta^{s-t}\left(\Pi_{j=t}^{s-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}} c_s^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \tag{34}$$

for some constant $K \geq 0$;

- $\lim_{t\to\infty} \pi_t = \beta^{\frac{1-\gamma}{\sigma-1}}$*: the recursion admits* $V_t = 0$ *for all* t *as a solution but we cannot conclude as to the existence of other solutions (34).*

Proposition 6, which is proved below, states that when there is maximal lifespan, or when survival rates become low at large ages, the only solution to the recursive equation (33) is zero utility, that is $V_t = 0$ for all t.[24] The recursive equation (33) 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 $\frac{1}{1-\beta^{\frac{1-\gamma}{\sigma-1}}}$, no matter their age.

The restrictions required to have a well-defined model are not mentioned in Córdoba and Ripoll (2017), where only $\pi_t > 0$ is explicitly assumed (while one needs $\pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$). In Hugonnier et al. (2013, 2020, 2022), the corresponding restriction is introduced as a transversality condition with no discussion regarding its implications for mortality rates or life expectancy.[25]

How demanding is the condition $\lim_{t\to\infty} \pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$ **?** For standard calibrations, the condition $\lim_{t\to\infty} \pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$ imposes a lower bound on life expectancy that is incompatible with demographic evidence. For example, consider the case where $\gamma = 0.5$, $\sigma = 2$ (values that are close to those considered in Córdoba and Ripoll, 2017). For $\beta = 0.97$, then one should have $\pi_s > 0.984$, implying a life expectancy of at least 66 years – reached at age 11 in the US life tables (2021 HMD data for the general population). Playing around with the numbers does not change the pattern much. For example assuming a low β, equal to 0.9, implies a minimum life expectancy of 20 years – reached at age 63 in the US life tables. Consequently, the

[24] The convergence in the case $\lim_{t\to\infty} \pi_t = \beta^{\frac{1-\gamma}{\sigma-1}}$ may be decided upon the precise dynamics of $(\pi_t)_t$. But, in any case, this is anecdotal.

[25] In Hugonnier et al. (2013), this is condition (18) p. 675 commented as follows: "The restriction imposed by (18) is a transversality condition that limits the health sensitivity of the agent's income rate to guarantee that the corresponding present value is finite." In Hugonnier et al. (2020), this is mentioned in Theorem 1 (p. 141): "Assume that conditions (Equation 17) of Appendix S1 hold true". The condition is not made explicit nor commented further in the main text. In Hugonnier et al. (2022), this is mentioned in condition (9a) of Theorem 1 (p. 1102). This condition is commented (p. 1102) as "Conditions (9a)–(9c) encompass transversality restrictions for a finite shadow value of human capital."

utility function (34), solution to (33), could be used in a model of perpetual youth with a high life expectancy, but not in a lifecycle model that accounts for actual mortality profiles.

Finally, note that the issue that appears when $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$ cannot be solved by taking the limit $v_\dagger \to 0$ in the recursion (32). Indeed, one can show that when considering the limit $v_\dagger \to 0$ (with $v_\dagger > 0$), while $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$, then the value of the mortality risk reduction implied by (32) goes to ∞ which make the model unfit to study the value of life (see Bommier et al., 2021 for further detail).

A remark on model refinements. Various model refinements were considered: health index in Córdoba and Ripoll (2017), time-varying discount factors in Córdoba et al. (2023), or endogenous mortality and perpetual consumption growth in Hugonnier et al. (2013, 2020, 2022). These refinements do not fundamentally change the existence of a convergence condition similar to the one of Proposition 6. For instance, if one abandons stationarity (and model parsimony) with the introduction of time-varying discount factors $(\beta_t)_t$, the condition $\lim_{t\to\infty} \pi_t/\beta_t^{\frac{1-\gamma}{\sigma-1}} > 1$ still needs to hold to have non-constant solutions. If π_t becomes small at old ages, then β_t also needs to be small, which raises the question of the empirical counterparts. Similarly, allowing for endogenous mortality still requires a condition of the form: $\lim_{t\to\infty} \pi_t^*/\beta_t^{\frac{1-\gamma}{\sigma-1}} > 1$ where π_t^* is the predicted endogenous survival probability. This means that health investment should be sufficiently efficient to maintain a stringent lower bound on the remaining life-expectancy, which is counterfactual. Finally, if we allow for perpetual consumption growth, the condition for obtaining non-constant solutions becomes $\lim_{t\to\infty} \beta\pi_t^{\frac{1-\sigma}{1-\gamma}} \left(\frac{c_{t+1}}{c_t}\right)^{1-\sigma} < 1$. Even though the issue is partly mitigated, it is not solved for standard calibrations. For example, with $\beta = 0.95$ (which is already a low value) and a consumption growth rate of 5% per year (which would already be difficult to rationalize with a high rate of impatience corresponding to the low β), convergence would require maintaining a life expectancy that never falls below 20 years (reached at age 63 in US life tables, as mentioned above).

Proof of Proposition 6. First observe that if $V_t = 0$ for some t, then $V_\tau = 0$ for all $\tau \leq t$, independently of survival probabilities. This derives from successive applications of the recursion (33) and shows that $V_t = 0$ for all t is always a solution of (33).

Second, consider the case $\lim_{t\to\infty} \pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$. The ratio test then implies that the solution (34), that solves (33), is well-defined.

Third, consider the case $\lim_{t\to\infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$. Assume, for example, that there exists $t_0 \geq 0$ such that $V_{t_0} > 0$. We then must have from our first point: $V_t > 0$ for all $t > t_0$. Let us rewrite equation (33) as: $V_{t+1}^{1-\sigma} = \beta^{-1}\pi_t^{\frac{\sigma-1}{1-\gamma}}(V_t^{1-\sigma} - (1-\beta)c_t^{1-\sigma})$, which by iteration yields for all $t > t_0$:

$$V_t^{1-\sigma} = \beta^{-t}\left(\Pi_{j=t_0}^{t-1}\pi_j\right)^{\frac{\sigma-1}{1-\gamma}}\left(V_{t_0}^{1-\sigma} - (1-\beta)\sum_{s=t_0}^{t-1}\beta^s\left(\Pi_{j=t_0}^{s-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}}c_s^{1-\sigma}\right). \tag{35}$$

Let us assume $t \to \infty$, while holding t_0 constant. Assuming that consumption is bounded from above, the right-hand side of (35) becomes negative, since $\sum_{s=0}^{t-1}\beta^s\left(\Pi_{j=0}^{s-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}}c_s^{1-\sigma} \geq \bar{c}^{1-\sigma}\sum_{s=0}^{t-1}\beta^s\left(\Pi_{j=0}^{s-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}}$, which diverges (as we are considering the case $\lim_{t\to\infty}\pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$), while the left-hand side has to be positive. We thus obtain a contradiction, proving that there cannot exist a t_0 for which $V_{t_0} > 0$.

B.3 Implications of EZW homothetic specifications

In fictive settings, like the perpetual youth model with low mortality, the condition $\lim_{t\to\infty}\pi_t > \beta^{\frac{1-\gamma}{\sigma-1}}$ could hold even when $\gamma < 1 < \sigma$. In such a case, solution (34) can be used for the analysis. As we explain below the implications of such models are however at odds with basic economic intuition when considering the effective discount factor (and thus saving behaviors) and the willingness to pay for mortality risk reduction.

Effective discount factor. The marginal rate of substitution between consumption in period $t+1$ and consumption in period t implied by (34) is:

$$\frac{\partial V_t}{\partial c_{t+1}} \Big/ \frac{\partial V_t}{\partial c_t} = \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} c_{t+1}^{-\sigma} c_t^{\sigma}, \tag{36}$$

The effective discount factor is thus $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}}$ which raises a number of concerns. This effective discount factor does not increase but *decreases* with the survival probability π_t, which means that the less likely survival is at a given age, the more the agent wants to save resources for consumption at that age. As Córdoba et al. (2023, p.25) note "survival and future consumption are substitutes". However, common sense would suggest that complementarity should hold, since in a model with no bequest motive, there is no utility for consumption after death. Note also that this effective discount factor $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}}$ is greater than β, which means that in a consumption-saving model, a mortal agent (i.e. with $\pi_t < 1$) saves more than what they would do if they were immortal (i.e., with $\pi_t = 1$). This is in fact a dominated strategy, since saving as if $\pi_t = 1$ (i.e., saving less than what the model predicts for $\pi_t < 1$) would provide higher utility in the case where the agent survives as well as in the case where the agent dies (remember that savings are useless in the event of death).

Relatedly, when applied to realistic mortality patterns, an effective discount factor of $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}}$ implies counterfactual consumption-savings behavior. Consider the standard lifecycle consumption-saving problem of Section 2.4.3, with the same calibration parameters ($r = 4\%$ $\beta = 0.97$, $\sigma = 2.0$) except that we now consider an EZW specification where $\gamma = 0.5$. Lifetime wealth is normalized to \$ 1,000,000, such that the consumption at age 20 for the additive agent is close to \$ 40,000. This normalization is of little importance, because preferences are homothetic. From the Karush-Kuhn-Tucker theorem, we know that any interior solution should fulfill the first-order condition $\frac{\partial V_t}{\partial c_{t+1}} \Big/ \frac{\partial V_t}{\partial c_t} = 1/(1+r)$. Using (36) consumption should grow at the rate $(\beta(1+r)\pi_t^{\frac{1-\sigma}{1-\gamma}})^{\frac{1}{\sigma}}$, thus very rapidly at old age. As an illustration, we plot in Figure 7 the consumption implied by this growth rate, assuming that consumption after age 110 is artificially set to 0.[26] For comparison, we also plot the consumption path implied by the additive model. The profile obtained for the homothetic EZW specification exhibits a consumption level that remains extremely low until age 100, but that sky-rockets after that. It is very different from the one obtained with the additive specification. Note that the y-axis is truncated at \$ 50,000 for enhanced legibility but under the homothetic EZW model, consumption actually reaches \$ 98,000 at age 100 and \$ 14,000,000 at age 110.[27]

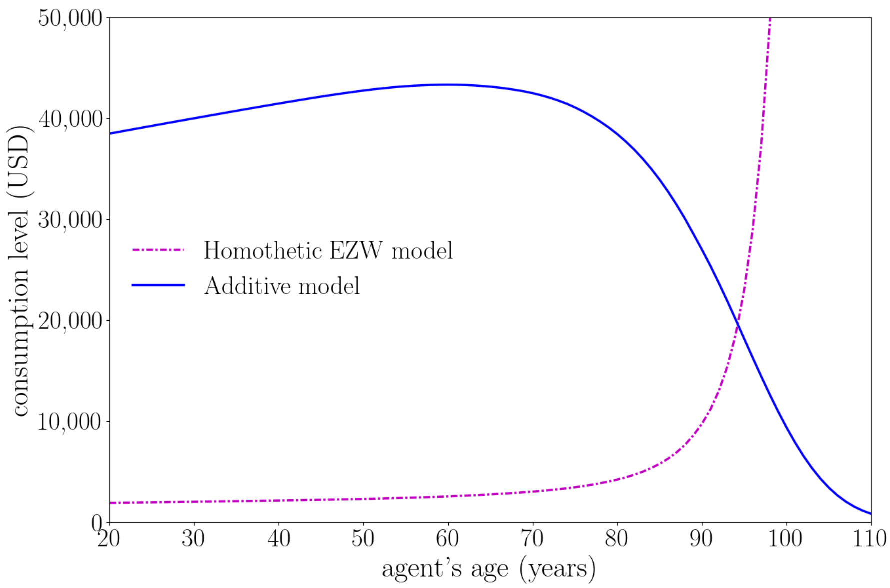

Figure 7: Consumption profiles implied by the additive model and the homothetic EZW model.

Willingness to pay for mortality risk reduction. From equation (33), we obtain the following expression for the willingness to pay for mortality risk reduction:

$$\frac{\partial V_t}{\partial \pi_t} \Bigg/ \frac{\partial V_t}{\partial c_t} = \frac{1}{(1-\gamma)(1-\beta)} c_t^{\sigma} \beta \pi_t^{\frac{\gamma-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}, \tag{37}$$

[26]Without such an assumption, there would be no solution fulfilling the budget constraint as consumption would rapidly grow to infinity at old ages.

[27]See also Zhang et al. (2018) for other issues related to EZW in consumption smoothing. They show that a low IES can amplify consumption variations (instead of smoothing them).

which is decreasing in the continuation utility V_{t+1}: $\frac{\partial}{\partial V_{t+1}}\left(\frac{\partial V_t}{\partial \pi_t}\Big/\frac{\partial V_t}{\partial c_t}\right) < 0$. This means that as the larger the loss in case of death (measured by the continuation utility V_{t+1}), the *less* (and not the more) the agent is willing to avoid the loss. At the extreme, the agent has an infinite willingness-to-pay to marginally increase their survival probability when they know that they will consume nothing and be miserable if they survive ($V_{t+1} = 0$). Conversely, this willingness-to-pay is zero if they know that they will consume infinite amounts and have an outstanding life if they survive ($V_{t+1} = \infty$). Also, the willingness to pay for mortality reduction does not increase but decreases with future survival probabilities π_{t+1},π_{t+2}, etc. (and hence with the expected length of life in case of survival). Again, these results are at odds with the most basic economic intuition (and at odds with results obtained using an additive expected utility specification or risk-sensitive preferences).

Better understanding the underlying mechanism. Let us start with a renormalization that will help clarify how these EZW specifications work. Set $U_t = \frac{V_t^{1-\sigma}}{(1-\sigma)(1-\beta)}$, which represent the same preferences as the V_t. From (33), we verify that the U_t fulfill the following recursion:

$$U_t = \frac{c_t^{1-\sigma}}{1-\sigma} + \beta\left(\pi_t^{\frac{1-\sigma}{1-\gamma}} U_{t+1} + (1 - \pi_t^{\frac{1-\sigma}{1-\gamma}})0\right), \tag{38}$$

This provides an insightful reading of the homothetic EZW model. Indeed, the above recursion is very similar to the standard additive specification where:

$$U_t = u(c_t) + \beta\left(\pi_t U_{t+1} + (1 - \pi_t)u_\dagger\right) \tag{39}$$

More precisely, (38) can be obtained from (39), by setting $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ for the instantaneous utility, $u_\dagger = 0$ for the utility of death and using transformed survival probabilities. Namely, (38) uses $\pi_t^{\frac{1-\sigma}{1-\gamma}}$ instead of π_t in (39). The EZW model can thus be seen as a (recursive) Rank-Dependent Expected Utility (hereafter RDEU, see Quiggin, 1982 and Yaari, 1987) that extends the standard (recursive) Expected Utility formulation. When $\frac{1-\sigma}{1-\gamma} > 0$, the probability transformation, $\pi \mapsto \pi^{\frac{1-\sigma}{1-\gamma}}$, corresponds to a standard way of changing risk aversion in RDEU. If $\frac{1-\sigma}{1-\gamma} < 0$, the probability transformation, $\pi \mapsto \pi^{\frac{1-\sigma}{1-\gamma}}$, is decreasing, a case excluded in RDEU, as it violates monotonicity with respect to first-order stochastic dominance. Here, non-monotonicity is reflected in the fact that the survival probability does not increase but decreases the weight given to continuation utility. This is what drives the substitutability between consumption and survival. This is also the reason why the willingness to pay for mortality risk reduction decreases with continuation utility. A model in which the weight given to welfare in a given state of the world decreases with the likelihood of occurrence of that state is not fully consistent with standard notions of rationality. However, this feature is necessary in this model to have a positive willingness to pay for mortality risk reduction. Indeed when $\sigma > 1$ the instantaneous utility $\frac{c_t^{1-\sigma}}{1-\sigma}$ in (38) is negative, and so are the U_t, which need to be "counterbalanced" in some way to get a positive value of life. Finally, we understand from (38) that, when $\gamma < 1 < \sigma$, the positive willingness to pay for mortality risk reduction results from two features: (i) the "good" state is death (utility 0), while being alive is the "bad" state (utility $U_{t+1} < 0$); (ii) the weight given to the utility obtained in the case of survival (U_{t+1}) does not increase but *decreases* with the survival probability. These features offset each other to produce a positive willingness to pay for mortality risk reduction. However, as discussed above, this also leads to an implausible substitutability between consumption and survival (while, in the absence of bequest, consumption is only valuable when alive) and also to a negative relationship between the value of mortality risk reduction and the quality and length of life expected in the case of survival.

References

Aczél, J. (1966): *Lectures On Functional Equations and Their Applications*, Academic Press, New York and London, reprinted in 2006 by Dover Publications.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2018): "Multiattribute Utility Theory, Intertemporal Utility, and Correlation Aversion," *International Economic Review*, 59, 537–555.

Bernheim, B. D. (2009): "Behavioral welfare economics," *Journal of the European Economic Association*, 7, 267–319.

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.

Bommier, A. (2006): "Uncertain Lifetime and Intertemporal Choice: Risk Aversion As a Rationale for Time Discounting," *International Economic Review*, 47, 1223–1246.

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

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," Working 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.

Calvet, L. E. and P. Sodini (2014): "Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios," *The Journal of Finance*, 69, 867–906.

Camerer, C. F. and T.-H. Ho (1994): "Violations of the Betweenness Axiom and Nonlinearity in Probability," *Journal of Risk and Uncertainty*, 8, 167–196.

Chew, S. H. and L. G. Epstein (1990): "Nonexpected Utility Preferences in a Temporal Framework with An Application to Consumption-Savings Behaviour," *J. Econ. Theory*, 50, 54–81.

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 (2023): "The Full Recession: Private Versus Social Costs of Covid-19," *International Economic Review*, Forthcoming.

Dionne, G. and L. Eeckhoudt (1985): "Self-Insurance, Self-Protection and Increased Risk Aversion," *Economics Letters*, 17, 39–42, (Errata, Volume 19, 1985, 100).

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

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004): "Do the Rich Save More?" *Journal of Political Economy*, 112, 397–444.

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," Working Paper 9, London: Imperial College COVID-19 Response Team.

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.

Hammond, P. J. (1988): "Consequentialist Foundations for Expected Utility," *Theory and Decision*, 25, 25–78.

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

Harrison, G. W., A. Hofmeyr, H. Kincaid, B. Monroe, D. Ross, M. Schneider, and J. T. Swarthout (2022): "Subjective Beliefs and Economic Preferences During the Covid-19 Pandemic," *Experimental Economics*, 25, 795–823.

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.

——— (2022): "Valuing Life As An Asset, As a Statistic and At Gunpoint," *The Economic Journal*, 132, 1095–1122.

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.

Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.

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.

Pashchenko, S. and P. Porapakkarm (2022): "Value of Life and Annuity Demand," *Journal of Risk and Insurance*, 89, 371–396.

Quiggin, J. C. (1982): "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, 3, 323–343.

Richard, S. F. (1975): "Multivariate Risk Aversion, Utility Independence and Separable Utility Functions," *Management Science*, 22, 12–21.

Stanca, L. (2021): "Recursive Preferences, Correlation Aversion, and the TemporalResolution of Uncertainty," Working paper, Northwestern University.

Starmer, C. (2000): "Developments in Non-expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk," *Journal of Economic Literature*, 38, 332–382.

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.

Viscusi, W. K. (2021): "Economic Lessons for COVID-19 Pandemic Policies," *Southern Economic Journal*, 87, 1064–1089.

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.

——— (1987): "The Dual Theory of Choice Under Risk," *Econometrica*, 55, 95–105.

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, `http://dx.doi.org/10.2139/ssrn.3245960`.