Recursive Preferences and the Value of Life: A Clarification

Antoine Bommier Daniel Harenberg François Le Grand*

January 11, 2021

Abstract

Two articles (Hugonnier et al., 2013; Córdoba and Ripoll, 2017) have proposed a recursive formulation of utility functions combining a positive value of life, preference homotheticity, and a constant elasticity of substitution. However, when the elasticity of substitution is below one and mortality rates take plausible values, this recursive formulation admits only a unique solution where utility is constant and equals zero everywhere. Non-constant solutions may only exist if mortality rates are assumed to remain low at all ages, namely in a world of perpetually young agents. Such solutions are therefore unsuitable for studying the value of life in demographically relevant settings and yield counterfactual predictions for saving behavior. We conclude this clarifying paper by reviewing various recursive specifications that can be used to study the value of life with no such problems.

Keywords: value of life, recursive utility, life-cycle models.

JEL codes: G11, J17.

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

1 Introduction

Following the seminal works of Epstein and Zin (1989) and Weil (1989), recursive utility models have become a workhorse of economic modeling, with applications in numerous fields. While recursive preferences were initially developed to address long-standing puzzles in the macro-finance literature, they are now increasingly employed in other fields such as the economics of climate change, health economics, or household finance. Two studies published in the Review of Economic Studies, Hugonnier et al. (2013, henceforth HPSA) and Córdoba and Ripoll (2017, henceforth CR) argue in favor of using homothetic recursive preferences to discuss questions relating to the value of life.[1] These papers are becoming increasingly popular, with citations in many prestigious journals, including Econometrica, and the Review of Economic Studies. They were suggested to reassess fundamental questions such as the welfare impact of the longevity crisis stemming from AIDS and wars. They are now being used to contribute to the growing literature that intends to price the mortality risk caused by the Covid-19 pandemic. For example, Córdoba et al. (2020) suggest to use for that purpose the recursive model of CR, rather the standard additive one – and thus to follow a route that diverges from that of Hall et al. (2020), Hammitt (2020), Greenstone and Nigam (2020), or Robinson et al. (2020) among others. The central argument of CR and HPSA is that the so-called Epstein-Zin-Weil (EZW, henceforth) preferences provide recursive utility representations featuring homotheticity, a constant elasticity of substitution (which may be smaller or larger than one), and a positive value of life independent of the level of consumption. The latter feature means that life is preferable to death, no matter the consumption level. With the standard additive model, combining these three properties is only possible when the intertemporal elasticity of substitution (IES, henceforth) is larger than one. HPSA and CR claim that a significant advantage of the EZW formulation is that it can also cover the case where the IES is below one, which is usually seen as the

[1]The two articles differ in their modeling approach. HPSA consider a continuous-time model, whereas CR employ a discrete-time model.

empirically relevant case (see, for example, Havránek, 2015).

The current paper makes three contributions. First, we point out the problematic features that emerge when the IES is assumed to be below one in the CR and HPSA specifications. Our remarks relate to both the mathematical properties of these models and the economic predictions they provide. These predictions are at odds with those of the standard additive model and with economic intuition. Second, we show that correcting the mathematical analysis in CR helps to better understand the origins of these problematic features. Last, we explain that there are recursive specifications, different from those of CR and HPSA, that provide a consistent extension of the additive model in the presence of mortality risk and can be most useful for economic analysis.

In our first contribution, we explain that if the IES is below one, the recursive models of HPSA and CR, when applied to actual mortality patterns, admit a unique solution where utility is constant and equals zero everywhere. Non-constant solutions only exist when mortality is constrained to remain small at all ages. Such a restriction involves assuming that agents are perpetually young, with a life expectancy that remains large at all ages. This makes such models unsuitable for studying value-of-life issues with realistic demographic data. In addition, these non-constant solutions imply that consumption at a given age and survival at that same age are substitutes (and not complements), leading to counterfactual predictions regarding life-cycle consumption profiles.

The constant zero utility in HPSA and CR arises when death is assumed to provide a utility equal to zero (just like being alive and consuming nothing, with their normalization). Looking at well-defined specifications where death is assumed to provide a positive utility level helps shed light on the contributions of HPSA and CR. In particular, useful insights may be gained by examining the limit-model that is obtained when the utility of death is assumed to be positive but infinitesimally close to zero. Such limit is actually mentioned in CR and is also discussed in Córdoba et al. (2020). Our second contribution involves demonstrating that, contrary to what is stated in those papers, when the IES is below one, this limit-model suggests consumption profiles that substantially differ from those predicted by the homothetic specifications of CR and HPSA. The implications regarding the impact of mortality risk on saving behaviors are in fact contrary to those found in CR. Moreover, we find that when the utility of death tends to zero, the value of life becomes infinite in the limit-model, meaning that this model is unable to deal with endogenous mortality choices.

This paper does not claim that all recursive preferences are inadequate for studying value-of-life questions. On the contrary, and this is the third contribution of our paper, we point to the risk-sensitive preferences of Hansen and Sargent (1995) that enable to consistently extend the additive model. With such a recursive framework, the value of life is mostly driven by a parameter that determines the utility gap between life and death. The IES and the risk aversion parameter (which can be varied independently from one another) can take any positive value without generating a discontinuity. This contrasts with homothetic EZW preferences which imply a negative value of life when the coefficient of risk aversion is set above 1 or a constant zero utility when the IES and the risk aversion coefficient are below 1. Moreover, risk-sensitive preferences are monotone with respect to first-order stochastic dominance, avoiding the choice of dominated strategies that may occur when working with EZW preferences (see Bommier et al., 2017 or Bommier et al., 2020 for illustrations). Risk-sensitive preferences therefore provide an appealing theoretical framework for extending the literature on the relationship between the value of life and risk aversion, offering a way to complement the analyses of Eeckhoudt and Hammitt (2004), Kaplow (2005), Andersson and Treich (2011) and Bommier and Villeneuve (2012). We also explain that the risk-sensitive model can be easily extended to account for ambiguity aversion. Although we leave it for future work, this recursive model could provide a pathway for extending to a multi-period setting the contributions of Treich (2010) and Bleichrodt et al. (2019) regarding the impact of ambiguity and ambiguity aversion.

The structure of our paper is as follows. In Section 2, we explain the shortcomings of the EZW homothetic specifications that account for the possibility of death and assume an IES below one. Section 3 clarifies the origin of such shortcomings. Section 4 then takes a constructive approach and points to a well-behaved recursive framework that can be used to discuss value-of-life issues.

For the sake of simplicity, the analysis in the main body of our paper is developed in a discrete-time framework, as is the case in CR. To relate our work to that of HPSA, we need to consider a continuous-time setting. This is is done in Appendix C, where we show that the arguments developed in the main body of the paper also apply to continuous-time modeling.

2 Utility functions in Córdoba and Ripoll (2017) when the IES is smaller than one

2.1 Recursive models

To understand the problems relating to the recursive formulation employed in CR, we briefly review the foundations of recursive approaches. In discrete time, recursive models state that at time t, agents maximize a utility function U_t, which is the solution of a backward recursive equation:

$$\forall t \geq 0,\ U_t = W(x_t, \mu(U_{t+1})), \tag{1}$$

where x_t is a vector collecting an agent's choices at time t (such as consumption or labor effort), U_{t+1} is the continuation utility at date $t+1$, which is a random variable from the date t point of view, $\mu(\cdot)$ is a certainty equivalent, mapping random variables into a scalar, and $W(\cdot,\cdot)$ is an aggregator that combines the certainty equivalent of the continuation utility with date-t choices.

Not all recursions admit a solution, of course, and it is therefore important to check that the recursive equations (1) characterize a well-defined sequence of utilities $(U_t)_{t\geq 0}$. The easiest way to guarantee the existence of the sequence $(U_t)_{t\geq 0}$ is to assume an exogenous finite horizon $T < \infty$ and an exogenous terminal value U_T. Applying equation (1) at $t = T-1$ provides utility U_{T-1} and by backward induction all utilities U_t from $t = T-1$ to $t = 0$. In infinite horizon settings, the problem is technically more involved. Existence of a solution can be, for instance, addressed either by using fixed-point theorems or by showing that the infinite-horizon specification can be obtained as the limit of a sequence of finite-horizon specifications (see, e.g., the discussion in Boyd, 1990).

In many cases, the existence of well-behaved solutions is well-known and it is possible to work directly with equation (1) to compute first-order conditions and optimal strategies. But checking that there exists a well-defined solution to the recursive equation is nevertheless central, otherwise conclusions could be inaccurate. Consider, for example, in an infinite horizon setting with no uncertainty, the recursion:

$$U_t = (z_t^{1-\sigma} + U_{t+1}^{1-\sigma})^{\frac{1}{1-\sigma}}, \tag{2}$$

where z_t denotes time-t consumption, assumed to take values in a compact interval $[z_{min}, z_{max}] \subset \mathbb{R}_{++}$, and σ a scalar assumed to be greater than 1. We can easily check that the only solution of this recursion is $U_t = 0$, for all $t \geq 0$, independently of consumption levels.[2,3] Utility being constant, all consumption strategies are equally good. If we set such considerations aside, it might seem natural to derive from (2) that:

$$\frac{\partial U_t}{\partial z_t} = z_t^{-\sigma} U_t^{\sigma} \text{ and } \frac{\partial U_t}{\partial z_{t+1}} = z_{t+1}^{-\sigma} U_t^{\sigma}, \tag{3}$$

and to conclude that the marginal rate of substitution $\frac{\partial U_t}{\partial z_t} / \frac{\partial U_t}{\partial z_{t+1}}$ is given by $(z_{t+1}/z_t)^{\sigma}$. This would typically lead to stringent restrictions on optimal consumption plans being derived, while in fact all strategies are optimal. The problem, which only becomes

[2]Throughout the paper, we follow the convention that for any real number $\kappa < 0$, we have $0^{\kappa} = \infty$ and $(\infty)^{\kappa} = 0$.

[3]This result is indirectly related to the findings of Koopmans (1960), which show that time preferences are necessary for building a sound recursive model of intertemporal choice in infinite-horizon settings.

apparent when solving for U_t, is that the marginal utilities in (3) are actually both equal to zero, their ratio thus being undefined. We mention this example because it shares a number of similarities with the recursive approaches we discuss below.

2.2 Córdoba and Ripoll (2017)'s formulation

In order to ease the discussion, we will focus on the case where the only uncertainty at play relates to mortality. The most widely used model in the value of life literature relies on the standard additive expected utility model initially introduced by Yaari (1965). In a discrete-time setting, the utility is recursively defined by:

$$V_t^{add} = u(z_t) + \beta \pi_t V_{t+1}^{add}, \tag{4}$$

where $z_t > 0$ is the consumption level, π_t is the probability of surviving from period t to period $t+1$ and $\beta < 1$ is a discount factor. The period utility function is most often assumed to be $u(z_t) = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l$, so that the IES, $\frac{u'(z_t)}{-z_t u''(z_t)}$, is constant and equal to $\frac{1}{\sigma}$. The constant u_l fixes the utility gap between life and death (whose utility is normalized to zero). While the constant u_l can be ignored when considering choices under an exogenous mortality patterns it plays a key role when discussing the welfare impacts of mortality risk. It can be checked that, if consumption takes value in a compact interval $[z_{min}, z_{max}] \subset \mathbb{R}_{++}$, such a recursion has for solution:

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

The utility V_t^{add} is thus a sum of instantaneous utilities weighted by the probability of being alive in the corresponding period and a discount factor.

CR suggest to depart from this specification and to use EZW preferences. In a first step, CR formalize the problem without constraining the utility of death to a specific value. Denoting the utility of a dead agent by $\underline{V}$, their recursive model is written as:

$$V_t = \left[z_t^{1-\sigma} + \beta \left(\pi_t V_{t+1}^{1-\gamma} + (1-\pi_t)\underline{V}^{1-\gamma}\right)^{\frac{1-\sigma}{1-\gamma}}\right]^{\frac{1}{1-\sigma}}, \tag{5}$$

where γ is a parameter driving risk aversion and other notation is the same as in the additive specification (4).

The problem arises when CR require the utility of dead agents to be equal to zero, $\underline{V} = 0$. Setting $\underline{V} = 0$ is central to their analysis, since they claim that this assumption guarantees that:

1. the value of life is positive for all consumption levels, avoiding the convexity issues raised by Rosen (1981) and the taste for a "Russian-roulette type of lottery" (CR, pp. 1473 and 1480–81);

2. the representation is homothetic (CR, Section 2.3); and

3. the model is well-ordered in terms of risk aversion, which would not be the case otherwise (CR, Section 5.2).

CR emphasize that "What makes EZW utility more flexible than EU is the possibility of setting $\underline{z} = 0$ [i.e., $\underline{V} = 0$] when $\sigma > 1$, so that non-convexities are eliminated and life is valued by all." (CR, p. 1482). Most of their subsequent analysis is restricted to the case where $\underline{V} = 0$. The same is true for Córdoba et al. (2020).

Let us therefore investigate the implications of setting $\underline{V} = 0$. CR explain that this requires $\gamma < 1$, so that $\underline{V}^{1-\gamma} = 0$. With $\underline{V} = 0$ and $\gamma < 1$, equation (5) reduces to:

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

We now analyze the utility representation that solves the recursive equation (6). We focus on the empirically relevant case where $\sigma > 1$. Note that with $\gamma < 1$, which is indeed needed for the value of life to be positive, this implies that the ratio $\frac{1-\sigma}{1-\gamma}$, which appears in (6) and in many instances in the following, is negative.

2.3 Solutions to Córdoba and Ripoll's recursive model

We start with the demographically relevant case, featuring a maximal and finite lifespan T. After age T, the agent is definitely dead and survival probabilities are zero. From an economic standpoint, the value of age T is of no importance, provided that it is larger than the number of years that a human may realistically live. It could be for example 200 years or more. Formally, we have $\pi_T = 0$ and $V_{T+1} = 0$, where the last equality comes from the assumption of zero utility for death, $\underline{V} = 0$. We can compute V_0 by backward induction. First, we compute V_T from equation (6), for all $z_T \geq 0$ and obtain:

$$V_T = \left[z_T^{1-\sigma} + \beta 0^{\frac{1-\sigma}{1-\gamma}} 0^{1-\sigma} \right]^{\frac{1}{1-\sigma}} = 0,$$

since $\sigma > 1 > \gamma$. Similarly, at date $T-1$, we obtain for all $z_{T-1} \geq 0$ and $\pi_{T-1} \in [0, 1]$:

$$V_{T-1} = \left[z_{T-1}^{1-\sigma} + \beta \pi_{T-1}^{\frac{1-\sigma}{1-\gamma}} 0^{1-\sigma} \right]^{\frac{1}{1-\sigma}} = 0.$$

By induction we then have $V_t = 0$ for all t, irrespective of consumption levels and mortality rates. Assuming a finite upper bound on lifespan therefore implies that the only solution to recursion (6) is the constant zero utility.

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

Proposition 1 *Consider the utility function defined by the recursion (6) with $\gamma < 1 < \sigma$.*

1. *If there is a maximal lifespan (i.e., there exists T such that $\pi_T = 0$), the only solution to (6) is $V_t = 0$ for all t.*

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

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

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

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

$$V_t = \left[\beta^{-t}\left(\Pi_{j=0}^{t-1}\pi_j\right)^{\frac{\sigma-1}{1-\gamma}} K + \sum_{s=t}^{\infty}\beta^{s-t}\left(\Pi_{j=t}^{s-1}\pi_j\right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \tag{7}$$

for some constant $K \geq 0$.

Proposition 1, which is proved in Appendix A, states that whenever survival rates become low at large ages, the only solution to the recursive equation (6) is zero utility, that is $V_t = 0$ for all t.

The recursive equation (6) 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.[4] Consequently, the utility function (7), solution to (6), could be used in a model of perpetual youth, but not in a life-cycle model that accounts for actual mortality profiles.

2.4 Implications of Córdoba and Ripoll's approach

As mentioned in our introductory example in Section 2.1, it can be tempting to compute the first-order conditions implied by a recursive utility function, without addressing convergence issues. We explore below the implications of such a computation when starting from equation (6).

Survival probability and the utility of consumption. First, let us compute the marginal rate of substitution (MRS, henceforth) between consumption in period $t+1$ and consumption in period t. Ignoring definitional issues, equation (6) implies:

$$\frac{\partial V_t}{\partial z_{t+1}} \bigg/ \frac{\partial V_t}{\partial z_t} = \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} z_{t+1}^{-\sigma} z_t^{\sigma}. \tag{8}$$

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

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

$$\frac{\partial}{\partial \pi_t}\left(\frac{\partial V_t}{\partial z_{t+1}} \Big/ \frac{\partial V_t}{\partial z_t}\right) < 0. \tag{9}$$

Equation (9) states that the less likely survival is at a given age, the more the agent wants to save resources for consumption at that age. The approach therefore involves assuming that consumption at a given age and survival at that same age are substitutes – while complementarity could be expected given the absence of bequest motive and thus of "utility of consumption" after death. In the limit case where $\pi_t \to 0$ (death is almost sure at the end of period t), we find that $\frac{\partial V_t}{\partial z_{t+1}} \Big/ \frac{\partial V_t}{\partial z_t} \to \infty$. Thus, a marginal increase in consumption would be infinitely more valuable in period $t+1$ than in period t, in spite of the fact that survival in period $t+1$ is extremely unlikely. This seems counter-intuitive, and is at odds with what the additive expected utility model suggests.[5] As is discussed below, such a MRS eventually yields counterfactual predictions.

Optimal life-cycle consumption profiles. Consider the standard life-cycle consumption-saving problem studied in Section 3 of CR. To simplify the problem, assume that there is no leisure or health in the analysis, no binding borrowing constraint, and no annuity market, so that the agent's program can be written as:[6]

$$\begin{aligned} V_t(w_t, \pi_t) &= \max_{z_t, w_{t+1}} \left(z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}(w_{t+1}, \pi_{t+1})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \\ \text{s.t. } y_t + w_t &= z_t + \frac{1}{1+r} w_{t+1}, \end{aligned} \tag{10}$$

where y_t and w_t denote income and wealth in period t, respectively. Savings are assumed to pay a constant riskless return $1+r$.

[5] With the additive utility model (4), $\frac{\partial V_t^{add}}{\partial z_{t+1}} \Big/ \frac{\partial V_t^{add}}{\partial z_t} = \beta \pi_t z_{t+1}^{-\sigma} z_t^{\sigma}$. The MRS thus increases with survival probability and tends to zero when the survival probability π_t becomes infinitesimally small.

[6] To make the link with CR, set $z_t = c_t$, $w_t = a_t$, and $H_t = 1$ in their equation (11). We study the impact of health in the next paragraph.

Using first-order conditions, CR derive:[7]

$$\frac{z_{t+1}}{z_t} = \left(\beta(1+r)\pi_t^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{\sigma}}. \tag{11}$$

CR comment on equation (11), stating that "the effect of higher survival on consumption growth under EZW preferences can be negative, which is not possible under EU". This means that if $\sigma > 1 > \gamma$, mortality reduces impatience instead of contributing to it.[8] This is a consequence of the substitutability between consumption and survival exhibited in equation (9).

We now examine the quantitative implications of equation (11) for consumption using benchmark parameters and realistic mortality profiles. We assume that the interest rate is $r = 4\%$. Preference parameters are the same as in Footnote 4, $\beta = 0.97$, $\sigma = 2.0$, and $\gamma = 0.5$. Mortality rates are those of the total US population in 2018, as reported in the Human Mortality Database. In Figure 1, we plot the consumption path implied by the first-order equation (11). For the sake of comparison, we also plot the consumption path implied by standard additive expected utility model (also assuming $\beta = 0.97$ and $\sigma = 2.0$). Lifetime wealth is normalized to 1,000,000 USD, such that the consumption at age 20 for the additive agent is close to 40,000 USD. This normalization is of little importance, because preferences are homothetic. The CR profile exhibits a consumption level that remains extremely low until age 100, but that sky-rockets after that. It is extremely different from the one obtained with the additive specification. Note that the y-axis is truncated at 50,000 USD for readability reason but under the CR model, consumption in fact reaches 98,000 USD at age 100 and 14,000,000 USD at age 110.

In a complementary analysis, Zhang et al. (2018) also remark that with the solution (11), the parameter σ, which is the inverse of the IES, may end up playing an unexpected role, qualitatively inconsistent with what would be obtained with an

[7]Cf. equation (15) in their paper.

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

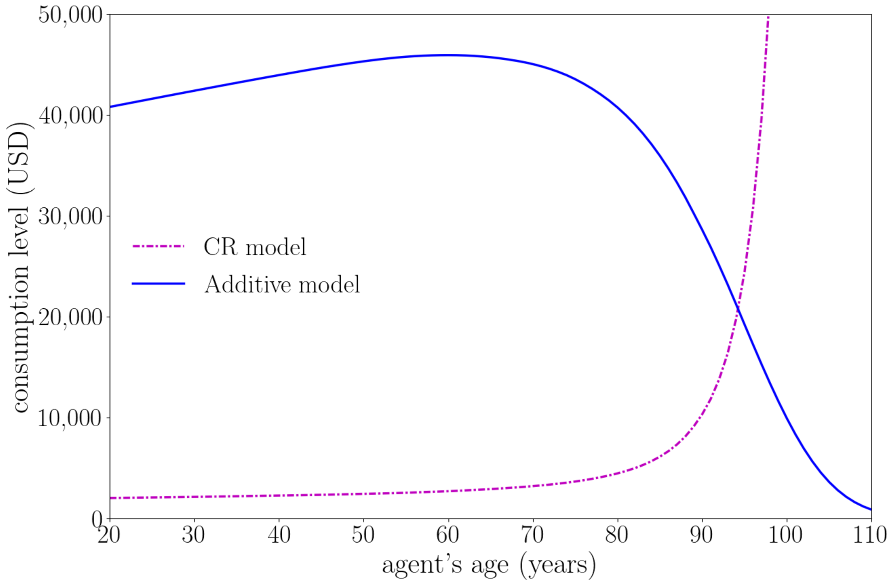

Figure 1: Consumption profiles implied by the additive and the CR models.

additive model. For example, when $\beta(1+r) = 1$, (11) implies a rate of consumption growth equal to $\frac{1-\sigma}{\sigma}\log(\pi_t)$. This growth rate is always positive when $\sigma > 1$ and increases with σ. Thus, instead of moderating the variations of consumption over time, a low IES would in fact amplify such variations.

Age-dependent health. In order to circumvent the difficulties highlighted above, CR introduce an age-dependent variable H_t, which is interpreted as health. The recursive equation (6) becomes:

$$V_t = \left[H_t z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}\right]^{\frac{1}{1-\sigma}}. \tag{12}$$

Since H_t depends on t, there are enough degrees of freedom to match any possible consumption profile. In principle, a rapidly declining profile H_t could potentially alter the counterfactual implications discussed above. However, with $\sigma > 1$, specification (12) assumes that utility *decreases* (rather than increases) with H_t, which hinders the interpretation of H_t as health.[9] In a similar vein, Córdoba et al. (2020) introduce age-dependent discount factors $(\beta_t)_{t\geq 0}$ which, as they explain, should decrease with age (when $\sigma > 1$) to counteract the term $\pi_t^{\frac{1-\sigma}{1-\gamma}}$. In other words, to avoid the pattern shown in Figure 3, one would have to assume that pure time preferences (which would govern impatience in absence of mortality risk) become rapidly stronger with age, so as to compensate for the implausible mortality effect discussed above.

The calibration of the "health parameters" (or of the time-varying discount factors β_t in Córdoba et al., 2020), and issues related to their potential endogeneity, may play a decisive role when discussing the impact of (exogenous or endogenous) mortality changes. Compare, for example, the utility functions in the CR model and the standard additive specification of Murphy and Topel (2006):

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

CR calibrate the health profile so as to obtain the same consumption profile as Murphy and Topel (2006). In other words, they first set $H_t = \left(\Pi_{j=0}^{t-1}\pi_j\right)^{\frac{\sigma-\gamma}{1-\gamma}} H_t^{MT}$, where H_t^{MT} is the health profile chosen by Murphy and Topel (2006). A key point, however, is that they then assume the profile H_t to be exogenous and independent of survival rates when looking at the impact of mortality changes. This ultimately results in the two models forming radically different conclusions regarding the consequences of mortality decline. Murphy and Topel's model predicts that a decline in mortality would significantly increase the propensity to save (agents become more patient when survival probabilities increase), while the effect is much smaller or even opposite in the CR model (agents become more impatient when survival probabilities increase).[10]

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

[10]Longevity extension also generates an income effect that adds to the impatience effect we

Opting for one specification over the other will thus provide very different views regarding the impact of population aging. While both models match the same calibration targets, the implicit assumptions they make about the role of health (considered as "good" in Murphy and Topel and as "bad" in CR) cause them to reach opposite conclusions.

Willingness to pay for mortality risk reduction. We now turn to the MRS between survival and consumption, which we will refer to as the willingness to pay for mortality risk reduction.[11] From equation (6), we obtain:

$$\frac{\partial V_t}{\partial \pi_t} \Big/ \frac{\partial V_t}{\partial z_t} = \frac{1}{1-\gamma} z_t^{\sigma} \beta \pi_t^{\frac{\gamma-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}. \tag{13}$$

With $\gamma < 1 < \sigma$, the willingness to pay for mortality risk reduction is decreasing in the continuation utility V_{t+1}:

$$\frac{\partial}{\partial V_{t+1}} \left(\frac{\partial V_t}{\partial \pi_t} \Big/ \frac{\partial V_t}{\partial z_t} \right) < 0. \tag{14}$$

Equation (14) means that as the possible future loss (measured by continuation utility V_{t+1}) increases, the agent will be less willing to avoid the loss. At the extreme, the agent has an infinite willingness-to-pay to marginally increase survival probability when she knows that she will consume nothing and be miserable if she survives ($V_{t+1} = 0$). Conversely, this willingness-to-pay is zero when she knows that she will consume huge amounts and have a superb life if she survives ($V_{t+1} = \infty$). Here again, the results are at odds with those obtained when using an additive expected utility specification, and with economic intuition. For example with the model of Murphy

emphasize, which explains why the overall impact can be ambiguous.

[11] The Environmental Protection Agency advises using this terminology, rather than the "value of a statistical life" (VSL). Note, moreover, that the literature generally defines VSL as the marginal rate of substitution between survival and wealth (and not consumption). In the absence of an annuity, this makes no difference and VSL, as usually defined, also equals the MRS shown in (13). When annuities are available, we have to account for the fact that lowering mortality rates may reduce the return of annuities. An adjustment must therefore be made. For example, with a perfect annuity market, VSL is the difference between the MRS (13) and the amount invested in annuities. The adjustment is independent of the agent's preferences, and would therefore have no impact on the subsequent discussion.

and Topel shown above, we have:

$$\frac{\partial V_t^{MT}}{\partial \pi_t} \Big/ \frac{\partial V_t^{MT}}{\partial z_t} = \frac{z_t^\sigma}{H_t^{MT}} \sum_{\tau=t+1}^{\infty} \beta^t \left(\Pi_{j=t+1}^{\tau-1} \pi_j\right) H_\tau^{MT} \left(\frac{z_\tau^{1-\sigma}}{1-\sigma} + u_l\right),$$

implying that the value of mortality risk reduction increases with future consumption levels: expecting a nice future makes survival more valuable.

3 Investigating the limit-model when $\underline{V} \to 0$

A way to (partially) fix the convergence issues discussed in Section 2 involves using the non-homothetic recursion (5) and taking the limit where $\underline{V} \to 0$. Such a limit is actually considered in Section 5 of CR, but as we explain below, the mathematical statements when considering this limit are incorrect.

For any $\underline{V} > 0$, if we assume that there is a finite maximal length of life, the recursion (5) provides a well-defined and non-zero sequence of utilities $(V_t)_{t\geq 0}$, from which we can compute optimal life-cycle consumption profiles and the value of mortality risk reduction. Looking at the limit where $\underline{V} \to 0$ (if it exists) may then provide a robust theoretical foundation for the study of the case where $\underline{V} = 0$. We explore this possibility below, focusing again on the case where $\gamma < 1 < \sigma$.

Formally, starting from recursion (5), we renormalize the utility representation by setting $W_t = V_t/\underline{V}$. The utilities $(W_t)_{t\geq 0}$ fulfill the recursion:

$$W_t = \left[\left(\frac{z_t}{\underline{V}}\right)^{1-\sigma} + \beta\left(\pi_t W_{t+1}^{1-\gamma} + 1 - \pi_t\right)^{\frac{1-\sigma}{1-\gamma}}\right]^{\frac{1}{1-\sigma}}. \tag{15}$$

The mistake in CR involves assuming that when $\underline{V} \to 0$, we necessarily have $W_{t+1} \to \infty$ since $W_{t+1} = V_{t+1}/\underline{V}$.[12] With such an assumption recursion (15) would

[12] See page 1503 in CR, where it is explained that the ratio $\underline{V}/V_{t+1}$ which appears in the first equation of their Section 5.1 collapses to zero when $\underline{V} \to 0$. This, however, does not take into account that V_{t+1} also tends to zero when $\underline{V} \to 0$, yielding a ratio $\underline{V}/V_{t+1}$ whose limit is of an indeterminate form $\frac{0}{0}$. In fact, as we will explain below, $\lim_{\lim_{\underline{V}\to 0}} \frac{\underline{V}}{V_{t+1}} \neq 0$, implying that $\lim_{\lim_{\underline{V}\to 0}} W_{t+1}$ is actually finite. Note that the limit $\underline{V} \to 0$ is also considered in Córdoba et al. (2020), but only when $\sigma = 1$. For $\sigma = 1$ the statement $\underline{V}/V_{t+1} \to 0$ is correct. That would also be true for $\sigma < 1$, but the issues for $\sigma > 1$ are no different in Córdoba et al. (2020) than in CR.

yield:

$$W_t \simeq \left[\left(\frac{z_t}{\underline{V}} \right)^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} W_{t+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{16}$$

with solution $W_t \simeq \frac{1}{\underline{V}} \sum_{s=t}^{\infty} \beta^{s-t} \left(\Pi_{j=t}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma}$. As explained in CR, this would lead us back to the formulation they use (up to the scaling factor $\frac{1}{\underline{V}}$) to deal with the case where $\underline{V} = 0$. Said differently, their analysis for the case where $\underline{V} = 0$ could be obtained by continuity from the cases where $\underline{V} > 0$. The definitional issues highlighted in Section 2.2 of the current paper would not be worrisome since models with $\underline{V} > 0$ are well-defined and would converge towards their preferred specification.

However, when $\underline{V} \to 0$, V_t also tends to zero, and the ratio $V_t/\underline{V}$ converges to a finite limit, and not to ∞. Indeed since $\sigma > 1$, we have $\lim_{\underline{V}\to 0} \left(\frac{z_t}{\underline{V}} \right)^{1-\sigma} = 0$. Equation (15) therefore implies that $\lim_{\underline{V}\to 0} W_t = \chi_t$ where $(\chi_t)_{t\geq 0}$ is defined by:

$$\chi_t = \beta^{\frac{1}{1-\sigma}} \left(\pi_t \chi_{t+1}^{1-\gamma} + 1 - \pi_t \right)^{\frac{1}{1-\gamma}}. \tag{17}$$

If $\beta^{\frac{1-\gamma}{1-\sigma}} \pi_t < 1$ for large t (which holds in case of a finite life), the forward recursion admits a unique solution, which is finite for all t. The approximation (16) does not hold, meaning that the specification CR attributes to the case $\underline{V} = 0$ is different from the limit obtained when considering the limit $\underline{V} \to 0$.

This matters for economic implications. To illustrate this, let us compute marginal rates of substitutions implied by the recursive representation (15) when $\underline{V} \to 0$. Since $\lim_{\underline{V}\to 0} W_t = \chi_t$, we obtain that when $\underline{V} \to 0$:

$$\frac{\frac{\partial W_t}{\partial z_{t+1}}}{\frac{\partial W_t}{\partial z_t}} = \beta \pi_t \left(\pi_t + (1-\pi_t) W_{t+1}^{\gamma-1} \right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t} \right)^{-\sigma}, \tag{18}$$

$$\to_{\underline{V}\to 0} \beta \pi_t \left(\pi_t + (1-\pi_t) \chi_{t+1}^{\gamma-1} \right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t} \right)^{-\sigma}, \tag{19}$$

and

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

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

Note that the MRS in (19) tends to zero when π_t tends to zero. This means that if survival from period t to period $t+1$ is very unlikely, the agent values an increase in consumption in period t much more than an increase in consumption in period $t+1$. While consistent with the additive model and economic intuition, this is contrary to the conclusions of CR, who consider that the MRS in the case where $\underline{V}=0$ is $\beta\pi_t^{\frac{1-\sigma}{1-\gamma}}\left(\frac{z_{t+1}}{z_t}\right)^{-\sigma}$, implying that the MRS would actually tend to ∞ (and not 0) when $\pi_t \to 0$. It can be observed that CR expression for the MRS can be obtained as the limit of equation (18) for $\underline{V} \to 0$, if it is assumed that $\lim_{\underline{V}\to 0} W_{t+1} = \infty$, which as we explain above does not hold.

When employed in applications, the limit model where $\underline{V}$ is taken to be infinitesimally small yields predictions that are relatively close to those of the additive model, but very different from the predictions obtained with CR specification. These aspects are illustrated in Figure 2 where we see that the limit- and additive models provide decreasing or hump-shaped consumption profiles (because mortality increase with age), in sharp contrast to the increasing consumption profile predicted by CR model. Note that, as in Figure 1, the y-axis is still truncated at 50,000 USD for readability reason.

Overall, we find that solving the issues outlined in Section 2.2 by considering the limit $\underline{V} \to 0$ yields a model where consumption and survival are complementary, but where the willingness to pay for mortality risk reduction is infinite (thus unsuitable for dealing with endogenous mortality choice). If this discussion helps to clarify the issues in CR and HPSA, it does not provide a convenient recursive framework for dealing with endogenous mortality choices.

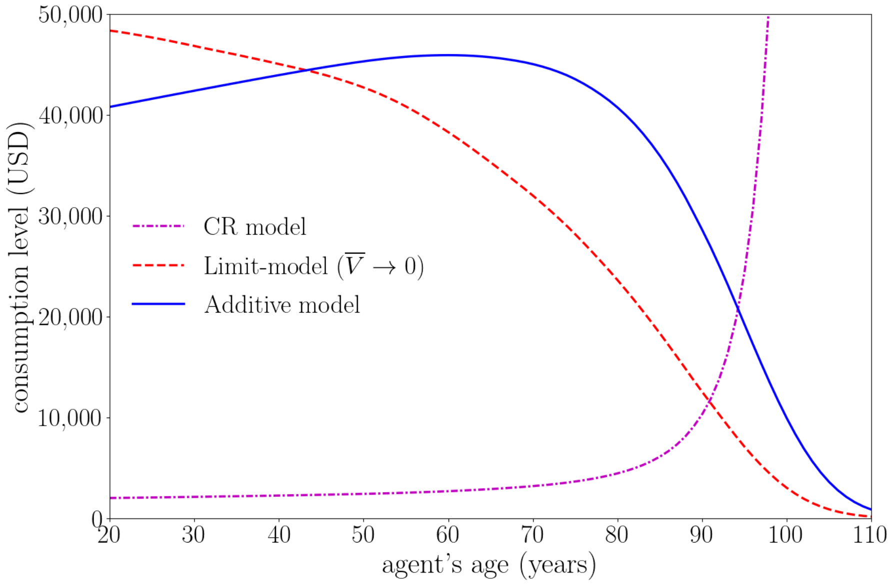

Figure 2: Consumption profiles implied by the additive, the CR, and the $\underline{V} \to 0$-limit models.

4 Studying the value of life with recursive models

The conclusion of our discussion should not be that recursive models are unsuitable for studying the value of mortality risk reduction. Recursive methods can be excellent tools for studying the value of life and intertemporal choice in a context of uncertain lifetime. Some theoretical aspects need however to be considered thoroughly.

4.1 An infinite horizon for finitely-lived agents

Recursive models usually assume an infinite horizon or a fixed finite horizon. However, to model choices under uncertain lifespans, we need to consider lives of unequal lengths, which is fundamentally different from an infinite or fixed horizon setting. There is a simple way to circumvent this difficulty. Instead of describing a life as a finite sequence of consumption periods, we can view a life as an infinite sequence of per-period "realizations", where each realization is either "being dead" or "being alive, with some positive consumption". Mathematically speaking, the set of possible per-period realizations is $\mathbb{R}_+ \cup \{d\}$, where d is a symbol used to denote "being dead".

A life is then represented by a sequence in the form:

$$(z_0, \ldots, z_T, d, d, \ldots) \in (\mathbb{R}_+ \cup \{d\})^\infty, \tag{20}$$

for some T. Sequence (20) corresponds to a lifespan of $T+1$ with consumption profile $(z_0, \ldots, z_T) \in \mathbb{R}_+^{T+1}$. Formally, if one excludes resurrection and immortality, the set of possible lives, denoted by L, is defined as follows:

$$L = \left\{(z_t)_{t\geq 0} \in (\mathbb{R}_+ \cup \{d\})^\infty : \ \exists T \in \mathbb{N}, \forall t \geq T, z_t = d \text{ and } \forall t < T, z_t \in \mathbb{R}_+\right\}. \tag{21}$$

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

4.2 Recursive models with mortality

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

$$U_t = f^{-1}((1-\beta)u(z_t) + \beta\phi^{-1}(E[\phi f(U_{t+1})]), \tag{22}$$

where $\beta \in (0,1)$ is a time preference parameter, ϕ is an increasing function representing risk preferences, $E[\cdot]$ denotes the expectation operator, and $u : \mathbb{R}_+ \cup \{d\} \to \mathbb{R} \cup \{-\infty, \infty\}$ is the period utility function.[13] The function f is only a normalization device and can be any increasing function, with no impact on preferences. For example, CR use $f(x) = \frac{x^{1-\sigma}}{1-\sigma}$, but it is also possible to use $f(x) = \phi^{-1}(x)$, as in Kreps and Porteus (1978), or simply $f(x) = x$, which we will use in the following.

We now discuss how to further parameterize the functions u and ϕ. It is important

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

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

For our discussion, it is useful to consider conditional utilities, which are the utilities obtained conditionally on being dead or on being alive. For a dead agent, we simply get $U_t(d, d, \ldots) = u(d)$, which is the value of the period utility function, u, in the death state, d. Plugging this into (22), we find that the utilities of alive agents, denoted by V_t, are linked through the following recursion:

$$V_t = (1-\beta)u(z_t) + \beta\phi^{-1}\left(\pi_t E[\phi(V_{t+1})] + (1-\pi_t)\phi(u(d))\right), \qquad (23)$$

where, as before, $z_t \in \mathbb{R}_+$ is the consumption at time t and π_t is the survival probability between dates t and $t+1$. Note that – with a slight abuse of notation – we still denote the expectation operator by $E[\cdot]$, although mortality risk is now treated separately.[14]

4.3 The parametrization of the functions u and ϕ

We now discuss popular specifications of the functions u and ϕ. Obviously, other specifications are possible, as long as the domain of ϕ is carefully chosen.

4.3.1 The period utility function u

A common specification is the case where preferences exhibit a constant IES. Formally, this means that $\frac{-zu'(z)}{u''(z)}$, for $z \in \mathbb{R}_+$, is independent of z; or equivalently, by integration, that for $z \in \mathbb{R}_+$, we have $u(z) = K\frac{z^{1-\sigma}}{1-\sigma} + u_l$, where σ is the inverse of the IES and

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

K and u_l are two constants. The function u, defined over $\mathbb{R}_+ \cup \{d\}$, is then:

$$\begin{cases} u(z) = K\frac{z^{1-\sigma}}{1-\sigma} + u_l \text{ for } z \in \mathbb{R}_+, \\ u(d) = u_d, \end{cases} \tag{24}$$

for some $K > 0$, $u_l \in \mathbb{R}$, and $u_d \in \mathbb{R} \cup \{-\infty, \infty\}$. Setting $K = 1$ corresponds to a normalization that is always possible. We can further normalize the function u by setting either $u_l = 0$ or $u_d = 0$. But imposing an additional relation between u_l and u_d (such as $u_l = u_d$) goes beyond a mere normalization. In particular, this would constrain the value of life. In other words, one can freely constrain u_l, or alternatively u_d, but constraining both simultaneously is not without loss of generality.

4.3.2 The risk aversion function ϕ

The function ϕ needs to be properly defined on the convex hull of $Im(u) = u(\mathbb{R}_+) \cup u_d$, and it needs to be increasing.[15] We discuss below three of the most common functional forms that can be found in the literature for ϕ: (i) affine, (ii) isoelastic, and (iii) exponential. We further assume specification (24) for the period utility function u.

Affine ϕ. Defining ϕ as an affine function is probably the most common choice in the literature on the value of life, and corresponds to the standard additive expected utility model of Yaari (1965). It is most often associated with the normalization $u_d = 0$, but renormalizing u by adding the same constant to both u_l and u_d, while keeping ϕ unchanged, would have no impact on preferences. With $u_d = 0$ the recursion (23) defining the utility function of alive agents can be expressed as follows:

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

[15]The function ϕ, which governs both risk aversion and preference for the timing of resolution of uncertainty, does not have to be concave. It follows from Theorem 3 of Kreps and Porteus (1978) that preferences exhibit preference for early (resp. late) resolution if the function $x \mapsto \phi((1-\beta)u(c) + \beta\phi^{-1}(x))$ is convex (resp. concave).

Additionally requiring that $u_l = 0$ might seem appealing for tractability reasons, but the value of life would then be mostly pinned down by β and σ. Moreover, if $\sigma > 1$ and $u_l = u_d = 0$, the value of life is always negative, independent of the consumption level, which is a very undesirable property when studying value of life matters. Independently of the value of u_l, a well-known limitation of the affine specification is that risk aversion and IES are intertwined.

Isoelastic ϕ. This corresponds to the EZW preferences. One difficulty with isoelastic functions relates to their definition sets. These functions are never defined on the whole set $\mathbb{R}$ but either only on $\mathbb{R}_+$ (e.g., $\phi(x) = \frac{x^{1-\alpha}}{1-\alpha}$) or only on $\mathbb{R}_-$ (e.g., $\phi(x) = -\frac{(-x)^{1-\alpha}}{1-\alpha}$). Since ϕ needs to be defined on the convex hull of $u(\mathbb{R}_+) \cup u_d$, the model is well-defined if and only if $\frac{z_t^{1-\sigma}}{1-\sigma} + u_l$ and u_d always have the same sign. This implies that u_l and u_d must have the same sign as $1-\sigma$. All the studies using isoelastic functions ϕ that we are aware of set $u_l = 0$. Using the notation $\epsilon = sign(1-\sigma)$, the recursion (23) defining the utility function V_t of alive agents can be expressed as follows:

$$V_t = (1-\beta)\frac{z_t^{1-\sigma}}{1-\sigma} + \epsilon\beta\left(\pi_t E[(\epsilon V_{t+1})^{1-\alpha}] + (1-\pi_t)(\epsilon u_d)^{1-\alpha}\right)^{\frac{1}{1-\alpha}}. \qquad (25)$$

In such a model, setting $u_d = 0$ or $u_d = \epsilon\infty$ is technically possible. Both options are actually found in the literature, as this is the only way to recover homotheticity. For example, CR specification involves setting $u_d = -\infty$ if $\sigma > 1$ and $u_d = 0$ if $\sigma < 1$.[16] But setting $u_d = 0$ or $u_d = \epsilon\infty$ represents a severe restriction on preferences. Indeed, apart from the case where $\sigma < 1$, $\alpha < 1$, and $u_d = 0$, this generates models that either imply constant utilities (independent of consumption choice) or exhibit a value of life that is always negative.

Exponential ϕ. As shown in Bommier et al. (2017), assuming that preferences are monotone with respect to first-order stochastic dominance implies that the

[16]Remember that CR use a different normalization, so that their $\underline{V}$ is related to u_d by $u_d = \frac{\underline{V}^{1-\sigma}}{1-\sigma}$.

function ϕ has an exponential form, with $\phi : x \mapsto \frac{1-e^{-kx}}{k}$, $k \neq 0$,[17] yielding so-called risk-sensitive preferences. Such preferences were initially introduced by Hansen and Sargent (1995) in an infinite horizon setting and later adapted to the problem of intertemporal choice under uncertain lifespan in Bommier (2014) and Bommier et al. (2020).[18] Since the function $\phi : x \mapsto \frac{1-e^{-kx}}{k}$ is well-defined and increasing on the whole set $\mathbb{R}$, there is no domain problem. Indeed, ϕ is defined on the convex hull of $u(\mathbb{R}_+) \cup u_d$, regardless of the choice of u and u_d. Moreover, we can easily check that, when ϕ is exponential, re-normalizing u by adding the same constant to both u_l and u_d, while keeping ϕ unchanged, has no impact on preferences. We can therefore set $u_d = 0$ without loss of generality. The utility V_t is then obtained by the following recursion:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k} \log(\pi_t E[e^{-kV_{t+1}}] + 1 - \pi_t), \tag{26}$$

where the parameter k drives risk aversion.

The risk-sensitive specification (i.e., the case of an exponential function ϕ) offers a theoretically appealing framework, in which preferences are always well-defined and monotone. Flexibility is afforded by four degrees of freedom: σ, k, β, and u_l, which determine the IES, risk aversion, time preferences, and the utility gap between life and death (and thereby the value of life), respectively. The value of life can then be calibrated by choosing u_l.

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

[17] The case $k = 0$ corresponds to an affine ϕ.

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

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

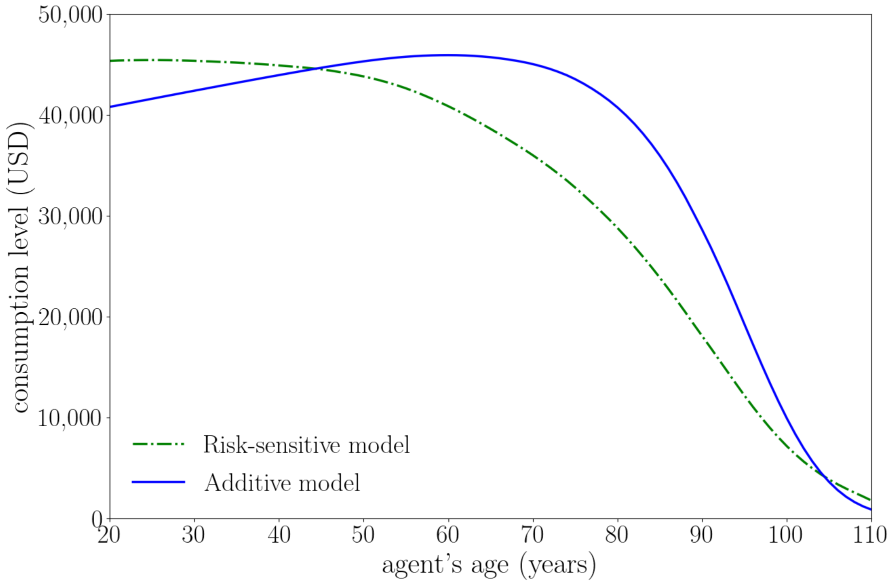

Figure 3: Consumption profiles implied by the additive and the RS models.

vanish if assuming k to be very small.

Further implications regarding portfolio choice and annuity demand are discussed in Bommier et al. (2020). In such a framework the value of mortality risk reduction has the following expression:

$$\frac{\frac{\partial V_t}{\partial \pi_t}}{\frac{\partial V_t}{\partial z_t}} = \frac{\beta}{k} z_t^{\sigma} \frac{1 - E[e^{-kV_{t+1}}]}{\pi_t E[e^{-kV_{t+1}}] + 1 - \pi_t}. \tag{27}$$

We observe in particular that the willingness to pay for mortality risk reduction is increasing in the continuation utility V_{t+1}. This means that people will make higher investment to reduce their mortality risk if they expect a nice future rather than a poor one. This is in line with what additive models suggest and with economic intuition but contrasts with the predictions discussed in Section 2.4.

An extension to account for ambiguity and ambiguity aversion. Interestingly, the risk-sensitive framework can easily be extended to account for ambiguity aversion while preserving monotonicity and differentiability (and hence tractability). A detailed presentation can be found in Bommier and LeGrand (2014) and an application to annuity investment features in André et al. (2020).

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

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k_A} \log\left(E_{\widetilde{\pi}_t}\left[\exp\left(\frac{k_A}{k_R} \log\left(\widetilde{\pi}_t E[e^{-k_R V_{t+1}}] + 1 - \widetilde{\pi}_t \right)\right)\right]\right), \quad (28)$$

where k_A is the ambiguity aversion coefficient and k_R is the risk aversion coefficient. The operator $E_{\widetilde{\pi}_t}[\cdot]$ is the expectation over $\tilde{\pi}_t$. With $k_A = 0$ (ambiguity neutrality), the model reduces to the risk-sensitive specification (26) with $k = k_R$ and $\pi_t = E[\widetilde{\pi}_t]$ (under ambiguity neutrality only the average survival probability matters). The additive model (4) is obtained when $k_A = k_R = 0$. When $k_R = 0$ and $k_A > 0$ we get:

$$V_t = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l - \frac{\beta}{k_A} \log\left(E_{\widetilde{\pi}_t}\left[\exp(-k_A \widetilde{\pi}_t E[V_{t+1}]) \right]\right). \quad (29)$$

To connect this specification to the previous literature, note that (29) can be rewritten:

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

where $\phi(x) = \exp(-\frac{k_A}{\beta}x)$ and $u(z_t) = \frac{z_t^{1-\sigma}}{1-\sigma} + u_l$. Written in such a way, the model appears to be a multi-period version of the Treich (2010) and Bleichrodt et al. (2019) models, while assuming constant absolute ambiguity aversion (which is necessary for preference monotonicity in multi-period settings), and additive separability of preferences under risk.[20] The more general monotone specification, given in (28), relaxes the assumption of additive separability of preferences under risk. As illustrated in André et al. (2020), this model can be used to jointly analyze the impacts of risk and ambiguity aversion.

In our view, the specification (28) is the best way to extend the standard ad-

[20] The Treich (2010) and Bleichrodt et al. (2019) models also feature a bequest motive, which we have not considered so far. To account for bequest motives specification (28) should be replaced by:

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

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

ditive specification, affording flexibility to account for risk and ambiguity aversion. Importantly, it maintains two fundamental features of the additive specification: recursivity and monotonicity. The former, recursivity, is key for tractability. The latter, monotonicity, has long been considered as inherent to rationality (see e.g., Arrow, 1951) and helps to get an intuitive understanding of the role of risk and ambiguity aversion (Bommier et al. 2017). Whether one should use such the specification (28), or focus on simpler versions with k_A or/and k_R set to zero involves considering a very common trade-off: Additional flexibility brings new insights, but also makes the calibration more challenging.

Appendix

A Proof of Proposition 1

Assume, for example, that there exists $t_0 \geq 0$ such that $V_{t_0} > 0$. It is then necessarily the case that $V_t > 0$ for all $t > t_0$.[21] Now, let us rewrite equation (6) as:

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

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

$$V_t^{1-\sigma} = \beta^{-t} \left(\Pi_{j=t_0}^{t-1} \pi_j \right)^{\frac{\sigma-1}{1-\gamma}} \left(V_{t_0}^{1-\sigma} - \sum_{s=t_0}^{t-1} \beta^s \left(\Pi_{j=t_0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right). \tag{30}$$

Consider now the case where $t \to \infty$, while holding t_0 constant. Assuming that consumption is bounded from above, the right-hand side of (30) becomes negative, since $\sum_{s=0}^{t-1} \beta^s \left(\Pi_{j=0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma}$ diverges (as we are considering the case $\lim_{t\to\infty} \pi_t = 0$), while the left-hand side has to be positive. We thus obtain a contradiction, proving that there cannot exist a t_0 for which $V_{t_0} > 0$. This impossibility result holds when $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}} > 1$ for large t.[22]

B Hugonnier, Pelgrin and St-Amour (2013) specification

The contribution of HPSA is more involved than that of CR because continuous-time modeling requires more advanced mathematics. But if we restrict our attention

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

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

to mortality risk while ignoring other risks, the HPSA model can be seen as a continuous-time limit of the CR model, with the addition of a minimum consumption level. As such, the model faces the same difficulties when the IES is below one.

To be more explicit, we use the notation of HPSA and build on their Appendix C.1 "Construction of the utility index" where they use the limit of discrete time models to derive their continuous-time model. As is explained after equation (C.1), "the agent's utility [is required] to drop to zero after death". Their equation (C.2) defines utility U_t by recursion over a time interval $\Delta > 0$. If Δ is small enough and in absence of risks other than mortality, U_t can be expressed as:

$$U_t = \left[\left(1 - e^{-\rho\Delta}\right) (c_t - a)^{1-\frac{1}{\varepsilon}} + e^{-\rho\Delta} \pi_t^{\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}} U_{t+\Delta}^{1-\frac{1}{\varepsilon}} \right]^{\frac{1}{1-\frac{1}{\varepsilon}}}, \tag{31}$$

where $\varepsilon > 0$ is the IES, $0 \leq \gamma_m < 1$ is a risk aversion parameter, $\rho > 0$ the rate of time-discounting, $a \geq 0$ a subsistence consumption level, and π_t the survival probability.[23] We consider again the case where the IES is below 1 ($\varepsilon < 1$). Note that – apart from notation and subsistence consumption a – expression (31) is very close to expression (6) of CR. It is straightforward to deduce that the conclusions we derived in Proposition 1 for the CR model (6) also apply to the discrete-time formulation shown in equation (31). Since these conclusions hold for any Δ and since the continuous-time utility expression in HPSA is the limit of the discrete-time version, the continuous-time expression suffers from the same drawbacks that we discussed in Sections 2.3 and 2.4.

These issues are also visible in the continuation utility that HPSA provide as the starting point of their paper (equation 10 in their paper). Let us use their equation for the case where mortality is the only risk and where consumption and mortality rates are independent of age (a more general analysis can be found in Appendix C). In that case, continuation utility is a constant U, and the distribution of age at death T_m, conditional on being alive at age t, has density function $\lambda_m e^{-\lambda_m (T_m - t)}$. Thus,

[23] Note that the survival probability π_t is not explicitly visible in their paper, but embedded in the certainty equivalent $m_t(\Delta)$ in their equation (C.2).

combining equations (10), (11), and (13) of HPSA, we find that the continuation utility U must fulfill:

$$U = \int_t^{\infty} \lambda_m e^{-\lambda_m (T_m - t)} \left(\int_t^{T_m} \left(\frac{\rho U}{1-\frac{1}{\varepsilon}} \left(((c-a)/U)^{1-\frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m \gamma_m}{1-\gamma_m} U \right) d\tau \right) dT_m.$$

Since there is no dependence on τ inside the integral, this equation simplifies to:

$$\left(\frac{\lambda_m}{1-\gamma_m} - \frac{\rho}{\frac{1}{\varepsilon} - 1} \right) U = \frac{\rho}{1-\frac{1}{\varepsilon}} (c-a)^{1-\frac{1}{\varepsilon}} U^{\frac{1}{\varepsilon}}.$$

In the case where $\varepsilon < 1$, this equation admits a unique (constant) solution $U = 0$ if $\lambda_m \geq \frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}\rho$, and two solutions, $U = 0$ and $U = \left(1 - \frac{\lambda_m(\frac{1}{\varepsilon}-1)}{\rho(1-\gamma_m)}\right)^{\frac{\varepsilon}{1-\varepsilon}} (c-a)$, otherwise. Again, this implies that utility is systematically equal to 0 when mortality is not constrained to be small, and that a non-trivial solution only exists when mortality rates are constrained to be low. Note that the upper bound on the mortality rate to have a non-constant solution is similar to the bound that appears in Proposition 1, except that it is expressed in continuous time.[24]

HPSA circumvent the problem by focusing on the case where the mortality rate remains low enough so that a positive solution exists. Indeed, their first theorem is stated with a condition (equation 24 in their paper) which, when $\varepsilon < 1$ and the only uncertainty at play is mortality, imposes that:

$$\frac{\lambda_m}{1-\gamma_m} < r + \frac{\varepsilon}{1-\varepsilon}\rho. \tag{32}$$

In such a case, their model predicts that excess consumption, $c - a$, grows at a rate equal to $\varepsilon(r - \rho - \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m)$. In order to make the link with our Proposition 1, let us consider the case of a flat optimal consumption (i.e., $r = \rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m$). Then condition (32) can be rewritten as:

$$\lambda_m < \frac{1-\gamma_m}{\frac{1}{\varepsilon} - 1}\rho,$$

which is also the condition we elicited above for the existence of a positive solution

[24] If, for example, we set $\gamma_m = 0.5$, $\varepsilon = \frac{2}{3}$, and $\rho = 0.03$, we get that the hazard rate of death λ_m has to be below 0.03, implying a life expectancy above 33 years.

in the stationary case.

HPSA comment the restriction (32) as being "entirely standard [...] except for the presence of the constant $\frac{\lambda_m}{1-\gamma_m}$ that reflects the combined impact of mortality risk and the agent's aversion to that risk on the optimal consumption schedule" (HPSA, p. 677). It is, however, noteworthy that this restriction imposes an upper bound on mortality rates which restricts the model to perpetually young agents, thereby inhibiting applications with realistic demographic data. Moreover, working with the positive solution of HPSA leads to conclusions that are in line with those discussed in Section 2.4. Indeed, as shown in equation (38) in our Appendix C, when a positive solution exists, the continuation utility at time t is (up to a normalization term) provided by:

$$U_t = \left(\rho \int_t^{\infty} (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_t^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s)\right) ds} d\tau \right)^{\frac{1}{1-\frac{1}{\varepsilon}}}.$$

This utility representation features an instantaneous discount rate at age s equal to $\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)$. Thus, with $\varepsilon < 1$, mortality reduces impatience instead of contributing to it, and consumption and survival display a substitutability property.[25] Moreover, the marginal rate of substitution between survival rate and consumption at time t (i.e., the value of life) is given by:[26]

$$-\frac{\frac{\partial U_t}{\partial \lambda_m(t)}}{\frac{\partial U_t}{\partial c_t(t)}} = \frac{1}{\rho} \frac{1}{1-\gamma_m} (c_t - a)^{\frac{1}{\varepsilon}} U_t^{1-\frac{1}{\varepsilon}}.$$

Since $\varepsilon < 1$ and $\gamma_m < 1$, the formula shows that the willingness-to-pay for mortality risk reduction is decreasing with continuation utility. As in the discrete-time case of CR, we find that the agent who expects to have a miserable life in the future ($c_\tau \simeq a$ for all $\tau > t$) would be willing to pay a lot to survive, but the one who expects to have an extraordinary life in the future ($c_\tau \simeq \infty$ for all $\tau > t$) would not be willing

[25]This is consistent with their analysis which explains that, when the IES is below one, the propensity to consume is decreasing with the mortality rate.

[26]As HPSA use a continuous time model, we use Volterra derivatives (see Ryder and Heal, 1973) to compute the marginal rate of substitution.

to pay anything to increase her survival probability.

To sum up, just like in the discrete-time setup, focusing on the positive solution requires to restrict the analysis to perpetually young agents with unrealistically low mortality rates and delivers conclusions which are at odds with economic intuition and observations.

C HPSA's utility function in the non-stationary case

We prove that the solution to the HPSA recursive formulation is still $U_t = 0$ for all t, when we relax the assumption of a constant mortality and allow for mortality rates to change with age. In the case where the only risk is mortality, HPSA recursive formulation is given by:

$$U_t = 1_{\{T_m > t\}} E_t \int_t^{Tm} \left(\frac{\rho U_\tau}{1 - \frac{1}{\varepsilon}} \left(((c_t - a)/U_\tau)^{1-\frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m(\tau)\gamma_m}{1-\gamma_m} U_\tau \right) d\tau. \tag{33}$$

Note that this specification presupposes that $U_t \geq 0$. Denoting by $\lambda_m(\tau)$ the hazard rate of death at time τ, the distribution of the age at death, T_m, conditional on being alive at age t, has a density function $\lambda_m(T_m)e^{-\int_t^{T_m} \lambda_m(a)da}$. Thus, we find that the continuation utility, U_t, must fulfill:

$$\begin{aligned} U_t = & \int_t^\infty \lambda_m(T_m) e^{-\int_t^{T_m} \lambda_m(s)ds} \\ & \times \left(\int_t^{T_m} \left(\frac{\rho U_\tau}{1 - \frac{1}{\varepsilon}} \left(((c_t - a)/U_\tau)^{1-\frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m(\tau)\gamma_m}{1-\gamma_m} U_\tau \right) d\tau \right) dT_m \end{aligned} \tag{34}$$

or, equivalently, after an integration by parts:

$$U_t = \int_t^\infty e^{-\int_t^\tau \lambda_m(s)ds} \left(\frac{\rho U_\tau}{1 - \frac{1}{\varepsilon}} \left(((c_t - a)/U_\tau)^{1-\frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m(\tau)\gamma_m}{1-\gamma_m} U_\tau \right) d\tau.$$

By differentiation, this gives:

$$\frac{d}{dt} U_t = \left(\frac{\lambda_m(t)}{1-\gamma_m} + \frac{\rho}{1-\frac{1}{\varepsilon}} \right) U_t - \frac{\rho U_t^{\frac{1}{\varepsilon}}}{1 - \frac{1}{\varepsilon}} (c_t - a)^{1-\frac{1}{\varepsilon}}.$$

One solution is $U_t = 0$. The question is about the existence of solutions that are not always equal to zero. Remember that the recursive definition presupposes that $U_t \geq 0$ (otherwise equation 33 does not make sense).

Let us assume therefore that there exists a $t_0 > 0$ such that $U_{t_0} > 0$. We must have $U_t > 0$ on a neighborhood of t_0. On this neighborhood the differential equation can be rewritten as:

$$\frac{d}{dt} U_t^{1-\frac{1}{\varepsilon}} = \left(\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(t) + \rho \right) U_t^{1-\frac{1}{\varepsilon}} - \rho \left(c_t - a \right)^{1-\frac{1}{\varepsilon}}, \tag{35}$$

where the consumption process $(c_t)_{t \geq 0}$ is assumed to be bounded and such that $c_t > a$ for all t.

The Cauchy-Lipschitz theorem implies that the linear differential equation (35) together with the initial condition $U_{t=t_0} = U_{t_0} > 0$ admits a unique solution on a maximal interval of existence containing t_0 in its interior. This maximal interval of existence is of the form $(t_-, t_+) \cap \mathbb{R}^+$, with $t_- \in [-\infty, \infty)$ and $t_+ \in (-\infty, \infty]$. The solution of the linear differential equation (35) is given by:

$$U_t^{1-\frac{1}{\varepsilon}} = e^{\int_{t_0}^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} \Big(U_{t_0}^{1-\frac{1}{\varepsilon}} - \rho \int_{t_0}^t (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_{t_0}^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} d\tau \Big), \tag{36}$$

for all $t \in (t_-, t_+)$, which has to be strictly positive. From equality (36), we deduce that $e^{-\int_{t_0}^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(a) \right) da} U_t^{1-\frac{1}{\varepsilon}}$ is decreasing (when defined). Therefore, $U_t^{1-\frac{1}{\varepsilon}}$ is well-defined and strictly positive for all $t \leq t_0$, and 0 belongs to the maximal interval, which can thus be written as $[0, t_+)$. Without loss of generality, we can rewrite $U_t^{1-\frac{1}{\varepsilon}}$ using the initial condition $U_{t=0} = U_0 > 0$. We have for all $t \in [0, t_+)$:

$$U_t^{1-\frac{1}{\varepsilon}} = e^{\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} \left(U_0^{1-\frac{1}{\varepsilon}} - \rho \int_0^t (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_0^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} d\tau \right). \tag{37}$$

Remember that the issue is whether $t_+ = \infty$ (existence of a non-constant global solution) or not (the only global solution is $U_t = 0$). From now on, we restrict our attention to the case where consumption is bounded from above, and the hazard rate of death $\lambda_m(s)$ is increasing with age s after a given age. There are then two possibilities depending on the value of $\lim_{t\to\infty} e^{-\int_0^t \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds}$.

Case 1. If $\lim_{t\to\infty} e^{-\int_0^t \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} = \infty$ (for instance implied by $\lim_{t\to\infty}\lambda_m(t) > \rho\frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$, remember that $\varepsilon < 1$), since (c_t) is bounded, there exists t_m, such that we have $\rho\int_0^{t_m}(c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_0^\tau \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} d\tau = U_0^{1-\frac{1}{\varepsilon}}$ and $U_{t_m}^{1-\frac{1}{\varepsilon}} = 0$ and for all $t \geq t_m$, $U_t \leq 0$. As a consequence, $t_+ \leq t_m$ and there exists no global strictly positive solution on $\mathbb{R}_+$.

It is noteworthy that from (37), we can deduce that if $\lim_{t\to\infty}\lambda_m(t) > \rho\frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$, the only nonnegative global solution corresponds to $U_0^{1-\frac{1}{\varepsilon}} = \infty$, or equivalently, $U_0 = 0$, which in turn implies $U_t = 0$ for all $t \geq 0$.

Case 2. If $\lim_{t\to\infty} e^{-\int_0^t \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} < \infty$ (for instance implied by $\lim_{t\to\infty}\lambda_m(t) < \rho\frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$), setting $U_0^{1-\frac{1}{\varepsilon}} = K + \rho\int_0^\infty (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_0^\tau \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} d\tau$, with $K \geq 0$, guarantees that $U_t^{1-\frac{1}{\varepsilon}} \geq 0$ for all $t \geq 0$. We have then for all $t \in [0, t_+)$:

$$U_t = \left(K e^{\int_0^t \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} + \rho\int_t^\infty (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_t^\tau \left(\rho+\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}\lambda_m(s)\right)ds} d\tau \right)^{\frac{1}{1-\frac{1}{\varepsilon}}}, \quad (38)$$

which is a global solution on $\mathbb{R}_+$.[27] We can set $t_+ = \infty$. Note that $U_t = 0$ for all $t \geq 0$ is still another global solution on $\mathbb{R}_+$.

[27] As in the CR case, the choice of K can be seen as a normalization with no impact on utility maximization.

References

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

ANDRÉ, E., A. BOMMIER, AND F. LEGRAND (2020): "The Impact of Risk and Ambiguity Aversion on Annuity and Saving Choices," Working Paper, ETH Zurich.

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

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

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

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

BOMMIER, A., D. HARENBERG, F. LEGRAND, AND C. O'DEA (2020): "Recursive Preferences, the Value of Life, and Household Finance," Discussion Paper 2231, Cowles Foundation.

BOMMIER, A., A. KOCHOV, AND F. LEGRAND (2017): "On Monotone Recursive Preferences," *Econometrica*, 85, 1433–1466.

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

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

BOYD, J. H. I. (1990): "Recursive Utility and the Ramsey Problem," *Journal of Economic Theory*, 50, 326–345.

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

CÓRDOBA, J. C., M. RIPOLL, AND S. YANG (2020): "The Full Recession: Private Versus Social Costs of Covid-19," Working Paper, Nankai University.

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

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

GREENSTONE, M. AND V. NIGAM (2020): "Does Social Distancing Matter?" *COVID Economics*, 7, 1–23.

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

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

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

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

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

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

Koopmans, T. C. (1960): "Stationary Ordinal Utility and Impatience," *Econometrica*, 28, 287–309.

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

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

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

Rosen, S. (1981): "Valuing Health Risk," *American Economic Review*, 71, 241–245.

Ryder, H. E. and G. Heal (1973): "Optimal Growth with Intertemporally Dependent Preferences," *Review of Economic Studies*, 40, 1–31.

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

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

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

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