

Recursive Preferences and the Value of Life: A Clarification

Antoine Bommier Daniel Harenberg François Le Grand*

First version: November 4, 2016

This version: September 13, 2018

Abstract

Two recent articles (Córdoba and Ripoll, 2017; Hugonnier, Pelgrin, and St-Amour, 2013) 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, the recursive formulation admits only a unique, constant solution where utility equals zero everywhere. Non-constant solutions may only exist if mortality rates are assumed to remain low at all ages, that is, in a world of perpetually young agents. Such solutions are therefore unsuitable for studying the value of life in realistic settings. In addition, these non-constant solutions exhibit the questionable property that consumption at a given age and survival at that same age are substitutes instead of complements. We conclude this clarifying paper by reviewing various recursive specifications that can be used to study the value of life without facing such problems.

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

JEL codes: G11, J17.

*Bommier: ETH Zurich, abommier@ethz.ch; Harenberg: ETH Zurich, dharenberg@ethz.ch; Le Grand: EMLyon Business School and ETH Zurich, legrand@em-lyon.com. The authors gratefully acknowledge financial support from 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 comment.

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 adopted in other fields such as the economics of climate change, health economics, or household finance. In two recent studies published in the Review of Economic Studies, Hugonnier, Pelgrin, and St-Amour (2013, henceforth HPSA) and Córdoba and Ripoll (2017, henceforth CR) make a strong case in favor of using homothetic recursive preferences to discuss questions related to the value of life. Both argue that the so-called Epstein-Zin-Weil (EZW, henceforth) preferences provide an appealing framework to discuss matters related to the value of life.¹ In particular, they propose a recursive formulation that would yield a utility representation 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. It is noteworthy that in the standard additive model, combining these three properties is only possible when the elasticity of substitution 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 elasticity of substitution is below one, which is usually seen as the empirically more relevant case (see, for example, Havránek 2015).

In the current paper we explain that when the elasticity of substitution is below one, the recursive models proposed by HPSA and CR are not well-behaved. When applied to realistic mortality patterns, these models admit a unique (constant) solution where utility 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 is of course at odds with demographic facts and makes such models completely inadequate for studying the value of life. In addition, these non-constant solutions exhibit the questionable property

¹Note that the two articles differ in their modeling approach, since HPSA consider a continuous time model, whereas CR consider a discrete time model.

that consumption at a given age and survival at that same age are substitutes instead of complements, leading to counterfactual predictions.

One possible conclusion is that the framework introduced by HPSA and CR should be restricted to the case where the elasticity of substitution is larger than one. As we discuss in Section 4, the above-mentioned problems are then avoided and the recursive models of HPSA and CR are well-defined. However, in that case, matching the value of life with homothetic EZW preferences imposes constraints on the risk aversion parameter. This has the drawback of simultaneously imposing a specific relation between the value of life at a specific age and the age profile of the value of life. It is worth noting that the ability to combine homotheticity, an elasticity of substitution greater than one, and a positive value of life is not restricted to the EZW framework since it is, for instance, afforded by the standard additive expected utility model. We conclude the paper with a section that reviews various recursive specifications that can be used to properly study the value of life.

2 Utility functions in CR and HPSA when the elasticity of substitution is smaller than one

2.1 Recursive models

To understand the problem raised by the recursive formulations proposed by CR and HPSA, it is useful to recall 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:

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

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.

Obviously, not all recursions admit a solution and it is therefore important to check that the recursive equations (1) properly define a sequence of utilities $(U_t)_{t \geq 0}$. The easiest

way to derive the existence of the U_t is to assume an exogenous finite horizon $T < \infty$ and an exogenous terminal value U_T . Applying equation (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, and 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, existence of well-behaved solutions is well known and one may work directly with equation (1) to compute first-order conditions and optimal strategies. But checking that there exists a well-behaved solution to the recursive equation is nevertheless central, otherwise conclusions could be seriously flawed. Consider, for example, the simple case with no uncertainty:

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

where z_t denotes time- t consumption, assumed to take values in a compact interval $[z_{min}, z_{max}] \in \mathbb{R}_+$, and σ a scalar assumed to be greater than 1. One 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 degenerate, all consumption strategies are equally good. If one overlooks this, it might seem natural to derive from (2) that

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

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 derive stringent restrictions on optimal consumption plans, while in fact all strategies are optimal. The problem, which only becomes apparent when solving for U_t , is that the marginal utilities in (3) are both equal to zero, their ratio thus being undefined. The reason for mentioning this example is that it shares a number of similarities with the recursive approaches we discuss below.

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

³This result is indirectly related to the findings of Koopmans (1960) stating that time preferences are needed to have a sound recursive model of intertemporal choice in infinite horizon settings.

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 is the one related to mortality and use the same notation as CR. In a first step, CR formalize the problem without constraining the utility of death to a specific value. Denoting by \underline{V} the utility of a dead agent, their recursive model writes as:

$$V_t = \left[z_t^{1-\sigma} + \beta \left(\pi_t V_{t+1}^{1-\gamma} + (1 - \pi_t) \underline{V}^{1-\gamma} \right)^{\frac{1-\sigma}{1-\gamma}} \right]^{\frac{1}{1-\sigma}}, \quad (4)$$

where $z_t > 0$ is the consumption level, π_t is the probability to survive from period t to period $t + 1$, σ is the inverse of the elasticity of substitution, and γ a parameter driving risk aversion.

The problem arises when CR impose the utility of dead agents to be equal to zero, $\underline{V} = 0$. Setting $\underline{V} = 0$ is central to their analysis, since—according to them—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);
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 $\underline{V} = 0$.

Let us therefore investigate the implications of setting $\underline{V} = 0$. CR explain that this requires $\gamma < 1$, so that $\underline{V}^{1-\gamma} = 0$.⁴ With $\underline{V} = 0$ and $\gamma < 1$, equation (4) reduces to

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

⁴Homotheticity could be also be obtained by setting $\gamma > 1$ and $\underline{V}^{1-\gamma} = 0$. This, however, leads to a model where the value of life is always negative, cf. Bommier, Harenberg, and LeGrand (2017a).

But does that recursive equation define a reasonable utility representation? As will be discussed below, the answer depends on whether the elasticity of substitution is below one ($\sigma > 1$) or not. The next section focuses on the case $\sigma > 1$, while Section 4 discusses the case $\sigma < 1$.

2.3 Solutions to Córdoba and Ripoll's recursive model

A simple way to solve the question of existence and uniqueness of the sequence V_t defined by (5) consists of assuming a maximal date T after which survival probabilities are zero. Formally, we have $\pi_T = 0$ and $V_{T+1} = 0$, where the last equality comes from the assumption of the zero utility for death, $\underline{V} = 0$. Now we can compute V_0 by backward induction. First, we compute V_T from equation (5), for all $z_T \geq 0$:

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

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 that $V_t = 0$ for all t , no matter the consumption levels and the mortality rates. The only solution to recursion (5) is therefore the constant zero utility.

The mere assumption of a finite upper bound on life duration is sufficient to generate this zero utility. Of course, one could relax the assumption of a maximal lifetime and approximate the distribution of observed life-durations by an “unlimited” survival pattern, where survival rates become low at large ages ($\lim_{t \rightarrow \infty} \pi_t = 0$). However, this assumption does not solve the issue. 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 $t > t_0$.⁵ Now, let us rewrite equation (5) 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 $t > t_0$:

$$V_t^{1-\sigma} = \beta^{-t} \left(\prod_{j=t_0}^{t-1} \pi_j \right)^{\frac{\sigma-1}{1-\gamma}} \left(V_{t_0}^{1-\sigma} - \sum_{s=t_0}^{t-1} \beta^s \left(\prod_{j=t_0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right). \quad (6)$$

⁵We previously established that if $V_t = 0$, then $V_\tau = 0$ for all $\tau \leq t$.

Consider now the case where $t \rightarrow \infty$, while holding t_0 constant. Assuming that consumption is bounded from above, the right-hand side of (6) becomes negative, since $\sum_{s=0}^{t-1} \beta^s \left(\prod_{j=0}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma}$ diverges (as we are considering the case $\lim_{t \rightarrow \infty} \pi_t = 0$), 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 $\lim_{t \rightarrow \infty} \pi_t = 0$ and, more generally, when $\beta \pi_t^{\frac{1-\sigma}{1-\gamma}} > 1$ for large t .⁶

In their paper, CR suggest that the solution of the recursive equation is:⁷

$$V_t = \left[\sum_{s=t}^{\infty} \beta^{s-t} \left(\prod_{j=t}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (7)$$

This corresponds to the solution of Epstein and Zin (1989) with the assumption that death provides the same utility as zero consumption. However, note that if survival probabilities get small at old ages, the sum in (7) diverges, yielding $V_t = 0$.

If consumption is bounded from above, the only way to have positive and non-constant solutions would be to assume that $\sum_{t=0}^{\infty} \beta^t \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} < \infty$. If one assumes that survival probabilities decrease with age after a given age (as is observed in adulthood), then the positive sequence $(\pi_t)_{t \geq 0}$ must admit a limit $\pi_{\infty} \geq 0$. The condition $\sum_{t=0}^{\infty} \beta^t \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} < \infty$ would require that $1 - \beta \pi_{\infty}^{\frac{1-\sigma}{1-\gamma}} \geq 0$ or, equivalently, $1 - \pi_{\infty} \leq 1 - \beta^{\frac{1-\gamma}{\sigma-1}}$. The validity of the model would be restricted to agents whose mortality rates are not greater than $1 - \beta^{\frac{1-\gamma}{\sigma-1}}$, which seems to be a rather strong limitation. To view it differently, the model would only be valid for agents whose life expectancy never goes below $\frac{1}{1 - \beta^{\frac{1-\gamma}{\sigma-1}}}$, no matter their age.⁸ Consequently, such a model could only consider perpetually young agents, preventing it from accounting for realistic mortality profiles, which is a crucial defect when studying the value of life.

The following proposition summarizes our findings.⁹

⁶A 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 $\gamma < 1 < \sigma$.

⁷See equation (9) in their paper.

⁸Just for illustration, taking $\beta = 0.97$, $\sigma = 1.5$, and $\gamma = 0.5$ implies that the model is only valid for agents whose life expectancy remains above 33 years, independently of their age.

⁹The proposition also applies to the setting of Epstein and Zin (1989), in which the risk of death is

Proposition 1 Consider the utility function defined by the recursion (5) with $\gamma < 1 < \sigma$.

1. If there is a maximal length of life (i.e., there exists T such that $\pi_T = 0$), the only solution to (5) is $V_t = 0$ for all t .
2. When death is never certain and survival probability decreases with age, there are two cases:
 - $\lim_{t \rightarrow \infty} \pi_t < \beta^{\frac{1-\gamma}{\sigma-1}}$: the recursion admits a unique solution $V_t = 0$ for all t ;
 - $\lim_{t \rightarrow \infty} \pi_t \geq \beta^{\frac{1-\gamma}{\sigma-1}}$: the recursion admits multiple solutions: one being $V_t = 0$ for all t , the other ones being given by

$$V_t = \left[\beta^{-t} \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{\sigma-1}{1-\gamma}} K + \sum_{s=t}^{\infty} \beta^{s-t} \left(\prod_{j=t}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} z_s^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (8)$$

for some constant $K \geq 0$.

Working with degenerate solutions, where utility is always equal to zero, is definitely not an interesting option. The non-degenerate solutions of equation (8) only exist when mortality rates are constrained to be low—or equivalently, when life expectancy must remain high at all ages.¹⁰ Besides being unable to cope with realistic mortality data, these non-degenerate solutions have questionable implications, as we discuss in the next section.

2.4 Implications of Córdoba and Ripoll's approach

In this section, we investigate the implications of using equation (7) without taking care of convergence issues.

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 that V_t may be equal to zero because of convergence issues, equation (5) implies

$$\frac{\partial V_t}{\partial z_{t+1}} \bigg/ \frac{\partial V_t}{\partial z_t} = \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} z_{t+1}^{-\sigma} z_t^{\sigma}. \quad (9)$$

replaced by that of having zero consumption forever (i.e., the risk of an irreversible ruin).

¹⁰The constant K is a normalization with no impact on utility maximization.

In this model, when $\gamma < 1 < \sigma$, the survival probability reduces this MRS. Formally:

$$\frac{\partial}{\partial \pi_t} \left(\frac{\partial V_t}{\partial z_{t+1}} / \frac{\partial V_t}{\partial z_t} \right) < 0. \quad (10)$$

Equation (10) 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 would be the natural property, given that there is no bequest in this model, and therefore no “useful consumption” after death. To better understand this feature, let us consider the limit case where $\pi_t \rightarrow 0$ (death is almost sure at the end of period t). In that case, we find that $\frac{\partial V_t}{\partial z_{t+1}} / \frac{\partial V_t}{\partial z_t} \rightarrow \infty$, implying that the agent consumes almost nothing in period t , so as to maximize consumption in period $t + 1$ (independently of whether annuities exist or not). The motivation for her consumption pattern would precisely be that the agent is aware that survival in period $t + 1$ is extremely unlikely.

Optimal life-cycle consumption profiles. Another way to illustrate the previous point is to look at the solution to the life-cycle consumption-saving problem that is derived in Section 3 of CR. To simplify the problem, assume that there is no leisure or health in the analysis, and no annuity market, so that the agent’s program writes as:¹¹

$$\begin{aligned} V_t(a_t, \pi_t) &= \max_{z_t, a_{t+1}} \left(z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}(a_{t+1}, \pi_{t+1})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \\ \text{s.t.} \quad & y_t + a_t = z_t + \frac{1}{1+r} a_{t+1}. \end{aligned} \quad (11)$$

Using first-order conditions, CR derive:¹²

$$\frac{z_{t+1}}{z_t} = \left(\beta(1+r)\pi_t^{\frac{1-\sigma}{1-\gamma}} \right)^{\frac{1}{\sigma}}. \quad (12)$$

CR comment equation (12) by saying 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.¹³ This

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

¹²Cf. equation (15) in their paper.

¹³In 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.

is a consequence of the substitutability exhibited in equation (10) and the unappealing features mentioned above.

Let us have a look at the quantitative implications of equation (12) for consumption using realistic data. For example, set $r = 3\%$, $\beta = \frac{1}{1+r} \simeq 0.97$, $\sigma = 1.5$, $\gamma = 0.5$, and consider two dates t_1 and t_2 with $t_2 > t_1$. Equation (12) implies:

$$\frac{z_{t_2}}{z_{t_1}} = \left(\frac{1}{S(t_1, t_2)} \right)^{\frac{2}{3}}, \quad (13)$$

where $S(t_1, t_2) = \prod_{s=t_1}^{s=t_2-1} \pi_s$ is the probability of being alive at age t_2 , conditional on being alive at age t_1 . In the USA, $S(20, 100) \simeq 1/44$ and $S(20, 110) \simeq 1/9000$.¹⁴ Equation (13) implies that people would consume about 12 times more at age 100 than at age 20 and about 430 times more at age 110 than at age 20, which is clearly counterfactual.

One may develop the analysis further, combining the first-order condition (12) with the budget constraint (11). This yields:

$$z_t = \frac{(1+r)^{\frac{t}{\sigma}} \beta^{\frac{t}{\sigma}} S(0, t)^{\frac{1-\sigma}{\sigma(1-\gamma)}}}{\sum_{\tau=0}^{\infty} (1+r)^{\frac{1-\sigma}{\sigma}\tau} \beta^{\frac{\tau}{\sigma}} S(0, \tau)^{\frac{(1-\sigma)}{\sigma(1-\gamma)}}} \left(\sum_{\tau=0}^{\infty} \frac{y_{\tau}}{(1+r)^{\tau}} \right). \quad (14)$$

But with $\sigma > 1$ and realistic demographic data—such that $S(0, \tau) = 0$ for τ sufficiently large—the denominator is equal to infinity, providing $z_t = 0$ for all finite t . Intuitively, the agent prefers to consume nothing in order to keep all available resources for periods she would never reach ($t = \infty$). These predictions are clearly at odds both with economic logic and with observations.¹⁵

Age-dependent health. In order to circumvent the difficulties raised above, CR propose to introduce an age-dependent health variable H_t and to assume the recursive

¹⁴The mortality data comes from the Human Mortality Database. We use mortality rates for the total US population in 2014.

¹⁵Note that the problem with perfect annuities is not better behaved. Indeed, with perfect annuities and the same parameters as above, people would consume about 150 times more at age 100 than at age 20 and about 190.000 times more at age 110 than at age 20. Equation (14) would become:

$$z_t = \frac{(1+r)^{\frac{t}{\sigma}} \beta^{\frac{t}{\sigma}} S(0, t)^{\frac{\gamma-\sigma}{\sigma(1-\gamma)}}}{\sum_{\tau=0}^{\infty} (1+r)^{\frac{1-\sigma}{\sigma}\tau} \beta^{\frac{\tau}{\sigma}} S(0, \tau)^{\frac{\gamma(1-\sigma)}{\sigma(1-\gamma)}}} \left(\sum_{\tau=0}^{\infty} S(0, \tau) \frac{y_{\tau}}{(1+r)^{\tau}} \right),$$

which also provides $z_t = 0$ for all t .

equation:

$$V_t = \left[H_t z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (15)$$

Since H_t depends on t , there are sufficiently many degrees of freedom to match any possible consumption profile. In principle, a rapidly declining profile H_t could potentially rectify the defects discussed above. There are, however, two fundamental problems:

1. First, with $\sigma > 1$, the specification (15) assumes that utility *decreases* (instead of increases) with health. This is of course completely implausible.¹⁶
2. CR's calibration procedure (see Section 4.2.1 of their paper) leads to implicitly setting $H_t = \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{\sigma-\gamma}{1-\gamma}} H_t^{MT}$, where H_t^{MT} is the health profile used by Murphy and Topel (2006). The assumption that H_t is exogenous and independent from survival probabilities π_t is thus highly questionable.

The question of health endogeneity is crucial for understanding the impact of (exogenous or endogenous) mortality changes. Compare, for example, the utility functions in CR's model and Murphy and Topel's model:

$$V_0 = \left[\sum_{t=0}^{\infty} \beta^t \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} H_t z_t^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \quad (\text{CR's model}),$$

$$V_0^{MT} = \left[\sum_{t=0}^{\infty} \beta^t \left(\prod_{j=0}^{t-1} \pi_j \right) (H_t^{MT} z_t^{1-\sigma} + u_0) \right]^{\frac{1}{1-\sigma}} \quad (\text{Murphy and Topel's model}).$$

When $H_t = \left(\prod_{j=0}^{t-1} \pi_j \right)^{\frac{\sigma-\gamma}{1-\gamma}} H_t^{MT}$, both models have exactly the same implications when mortality is fixed and exogenous. Moreover, with a proper calibration of u_0 , they can match the same average VSL between ages 25 and 55, which is the calibration target used in CR. However, the models lead to radically different conclusions regarding the consequences of mortality decline. Murphy and Topel's model predicts that mortality decline significantly increases the propensity to save (agents become more patient when survival probabilities increase), while the effect is much smaller or even opposite in CR's

¹⁶Another option would be to consider $V_t = \left[h_t^{1-\sigma} z_t^{1-\sigma} + \beta \pi_t^{\frac{1-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}$ and to interpret h_t as a health index. In such a case, h_t would positively contribute to utility, but matching realistic consumption profiles would require that this health profile h_t increases and converges to $+\infty$ at old ages. Again, this does not seem to be an appealing modeling choice.

model (agents become more impatient when survival probabilities increase).¹⁷ Opting for one or the other specification will then provide very different views about the impact of population aging. While both models match the same calibration targets, the implicit assumption they make about the role of health (considered as a “good” in Murphy and Topel and as a “bad” in CR) ends up having significant consequences.

Willingness to pay for mortality risk reduction. Another questionable behavioral pattern emerges when computing the MRS between survival and consumption, i.e., the value of life. From equation (5), we obtain:

$$\frac{\partial V_t}{\partial \pi_t} \bigg/ \frac{\partial V_t}{\partial z_t} = \frac{1}{1-\gamma} z_t^\sigma \beta \pi_t^{\frac{\gamma-\sigma}{1-\gamma}} V_{t+1}^{1-\sigma}. \quad (16)$$

With $\gamma < 1 < \sigma$, the value of life is decreasing in the continuation utility V_{t+1} :

$$\frac{\partial}{\partial V_{t+1}} \left(\frac{\partial V_t}{\partial \pi_t} \bigg/ \frac{\partial V_t}{\partial z_t} \right) < 0. \quad (17)$$

This means that the greater the possible future loss (measured by continuation utility V_{t+1}), the less the agent is 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 ($V_{t+1} = 0$), while this willingness-to-pay is zero when she knows that she will consume huge amounts and have a superb life in case of survival ($V_{t+1} = \infty$). Again, this is at odds with standard economic theory, where the willingness-to-pay for a risk reduction rises with the utility loss induced by the risk.

2.5 Hugonnier, Pelgrin and St-Amour (2013)’s 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 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 elasticity of substitution 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

¹⁷Longevity extension also generates an income effect that adds to the impatience effect we emphasize, which explains why the overall impact can be ambiguous.

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

$$U_t = \left[(1 - e^{-\rho\Delta})(c_t - a)^{1-\frac{1}{\varepsilon}} + e^{-\rho\Delta}\pi_t^{\frac{1-\frac{1}{\varepsilon}}{1-\gamma_m}} U_{t+\Delta}^{1-\frac{1}{\varepsilon}} \right]^{\frac{1}{1-\frac{1}{\varepsilon}}}, \quad (18)$$

where $\varepsilon > 0$ is the elasticity of substitution, $0 \leq \gamma_m < 1$ 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.¹⁸ We consider again the case where the elasticity of substitution is below 1 ($\varepsilon < 1$). Note that—apart from notation and subsistence consumption a —expression (18) is very close to expression (5) of CR. It is straightforward to deduce that the conclusions we derived in Proposition 1 for the CR model (5) also apply to the discrete-time formulation shown in equation (18). 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 shortcomings 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 our Appendix A). 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, 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.

¹⁸Note 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).

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.¹⁹

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:

$$\frac{\lambda_m}{1 - \gamma_m} < r + \frac{\varepsilon}{1 - \varepsilon} \rho. \quad (19)$$

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 (19) 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 in the stationary case.

HPSA comment the restriction (19) 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 (39) in our Appendix A, when a positive solution exists, the continuation utility at time t is (up to a normalization term) provided by:

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

¹⁹If, 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.

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.²⁰ Moreover, the marginal rate of substitution between survival rate and consumption at time t (i.e., the value of life) is given by:²¹

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

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 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.

3 Fixing the problem by considering the limit $\underline{V} \rightarrow 0$?

A way to (partially) fix the problems discussed in Section 2 involves using the non-homothetic recursion (4) and to take the limit where $\underline{V} \rightarrow 0$. For any $\underline{V} > 0$, if one assumes that there is a finite maximal length of life (which seems quite reasonable), the recursion (4) provides a well-defined and non-degenerate sequence of utilities $(V_t)_t$, from which one can compute optimal life-cycle consumption profiles and the VSL. Looking at the limit where $\underline{V} \rightarrow 0$ (if it exists) may then provide a robust theoretical foundation to the case where $\underline{V} = 0$. We explore this possibility below, focusing again here on the case where $\gamma < 1 < \sigma$ for which the model of CR was shown to be degenerate when $\underline{V} = 0$.

²⁰This is consistent with their analysis which explains that, when the elasticity of substitution is below one, the propensity to consume is decreasing with the mortality rate.

²¹As HPSA use a continuous time model, we use Volterra derivatives (see Ryder and Heal, 1973) to compute the marginal rate of substitution.

Formally, starting from the recursion (4), we renormalize the utility representation by setting $W_t = V_t/\underline{V}$. The utilities $(W_t)_t$ fulfill the recursion:

$$W_t = \left[\left(\frac{z_t}{\underline{V}} \right)^{1-\sigma} + \beta \left(\pi_t W_{t+1}^{1-\gamma} + 1 - \pi_t \right)^{\frac{1-\sigma}{1-\gamma}} \right]^{\frac{1}{1-\sigma}}. \quad (20)$$

After some algebra, we derive from (20) the following marginal rates of substitution:

$$\frac{\frac{\partial W_t}{\partial z_{t+1}}}{\frac{\partial W_t}{\partial z_t}} = \beta \pi_t \left(\pi_t + (1 - \pi_t) W_{t+1}^{\gamma-1} \right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t} \right)^{-\sigma}, \quad (21)$$

and

$$\frac{\frac{\partial W_t}{\partial \pi_t}}{\frac{\partial W_t}{\partial z_t}} = \underline{V}^{1-\sigma} \frac{\beta}{1-\gamma} z_t^\sigma \left(W_{t+1}^{1-\gamma} - 1 \right) \left(\pi_t W_{t+1}^{1-\gamma} + 1 - \pi_t \right)^{\frac{\gamma-\sigma}{1-\gamma}}. \quad (22)$$

Since $\sigma > 1$, we have $\lim_{\underline{V} \rightarrow 0} \left(\frac{z_t}{\underline{V}} \right)^{1-\sigma} = 0$. Equation (20) therefore implies that $\lim_{\underline{V} \rightarrow 0} W_t = \chi_t$ where $(\chi_t)_{t \geq 0}$ is a sequence of numbers in $[\beta^{\frac{1}{1-\sigma}}, +\infty)$ defined by:

$$\chi_t = \beta^{\frac{1}{1-\sigma}} \left(\pi_t \chi_{t+1}^{1-\gamma} + 1 - \pi_t \right)^{\frac{1}{1-\gamma}}. \quad (23)$$

Plugging $\lim_{\underline{V} \rightarrow 0} W_t = \chi_t$ into equation (21), we obtain:

$$\lim_{\underline{V} \rightarrow 0} \left(\frac{\frac{\partial W_t}{\partial z_{t+1}}}{\frac{\partial W_t}{\partial z_t}} \right) = \beta \pi_t \left(\pi_t + (1 - \pi_t) \chi_{t+1}^{\gamma-1} \right)^{\frac{\gamma-\sigma}{1-\gamma}} \left(\frac{z_{t+1}}{z_t} \right)^{-\sigma}, \quad (24)$$

while equation (22) leads to $\lim_{\underline{V} \rightarrow 0} \frac{\frac{\partial W_t}{\partial \pi_t}}{\frac{\partial W_t}{\partial z_t}} = +\infty$.

The limit when $\underline{V} \rightarrow 0$ corresponds to a setting where the VSL tends to infinity, while the MRS (marginal rate of substitution between consumption in period $t+1$ and consumption in period t) converges to a well-defined limit that depends on the survival pattern and preference parameters. What is interesting though, is that the equation (24) characterizing the MRS for the non-recursive model (4) when $\underline{V} \rightarrow 0$ looks very different from the MRS expression of equation (9), derived from the problematic recursion (5) and used by CR to compute optimal consumption profiles. In particular, when π_t gets very small, equation (24) indicates that the MRS $\frac{\partial W_t}{\partial z_{t+1}} / \frac{\partial W_t}{\partial z_t}$ also becomes very small, reflecting that an agent cares little for consumption in periods she will almost surely never see. This has to be contrasted with the result implied by equation (9), which states that instead of getting small the MRS becomes infinitely large.

In fact, one can easily check that the numerical solutions obtained when using equation (24) provide reasonable consumption patterns instead of the implausible findings discussed in Section 2.4. Therefore, properly taking the limit $\underline{V} \rightarrow 0$ of the non-homothetic model solves part of the problem raised by recursion (5). We need however to emphasize three points.

1. Although the limit case $\underline{V} \rightarrow 0$ may be used to find optimal life-cycle consumption profiles under exogenous mortality patterns,²² it is inadequate to discuss value of life issues. Indeed at the limit $\underline{V} \rightarrow 0$, the value of life tends to infinity, meaning that agents would be willing to spend all their wealth to improve their longevity.
2. The whole analysis in CR and HPSA is based on equation (9)—or on its continuous-time version in HPSA—and not on (24). This is why they have to introduce some unrealistic assumptions (such that an upper bound on mortality rates in HPSA or an ad-hoc health profile which reduces utility in CR) for their results to look not too absurd. Using (24) instead of (9) would totally change their analyses.
3. The contrast between the limit case $\underline{V} \rightarrow 0$ and the approach adopted by CR to represent the case $\underline{V} = 0$ is completely overlooked in their paper. In fact, in the Section 5.1 of CR it is claimed that taking the limit $\underline{V} \rightarrow 0$ brings us back to the approach they use for $\underline{V} = 0$. This is because it is mistakenly taken for granted that if $\underline{V} \rightarrow 0$ then the ratio $\frac{V}{\underline{V}_{t+1}}$ that appears in the Euler equation provided in their Section 5.1 also converges towards zero. As is explained above, however, V_{t+1} also tends to zero when $\underline{V} \rightarrow 0$ and, instead of converging towards zero, the ratio $\frac{V_{t+1}}{\underline{V}}$ converges towards χ_{t+1} , where χ_{t+1} is defined by (23).²³

²²This is similar to the first-order approximation suggested in Section 4.3 of Bommier 2006.

²³One can also check that when taking the limit $\underline{V} \rightarrow 0$ properly (i.e., when taking into account that the ratio $\frac{V_{t+1}}{\underline{V}}$ converges towards χ_{t+1} and not toward zero) the VSL computed in the Section 5.1 of CR converges to $+\infty$. Indeed, in the perpetual youth model that is considered, when $\underline{V} \rightarrow 0$, one has $\pi + (1 - \pi)(\underline{V}/V)^{1-\gamma} \rightarrow \beta^{\frac{1-\gamma}{\sigma-1}}$. Moreover, to maintain the assumption of a flat consumption, one has to further assume that $(1 + r) \rightarrow \beta^{\frac{1-\gamma}{\sigma-1}}$. Therefore the denominator in the formula that provides the VSL gets infinitesimally small when $\underline{V} \rightarrow 0$, implying that the VSL converges to $+\infty$.

4 Utility functions in CR and HPSA when the elasticity of substitution is greater than one

Another possibility to overcome the problem discussed in Section 2, while keeping $\underline{V} = 0$, is to restrict the analysis to an elasticity of substitution greater than one. In this case, the recursive specification (5) admits a unique non-trivial solution given by (7). The utility function features homotheticity as well as a value of life that is always finite and positive.

Setting $\underline{V} = 0$ involves assuming that consuming zero is as bad as death. A potentially problematic consequence of this constraint is that the value of life is pinned down by preference parameters (time discounting parameter, elasticity of substitution, and risk aversion coefficient) that are usually calibrated to reflect other preference traits.

Before getting to CR's specification, it is insightful to look at the standard additive homothetic model where utility U_t at date t can be expressed as

$$U_t = \frac{z_t^{1-\sigma}}{1-\sigma} + \beta\pi_t U_{t+1}. \quad (25)$$

When the consumption path is constant and equal to z over the life-cycle, we obtain that the MRS between survival and consumption is given by

$$\frac{\partial U_t}{\partial \pi_t} \bigg/ \frac{\partial U_t}{\partial z_t} = z \times \frac{1}{1-\sigma} \sum_{s=t+1}^{\infty} \beta^{s-t} \left(\prod_{j=t+1}^{s-1} \pi_j \right). \quad (26)$$

The willingness-to-pay for mortality risk reduction (i.e., the value of life) is thus proportional to z (a consequence of preference homotheticity) and to a term $e_t = \sum_{k=0}^{\infty} \beta^k \left(\prod_{j=t+1}^{t+k} \pi_j \right)$ which can be interpreted as an adjusted demographic factor. With $\beta = 1$, e_t is the life expectancy at age t (roughly 36 years at age 45). With $\beta = 0.97$, e_t would be approximately equal to 21 at age 45. Empirical estimates, mostly derived from US data, tend to indicate that a reasonable order of magnitude for the ratio $\frac{\partial V_t}{\partial \pi_t} \bigg/ \frac{\partial V_t}{\partial z_t}$ is 300.²⁴ For the theoretical result shown in equation (26) to match the empirical estimates, one has to play with the parameter σ , which is constrained to be smaller than 1. The closer σ gets to one, the larger the ratio of the value of life over consumption. Calibrating the model to fit empirical estimates of the value of a statistical life therefore consists in choosing the elasticity of

²⁴This would imply a value of a statistical life of 6 million USD for people consuming 20 000 USD per year. Such numbers are roughly in line with those discussed in Viscusi and Aldy (2003).

substitution. Given the empirical consensus about the elasticity of substitution, this raises the question whether a single parameter can simultaneously match the aversion for fluctuations in consumption over time and the value of life.

The above problem is not related to the lack of flexibility of the standard additive model, i.e., to the fact that elasticity of substitution and risk aversion coefficients are intertwined. Indeed, a similar issue appears when considering EZW preferences. When focusing on a constant consumption path z , equation (16) implies that the MRS between survival and consumption becomes

$$\frac{\partial V_t}{\partial \pi_t} \bigg/ \frac{\partial V_t}{\partial z_t} = z \times \frac{\pi_t^{\frac{\gamma-\sigma}{1-\gamma}}}{1-\gamma} \left[\sum_{s=t+1}^{\infty} \beta^{s-t} \left(\prod_{j=t+1}^{s-1} \pi_j \right)^{\frac{1-\sigma}{1-\gamma}} \right], \quad (27)$$

which is similar to expression (26), except that survival probabilities are transformed and that the factor $\frac{1}{1-\sigma}$ is now replaced by $\frac{1}{1-\gamma}$. From equation (27), we see that fitting empirical estimates of the value of life requires adjusting the degree of risk aversion—which has to be below one. Again a single parameter (γ) is pinning down two different aspects of preferences: the degree of risk aversion—generally calibrated by looking at portfolio choices—and the value of life.

CR suggest to solve the problem by using different coefficients of risk aversion for each risk and, in particular, a specific coefficient of risk aversion for mortality risk and another one for financial risk. This is also the approach suggested in the continuous-time setup of HPSA. Working in that direction amounts to, first, artificially constraining preferences over deterministic outcomes by requiring the utility of death to be the same as that of zero consumption and, second, compensating the previous constraint by considering several specific risk aversion coefficients. It is then possible to fit the value of life (provided that the elasticity of substitution is above one and the coefficient of mortality risk aversion is below one). However, Bommier and Villeneuve (2012) emphasize that mortality risk aversion is key to determine how the value of life changes with age. Fitting the value of life by the risk aversion parameter therefore imposes a specific relationship between the value of life (at a specific age) and the age-profile of the value of life.

5 Studying the value of life with recursive models

The conclusion of our discussion should not be that recursive models are inadequate to study the value of life. Recursive methods provide possibly excellent tools to study the value of life and intertemporal choice under uncertain lifetime. They just need to be carefully implemented, accounting for a couple of specific issues as discussed below.

5.1 An infinite horizon for finitely-lived agents

Recursive models usually assume an infinite horizon or a fixed finite horizon. However, modeling choices under uncertain lifetime requires considering lives of different lengths, which is fundamentally different from the infinite or fixed horizon setting. The difference may be circumvented in a simple way. Instead of describing a life as a finite sequence of per-period consumption choices, one may view a life as an infinite sequence of per-period “realizations”, where each realization is either “being in the death state” 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 the “death state”. A life is then represented by a sequence of the form:

$$(c_0, \dots, c_T, d, d, \dots) \in (\mathbb{R}_+ \cup \{d\})^\infty, \quad (28)$$

for some T . The sequence (28) corresponds to a life duration until age T with consumption profile $(c_0, \dots, c_T) \in \mathbb{R}_+^{T+1}$. Formally, if one excludes resurrection and immortality, the set of possible lives, denoted by L , is defined as follows:

$$L = \{(z_t)_{t \geq 0} \in (\mathbb{R}_+ \cup \{d\})^\infty : \exists T \leq T_{max}, \forall t \geq T, z_t = d \text{ and } \forall t < T, z_t \in \mathbb{R}_+\}, \quad (29)$$

where $T_{max} \geq 0$ is the maximum age of death. Since the upper bound T_{max} can be any positive number, we do not consider it a serious limitation for the modeling of human beings. The set L being a subset of an infinite product space, recursive methods can be used in the standard way. One has, however, to keep in mind that the underlying product space is not \mathbb{R}_+^∞ , as is usually the case, but $(\mathbb{R}_+ \cup \{d\})^\infty$. This means that instantaneous 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 will raise some (minor) technical issues.

5.2 Recursive models with mortality

Nearly all applied studies using recursive methods assume a parametric form that can be expressed as:

$$V_t = f^{-1}((1 - \beta)u(z_t) + \beta\phi^{-1}(E[\phi f(V_{t+1})])), \quad (30)$$

where $\beta \in (0, 1)$ is a time preference parameter, ϕ an increasing function representing risk preferences, $E[\cdot]$ an expectation operator, and $u : \mathbb{R}_+ \cup \{d\} \rightarrow \mathbb{R} \cup \{-\infty, \infty\}$ is a function representing periodic preferences.²⁵ 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 one may as well use $f(x) = \phi^{-1}(x)$, as in Kreps and Porteus (1978), or simply $f(x) = x$, which we will use in the following. The assumption of an upper bound on possible life durations ($T_{max} < \infty$) avoids issues related to existence or multiplicity.

We now discuss how to parametrize further the functions u and ϕ . It is important to note that, for the model to be well-defined, the function ϕ must have a domain that includes not only $Im(u) = u(\mathbb{R}_+ \cup \{d\})$ but 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 proves useful to consider conditional utilities, which are the utilities obtained conditionally on being dead or on being alive. In case of a dead agent, we simply get $V_t(d, d, \dots) = u(d)$, which is the value of the periodic utility function, u , in the death state, d . Plugging that into (30), we obtain that the utilities of alive agents, denoted by U_t , are linked through the following recursion:

$$U_t = (1 - \beta)u(c_t) + \beta\phi^{-1}(\pi_t E[\phi(U_{t+1})] + (1 - \pi_t)\phi(u(d))), \quad (31)$$

where $c_t \in \mathbb{R}_+$ is the consumption at time t and π_t the survival probability between dates t and $t + 1$. Note that—with a slight abuse of notation—we still denote the expectation operator by $E[\cdot]$, even though the mortality risk is now treated separately.²⁶

²⁵Cases where the function u may take infinite values ($-\infty$ or ∞) are perfectly possible if the function ϕ is defined for these.

²⁶Treating mortality separately is possible as long as mortality is independent of other risks.

5.3 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.

5.3.1 The period utility function u

A common specification is the case where preferences exhibit a constant elasticity of substitution. Formally, this means that $\frac{-cu'(c)}{u''(c)}$ is independent of c or equivalently, after integration, that for $c \in \mathbb{R}_+$, we have $u(c) = K \frac{c^{1-\sigma}}{1-\sigma} + u_0$, where σ is the inverse of the elasticity of substitution and K and u_0 are two constants. The function u , defined over $\mathbb{R}_+ \cup \{d\}$, is then:

$$\begin{cases} u(c) = K \frac{c^{1-\sigma}}{1-\sigma} + u_0 \text{ for } c \in \mathbb{R}_+, \\ u(d) = u_d, \end{cases} \quad (32)$$

for some $K > 0$, $u_0 \in \mathbb{R}$, 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 $u_0 = 0$ or $u_d = 0$. But imposing an additional relation between u_0 and u_d (such as $u_0 = u_d$) goes beyond a mere normalization. In particular, this would constrain the value of life. In other words, one can freely constrain u_0 , or alternatively u_d , but constraining both simultaneously is not without loss of generality.

5.3.2 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.²⁷ 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 the specification (32) for the period utility function u .

²⁷The 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).

Affine ϕ . Defining ϕ as an affine function is probably the most common choice in the literature on the value of life. It is most often associated with the normalization $u_d = 0$, but renormalizing u by adding the same constant to both u_0 and u_d , while keeping ϕ unchanged, would have no impact on preferences. With $u_d = 0$ the recursion (31) defining the utility function of an alive agent can be expressed as follows:

$$U_t = (1 - \beta) \left(\frac{c_t^{1-\sigma}}{1-\sigma} + u_0 \right) + \beta \pi_t (E[U_{t+1}]).$$

Imposing, in addition, that $u_0 = 0$ might seem appealing for tractability reasons, but the value of life would then be pinned down by β and σ mostly (see the discussion in Section 4). Actually, if $\sigma > 1$ and $u_0 = u_d = 0$, the value of life is always negative, independently of the consumption level, which is a very undesirable property when studying the value of life.

A well-known drawback of the affine specification is that it imposes risk aversion and elasticity of substitution to be intertwined.

Isoelastic ϕ . This corresponds to EZW preferences. A difficulty that arises with isoelastic functions is related 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 $\frac{c_t^{1-\sigma}}{1-\sigma} + u_0$ and u_d always have the same sign. This implies that u_0 and u_d must have the same sign as $1 - \sigma$. All the studies using isoelastic functions ϕ that we are aware of set $u_0 = 0$.²⁸ Denoting $\epsilon = \text{sign}(1 - \sigma)$, the recursion (31) defining the utility function U_t of an alive agent can be expressed as follows:

$$U_t = (1 - \beta) \frac{c_t^{1-\sigma}}{1-\sigma} + \epsilon \beta \left(\pi_t E[(\epsilon U_{t+1})^{1-\alpha}] + (1 - \pi_t) (\epsilon u_d)^{1-\alpha} \right)^{\frac{1}{1-\alpha}}. \quad (33)$$

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's specification involves setting $u_d = -\infty$ if $\sigma > 1$ and $u_d = 0$ if $\sigma < 1$.²⁹ But setting $u_d = 0$ or $u_d = \epsilon \infty$ represents a severe restriction on preferences. Indeed, apart

²⁸An exception is our companion paper Bommier, Harenberg, and LeGrand (2017a).

²⁹Remember that CR use a different normalization, so that what they denote \underline{V} is related to u_d by $u_d = \frac{V^{1-\sigma}}{1-\sigma}$.

from the case where $\sigma < 1$, $\alpha < 1$, and $u_d = 0$, this results in models that either imply constant utilities (independent of consumption choice) or exhibit a value of life that is always negative.

Using an isoelastic function ϕ makes it possible to disentangle risk aversion and elasticity of substitution but has three drawbacks. First, it constrains the parameter u_d to be of a given sign (for domain reasons), which imposes a restriction on the value of life. Second, since the restriction depends on the sign of $(1 - \sigma)$, these models are often specified such that discontinuities arise when the elasticity of substitution varies from values below one to values above one. The case where the elasticity of substitution is equal to one is generally not considered. Last, and this is not related to the possibility of death, recursive preferences with isoelastic functions ϕ are not monotone with respect to first order stochastic dominance.³⁰

Exponential ϕ . As shown in Bommier, Kochov, and LeGrand (2017b), assuming that preferences are monotone with respect to first order stochastic dominance implies that the function ϕ has an exponential form, with $\phi : x \mapsto \frac{1-e^{-kx}}{k}$, $k \neq 0$,³¹ yielding the 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 lifetime in Bommier (2014) and Bommier, Harenberg, and LeGrand (2017a).³² Since the function $x \mapsto \frac{1-e^{-kx}}{k}$ 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 $u(\mathbb{R}_+) \cup u_d$, no matter the choice of u and u_d . Moreover, one can easily check that, when ϕ is exponential, re-normalizing u by adding the same constant to both u_0 and u_d , while keeping ϕ unchanged, has no impact on preferences. One can therefore set $u_d = 0$ without

³⁰When preferences are not monotone, agents may choose an action while a different action would provide better outcomes in all circumstances. This is similar to choosing dominated strategies in game theory. An example of dominated saving behavior with Epstein-Zin-Weil preferences can be found in Bommier, Kochov, and LeGrand (2017b).

³¹The case $k = 0$ corresponds to an affine ϕ .

³²The “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 respectively study the value of life and the demand for annuities.

loss of generality. The utility U_t is then obtained by the following recursion:

$$U_t = \frac{c_t^{1-\sigma}}{1-\sigma} + u_0 - \frac{\beta}{k} \log(\pi_t E[e^{-kU_{t+1}}] + 1 - \pi_t),$$

where the parameter k drives risk aversion. The value of life can then be calibrated by choosing u_0 .

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_0 that respectively determine the elasticity of substitution, risk aversion, time preferences, and the utility gap between life and death (and thereby the value of life).

The main drawback of the risk-sensitive specification is that it is generally not homothetic. But—as we have argued in this clarification paper—homotheticity needs to be given up anyway if one wants to study the value of life in a recursive utility model with realistic mortality rates and an elasticity of substitution below one.

Appendix

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

In the case where the only risk is mortality, HPSA's recursive formulation is given by:

$$U_t = 1_{\{T_m > t\}} E_t \int_t^{T_m} \left(\frac{\rho U_\tau}{1 - \frac{1}{\varepsilon}} \left(((c_t - a)/U_\tau)^{1 - \frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m(\tau) \gamma_m U_\tau}{1 - \gamma_m} \right) d\tau. \quad (34)$$

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 $\lambda_m(T_m) e^{-\int_t^{T_m} \lambda_m(s) ds}$. Thus, we find that the continuation utility, U_t , must fulfill:

$$U_t = \int_t^\infty \lambda_m(T_m) e^{-\int_t^{T_m} \lambda_m(s) ds} \left(\int_t^{T_m} \left(\frac{\rho U_\tau}{1 - \frac{1}{\varepsilon}} \left(((c_t - a)/U_\tau)^{1 - \frac{1}{\varepsilon}} - 1 \right) - \frac{\lambda_m(\tau) \gamma_m U_\tau}{1 - \gamma_m} \right) d\tau \right) dT_m \quad (35)$$

or, equivalently, after an integration by parts:

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

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 34 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:

$$\frac{d}{dt} U_t^{1 - \frac{1}{\varepsilon}} = \left(\frac{1 - \frac{1}{\varepsilon}}{1 - \gamma_m} \lambda_m(t) + \rho \right) U_t^{1 - \frac{1}{\varepsilon}} - \rho (c_t - a)^{1 - \frac{1}{\varepsilon}}, \quad (36)$$

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 (36) 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 (36) is given by:

$$\forall t \in (t_-, t_+), U_t^{1-\frac{1}{\varepsilon}} = e^{\int_{t_0}^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} \left(U_{t_0}^{1-\frac{1}{\varepsilon}} - \rho \int_{t_0}^t (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_{t_0}^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} d\tau \right), \quad (37)$$

which has to be strictly positive. From equality (37), 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:

$$\forall t \in [0, t_+), U_t^{1-\frac{1}{\varepsilon}} = e^{\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} \left(U_0^{1-\frac{1}{\varepsilon}} - \rho \int_0^t (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_0^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} d\tau \right). \quad (38)$$

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 $\lim_{t \rightarrow \infty} e^{-\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds}$.

Case 1. If $\lim_{t \rightarrow \infty} e^{-\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} = \infty$ (for instance implied by $\lim_{t \rightarrow \infty} \lambda_m(t) > \rho \frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$, 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 (38), we can deduce that if $\lim_{t \rightarrow \infty} \lambda_m(t) > \rho \frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$, 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 $\lim_{t \rightarrow \infty} e^{-\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} < \infty$ (for instance implied by $\lim_{t \rightarrow \infty} \lambda_m(t) < \rho \frac{1-\gamma_m}{\frac{1}{\varepsilon}-1}$), 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 $K \geq 0$, guar-

antees that $U_t^{1-\frac{1}{\varepsilon}} \geq 0$ for all $t \geq 0$. We have then:

$$\forall t \in [0, t_+), U_t = \left(K e^{\int_0^t \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} + \rho \int_t^\infty (c_\tau - a)^{1-\frac{1}{\varepsilon}} e^{-\int_t^\tau \left(\rho + \frac{1-\frac{1}{\varepsilon}}{1-\gamma_m} \lambda_m(s) \right) ds} d\tau \right)^{\frac{1}{1-\frac{1}{\varepsilon}}}, \quad (39)$$

which is a global solution on \mathbb{R}_+ .³³ We can set $t_+ = \infty$. Note that $U_t = 0$ for all $t \geq 0$ is still another global solution on \mathbb{R}_+ .

References

- Antoine Bommier. Uncertain Lifetime and Intertemporal Choice: Risk Aversion as a Rationale for Time Discounting. *International Economic Review*, 47(4):1223–1246, November 2006.
- Antoine Bommier. Life Cycle Preferences Revisited. *Journal of European Economic Association*, 11(6):1290–1319, December 2013.
- Antoine Bommier. Mortality Decline, Impatience and Aggregate Wealth Accumulation with Risk-Sensitive Preferences. CER-ETH Working Paper 14/194, ETH Zurich, 2014.
- Antoine Bommier and François LeGrand. Too Risk Averse to Purchase Insurance? A Theoretical Glance at the Annuity Puzzle. *Journal of Risk and Uncertainty*, 48(2):135–166, April 2014.
- Antoine Bommier and Bertrand Villeneuve. Risk Aversion and the Value of Risk to Life. *Journal of Risk and Insurance*, 79(1):77–103, March 2012.
- Antoine Bommier, Daniel Harenberg, and François LeGrand. Recursive Preferences, Household Finance, and the Value of Life. Working paper, ETH Zurich, 2017a.
- Antoine Bommier, Asen Kochov, and François LeGrand. On Monotone Recursive Preferences. *Econometrica*, 85(5):1433–1466, September 2017b.
- John H. Boyd III. Recursive Utility and the Ramsey Problem. *Journal of Economic Theory*, 50(2):326–345, April 1990.

³³As in the CR case, the choice of K can be seen as a normalization with no impact on utility maximization.

- Juan Carlos Córdoba and Maria Ripoll. Risk Aversion and the Value of Life. *Review of Economic Studies*, 84(4):1472–1509, October 2017.
- Larry G. Epstein and Stanley E. Zin. Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica*, 57(4):937–969, July 1989.
- Lars Peter Hansen and Thomas J. Sargent. Discounted Linear Exponential Quadratic Gaussian Control. *IEEE Transactions on Automatic Control*, 40(5):968–971, May 1995.
- Thomáš Havránek. Measuring Intertemporal Substitution: The Importance of Method Choices and Selective Reporting. *Journal of the European Economic Association*, 13(6): 1180–1204, December 2015.
- Julien Hugonnier, Florian Pelgrin, and Pascal St-Amour. Health and (Other) Asset Holdings. *Review of Economic Studies*, 80(2):663–710, October 2013.
- Tjalling C. Koopmans. Stationary Ordinal Utility and Impatience. *Econometrica*, 28(2): 287–309, April 1960.
- David M. Kreps and Evan L. Porteus. Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica*, 46(1):185–200, January 1978.
- Kevin M. Murphy and Robert H. Topel. The Value of Health and Longevity. *Journal of Political Economy*, 114(5):871–904, October 2006.
- Sherwin Rosen. Valuing Health Risk. *American Economic Review*, 71(2):241–245, May 1981.
- Harl E. Ryder and Geoffrey Heal. Optimal Growth with Intertemporally Dependent Preferences. *Review of Economic Studies*, 40(1):1–31, January 1973.
- W. Kip Viscusi and Joseph E. Aldy. The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World. *Journal of Risk and Uncertainty*, 27(1):5–76, August 2003.
- Philippe Weil. The Equity Premium Puzzle and the Risk-free Rate Puzzle. *Journal of Monetary Economics*, 24(3):401–421, November 1989.