RECURSIVE PREFERENCES, THE VALUE OF LIFE, AND HOUSEHOLD FINANCE

By

Antoine Bommier, Daniel Harenberg, François Le Grand, and Cormac O'Dea

May 2020

COWLES FOUNDATION DISCUSSION PAPER NO. 2231

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

Recursive Preferences, the Value of Life, and Household Finance*

Antoine Bommier Daniel Harenberg François Le Grand
Cormac O'Dea†

May 1, 2020

Abstract

We analyze lifecycle saving strategies using a recursive utility model calibrated to match empirical estimates for the value of a statistical life. We show that, with a positive value of life, risk aversion reduces savings and annuity purchase. Risk averse agents are willing to make an early death a not-so-adverse outcome by enjoying greater consumption when young and bequeathing wealth in case of death. We also find that greater risk aversion lowers stock market participation. We show that this model can rationalize low annuity demand while also matching empirically documented levels of wealth and private investments in stocks. Our findings stand in contrast to studies that implicitly assume a negative value of life.

Keywords: recursive utility, lifecycle model, value of life, risk aversion, saving choices, portfolio choices, annuity puzzle.

JEL codes: D91, G11, J14, J17.

---

*We would like to thank Francisco Gomes, Michael Haliassos, Felix Kubler, Alexander Michaelides, Stephanie Weber and seminar participants at the University of Zurich, Goethe University Frankfurt, Bielefeld University, University of Cologne, the Paris School of Economics, Luxembourg School of Finance, Collegio Carlo Alberto in Turin, the QSPS workshop in Utah, the CEPR household finance workshop, the Netspar International Pension Workshop, ESEM, and the annual meetings of SAET and the Verein für Socialpolitik for their very helpful comments. We are also very grateful to John Bailey Jones for sharing data underlying graphs in his previous published work. A previous version of this paper circulated under the title "Household Finance and the Value of Life".

†Bommier: ETH Zurich, Zürichbergstr. 18, 8092 Zurich, Switzerland, e-mail: abommier@ethz.ch; Harenberg: ETH Zurich, Switzerland, and Oxford Economics Ltd., Germany, e-mail: dan.harenberg@gmail.com; Le Grand: EMLyon Business School, 23 avenue Guy de Collongue, CS 40203, 69134 Ecully Cedex, France, and ETH Zurich, e-mail: legrand@em-lyon.com; O'Dea: Yale University and Institute for Fiscal Studies, 37 Hillhouse Ave., New Haven, CT, USA, e-mail: cormac.odea@yale.edu. Bommier, Harenberg, and Le Grand gratefully acknowledge support from Swiss-Re Foundation and ETH Zurich Foundation. O'Dea gratefully acknowledges funding from the Economic and Social Research Council (grant number ES/P001831/1).

1 Introduction

Household finance and the economic appraisal of the value of life are central issues in the economics of aging. Although both topics make use of similar theoretical foundations based on micro-economic lifecycle models, they have been studied in separate strands of the literature with limited exchange between strands.

On the one hand, the literature on "household finance" (hereafter HF) tackles questions related to consumption, saving and financial portfolio choices over the lifecycle. These questions are of first-order importance for analyzing and designing efficient pension systems, and for thinking more generally about saving incentives. Mortality is usually taken into account, but since it is considered to be exogenous, there tends to be no consideration of the willingness to pay for mortality risk reduction.

On the other hand, the literature on the value of life (hereafter VoL) is centered around questions related to endogenous mortality risk reduction, though has paid limited attention to saving behaviors or portfolio choices. This literature has aimed at providing insights into the trade-off that people face between wealth and mortality, which is key to the evaluation of public policies aiming at lowering mortality risk – such as road safety investments or public health spending.

Both literatures were initiated with contributions relying on the same decision model: the standard additive expected utility model, as introduced by Yaari (1965) in the HF literature and used by Shepard and Zeckhauser (1984) and Rosen (1988) in the VoL literature. There was therefore a single model of rational behavior that could be used to discuss both kinds of issues. This additive model was criticized for several reasons, however. The HF literature focused on the drawback of having risk aversion and the intertemporal elasticity of substitution (IES) intertwined, making it impossible to consider agents with both a low IES and a high degree of risk aversion. The VoL literature further criticized the additive model for suggesting that agents unavoidably prefer death to life when their consumption gets small enough when the IES is below one (see e.g., Marshall, 1984 or Rosen, 1988).

Both lines of literature tried to circumvent these difficulties by adopting recursive specifications inspired by the framework of Epstein and Zin (1989) and Weil (1990)

– which are usually referred to as EZW preferences. However, both branches of the literature suggested the radically different and mutually incompatible choices regarding preference parameters. On the one hand, papers in the HF literature typically assume a coefficient of risk aversion above one. This, however, yields EZW specifications that are inadequate for VoL matters, since they suggest that people would prefer to have shorter lives (see our discussion in Section 2.3). On the other hand, recent contributions in the VoL literature tackled the issues that appear when using the additive model with a low elasticity of substitution by assuming a degree of risk aversion below one (Hugonnier et al., 2013 and Córdoba and Ripoll, 2017). However, EZW specifications that assume both the elasticity of substitution and the coefficient risk aversion to be below one turn out to be ill-defined when applied to situations with realistic mortality patterns and yield counterfactual predictions for lifecycle behaviors unless some implausible "patches" are added to correct for anomalies of the model. See Section 2.3 or Bommier et al. (2018) for an extensive discussion.

Considered together, these recent developments in the HF and VoL literature represent a fragmented approach where a given form of rationality is used when dealing with some issues and another one, incompatible with the former, is used for other issues. There is a need, therefore, for a consistent framework that can be used for jointly modelling savings and mortality choices.

In the current paper, we introduce a well-behaved recursive framework that can be used in both lines of literature, without suffering any of the drawbacks mentioned above. In particular, our model makes it possible to disentangle the IES and risk aversion, allowing for a high degree of risk aversion and a low IES, while predicting plausible (and positive) levels for the value of mortality risk reduction and realistic lifecycle profiles of consumption, savings, portfolio choice and annuity demand.

Central to our approach is that we restrict our attention to models which are monotone with respect to first-order stochastic dominance. The property of monotonicity is essential as it rules out the choice of dominated strategies, which, notwithstanding the fact that they are dominated, can represent optimal strategies in EZW models. In Section 2.4, we provide two examples of this – one in a setting where an individual is deciding how much to invest in some risk-reducing factor (e.g., health) and one in which an individual is deciding how much to save. As shown in Bommier et al. (2017), imposing preference monotonicity leads to a focus on risk-sensitive preferences, initially introduced by Hansen and Sargent (1995), which we will use in the current paper.

The main finding of the paper is that once we constrain the value of mortality risk reduction to be positive and preferences to be monotone, risk aversion is found to reduce saving, annuity demand and stock purchases. These results stand in contrast with those derived in previous studies based on recursive preferences featuring a negative value of mortality risk reduction. Our model can further explain why people purchase very few annuities (the so-called 'annuity puzzle'; see e.g., Mitchell et al., 1999) while matching empirically documented levels of wealth accumulation and participation in stock markets.

The paper has three parts. In the first one, we provide a theoretical review of recursive preferences, including additive, EZW and risk-sensitive preferences. We highlight, in particular, the meaning of preference monotonicity. In the second part, we work with a simple two-period model to emphasize why it is important to properly calibrate the value of life for exploring the role of risk aversion. The last part of the paper is a quantitative exercise that emphasizes the insights that can be gained by using monotone preferences that feature a plausible value of life. We discuss the predictions obtained with a rich calibrated model in which agents with risk-sensitive preferences face different risks (mortality, income and asset return risks), have bequest motives and have access to an annuity market.

2 Theory

2.1 The additive model

The most popular model in both the HF and the VoL literature is the time-additive expected utility model. To link with the rest of the paper, we provide its recursive definition:

$$U_t = (1-\beta)u(c_t) + \beta E[U_{t+1}], \tag{1}$$

where U_t is utility at time t, $\beta \in (0,1)$ is a time preference parameter and $u(c_t)$ is the instantaneous utility derived from consumption at time t.[1] Accounting for mortality is achieved by assigning a utility level, u_d, obtained when death occurs. As it is generally assumed that mortality risk is independent of other risks, the expectation can be decomposed in two stages, one that accounts for the risk of mortality and one that accounts for other risks. Formally denoting the utility conditional on being alive at time t by V_t and the probability of surviving from period t to $t+1$ by π_t the recursion (1) yields:

$$V_t = (1-\beta)u(c_t) + \beta\left(\pi_t E[V_{t+1}] + (1-\pi_t)u_d\right). \tag{2}$$

Additive preferences, as defined in the above equation, are invariant when changing u and u_d (and V_t) by the same positive affine transformation. It is therefore possible to assume, without loss of generality, that the utility of death is set to zero ($u_d = 0$). If one assumes a constant IES, the function u has to be of the form $u(c) = u_l + K\frac{c^{1-\sigma}}{1-\sigma}$ for some constant u_l (the subscript l stands for "life") and a positive scalar K. The IES is then given by $\frac{1}{\sigma}$. By (multiplicative) normalization, the scalar K can be set equal to 1. It is not, however, possible to make a second normalization and set u_l to zero, since the zero utility level is already pinned down as representing the utility of death. The parameter u_l, which determines the utility gap between life and death, is thus an important preference parameter. It is noteworthy that when mortality is exogenous, the constant u_l only contributes to an exogenous additive term that has no impact on the ordering of consumption profiles. This explains why the constant u_l is generally ignored in all studies that assume an exogenous mortality pattern but is explicitly mentioned in the VoL literature (see, for example, the discussion in Hall and Jones, 2007).

The additive specification has been criticized on several grounds. Papers in the HF literature – such as Gomes and Michaelides (2005), among others – expressed concerns regarding the lack of flexibility of the additive specification, especially

[1] The additive model is frequently defined using the recursion $U_t = u(c_t) + \beta E[U_{t+1}]$, which is, of course, equivalent to (1), up to a multiplicative renormalization of the function u.

for its inability to disentangle the IES from risk aversion. The VoL literature, in addition to this lack of flexibility, also emphasized the fact that when the IES is smaller than one (i.e., $\frac{1}{\sigma} < 1$), life becomes unavoidably worse than death when consumption gets small enough. This creates an incentive to enter into Russian-roulette games, so as to avoid having to live in a state which is worse than death. As we discuss in Appendix A, this feature is actually not related to the additive structure but to the assumption of a constant IES and could be avoided by assuming that the IES gets higher when consumption tends to zero.

2.2 Recursive models with mortality

The search for greater flexibility led researchers contributing to each of the HF and the VoL literatures to adopt non-additive recursive models for preferences. These models assume that utility U_t at any date t is defined by the following recursion:

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

where ϕ is an increasing function representing risk preferences and f is a normalization device that can be any increasing function. As in the additive model of Section 2.1, $\beta \in (0,1)$ is a time preference parameter and $u(c_t)$ is the instantaneous utility function. The function f has no impact on preferences and its role is to facilitate convenient representations of the recursion. A common choice is $f(x) = x$, as in the additive specification (1), which we will also use for the representation of risk-sensitive preferences in Section 2.5. Another common choice for f, most often introduced in the case of EZW preferences – see Section 2.3 – involves using the same CRRA function as that used for the instantaneous utility: $f(x) = u(x) = \frac{x^{1-\sigma}}{1-\sigma}$.[2]

Recursive preferences defined by (3) may exhibit non trivial preference for the timing of resolution of uncertainty. From Kreps and Porteus (1978), we know that they exhibit preference for early resolution of uncertainty when, for any $c > 0$, the

[2]Such a normalization implies that an infinitely long consumption path providing the same consumption c in all periods yields a lifetime utility equal to c.

function g_c defined by:

$$x \in Im(\phi) \mapsto g_c(x) = \phi((1-\beta)u(c) + \beta\phi^{-1}(x)) \tag{4}$$

is convex, and for late resolution if g_c is concave. Moreover, as explained by Chew and Epstein (1989), the degree of convexity of g_c, which is measured by $x \mapsto \frac{g_c''(x)}{g_c'(x)}$, quantifies the intensity of preference for the timing, that is, the more convex g_c, the stronger the preference for early resolution.

Finally, as in the additive case, we can derive from equation (3) the recursion defining the utility conditional on being alive at time t:

$$V_t = f^{-1}\left((1-\beta)u(c_t) + \beta\phi^{-1}\left(\pi_t E[\phi f(V_{t+1})] + (1-\pi_t)\phi f(u_d)\right)\right). \tag{5}$$

Here, again, u_d denotes the utility level assigned to death and mortality risk is assumed to be independent of other risks.

2.3 Homothetic EZW preferences

A popular recursive specification is that of Epstein and Zin (1989) and Weil (1990), which is obtained from equation (3), with $u(c) = f(c) = \frac{c^{1-\sigma}}{1-\sigma}$, and $\phi(x) = \frac{1}{1-\gamma}\left((1-\sigma)x\right)^{\frac{1-\gamma}{1-\sigma}}$. This yields:

$$U_t = \left((1-\beta)c_t^{1-\sigma} + \beta\left(E[U_{t+1}^{1-\gamma}]\right)^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{1-\sigma}}. \tag{6}$$

When $\sigma = \gamma$, we obtain the additive model – up to the transformation by f. With $\gamma > \sigma$, as is usually assumed in the HF literature, EZW exhibit preference for early resolution of uncertainty. The magnitude of the implied timing premia are discussed in Epstein et al. (2014).

The recursion relating the utilities conditional of being alive, shown in (5) is now:

$$V_t = \left((1-\beta)c_t^{1-\sigma} + \beta\left(\pi_t E[V_{t+1}^{1-\gamma}] + (1-\pi_t)u_d^{1-\gamma}\right)^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{1-\sigma}}. \tag{7}$$

It is worth emphasizing that in specification (7), choosing a specific value for u_d is not a normalization choice, since changing u_d would impact individual preferences.

The HF and VoL literatures then follow two different (and incompatible) routes.

In the HF literature, the standard assumption is to choose a coefficient of risk aversion $\gamma > 1$. In order to get a homothetic specification, it is moreover assumed that $u_d^{1-\gamma} = 0$ and thus, implicitly, that the utility of death is $u_d = +\infty$. The recursion (7) reduces then to:

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

as is the case, for example, in Gomes and Michaelides (2008). Clearly, setting $u_d = +\infty$ involves assuming that death is preferable to life, regardless of consumption levels. This has the consequence that utility is declining with survival probability ($\frac{\partial V_t}{\partial \pi_t} < 0$).[3] Since the HF literature focuses on cases where mortality is exogenous, one could think that this assumption is harmless. However, as we will see in Section 3, this is not the case, and this assumption drives some key results in the field.

Assuming that death is preferable to life precludes the use of the model to study VoL issues. For this reason, rather than assuming $\gamma > 1$ and $u_d = +\infty$, Hugonnier et al. (2013) and Córdoba and Ripoll (2017) suggest the imposition of $\gamma < 1$ and $u_d = 0$, which again yields the recursion (8). An additional issue arises, however, when considering the case where the IES is smaller than one ($\frac{1}{\sigma} < 1$), which is the relevant case in most of the literature (see the meta-analysis of Havránek, 2015). When applied to realistic mortality patterns, such models are typically ill-defined, since the unique solution to the above recursion is the null function: $V_t = 0$ independent of the consumption profile. Moreover, working with this model without a consideration of definitional issues yields extremely odd predictions where agents consume nothing early in the lifecycle to sustain a consumption that sky-rockets towards $+\infty$ at old ages. Such counterfactual implications are avoided in Hugonnier et al. (2013) by assuming a model of perpetual youth, where mortality risk remains small at all ages. In Córdoba and Ripoll (2017), the anomaly is accommodated by the introduction of an ad hoc (exogenous) health profile, which aims at decreasing the marginal utility of consumption in order to counterbalance the effect mentioned above. However, for the correction to work, one has to assume

[3]Equation (8) yields $\frac{\partial V_t}{\partial \pi_t} = \frac{\beta}{1-\gamma} \left(E[V_{t+1}^{1-\gamma}] \right)^{\frac{1-\sigma}{1-\gamma}} V_t^{\sigma}$, which has the same sign as $1-\gamma$. See Section 5 for further discussion.

that greater health implies a lower utility level. The problem of a negative value of life is solved by introducing another one: that of a negative value of health.[4] We refer to Bommier et al. (2018) and Zhang et al. (2018) for an extensive discussion on these aspects.

Assuming a negative value of life, as in Gomes and Michaelides (2008) and in many other papers in the HF literature, or using a framework that is ill-behaved when the IES is smaller than one, as in Hugonnier et al. (2013) and Córdoba and Ripoll (2017), will not be appropriate in many applications. There is therefore a need to introduce another approach. Rather than bringing a new specification chosen for its tractability, we suggest the imposition of a natural property: preference monotonicity.

2.4 Preference monotonicity

An appealing property of preferences is monotonicity with respect to first-order stochastic dominance (see Bommier and LeGrand, 2014a for a formal definition). In words, monotonicity requires that an agent will not take an action if another one would seem preferable in all future circumstances. This property, which is similar to the elimination of dominated strategies in game theory, seems to be a reasonable requirement when modeling individual rationality. The standard additive model is monotone, but only few recursive extensions are monotone. Bommier et al. (2017) show that EZW preferences are non-monotone, except when an IES equal to one is assumed.[5]

The non-monotonicity feature of EZW preferences can be illustrated in various

[4]Córdoba and Ripoll (2017) extend (8) by assuming that:

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

where H_t is an exogenous health profile. By choosing an appropriate health profile, it is then possible to match any possible consumption profile. The calibration in their paper leads to the choice of a profile H_t that rapidly declines with age – so as to mechanically counterbalance the counterfactual increase in consumption arising in that model when the IES is smaller than one. When the IES is smaller than one (i.e., $\frac{1}{\sigma} < 1$), the model implies that health has a negative impact on welfare: $\frac{\partial V_t}{\partial H_t} < 0$.

[5]When the IES equals one (i.e., the limit case where $\sigma \to 1$), EZW preferences enter in the class of risk-sensitive preferences that we consider in Section 2.5.

settings. For example, Bommier et al. (2017) consider the case of saving behavior with uncertain income and asset returns, in the absence of mortality risk. To get closer to the problem considered by the current paper, we give below two examples where EZW preferences yield dominated choices in settings where the uncertainty is solely related to mortality.

First example: Investment in mortality risk reduction. In our first example, we consider an agent who may live for one or two periods. We assume that consumption in the second period, in the case of survival, is exogenous and equal to c_1. The only decision to be made by the agent is to dedicate an amount e in the first period to some investment that reduces mortality risk (e.g., health expenditures). We assume that such investments have a unit price. Letting w_0 be the agent's initial endowment of wealth, a mortality-risk reducing investment equal to e leads to consumption in the first period of $w_0 - e$.[6] This choice impacts the survival probability at the end of the first period. We further assume that the survival probability also depends on another factor that the agent cannot observe. More precisely, there are two possible states of nature h and s – that can be thought of as "healthy" and "sick" – that occur with probability p_h and $p_s = 1 - p_h$, respectively. Investing e generates a survival probability $\pi_h(e)$ in state h and survival probability $\pi_s(e)$ in state s. The agent has to select the level of investment, e, without knowing the actual state of nature. A given effort level e^* is said to be dominated if another level $\hat{e}$ would be preferable in both states h and s. By definition, monotone preferences never yield dominated choices. This is not the case for EZW preferences whenever $\gamma \neq \sigma$ and $\sigma \neq 1$. With EZW preferences, the agent's problem can be written as:

$$\max_{e \geq 0} \left((1-\beta)(w_0 - e)^{1-\sigma} + \beta(1-\beta)(E[\pi(e)])^{\frac{1-\sigma}{1-\gamma}} c_1^{1-\sigma} \right)^{\frac{1}{1-\sigma}},$$

[6] The wealth w_0 aggregates all sources of agent's revenues. Bequests received from parents, if any, are supposed to be redistributed through a lump-sum transfer that is included in w_0. The agent is therefore assumed not internalize possible bequest she may receive in her consumption-saving decisions.

where $E[\pi(e)] = p_h \pi_h(e) + p_s \pi_s(e)$. If an interior optimum effort level $e^* > 0$ exists, it has to fulfill the following first-order condition:

$$(w_0 - e^*)^{-\sigma} = \frac{1}{1-\gamma} \beta E[\pi'(e^*)] \left(E[\pi(e^*)]\right)^{\frac{\gamma-\sigma}{1-\gamma}} c_1^{1-\sigma}.$$

Two remarks can be made. First if $\gamma > 1$ an interior solution can be obtained only if $E[\pi'(e^*)] < 0$, meaning that investing further in precaution would reduce life expectancy. This is in line with the fact that when $\gamma > 1$, the agent is willing to reduce her survival probability (i.e., the value of mortality risk reduction is negative) – see the discussion about $\frac{\partial V_t}{\partial \pi_t} < 0$ in Section 2.3. The second point, which relates to our discussion of monotonicity, is that whenever $\gamma \neq \sigma$ and $\sigma \neq 1$, it is possible to find functions $\pi_h(\cdot)$ and $\pi_s(\cdot)$ such that the optimal effort is obtained for some $e^* > 0$ while another optimal effort level $\widehat{e} \neq e^*$ would be preferable in both states h and s. The formal proof is given in Section B.2 of the Appendix.

Second example: A cake-eating problem. Our second example also considers a two-period problem using EZW preferences. Here, we study the savings decision of an agent who has an exogenous probability ($\pi_0 \in (0,1)$) of surviving to the second period. Formally, we consider a simple "cake-eating" problem where the agent endowed with an initial amount of wealth w_0, and no additional source of revenue, has to decide in period 0 how much to keep for period 1. For simplicity, we assume that the agent has to choose an amount b to be invested in a risk-free bond whose return is equal to $\frac{1}{\beta}$. The decision problem is therefore:

$$\max_{b \in \mathbb{R}_+} \left((1-\beta)(w_0 - b)^{1-\sigma} + \beta \pi_0^{\frac{1-\sigma}{1-\gamma}} (1-\beta) (\frac{b}{\beta})^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{9}$$

Ex ante, the agent foresees two possible states of the world. Either she lives only one period, and optimal behavior would be to consume everything in that period (zero savings: $b = 0$). Or she lives two periods and the optimal strategy involves saving $b = \frac{w_0}{1+\beta}$ in order to have the same consumption level in both periods. The agent however, has to make a saving decision before knowing which state realizes. Choosing a saving level above $\frac{w_0}{1+\beta}$ would be a dominated choice since saving exactly $\frac{w_0}{1+\beta}$ would provide a higher ex-post utility whether she survives or not.

Simple derivations show that the solution to (9) is $b^* = \frac{w_0}{1+\beta\pi_0^{\frac{1}{\sigma}\frac{\sigma-1}{1-\gamma}}}$. The interesting case is the one where the risk aversion is below 1 ($\gamma < 1$) and the IES is also below 1 ($\frac{1}{\sigma} < 1$). This implies that $b^* > \frac{w_0}{1+\beta}$, indicating that the agent opts for a dominated strategy. Even more surprising, when π_0 becomes very small, then b^* gets close to w_0. In other words, when the agent is almost sure of dying at the end of the first period, she decides to consume almost nothing in the first period, to keep all her resources for a second period that she will almost surely never see. The non-monotonic feature of the agent's behavior here is quite extreme and illustrates the drawbacks of opting for non-monotone preferences.

Such drawbacks extend to multi-period settings, yielding similar conclusions regarding saving behaviors. We refer again to Bommier et al. (2018) for a detailed analysis.

2.5 Risk-sensitive preferences

Imposing monotonicity does not preclude the use of non-additive recursive models. It does, however, dramatically reduce the set of recursive specifications that can be used. It was shown in Bommier et al. (2017) that the only class of Kreps and Porteus (1978) recursive preferences that is monotone and flexible enough to disentangle risk aversion from the IES is the one provided by risk-sensitive preferences. Such preferences correspond to the recursive setting where:

$$U_t = (1-\beta)u(c_t) - \frac{\beta}{k}\log\left(E_t[e^{-kU_{t+1}}]\right). \tag{10}$$

As was the case for EZW preferences, recursion (10) is a particular case of equation (3), where $f(x) = x$, and $\phi(x) = \frac{1-e^{-kx}}{k}$. When $k \to 0$, recursion (10) converges toward the standard additive model of Section 2.1.

Intuitively, preferences represented by recursion (10) are monotone, as the term $(1-\beta)u(c_t)$ can be included inside the expectation. That is, recursion (10) can also be written as $U_t = -\frac{\beta}{k}\log\left(E[e^{-\frac{k}{\beta}((1-\beta)u(c_t)+\beta U_{t+1})}]\right)$, implying that the choices at time t involve maximizing an expectation, just like in the expected utility framework.

With risk-sensitive preferences, the function g_c defined in (4) takes the form:

$$g_c(x) = \frac{1 - (1 - kx)^{\beta} e^{-k(1-\beta)u(c)}}{k}, \text{ for } x \in Im(\phi) = (-\infty, \frac{1}{k}).$$

Simple calculation show that $\frac{g_c''(x)}{g_c'(x)} = \frac{k(1-\beta)}{1-kx}$, which is independent of c. Risk-sensitive preferences with $k > 0$, therefore exhibit a preference for early resolution of uncertainty, the degree of which depends on both β and k.[7] A feature that is worth emphasizing is that preferences defined by (10) are invariant when adding a constant to the instantaneous utility function u. Indeed, equation (10) directly implies that for any $\lambda \in \mathbb{R}$, we have $U_t + \lambda = (1 - \beta)(u(c_t) + \lambda) - \frac{\beta}{k} \log \left(E_t[e^{-k(U_{t+1}+\lambda)}] \right)$. Thus adding a constant λ to the instantaneous utility amounts to adding the same constant λ to the utilities U_t, with no impact on agent's preferences. The risk-sensitive model thus preserves one of the invariance properties of the additive specification.[8] This is convenient for normalization matters, since just like in the additive model, it can be assumed that $u_d = 0$ with no generality loss. The recursion (5), relating the utilities conditional on being alive then becomes:

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

Interestingly risk-sensitive preferences do not constrain the IES to be constant, so that the simple solutions suggested in Appendix A could possibly be adopted to avoid tastes for Russian-roulette lotteries.

Note also that risk-sensitive preferences could serve as a basis for extensions that would combine recursivity, monotonicity and recent developments in decision theory – in particular the models, building upon the recursive structure of Kreps and Porteus (1978), which account for source-dependent risk aversion. See Skiadas (2013, 2015), as well as Gârleanu and Panageas (2015, Appendix D).

[7] This reflects a relationship between time preferences, risk aversion and preference for the timing, which necessarily holds for monotone recursive preferences. See Bommier et al. (2017). Note that when β gets close to one, preference for the timing vanishes.

[8] The additive specification is also invariant to a (multiplicative) rescaling of the instantaneous utility. Here, rescaling instantaneous utility without impacting preferences would require to rescale the risk aversion parameter, k.

3 Why does the value of life matter?

With monotone preferences, a utility maximizing choice under uncertainty can be seen as making trade-offs between the ex-post utilities obtained in each state of the world.[9] Moreover, increasing risk aversion has the effect of putting greater weight on "bad states" of the world when considering such trade-offs. In a number of cases, this leads to intuitive predictions, as illustrated with a precautionary saving problem in Bommier et al. (2017). Similarly, the two-period problem we explore below shows that this ability to derive simple predictions also applies to intertemporal choice under uncertain lifetime, with clear-cut predictions on saving behaviors, bequests, and annuity purchase. Moreover, this example will emphasize that in order to understand the role of risk aversion, it is essential to correctly identify what are good states and bad states of the world. The predictions will be completely different if one assumes that having a long life is a good outcome (i.e., if the value of mortality risk reduction is positive) or if one assumes that having a long life is a bad outcome (i.e., if the value of mortality risk reduction is negative), even if mortality is assumed to be exogenous.

We consider here an agent who lives for at most two periods. In period 0, the agent is endowed with a level of wealth w_0 and has to make saving and annuity purchase decisions. More precisely, the agent may invest in bonds, which yield a safe return of R^f and which are bequeathed in case of death, or in annuities, with return $\frac{R^f}{\pi_0}$ but are not bequeathable. Formally, denoting by c_0 and c_1 the consumption in the first and second periods, by b the amount invested in bonds, by a the amount invested in annuities, and by x the amount bequeathed in case of death, we have:

$$w_0 = c_0 + b + a; \quad c = R^f(b + \frac{a}{\pi_0}); \quad x = R^f b.$$

The agent is endowed with risk-sensitive preferences, as represented by equation (10). However, the utility associated with death is not constant and equal to u_d anymore, but depends on the size of the bequest left by the agent to her heirs.

[9]This is not the case with non-monotone preferences, since the agent may take actions that lower the ex-post utilities in all states.

We denote by $v(x)$ the utility associated with bequeathing the amount x. The recursion (5) defining the utility V_0 representing the agent's preferences becomes, in the presence of bequests:

$$V_0 = (1-\beta)u(c_0) - \frac{\beta}{k}\log\left(\pi_0 e^{-k(1-\beta)u(c_1)} + (1-\pi_0)e^{-k(1-\beta)v(x)}\right). \quad (11)$$

We further assume here that $u(c) = u_l + \frac{c^{1-\sigma}}{1-\sigma}$ and $v(x) = u_d + \theta\frac{x^{1-\sigma}}{1-\sigma}$, where the scalar $\theta \geq 0$ quantifies the strength of the bequest motives. There is no consensus on the formulation of v, but the form we have chosen is the same as in Cocco et al. (2005), Inkmann et al. (2011) and Yogo (2016).[10] With no loss of generality, we can normalize the utility with $u_d = 0$. Note that further constraining utility by setting, for instance $u_l = 0$ would not be a mere normalization and would impose constraints on the value of mortality risk reduction. The larger u_l, the greater the utility gap between life and death and the larger the value of mortality risk reduction. In particular, it follows from (11) that:

$$\frac{\partial V_0}{\partial \pi_0} = \frac{\beta}{\pi_0 e^{-k(1-\beta)u(c_1)} + (1-\pi_0)e^{-k(1-\beta)v(x)}} \frac{e^{-k(1-\beta)(u_l + \frac{c_1^{1-\sigma}}{1-\sigma})} - e^{-k(1-\beta)\theta\frac{x^{1-\sigma}}{1-\sigma}}}{-k}, \quad (12)$$

implying that, for a given c_1 and w, we have $\frac{\partial V_0}{\partial \pi_0} > 0$ if u_l is above $\frac{\theta x^{1-\sigma} - c_1^{1-\sigma}}{1-\sigma}$ (people then prefer longer lives) and $\frac{\partial V_0}{\partial \pi_0} < 0$ if u_l is below that threshold (people then prefer shorter lives).

Let b_k and a_k be the optimal saving and annuity choices of an agent with risk aversion k. Formally, the consumption-saving program can be written as:

$$\begin{aligned}(b_k, a_k) = \arg\max_{(b,a)\in\mathbb{R}_+^2} & (1-\beta)u(w_0 - b - a) \\ & - \frac{\beta}{k}\log\left(\pi_0 e^{-k(1-\beta)u(R^f b + R^f \frac{a}{\pi_0})} + (1-\pi_0)e^{-k(1-\beta)v(R^f b)}\right)\end{aligned} \quad (13)$$

We denote by $c_{0,k} = w_0 - b_k - a_k$ and $c_{1,k} = R^f(b_k + \frac{a_k}{\pi_0})$ the corresponding optimal first- and second-period consumption levels. Here again, the case of additively separable preferences is obtained by taking the limit as $k \to 0$.

[10] Other papers, such as De Nardi (2004), Lockwood (2012) or Bommier and LeGrand (2014b) consider a bequest utility of the form $\theta\frac{(\overline{x}+x)^{1-\sigma}}{1-\sigma}$ that is not homothetic but enables bequests to be modeled as a luxury good. We will use such a specification in the calibrated multi-period quantitative model presented in Section 4.

Proposition 1 *Consider the consumption-saving problem of equation (13).*

If $k = 0$, then the choices a_0 and b_0, and hence the consumption levels $c_{0,0}$ and $c_{1,0}$, are independent from u_l.

If $k > 0$, we have:

- *if u_l such that $\frac{\partial V_0}{\partial \pi_0} > 0$ at the optimum (i.e., if the value of mortality risk reduction is positive), then $\frac{\partial a_k}{\partial k} < 0$, $\frac{\partial b_k}{\partial k} > 0$, $\frac{\partial c_{0,k}}{\partial k} > 0$ and $\frac{\partial c_{1,k}}{\partial k} < 0$;*
- *if u_l such that $\frac{\partial V_0}{\partial \pi_0} < 0$ at the optimum (i.e., if the value of mortality risk reduction is negative), then $\frac{\partial a_k}{\partial k} > 0$, $\frac{\partial b_k}{\partial k} < 0$, $\frac{\partial c_{0,k}}{\partial k} < 0$ and $\frac{\partial c_{1,k}}{\partial k} > 0$.*

The proof of Proposition 1 can be found in Section B.1 of the Appendix. Here we will comment on the results of Proposition 1, focusing on the case where $k > 0$. We begin with the case where $\frac{\partial V_0}{\partial \pi_0} > 0$ at the optimum, that is, when the agent would prefer to have a greater survival probability. We find that increasing risk aversion increases first-period consumption ($\frac{\partial c_{0,k}}{\partial k} > 0$), implying therefore that overall savings are reduced ($\frac{\partial}{\partial k}(b_k + a_k) < 0$). Moreover, savings shift towards safe assets ($\frac{\partial b_k}{\partial k} > 0$) at the expense of annuity purchases ($\frac{\partial a_k}{\partial k} < 0$). In fact, the greater the risk aversion, the more the agent is concerned about the bad state, that is, the case where she would die after the first period. To increase lifetime utility derived in this bad state, she increases her first period consumption and the amount she leaves as bequest ($\frac{\partial b_k}{\partial k} > 0$). At the same time, she purchases a smaller amount of annuities. These choices make a short life a less adverse outcome. Of course, this comes at the cost of a lower second-period consumption ($\frac{\partial c_{1,k}}{\partial k} < 0$) and of a lower lifetime utility in case of survival ($\frac{\partial}{\partial k}(u(c_{0,k}) + \beta u(c_{1,k})) < 0$).

Interestingly, the findings look very different in the counterfactual case where $\frac{\partial V_0}{\partial \pi_0} < 0$, that is, when the agent would prefer to have a lower survival probability. Indeed, as shown in Proposition 1, the impact of an increase in risk aversion is systematically reversed compared to the case where the value of mortality risk reduction is positive. This is understandable since the impact of an increase in risk aversion is to put greater focus on the bad state, and thus mechanically depends on which state is actually the worst. When $u(c_1) < v(x)$, the bad state is the one where the agent lives for two periods, and an agent with $k > 0$ purchases a larger amount of annuities than an agent with additive preferences ($k = 0$), precisely to make such a "bad outcome" not so bad. As we will see in Section 5, such matters related to the sign of the value of life are key to understand how our results contrast with those of previous contributions.

A final remark is that since preference for the timing and risk aversion are intertwined for risk-sensitive preferences, one may be concerned that the results of Proposition 1 are driven by preference for the early resolution of uncertainty and not by risk aversion. The case $\beta = 1$ can help disentangle the two effects. First, when $\beta = 1$, the RS model features indifference with respect to the timing of the resolution of uncertainty (see Footnote 7 in Section 2.5). Second, it can be checked that the results of Proposition 1 are still valid – with strict inequalities – when $\beta = 1$.[11] Together, these two facts show that the force driving Proposition 1 is risk aversion and not preference for the timing.

4 A quantitative lifecycle model

Having argued why the value of life matters even in models with exogenous mortality risk, this section outlines a quantitative lifecycle model which can be used to study the interplay between that risk, saving behavior, portfolio choice and annuity purchases. The model's innovation is to use, in an otherwise standard lifecycle model, risk-sensitive preferences to ensure preference recursivity and monotonicity while matching empirical estimates of the value of mortality risk reduction. .

4.1 The setup

We consider an economy populated by agents endowed with risk-sensitive preferences who face risks over mortality, income and asset returns. The agent may save through a bond, a risky asset and may insure against longevity risk by purchasing an annuity.

[11] In order to deal with the case $\beta = 1$, instantaneous utility has to be normalized differently (to avoid the factor $1 - \beta$ which is then null and no longer a mere normalization – see also Footnote 1 in Section 2.1). With $\beta = 1$, the problem (13) should be written as:

$$(b_k, a_k) = \arg \max_{(b,a)\in\mathbb{R}^2_+} u(w_0 - b - a) - \frac{1}{k} \log \left(\pi_0 e^{-ku(R^f b + R^f \frac{a}{\pi_0})} + (1 - \pi_0) e^{-kv(R^f b)} \right).$$

Time is discrete, a model period is a year, and time t denotes model age, which is biological age minus 20. The agent enters the model at the start of working life, at $t = 0$. There is a single consumption good, whose price serves as a numeraire.

Mortality risk. The agent faces mortality risk, which is assumed to be exogenous and independent of all other risks. If alive at date t, the agent survives to date $t + 1$ with probability π_t. There exists a date T_M, such that the probability of living after T_M is $\pi_{T_M} = 0$.

Labor income risk. At any age, when alive, the agent receives an income denoted y_t. The agent exogenously retires at age T_R. During retirement ($t \geq T_R$), the agent receives an annual pension income $y_t = y^R$. During working life ($t < T_R$), the agent earns a risky labor income $y_t = y_t^L$,

$$\ln y_t^L = \mu_t + \zeta_t, \tag{14}$$

$$\zeta_t = \rho \zeta_{t-1} + \upsilon_t. \tag{15}$$

where $(\mu_t)_{t \geq 0}$ is a deterministic process that depends on age, and ζ_t is an AR(1) stochastic component, with persistence parameter ρ and innovation $(\upsilon_t)_{t \geq 0}$, which is IID normally distributed with mean 0 and variance σ_υ^2. We denote average earnings over working life as $\bar{y}$.

Financial risk and security markets. The agent can save through bonds and stocks and can purchase an annuity. The bond pays a constant risk-free gross return, R^f. The stock yields a risky return, defined as: $R_t^s = R^f + \omega + \nu_t$, where ω represents the average risk premium of stocks over bonds, while the financial risk $(\nu_t)_{t \geq 0}$ is an IID normally distributed process with mean 0 and variance σ_ν^2.

The agent must pay a cost $F \geq 0$ to participate in the stock market, which may be interpreted as the opportunity cost of discovering how the stock market works. In our baseline model, we assume it is a flat once-in-a-lifetime cost: if the cost is paid at a given date t, the agent can freely trade stocks at date t and at any date afterwards.[12]

[12] In Appendix F.3, we investigate another participation structure, where the cost must be paid in every period the agent wants to buy stocks.

Finally, an annuity can be purchased in the period before retirement age ($T_R - 1$). The annuity is a financial asset that pays one unit of income every period, as long as its holder is alive. The price of a single unit of annuity income, q, is:

$$q = (1+\delta) \sum_{\tau=1}^{T_M - T_R} \frac{\prod_{s=0}^{\tau} \pi_{T_R - 1 + s}}{(R^f)^{\tau}}, \tag{16}$$

where the parameter $\delta \geq 0$ is a loading factor on annuity. When $\delta = 0$, the annuity is actuarially fair and its price equals the discounted present value of future payoffs. The larger δ, the further is annuity pricing from actuarial fairness. There is one final annuity market imperfection. Following Pashchenko (2013), annuity purchases below a minimum threshold ($\underline{a} > 0$) are not allowed.

Timing and choices. At the start of every period, the agent first learns the realizations of financial and labor income shocks and whether she is alive or not.

If she is alive, her resources at the beginning of the period consist of her wealth, comprising bond, annuity and stock payoffs plus labor income earned, or public pension income received, in the period. Resources are used to cover consumption as well as the purchase of bonds, annuities, and stocks. Formally, the budget constraint of the living agent at date t can then be expressed as follows:

$$c_t + q a_t + b_t + s_t + \mathbb{1}_{\eta_t=1} \mathbb{1}_{\eta_{t-1}=0} F = y_t + w_t, \tag{17}$$

$$\text{with: } w_t = a_{T_R-1} \mathbb{1}_{t \geq T_R} + R^f b_{t-1} + R_t^s s_{t-1}, \tag{18}$$

where c_t and w_t are consumption and wealth in period t and b_t, s_t and a_t are, respectively, the quantity of bonds, stocks and annuities purchased in period t. The index η_t reflects market participation status and is equal to 0 if she has never paid the participation cost before and therefore never held stocks. The term $\mathbb{1}_{\eta_t=1} \mathbb{1}_{\eta_{t-1}=0} F$ in equation (17) represents the fixed cost of participation.[13] Annuity returns are received from age T_R and are therefore equal to $a_{T_R-1} \mathbb{1}_{t \geq T_R}$. No asset,

[13]Throughout the paper, $\mathbb{1}_A$ denotes an indicator function equal to 1 if A is true and 0 otherwise.

including annuities, can be sold short. These constraints are summarized here:

$$s_t = 0 \text{ if } \eta_t = 0, \tag{19}$$

$$a_t = 0 \text{ if } t \neq T_R - 1, \tag{20}$$

$$a_{T_R-1} = 0 \text{ or } a_{T_R-1} \geq \underline{a}, \tag{21}$$

$$b_t \geq 0, s_t \geq 0 \text{ and } c_t > 0 \tag{22}$$

If the agent is dead at date t, she bequeaths bonds and stocks, but not annuities. The bequest x_t amounts to:

$$x_t = R^f b_{t-1} + R_t^s s_{t-1}. \tag{23}$$

A feasible allocation is a sequence of choices $(c_t, b_t, a_t, s_t, x_t, \eta_t)_{t\geq 0}$ satisfying the constraints (17)–(23). The set of feasible allocations is denoted $\mathcal{A}$. We assume exogenous initial conditions: $\eta_{-1} = 0$, $b_{-1} \geq 0$, $a_{-1} \geq 0$, and $s_{-1} \geq 0$.

4.2 Preferences and agent's program

Intertemporal preferences. Agents have risk-sensitive preferences. The utility of the alive agent at age t, V_t, is defined through the following recursion:

$$V_t = (1-\beta)u(c_t) - \frac{\beta}{k} \log\left(\pi_t E_t\left[e^{-kV_{t+1}}\right] + (1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]\right), \tag{24}$$

where $v(x_t)$ is the instantaneous utility obtained upon dying and bequeathing x_t. Utility is normalized here so that being dead while leaving no bequest provides utility $v(0) = 0$.[14]

Instantaneous utility function specification. We assume that the agent has a constant IES. Formally,

$$u(c) = \begin{cases} u_l + \frac{c^{1-\sigma}-1}{1-\sigma} & \text{if } \sigma \neq 1, \\ u_l + \log(c) & \text{if } \sigma = 1, \end{cases} \tag{25}$$

[14]With such a normalization the continuation utility when dead, $(1-\beta)v(x_{t+1}) + \beta v(0)$, simplifies to $(1-\beta)v(x_{t+1})$.

where the parameter $\sigma > 0$ is the inverse of the IES, and u_l is a parameter that provides the instantaneous utility derived when alive and consuming one unit of consumption ($u(1) = u_l$). It can also be interpreted as the difference in utility between being alive and consuming one unit and being dead and bequeathing nothing ($u_l = u(1) - v(0)$). Since utility has already been normalized when assuming $v(0) = 0$, we cannot further set u_l to an arbitrary value. This parameter must, therefore, be carefully calibrated. Note that when $\sigma > 1$, there necessarily exists a threshold $\underline{c}$ below which $u(c) < 0$, implying that the agent would prefer to die and leave no bequest rather than staying alive. In theory, this threshold could be used to calibrate u_l. However, this would involve basing the calibration on extreme cases (suicides), for which our model is surely ill-suited (a suicidal decision is a complex multi-dimensional decision that, needless to say, involves more than poverty). A better calibration strategy, which we will use in Section 4.4, relies on using agents' decisions that relate to safety. Since a higher value of u_l means a higher valuation of being alive relative to being dead, the value of u_l should be reflected in the financial decisions that agents make to lower their mortality risks, for example when investing in safer (but more expensive) cars or opting for safer (but lower paid) jobs.

The utility derived from bequests, $v(x)$, is assumed to be continuous, increasing in the amount of bequest, and to exhibit bounded and decreasing marginal utility. The functional form we use has been widely applied (see e.g., De Nardi, 2004, De Nardi et al., 2010, Ameriks et al., 2011, and Lockwood, 2012 and 2018). Formally,

$$v(x) = \begin{cases} \frac{\theta}{1-\sigma}\left[(\overline{x}+x)^{1-\sigma} - \overline{x}^{1-\sigma}\right] & \text{if } \sigma \neq 1, \\ \theta \log\left(\frac{\overline{x}+x}{\overline{x}}\right) & \text{if } \sigma = 1, \end{cases} \tag{26}$$

where σ is the inverse of the IES used in the expression (25) defining the function u, while $\theta \geq 0$ is the strength of the bequest motive. With $\overline{x} > 0$, bequests are a luxury good, as has been shown by, for example, Hurd and Smith (2002). The derivative $v'(0)$ is finite, so that agents bequeath only when their wealth is large enough (an empirical regularity documented by e.g., De Nardi, 2004).

Agent's program. The agent's problem involves determining the feasible allocation in the set $\mathcal{A}$ that maximizes her utility defined in (24). There is no analytical solution to this problem. We thus solve the model numerically. In short, state variables are discretized and decision rules are obtained for points on the grid by backwards induction from the last period. Linear approximation is used to evaluate the value function at points off the grid and integration over earnings and asset price shocks is carried out using Tauchen (1986). Further details are given in Appendix D.

4.3 Value of mortality risk reduction

To properly calibrate u_l, we will consider the marginal rate of substitution between survival probability and wealth, which quantifies how much a given agent is willing to pay – in terms of wealth – to reduce her mortality risk. This marginal rate of substitution is most often called the *value of a statistical life* (VSL, henceforth), even though there are recent recommendations to use the terminology "value of mortality risk reduction".[15] Formally, the VSL at date t, denoted VSL_t, is defined as:

$$VSL_t = \frac{\frac{\partial V_t}{\partial \pi_t}}{\frac{\partial V_t}{\partial w_t}}, \tag{27}$$

where we recall that expression of the wealth w_t is provided in equation (18). This definition is standard and is used in Rosen (1988), for instance. The formal expressions for the risk-sensitive and additive models – that will be used to calibrate u_l – can be found in Section B.3 of the Appendix.

Empirical literature on the value of a statistical life. The value of mortality risk reduction is a central parameter for cost-benefit analyses in many policy realms. This includes evaluating environmental policy (see U.S. Environmental Protection Agency 2011 where the value of mortality risk reduction is central in estimating the benefits of the Clean Air Act), transport policy (see US Department of Transportation, 2016 on quantifying the benefit of road safety rules) and decisions around the availability of medical procedures in countries where the choice of which

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

(expensive) medical treatments to offer is made by the state (see Rawlins and Culyer, 2004 for a discussion in the context of the UK's National Health Service).

There are two distinct approaches that have been used to estimate the value of mortality risk reduction. The first is a revealed preference approach which estimates it from observed decisions by individuals (e.g., from compensating differentials associated with risky jobs or willingness to pay for safety features on vehicle purchases). The second is a stated preference approach, where individuals' valuation is explicitly elicited by a survey. Both approaches provide a relatively broad range of estimates. This is not surprising, as any estimate of willingness-to-pay for mortality risk reduction will depend on individual preferences and individual financial and demographic characteristics. A broad overview of the literature and details on the range of estimates that have been reported can be found in Viscusi and Aldy (2003) and Kniesner and Viscusi (2019).

4.4 Calibration

Demographics, endowments and asset market parameters. We describe here how exogenous parameters, relating to demographics, endowments and asset markets, are set. Table 1 provides a summary.

Taking demographics first, agents retire at the age of 65, which for many years was the Normal Retirement Age for Social Security in the U.S. Mortality rates are taken from the Human Mortality Database for the USA for 2016. We assume that all individuals die at the age of 100 if not before.

Turning to endowments, agents earn a wage in each period up to the age of 64 and receive Social Security payments from the age of 65. The shape of the deterministic age-productivity profile is taken from Harenberg and Ludwig (2019), who compute it from PSID data using the method in Huggett et al. (2011). This series is transformed so that average earnings are set at \$46,640, average US net earnings in 2016 (Social Security Administration, 2020). The values for the stochastic component of the income process are taken from Guvenen (2009), who reports an autocorrelation $\rho = 0.988$ and a variance of persistent shocks of $\sigma_v^2 = 0.015$. Further details on the earnings process are given in Appendix E. Public

Table 1: Exogenous parameters in baseline economy

Parameter	Value	Source
Demographics		
Age at retirement, T_R	45 (= 65 − 20)	SSA Historical Normal Retirement Age in US
Maximum age, T_M	80 (= 100 − 20)	
Cond. survival rates, $\{\pi_t\}$		Human Mortality Database, U.S. 2016
Endowments		
Average wage, $\bar{y}$	US\$ 46,640	Average net compensation 2016, SSA
Age productivity, $\{\mu_t\}$	cf. App. E	Harenberg and Ludwig (2019)
Public pension, y^R	$40\%\times\bar{y}$	Average SS replacement rate (Biggs and Springstead, 2008)
Labor income autocorr., ρ	0.988	Guvenen (2009)
Var. of persistent shocks, σ_v^2	0.015	Guvenen (2009)
Asset Markets		
Gross risk-free return, R^f	1.02	Campbell and Viceira (2002)
Equity premium, ω	4%	Campbell and Viceira (2002)
Stock volatility, σ_ν	15.7%	Campbell and Viceira (2002)
Annuity Market		
Min. annuity purchase, $\underline{a}$	US\$ 3,680	Pashchenko (2013)
Administrative load, δ	10%	Pashchenko (2013)

pensions, y^R, are set at 40 percent of the average earnings, which is approximately the average replacement rate afforded by Social Security (Biggs and Springstead, 2008). We assume that the agent enters the model with no assets, $s_{-1} = 0$ and $b_{-1} = 0$.

Turning finally to our parameterization of asset market features, we follow Campbell and Viceira (2002). The gross risk-free return is set at $R^f = 1.02$. We set the equity premium to $\omega = 4\%$ and stock volatility is $\sigma_\nu = 15.7\%$. These values represent common choices in the lifecycle literature (see, for example, Lusardi et al., 2017). We set the administrative load (δ) to 10% (the value used by Pashchenko, 2013, Lockwood, 2012, and in range of Brown, 2007). We set the minimum annuity purchase to \$3,680, calculated by converting the value in Pashchenko (2013) to 2016 dollars. The final feature of the asset market, F, the participation cost in the risky asset, is calibrated to match participation in stock markets – we defer discussion of that to below.

Preference parameters. We set the intertemporal elasticity of substitution to 0.5, so that its inverse is $\sigma = 2$, which is a common value in the literature. The time preference and risk aversion parameters, β and k, and the utility gap between life and death, u_l, are set so that lifecycle behaviors match our calibration targets, as we will detail later on.

To set the bequest function parameters (θ, which governs the intensity of the bequest motive, and $\bar{x}$, which governs the extent to which bequests are a luxury good), we use the estimates of De Nardi et al. (2010), who study the problem of bequests in detail. Their model, which features additive preferences, implies different values for IES and the discount factor to ours, and we therefore cannot directly import their bequest parameters. Our approach is to replicate two targets implied by their model: (i) the maximal wealth with zero bequest; (ii) the marginal propensity toconsume wealth, both of which are computed for an alive agent at the maximal age – such that she dies for sure in the next period – who can only save in the riskless asset. This removes any risk in the model and the risk-sensitive model reduces to the additive one. The parameters of De Nardi et al. (2010) imply a value of \$36,000 (in 1998 dollars) for the maximal no-bequest wealth and of 0.88 for the marginal propensity to bequeath wealth in the last period of life.[16] To match these parameters, we set $\bar{x} = 8.50$ and $\theta = 56.55$.[17] See Appendix C for the formal expressions of the quantities we match and further details on the calculation of these numbers.

[16]These numbers are implied by the estimates in column 3 of Table 3 of their paper. See the discussion in Appendix D of their paper.

[17]$\bar{x}$ is expressed in units of average income, which is the unit used in solving the model. The dollar equivalent is \$396,477.

Calibration of remaining parameters to match lifecycle behavior. Four parameters are calibrated to match features of behavior over the lifecycle: the utility gap between life and death (u_l), the discount factor (β), the parameter governing risk aversion (k), and the cost of fully participating in financial markets (F). We calibrate them by matching four targets: the VSL at age 45, asset holdings at age 65, annuity holdings at 65 and the proportion of people at 65 who hold risky assets.[18]

For the VSL at age 45, we target a value of \$10m, which is suggested as a central estimate for the US by the recent review article by Kniesner and Viscusi (2019). This is close to the value of \$9.6m used by the US Department of Transportation (2016) and in the range of \$8.8m–\$10.6m, used by the Environmental Protection Agency (2011) in evaluating the Clean Air Act.

To obtain a target to match mean modeled wealth at the age of 65, we use the 2016 wave of the Survey of Consumer Finances. We select a sample of single individuals without children. There are only approximately 100 such individuals at each age – and so to increase the sample size we calculate the mean for those between 60 and 69. We winsorize at the 1st and 99th percentiles. This yields a mean of \$366,000. We also show in Appendix F.2 the robustness of our results to matching simulated wealth at 45 to a sample defined similarly to above of those between 40 and 49 (the mean for this group is \$169,000).

We target the proportion of households holding annuities. This proportion in the U.S. is extremely low (this fact is, of course, the source of the annuity 'puzzle'). Lockwood (2012) reports that 3.6% of single retirees hold annuities; Pashchenko (2013) reports 5%, also for a sample of single retirees. We target the latter, slightly higher, number.

Finally, participation cost is calibrated (jointly with the preference parameters) to deliver a stock market participation rate of 50% at age 65, which is in the range reported by Alan (2006).

In our results, we will compare our simulated profiles to an additive model that will also be calibrated to meet these targets. As the risk aversion parameter

[18]For the value of life we choose its estimated value at age 45 (and not 65) as a calibration target, as most empirical studies rely on wage-risk trade-offs, estimated on samples of workers.

k is constrained to equal zero in the additive specification, matching the targets requires us to add one degree of freedom. To do this, rather than imposing the 10% administrative load (δ), informed by the literature, we find the administrative load that would rationalize annuity demand.

Table 2: Calibration targets

Target	Value	Source
VSL	\$10m	Kniesner and Viscusi (2019)
Mean wealth at age 65	\$366,000	Survey of Consumer Finances
Proportion holding annuities at age 65	5%	Pashchenko (2013)
Stock market participation rate at age 65	50%	Alan (2006)

Notes: See the text for calculation and further details on source.

4.5 Results

Parameter estimates. Table 3 gives our baseline parameter estimates and also shows, for comparison, the parameter estimates for the additive model. The estimate for our risk-sensitive model of β, 0.966, is close to values typically estimated in lifecycle models and additionally is very close to the value (0.97) assumed by Pashchenko (2013) and Lockwood (2012), two papers which try to rationalize low annuity-demand. De Nardi et al. (2010), on whose estimates we base our parameterization of the bequest function, also use a value of 0.97. The specific values we obtain for u_l and k depend on normalization choices we made, and are best understood when thinking of trade-off between consumption and life duration (for u_l) or when looking at how agents compare different lotteries in life duration (for k). Regarding u_l, consider a setting with an agent whose consumption equals the average yearly income. Such an agent endowed with the calibrated value of u_l would be willing to give up about 7% of her consumption during her last year of life in exchange for one extra week of life. As for an interpretation of the magnitude of k, which impacts risk aversion with respect to life duration, consider the following situation. A 65 year-old agent, endowed with the set of preferences estimated in our baseline model is faced with a previously unanticipated option to undergo some surgery, knowing that the surgery would increase her life expectancy, but would involve taking a 5% risk of an immediate death. With our estimated value of k, the agent would opt for surgery only if the increase in life expectancy (taking account of the risk of dying in the 'operation') exceeds 12 months.[19] Had we set $k = 0$ (and kept all other parameters the same), the agent would opt for the operation when the increase in life expectancy is as short as 4 months.

The participation cost, at 170% of average annual income, is necessarily large to match observed (non-)participation in risky assets when the equity premium is 4%, (see, for example, Mehra and Prescott, 1985 and Kocherlakota, 1996). However, note that this is paid only once.[20] To place this quantity in perspective, for those who pay the participation cost, it represents 2.4% of lifetime consumption (discounted by the risk-free rate).

Table 3: Estimated parameters in baseline economy

Parameter	Risk-Sensitive	Additive Model
Preferences		
Inverse of IES, σ	$2.0^{\dagger}$	$2.0^{\dagger}$
Risk aversion parameter, k	0.867	$0.000^{\dagger}$
Life-death utility gap, u_l	3.570	12.020
Discount factor, β	0.966	0.953
Bequest motive strength, θ	$56.55^{\dagger}$	$56.55^{\dagger}$
Bequest luxury good, $\overline{x}$	$8.50^{\dagger}$	$8.50^{\dagger}$
Annuity administrative load, δ	$10\%^{\dagger}$	23%
Asset Markets		
Participation cost, F	170% of $\bar{y}$	175% of $\bar{y}$

Notes: One unit of consumption is equal to $\bar{y}$. Quantities indicated by a † are imposed rather than estimated.

[19]To calculate this, we find the scalar α that, when multiplied by all survival probabilities for ages after 65, exactly compensates agents, in expectation, for the loss in utility associated with the additional 5% chance of dying at the age of 65. We can then use this quantity to calculate new life expectancies taking into account both the risk of dying in the operation and the greater survival in each period if she survives.

[20]In Section F, we study participation costs paid every time the agent invests in stock, instead of once per life. In that calibration, the costs fall to 14% of average income.

Estimated profiles. Figure 1 shows profiles of mean consumption, wealth, participation in stock market and VSL over the lifecycle for a simulated sample of individuals. Results are reported for three specifications. The first two correspond to the calibrated versions of the risk-sensitive (referred to as RS) and additive (referred to as "calibrated-additive") models, whose parameters are shown in Table 3. Both models, by construction, match the quantitative targets which are marked on the graphs and predict that the proportion of individuals holding annuities is 5%. The third specification, labeled as "uncalibrated-additive" corresponds to the model obtained while keeping all parameters of the risk-sensitive case fixed, but setting $k = 0$ to recover an additive specification. We show these results to facilitate comparative statics with respect to the risk aversion parameter k.

Figure 1: Lifecycle profiles for the baseline calibration

(a) Consumption (b) Mean wealth

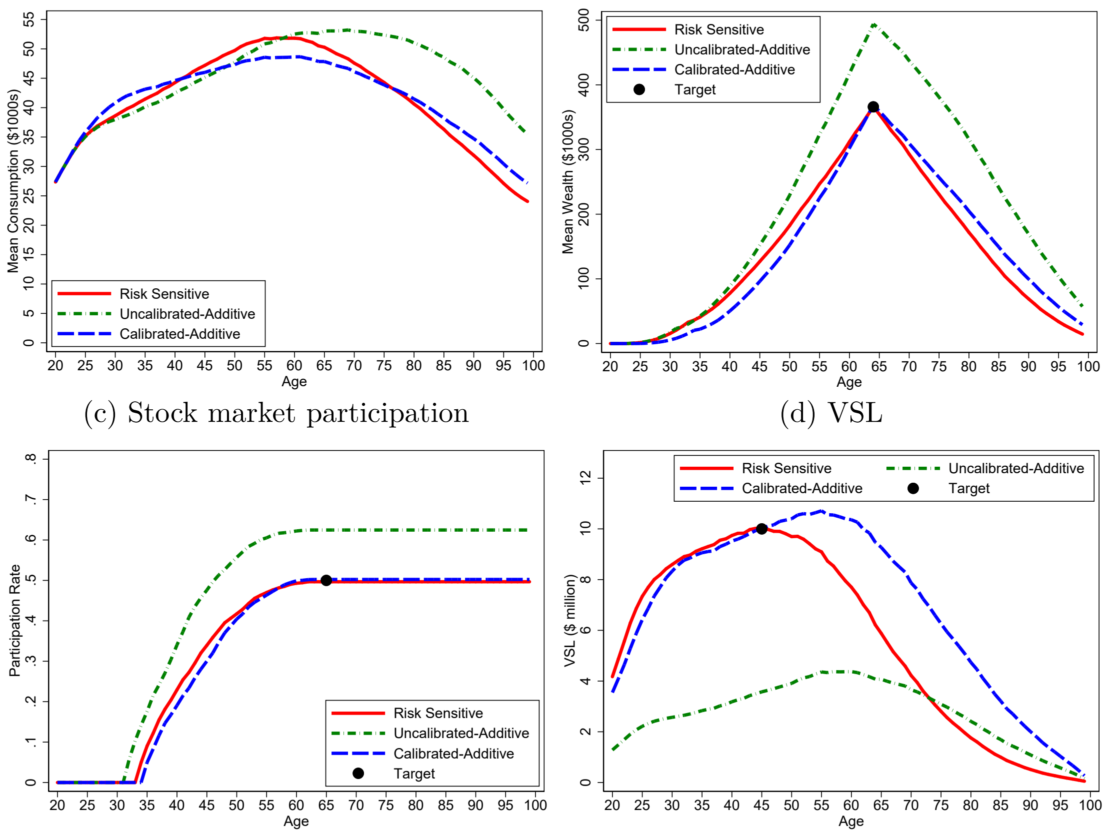

(c) Stock market participation (d) VSL

Comparative risk aversion. The RS agent is more risk averse than the "uncalibrated-additive" agent, but is identical in all other aspects. Comparing the predictions of the RS model with those of the "uncalibrated-additive" model therefore reflects the theoretical predictions of Section 3. We note however, that the framework here is a richer one than that outlined in Section 3 – in particular with the addition of income and asset return uncertainty. These additions are likely to have important effects: income uncertainty generates a precautionary savings effect that is amplified with risk aversion.[21] The negative relationship between risk aversion and savings derived in Section 3 will thus be complemented by a precautionary effect that would imply an opposed relation. Uncertainty in asset returns may also contribute to a positive or negative relationship between risk aversion and savings, depending on the intertemporal elasticity of substitution and lifecycle income profile. The overall impact of risk aversion is therefore theoretically ambiguous, with its sign depending on the magnitude of the different risks at play. Figure 1 shows, however, that the RS agent saves less than the "uncalibrated-additive" agent. Moreover the RS agent is less likely to purchase annuities (the annuity participation rate is 5% for RS agents compared to 33% for additive agents). These are in line with the predictions shown in Proposition 1.

From a quantitative standpoint, this means that the effects of income and financial risks that we added in this quantitative investigation turn out to be too small to offset the effect of mortality risk highlighted in our theory section. The effect of the mortality risk tends therefore to dominate those of other risks. We can interpret this as an indication that mortality risks loom larger for individuals in their decision making than do the other risks they face.[22]

Figure 1 shows also that RS agents have a higher VSL at all ages and are less likely to invest in stocks than the "uncalibrated-additive" agent. This simply reflects that risk aversion increases the willingness to reduce exposure to mortality and financial risks.

Overall, our results regarding the impact of risk aversion highlight that the more risk averse the individual, the more she dislikes taking risks of any kind, whether they are related to mortality, income or financial matters. It is worth

[21]See Bommier and LeGrand (2019), who show, using risk-sensitive preferences in an infinite horizon setting that there is a positive relationship between risk aversion and precautionary savings.

[22]This fact is also reflected in the very high willingness to pay for mortality risk reduction revealed by empirical studies, which we used to calibrate our model.

noting, however, that our findings strongly contrast with those of well-known studies in the HF literature, such as Gomes and Michaelides (2005, 2008), who find a positive relationship between risk aversion and stock market participation, and Inkmann et al. (2011), who find that risk aversion increases the demand for annuities. Explanations for these differences are provided in Section 5.

Comparison of the calibrated models. Let us now compare the predictions of RS and "calibrated-additive" specifications. By construction, both specifications predict the same accumulated wealth, stock market participation and annuity market participation at age 65, and the same value of life at age 45. The lifecycle profiles for consumption, wealth and stock market participation are therefore similar, though with a divergence in the VSL over the second half of the lifecycle.[23]

A fundamental difference between the specifications is in how they confront the 'annuity puzzle', that is, how they rationalize realistically low annuity demand. Low annuity demand is rationalized for the additive agent with a counterfactually high administrative load (23% compared to 10% for the RS agent) and a relatively low discount factor (0.953 compared to 0.966 in the RS case). The RS agent, on the other hand, is concerned that purchasing annuities may incur a loss in case of an early death. Although she values the benefits of holding annuities to insure against the consumption needs in the case of a long life, she also wants to have significant investments in bonds or in stocks so that the early death adverse event is mitigated by the transmission of a bequest to her heirs. A low level of annuity demand can be rationalized, even if annuities are priced at close to actuarially fair levels.

4.6 Does the precise VSL value matter?

In Appendix F, we provide an extended sensitivity analysis where we investigate the robustness of the results to the structure of the participation cost, the calibration age for the wealth and the target VSL. Here, we provide a brief discussion of the

[23]Increasing risk aversion leads to a greater willingness to pay to avoid dramatic outcomes, such as death at young age, as compared to adverse but less dramatic outcomes, such at death at old age. Increasing risk aversion therefore tends to amplify the relationship between age and VSL.

last of these – the impact of the VSL estimate, which is the main innovation of this paper. We know from our theoretical section that the choice of a plausible (and hence positive) value of the VSL is key to get the impact of risk aversion right. A natural question, given that the empirical literature on VSL does not contain a consensus estimate of its value, is how much our results would be impacted by choosing a VSL target in the upper or lower range of empirical estimates.

The value that we choose, \$10m, is approximately the level currently used by US policy-makers. To assess whether this particular choice is instrumental in the results, we estimate the parameters using the same calibration strategy and the same calibration targets as in Section 4.4, except that we consider two alternative targets for the VSL, \$7.5m and \$12.5m.

The calibrated parameters for the two VSL targets, as well as the plots of the related lifecycle profiles for consumption, wealth, VSL, and stock market participation can be found in Appendix F.1. These graphs show clearly that changing the VSL target has very little impact on individual decisions (consumption, wealth, asset market participation, and annuity purchase). The only significant difference between calibrations is on the VSL lifetime profiles.

The reason that the precise VSL plays little role in agents' choices when the VSL is large is that the RS model admits a well-defined limit when the VSL becomes infinite. To show this, we take the recursive equation (24) that defines the utility V_t, representing risk-sensitive preferences, and consider the limit where u_l tends to infinity while maintaining the product ku_l equal to a constant, denoted by κ in the following. This involves taking the limit where VSL goes to infinity while maintaining risk aversion with respect to the life duration constant. It can easily be shown through a first-order Taylor expansion that the risk-sensitive utility, V_t, can be approximated by $V_t \simeq \frac{1}{k} Y_t + X_t$, where X_t and Y_t are recursively defined by:

$$\begin{aligned} X_t = {} & (1-\beta)\frac{c_t^{1-\sigma}-1}{1-\sigma} && (28) \\ & + \frac{\beta}{\pi_t + (1-\pi_t)e^{Y_{t+1}}}\Big(\pi_t E_t[X_{t+1}] + (1-\pi_t)e^{Y_{t+1}} E_t[(1-\beta)v(x_{t+1})]\Big), \\ Y_t = {} & (1-\beta)\kappa - \beta \log\Big(\pi_t E_t\left[e^{-Y_{t+1}}\right] + (1-\pi_t)\Big). && (29) \end{aligned}$$

Under exogenous mortality, Y_t is exogenous, and maximizing V_t is thus equivalent to maximizing X_t. For a large VSL, the RS model therefore converges to the model represented by X_t. Notice that the recursive equation that defines X_t has an additive structure, similar to the usual additive specification. However, it embeds an (exogenous) age-dependent discount factor $\frac{\beta}{\pi_t+(1-\pi_t)e^{Y_{t+1}}}$ that reflects the impact of mortality and risk aversion on impatience.[24]

Since the results we obtain depend very little on the VSL calibration target, we deduce that our calibrated model is in fact relatively close to the infinite VSL limit. In particular, as long as the VSL remains large, the uncertainty regarding the precise value of the VSL is of limited concern. It also means that, to a first-order approximation, the quantitative model we use could be replaced by a standard additive model with an age-dependent discount factor. This may be helpful for further applications.

5 Relation to previous studies in the HF literature

Most papers in the HF literature rely on the additive specification, which lacks the flexibility to fully study the role of risk aversion in decision-making over the lifecycle. In such papers, the expression "risk aversion" is most often used to refer to a parameter (σ in our paper) that governs both intertemporal substitutability and risk preferences. It is, however, well understood that models that assume different IES are not comparable in terms of risk aversion (see Kihlstrom and Mirman, 1974, for instance). The findings of those papers cannot therefore be compared to ours.

The relevant comparison is with the subset of papers, such as Gomes and Michaelides (2005, 2008), Inkmann et al. (2011) and many others, which, like us, use recursive preferences to study the role of risk aversion in isolation. A key difference is that our model was designed to fit empirical estimates of the VSL, without imposing preference homotheticity, while the other papers typically use EZW preferences to obtain homothetic (and tractable) specifications without

[24] While the model obtained when maximizing X_t features age-dependent time discounting, it is time consistent. Preferences are not stationary (they depend on age, reflecting the relationship between age and mortality), but they do not exhibit preference reversals.

paying attention to the implications for the VSL.

Formally, for EZW preferences in the presence of bequests, the recursion (8) defining the utility conditional on being alive becomes:

$$V_t = \left((1-\beta)c_t^{1-\sigma} + \beta\left(E_t\left[\pi_t V_{t+1}^{1-\gamma} + (1-\pi_t)\theta x_{t+1}^{1-\gamma}\right]\right)^{\frac{1-\sigma}{1-\gamma}}\right)^{\frac{1}{1-\sigma}}, \qquad (30)$$

where, as in (26), θ determines the intensity of the bequest motive.[25] There is typically no discussion of the value of mortality risk reduction in HF papers, as they assume that mortality is exogenous. Equation (30) nevertheless implicitly assumes a specific sign for the value of mortality risk reduction. Indeed:

$$\frac{\partial V_t}{\partial \pi_t} = \beta \frac{E_t[V_{t+1}^{1-\gamma}] - \theta E_t[x_{t+1}^{1-\gamma}]}{1-\gamma}\left(E_t\left[\pi_t V_{t+1}^{1-\gamma} + (1-\pi_t)\theta x_{t+1}^{1-\gamma}\right]\right)^{\frac{\gamma-\sigma}{1-\gamma}} V_t^{\sigma},$$

which can be positive or negative. If $\gamma > 1$, as is assumed in the papers referenced above, a positive value of mortality risk reduction is obtained only if $\theta > \frac{E_t[V_{t+1}^{1-\gamma}]}{E_t[x_{t+1}^{1-\gamma}]}$. The results of Gomes and Michaelides (2005, 2008) and those of Inkmann et al. (2011) indicate that this condition does not hold (at least not always) in their simulations.[26] In particular, a negative value of mortality risk reduction is systematically obtained when there is no bequest motive ($\theta = 0$), a case considered in several instances in those papers. With a negative value of mortality risk reduction, the rate of time discounting is underestimated, and the underestimation is amplified by risk aversion. Risk aversion is then found to amplify savings. The difference in saving behaviors eventually generates differences in the propensity to pay the stock market participation cost. This explains why Gomes and Michaelides (2005, 2008) find that more risk averse agents tend to participate more frequently in the stock market. Moreover, with a negative value of mortality risk reduction, the risk of losing annuitized wealth in case of an early death is not seen as a major concern, as short lives are seen as good outcomes. This impacts the willingness to

[25] A formal derivation of equation (30) can be found in the appendix of Gomes et al. (2009), for example.

[26] One should notice, moreover, that if specification (30) were to be used with $\gamma > 1$ and a parameter θ large enough to generate positive values of mortality risk reduction, we would obtain a framework where the intensity of the bequest motive would increase the willingness to pay for mortality risk reduction: $\frac{\partial^2 V_t}{\partial \pi_t \partial \theta} > 0$. However, this would go against intuition, since deriving utility from bequest reduces the welfare gap between life and death. In well-behaved models, such as the RS or additive models, altruism has a negative impact on the value of mortality risk reduction (see equations (36) and (37) for instance) – as one would expect.

purchase annuities, which is found to increase with risk aversion (Inkmann et al., 2011). Overall, with a negative VSL, risk aversion is found to increase savings, stock market participation and annuity purchases, providing conclusions which are opposite to ours. This is fully in line with the theoretical results developed in Section 3, where we discussed the case of a counterfactual negative value of mortality risk reduction.

6 Conclusion

Inspired by Samuelson (1937), economic contributions on intertemporal choice have most often relied on models that assume time-additive preferences. While time-additive preferences offered an elegant framework to formalize insightful theories, such as Modigliani's lifecycle consumption theory, they have some serious limitations. One of the caveats associated with the time-additive model is its lack of flexibility, in particular, the fact that its use means that risk aversion and intertemporal substitutability cannot be disentangled. This was underlined both by theoreticians (Epstein and Zin, 1989) and experimentalists (Andersen et al., 2008). Theoretical contributions, such as those of Epstein and Zin (1989) and Weil (1990), have addressed this limitation by introducing a recursive framework that is both tractable and flexible. These EZW preferences have met with a remarkable success. While EZW preferences were initially developed to deal with infinitely-lived agents, their adaptation to finite and random horizon settings encountered some serious difficulties. Some papers in the VoL literature introduced specifications that are ill-defined and yield counterfactual implications. Others, in the HF literature, implicitly assumed that all agents have a preference for death over survival. Papers in both literatures imported the non-monotonicity features that are inherent to EZW preferences. The consequences of these features were discussed in Section 2.

In the current paper, we propose a framework that can model lifecycle behaviors with recursive preferences that are well-defined, flexible, monotone and can simultaneously match realistic (positive) values for mortality risk reduction and plausible lifecycle profiles for consumption and wealth. We outline how using such a specification facilitates new and intuitive insights on the role of risk aversion, affording a potential explanation for several documented – though imperfectly understood – household behaviors, such as the low demand for annuities. We believe that in an era of on-going demographic changes, the economics of aging will remain a key research area, with new forms of risks becoming increasing sources of concern (not least the increased likelihood of expensive long-term care at the oldest ages). While time-additive preferences have facilitated valuable insights, the field would benefit from a framework that allows a full investigation of the role of risk aversion. The approach put forward in this paper could serve as a foundation of such a framework for the HF or the VoL literature.

Appendix

A Accounting for the taste for Russian-roulette lotteries

A point regularly made against the additive model is that in the usual case where the intertemporal elasticity of substitution is below one, there necessarily exists a consumption level such that the agent prefers being dead to being alive with that level of consumption. Formally, if one assumes $\frac{-cu(c)'}{u''(c)}$ to be constant (say equal to $\frac{1}{\sigma}$), one must have $u(c) = u_0 + K\frac{c^{1-\sigma}}{1-\sigma}$ for some constant K, which has to be positive for utility to be increasing with consumption. Then, if $\sigma > 1$, one necessarily has $u(c) < 0$ for small c, meaning that life becomes worse than death. Agents provided with a wealth endowment that would only allow them to sustain such "worse than death" consumption levels, would prefer to enter Russian-roulette games, whose outcome is either "death" or a "better than death" consumption level.

But why do we assume that $\frac{-cu(c)'}{u''(c)}$ is constant over the whole set $\mathbb{R}_+$? Empirical studies typically consider a setting in which consumption lies on a compact interval $[c_{min}, c_{max}]$, with $c_{min} > 0$. Thus, if we know something about the IES, it is at most that it is constant on an interval $[c_{min}, c_{max}]$. There is no difficulty in constructing a well-behaved utility function such that $\frac{-cu(c)'}{u''(c)}$ is constant over $[c_{min}, c_{max}]$ and $u(c) > 0$ over $\mathbb{R}_+$. Consider for example $\sigma > 1$, $\tilde{\sigma} < 1$, a constant $u_l > \frac{c_{min}^{1-\sigma}}{1-\tilde{\sigma}} - \frac{c_{min}^{1-\sigma}-1}{1-\sigma}$ and the utility given by:

$$u(c) = \begin{cases} u_l + \frac{c^{1-\sigma}-1}{1-\sigma} & \text{for all } c \geq c_{min}, \\ \tilde{u}_l + c_{min}^{\tilde{\sigma}-\sigma}\frac{c^{1-\tilde{\sigma}}-1}{1-\tilde{\sigma}} & \text{for all } c \leq c_{min}, \end{cases} \tag{31}$$

where $\tilde{u}_l = u_l + \frac{c_{min}^{1-\sigma}-1}{1-\sigma} + \frac{c_{min}^{\tilde{\sigma}-\sigma}-c_{min}^{1-\sigma}}{1-\tilde{\sigma}}$. The utility function is continuous and positive on $\mathbb{R}_+$ (including in 0). Moreover, we have:

$$\begin{cases} u'(c) = c^{-\sigma} & \text{for all } c \geq c_{min}, \\ u'(c) = c_{min}^{-\sigma}\left(\frac{c}{c_{min}}\right)^{-\tilde{\sigma}} & \text{for all } c \leq c_{min}, \end{cases}$$

implying that the marginal utility function is also continuous, positive and decreasing (which implies in turn that u is concave). Such a utility function is therefore "well-behaved" (increasing, concave and continuously differentiable). It would be consistent with the empirical evidence that the IES is constant and equal to $\frac{1}{\sigma}$ (thus below 1) when consumption lies in $[c_{min}, c_{max}]$, as well as with the evidence that nobody likes to play the Russian-roulette-style lotteries.

Finally, note that given that allowing for a non-constant elasticity of substitution as suggested above makes it possible to avoid the existence of consumption levels for which utility becomes negative, and therefore the taste for Russian-roulette lotteries, we consider it to be a theoretical curiosity with limited implications for applied work. That is due to the fact that, if one restricts attention to settings in which consumption always remain above c_{min}, what is assumed regarding the IES below c_{min} is mostly irrelevant. This explains why the simulations developed in the paper simply assume an instantaneous utility of the form $u_l + \frac{c^{1-\sigma}-1}{1-\sigma}$ without introducing the more complex form shown in (31).

B Proofs

B.1 Proof of Proposition 1

Denote $\tilde{\theta} = \frac{\theta}{\beta}$, $\frac{u}{1-\beta} = u_l + \frac{\left(R^f(b+\frac{a}{\pi_0})\right)^{1-\sigma}}{1-\sigma}$, $\frac{v}{1-\beta} = \theta\frac{\left(R^f b\right)^{1-\sigma}}{1-\sigma}$ and define $\tilde{u}$ by:

$$e^{-k\tilde{u}} = \pi e^{-ku} + (1-\pi_0)e^{-kv}. \tag{32}$$

The solution of the optimization problem fulfills:

$$a = w_0 \frac{1-\tilde{\theta}^{\frac{1}{\sigma}} e^{\frac{k}{\sigma}(u-v)}}{1+\frac{(\beta R^f)^{-\frac{1}{\sigma}}}{\pi_0} e^{\frac{k}{\sigma}(u-\tilde{u})}}, \text{ and } b = \frac{w_0}{\pi_0} \frac{\tilde{\theta}^{\frac{1}{\sigma}} e^{\frac{k}{\sigma}(u-v)}}{1+\frac{(\beta R^f)^{-\frac{1}{\sigma}}}{\pi_0} e^{\frac{k}{\sigma}(u-\tilde{u})}}. \tag{33}$$

Equations in (33) imply:

$$b + \frac{a}{\pi_0} = \frac{w_0}{\pi_0} \frac{1}{1+\frac{(\beta R^f)^{-\frac{1}{\sigma}}}{\pi_0} e^{\frac{k}{\sigma}(u-\tilde{u})}}.$$

Now consider the impact of an increase in k. From (32) we have $e^{k(u-\tilde{u})} = \pi_0 + (1-\pi_0)e^{k(u-v)}$, such that $\frac{\partial}{\partial k}\left(k(u-v)\right)$ and $\frac{\partial}{\partial k}\left(k(u-\tilde{u})\right)$ have the same sign. Similarly, $e^{k(v-\tilde{u})} = \pi_0 e^{k(v-u)} + 1 - \pi_0$, implying that $\frac{\partial}{\partial k}\left(k(\tilde{u}-v)\right)$ also has the same sign as $\frac{\partial}{\partial k}(k(u-v))$.

Assume now that $u-v>0$ and $k \geq 0$. Imagine that $k(u-v)$ decreases when k increases. We would have $b+\frac{a}{\pi_0}$ (thus u) increasing, and b (thus v) decreasing. Thus $(u-v)$ would increase, yielding a contradiction, since $\frac{\partial}{\partial k}(k(u-v)) = (u-v)+k\frac{\partial}{\partial k}(u-v)$ with both terms that would be positive. Thus $k(u-v)$ increases with k, implying that a decreases with k, b increases, $b+\frac{a}{\pi_0}$ and $b+a$ decreases.

Assume now that $u-v<0$ and $k \geq 0$. Imagine that $k(u-v)$ increases when k increases. We would have $b+\frac{a}{\pi_0}$ (thus u) decreasing, and b (thus v) increasing. Thus $u-v$ would decrease, yielding a contradiction, since $\frac{\partial}{\partial k}(k(u-v)) = (u-v)+k\frac{\partial}{\partial k}(u-v)$ with both terms that would be negative. Thus $k(u-v)$ decreases with k, which implies that a increases with k, b decreases, $b+\frac{a}{\pi_0}$ and $b+a$ increase.

B.2 Proof of the example about "investment in precaution"

In each state $i=h,s$, the best response e_i^* maximizes:

$$e_i^* = \arg\max_{e\in[0,w_0]} \left((1-\beta)(w_0-e)^{1-\sigma} + (1-\beta)\beta\pi_i(e)^{\frac{1-\sigma}{1-\gamma}} c_1^{1-\sigma}\right)^{\frac{1}{1-\sigma}},$$

which implies the following first-order condition:

$$(w_0-e_i^*)^{-\sigma} = \frac{\beta}{1-\gamma}\pi_i'(e_i^*)\pi_i(e_i^*)^{\frac{\gamma-\sigma}{1-\gamma}} c_1^{1-\sigma}.$$

When the value of mortality risk reduction is negative, or equivalently when $\gamma>1$, a necessary condition for an interior solution to exist is $\pi_i'(e)<0$: in other words, extra precaution should diminish the survival probability. If the opposite holds, i.e., if $\pi_i'(e)>0$, the optimal choice is always zero precaution. Let us assume that conditions hold for an interior solution to exist in each state.

We consider a case in which best responses e_h^* and e_s^* are identical in both states. This holds if:

$$\pi_h'(e_h^*)\pi_h(e_h^*)^{\frac{\gamma-\sigma}{1-\gamma}} = \pi_s'(e_s^*)\pi_s(e_s^*)^{\frac{\gamma-\sigma}{1-\gamma}}. \tag{34}$$

In the absence of perfect foresight, the agent solves the following program:

$$\max_{e\geq 0}\left((1-\beta)(w_0-e)^{1-\sigma} + \beta(1-\beta)\left(p_h\pi_h(e)+p_s\pi_s(e)\right)^{\frac{1-\sigma}{1-\gamma}} c_1^{1-\sigma}\right)^{\frac{1}{1-\sigma}}.$$

If it exists, the interior optimal effort level e^* solves:

$$(w_0 - e^*)^{-\sigma} = \frac{\beta}{1-\gamma} \left(p_h \pi_h'(e^*) + p_s \pi_s'(e^*)\right) \left(p_h \pi(e^*) + p_s \pi_s(e^*)\right)^{\frac{\gamma-\sigma}{1-\gamma}} c_1^{1-\sigma}.$$

For the optimum to be the same as the (identical) perfect-foresight best-responses we should have, using (34), $E[\pi^{\frac{\sigma-\gamma}{1-\sigma}}]E[\pi]^{-\frac{\sigma-\gamma}{1-\sigma}} = 1$, where $E[\cdot]$ is the expectation with respect to p_h and p_s. When $p_h \in (0,1)$, Jensen's Inequality implies that the previous inequality cannot hold (unless $\gamma = \sigma$, which corresponds to the risk-sensitive or the additive model). In conclusion, since $e^* \neq e_h^* = e_s^*$, the agent takes a dominated action as choosing e_h^* or e_s^* would provide a higher welfare in states h and s.

B.3 Deriving VSL expressions

We rewrite the problem using wealth and portfolio shares. Let $w_t = A_{t-1} + R^f b_{t-1} + R_t^s s_{t-1}$ be the beginning-of-period wealth at date t. We also denote by $\omega_t = q_t a_t + b_t + s_t$ the total saving choice, by $\alpha_t^b = \frac{b_t}{\omega_t}$ the share in bonds and by $\alpha_t^s = \frac{s_t}{\omega_t}$ the share in stocks. The program of the alive agent can be rewritten as:

$$\begin{aligned}
V_t(w_t, A_{t-1}, \eta_{t-1}, \zeta_{t-1}) = &\max_{c_t \geq 0, \omega_t \geq 0, (\alpha_t^b, \alpha_t^s) \in [0,1]^2} (1-\beta) u(c_t) \\
&- \frac{\beta}{k} \log\left(\pi_t E_t\left[e^{-kV_{t+1}(w_{t+1}, A_t, \eta_t, \zeta_t)}\right] + (1-\pi_t) E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]\right),
\end{aligned}$$

subject to the following constraints:

$$\begin{aligned}
y_t + w_t &= c_t + \omega_t + 1_{\eta_t=1} 1_{\eta_{t-1}=0} F, \\
w_t &= A_{t-1} + \omega_t \left(\frac{1}{q_t} + \left(R^f - \frac{1}{q_t}\right)\alpha_t^b + \left(R_{t+1}^s - \frac{1}{q_t}\right)\alpha_t^s\right), \\
x_t &= \omega_t \left(R^f \alpha_t^b + R_{t+1}^s \alpha_t^s\right).
\end{aligned}$$

The envelope theorem yields $\frac{\partial V_t}{\partial w_t} = \frac{\partial V_t}{\partial c_t} = (1-\beta)u'(c_t)$. Furthermore:

$$\begin{aligned}
\frac{\partial V_t}{\partial \pi_t} = &-\frac{\beta}{k} \frac{\left(E_t\left[e^{-kV_{t+1}(w_t, A_t, \eta_t, \zeta_t)}\right] - E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]\right)}{\pi_t E_t\left[e^{-kV_{t+1}(w_{t+1}, A_t, \eta_t, \zeta_t)}\right] + (1-\pi_t) E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]} \\
&+ \beta \frac{\pi_t E_t\left[(1-\beta)u'(c_{t+1})\omega_t \left(1 - \alpha_t^b - \alpha_t^s\right) \frac{\partial(1/q_t)}{\partial \pi_t} e^{-kV_{t+1}(w_t, A_t, \eta_t, \zeta_t)}\right]}{\pi_t E_t\left[e^{-kV_{t+1}(w_{t+1}, A_t, \eta_t, \zeta_t)}\right] + (1-\pi_t) E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]}.
\end{aligned} \tag{35}$$

We have $\frac{\partial(1/q_t)}{\partial\pi_t} = -\frac{1}{q_t^2}\frac{\partial q_t}{\partial\pi_t}$, with $\frac{\partial q_t}{\partial\pi_t} = (1+\delta)\frac{\partial}{\partial\pi_t}\sum_{\tau=1}^{\infty}\frac{\prod_{s=t}^{t+\tau-1}\pi_s}{\left(R^f\right)^{\tau}} = \frac{q_t}{\pi_t}$. This gives $\frac{\partial(1/q_t)}{\partial\pi_t} = -\frac{1}{\pi_t q_t}$. We deduce from (35):

$$
\begin{aligned}
\frac{\partial V_t}{\partial\pi_t} = & -\frac{\beta}{k}\frac{\left(E_t\left[e^{-kV_{t+1}(w_t,A_t,\eta_t,\zeta_t)}\right]-E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]\right)}{\pi_t\pi_t E_t\left[e^{-kV_{t+1}(w_t,A_t,\eta_t,\zeta_t)}\right]+(1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]} \\
& -\frac{\beta(1-\beta)E_t\left[u'(c_{t+1})\frac{\omega_t}{q_t}\left(1-\alpha_t^b-\alpha_t^s\right)e^{-kV_{t+1}(w_t,A_t,\eta_t,\zeta_t)}\right]}{\pi_t E_t\left[e^{-kV_{t+1}(w_t,A_t,\eta_t,\zeta_t)}\right]+(1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]}.
\end{aligned}
$$

With $\frac{\omega_t}{q_t}\left(1-\alpha_t^b-\alpha_t^s\right) = a_t$, we finally obtain:

$$
\begin{aligned}
VSL_t = & \frac{\beta}{1-\beta}\frac{1}{c_t^{-\sigma}}\frac{\left(-\frac{E_t\left[e^{-kV_{t+1}(b_t,A_t,s_t,\eta_t,\zeta_t)}\right]-E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]}{k}\right)}{\pi_t E_t\left[e^{-kV_{t+1}(b_t,A_t,s_t,\eta_t,\zeta_t)}\right]+(1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]} \\
& -\beta a_t\frac{E_t\left[\left(\frac{c_{t+1}}{c_t}\right)^{-\sigma}e^{-kV_{t+1}(b_t,A_t,s_t,\eta_t,\zeta_t)}\right]}{\pi_t E_t\left[e^{-kV_{t+1}(b_t,A_t,s_t,\eta_t,\zeta_t)}\right]+(1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]},
\end{aligned}
\tag{36}
$$

and in the additive case, by continuity for $k \to 0$:

$$
\begin{aligned}
VSL_t^{add} = & \frac{\beta}{1-\beta}\frac{1}{c_t^{-\sigma}}\left(V_{t+1}^{add}(b_t,A_t,s_t,\eta_t,\zeta_t)-k(1-\beta)v(x_{t+1})\right) \\
& -\beta a_t E_t\left[\frac{u'(c_{t+1})}{u'(c_t)}\right].
\end{aligned}
\tag{37}
$$

Note that, in both the RS and additive cases, as well as in real life, the VSL is impacted by uncertainty regarding future consumption. Such aspects are usually ignored in the VSL literature, but can be taken into account with our approach. This makes the calibration more realistic, even though some other sources of uncertainty (as those related to health) which could also be taken into account are not included.

C Calibration of bequest parameters

To calibrate bequest parameters, we compute the marginal propensity to consume (MPC) and the maximal no bequest wealth for an agent in the last period before certain death. We further assume that she can only invest in a riskless bond. Up to a normalization by $(1-\beta)^{-1}$, her program can be expressed as follows:

$$
\max_{s\in\mathbb{R}} \; u(W-b)+\beta v(R^f b),
$$

where b is the riskless saving. There are two cases for first-order conditions.

1. No bequest is left if and only if $u'(W) \geq \beta R^f v'(0)$.

2. A positive bequest is left otherwise, and is determined by:

$$u'(W-b) = \beta R^f v'(R^f b). \tag{38}$$

   Equation (38) implies that, if the maximal wealth that yields no bequest is denoted as W_0, it follows that $u'(W_0) = \beta R^f v'(0)$.

Using the expressions (25) and (26) for instantaneous utility functions, the previous equation becomes:

$$W_0 = \kappa \overline{x}, \text{ with } \kappa = \left(\beta R^f \theta\right)^{-\frac{1}{\sigma}}. \tag{39}$$

Now, let us assume that $W \geq W_0$. The FOC (38) implies the following expression for the optimal bequest:

$$b = \frac{W - \kappa \overline{x}}{1 + \kappa R^f}, \tag{40}$$

as well the following MPC $\frac{1}{W-b}\frac{\partial c}{\partial W} = R^f \frac{1}{R^f b}\left(1 - \frac{\partial c}{\partial W}\right)$. Using (40), we finally get:

$$\frac{\partial c}{\partial W} = \frac{R^f \kappa}{1 + R^f \kappa}, \text{ with: } \kappa = \left(\beta R^f \theta\right)^{-\frac{1}{\sigma}}. \tag{41}$$

D Details on the computational implementation

There is no analytical solution to the agents' problem outlined in Section 4. We solve the model and obtain decision rules numerically and then use those decision rules to simulate the agents' behavior. The next two subsections describe, respectively, the solution of the model and the simulation of decision rules.

D.1 Model solution

While alive, the agent maximizes her intertemporal utility by choosing a feasible allocation $(c_t, b_t, a_t, s_t, \eta_t)_{t \geq 0}$ in the set $\mathcal{A}$. The utility V_t of the alive agent at age t depends on five state variables: the beginning-of-period holdings in bonds b_{t-1}, annuities A_{t-1} and stocks s_{t-1}; the stochastic component of labor income, ζ_{t-1}; and the stock market participation status, $\eta_{t-1} \in \{0,1\}$. The last of these is discrete, while the first four are continuous. Given that annuity purchase may only occur in period $T_R - 1$, we have $A_t = 0$ for all $t < T_R$ and $A_t = a_{T_R-1}$ for $t \geq T_R$. Since there exists a maximal age for the agent, T_M, we solve the model by iterating on the value function, starting from the last period of life. Utility maximization involves solving:

$$V_{T_M}(b_{T_m-1}, A_{T_m-1}, s_{T_m-1}, \eta_{T_m-1}, \zeta_{T_m-1}) = \max_{(c_{T_M}, b_{T_M}, a_{T_M}, s_{T_M}, w_{T_M}, \eta_{T_M}) \in \mathcal{A}} (1-\beta)u(c_{T_m}) + E_t\left[e^{-k(1-\beta)v(x_{T_m+1})}\right], \tag{42}$$

subject to the constraints, which we don't restate here, outlined in Section 4.1. Due to the presence of stocks in bequest, the continuation utility in case of death is uncertain. Both the instantaneous utility function for period $T_m(u(x_{T_m}))$ and the utility obtained when dying and bequeathing $x_{T_{M+1}}(v(x_{T_{M+1}}))$ are known and the model is solved for a discrete set of points on a grid. This gives us knowledge of V_{T_m} at a subset of the points in the state space and allows us to approximate V_{T_m} as $\hat{V}_{T_m}$ at all points. With this approximation in hand we solve an approximation to problem (42) for period $T_M - 1$ and then, iteratively, for all preceding periods t:

$$V_t(b_{t-1}, A_{t-1}, s_{t-1}, \eta_{t-1}, \zeta_{t-1}) = \max_{(c_t, b_t, a_t, s_t, w_t, \eta_t) \in \mathcal{A}} (1-\beta)u(c_t) - \frac{\beta}{k}\log\left(\pi_t E_t\left[e^{-k\hat{V}_{t+1}(b_t, A_t, s_t, \eta_t, \zeta_t)}\right] + (1-\pi_t)E_t\left[e^{-k(1-\beta)v(x_{t+1})}\right]\right). \tag{43}$$

Note that the only difference between the maximand in (43) and that in our household's problem is that the continuation value function is an approximate value function. We will now, briefly, discuss four features of the numerical procedure: i) the discretization of the continuous variables, ii) the integration of the value function, iii) the approximation method for evaluating the value function at points outside the discretized state space and iv) how the optimization is carried out.

Discretization of the continuous variables. We define a variable, total liquid wealth, which is the sum of bond and stock holdings at the start of a period. We define a grid of 54 points from \$0 to \$10m such that the gaps between successive grid points are smaller at lower levels of wealth, where the curvature of the value function will be greater. We define a grid of 36 points for annuity income from \$0 to \$800,000 such that the gaps between successive grid points are smaller at lower levels of income. Earnings are placed on a grid of 7 points each year using the procedure introduced by Tauchen (1986).

Integration. There are three risks facing households: mortality, earnings and financial risks. The risks are independent. Realizations for the first of these are naturally discrete and integration involves a simple weighted average. For the latter two, we define a discrete set of possible realizations and integrate over outcomes using the Tauchen (1986) procedure.

Approximation. To evaluate the value function at points other than those in the discrete sub-set of points, we use linear interpolation in multiple dimensions.

Optimization. Every period households make up to four choices. They decide on how much to consume, how much to save in each of bonds and the risky asset, whether to pay the participation charge (if they have not previously done so) and how much of an annuity income stream to purchase (in the period before retirement). Our problem is not globally concave, so our optimization of the household's decision problem cannot fully rely on local approaches. We therefore start by discretizing the choice variables. We define three grids on the unit interval, S_a^t, S_s, S_c, which will represent shares of available resources dedicated to purchases of annuities, portfolio shares in stocks and consumption respectively. All grids have equally spaced nodes from 0 to 1 (except that in periods other than period T_{R-1}, the grid for annuities has only one element – 0 – as no annuity purchase is possible in those periods). We evaluate the household's objective function at each combination of $(s_{a,i}^t)_{i=1,\ldots,I^t}$ in S_a^t, $(s_{s,j})_{j=1,\ldots,J}$ in S_s, and $(s_{c,k})_{k=1,\ldots,K}$ in $S_{c,k}$. In particular, defining the total available resources as $R_t = y_t + w_t$ (the sum of income and wealth) in time t, we set the annuity choice to $s_{a,i}^t R_t$, the level of saving in the risky asset set to $s_{s,j}(1 - s_{a,i}^t)R_t$, the level of consumption as $s_{c,k}(1 - s_{s,j})(1 - s_{a,i}^t)R_t$ and the savings in the bond as $(1 - s_{c,k})(1 - s_{s,j})(1 - s_{a,i}^t)R_t$. We find the combination of the points that yields the maximum value to the household. This is our candidate optimal decision. We then do a further local search for the split between consumption and the bond using golden section search. Formally, taking the candidate maximum as indexed by s^t_{a,i^*} in S^t_a, s_{s,j^*} in S_s, s_{c,k^*} in $S_{c,k}$, we look for the utility-maximizing split between consumption and the bond in the interval for consumption of $[s_{c,\max\{1,k^*-1\}}(1-s_{s,j^*})(1-s^t_{a,i^*})R, s_{c,\min\{K,k^*+1\}}(1-s_{s,j^*})(1-s^t_{a,i^*})R]$.

For points in the state space where individuals have not already paid the participation charge we implement this procedure twice, once assuming they pay the charge, once assuming they do not. The maximum of these indicates the decision rule.

This procedure yields decision rules as a function of the vector of state variables $\boldsymbol{X}_t$: $\hat{c}(\boldsymbol{X}_t)_t, \hat{b}(\boldsymbol{X}_t)_t, \hat{s}(\boldsymbol{X}_t)_t, \hat{a}(\boldsymbol{X}_t)_t, \hat{\eta}(\boldsymbol{X}_t)_t$.

D.2 Simulation of Profiles

Once decision rules $\hat{c}(\boldsymbol{X}_t)_t, \hat{b}(\boldsymbol{X}_t)_t, \hat{s}(\boldsymbol{X}_t)_t, \hat{a}(\boldsymbol{X}_t)_t, \hat{\eta}(\boldsymbol{X}_t)_t$ are obtained, we simulate a data set for 3,000 individuals. We do this as follows:

1. Initial values for wealth are set to 0.

2. Earnings draws for the first period of economic life are drawn randomly for each individual. Using these values of the state variables and the decision rules we can obtain optimal behavior in the first period.

3. We draw a equity price shock to apply to any equity holdings held at the end of the period.

4. Optimal behavior, the rate of return and the inter-temporal budget constraint yield the state variables for period 2.

5. We repeat steps (2) to (4) to obtain optimal behavior and subsequent state variables for each age up to 100. In most time periods, individuals will have realizations of the continuous state variables that are off this grid. Our approach here is to solve the individual's problem for decision rules as we do in the solution stage.

E Lifecycle productivity

The shape of the deterministic age-productivity profile is taken from Harenberg and Ludwig (2019), who compute it from PSID data using the method in Huggett et al. (2011) US$.[27] Figure 2 displays our average earnings profile.

Figure 2: Average earnings over the lifecycle

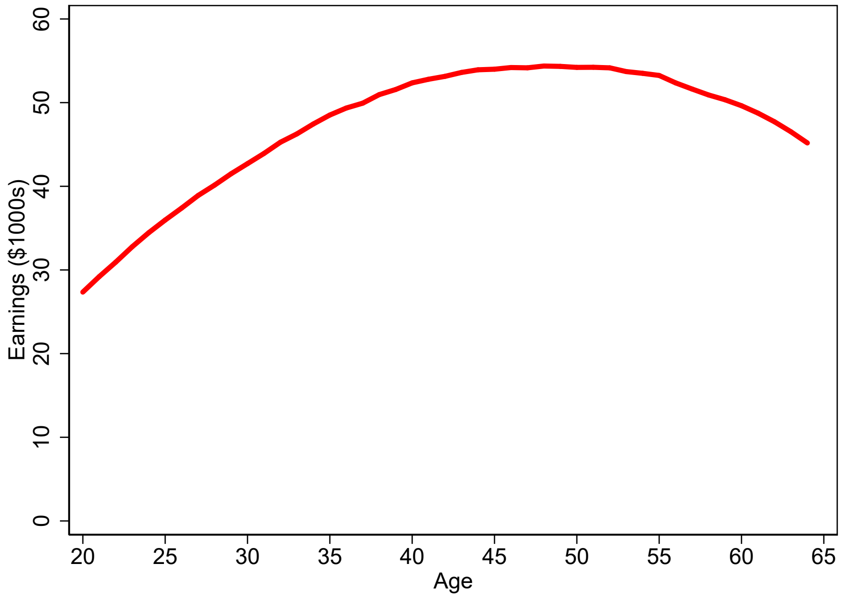

F Sensitivity analysis

F.1 High and low targets for the VSL

In this section we discuss the sensitivity of our results with respect to the VSL target. We calibrate the model using the same targets as in the baseline calibration (see Section 4.4), except that we consider two different targets for the VSL: a high VSL target of 12.5 million USD (+25% compared to the baseline target) and a low VSL target of 7.5 million (−25% compared to the baseline target). Calibration values can be found in Table 4. We report the parameter values for the additive and the RS models, for the two VSL targets.

For the additive model, Table 4 shows, as expected, that the calibration is not sensitive to the VSL target, except for the life-death utility gap u_l. This results

[27]Note that we multiply labor income with $\bar{y}$, thereby making sure that average labor income in 2016 is matched, cf. equation 14.

Table 4: Parameter calibration with low and high VSL targets

Parameter	Low VSL RS	Low VSL Additive	High VSL RS	High VSL Additive
Preferences				
Risk aversion parameter, k	1.49	$0.00^{\dagger}$	0.60	$0.00^{\dagger}$
Life-death utility gap, u_l	2.30	9.03	4.81	15.01
Discount factor, β	0.965	0.953	0.966	0.953
Bequest motive strength, θ	$56.55^{\dagger}$	$56.55^{\dagger}$	$56.55^{\dagger}$	$56.55^{\dagger}$
Bequest luxury good, $\overline{x}$	$8.50^{\dagger}$	$8.50^{\dagger}$	$8.50^{\dagger}$	$8.50^{\dagger}$
Annuity admin. load, δ	$10\%^{\dagger}$	23%	$10\%^{\dagger}$	23%
Asset markets				
Participation cost, F	161% of $\bar{y}$	175% of $\bar{y}$	167% of $\bar{y}$	175% of $\bar{y}$

Notes: One unit of consumption is equal to $\bar{y}$. Quantities indicated by a † are imposed rather than estimated.

from the invariance of the additive model up to a positive affine transformation. Therefore, individual choices with either a low or a high VSL target remain unchanged from those in the baseline model.

For the RS model, the calibration does depend on the VSL target and Table 4 shows that parameter values other than u_l differ when the calibration target changes. However, despite these different parameter values, individual lifetime profiles are almost indistinguishable across calibrations, as can be seen in Figure 3. These three profiles in each graph correspond to the RS model for the three VSL targets (the baseline, the low and the high values). As explained in the main text, except for the VSL lifecycle profiles that differ by construction, the lifecycle profiles for the consumption, wealth and stock market participation are basically the same for the three calibrations. As long as the VSL target is sufficiently high, the exact VSL value has little impact on individual choices, since the RS model can be approximated by a limit model, which is additive with an age-dependent discount factor (see its formal definition in equations (28) and (29)).

Figure 3: Lifecycle profiles for different VSL targets

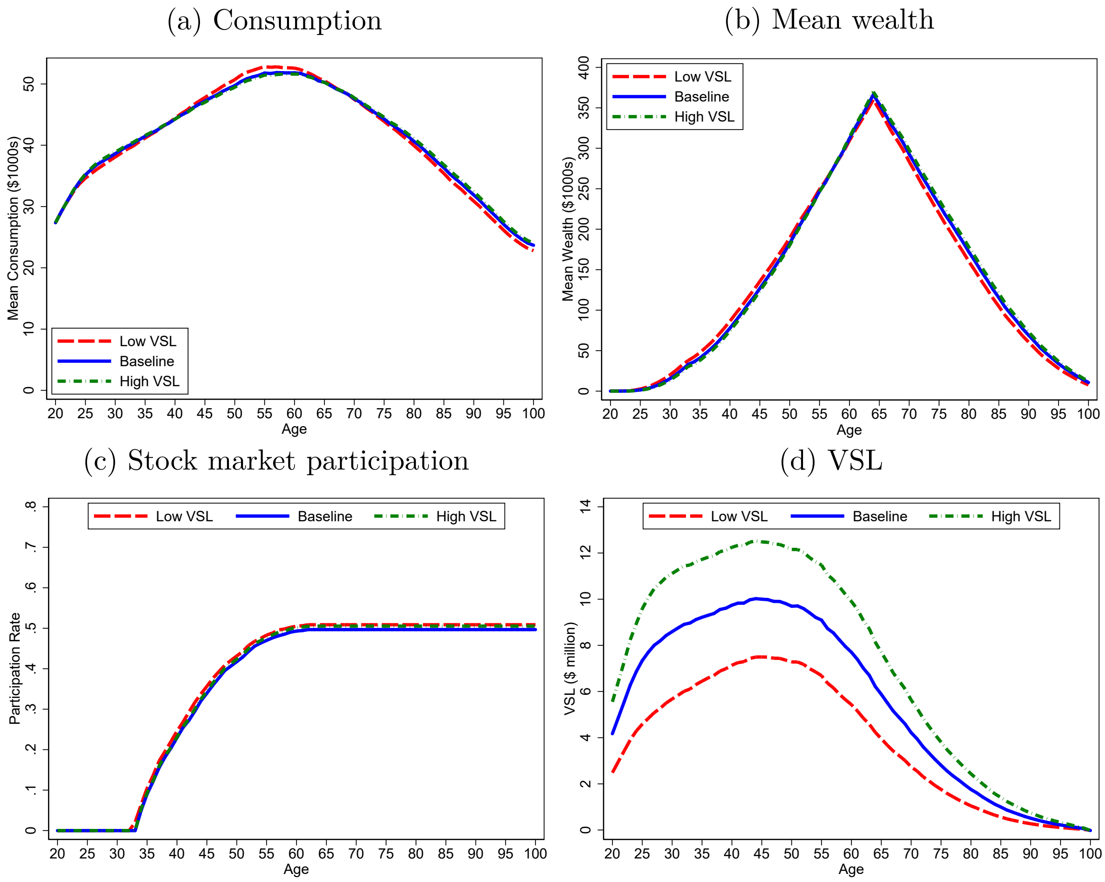

F.2 Changing the age for the wealth target from age 65 to age 45

In our baseline calibration, age 65 plays a key role: our calibration targets for wealth, the stock market participation, and the annuity holdings are matched at that age. This choice is motivated by the fact that age 65 is the exogenous retirement age in the model, and the annuity decision is only made once, in the period before retirement. Here we show the sensitivity of our results to the calibration age for the wealth target, which we change from age 65 to age 45. All other targets remain unchanged (including VSL), and both the RS and the additive models are still calibrated following the procedure described in Section 4.4.

The details of the calibration parameters for the two models can be found in Table 5. The main differences in this specification's calibration concern the participation cost, F, and the annuity administrative load, δ. First, the administrative load, which remains 10% by construction for the RS model, is now required to be

Table 5: Parameter calibration with age 65 for the wealth target

Parameter	Risk-Sensitive	Additive Model
Preferences		
Risk aversion parameter, k	1.90	0.00†
Life-death utility gap, u_l	2.49	9.68
Discount factor, β	0.977	0.976
Bequest motive strength, θ	56.55†	56.55†
Bequest luxury good, $\overline{x}$	8.50†	8.50†
Annuity administrative load, δ	10%†	54%
Asset Markets		
Participation cost, F	210% of $\bar{y}$	330% of $\bar{y}$

54% in the additive model in order to match annuity demand. This compares to a level of 23% in the age 65 calibration (see Table 3). Second, while the stock market participation cost slightly increases for the RS model to 210%, from 170% for the age 65 calibration (Table 3, again), for the additive model, the participation cost increase is much bigger, to 330% from 175% for the age 65 calibration.

In conclusion, targeting wealth at age 45 requires more substantial financial market imperfections – a higher participation cost and a higher annuity load – than in the baseline calibration when targeting wealth at age 65. This sensitivity exercise reinforces our conclusions in Section 4.5 regarding the comparison between RS and additive models.

Figure 4 reports the lifetime profiles for consumption, average wealth, VSL and stock market participation for the baseline and for the case when age 45 is targeted. Quantitative differences are modest.

F.3 Participation cost every period

We here assume that the stock market participation cost is paid in every period in which the agent trades stocks. This follows Fagereng et al. (2017), among others. The budget constraint (17) when alive becomes:

$$c_t + b_t + s_t + F1_{s_t>0} = y_t + R^f b_{t-1} + R^s_t s_{t-1}.$$

Figure 4: Lifecycle profiles for different targets regarding median wealth age

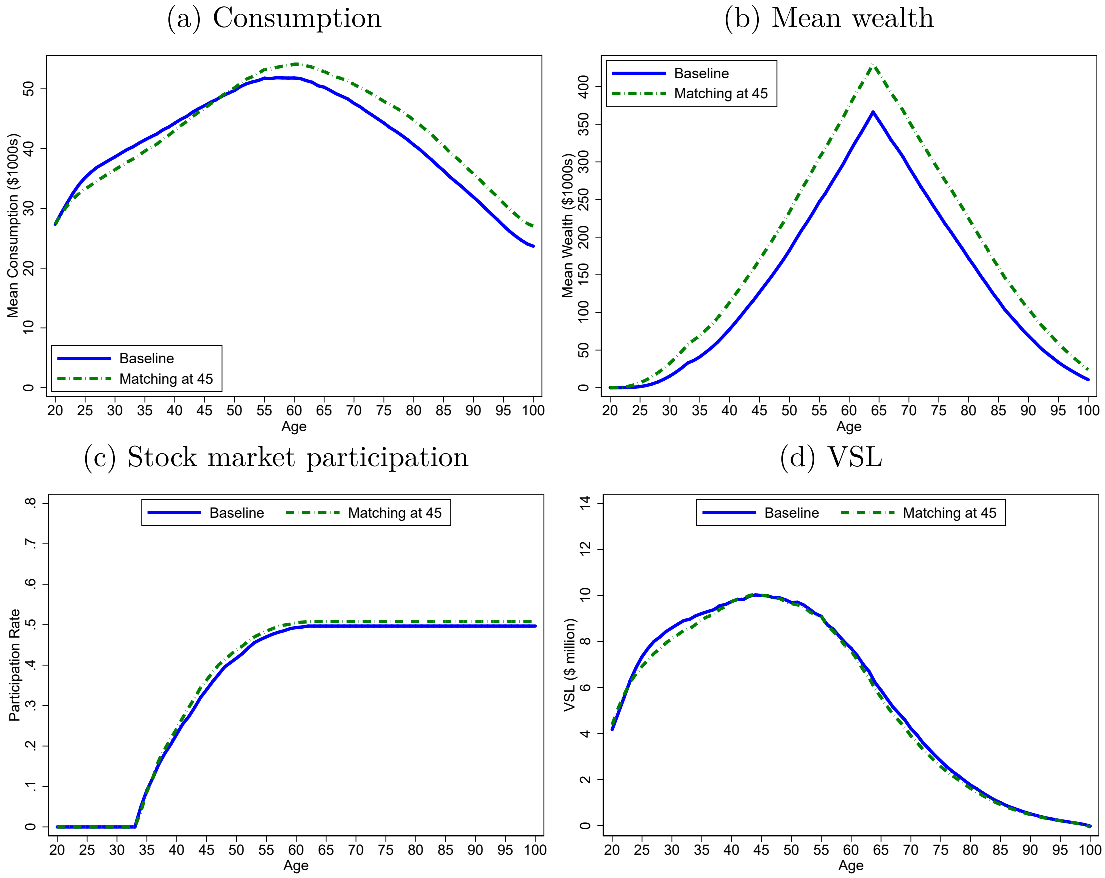

The rest of the agent's program remains unchanged. Note, however, that since the stock market participation is decided in every period, the participation status, denoted by η in the original program, is not needed as a state variable any more. The new values for the calibrated parameters are displayed in Table 6. The main difference compared to the baseline is the participation cost that now falls from a one-off charge of 170% of annual income to 14.2% of average income $\overline{y}$. Our substantive conclusions are unaffected by this choice.

Table 6: Parameter calibration with age 65 for the wealth target

Parameter	Risk-Sensitive	Additive Model
Preferences		
Risk aversion parameter, k	0.95	0.00†
Life-death utility gap, u_l	3.50	12.273
Discount factor, β	0.967	0.953
Bequest motive strength, θ	56.55†	56.55†
Bequest luxury good, $\overline{x}$	8.50†	8.50†
Annuity administrative load, δ	10%†	23%
Asset Markets		
Annual Participation cost, F	14.2% of $\bar{y}$	14.4% of $\bar{y}$

References

Alan, S. (2006): "Entry Costs and Stock Market Participation over the Life Cycle," *Review of Economic Dynamics*, 9, 588–611.

Ameriks, J., A. Caplin, S. Laufer, and S. Van Nieuwerburgh (2011): "The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Bequest and Precautionary Motives," *Journal of Finance*, 66, 519–561.

Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008): "Eliciting Risk and Time Preferences," *Econometrica*, 76, 583–618.

Biggs, A. J. and G. R. Springstead (2008): "Alternate Measures of Replacement Rates for Social Security Benefits and Retirement Income," *Social Security Bulletin.*

Bommier, A., D. Harenberg, and F. LeGrand (2018): "Recursive Preferences and the Value of Life: A Clarification," Working Paper, ETH Zurich.

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

Bommier, A. and F. LeGrand (2014a): "A Robust Approach to Risk Aversion," Working Paper, ETH Zurich.

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

——— (2019): "Risk Aversion and Precautionary Savings in Dynamic Settings," *Management Science*, 65, 1386–1397.

Brown, J. R. (2007): "Rational and Behavioral Perspectives on the Role of Annuities in Retirement Planning," NBER Working Paper 13537, National Bureau of Economic Research.

Campbell, J. Y. and L. M. Viceira (2002): *Strategic Asset Allocation*, Oxford: Oxford University Press.

Chew, S. H. and L. G. Epstein (1989): "The Structure of Preferences and Attitudess towards the Timing of the Resolution of Uncertainty," *International Economic Review*, 30, 103–117.

Cocco, J. F., F. J. Gomes, and P. J. Maenhout (2005): "Consumption and Portfolio Choice over the Life Cycle," *Review of Financial Studies*, 18, 491–533.

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

De Nardi, M. (2004): "Wealth Inequality and Intergenerational Links," *Review of Economic Studies*, 71, 743–768.

De Nardi, M., E. French, and J. B. Jones (2010): "Why do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 118, 39–75.

Epstein, L. G., E. Farhi, and T. Strzalecki (2014): "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review*, 104, 2680–2697.

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.

Fagereng, A., C. Gottlieb, and L. Guiso (2017): "Asset Market Participation and Portfolio Choice over the Life-Cycle," *Journal of Finance*, 72, 705–750.

Gârleanu, N. and S. Panageas (2015): "Young, Old, Conservative, and Bold: The Implications of Heterogeneity and Finite Lives for Asset Pricing," *Journal of Political Economy*, 123, 670–685.

Gomes, F. and A. Michaelides (2005): "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence," *Journal of Finance*, 60, 869–904.

Gomes, F., A. Michaelides, and V. Polkovnichenko (2009): "Optimal Savings with Taxable and Tax-Deferred Accounts," *Review of Economic Dynamics*, 12, 718–735.

Gomes, F. J. and A. Michaelides (2008): "Asset Pricing with Limited Risk Sharing and Heterogeneous Agents," *Review of Financial Studies*, 21, 415–448.

Guvenen, F. (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12, 58–79.

Hall, R. E. and C. I. Jones (2007): "The Value of Life and the Rise in Health Spending," *Quarterly Journal of Economics*, 122, 39–72.

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

Harenberg, D. and A. Ludwig (2019): "Idiosyncratic Risk, Aggregate Risk, and the Welfare Effects of Social Security," *International Economic Review*, 60, 661–692.

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

Huggett, M., G. Ventura, and A. Yaron (2011): "Sources of Lifetime Inequality," *American Economic Review*, 101, 2923–2954.

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

Hurd, M. D. and J. P. Smith (2002): "Expected Bequests and their Distribution," NBER Working Paper 9142, National Bureau of Economic Research.

Inkmann, J., P. Lopes, and A. Michaelides (2011): "How Deep Is the Annuity Market Participation Puzzle?" *Review of Financial Studies*, 24, 279–319.

Kihlstrom, R. E. and L. J. Mirman (1974): "Risk Aversion with many Commodities," *Journal of Economic Theory*, 8, 361–388.

Kniesner, T. J. and W. K. Viscusi (2019): "The Value of a Statistical Life," *Forthcoming, Oxford Research Encyclopedia of Economics and Finance*, 19–15.

Kocherlakota, N. R. (1996): "The Equity Premium: It's Still a Puzzle," *Journal of Economic Literature*, 34, 42–71.

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

Lockwood, L. M. (2012): "Bequest Motives and the Annuity Puzzle," *Review of Economic Dynamics*, 15, 226–243.

——— (2018): "Incidental Bequests and the Choice to Self-Insure Late-Life Risks," *American Economic Review*, 108, 2513–2550.

Lusardi, A., P.-C. Michaud, and O. S. Mitchell (2017): "Optimal Financial Knowledge and Wealth Inequality," *Journal of Political Economy*, 125, 431–477.

Marshall, J. M. (1984): "Gambles and the Shadow Price of Death," *The American Economic Review*, 74, 73–86.

Mehra, R. and E. C. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of monetary Economics*, 15, 145–161.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown (1999): "New Evidence on the Money's Worth of Individual Annuities," *American Economic Review*, 89, 1299–1318.

Pashchenko, S. (2013): "Accounting for Non-Annuitization," *Journal of Public Economics*, 98, 53–67.

Rawlins, M. D. and A. J. Culyer (2004): "National Institute for Clinical Excellence and its Value Judgments," *BMJ*, 329, 224–227.

Rosen, S. (1988): "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty*, 1, 285–304.

Samuelson, P. A. (1937): "A Note on Measurement of Utility," *Review of Economic Studies*, 4, 155–161.

Shepard, D. S. and R. Zeckhauser (1984): "Survival versus Consumption," *Management Science*, 30, 423–439.

Skiadas, C. (2013): "Scale-Invariant Uncertainty-Averse Preferences and Source-Dependent Constant Relative Risk Aversion," *Theoretical Economics*, 8, 59–93.

——— (2015): "Dynamic Choice with Constant Source-Dependent Relative Risk Aversion," *Economic Theory*, 60, 393–422.

Social Security Administration (2020): "Measures of Central Tendency For Wage Data," `https://www.ssa.gov/OACT/COLA/central.html`, Accessed: April 3 2020.

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181.

US Department of Transportation (2016): *Revised Departmental Guidance on Valuation of a Statistical Life in Economic Analysis.*

US Environmental Protection Agency Office of Air and Radiation (2011): "The Benefits and Costs of the Clean Air Act from 1990 to 2020," Second prospective report, Environmental Protection Agency.

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

Weil, P. (1990): "Non-Expected Utility in Macroeconomics," *Quarterly Journal of Economics*, 105, 29–42.

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

Yogo, M. (2016): "Portfolio Choice in Retirement: Health Risk and the Demand for Annuities, Housing, and Risky Assets," *Journal of Monetary Economics*, 80, 17–34.

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.