Household Finance and the Value of Life*

Antoine Bommier Daniel Harenberg François Le Grand†

February 1, 2017

Abstract

We analyze life-cycle saving strategies with a recursive model that is designed to provide reasonable positive values for the value of a statistical life. With a positive value of life, risk aversion amplifies the impact of uncertain survival on the discount rate, and thus reduces savings. Our model also predicts that risk aversion lowers stock market participation and leads to choose more conservative portfolios.

Keywords: recursive utility, life-cycle model, risk aversion, saving choices, portfolio choices, value of life.

JEL codes: D91, G11, J17.

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 life-cycle models, they belong to and develop in quite separated spheres of the economic literature with apparently little exchange between them.

On the one hand, the household finance literature tackles questions related to optimal lifecycle consumption-saving and financial portfolio choices. These questions are of first-order importance

*We would like to thank Francisco Gomes, Michael Haliassos, Felix Kubler, Alexander Michaelides, and seminar participants at the University of Zurich, the Paris School of Economics, Luxembourg School of Finance, the QSPS workshop in Utah, the CEPR household finance workshop, and the Netspar International Pension Workshop for their very helpful comments. A previous version of this paper was circulated under the title "A Bird in the Hand is Worth Two in the Grave: Risk Aversion and Life-Cycle Savings".

†Bommier: ETH Zurich, abommier@ethz.ch; Harenberg: ETH Zurich, dharenberg@ethz.ch; Le Grand: EMLyon Business School and ETH Zurich, legrand@em-lyon.com. The authors gratefully acknowledge support from the Swiss-Re Foundation and the ETH Zurich Foundation.

for analyzing and designing efficient pension systems, and more generally for thinking about saving incentives. To answer such questions, many sophisticated techniques have been adopted, such as recursive models representing preferences à la Kreps and Porteus (1978). These preferences enable to disentangle risk aversion from intertemporal elasticity of substitution (IES), which has been shown to help explaining lifecycle consumption and risky saving patterns. From a modeling standpoint, most of this literature assumes that individuals die with exogenous mortality rates. Mortality being exogenous, there is no explicit trade-off between consumption and life duration. The notion of the value of life, that precisely measures the individual willingness to give up consumption for living longer, is generally not considered.

On the other hand, the literature on the value of life is centered around questions related to mortality risk reduction. Public policies aiming at lowering mortality risk—such as road safety investments or public health campaigns—are typically quite expensive and evaluated through cost-benefit analyses. This requires comparing the monetary amounts for financing the new policy to human lives that have been saved or improved (less severe physical injury for instance). Therefore, having an estimate of the statistical value of life is key, and a large share of the literature has focused on obtaining robust empirical evidence. However, little use has been made of the theoretical advances that were derived in household finance, and especially of the recursive models that allow for disentangling risk aversion from the IES.

In this paper, we investigate the benefits of importing knowledge from the value of life literature into household finance. In particular, we show that a proper calibration of the value of life brings new insights on the determinants of savings and portfolio choice. We motivate our work by a basic two-period model, which is sufficient to show that the value of life may significantly affect time discounting and agents' decisions, even though mortality is exogenous. The main contribution of the paper consists of the development and analysis of a large-scale, quantitative lifecycle model that is standard in most aspects but adds an additional parameter offering the flexibility to match empirical estimates of the value-of life. It thereby nests the models in the previous literature as a limit case.

When the model is calibrated to match the empirical estimates of the value of life, the conclusions differ from those derived in the literature with models that implicitly assumed a negative value of life. In particular, in contrast to the seminal contributions of Gomes and Michaelides (2005, 2008), our model predicts that more risk averse households tend to save less and to participate less in the stock market. The reason is that, with non-additive preferences and a positive value of life, risk aversion amplifies the role of survival risk in the discount rate. The resulting higher discount rate reduces the agent's propensity to save. Quantitatively, this effect turns out to be much stronger than the wealth accumulation induced by a higher prudence. Our result that

more risk averse households participate less in the stock market and opt for more conservative portfolios re-aligns the model prediction with the common understanding of risk-taking behavior as well as with empirical evidence, see, e.g., Dohmen, et al. (2011).

Regarding the relationship between risk aversion and savings, we provide a small empirical study suggesting that more risk averse agents do indeed save less, as predicted by our model. Specifically, we use the German Socio-economic panel, because it has a measure of general risk aversion that has been shown to be a good predictor for a variety of risky behavior (Dohmen, et al. (2011)). In our estimation strategy, we closely follow Fuchs-Schündeln and Schündeln (2005) and find a highly significant negative relationship between risk aversion and savings. Moreover, by estimating interaction terms, we find that it is the interplay of risk aversion and mortality risk that decreases savings, while conjunction of risk aversion and income risk has an opposing effect. This affords further support to our model.

The rest of the paper is organized as follows. In Section 2, we explain the concept of the value of a statistical life and, with the help of a basic two-period model, demonstrate why it is important for household finance. We then present our quantitative lifecycle model in Section 3. We specify utility functions in Section 4 and describe the calibration in Section 5. Section 6 presents the results and discusses how they relate to the previous literature. Section 7 provides a short discussion of our supporting empirical evidence. We conclude in Section 8.

2 The value of a statistical life

2.1 A marginal rate of substitution between survival probability and consumption

We present the definition for the value of a statistical life (VSL), as well as the rationale behind it. As a starting point for this definition we need the concept of lifetime utility. In standard lifecycle models with uncertain life duration, the lifetime utility of an agent depends on consumption levels at different ages and survival probabilities, as well as possibly on other factors, such as health, or leisure for instance. Denoting by c_t the consumption at age t, by p_t the survival probability between ages t and $t+1$, and x_t the vector gathering other aspects, such as health and leisure, the lifetime utility of an agent can be expressed as $U(c_0, \ldots, c_t, \ldots, p_0, \ldots, p_t, \ldots, x_0, \ldots, x_t, \ldots)$. The functional form of the function U is generally further specified, for instance through expected utility models. However, the concept of VSL is probably easier to understand if we abstract from specific functional forms. The VSL can be defined as a marginal rate of substitution between

survival probability and consumption.[1] Formally, the VSL at date t, denoted VSL_t, is defined as

$$VSL_t = \frac{\frac{\partial U}{\partial p_t}}{\frac{\partial U}{\partial c_t}} \tag{1}$$

It is worth explaining why the marginal rate of substitution between survival probability and consumption is called "Value of a Statistical Life". By definition of the marginal rate of substitution, an agent would be willing to exchange εVSL_t units of period-t consumption for an increase of ε in her survival probability (where $0 < \varepsilon \ll 1$ is infinitesimally small). Considering a population of $\frac{1}{\varepsilon}$ identical agents, they would in aggregate be willing to pay $\frac{1}{\varepsilon} \times \varepsilon VSL_t = VSL_t$ to increase the expected number of survivors by $\frac{1}{\varepsilon} \times \varepsilon = 1$. In other words, VSL_t represents the aggregate willingness-to-pay to save (on average) one life. This explains why the terminology of the value of a statistical life has been coined.

2.2 Empirical literature on the VSL

As mentioned in the introduction, the VSL is a key parameter for cost-benefit analysis in policy designs. For example, about 85% of the benefits of the Clean Air Act are related to mortality risk reduction, as computed by the US Environmental Protection Agency (2011). In other words, using VSL estimates that would be off by a factor of 2 would also lead to under- or overestimate the benefits of the Clean Air Act by a factor of 1.8. This shows how much good estimates of the VSL are needed. This need has long been acknowledged and many researchers (e.g., see the review in Viscusi, 2003) and institutes (e.g., EPA and FDA in the USA) put significant effort into obtaining estimates from observed behavior. Essentially, one wants to find empirical evidence on the willingness to pay for mortality risk reduction. One approach is too look at wage-risk trade-offs. Another is to look at the willingness to pay to get safer cars, safer homes, etc. Direct questionnaires can also be informative.

Although empirical estimation proves difficult, there is a wide consensus that for most people the value of life is positive and large. There is of course some heterogeneity across studies and social context, but for a country like the USA, a VSL of 6 to 7 million dollars is considered a reasonable estimate.

2.3 Why should we care for the VSL even if mortality is exogenous?

In this section we want to emphasize that even if mortality is exogenous, assumptions made regarding the value of life have major consequences on savings behavior. In order to make this point we focus on a simple setting where people live at most two periods (0 and 1), and only care for

[1] This definition is standard, see, e.g., Johansson (2002).

consumption and survival.[2] Their utility is therefore a function $U(c_0, c_1, p_0)$. A popular solution to further specify this utility function is to consider a recursive approach. This approach requires to first compute the agent utility in period 1 (depending on her consumption and whether she is alive or not), and then to derive the period-0 utility using a recursive expression à la Kreps and Porteus (1978). Period-1 utility, which depends on consumption c_1 and on survival status $\xi_1 \in \{\text{alive, dead}\}$, is denoted by $U_1(c_1, \xi_1)$. The period-0 utility can then be expressed as

$$U(c_0, c_1, p_0) = u(c_0) + \beta\phi^{-1}E_{p_0}[\phi(U_1)],$$

where $\beta \in (0, 1)$ is a discounting factor and $\phi : Im(U) \to \mathbb{R}$ is an increasing concave function governing risk aversion. We assume that age effects are discarded and the periodic utilities when being alive at dates 0 and 1 are assumed to be the same, so that $U_1(c_1, \text{alive}) = u(c_1)$. The agent has no bequest motive and does not care about her consumption after death. Consequently, $U_1(c_1, \text{dead})$ is independent of c_1 and we set $U_1(c_1, \text{dead}) = u_d \in [-\infty, +\infty]$, since infinite evaluations are theoretically possible. The value of u_d governs the utility gap between life and death. Lifetime utility is given by

$$U(c_0, c_1, p_0) = u(c_0) + \beta\phi^{-1}\left(p_0\phi(u(c_1)) + (1 - p_0)\phi(u_d)\right). \tag{2}$$

The value of life. Applying definition (1) we get

$$VSL_0 = \frac{\beta\left(\phi(u(c_1)) - \phi(u_d)\right)}{u'(c_0)\phi'\left(\phi^{-1}\left(p_0\phi(u(c_1)) + (1 - p_0)\phi(u_d)\right)\right)}. \tag{3}$$

The VSL is therefore connected to the value of u_d, being positive if $u_d < u(c_1)$ and negative if $u_d > u(c_1)$. It is standard in the household finance literature to pick values of u_d that provide tractable specifications without paying attention to the implied value of life. A commonly used specification is the Epstein-Zin utility function, given by $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$, $\phi(u) = \frac{((1-\sigma)u)^{\frac{1-\gamma}{1-\sigma}}}{1-\gamma}$, where $\frac{1}{\sigma}$ is the IES and γ controls risk aversion. Typically both σ and γ are assumed to be larger than 1. Maximal tractability is then obtained by setting $u_d = 0$ so that $\phi(u_d) = 0$. Indeed, in such a case equation (2) provides

$$\begin{aligned} U(c_0, c_1, p_0) &= \frac{c_0^{1-\sigma}}{1-\sigma} + \frac{\beta}{1-\sigma}\left(p_0 c_1^{1-\gamma}\right)^{\frac{1-\sigma}{1-\gamma}} \\ &= \frac{c_0^{1-\sigma}}{1-\sigma} + \frac{\beta}{1-\sigma} p_0^{\frac{1-\sigma}{1-\gamma}} c_1^{1-\sigma}, \end{aligned}$$

which is homothetic, a very helpful property for deriving solutions of utility maximization problems. This convenient specification, however, implicitly assumes a negative value of life, because γ and

[2] Most of the simplifying assumptions in this subsection will be relaxed in the large-scale quantitative model of Section 3.1.

σ are assumed to be greater than one.[3] It is noteworthy that a different utility representation is sometimes considered: $V(c_0, c_1, p_0) = \left(c_0^{1-\sigma} + \beta p_0^{\frac{1-\sigma}{1-\gamma}} c_1^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$. However this is fully equivalent since the function V and U are related through an increasing transformation: $V(c_0, c_1, p_0) = ((1-\sigma)U(c_0, c_1, p_0))^{\frac{1}{1-\sigma}}$. While U is negative, V is always positive. However, the utility function V, despite being positive, also implies a negative value of life, just like U.

The value of life and time discounting. Even with exogenous mortality, using a good estimate for the VSL is important. Let us for example consider the implications on the consumption discount rate between periods 0 and 1. This discount rate can be defined as

$$\delta_0 = \left.\frac{\frac{\partial U}{\partial c_0}}{\frac{\partial U}{\partial c_1}}\right|_{c_0=c_1} - 1, \tag{4}$$

which is the rate of change of marginal utility for a constant consumption stream. With preferences represented by (2), we obtain the following expression:

$$\delta_0 = \frac{1}{\beta p_0} \frac{\phi'(\tilde{u})}{\phi'(u(c_1))} - 1, \tag{5}$$

where

$$\tilde{u} = \phi^{-1}\left(p_0 \phi(u(c_1)) + (1-p_0)\phi(u_d)\right) \tag{6}$$

is the certainty equivalent (in utility levels) of a lottery that gives utility $u(c_1)$ with probability p_0 and u_d with probability $1-p_0$. In the case where ϕ is affine, ϕ' is constant, and equation (5) reduces to $\delta_0 = \frac{1}{\beta p_0} - 1$. Mortality contributes to impatience in a very simple way, and in particular independently of the constant u_d. This however does not extend to non-additive models with strictly concave function ϕ. To see the impact of ϕ, rewrite equation (5) as

$$\delta_0 = \frac{1}{\beta p_0} \exp\left(\int_{\tilde{u}}^{u(c_1)} \lambda_\phi(u) du\right) - 1, \tag{7}$$

where $\lambda_\phi(u) = -\frac{\phi''(u)}{\phi'(u)}$ is a measure of the concavity of ϕ and thus a measure of risk aversion. When ϕ is concave, the factor $\frac{1}{\beta p_0}$ is multiplied by the quantity $\exp\left(\int_{\tilde{u}}^{u(c_1)} \lambda_\phi(u) du\right)$, which intertwines risk aversion, mortality, and impatience. This exponential term depends on $\tilde{u}$, defined in equation (6), and therefore on how u_d compares to $u(c_1)$.

Consider first the case where $u_d < u(c_1)$, which means a positive value of life. Then, the greater the concavity of ϕ and the higher the mortality rate $1-p_0$, the larger the expression $\exp\left(\int_{\tilde{u}}^{u(c_1)} \lambda_\phi(u) du\right)$ and the discount rate δ_0. In other words, if the value of life is positive, risk aversion amplifies the impact of mortality on the discount rate: in the presence of mortality

[3] Indeed, $\frac{\partial U(c_0, c_1, p_0)}{\partial p_0} = \frac{\beta}{1-\gamma} p_0^{\frac{\gamma-\sigma}{1-\gamma}} c_1^{1-\sigma} < 0$, no matter p_0 and c_1.

risk, more risk averse agents are more impatient. It is worth contrasting this case with the one where $u_d > u(c_1)$, implying a negative value of life. We now have $\tilde{u} > u(c_1)$, so that the factor $\exp\left(\int_{\tilde{u}}^{u(c_1)} \lambda_\phi(u)du\right)$ is smaller than one, thus reducing—instead of amplifying—the impact of mortality on the discount rate .

Overall, for given intertemporal elasticity of substitution and risk aversion (i.e., given functions u and ϕ, respectively), a model with a negative value of life will predict a lower rate of discount than the same model with a positive value of life. As a lower discount rate means larger savings, we find that underestimating the value of life (and in particular using a negative value) is likely to exaggerate the propensity to save.

3 A quantitative life-cycle model

Having explained why it is important to carefully calibrate the value of life, we present in this section a life-cycle model that is mostly standard, but in which the utility level of the death state (u_d in the previous Section) will be fixed to fit empirical estimates on the VSL.

3.1 The setup

We consider a partial equilibrium economy populated by an agent endowed with recursive preferences and facing several risks: a mortality risk, an income risk and an investment risk through risky financial returns. The agent may save through bonds and a risky asset (similar to a stock). Time is discrete, and the agent's age is denoted by t. The agent enters the model at working age, $t = 0$. There is a single consumption good, whose price serves as a numeraire.

Mortality risk. The agent faces mortality risk, which is assumed to be independent of any other risk in the economy. If alive at date t, the agent survives to date $t+1$ with probability p_t. There exists a date T_M, such that the probability to live after T_M is $p_{T_M} = 0$.

Labor income risk. At any age, when alive, the agent receives an income denoted y_t. The agent exogenously retires at age T_R and during retirement ($t \geq T_R$), the agent receives a constant pension income $y_t = y^R$. During working age $t < T_R$, the agent earns a risky labor income $y_t = y_t^L$, which is subject to both persistent shocks, π_t, and transitory shocks, ϑ_t:

$$y_t^L = \bar{y} \exp\left(\mu_t + \pi_t + \vartheta_t\right), \tag{8}$$

$$\pi_t = \rho \pi_{t-1} + \upsilon_t. \tag{9}$$

The two independent processes $(\upsilon_t)_{t\geq 0}$ and $(\vartheta_t)_{\geq 0}$ are IID and normally distributed with mean 0 and respective variance σ_υ^2 and σ_ϑ^2. The quantity $\bar{y}$ in (8) represents an average, constant wage rate, while $(\mu_t)_{t\geq 0}$ is a deterministic process that contributes to fit the wage process to the data and in particular the humped-shape pattern of income during active age. The parameter ρ in (9) drives the persistence of the process π.

Financial risk and security markets. The agent has the opportunity to save through a riskless one period asset (similar to a T-Bill) and a risky asset (similar to a stock). The bond is a security of price 1 which pays R^f as a riskless gross return in the subsequent period. The rate of interest R^f is constant and exogenous.The risky return is:

$$\ln R_t^s = \ln(R^f + \omega) + \nu_t,$$

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 financial risk is assumed to be possibly correlated to both income shocks, (υ_t) and (ϑ_t). The correlation with each income process is assumed to be constant and is denoted respectively by κ_υ and κ_ϑ.

The agent must pay a fixed cost $F \geq 0$ to participate to the stock market, which may be interpreted as an opportunity cost to discover how the stock market works. We assume it is a 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.

Timing and notation. At the beginning of every period, the agent first learns the realizations of financial and labor income shocks and whether she is alive or not. She thus knows the amount of her current savings and, if she is alive, her current income. More precisely, at any date t, we assume that the agent knows the entire history of all shocks up to date t, which is formalized by the natural filtration $(\mathcal{F}_t)$ generated by the processes (υ_t), (ϑ_t), and (ν_t). The alive agent then decides her consumption level c_t, her savings in bonds b_t and stocks s_t, and her stock market participation status η_t (equal to 0, if she has never paid the participation cost before and therefore never participated).

Constraints. If the agent is dead at date t, she bequeaths all her wealth w_t:

$$w_t = R^f b_{t-1} + R_t^s s_{t-1}. \tag{10}$$

The stock holding s_{t-1} may be null if the agent has never participated to the stock market.

If the agent is alive, her resources at the beginning of the period consist of stock and bond payoffs plus the labor income y_t of the period. Resources cover consumption as well as the purchase of

bonds and stocks. The agent can only invest in stocks if the participation cost, F, has been paid at some date prior to t, i.e., if $\eta_t = 1$. Moreover, the agent may also have to pay the participation cost F at date t if she participates at date t in the stock market for the first time, i.e., if $\eta_t = 1$ and $\eta_{t-1} = 0$. The budget constraint at date t of the alive agent can then be expressed as follows:

$$c_t + b_t + s_t 1_{\eta_{t=1}} + F 1_{\eta_{t=1}} 1_{\eta_{t-1=0}} = y_t + R^f b_{t-1} + R_t^s s_{t-1}, \tag{11}$$

where $1_{\eta_s=1}$ is an indicator function equal to 1 if $\eta_s = 1$ and 0 otherwise. Neither asset can be sold short and consumption must be strictly positive:

$$b_t \geq 0 \text{ and } s_t \geq 0, \tag{12}$$

$$c_t > 0. \tag{13}$$

A feasible allocation is a sequence of choices $(c_t, b_t, s_t, \eta_t)_{t\geq 0}$ satisfying the constraints (10)–(13). The set of feasible allocations is denoted $\mathcal{A}$.

Regarding initial conditions, we assume without loss of generality that $\eta_{-1} = 0$ and $b_{-1} \geq 0$ and $s_{-1} \geq 0$ are exogenous values. Equations (10) and (11) can be assumed to hold at date $t = 0$.

3.2 Preferences

We denote by $u(c_t) : \mathbb{R}^+ \to I$ the instantaneous felicity the agent gets when being alive and consuming c_t and by $v(w_t)$ the utility she derives when being dead and bequeathing the amount w_t. Preferences are separable over time and future instantaneous utilities are discounted by a factor $\beta \in (0, 1)$ representing the agent's exogenous time preference.

Regarding risk preferences, we consider recursive utilities *à la* Kreps and Porteus (1978). Agents value the certainty equivalent of future utility streams. More precisely, for an increasing concave function $\Phi : \mathbb{R} \to \mathbb{R}$, the utility U_t at date t expresses as follows:[4]

$$U_t = (1-\beta)u_t + \beta\Phi^{-1}\left(E_t^{\mathcal{F}\times\mathcal{G}}\left[\Phi(U_{t+1})\right]\right), \text{ with } u_t = \begin{cases} u(c_t), & \text{if the agent is alive at } t, \\ v(w_t), & \text{if the agent is dead at } t. \end{cases} \tag{14}$$

In the above equation, $E_t^{\mathcal{F}\times\mathcal{G}}[\cdot]$ is the conditional expectation operator with respect to the information available at date t. Formally, the information is the filtration $(\mathcal{F}_t \otimes \mathcal{G}_t)_{t\geq 0}$, where $(\mathcal{G}_t)_{t\geq 0}$ is the filtration generated by the independent mortality process. The factor $1-\beta$ multiplying u_t in equation (14) is a normalization that simplifies the expression of the utility of a dead agent (see equation (15) below).

[4] Formally speaking, preferences are defined over the set of temporal lotteries, allowing for preferences for late or early uncertainty resolution. See Epstein and Zin (1989) or Wakai (2007) for a formal treatment.

In such models, if there were no uncertainty, the utility U_t would be independent of the function Φ and the recursion (14) would reduce to $U_t = (1-\beta)u_t + \beta U_{t+1}$. We thus have a possible separation between preferences over certain consumption streams (determined by the functions u, v and the scalar β) and risk preferences driven by the function Φ. A more concave Φ implies lower certainty equivalents $\Phi^{-1}\left(E_t^{\mathcal{F}\times\mathcal{G}}\left[\Phi(U_{t+1})\right]\right)$ and therefore greater risk aversion.

Our specification of recursive preferences nests some of the most standard cases, including the additive specification, the Epstein and Zin (1989) isoelastic specification, and the risk-sensitive specification introduced by Hansen and Sargent (1995) in their work on robustness. In Section 4, we make precise the functions Φ which correspond to these different popular specifications. Our results will make it possible to discuss the impact of these specifications on saving decisions.

3.3 Agent's program

We can now write the agent's program recursively by taking advantage of the structure of preferences. We denote by U_t^D the intertemporal utility at date t of a dead agent and by U_t^A that of an alive agent. Regarding the dead agent, note that the instantaneous utility of an agent dead for more than one period is constant (all bequests take place in the first period after death) and equal to $v(0)$. From the recursive formulation (14), we deduce that there is no actual optimization and that the program of a dead agent can be expressed as:

$$U_t^D(w_t) = (1-\beta)v(w_t) + \beta v(0). \tag{15}$$

While alive, the agent maximizes her intertemporal utility by choosing a feasible allocation $(c_t, b_t, s_t, \eta_t)_{t\geq 0}$ in the set $\mathcal{A}$. The utility U^A of the alive agent depends on four state variables: the beginning-of-period holdings in stocks s_{t-1} and bonds b_{t-1}, the permanent shock π_{t-1} of labor income and the stock market participation $\eta_{t-1} \in \{0,1\}$. The latter is discrete, while the three former ones are continuous. From the recursive formulation (14) and feasibility constraints (10)–(13) and using the fact that the mortality risk is assumed to be independent of other risks, the program of an alive agent at date t can be expressed as follows:

$$\begin{aligned} U_t^A(s_{t-1}, b_{t-1}, \eta_{t-1}, \pi_{t-1}) = & \max_{(c_t, s_t, b_t, \eta_t)\in\mathcal{A}} (1-\beta)u(c_t) \\ & + \beta\Phi^{-1}\left(p_t E_t[\Phi(U_{t+1}^A(s_t, b_t, \eta_t, \pi_t))] + (1-p_t)E_t[\Phi(U_{t+1}^D(w_{t+1}))]\right), \end{aligned} \tag{16}$$

where $E_t[\cdot]$ is the expectation for an alive agent with respect to the filtration $\mathcal{F}$ (i.e., made of all past shock realizations but death). Note that we distinguish it from the expectation $E_t^{\mathcal{F}\times\mathcal{G}}[\cdot]$ with respect to the whole information $(\mathcal{F}_t \otimes \mathcal{G}_t)_{t\geq 0}$ (including death information). It should be noted that the program (16) has a finite-horizon since there exists a maximal age for the agent, T_M.

3.4 Value of life

Applying the definition in equation (1), we derive the following expression for the value of a statistical life at age t:

$$VSL_t = \frac{\frac{\partial U_t^A}{\partial p_t}}{\frac{\partial U_t^A}{\partial c_t}} = \frac{\beta}{1-\beta} \frac{E_t\left[\Phi\left(U_{t+1}^A\right) - \Phi\left((1-\beta)v(w_{t+1}) + \beta v(0)\right)\right]}{u'(c_t)\Phi'\left(\frac{1}{\beta}\left(U_t^A - (1-\beta)u(c_t)\right)\right)}. \tag{17}$$

We will use this expression to calibrate the model to empirical estimates of the VSL, cf. Section 5.

4 Specifications of utility functions

We now specify the functional forms for felicity functions u and v and for the aggregator Φ.

4.1 Felicity function specification

We begin with specifying u and v. We assume that the agent has a constant IES, which means that $-\frac{u'(c)}{c\,u''(c)}$ is constant. This implies that u is equal, up to an affine transformation, to:

$$u(c) = \begin{cases} \frac{c^{1-\sigma}}{1-\sigma} - \frac{1}{1-\sigma} & \text{if } \sigma \neq 1, \\ \ln(c) & \text{if } \sigma = 1. \end{cases} \tag{18}$$

where the parameter $\sigma > 0$ is the inverse of the IES. The above specification is such that $u(1) = 0$. It embeds therefore a normalization assumption, which is without generality loss. The case $\sigma = 1$ is obtained by continuity from the general case.

The felicity derived from bequeathing wealth w is assumed to have the following functional form:

$$v(w) = \begin{cases} -\bar{v} + \frac{\theta}{1-\sigma}\left[(\overline{w}+w)^{1-\sigma} - \overline{w}^{1-\sigma}\right] & \text{if } \sigma \neq 1, \\ -\bar{v} + \theta \ln\left(\frac{\overline{w}+w}{\overline{w}}\right) & \text{if } \sigma = 1, \end{cases} \tag{19}$$

where $\overline{v} = -v(0) \in \mathbb{R}$, $\theta \geq 0$, σ is the inverse of the IES used in the expression (18) of the felicity of u, and $\overline{w} \geq 0$ is a parameter that we discuss below. As for u, the case $\sigma = 1$ in (19) is obtained by continuity from the general case.

We can distinguish two components in the specification of v in equation (19). The first one is the term (-$\overline{v}$), which corresponds to the difference in utility between being dead and bequeathing nothing and being alive and consuming one unit. This difference is negative if death is not as good as being alive with one unit of consumption, the constant $\overline{v}$ being then positive. A higher (resp. lower) value of $\overline{v}$ will be associated with a higher (resp. lower) valuation of being alive, compared to being dead. The value of $\overline{v}$ thus strongly connects to the value of life. The second

part, $\frac{1}{1-\sigma}\left[(\overline{w}+w)^{1-\sigma}-\overline{w}^{1-\sigma}\right]$ measures the contribution of bequest to post-mortem felicity. This extra felicity derived from bequest is assumed to be continuous in zero, increasing in the amount of bequest and exhibiting bounded and decreasing marginal felicity. The rationale for this functional expression is the following one. Heirs may already have individual resources at their disposition, summarized by the quantity $\overline{w}$, and they enjoy bequest in addition to these resources $\overline{w}$. The felicity derived by heirs from bequest is proxied by the quantity $\frac{1}{1-\sigma}\left[(\overline{w}+w)^{1-\sigma}-\overline{w}^{1-\sigma}\right]$.[5] The agent values the felicity of her heirs with the weight θ that can therefore be interpreted as the intensity of the altruistic bequest motive. With $\overline{w}>0$, bequests are a luxury good, as reported in the data (e.g., in Hurd and Smith, 2002). Indeed, the derivative $v'(0)$ is finite, so that agents bequeath only when their wealth is large enough. This functional form has been chosen for example in De Nardi (2004), De Nardi et al. (2010), Ameriks et al. (2011), and Lockwood (2012, 2014).

It is sometimes assumed in the literature that $\overline{v}=\frac{1-\theta\overline{w}^{1-\sigma}}{1-\sigma}$ so that $v(w)=\frac{1}{1-\sigma}\left[\theta\left(\overline{w}+w\right)^{1-\sigma}-1\right]$, which has some advantage in terms of tractability.[6] However, this constraint on $\overline{v}$ implies a non-trivial relationship between the utility of bequest and the value of life. In particular, if θ is set to zero (no altruism) and $\sigma>1$, then the utility of being dead is always higher to that of being alive, implying a negative value of life. We will not make assumptions of these kinds as we want our model to match standard empirical estimates for the value of life.

We now discuss discuss the two functional forms we consider for the function Φ.

4.2 Risk-sensitive preferences

First, we consider risk-sensitive preferences:

$$\Phi(u)=\begin{cases}-\frac{1}{k}\left(\exp(-ku)-1\right) & \text{if } k\neq 0,\\ u & \text{if } k=0,\end{cases} \tag{20}$$

where k is a constant driving risk aversion. The case $k=0$ corresponds to the usual additive model and is obtained by continuity of the general case. For an agent endowed with risk-sensitive preferences to be more risk averse than in the usual additive model, we need to assume that $k>0$. Risk-sensitive preferences have been introduced in Hansen and Sargent (1995) and axiomatized in Strzalecki (2011). As shown in Bommier, Kochov and LeGrand (2016), this is the only functional form Φ for which preferences represented by the utility function in recursion (14) are monotone. The monotonicity of preferences has to be understood as monotonicity with respect to first-order stochastic dominance. Such a monotonicity means that if two uncertain consumption streams are

[5] The proxy is exact if (i) heirs have the same IES as the donator and (ii) heirs fully annuitize their wealth.

[6] Additional tractability can then be obtained by setting $\overline{w}=0$, in order to have homogeneous specifications, as in Inkmann, Lopes and Michaelides (2011), for example.

available and the first one is preferred to the other one in any possible state of the world, the former will always be preferred to the latter. This is distinct from and not implied by the fact that more certain consumption is preferred to less, which in our setup is equivalent to an increasing felicity function u. Moreover, as proved in Bommier and LeGrand (2014), risk-sensitive preferences are well-ordered with respect to risk aversion, both "in the large" (i.e., in terms of willingness-to-pay to eliminate all risks) but also "in the small" (i.e. in terms of willingness-to-pay for marginal risk reductions). This last aspect is important when addressing problems where complete risk elimination is not possible, or simply not optimal, as is the case in the portfolio choice that we study.

4.3 Epstein-Zin preferences

Second, we consider Epstein-Zin isoelastic preferences, which correspond to the following functional form for Φ:

$$\Phi(u) = \begin{cases} \frac{1}{1-\gamma}(1+(1-\sigma)u)^{\frac{1-\gamma}{1-\sigma}} - \frac{1}{1-\gamma}, & \text{if } \gamma \neq 1 \text{ and } \sigma \neq 1, \\ \frac{1}{1-\sigma}\ln(1+(1-\sigma)u), & \text{if } \gamma = 1 \text{ and } \sigma \neq 1, \\ \frac{1}{1-\gamma}e^{(1-\gamma)u} - \frac{1}{1-\gamma}, & \text{if } \gamma \neq 1 \text{ and } \sigma = 1, \\ u, & \text{if } \gamma = 1 \text{ and } \sigma = 1, \end{cases} \tag{21}$$

where $\gamma \in \mathbb{R}$ and $1+(1-\sigma)u \geq 0$.[7] Whenever $\gamma = \sigma$ (but possibly different from 1), we get $\Phi(u) = u$ and Epstein-Zin preferences are additive. It is also well-known from Tallarini (2000), and directly visible from the last two lines of (21), that when $\sigma = 1$ Epstein-Zin preferences coincide with risk-sensitive preferences. Thus, the cases where $\sigma = 1$ are already addressed with risk-sensitive preferences and do not need further consideration. We will therefore exclude them whenever we refer to Epstein-Zin preferences below. For $\sigma \neq 1$, the constraint $1+(1-\sigma)u \geq 0$ is not trivial. It holds whenever the agent is alive, since we have $1+(1-\sigma)u(c) = c^{1-\sigma}$, but imposes constraints on the felicity of bequest defined in equation (19). The constraint $1+(1-\sigma)u \geq 0$ is equivalent to

$$\begin{cases} \overline{v} \leq \frac{1}{1-\sigma} & \text{if } \sigma < 1, \\ \overline{v} \geq \frac{1-\theta\overline{w}^{1-\sigma}}{1-\sigma} & \text{if } \sigma > 1. \end{cases}$$

[7] Epstein Zin preferences are often introduced with a different but equivalent normalization for the function u (e.g. using $u(c) = \frac{c^{1-\sigma}}{1-\sigma}$ instead $u(c) = \frac{c^{1-\sigma}}{1-\sigma} - \frac{1}{1-\sigma}$), and therefore different functions Φ (the constant 1 being no longer needed). Our (equivalent) approach has however the advantage to have the cases $\sigma = 1$ or $\gamma = 1$ directly obtained as limit cases of the others, while keeping $\overline{v}$ or $\overline{w}$ independent of σ and γ. This normalization choice has no impact on our results.

Isoelastic Epstein-Zin preferences are very popular in macroeconomics and finance, one of their main advantage being that they usually provide a homogeneous specification, which is key for tractability. Note that this is not the case in our setup, where Epstein-Zin preferences are not homothetic. The reason is not the normalization that we made in equation (21), but stems from our choice of imposing a plausible value of statistical life.

An inconvenient aspect of Epstein-Zin preferences is that they are not monotone.[8] This may yield unintuitive conclusions in some cases, even though this is not the case in the current paper.

5 Calibration and computation

In this section, we first give an overview of our calibration strategy. We then discuss the resulting parametrization in detail. The last section discusses the main aspects of solving the model computationally.

5.1 Calibration strategy

Our calibration shares many common aspects with the related literature but differs mainly in that we target the value of a statistical life explicitly. Given a realistic value of life, the objective of the calibration exercise is to highlight the impact of risk aversion. To this aim, we consider three agents: one with standard additively separable preferences, one with Epstein-Zin preferences corresponding to the aggregator (21), and one with risk-sensitive preferences corresponding to the aggregator (20).[9] We will henceforth refer to the three agents as the additive, the Epstein-Zin and the risk-sensitive agent, respectively. Importantly, we calibrate only the additive agent to the data. The other two are built using the parameters of the additive specification but assuming a higher degree of risk aversion ($k > 0$ for risk-sensitive preferences and $\gamma > \sigma$ for Epstein-Zin preferences).

We now describe our strategy for calibrating the additive agent; the resulting parameter values are discussed in the following sections. We set the intertemporal elasticity of substitution, $\frac{1}{\sigma}$, to a standard value. We then jointly calibrate the discount factor, β, the bequest motive, θ, and the life-death utility gap, $\overline{v}$, to match the following three targets. The first target is an estimate of the value of a statistical life at age 45, VSL_{45}, as defined in equation (17). Targeting this is central to our exercise. The second and third targets are average assets at age 45 and average bequests at age 90, respectively. All other parameters are set to values that are taken directly from available data or related studies. Table 1 summarizes the parameter values, displaying endogenously calibrated

[8] Cf. Bommier, Kochov, LeGrand (2016).

[9] Recall from Section 4 that the additively separable case is nested in Epstein-Zin preferences when $\gamma = \sigma$, and in risk-sensitive preferences when $k = 0$.

values in italics.

After having calibrated the additive agent, we set the risk aversion of the Epstein-Zin agent to a slightly higher value $\gamma^{EZ} > \sigma$, while keeping all other parameters the same. This isolates how risk aversion impacts lifecycle savings and portfolio choices. Similarly, we increase risk aversion k of the risk-sensitive agent by setting $k > 0$. More precisely, we calibrate k to produce the same average savings at age 45 as the Epstein-Zin agent, i.e., such that $E_0[s_{45}^{RS} + b_{45}^{RS}] = E_0[s_{45}^{EZ} + b_{45}^{EZ}]$. We do this, because we want the increases in risk-aversion to be of similar magnitude, whether it is achieved with risk-sensitive or Epstein-Zin preferences. Both the Epstein-Zin and risk-sensitive agents are more risk averse than the additive agent but are not comparable with each other in terms of risk aversion.

5.2 Demographics

A model period corresponds to one year. Agents start being economically active in the model at the working age of 20. They exogenously retire at the fixed age of 65, which corresponds to the statutory retirement age in the U.S. Mortality rates are taken from the Human Mortality Database for the USA for 2007. The maximum biological age is capped at 100, since mortality estimates become inaccurate after that.

5.3 Preferences

The IES is set to 0.5, a common value in the literature, so that its inverse is $\sigma = 2$. For the Epstein-Zin agent we increase the risk aversion parameter moderately to $\gamma^{EZ} = 3$, since we do want not deviate too much from the additive agent. Last, for the risk-sensitive agent, we calibrate the risk aversion parameter to match the same average savings as the Epstein-Zin agent at the age of 45, which yields $k = 0.09$.

We then calibrate $\overline{v}$, β, and θ jointly so that the additive agent has a value of a statistical life at age 45, assets at age 45, and bequest at age 90, that match their empirical counterparts. For VSL we target US$ 6.5 million, which is in the middle of available estimates.[10] For average individual assets we target US$ 100 000, which is consistent with Census data, while for bequest we match US$ 50 000 at age 90. This yields $\overline{v} = 1.48 \times 10^{-3}$, $\beta = 0.93$, and $\theta = 5.3$ which we keep constant for the three agents.[11]

[10] Viscusi and Aldy (2003) provide a discussion of available estimates of the value of life.

[11] Note that the value of $\overline{v}$ looks small because it corresponds to an income $\bar{y}$ approximately equal to 22000 USD. Indeed, $\bar{v}$ would amount to 32.3, if the income $\bar{y}$ would have been normalized to 1.

Table 1: Parameterization in baseline economy

Parameter	Value	Source, empirical counterpart, or target
Demographics		
Biological age at $t = 0$	20	age at labor force entry (college)
Model age at retirement, T_R	45	S.S.A. statutory retirement age of 65
Model age maximum, T_M	80	Biological maximum age of 100
Cond. survival rates, $\{p_t\}$		Human Mortality Database, U.S. 2007
Preferences		
Inverse IES, σ	2.0	
Risk aversion, Epstein-Zin, γ^{EZ}	3.0	
Risk aversion, risk-sensitive, k	*0.09*	Assets of EZ agent at age 45
Life-death utility gap, $\overline{v}$	1.48×10^{-3}	VSL_{45} = US$ 6.5m (add. pref.)
Discount factor, β	*0.93*	$Assets_{45}$ = US$ 100000 (add. pref.)
Strength of bequest motive, θ	*5.3*	$Bequests_{90}$ = US$ 50000 (add. pref.)
Exogenous offspring endowment, $\overline{w}$	$\overline{y}$	
Endowments		
Average wage, $\bar{y}$	21 756 USD	Net compensation 2007, SSA
Pension, y^R	$33\% \times \bar{y}$	Replacement rate
Age productivity, $\{\mu_t\}$	cf. app.	Earnings profiles 2007, PSID
Labor income autocorrelation, ρ	0.95	Storesletten, et al. (2004)
Var. of persistent innovation, σ_υ^2	0.30	Storesletten, et al. (2004)
— Correlation with stock return, κ_υ	0.15	Gomes and Michaelides (2005)
Var. of transitory innovation, σ_ϑ^2	0.12	Storesletten, et al. (2004)
— Correlation with stock return, κ_ϑ	0.30	Gomes and Michaelides (2005)
Asset Markets		
Gross risk-free return, R^f	1.01	Bond return (Shiller)
Equity premium, ω	0.03	
Stock volatility, σ_ν	0.18	Shiller data
Participation cost, F	3620 USD	

Notes: Values in italics have been calibrated to their respective targets.

5.4 Endowments

Pensions are set to one third of the average wage, in line with the U.S. social security replacement rate. The deterministic age-productivity profile is taken from Harenberg and Ludwig (2015), who compute it from PSID data using the method of Huggett, Ventura and Yaron (2011). The values are displayed in the computational appendix.[12]

The values for the persistent income process are taken from Storesletten, Telmer and Yaron (2004). Using PSID data, they find an autocorrelation $\rho = 0.95$ and a variance of shocks of $\sigma_\upsilon^2 = 0.30$. We also use their value for the variance of the temporary income shock, which is set to $\sigma_\vartheta^2 = 0.12$. Regarding correlations with stock returns, we take the values from the baseline case of Gomes and Michaelides (2005) and set for the permanent shocks $\kappa_\upsilon = 0.15$ and for the transitory one $\kappa_\vartheta = 0.30$.

5.5 Asset markets

The parameter values for asset markets are mostly preliminary. The gross risk-free return is set to the average bond return of the last 50 years in the data of Robert Shiller, $R^f = 1.01$ percent.[13] The equity premium takes a value of $\omega = 3\%$. Stock volatility is $\sigma_\nu = 0.18$, again as measured from Robert Shiller's data over the last 50 years. Participation cost is set to a preliminary value of $F = 3620$ USD to get a reasonable stock market participation rate for the additive agent.

5.6 Computational solution

From a computational perspective, the solution is mostly standard, but there are two issues when solving the model that are worth mentioning. The first issue is that we want to numerically approximate the risks as precisely as possible. Typically, the autoregressive process driving the persistent income shock is discretized as a finite Markov chain, e.g., using the method of Tauchen and Hussey (1991). However, this has been shown to be sensitive both in statistical as well as economic terms and it cannot handle cross-correlated processes.[14] Instead of relying on a finite-state approximation, we keep the continuous representation in equations (8) and (9) and treat π_t as an additional, continuous state variable. We use 24 gridpoints to approximate this continuous state and evaluate the expectations with Gauss-Hermite quadrature, for which convergence is well-known. To evaluate continuation utility at points off the grid, we use cubic two-dimensional B-splines. Details are provided in the computational appendix.

[12] The computational appendix is available on request.

[13] Robert Shiller's data are freely available at `http://www.econ.yale.edu/~shiller/data.htm`.

[14] See, e.g., Flodén (2008) and Galindev and Lkhagvasuren (2010). Also more recent methods like Kopecky and Suen (2010) can't directly handle cross-correlated processes.

The second issue is that the model has a discrete choice—the stock market participation decision—which implies that the agent's problem is not (globally) differentiable in the continuous savings and portfolio choices.[15] As a consequence, we cannot rely on Euler equations and Newton-like nonlinear equation solvers. Even if, as is standard in the literature, we rewrite the model in terms of cash-at-hand, thereby reducing the state space by one dimension, a brute-force maximization using discretization of the state space and the choices is also infeasible. Indeed, we still have two continuous state variables, one binary state variable, 80 generations, along with two continuous and one discrete choice and want to calibrate the model to the data.[16] We solve this with a novel solution algorithm that is robust, fast, and generally applicable to finite-horizon problems. The main idea is to interpolate the expected continuation utility, $E_t U_{t+1}$, with a multi-dimensional cubic B-spline, because it can be proven that $E_t U_{t+1}$ is twice differentiable. The divide and conquer algorithm of Gordon and Qiu (2015) is then used to quickly find a bracket for a global maximum on a fine grid. Given the bracket, the maximum is then computed with high precision using a Newton-like maximizing routine, which can be defended with the result of Clausen and Strub (2012) on the local differentiability around an optimum.

On top of that, we speed up the algorithm by making use of the fact that, after minor transformations, the optimal stock choice can be represented and computed as a function of the optimal savings choice. Programmed in Fortran 2008, the code is parallelized and runs on 24 cores. Further details are provided in the computational appendix.

6 Results

We first describe the outcomes of the model, as calibrated in Section 5. Then we provide further explanation by relating our findings to the previous literature.

6.1 Lifecycle profiles

To present our results, we focus on average lifecycle profiles for agent choices. Each profile corresponds to the profile conditional on the agent surviving until the maximal age, averaged over all possible realizations for the income and investment risks. For agents who die before the maximal

[15] Even recent, more general envelope theorems are of only very limited use in a computational application. E.g., the very powerful result in Clausen and Strub (2012) is not directly applicable, because in a numerical solution we search for an optimal choice and need to evaluate continuation utility also at points that are not optimal and may therefore not be differentiable. The computational appendix provides more details.

[16] The two continuous states are cash-at-hand, x_t, and the stochastic state, π_t, the discrete states are the 80 generations and the stock market participation indicator. The continuous choices are bond and stock investments and the discrete choice is stock market participation.

age, the savings-consumption profiles are simply truncated at the age of death. Lifecycle profiles –for savings for example– are computed as follows. For a given age, we compute the optimal saving response as a function of cash-at-hand as well as the distribution (conditional on surviving) of agents in terms of cash-at-hand.[17] From both optimal saving responses and conditional distribution, we directly get the average saving at each age.

The panels in Figure 1 display the lifecycle profiles for the additive, risk-sensitive and Epstein-Zin agent. Let us first focus on total lifecycle savings, which are shown in panel (a). Overall, the shape of the saving profile is very similar for the three agents. As is typical of such lifecycle models, the agents build up savings during their working age, until they reach the exogenous retirement age of 65, and then gradually decumulate their savings. However, the agents differ markedly in the level of savings they accumulate. The additive agent saves much more on average than both the Epstein-Zin and risk-sensitive agents. The reason is the one explained in Section 2.3, namely that risk aversion amplifies the role of mortality risk in the discount rate, as long as the value of a statistical life is positive. Since both the Epstein-Zin and the risk-sensitive agent are more risk averse than the additive one, and that we assume a positive VSL, their discount rates are higher than that of the additive agent, and therefore they save less. Crucially, this effect is strong enough to overturn the higher savings due to higher prudence of more risk-averse agents. Note that the impact of risk aversion is not strictly identical for risk-sensitive and Epstein-Zin agents, except at age 45, where savings are the same because of our calibration strategy, cf. Section 5.1. In particular, savings decrease more for the risk-sensitive agent at earlier age (before 45), but the pattern is reversed for later ages (after 45).

The corresponding lifecycle consumption profiles are shown in panel (b). They are consistent with the lifecycle saving profiles. Note that consumption profiles for the three agents are hump-shaped. The risk averse agents consume more at earlier ages (between ages 30 and 60) than the additive agent. The opposite holds at older ages, greater than 60. A greater risk aversion tends therefore to increase consumption at earlier age and to decrease it a later age. Again, this is consistent with the fact that risk aversion amplifies the role of mortality in the discount rate. A more risk averse agent will be more impatient and therefore consume early.

Let us now turn to the lifecycle stock market participation rates, displayed in panel (c). The participation rates are an increasing function of age, because the cost to access the stock market is paid only once in life. Importantly, the more risk averse agents, i.e., the Epstein-Zin and risk-sensitive agents, on average participate less in the stock market. While it is intuitive that the

[17] Since the shocks are continuous, we get a continuous distribution over cash-at-hand. We approximate this distribution with a piecewise linear function over 3600 points in the cash at hand grid. For details, see the computational appendix.

Figure 1: Lifecycle profiles for benchmark economy

(a) Total savings

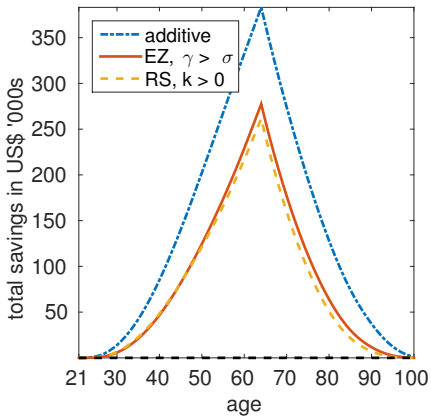

(b) Consumption

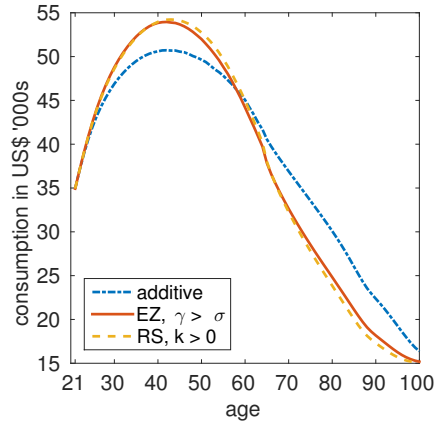

(c) Stock market participation

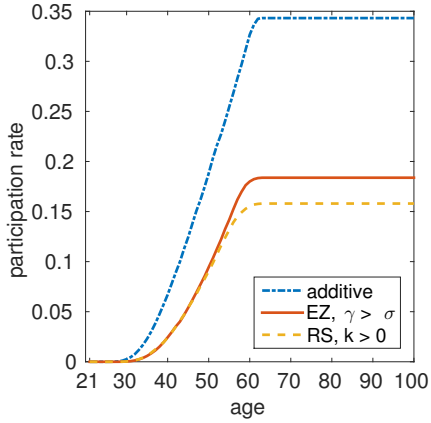

(d) Conditional share in stock

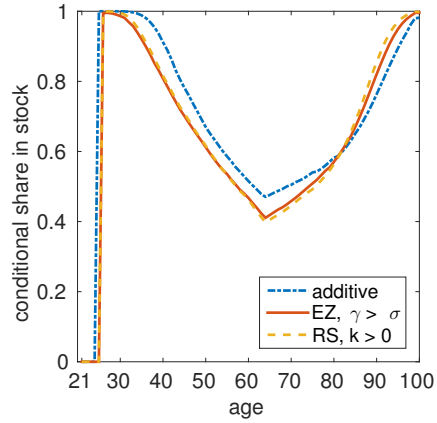

(e) VSL

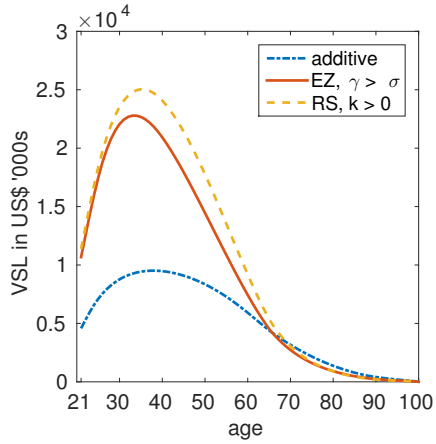

more risk-averse agents would choose to have a smaller exposure to the investment risk, this result stands in contrast to the previous literature. The reason for the different finding is that we assume a positive VSL, implying that the more risk averse agents accumulate less assets and are therefore less willing to pay the participation cost.

Conditional on participation, more risk averse agents hold a smaller share of their savings in stocks until age 80, cf. panel (d). After age 80 the wealth of agents becomes very low, which explains the increasing share in stock. Thus, more risk averse agents consistently choose to take less risks: first, by participating less in the stock market, and second, by holding a smaller share of wealth in risky assets if they participate.

Last, but not least, Panel (e) of Figure 1 plots the lifecycle profile of the VSL. It is positive at all ages, and has a inverse U-shape. As explained above, it is calibrated such that the additive agent has a VSL of approximately US $ 6.5 million at age 45. The VSL of the Epstein-Zin and the risk-sensitive agent are substantially larger until age 68 because of their higher risk aversion. After age 68, the VSL of the additive agent is slightly higher, because of his higher consumption stream at older ages.

6.2 Relation to previous literature

Our calibrated life-cycle model provides news insights on life-cycle behavior. In particular, we find that risk aversion decreases savings, stock market participation, and the share of wealth invested in stocks. This contrasts with the predictions of Gomes and Michaelides (2005, 2008) and many related papers in the household finance literature. The divergence in predictions is directly related to value of life matters. In the current paper, our model was designed to fit empirical estimate of the VSL. Therefore, by construction the model provides a large and positive VSL. In that respect, our approach differs from the standard one in household finance, which consists in focusing on a homothetic version of Epstein-Zin preferences, without paying attention to the implications regarding the value of life. According to the homothetic version of the Epstein-Zin specification, the utility conditional on being alive at time t, denoted V_t, can then be expressed as

$$V_t = \left((1-\beta) c_t^{1-\sigma} + \beta \left(E_t \left[p_t V_{t+1}^{1-\gamma} + (1-p_t) b w_{t+1}^{1-\gamma} \right] \right)^{\frac{1-\sigma}{1-\gamma}} \right)^{\frac{1}{1-\sigma}}, \tag{22}$$

where, as before, c_t is consumption in period t, p_t is the probability of remaining alive in period $t+1$, and w_{t+1} is the amount bequeathed in case of death.[18] The parameter σ is the inverse of the elasticity of substitution, γ the coefficient of relative risk aversion and b the intensity of the

[18] A formal derivation of equation (22) can be found in the appendix of Gomes, Michaelides, and Polkovnichenko (2009).

bequest motives. This model was not designed for fitting a specific VSL but rather chosen for its tractability. Equation (22) nevertheless implicitly assumes a specific value of life given by:

$$VSL_t = \frac{\frac{\partial V_t}{\partial p_t}}{\frac{\partial V_t}{\partial c_t}} = \frac{\beta c_t^{\sigma}}{1-\beta} \frac{E_t[V_{t+1}^{1-\gamma}] - bE_t[w_{t+1}^{1-\gamma}]}{1-\gamma} \left(E_t \left[p_t V_{t+1}^{1-\gamma} + (1-p_t) b w_{t+1}^{1-\gamma} \right] \right)^{\frac{\gamma-\sigma}{1-\gamma}} .$$

It occurs that the VSL may be positive or negative. In particular, if $\gamma > 1$, as is assumed in the papers mentioned above, a positive value of life is obtained only if $b > \frac{E_t[V_{t+1}^{1-\gamma}]}{E_t[w_{t+1}^{1-\gamma}]}$. The results of Gomes and Michaelides (2005, 2008) tend to indicate that such an equality does not hold (at least not always) in their simulations. In particular, a negative value of life is systematically obtained when there is no bequest motives ($b = 0$), a case that Gomes and Michaelides consider in several instances, with no significant impact on their qualitative findings about the relationship between risk aversion and savings. As explained in Section 2.3, with a negative value of life, the rate of time discounting is underestimated, the bias being amplified by risk aversion. That is, for the case of a negative VSL, risk aversion is found to lower the discount rate, and hence to increase savings, providing a conclusion which is opposite to ours.[19] The difference in saving behavior eventually generates differences in the propensity to pay the fixed cost associated with stock market participation. This explains why they find that more risk averse agents tend to participate more frequently in the stock market, while we obtain the converse.

It has often been thought as intuitive that the more risk averse an agent, the more she will save. This naturally arises in models where there is no mortality risk (see, e.g., Bommier and LeGrand, 2016). But as soon as there is mortality risk, the driving force is typically mortality. When the VSL is negative, the effect of mortality will go in the same direction as that of prudence. With a positive VSL, in contrast, the effect of mortality will go in the other direction, and indeed overturn the effect of prudence.

7 Supporting evidence

In this section we present results from a short empirical study on the relationship between risk aversion and household savings. Our results indicate that savings comove negatively with risk aversion in the data and that the regression coefficient is highly significant. We first describe the data, then explain our estimation strategy and finally present the related results.

[19] One should moreover notice that if specification (22) were to be used with $\gamma > 1$ and a parameter b large enough to generate positive values of life, we would obtain a framework where the intensity of bequest motives would increase the willingness to pay for mortality risk reduction. This goes against intuition, since deriving utility from bequest reduces the welfare gap between life and death.

7.1 Data description

The data comes from the German Socio-economic Panel (SOEP), which is an annual panel of German households starting in 1984. It covers a wide range of information, including financial situation and personal attitudes. Importantly, it also has a measure of general risk aversion, which asks respondents to rate their willingness to take risks on a scale from 0 (not willing at all) to 10 (very willing). Dohmen, et al. (2011) translate the question from German as "How do you see yourself: are you generally a person who is fully prepared to take risks or do you try to avoid taking risks? Please tick a box on the scale, where the value 0 means: 'not at all willing to take risks' and the value 10 means: 'very willing to take risks'." The question is on purpose not focusing on a single risk and doesn't make use of a lottery. This is an advantage in our context, where we consider aversion towards very different risks. As shown by Dohmen, et al. (2011), this measure is a good predictor of risky behavior such as investing in stock or smoking. They also validate the measure by conducting a field experiment with an additional representative sample of 450 subjects. They conclude that there seems to be a general trait of risk aversion affecting behavior in different contexts and that the SOEP measure is a good approximation. We are therefore confident that we can use it to study the impact of risk aversion on savings.

We use the SOEP waves 2004 and 2008-2014 because risk aversion is available only in those. We keep only observations where the household head answered the household questionnaire and associate the personal characteristics of the head to that household. After dropping observations with missing or inadmissible values for risk aversion, the sample contains 24.603 households.

Household monthly net income is available in the data. Since financial wealth is not directly available for all waves, we follow Fuchs-Schündeln and Schündeln (2005) (henceforth, FSS) and impute it from interest and dividend income. As an alternative, we also use the reported value of the house as an indicator for financial wealth.

7.2 Regressions

Our model predicts that the more risk averse an agent, the more she saves. To test this prediction, we regress the logarithm of monthly savings on the risk aversion index, while controlling for several factors.[20] In particular, we need to control for income, mortality risk, wealth, as well as other saving motives. The first control we use is the logarithm of financial wealth, to account for the fact that savings are path-dependent and that the prediction of our model is for agents endowed with the same initial wealth. For the other controls, we closely follow FSS. More precisely, regarding

[20] Formally, the index in the data varies between 0 and 10 and the larger the index, the more the respondent is willing to take risks. We then define our risk aversion index as 10−this index.

the income process, we control for permanent income and a measure of the income risk. For permanent income, we replicate the construction of FSS.[21] To limit the effect of measurement errors and of bias due to the small number of observations for every household, we run a two-stage regression, by instrumenting the logarithm of permanent income by education and age variables (including the square of these variables and the interaction term). For every household, we measure income risk as the standard deviation of the difference between actual income and permanent income (at the household level). Second, we control mortality risk as in FSS by using age and age squared. Other saving motives are proxied by the following control variables: education, square of education, gender, household size, number of children, and dummy variables for marital status, current residence (1 if in former West Germany and 0 otherwise), and survey year. Finally, we cluster standard errors by household.

In order to further test the model predictions, we include two other variables besides risk aversion. The first one is an interaction term between risk aversion and income risk, capturing a precautionary savings motive. The model predicts this regression coefficient to be positive. The second additional variable is an interaction term between risk aversion and mortality risk. According to the model, the regression coefficient should be negative.

On the whole, we run four regressions. The first regression is the one including risk aversion, as well as all control variables. It does not include any of the two interaction terms. The second and third regressions include one single additional interaction term, between risk aversion and income risk and between risk aversion and mortality risk respectively. Finally, the last regression includes both interaction terms together.

7.3 Results

The results of our four regressions are displayed in Table 2.

Table 2: Regression results

Dependent variable = log(savings)	Risk aversion only (1)	(1) + interaction with income risk (2)	(1) + interaction with mortality risk (3)	(1) + both interactions (4)
risk aversion	-0.013^{***}	-0.020^{***}	-0.010^{**}	-0.016^{***}
	(0.004)	(0.006)	(0.004)	(0.006)

[21] See their Footnote 20, p. 1098. First, we compute detrended income every year as the ratio of household income divided by average income for all households in the corresponding survey year. Permanent income is then equal to average detrended household income (computed for every household over all survey years) multiplied by average income of all households within each survey year.

risk aversion		0.036		0.033
× income risk		(0.026)		(0.026)
risk aversion			-0.200^{*}	-0.191^{*}
× mortality (. 10^4)			(0.118)	(0.118)
log(wealth)	0.088^{***}	0.088^{***}	0.088^{***}	0.088^{***}
	(0.004)	(0.004)	(0.004)	(0.004)
permanent income	−0.005	−0.010	−0.014	−0.018
	(0.208)	(0.208)	(0.208)	(0.208)
income risk	0.285^{***}	0.098	0.285^{***}	0.111
	(0.081)	(0.156)	(0.081)	(0.157)
age	0.010^{*}	0.011^{*}	0.007	0.007
	(0.006)	(0.006)	(0.006)	(0.006)
age squared	−0.009	−0.009	−0.004	-0.004^{***}
(. 10^2)	(0.006)	(0.006)	(0.004)	(0.004)
educ.	0.196^{***}	0.196^{***}	0.197^{***}	0.197^{***}
	(0.042)	(0.042)	(0.042)	(0.042)
educ. squared	-0.004^{***}	-0.004^{***}	-0.004^{***}	-0.004^{***}
	(0.001)	(0.001)	(0.001)	(0.001)
gender (1 if male)	0.008	0.008	0.008	0.008
	(0.011)	(0.011)	(0.011)	(0.011)
married and	0.195^{***}	0.195^{***}	0.196^{***}	0.196^{***}
not separated	(0.030)	(0.030)	(0.030)	(0.030)
married	-0.216^{***}	-0.217^{***}	-0.213^{***}	-0.213^{***}
and separated	(0.066)	(0.066)	(0.066)	(0.066)
divorced	-0.167^{***}	-0.168^{***}	-0.164^{***}	-0.165^{***}
	(0.039)	(0.039)	(0.039)	(0.039)
widowed	0.186^{***}	0.186^{***}	0.194^{***}	0.194^{***}
	(0.047)	(0.047)	(0.047)	(0.047)
not married and	0.514	0.508	0.521	0.514
living together	(0.442)	(0.447)	(0.445)	(0.449)
household size	0.202^{***}	0.201^{***}	0.202^{***}	0.201^{***}
	(0.020)	(0.020)	(0.020)	(0.020)
number of children	-0.235^{***}	-0.234^{***}	-0.233^{***}	-0.232^{***}
	(0.022)	(0.022)	(0.022)	(0.022)

current residence	0.322***	0.322***	0.323***	0.323***
(1 if West)	(0.025)	(0.025)	(0.025)	(0.025)
constant	1.918	1.991	2.027	2.090*
	(1.268)	(1.268)	(1.269)	(1.269)
year dummies	yes	yes	yes	yes
observations	22973	22973	22973	22973
R^2	21.56	21.57	21.57	21.58

*** significant at the 1 percent level; ** significant at the 5 percent level; * significant at the 10 percent level.

The first conclusion that we can draw from our regression results is that risk aversion has a negative impact on savings. More risk averse agents, when controlling for wealth, income risk, and mortality risk, tend to save less. The regression coefficient is consistently highly significant in all regressions. Looking at the specifications that include interaction terms, we find that the interaction between risk aversion and income risk has a positive effect, though not highly significant (p-values vary between 15% and 20%). This finding, in spite of weak significance, is consistent with the predictions of our model. The interaction between risk aversion and mortality risk has a negative impact on savings. The regression coefficient is significant at the 10% level. The (negative) effect of risk aversion on savings is amplified by the level of mortality. This last finding is also predicted by our model.

8 Conclusion

The notion of value of life is generally not discussed in the household finance literature. Actually, there is a real cost in using models that are flexible enough to provide realistic estimates of value of life: preference homotheticity has to be relaxed, which of course translates into an increased difficulty at the optimization stage. With the current paper, we want to argue that it is definitely worth making this extra effort. The value of life plays indeed a key role in the relation between mortality, risk aversion and time discounting. It is therefore a key determinant of life-cycle behavior, even if mortality is exogenous.

Once the value of life is set at reasonable positive levels, we find that risk aversion has a negative impact on savings, stock market participation and the the share of wealth held in risky assets. The basic intuition is the following: saving involves keeping resources for periods that are only lived in favorable odds (i.e., in a case of a long life). Saving is thus like a bet that pays in good outcomes (i.e., survival) an therefore a risk increasing behavior. Our results indicate, rather intuitively, that risk aversion reduces the propensity to engage in risky behaviors and therefore to save. The results

regarding stock market participation and the the share of wealth held in risky assets reflect both direct risk aversion effect and indirect wealth effect. Let us however notice that if we were to assume that having a long life is a bad outcome, that is if assuming a negative value of life, the logic would be reversed. Savings would then imply keeping resources for adverse realizations (long lives) and would then be a risk reducing behavior. The result would then be that risk aversion increases savings, exactly because of the assumption of a negative value of life.

One of the implications of our study, is that relatively low level of savings could be explained by risk aversion. While the economic literature abounds of works arguing that observed saving behaviors have to reflect strongly myopic preferences or some form of irrationality, our analysis suggests on the contrary that saving little could just be a rational decision for risk averse agents who are well aware that life duration is uncertain. Of course, low saving levels typically result in having a majority of (surviving) elderly declaring that they failed to save enough. But this is not evidence of under-savings. If they could be questioned, those who died before retirement may well answer that they actually saved too much.

References

Ameriks, J., Caplin, A., Laufer, S. and Van Nieuwerburgh, S.: 2011, The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Bequest and Precautionary Motives, *Journal of Finance* **66**(2), 519–561.

Bommier, A., Kochov, A. and LeGrand, F.: 2016, On Monotone Recursive Preferences, *mimeo*, ETH Zurich.

Bommier, A. and Le Grand, F.: 2016, Risk aversion and precautionary savings in dynamic settings, *Ssrn working paper 2809553*.

Bommier, A. and LeGrand, F.: 2014, A Robust Approach to Risk Aversion, *Working paper*, ETH Zurich.

Clausen, A. and Strub, C.: 2012, Envelope Theorems for Non-Smooth and Non-Concave Optimization, *Working paper*, University of St. Gallen.

De Nardi, M.: 2004, Wealth Inequality and Intergenerational Links, *Review of Economic Studies* **71**(3), 743–768.

De Nardi, M., French, E. and Jones, J. B.: 2010, Why Do the Elderly Save? The Role of Medical Expenses, *Journal of Political Economy* **118**(1), 39–75.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J. and Wagner, G. G.: 2011, Individual risk attitudes: Measurement, determinants, and behavioral consequences, *Journal of the European Economic Association* **9**(3), 522–550.

Epstein, L. G. and Zin, S. E.: 1989, Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, *Econometrica* **57**(4), 937–969.

Flodén, M.: 2008, A Note on the Accuracy of Markov-Chain Approximations to Highly Persistent AR(1) Processes, *Economics Letters* **99**(3), 516–520.

Fuchs-Schuendeln, N. and Schuendeln, M.: 2005, Precautionary Savings and Self-Selection: Evidence from the German Reunification "Experiment", *Quarterly Journal of Economics* **120**(3), 1085–1120.

Galindev, R. and Lkhagvasuren, D.: 2010, Discretization of Highly Persistent Correlated AR(1) Shocks, *Journal of Economic Dynamics and Control* **34**(7), 1260–1276.

Gomes, F. J. and Michaelides, A.: 2008, Asset Pricing with Limited Risk Sharing and Heterogeneous Agents, *The Review of Financial Studies* **21**(1), 415–448.

Gomes, F. and Michaelides, A.: 2005, Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence, *The Journal of Finance* **60**(2), 869–904.

Gomes, F., Michaelides, A. and Polkovnichenko, V.: 2009, Optimal savings with taxable and tax-deferred accounts, *Review of Economic Dynamics* **12**(4), 718–735.

Gordon, G. and Qiu, S.: 2015, A Divide and Conquer Algorithm for Exploiting Policy Function Monotonicity, *CAEPR Working Paper 2015-002*, Center for Applied Economics and Policy Research, Indiana University.

Hansen, L. P. and Sargent, T. J.: 1995, Discounted Linear Exponential Quadratic Gaussian Control, *IEEE Transactions on Automatic Control* **40**(5), 968–971.

Harenberg, D. and Ludwig, A.: 2015, Social Security in an Analytically Tractable Overlapping Generations Model with Aggregate and Idiosyncratic Risk, *CER-ETH Working Paper 14/204*, ETH Zurich.

Huggett, M., Ventura, G. and Yaron, A.: 2011, Sources of Lifetime Inequality, *American Economic Review* **101**(7), 2923–2954.

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

Inkmann, J., Lopes, P. and Michaelides, A.: 2011, How Deep Is the Annuity Market Participation Puzzle?, *Review of Financial Studies* **24**(1), 279–319.

Johansson, P.-O.: 2002, On the Definition and Age-Dependency of the Value of a Statistical Life, *Journal of Risk and Uncertainty* **25**(3), 251–263.

Kopecky, K. A. and Suen, R. M.: 2010, Finite State Markov-Chain Approximations to Highly Persistent Processes, *Review of Economic Dynamics* **13**(3), 701–714.

Kreps, D. M. and Porteus, E. L.: 1978, Temporal Resolution of Uncertainty and Dynamic Choice Theory, *Econometrica* **46**(1), 185–200.

Lockwood, L. M.: 2012, Bequest Motives and the Annuity Puzzle, *Review of Economic Dynamics* **15**(2), 226–243.

Lockwood, L. M.: 2014, Bequest Motives and the Choice to Self-Insure Late-Life Risks, *NBER Working Paper 20745*, National Bureau of Economic Research.

Storesletten, K., Telmer, C. I. and Yaron, A.: 2004, Consumption and Risk sharing over the Life Cycle, *Journal of Monetary Economics* **51**(3), 609–633.

Strzalecki, T.: 2011, Axiomatic foundations of multiplier preferences, *Econometrica* **79**(1), 47–73.

Tallarini, T. D. J.: 2000, Risk-Sensitive Real Business Cycles, *Journal of Monetary Economics* **45**(3), 507–532.

Tauchen, G. and Hussey, R.: 1991, Quadrature-Based Methods for Obtaining Approximate Solutions to Nonlinear Asset Pricing Models, *Econometrica* **59**(2), 371–396.

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 Aldy, J. E.: 2003, The Value of a Statistical Life: A Critical Review of Market Estimates Throughout the World, *The Journal of Risk and Uncertainty* **27**(1), 5–76.

Wakai, K.: 2007, Risk Non-Separability without Force of Habit, *Working paper*, The State University of New York at Buffalo.