The Optimal Allocation of Longevity Risk with Perfect Insurance Markets

February 26, 2018

Antoine Bommier*
ETH Zurich
email: abommier@ethz.ch

Hélène Schernberg*
ETH Zurich
email: hschernberg@ethz.ch

Abstract

This paper discusses the allocation of aggregate longevity risk in the case of perfect insurance markets. We show that the optimal allocation transfers some risk to the pensioners, even if pension providers have access to a perfect insurance market. Individuals prefer contributions and benefits to depend on the evolution of aggregate mortality rates rather than being fixed. Indeed, this flexibility offers an interesting diversification strategy where the prospect of a shorter life (e.g. the emergence of new diseases) implies higher consumption levels and conversely, the prospect of a longer life (e.g. thanks to medical progress) implies lower consumption levels. The underlying mechanism only emerges when individuals are temporally risk averse. We illustrate it with risk-sensitive preferences.

Keywords: pensions, longevity risk, risk-sharing, risk-sensitive preferences.

JEL codes: D91, G28, H55, J32.

---

*We gratefully acknowledge financial support from the Swiss Re foundation and the ETH Zurich foundation.

1 Introduction

Life expectancy changed abruptly in the last century. In many societies, longevity increased at an unexpected pace while it strongly declined in others due to epidemics such as AIDS. It is hard to predict the direction, let alone the magnitude, of such changes in the future. This uncertainty is a major stumbling block when designing a pension system. In particular, when benefits are fixed, unexpected longevity gains directly translate into financial losses for pension providers, threatening the sustainability of pension systems. To avoid this outcome, one could use a tontine-like system where fixed aggregate benefits are divided among survivors. Fluctuations in survival would not affect aggregate pension payments, thus relieving providers from the aggregate longevity risk (see Milevsky 2015). Pensioners' incomes, however, would become highly unpredictable. The optimal design probably lies somewhere between these two extreme cases, with a sophisticated form risk sharing between pension providers and participants that relies on flexible levels of contributions and benefits.

The current paper contributes to the debate on the optimal allocation of longevity risk by exploring the simple case where pension providers can access a perfect insurance market. This assumption simplifies our analysis but does not reflect reality well. Nonetheless, a good understanding of this simple case is a prerequisite for the development of more complex analyses featuring imperfect insurability. We show that even if pension providers can access fair insurance, individuals retain some aggregate longevity risk at the optimum. Individuals prefer a system that negatively correlates pension benefits with life expectancy to one granting them fixed benefits. Such a system combines adverse demographic shocks (e.g. a shorter life expectancy triggered by an epidemic) with improved pension benefits and, conversely, combines improved longevity prospects with a reduction in benefits. This preference for negative correlation, that is for combining "goods" with "bads", only emerges under temporal risk aversion. We illustrate it with risk-sensitive preferences.

First, we develop our argument in a simple three-periods setting, where individuals die according to exogenous, uncertain probabilities. Importantly, we allow for some dependency in the evolution of survival probabilities. The information released over time through the observation of past survival rates thus informs about future survival prospects and can then be used to establish future pension levels. Our theoretical model highlights that dynamic adjustments are beneficial, even if pension providers can access fair insurance.

Second, we quantify these adjustments, up to an order of magnitude, by calibrating and numerically solving a multi-period life-cycle model. Our analysis covers both the active life and retirement, as the aforementioned dynamic adjustments affect contributions as well as benefits. In line with our theoretical findings, we find that contributions decrease and benefits increase when survival prospects worsen while contributions increase and benefits decrease when survival prospects improve. Importantly, these dynamic adjustments of contributions and benefits may relieve pension providers from about half of the risk of financial losses generated by this aggregate longevity risk, without lowering individuals' *ex-ante* utilities.

Our approach both differs from and complements the vast literature on optimal risk sharing, and in particular contributions focusing on longevity risk. Most of those studies emphasize the impossibility of insuring longevity risk. They explore several mechanisms for allocating the risk across different generations or different types of agents.[1] They take it for granted that the risk allocation would be trivial if the longevity risk could be insured at no cost. Instead, we emphasize that even in the case of perfect insurability, determining the optimal risk allocation is non-trivial. This is because life is an irreplaceable commodity. Hence, the mortality risk cannot be removed by a mere redistribution of wealth. In other words, pensions become devices for implementing the best allocation of wealth when agents face a non-insurable mortality risk. As we will show, the fact that mortality rates are themselves uncertain impacts individuals' demand for insurance, and therefore the risk allocation at the optimum.

The paper is structured as follows. In Section 2, we introduce the theoretical framework. Section 3 characterizes the optimal pension profile and provides our main findings, which are discussed in depth in Section 4. In Section 5, we numerically quantify the share of the aggregate longevity risk transferred to individuals. Section 6 concludes.

[1]Because future generations are not born yet, they cannot share the aggregate longevity risk with current ones. Gordon and Varian (1988) explore the optimal risk-sharing across generations when a government has the power to pre-commit future generations. Demange and Laroque (2000) compare the optimality of two pension systems: Mandatory contributions based on capital and labor income, and voluntary pay-as-you-go Social security. Bohn (2001) characterizes the risk-sharing property of four alternative Social security systems, and emphasizes that efficiency requires that all agents in the economy, including current retirees, share all risks. Krueger and Kubler (2002) suggest that, whenever financial markets are incomplete, the introduction of Social security in a competitive economy provides a Pareto-improving allocation of risks between generations. Ball and Mankiw (2007) examine the optimalit risk-sharing in a fully funded Social security system. If the Social security fund holds safe debt, they find that, at the optimum, , pension benefits must be negatively indexed to equity returns.

2 A simple theoretical model

2.1 Setting

The economy comprises a large number of identical individuals and a single pension provider. In period 0, individuals are active, consume an exogenous level c_0 and face no risk. They retire at the end of period 0 and rely on retirement benefits for the remaining two periods. Individuals may die at the end of period 0 (i.e. when reaching retirement), at the end of period 1, or at the end of period 2. Nobody survives after period 2. Individuals only derive utility from their own consumption, thus do not bequeathe.

Individuals face risks at the end of periods 0 and 1, since in both cases they may not survive until the following period. This standard mortality risk is combined with some aggregate longevity risk, which reflects an imperfect knowledge of the underlying mortality process. Henceforth, longevity risk refers to the fact that survival probabilities at the end of period 0 and 1 are themselves uncertain *ex-ante*. Interestingly, the number of survivors in period 0 is observed by individuals and pension providers alike at the end of period 0, hence revealing the actual survival rate at the end of period 0. Moreover, we assume that individual's subjective beliefs about their own survival coincide with the observed aggregate probabilities. The information acquired at the end of period 0 allows updating the probability distribution of the survival rates at the end of period 1. The question we address is how, if at all, this learning should be reflected in future pensions levels.

Notation Time is discrete, indexed by $t \in \mathcal{T} = \{0, 1, 2\}$. There is a finite set of states of the world Ω. For each $\omega \in \Omega$, we denote by p_ω the probability that the world is in that state. To any state of the world $\omega \in \Omega$ corresponds a vector of survival probabilities $\pi^\omega = (\pi_0^\omega, \pi_1^\omega) \in [0, 1]^2$, where π_0^ω is the probability to survive until the end of period 0 and π_1^ω is the probability to survive until the end of period 1, conditionally on being alive in period 1. *Ex-ante*, the true state of the world cannot be observed. However, π_0^ω is observed at the end of period 0.[2] This information leads individuals and pension providers to update beliefs in a standard Bayesian way. For any $\omega \in \Omega$,

[2]As we assume a large homogeneous population the fraction of individuals that survive provides an exact estimate of the survival probabilities.

we define $p_{\omega|\pi_0^\omega} \in [0,1]$ by:

$$p_{\omega|\pi_0^\omega} = \frac{p_\omega}{\sum_{\{\omega'|\pi_0^{\omega'}=\pi_0^\omega\}} p_{\omega'}}.$$

This is the probability to be in state ω, when the survival rate observed at the end of period 0 equals π_0^ω. We also define $\bar{\pi}_{1|\pi_0^\omega} \in [0,1]$ by:

$$\bar{\pi}_{1|\pi_0^\omega} = \sum_{\{\omega'|\pi_0^{\omega'}=\pi_0^\omega\}} p_{\omega'|\pi_0^\omega} \pi_1^{\omega'}.$$

This is the expected survival rate at the end of period 1, when the mortality rate observed at the end of period 0 equals π_0^ω. By the law of iterated expectations, $\bar{\pi}_{1|\pi_0^\omega}$ can also be interpreted as the probability that an individual survives until period 2, conditional on surviving until period 1 and on the past survival rate being π_0^ω.

Pensions and state-contingent consumption profile A well-designed pension system should allow its participants to reach the highest possible level of welfare given a budget constraint. In our setting, there is no utility for bequests. Thus, at the optimum, individuals do not save (see Davidoff et al. 2005). Searching for the optimal pension system therefore amounts to deriving the optimal consumption pattern. Moreover, given the uncertain survival rates and sequential learning, the optimal strategy may involve state-contingent consumption processes. These are formalized by a mapping $\omega \in \Omega \to (c_1^\omega, c_2^\omega) \in \mathbb{R}_+^2$, where c_1^ω and c_2^ω denote the amount consumed by (living) individuals in periods 1 and 2 in state ω. Such state-contingent consumption processes must be *adapted*: Information that is only revealed at the end of a given period cannot be used to determine the consumption level during that period. Thus, for any two distinct states $(\omega, \omega') \in \Omega^2$ such that $\pi_0^\omega = \pi_0^{\omega'}$ (i.e., states which cannot be distinguished before the end of period 1) one must have $c_1^\omega = c_1^{\omega'}$. Moreover, for any two states $(\omega, \omega') \in \Omega^2$ such that $\pi_0^\omega = \pi_0^{\omega'}$ and $\pi_1^\omega = \pi_1^{\omega'}$ (i.e., states which cannot be distinguished before the end of period 2), one must have $c_1^\omega = c_1^{\omega'}$ and $c_2^\omega = c_2^{\omega'}$. We denote by $\mathcal{C}$ the set of adapted state-contingent consumption processes.

We say that a state-contingent consumption process is deterministic if (c_1^ω, c_2^ω) is independent of ω, and adaptive if $(c_1^\omega, c_2^\omega) \neq (c_1^{\omega'}, c_2^{\omega'})$ for some states of the world $(\omega, \omega') \in \Omega^2$.

The present value of retirement consumption We assume perfect intertemporal markets with a discount factor δ. In state ω, a retiree receives her retirement income in period 1 with probability π_1^ω and in period 2 with probability $\pi_1^\omega \pi_2^\omega$. Thus, the present value of retirement consumption in state $\omega \in \Omega$ equals $\pi_1^\omega c_1^\omega + \delta \pi_1^\omega \pi_2^\omega c_2^\omega$.

The present value of retirement consumption is uncertain *ex-ante*, since ω is not known. Assuming that pension providers can access a perfect insurance market, they care about the expected cost of providing state contingent consumption (c_1^ω, c_2^ω), that is $\sum_\omega p_\omega (\pi_1^\omega c_1^\omega + \delta \pi_1^\omega \pi_2^\omega c_2^\omega)$.

The diagram below summarizes the timing of life or death occurrences, information release and consumption.

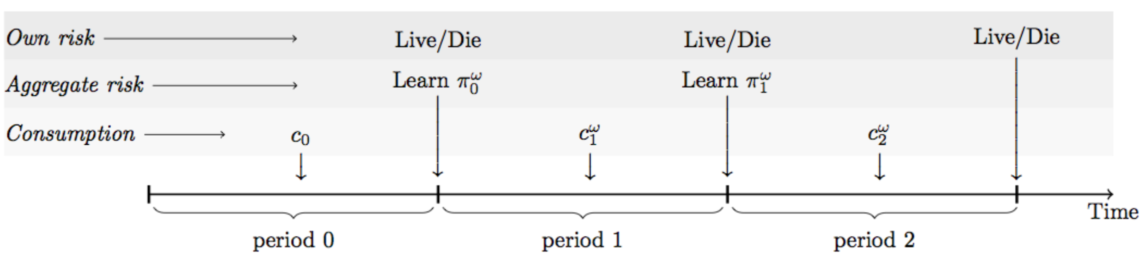

Figure 1: Timing of life or death, information release and consumption.

2.2 Preferences

We consider individuals endowed with risk sensitive-preferences. Utility is defined recursively by:

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

These preferences, introduced by Hansen and Sargent (1995), are similar to Epstein-Zin preferences but feature monotonicity (Bommier et al. 2017). They make it possible to disentangle ordinal preferences (determined by the instantaneous utility function u and the time preference parameter β) and risk aversion (driven by the parameter k). The higher k, the higher risk aversion. The case where $k = 0$, obtained by taking the limit $k \to 0$ in (1), corresponds to the standard additive model with exponential discounting. Most of the literature on life-cycle theory relies on this additive case, assuming temporal risk neutrality. This is particularly restrictive. Consider for example the following two lotteries in an infinite horizon setting:

$$L_A \to \begin{cases} (c, c, c, c, \cdots) & \text{with prob. } \frac{1}{2} \\ (C, C, C, C, \cdots) & \text{with prob. } \frac{1}{2} \end{cases} \text{ and } L_B \to \begin{cases} (c, C, c, C, \cdots) & \text{with prob. } \frac{1}{2} \\ (C, c, C, c, \cdots) & \text{with prob. } \frac{1}{2} \end{cases}$$

where c and C are consumption levels with $c < C$. When $k = 0$, the individual is indifferent between both lotteries. Intuitively, however, lottery L_B is less risky than lottery L_A, since instead of delivering the worst outcome $(c, c, c, \cdots)$ or the the best outcome $(C, C, C, C, \cdots)$ with equal probabilities, lottery L_B delivers the intermediate outcomes $(c, C, c, C, \cdots)$ and $(C, c, C, c, \cdots)$ with equal probabilities. When $k > 0$, L_B is preferred to L_A. This reflects a willingness to combine "goods" with "bads" (see Epstein and Tanny 1980 and Eeckhoudt et al. 2007).[3]

Additionally, recursive preferences imply non-trivial preferences for the timing of the resolution of uncertainty (Kreps and Porteus 1978). In the case of risk-sensitive preferences, individuals prefer an early resolution of uncertainty when $k > 0$ and $\beta < 1$, and are indifferent to the timing of the resolution of uncertainty when $k = 0$ or $\beta = 1$.[4]

Risk-sensitive preferences are most often used to model infinitely long-lived individuals. Yet, they are also relevant when individuals have a finite but possibly uncertain horizon. Bommier (2014) first applied risk-sensitive preferences in a context of uncertain lifetime. A wider theoretical discussion is provided Bommier et al. (2017). In practice, denote by U_t the lifetime utility at time t conditionally on being alive at time t, and normalize the utility of being dead to 0. The individuals' *ex-ante* utility U_0 is computed inductively using equation (1) and accounting for the survival probabilities in period $t+1$. Applying it to our three-period model, we obtain that a state-contingent survival profile $(\pi_0^\omega, \pi_1^\omega)$

[3] The temporally risk-neutral case where $k = 0$ is not supported by evidence (see Coble and Lusk 2010, Abdellaoui et al. 2013 or Leroux et al. 2016). In an experiment, Andersen et al. (2018) present subjects with a series of choices between lotteries such as L_A and L_B and find evidence of temporal risk aversion. Cheung (2015) and Miao and Zhong (2015) also reports evidence of a preference for a so-called "intertemporal diversification" motive.

[4] The case where $\beta = 1$ corresponds to the multiplicative model studied in Bommier (2013).

and an adapted consumption process (c_1^ω, c_2^ω) provide the utility U_0 given by:

$$U_0 = u(c_0) - \frac{\beta}{k} \log \left\{ \sum_{\omega \in \Omega} p_\omega \left[\pi_0^\omega e^{-kU_1^\omega} + (1 - \pi_0^\omega) \right] \right\}, \tag{2}$$

$$U_1^\omega = u(c_1^\omega) - \frac{\beta}{k} \log \left\{ \sum_{\{\omega' | \pi_0^{\omega'} = \pi_0^\omega\}} p_{\omega' | \pi_0^\omega} \left[\pi_1^{\omega'} e^{-kU_2^{\omega'}} + \left(1 - \pi_1^{\omega'}\right) \right] \right\}, \tag{3}$$

$$U_2^\omega = u(c_2^\omega). \tag{4}$$

Henceforth, we assume the instantaneous utility function u to be increasing and concave. Moreover, to avoid possible corner solutions, we assume that $\lim_{c \to 0} u'(c) = +\infty$. The parameter β reflecting pure time preference is assumed to take values in $(0, 1]$, while the risk aversion parameter k is constrained to be non-negative.

3 Optimal pension systems

We assume that pension providers have access to a perfect insurance market. Hence, the optimization problem is the following:

$$\max_{\{(c_1^\omega, c_2^\omega)\}_{\omega \in \Omega} \in \mathcal{C}} U_0$$

$$\text{s.t.} \quad \sum_\omega p_\omega (\pi_0^\omega c_1^\omega + \delta \pi_1^\omega \pi_2^\omega c_2^\omega) \le B$$

where B is the endowment of the pension provider. In this theoretical part, we assume that B is an exogenous constant.

3.1 First-order conditions

Because risk-sensitive preferences are convex, we address the problem by solving its first-order conditions. We show in Appendix A that these first-order conditions write as follows: There exist $\mu \in \mathbb{R}_+$ such that for all $\omega \in \Omega$:

$$e^{-kU_1^\omega} u'\left(c_1^\omega\right) = \mu \tag{5}$$

$$\beta e^{-\frac{k}{\beta} u\left(c_1^\omega\right)} e^{-ku\left(c_2^\omega\right)} u'\left(c_2^\omega\right) e^{k\frac{1-\beta}{\beta} U_1^\omega} = \mu\delta. \tag{6}$$

Before proceeding further, we introduce a technical result that will prove useful later on. It relies on the notion of *comonotonicity*, defined as follows. We say that two real-valued processes $\{x^\omega\}$ and $\{y^\omega\}$ are strictly comonotone if, for all $(\omega, \omega') \in \Omega^2$ one has $x^\omega > x^{\omega'} \Leftrightarrow y^\omega > y^{\omega'}$. Strict anti-comonotonicity corresponds to to $x^\omega > x^{\omega'} \Leftrightarrow y^\omega < y^{\omega'}$. Note that if $\{x^\omega\}$ and $\{y^\omega\}$ are both strictly comonotone and strictly anti-comonotone, then they must be constant (i.e. x^ω and y^ω must be independent of ω). The first-order conditions (5) and (6) yield the following property:

Lemma 1. *The optimal state-contingent consumption profile is such that* $\{c_1^\omega\}$ *and* $\{c_2^\omega\}$ *are strictly anti-comonotone.*

Proof. Combining (5) and (6) one gets:

$$\beta e^{-ku\left(c_2^\omega\right)} u'\left(c_2^\omega\right) = \mu^{\frac{1}{\beta}} \delta e^{\frac{k}{\beta} u\left(c_1^\omega\right)} \left(u'(c_1^\omega)\right)^{\frac{\beta-1}{\beta}}. \tag{7}$$

Anti-comonotonicity directly follows from the monotonicity and concavity of u, and from $k \geq 0$ and $\beta \leq 1$. □

3.2 Information and adapted consumption profiles

This section explores whether and how an optimal pension system should react to the incremental receipt of information about survival rates. A first result indicates that the information revealed at the beginning of the last period has no instrumental value.

Proposition 1. *The optimal pension profile does not use the information revealed at the end of period 1. Formally for any two states* $(\omega, \omega') \in \Omega^2$ *which are indistinguishable before the end of period 1 (i.e. states such that* $\pi_0^\omega = \pi_0^{\omega'}$*) one must have* $(c_1^\omega, c_2^\omega) = (c_1^{\omega'}, c_2^{\omega'})$.

Proof. Assume that $\pi_0^\omega = \pi_0^{\omega'}$. Because (c_1^ω, c_2^ω) is an adapted consumption process, $c_1^\omega = c_1^{\omega'}$. Then, $c_2^{\omega'} = c_2^\omega$ directly follows from equation (7). □

For consequentialist individuals, there is no value in learning about past realizations when all uncertainty has resolved. Here, all uncertainty resolves at the end of period 1 (see Figure 1). In period 2, either the individual is already dead, or she knows that she has exactly one period left to live. In both cases (whether she is dead or alive), she obtains a utility that is independent of $\bar{\pi}_{1|\pi_0^\omega}$, meaning that period 2 utility is state-independent. It is well known from Borch (1960) that if utilities are state-independent and if there exists a risk-neutral agent, then, at the optimum, the risk-neutral agent should bear all the risk in the economy. Thus, pension providers cover all the uncertainty about π_1^ω that is only revealed at the end of period 1.

Proposition 1 bears on the choice of c_2^ω. If we turn to the choice of c_1^ω, the argument of state-independence also applies in the case where $\bar{\pi}_{1|\pi_0^\omega}$ is independent of π_0^ω. That is when observing π_0^ω does not inform about future survival probabilities. Here again, risk neutral agents bears all risks:

Proposition 2. *Assume that for any $(\omega, \omega') \in \Omega^2$ one has $\bar{\pi}_{1|\pi_0^\omega} = \bar{\pi}_{1|\pi_0^{\omega'}}$. Then (c_1^ω, c_2^ω) must be independent from ω.*

Proof. As a consequence of Proposition 1, one has:

$$U_1^\omega = u\left(c_1^\omega\right) - \frac{\beta}{k}\log\left\{e^{-ku\left(c_2^\omega\right)}\bar{\pi}_{1|\pi_0^\omega} + \left(1 - \bar{\pi}_{1|\pi_0^\omega}\right)\right\}. \tag{8}$$

Combining (5) and (6) we have:

$$e^{-\frac{k}{\beta}\left(U_1^\omega - u\left(c_1^\omega\right)\right)} = \frac{\beta\delta^{-1}e^{-ku\left(c_2^\omega\right)}u'\left(c_2^\omega\right)}{u'\left(c_1^\omega\right)}. \tag{9}$$

From (8) and (9) one derives:

$$u'\left(c_1^\omega\right) = \frac{\beta\delta^{-1}u'\left(c_2^\omega\right)}{\bar{\pi}_{1|\pi_0^\omega} + \left(1 - \bar{\pi}_{1|\pi_0^\omega}\right)e^{ku\left(c_2^\omega\right)}} \tag{10}$$

Since we assumed $\bar{\pi}_{1|\pi_0^\omega}$ to be independent of ω, we conclude from (10) that $\{c_1^\omega\}$ and $\{c_2^\omega\}$ are strictly co-monotone. Given that $\{c_1^\omega\}$ and $\{c_2^\omega\}$ are also anti-comonotone (Lemma 1), both c_1^ω and c_2^ω must be independent of ω. □

Of course, assuming that $\bar{\pi}_{1|\pi_0^\omega}$ is independent of π_0^ω is extremely restrictive. More realistically, one may believe that a favorable outcome in terms of survival at the end of period 0 reflects the impact of, say, medical progress. Medical progress will also improve survival rates in future periods. All popular stochastic mortality models allow for such dependency between survival rates.

If survival rates in different periods are not independently distributed, some useful information is revealed at the end of period 0 and the utility function in period 1 is state-dependent. Formally this is reflected in equation (8) where U_1^{ω} may depend on π_0^{ω}. Therefore, we expect that the result of Proposition 2 may not extend to such a setting. This actually depends on whether individuals are temporally risk averse or not (i.e., whether $k > 0$ or $k = 0$). First let us consider the case where $k = 0$, which corresponds to the standard time-additive model of Yaari (1965).

Proposition 3. *If $k = 0$, then $(c_1^{\omega}, c_2^{\omega})$ must be independent from ω.*

Proof. If $k = 0$, the first-order conditions write $u'(c_1^{\omega}) = \mu$ and $\beta u'(c_2^{\omega}) = \mu\delta$, directly implying that $(c_1^{\omega}, c_2^{\omega})$ must be independent from ω. □

When $\bar{\pi}_{1|\pi_0^{\omega}}$ depends on ω, then the utility U_1^{ω} depends on π_0^{ω}. However, under temporal risk-neutrality, that is $k = 0$, the state-dependence takes an additive form. In this case, the risk-neutral agent takes all the risk.

We now turn to the case of temporal risk aversion, or $k > 0$, which leads to the main result of this theory part. We find that temporal risk aversion implies specific co-movements of consumption and aggregate survival probabilities:

Proposition 4. *Assume $k > 0$ and that the optimal profile is such that $u(c_2^{\omega}) > 0$ for all $\omega \in \Omega$.*[5] *Then:*

$$\bar{\pi}_{1|\pi_0^{\omega}} > \bar{\pi}_{1|\pi_0^{\omega'}} \Leftrightarrow \left(c_1^{\omega} < c_1^{\omega'} \text{ and } c_2^{\omega} > c_2^{\omega'}\right).$$

Proof. Assume that $\bar{\pi}_{1|\pi_0^{\omega}} > \bar{\pi}_{1|\pi_0^{\omega'}}$ and $c_2^{\omega} \leq c_2^{\omega'}$ for some $(\omega, \omega') \in \Omega^2$. From (10) we obtain $u'(c_1^{\omega}) > u'(c_1^{\omega'})$ and therefore $c_1^{\omega} < c_1^{\omega'}$. This is inconsistent with the anti-comononicity of $\{c_1^{\omega}\}$ and $\{c_2^{\omega}\}$ established in Lemma 1. Thus if $\bar{\pi}_{1|\pi_0^{\omega}} > \bar{\pi}_{1|\pi_0^{\omega'}}$ one must have $c_2^{\omega} > c_2^{\omega'}$, and due to the anti-comonotonicity property mentioned above, $c_1^{\omega} < c_1^{\omega'}$. □

[5] The inequality $u(c_2^{\omega}) > 0$ simply means that the individual is better off being alive and consuming c_2^{ω} than being dead.

Proposition 4 indicates that temporally risk averse individuals who interact with a risk-neutral pension provider voluntarily accept some of the aggregate longevity risks. The following section further clarifies the findings of Proposition 4. The numerical simulation provided in Section 5 illustrates this phenomenon in a multi-period model. It provides an order of magnitude of the effects at play when assuming realistic demographic uncertainty.

4 Underlying intuitions

Proposition 4 indicates that improved prospects of future survival lead to a lower pension in period 1 and a higher one in period 2. Conversely, lower survival prospects imply a higher pension in period 1 and a lower one in period 2. This results from the combination of three non-trivial effects that we now describe. In order to provide an intuitive understanding, we first discuss what would be obtained in constrained problems where consumption in one or the other period is exogenously fixed (Sections 4.1 and 4.2). We return to the unconstrained case in Section 4.3.

4.1 Temporal risk aversion

Assume here that consumption in period 2 is exogenously fixed to some level $c_2^\star$. Then, optimization is only a matter of choosing $\{c_1^\omega\}$. If $c_2^\omega = c_2^\star$ for all $\omega \in \Omega$ and $u(c_2^\star) > 0$, then an improvement in the probability of surviving to period 2, $\bar{\pi}_{1|\pi_0^\omega}$, translates into a higher lifetime utility U_1^ω. Because temporally risk averse individuals prefer to combine "goods" with "bads", they prefer a system where c_1^ω and $\bar{\pi}_{1|\pi_0^\omega}$ move in opposite directions to a system where survival probabilities and consumption are independent. Formally, this constrained problem yields the first-order condition (5) with c_2^ω replaced by $c_2^\star$. This directly implies that if $k > 0$, then $\bar{\pi}_{1|\pi_0^\omega} > \bar{\pi}_{1|\pi_0^{\omega'}} \Rightarrow c_1^\omega < c_1^{\omega'}$ for any $(\omega, \omega') \in \Omega^2$. In words, π_0^ω and c_1^ω are strictly anti-comonotone.

4.2 Preference for early resolution of uncertainty

We now turn to the case where consumption in period 1 equals some exogenous level $c_1^\star$. Then, optimization is only a matter of choosing $\{c_2^\omega\}$. Despite an apparent similarity with the constrained problem discussed above, choosing $\{c_2^\omega\}$ while consumption in period 1 is fixed is a rather different issue. Indeed, consumption in period 2 only provides utility when the individual is alive in the second period. Yet, once it is known that she is alive in period 2, the information about $\bar{\pi}_{1|\pi_0^\omega}$ is irrelevant. In fact, if $\bar{\pi}_{1|\pi_0^\omega}$ were only revealed at the end of period 1 (e.g. because of delays in building mortality statistics), we would be in the case discussed in Proposition 1, with c_2^ω being independent of $\bar{\pi}_{1|\pi_0^\omega}$.

In our setting, one learns π_0^ω at the end of the period 0. For individuals with a strict preference for early resolution of uncertainty (when $k > 0$ and $\beta < 1$), the timing of the information release matters. In particular, the earlier they obtain information, the higher their tolerance for risks on lifetime utility. In the case at hand, this greater risk tolerance leads to allocating more resources (i.e. higher c_2^ω) to good states (i.e. those with a high $\bar{\pi}_{1|\pi_0^\omega}$) when $\bar{\pi}_{1|\pi_0^\omega}$ is revealed at the end of period 0, compared to when $\bar{\pi}_{1|\pi_0^\omega}$ is revealed at the end of period 1. Recall, however, that in the case where uncertainty would resolve at the end of period 1, c_2^ω would be independent of ω. Thus, when uncertainty resolves at the end of period 0, the willingness to put more resources on good states eventually leads to choosing large values c_2^ω when $\bar{\pi}_{1|\pi_0^\omega}$ is high and low values c_2^ω when $\bar{\pi}_{1|\pi_0^\omega}$ is low.

Formal mathematical resolution confirms this intuition. This constrained problem yields the first-order condition (6), with c_1^ω replaced by $c_1^\star$. This directly implies the strict comonotonicity of c_2^ω and $\bar{\pi}_{1|\pi_0^\omega}$ when $k > 0$ and $\beta < 1$. Moreover, note that, if $k = 0$ or $\beta = 1$ in this constrained problem, then equation (6) indicates that c_2^ω should be independent of $\bar{\pi}_{1|\pi_0^\omega}$, thus deterministic. This is consistent with the fact that $k = 0$ or $\beta = 1$ implies indifference to the timing of the resolution of uncertainty, which cancels the effect we just discussed.

4.3 Impatience

The two constrained problems above help explain why we expect c_1^ω and $\bar{\pi}_{1|\pi_0^\omega}$ to be anti-comonotone as a consequence of temporal risk aversion, and c_2^ω and $\bar{\pi}_{1|\pi_0^\omega}$ to be comonotone as a consequence of a preference for an early resolution of uncertainty. In fact, these two effects are amplified by a third one that only emerges when $\{c_1^\omega\}$ and $\{c_2^\omega\}$ are chosen jointly. Indeed, temporal risk aversion amplifies the effect of lifetime uncertainty on impatience, as discussed further in Bommier et al. (2017). Consequently, when $k > 0$, individuals with low mortality prospects choose less steeply decreasing (or more steeply increasing) consumption paths than individuals with high mortality prospects. Therefore, this impatience effect leads to a lower first-period consumption c_1^ω and a higher second-period consumption c_2^ω when $\bar{\pi}_{1|\pi_0^\omega}$ is high than when it is low. Impatience thus reinforces the two effects discussed above. This holds true even in the presence of a perfect insurance market. In particular, note that Proposition 4 covers the case where $\beta = 1$. Although there is no preference for early resolution of uncertainty when $\beta = 1$, so that the effect discussed in Section 4.2 vanishes, we find that c_2^ω and $\bar{\pi}_{1|\pi_0^\omega}$ remain comonotone. This is due to impatience.

Table 1 below summarizes the joint movements of consumption in periods 1 and 2, c_1^ω and c_2^ω , when the outcome at the end of period 0 predicts an increase in survival in the following period.

Effects	Temporal risk aversion	Preference for the timing of information release	Impatience	Combined effects
c_1^ω	$\searrow$		$\searrow$	$\searrow$
c_2^ω		$\nearrow$	$\nearrow$	$\nearrow$
Holds when	$k > 0$	$k > 0$ and $\beta < 1$	$k > 0$	$k > 0$ and $\beta \leq 1$
Discussed in	Section 4.1	Section 4.2	Section 4.3	Proposition 4

Table 1: Impact of positive news regarding $\bar{\pi}_{1|\pi_0^\omega}$.

5 A calibrated multi-period model

In this section, we assess the order of magnitude of the variations in consumption discussed above. We calibrate and numerically solve a multi-period life-cycle model with a realistic dynamics of survival rates and temporal risk aversion. It covers the active life and retirement. Indeed, to the extent that the aggregate mortality rates experienced during one's active life inform about one's future life expectancy, contributions as well as benefits will vary across states of the world at the optimum.

We describe the optimization program in the multi-period setup in Section 5.1. We explain how to solve and calibrate it in Sections 5.2 and 5.3. We discuss our numerical results in Section 5.4. We proceed to a sensitivity analysis in Section 5.5.

5.1 The life-cycle setup

Individuals are active from period 0 until retirement in period T_{ret}. They can live up to period T at most, and may die in any period between 0 and T. Let $\mathcal{T} = \{0, \ldots, T\}$. For each ω in a set Ω of possible states of the world, p_ω denotes the probability that the world will be in that state. To any state of the world $\omega \in \Omega$ corresponds a vector $\pi^\omega = \{\pi_t^\omega\}_{t\in\mathcal{T}} \in [0,1]^T$ of survival probabilities and an adapted consumption process $c^\omega = \{c_t^\omega\}_{t\in\mathcal{T}}$. For any state $\omega \in \Omega$, π_t^ω denotes the probability of surviving to the end of period t, conditionally on being alive at the end of period $t-1$. For any $t \in \mathcal{T}$ and any $\omega \in \Omega$, we define $p_{\omega|\pi_0^\omega,\ldots,\pi_t^\omega} \in [0,1]$ by

$$p_{\omega|\pi_0^\omega,\ldots,\pi_{t-1}^\omega} = \frac{p_\omega}{\sum_{\{\omega'|\pi_0^{\omega'},\ldots,\pi_{t-1}^{\omega'}=\pi_0^\omega,\ldots,\pi_{t-1}^\omega\}} p_{\omega'}}.$$

This is the probability to be in state ω, conditionally on observing survival rates equal to $\pi_0^\omega, \ldots, \pi_{t-1}^\omega$ up to the end of period t. Let us recall that information revealed at the end of a given period cannot be used to determine the consumption level during that period. Hence, for any two states $(\omega, \omega') \in \Omega^2$ such that $\pi_0^\omega, \ldots, \pi_{t-1}^\omega = \pi_0^{\omega'}, \ldots, \pi_{t-1}^{\omega'}$, that is states which cannot be distinguished before the end of period t, one must have $c_1^\omega, \ldots, c_t^\omega = c_1^{\omega'}, \ldots, c_t^{\omega'}$. Denote by U_t^ω the lifetime utility at time t conditionally on being alive at time t in state ω and normalize the utility of being dead to 0. A state-contingent survival profile $\{\pi_t^\omega\}_{t\in\mathcal{T}}$ and an adapted consumption process $\{c_t^\omega\}_{t\in\mathcal{T}}$ provide the utility given by the end point U_0 of the following backward induction:

$$U_T^\omega = u\left(c_T^\omega\right), \tag{11}$$

and for all $t \leq T-1$,

$$U_t^\omega = u\left(c_t^\omega\right) - \frac{\beta}{k}\log\left(\sum_{\{\omega'|\pi_0^{\omega'},\ldots,\pi_{t-1}^{\omega'}=\pi_0^\omega,\ldots,\pi_{t-1}^\omega\}} p_{\omega'|\pi_0^{\omega'},\ldots,\pi_{t-1}^{\omega'}}\left[\pi_t^{\omega'} e^{-kU_{t+1}^{\omega'}} + \left(1-\pi_t^{\omega'}\right)\right]\right). \tag{12}$$

Pensions are assumed to be the sole source of income during retirement. Individuals earn an exogenous income $\{I_t\}_{0\leq t<T_{ret}}$ when active, and rely on their pension after

retirement. We assume the pension system to be actuarially fair so that:

$$\sum_{\omega\in\Omega}\sum_{t=0}^{T} p_\omega \delta^t \left(\prod_{\tau=0}^{t-1}\pi_\tau^\omega\right) c_t^\omega \leq \sum_{\omega\in\Omega}\sum_{t=0}^{T_{ret}-1} p_\omega \delta^t \left(\prod_{\tau=0}^{t-1}\pi_\tau^\omega\right) I_t = B. \tag{13}$$

Contributions in the active life are simply the difference between income and consumption, $I_t - c_t^\omega$. Pension benefits, when retired, equal c_t^ω.

5.2 Solving the extended model

The optimization problem is the following:

$$\begin{aligned} &\max_{\{c_t^\omega\}\in\mathcal{C}} U_0 \\ \text{s.t.} \quad &\textstyle\sum_{\omega\in\Omega}\sum_{t\in\mathcal{T}} p_\omega \left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)\delta^t c_t^\omega \leq B \end{aligned} \tag{14}$$

where by convention $\Pi_{\tau=0}^{-1} = 1$.

5.2.1 First-order conditions

Risk-sensitive preferences being recursive, we can solve the optimization problem using Bellman's principle of dynamic optimality. Formally, for any $t \in \mathcal{T}$ and $\omega \in \Omega$, the objective function in period t and state ω rewrites as

$$\max_{\{c_\tau^{\omega'}\}_{\tau\geq t,\omega'|\pi_0^{\omega'},\ldots,\pi_{t-1}^{\omega'}=\pi_0^\omega,\ldots,\pi_{t-1}^\omega}} U_t^\omega$$

$$\text{s.t.} \quad \sum_{\left\{\omega'|\pi_0^{\omega'},\ldots,\pi_{t-1}^{\omega'}=\pi_0^\omega,\ldots,\pi_{t-1}^\omega\right\}} \sum_{i=t}^{T} p_{\omega'|\pi_0^\omega,\ldots,\pi_{t-1}^\omega} \left(\Pi_{\tau=t}^{i-1}\pi_\tau^{\omega'}\right)\delta^{i-t} c_i^{\omega'} \leq B_t^\omega$$

where B_t^ω depends on past consumption up to date $t-1$. Maximizing U_t^ω amounts to maximizing $-\frac{1}{k}e^{-kU_t^\omega}$. The first-order conditions with respect to c_t^ω and c_{t+1}^ω write

$$e^{-kU_t^\omega} u'\left(c_t^\omega\right) = \lambda_t^\omega, \tag{15}$$

$$\beta e^{-\frac{k}{\beta}u(c_t^\omega)} e^{-kU_{t+1}^\omega} u'\left(c_{t+1}^\omega\right) e^{\frac{k(1-\beta)}{\beta}U_t^\omega} = \lambda_t^\omega \delta, \tag{16}$$

where $\{\lambda_t^\omega\}$ is an adapted process of positive Lagrange multipliers. From (15) and (16) one derives:

$$\frac{u'\left(c_t^\omega\right)}{u'\left(c_{t+1}^\omega\right)} = \beta\delta^{-1} e^{\frac{k}{\beta}\left(U_t^\omega - u\left(c_t^\omega\right)\right)} e^{-kU_{t+1}^\omega}. \tag{17}$$

This Euler equation is independent of past consumption.

5.2.2 Implementation

When $k \neq 0$, the discount factor $\beta\delta^{-1} e^{\frac{k}{\beta}\left(U_t^\omega - u\left(c_t^\omega\right)\right)} e^{-kU_{t+1}^\omega}$ in the right-hand side of equation (17) depends on future consumption choices through the utility levels U_t^ω and U_{t+1}^ω. This makes it difficult to find an analytical solution. Yet, it will become apparent later on, when discussing value of life matters in Section 5.3, that this endogenous discount factor varies relatively little with consumption levels. This provides an efficient iterative numerical solution method. The strategy involves selecting a first guess $\left\{c_{t,(0)}^\omega\right\}$, from which we compute the corresponding instantaneous utilities $\left\{u\left(c_{t,(0)}^\omega\right)\right\}$, the lifetime utilities $\left\{U_{t,(0)}^\omega\right\}$ and the discount factors $RD_{t+1,(0)}^\omega = \beta\delta^{-1} e^{\frac{k}{\beta}\left(U_{t,(0)}^\omega - u\left(c_{t,(0)}^\omega\right)\right)} e^{-kU_{t+1,(0)}^\omega}$. One can then easily compute an adapted consumption profile $\left\{c_{t,(1)}^\omega\right\}$ that fulfills both

$$\frac{u'\left(c_{t,(1)}^\omega\right)}{u'\left(c_{t+1,(1)}^\omega\right)} = RD_{t+1,(0)}^\omega \tag{18}$$

and the budget constraint (13). The profile $\left\{c_{t,(1)}^\omega\right\}$ does not solve (17) because the discount rate $RD_{t+1,(0)}^\omega$ computed from $\left\{c_{t,(0)}^\omega\right\}$ is different from the discount rate $RD_{t+1,(1)}^\omega$ implied by $\left\{c_{t,(1)}^\omega\right\}$. However, from $RD_{t+1,(1)}^\omega$ one can derive $\left\{c_{t,(2)}^\omega\right\}$ that fulfills an equation similar to (18) and the budget constraint (13), and then iterate the process to find $\left\{c_{t,(3)}^\omega\right\}$, etc. The exact solution to (14) is provided by $\lim_{n\to+\infty}\left\{c_{t,(n)}^\omega\right\}$ since this limit fulfills both the budget constraint and the first order condition (17). With a standard computer, convergence occurs in a fraction of a second.[6]

For the deterministic pension system, we show in Appendix B that the first order conditions is:

$$\frac{u'\left(c_t\right)}{u'\left(c_{t+1}\right)} = \beta\delta^{-1} e^{-\frac{k}{\beta}u(c_t)} \frac{K_{t+1}}{K_t}$$

[6] Matlab code is available upon request.

where for all $t \in \mathcal{T}$,

$$K_t = \frac{\sum_{\omega \in \Omega} p_\omega \left(\Pi_{\tau=0}^{t-1} \pi_\tau^\omega\right) e^{k \frac{1-\beta}{\beta} \sum_{\tau=1}^{t-1} U_\tau^\omega} e^{-kU_t^\omega}}{\sum_{\omega \in \Omega} p_\omega \left(\Pi_{\tau=0}^{t-1} \pi_\tau^\omega\right)}.$$

Here again the discount rate is endogenous and we lack an analytical solution. The same iterative procedure can be implemented to derive the optimal consumption path.

5.3 Calibration

First, we specify the instantaneous utility function u. We assume the intertemporal elasticity of substitution to be constant: $-\frac{u'(c)}{cu''(c)} = \sigma$, or equivalently:

$$u(c) = 1 + \lambda \frac{c^{1-\frac{1}{\sigma}} - 1}{1 - \frac{1}{\sigma}}$$

where λ is a constant. We normalize u to equal zero in case of death. The difference in instantaneous utility between being alive but consuming 1 unit of consumption, and being dead thus equals one. The parameter λ drives the level of the marginal utility of consumption when alive. The smaller λ, the less variations in consumption affect instantaneous utility compared to how death affects it. Arthur (1981) suggests that for individuals living in wealthy societies, mortality differentials have a greater impact on welfare than variations in consumption. The relatively high empirical estimates for the value of life confirm that this is indeed the case. Consequently, λ is typically found to be small, as in our calibration below. This explain why utility levels, and hence the stochastic discount factor in equation (17), vary little with consumption, making the method discussed in Section 5.2.2 particularly efficient.

5.3.1 Exogenous calibration

First, we define period 0 as corresponding to age 25. Retirement occurs at age 65 and the maximal life duration is artificially set at 100 years. Up to age 100, we model survival rates using the standard method of Lee and Carter (1992). A single index of the "mortality level", the state-contingent quantity $\{\kappa_t^\omega\}_{t \in \mathcal{T}}$, drives the dynamics of survival rates for all ages. It follows a random walk with drift. For any $t \in \mathcal{T}$ and

$\omega \in \Omega$,

$$\kappa_t^\omega = \kappa_{t-1}^\omega + m + \nu_t^\omega.$$

Because mortality rates historically tended to decline, $m < 0$. In the standard model of Lee and Carter (1992), the $\{\nu_t^\omega\}_{t \in \mathcal{T}}$ are independent and identically distributed white noises with volatility ν. The survival rate revealed in the end of period t in state ω for an individual aged x writes $\pi_t^\omega(x) = 1 - e^{\alpha_x + \beta_x \kappa_t^\omega}$, where $\beta_x > 0$. Due to the random walk modeling, a favorable evolution of survival probabilities in period t (i.e. a high κ_t^ω), has a positive impact on all future survival probabilities. An adverse evolution of survival probabilities has the opposite effect. Survival probabilities are thus positively correlated with (remaining) life expectancy in case of survival.

We estimate m and ν using historical mortality data for the United States population from 1945 to 2013 and for ages 25 to 100.[7] To better capture the mortality improvements that characterized the second half of the twentieth century, we don't include data points prior to 1945.

To reduce the time of our computations, we consider a five-year time step. Hence, $T_{ret} = 9$ and $T = 16$. Moreover, we approximate $\tilde{\kappa}$ by a simple random walk taking values $\{\nu, -\nu\}$ with probabilities $\left\{\frac{1}{2}, \frac{1}{2}\right\}$. This generates $2^{15} = 32,768$ different survival paths.

Figure 2 displays life expectancies and survival rates in the best-case and worst-case scenarios. The uncertainty about life expectancy decreases over time, a well-documented horizon effect.

[7] This data is available on the Human Mortality Database of the University of California, Berkeley. We fit the model using the singular value decomposition method to retrieve a least-squares solution for κ, as described in Lee and Carter (1992). The parameters m and ν are obtained by a least-squares procedure.

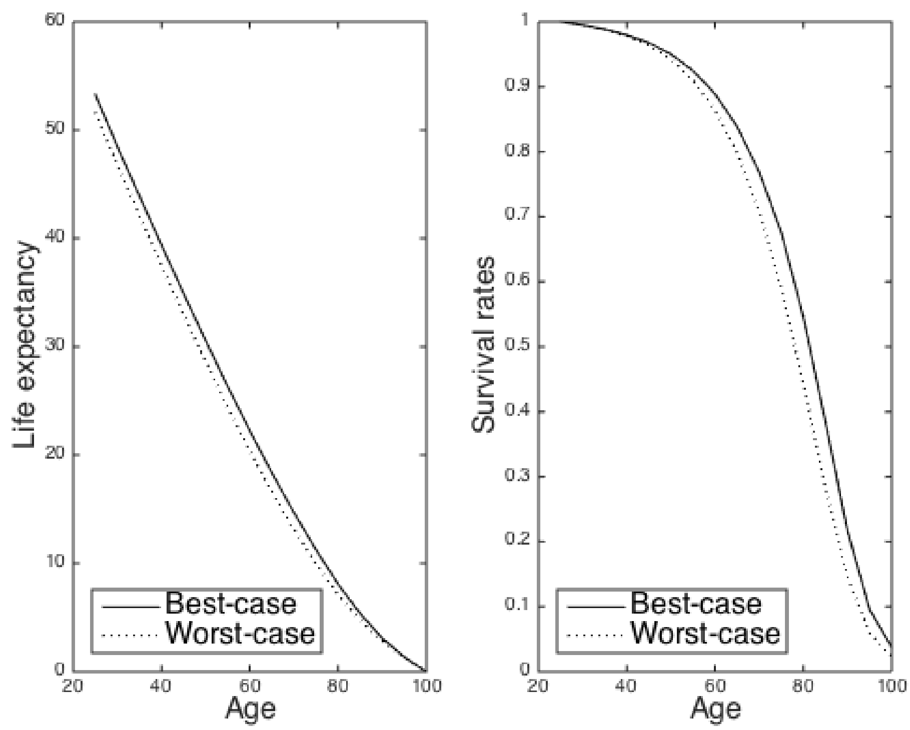

Figure 2: Life expectancy (left) and survival rates (right) as a function of age.

We assume that income increases over time, with a 50 percent increase between age 25 and age 65. We normalize it so that consumption in the last period of 5 years equals 1 unit. In our simulation, the average retirement income is 1.12 units per period of 5 years. To link it to the real world, one can think of one consumption unit being equal to 115,000 USD -for a period of five years- so that our average retirement income would then match the median retirement income measured in the U.S. in 2010 for individuals older than 65 (Trenkemp 2012 reports a level of about 25,800 USD per year). The non-discounted expected lifetime income, which equals 13.17 consumption units in our analysis, would correspond to slightly more than 1.5 million USD. This value is consistent with the findings of Tamborini et al. (2015). It is close to the lifetime earnings of a male with some college education.

We compute the 5-year discount factor $(1+r)^{-5}$, assuming a constant yearly interest rate $r = 1\%$. This matches the historical average of the risk-less short-term interest rate approximated by the 3-month T-bond, between 2007 and 2017.

We set two preference parameters exogenously. First, we assume that there is no pure preference for the present: $\beta = 1$. This case corresponds to the multiplicative model analyzed in depth in Bommier (2013). When $\beta = 1$ impatience arises from the combination of risk on the length of life and temporal risk aversion. As emphasized in Bommier (2013), such a model is actually better able to match empirical consumption profiles than the usual additive model with $\beta < 1$. In addition, $\beta = 1$ entails indifference to the timing of resolution of uncertainty. The assumption of preference for early resolution of uncertainty is standard in applied works using recursive preferences. Epstein et al. (2014), however, argue that, rather than an indisputable trait of "rational" preferences, this assumption might be viewed as a technical cost to pay for disentangling risk and time preferences in an infinite-time setting (where convergence requires that $\beta < 1$).[8] Consequently, the assumption of indifference to the timing of information may seem more desirable than costly. Second, we set $\sigma = 0.5$, a value of the elasticity of intertemporal substitution considered plausible in the literature (see Trabandt and Uhlig 2011 or Havranek 2015).

5.3.2 Endogenous calibration

We calibrate k and λ so as to replicate two targets. The first target is the value of a statistical life at age 65, VSL_{65}, defined as the marginal rate of substitution between statistical survival and consumption. The second target is the rate of time discounting at age 65, ρ_{65}, defined as the rate of change of marginal utility when controlling for the variations of consumption. It determines the slope of the consumption path at age 65.

In the present paper, consumption profiles are not deterministic. Therefore, we look at average values:

$$\rho_{65} = \sum_{\omega \in \Omega} p_\omega \rho_{65}(\omega) \quad \text{and} \quad VSL_{65} = \sum_{\omega \in \Omega} p_\omega VSL_{65}(\omega)$$

where, for all $\omega \in \Omega$,

$$1 + \rho_{65}(\omega) = \left. \frac{\frac{\partial \widetilde{U}_{65}^\omega}{\partial c_{65}^\omega}}{\frac{\partial \widetilde{U}_{65}^\omega}{\partial c_{70}^\omega}} \right|_{c_{65}^\omega = c_{70}^\omega} \quad \text{and} \quad VSL_{65}(\omega) = \frac{\frac{\partial \widetilde{U}_{65}^\omega}{\partial \pi_{65}^\omega}}{\frac{\partial \widetilde{U}_{65}^\omega}{\partial c_{65}^\omega}}.$$

[8] The timing premia implied by recursive utility models in the case of long-run risks may even be implausibly high (Epstein et al. 2014).

The $\left\{\widetilde{U}_{65}^{\omega}\right\}$ are computed assuming that $\pi_{65}^{\omega'} = \pi_{65}^{\omega}$ for all states ω' such that $\pi_{25}^{\omega'}, \ldots, \pi_{60}^{\omega'} = \pi_{25}^{\omega}, \ldots, \pi_{60}^{\omega}$ (states that are undistinguishable before age 65). That is, $\widetilde{U}_{65}^{\omega} = u(c) - \frac{1}{k}\log\left(\pi_{65}^{\omega}e^{-kU_{70}^{\omega}} + 1 - \pi_{65}^{\omega}\right)$ where $c = c_{65}^{\omega}$ for the value of a statistical life target and $c = c_{70}^{\omega}$ for the rate of time discounting since we control for variations in consumption at age 65. In both cases, the value of U_{70}^{ω} is computed without any adjustments, as in equation (12). Then,

$$1 + \rho_{65}(\omega) = \frac{e^{kU_{70}^{\omega}}e^{-k\left(\widetilde{U}_{65}^{\omega} - u\left(c_{70}^{\omega}\right)\right)}}{\pi_{65}^{\omega}} \text{ and } VSL_{65}(\omega) = \frac{\left(1 - e^{-kU_{70}^{\omega}}\right)e^{k\left(\widetilde{U}_{65}^{\omega} - u\left(c_{65}^{\omega}\right)\right)}}{ku'\left(c_{65}^{\omega}\right)}.$$

Our targets correspond to a United-States-based individual aged 65. Then, the corresponding target for the rate of time discounting equals 3.5 percent per year. It implies that, on average, consumption decreases by 0.2 percent per year at age 65. This is consistent with the findings of Fernández-Villaverde and Krueger (2007). Viscusi and Aldy (2007) suggest a target value of a statistical life equal to 500 times the average consumption between ages 65 and 100.

Table 2 sums up the values of the parameters.

Exogenous parameters	β	1
	σ	0.5
	r	1.00%
	T	100
Calibrated parameters	k	0.179
	λ	1.03×10^{-2}

Table 2: Parameter calibration.

Because k quantifies risk aversion "per util", it should be considered together with the level of instantaneous utility. Set c_m equal to the average consumption level in our model. Then $ku(c_m) = 0.177$ is the per period coefficient of risk aversion with respect to the length of life (for a 5-year period) for an individual whose consumption profile is flat with $c_t = c_m$ in all periods. This amounts to a coefficient of risk aversion with respect to the length of life of 3.54×10^{-2} per year. To provide an order of magnitude, an individual aged 65 with a life expectancy of 16.4 years and endowed with such preferences would undergo a surgical procedure with a five percent chance of death during surgery only if it provides more than one year and two months of additional life expectancy (accounting for the risk of dying during surgery). It is difficult to obtain precise empirical estimates of individuals' risk aversion with respect to the length of life. In the lab, this involves presenting subjects with the unusual exercise of choosing between hypothetical lotteries on the length of life, controlling for variations in consumption and pure time preferences. Yet, in one such experiment, Leroux et al. (2016) provide evidence of deviations from the standard assumption of risk neutrality with respect to the length of life.

5.4 Results

In this section, we discuss the consumption profiles delivered by the adaptive and deterministic systems and highlight an interesting consequence of letting contributions and benefits vary with life expectancy.

5.4.1 Longevity-dependent consumption profile

In Figure 3, we draw the consumption profile in the deterministic system (dotted line) as well as in the adaptive system for the best and worst survival scenarii (solid and dashed lines). In the best-case survival scenario, individuals only receive "good" news about survival over time. In the worst-case, they only receive "bad" news. All other scenarii comprise a succession of "good" and "bad" news. The optimal consumption profiles are realistically humped-shaped. The hump occurs between ages 55 and 60. This is consistent with the empirical findings of Fernández-Villaverde and Krueger (2007).

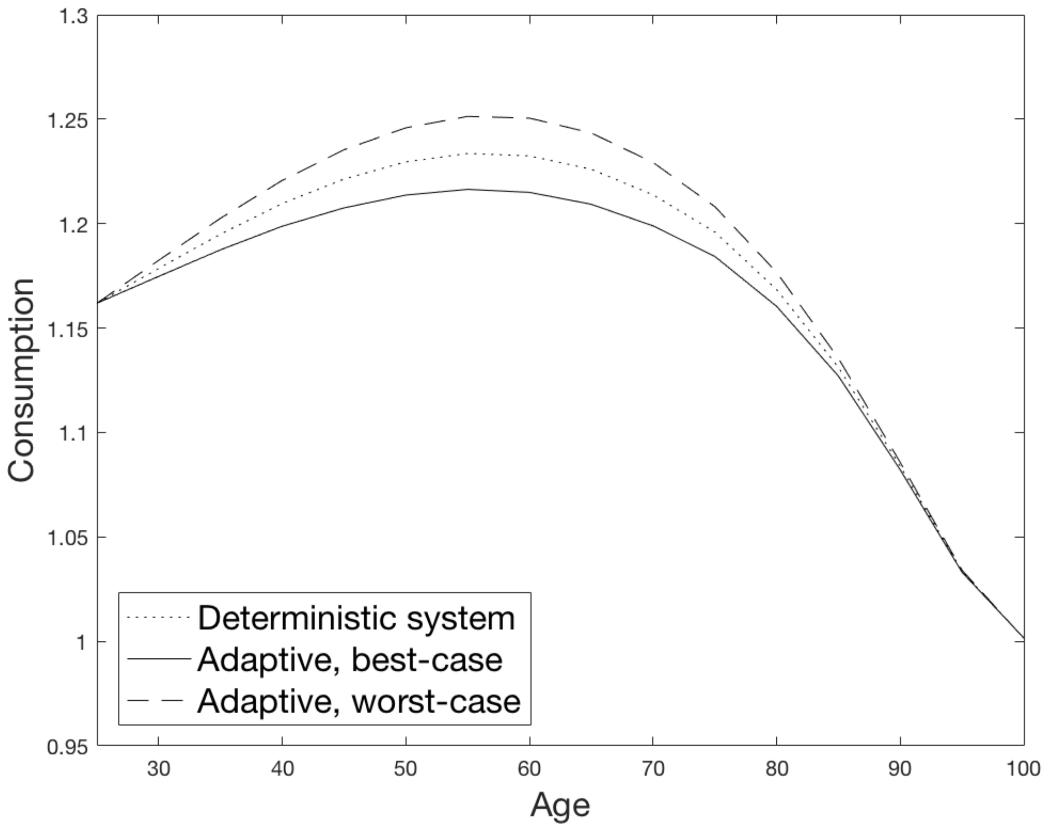

Figure 3: Consumption paths.

The dynamics of mortality rates allow for some learning over time. Lower mortality rates in a given period increase the chances of survival for all subsequent periods, and the other way around. Consistently with Proposition 4, the optimal state-dependent consumption profile is adaptive. Consumption and survival are anti-comonotone. Decreases in life expectancy entail a decrease in contributions and an increase in benefits. Increases in life expectancy have the opposite effect. Because we have assumed indifference to the timing of uncertainty resolution, we do not observe the effect discussed in subsection 4.2. Finally, the impatience effect translates into a steeper decline in consumption during retirement when life expectancy is lower.

The deterministic consumption profile lies in the interval drawn by the two extreme scenarios of only "good" or "bad" survival shocks. Consumption in the adaptive system varies only slightly compared to the deterministic one. Variations range from -1.22 to +1.25 percent of the deterministic consumption values. Therefore, the living standard of surviving individuals is only moderately at risk.

5.4.2 Welfare gains

Lifetime utility is higher in the adaptive system. The welfare gain, however, is almost negligible. It is comparable to an increase of 0.01% of income. Indeed, our calibration entails a relatively low degree of temporal risk aversion. Let us emphasize that one's own mortality risk (i.e. whether one lives or dies) has much larger consequences on welfare than the aggregate longevity risk. Combining this very salient individual mortality risk with a high degree of temporal risk aversion would lead to implausible consumption profiles. Pension systems that better cope with the aggregate longevity risk do make individuals better off. Yet, from an individual's point of view, the risk reduction is relatively small given that the idiosyncratic component is of much greater significance. As we will see below, however, the risk reduction is sizable for pensions providers.

5.4.3 Risk allocation

Pension providers are interested in the distribution of the net costs of a pension system. In a given state of the world ω, the net lifetime payments write as

$$\sum_{t=0}^{T} p_\omega \delta^t \left(\prod_{\tau=0}^{t-1} \pi_\tau^\omega \right) c_t^\omega - \sum_{t=0}^{T_{ret}-1} p_\omega \delta^t \left(\prod_{\tau=0}^{t-1} \pi_\tau^\omega \right) I_t.$$

With the Lee and Carter modeling, the variations in aggregate survival rates are small before retirement. Therefore, the impact of longevity improvements bears mostly on benefits, which increase with the number of survivors. In both systems, an unexpected increase in survival translates into net costs for pension providers. In an adaptive system, however, contributions and survival probabilities are comonotone, while benefits and survival probabilities are anti-comonotone. This shifts part of the financial cost of longer lives, and part of the financial benefits of shorter ones, to individuals. In words, compared to the deterministic system, the adaptive pension system offers a Pareto-improving risk-sharing mechanism for the aggregate longevity risk, in which individuals bear more longevity risk and pension providers less.

The distributions of the net lifetime payments in the deterministic pension system (dotted line) and the adaptive one (solid line) appear in Figure 4 as a percentage of the expected lifetime income B.

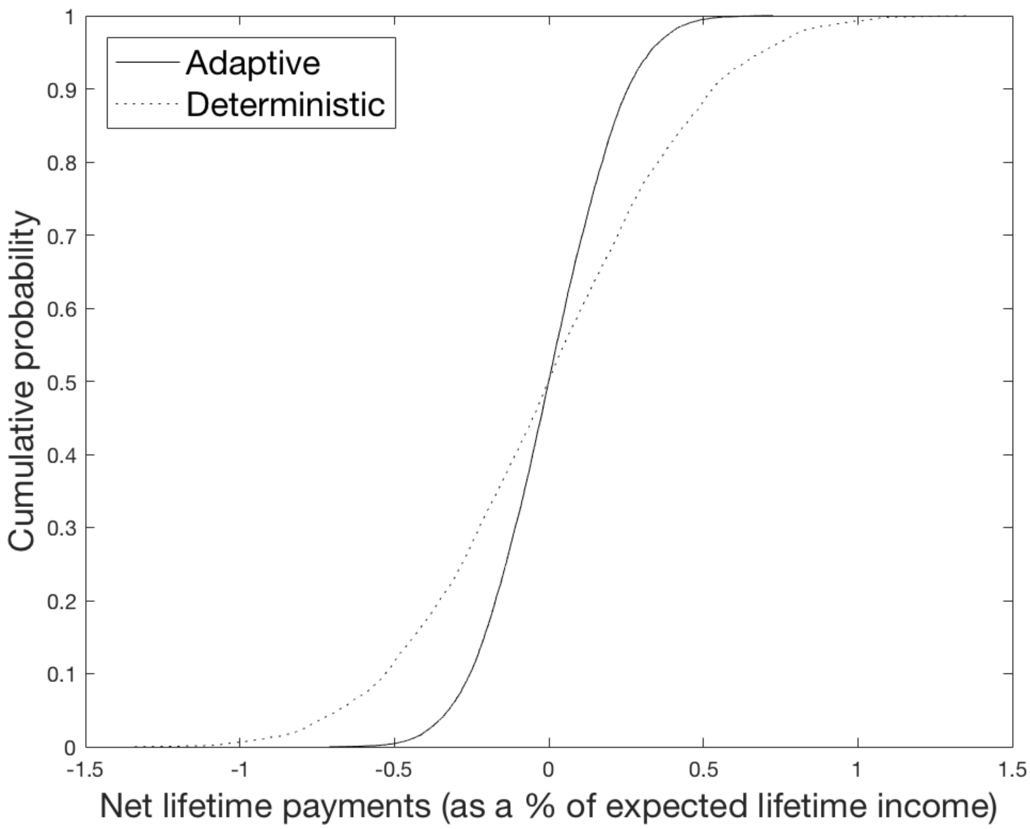

Figure 4: Cumulative density functions of net lifetime payments as a percentage of the expected lifetime income.

By construction, both systems are actuarially fair. Thus the expected net lifetime payments are identical and equal to zero. The volatility of the net lifetime payments is more than twice smaller with the adaptive system than with the deterministic one. Indeed, under temporal risk aversion, the adaptive system achieves a lower correlation between length of life and lifetime consumption. Individuals then bear more than half of the aggregate longevity risk. The first two moments appear in Table 3.

	Deterministic	Adaptive
Mean	0	0
Standard deviation	4.12×10^{-2}	1.99×10^{-2}

Table 3: Mean and standard deviation of the distribution of net lifetime payments.

5.5 Sensitivity analysis

We now study how sensitive the fraction of the aggregate longevity risk reallocated to individuals is to our calibration targets. We re-calibrate and re-estimate the model for a range of target values above and below the benchmarks. For the value of a statistical life at age 65, we consider calibration targets that range from 400 to 600 times the average consumption. As for the rate of time discounting at age 65, we vary the calibration target from 3% to 4%. Table 4 summarizes our findings. It provides the value of the calibrated parameters k and λ as well as the fraction of the aggregate longevity risk reallocated to individuals.

Sensitivity analysis

		$VSL = 400c_m$	$VSL = 500c_m$	$VSL = 600c_m$
ρ =3%	Parameters	$k = 0.124$ $\lambda = 1.20 \times 10^{-2}$	$k = 0.123$ $\lambda = 0.96 \times 10^{-2}$	$k = 0.123$ $\lambda = 0.80 \times 10^{-2}$
	Risk reduction	44.72%	44.74%	44.77%
ρ =3.5%	Parameters	$k = 0.179$ $\lambda = 1.29 \times 10^{-2}$	$k = 0.179$ $\lambda = 1.03 \times 10^{-2}$	$k = 0.179$ $\lambda = 0.87 \times 10^{-2}$
	Risk reduction	51.76%	51.74%	51.71%
ρ =4%	Parameters	$k = 0.230$ $\lambda = 1.37 \times 10^{-2}$	$k = 0.230$ $\lambda = 1.10 \times 10^{-2}$	$k = 0.230$ $\lambda = 0.90 \times 10^{-2}$
	Risk reduction	36.19%	36.05%	35.76%

Table 4: The risk reduction quantifies the reduction in the standard deviation of net lifetime payments obtained when switching from a deterministic system to an adaptive one.

Changes in the value of a statistical life: Fitting different values for the VSL while holding the rate of time discounting fixed involves changing the parameter λ while maintaining the other parameters (almost) unchanged. The effect on the optimal consumption profile, and hence the risk reduction, is negligible. In fact, when mortality is exogenous, there exists an asymptotic limit in the case where the VSL goes to infinity (λ converging to 0 in our specification). Empirical estimates imply large values of the VSL with predictions close to those obtained in the asymptotic limit.

Changes in the rate of time discounting: Time discounting results from the combination of risk aversion and uncertainty on the length of life. Thus, varying the time discounting calibration target involves varying the risk aversion parameter k, with greater discount rates obtained for greater k. Since increasing risk aversion also increases the value of life, keeping the same VSL target for different values of k requires adjusting the parameter λ (which negatively impacts the VSL) in a comonotonic way. More risk averse individuals prefer lower correlations between the risks on length of life and on lifetime consumption. Because survival is exogenous, this lower correlation is achieved by increasing the amplitude of the consumption fluctuations across states of the world. This increases the proportion of the aggregate longevity borne by individuals. The relationship between the target value for the rate of time discounting and the share of aggregate longevity risk, however, is non-monotonic. This is because the income profile is unchanged in all simulations, while the consumption profile, and thus the amount of pensions to be paid, varies in shape.

The share of the aggregate longevity risk that individuals retain in an adaptive system is clearly significant, even when using different calibration targets than ours. The risk allocation appears to be sensitive to the rate of time discounting, which is a major determinant of the consumption profile.

6 Conclusion

The literature on the allocation of longevity risk generally takes it for granted that risk-averse pensioners prefer pensions to be deterministic and, therefore, that the question of the optimal allocation of the longevity risk would be trivial if this risk could be insured at no cost. With the current paper, we emphasize that this conclusion is correct only if pensioners are temporally risk neutral, or if past survival rates convey no information about future survival rates. Those very restrictive assumptions are at odds with empirical facts. Experimental studies have found that individuals are not temporally risk neutral but temporally risk averse (Andersen et al. 2018) and that mortality rates exhibit a strong serial correlation.

Accounting for pensioners' temporal risk aversion and for the serial correlation of survival rates, we find that pensioners should bear some of the aggregate longevity at the optimum, even when this risk could be insured at no cost. According to our numerical simulation, pensioners would be willing to take about 50% of the overall risk (quantified by the variance of aggregate lifetime net payments) in order to implement a sort of diversification strategy where bad news in terms of survival are combined with good news in terms of living standards, and *vice versa*. Thus, only the remaining 50 % of the longevity risk would require insurance.

For pension providers, increases in life expectancy may entail enormous financial losses. We are well aware that insuring the remaining 50 % of the longevity risk at actuarially fair prices is unrealistic. Our results should thus be eventually combined with those of previous studies discussing the intergenerational risk-sharing of the longevity risk or its transfer to financial markets through securitization. The novelty is that the reference point should not be a deterministic pension system but an adaptive one, where pensions are periodically adjusted to account for the evolution of life expectancy. This presupposes that governments would be capable of implementing such adaptive pension systems.

Interestingly, many countries now link their pension system's parameters to longevity improvements. France and Canada increase contribution rates while Sweden or Italy revalue benefits according to the most recent demographic information. These adjustments, however, only bear either on contributors or on new cohorts of retirees, with no adjustment taking place during the retirement period. Our contribution suggests that it would be optimal to include retirees in the scope of longevity-linked adjustments. It is worth emphasizing that the adjustments are quite limited in their magnitude. We find that, in the most extreme scenario of sustained longevity improvements, benefits would decrease by less than 1.25 percent. Such small adjustments would nevertheless significantly reduce the risk borne by pension providers, without lowering pensioner's welfare.

References

[1] Abdellaoui, M., H. Bleichrodt, O. L'Haridon, and C. Paraschiv (2013): "Is There One Unifying Concept of Utility? An Experimental Compar-

ison of Utility Under Risk and Utility Over Time," *Management Science*, 59, 2153–2169.

[2] Aldy, J. E. and W. K. Viscusi (2007): "Age Differences in the Value of Statistical Life: Revealed Preference Evidence," *Review of Environmental Economics and Policy*, 1, 241–260.

[3] Andersen, S., G. W. Harrison, M. Lau, and E. E. Rutström (2018): "Multiattribute Utility Theory, Intertemporal Utility and Correlation Aversion," *International Economic Review*, forthcoming.

[4] Arthur, W. B. (1981): "The Economics of Risks to Life," *The American Economic Review*, 71, 54–64.

[5] Ball, L. and N. G. Mankiw (2007): "Intergenerational Risk Sharing in the Spirit of Arrow, Debreu, and Rawls, with Applications to Social Security Design," *Journal of Political Economy*, 115, 523–547.

[6] Bohn, H. (2001): "Social Security and Demographic Uncertainty, The Risk-Sharing Properties of Alternative Policies," in *Risk Aspects of Investment-Based Social Security Reform*, University of Chicago Press, 203–246.

[7] Bommier, A. (2013): "Life-Cycle Preferences Revisited," *Journal of the European Economic Association*, 11, 1290–1319.

[8] ——— (2014): "Mortality Decline, Impatience and Aggregate Wealth Accumulation with Risk-Sensitive Preferences," *Economics Working Paper Series, ETH Zurich*.

[9] Bommier, A., A. Kochov, and F. Le Grand (2017): "On Monotone Recursive Preferences," *Econometrica*, 85, 1433–1466.

[10] Borch, K. (1960): "The safety loading of reinsurance premiums," *Scandinavian Actuarial Journal*, 1960, 163–184.

[11] Cheung, S. L. (2015): "Comment on "Risk Preferences Are Not Time Preferences": On the Elicitation of Time Preference under Conditions of Risk," *American Economic Review*, 105, 2242–2260.

[12] Coble, K. H. and J. L. Lusk (2010): "At the Nexus of Risk and Time Preferences: An Experimental Investigation," *Journal of Risk and Uncertainty*, 41, 67–79.

[13] Davidoff, T., J. R. Brown, and P. A. Diamond (2005): "Annuities and Individual Welfare," *American Economic Review*, 95, 1573–1590.

[14] Delprat, G., M.-L. Leroux, and P.-C. Michaud (2016): "Evidence on Individual Preferences for Longevity Risk," *Journal of Pension Economics and Finance*, 15, 160–179.

[15] Demange, G. and G. Laroque (2000): "Social Security, Optimality, and Equilibria in a Stochastic Overlapping Generations Economy," *Journal of Public Economic Theory*, 2, 1–23.

[16] Eeckhoudt, L., B. Rey, and H. Schlesinger (2007): "A Good Sign for Multivariate Risk Taking," *Management Science*, 53, 117–124.

[17] 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.

[18] Epstein, L. G. and S. M. Tanny (1980): "Increasing Generalized Correlation: A Definition and Some Economic Consequences," *The Canadian Journal of Economics*, 13, 16–34.

[19] Fernández-Villaverde, J. and D. Krueger (2007): "Consumption over the Life Cycle: Facts from Consumer Expenditure Survey Data," *Review of Economics and Statistics*, 89, 552–565.

[20] Gordon, R. H. and H. R. Varian (1988): "Intergenerational Risk Sharing," *Journal of Public Economics*, 37, 185–202.

[21] Hansen, L. P. and T. J. Sargent (1995): "Discounted Linear Exponential Quadratic Gaussian Control," *IEEE Transactions on Automatic Control*, 4, 968–971.

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

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

[24] Krueger, D. and F. Kubler (2002): "Intergenerational Risk-Sharing via Social Security When Financial Markets Are Incomplete," *The American Economic Review*, 92, 407–410.

[25] Lee, R. D. and L. R. Carter (1992): "Modeling and Forecasting US Mortality," *Journal of the American Statistical Association*, 87, 659–671.

[26] Miao, B. and S. Zhong (2015): "Comment on "Risk Preferences Are Not Time Preferences": Separating Risk and Time Preferences," *American Economic Review*, 105, 2272–2286.

[27] Milevsky, M. A. and T. S. Salisbury (2015): "Optimal Retirement Income Tontines," *Insurance: Mathematics and Economics*, 64, 91–105.

[28] Tamborini, C. R., C. H. Kim, and A. Sakamoto (2015): "Education and Lifetime Earnings in the United States," *Demography*, 52, 1383–1407.

[29] Trabandt, M. and H. Uhlig (2011): "The Laffer Curve Revisited," *Journal of Monetary Economics*, 58, 305–327.

[30] Trenkamp, B. (2012): "Income of the Aged Chartbook, 2010," *SSA publication, Social Security Administration, Office of Retirement and Disability Policy.*

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

A First-order conditions in the two-periods model

Since c_0 is exogenous, maximization of U_0 is equivalent to that of

$$W_0 = -\frac{1}{k} \sum_{\omega \in \Omega} p_\omega \left[\pi_0^\omega e^{-kU_1^\omega} + (1 - \pi_0^\omega) \right].$$

For any $(\omega,\omega')\in\Omega^2$ one has

$$\frac{\partial U_1^{\omega'}}{\partial c_1^{\omega}}=\begin{cases}u'\left(c_1^{\omega}\right) & \text{if } \omega=\omega'\\ 0 & \text{otherwise.}\end{cases}$$

Thus:

$$\frac{\partial W_0}{\partial c_1^{\omega}}=p_\omega \pi_0^\omega e^{-kU_1^\omega}u'\left(c_1^\omega\right).$$

Similarly for any $(\omega,\omega')\in\Omega^2$ one has:

$$\frac{\partial U_2^{\omega'}}{\partial c_2^{\omega}}=\begin{cases}u'\left(c_2^{\omega}\right) & \text{if } \omega=\omega'\\ 0 & \text{otherwise.}\end{cases}$$

Moreover

$$e^{-kU_1^{\omega'}}=e^{-ku\left(c_1^{\omega'}\right)}\left(\sum_{\left\{\omega|\pi_0^\omega=\pi_0^{\omega'}\right\}}p_{\omega|\pi_0^\omega}\left[\pi_1^\omega e^{-kU_2^\omega}+(1-\pi_1^\omega)\right]\right)^\beta.$$

We thus have:

$$\frac{\partial e^{-kU_1^{\omega'}}}{\partial c_2^{\omega}}=\begin{cases}-\beta\pi_1^\omega p_{\omega|\pi_0^\omega}e^{-\frac{k}{\beta}u\left(c_1^{\omega'}\right)}e^{-kU_2^\omega}u'\left(c_2^\omega\right)e^{k\frac{1-\beta}{\beta}U_1^{\omega'}} & \text{if } \pi_0^{\omega'}=\pi_0^\omega\\ 0 & \text{otherwise.}\end{cases}$$

Therefore:

$$\frac{\partial W_0}{\partial c_2^{\omega}}=\beta\pi_0^\omega\pi_1^\omega\sum_{\left\{\omega'|\pi_0^{\omega'}=\pi_0^\omega\right\}}p_{\omega'}p_{\omega|\pi_0^\omega}e^{-\frac{k}{\beta}u\left(c_1^{\omega'}\right)}e^{-ku\left(c_2^\omega\right)}u'\left(c_2^\omega\right)e^{k\frac{1-\beta}{\beta}U_1^{\omega'}}.$$

Now recall that $\pi_0^\omega=\pi_0^{\omega'}$ implies that $c_1^\omega=c_1^{\omega'}$ and, from equation (3), we see that $U_1^\omega=U_1^{\omega'}$. Moreover $\sum_{\left\{\omega'|\pi_0^{\omega'}=\pi_0^\omega\right\}}p_{\omega'}p_{\omega|\pi_0^\omega}=p_\omega$. Thus:

$$\frac{\partial W_0}{\partial c_2^{\omega}}=\beta\pi_0^\omega\pi_1^\omega p_\omega e^{-\frac{k}{\beta}u\left(c_1^{\omega}\right)}e^{-ku\left(c_2^\omega\right)}u'\left(c_2^\omega\right)e^{k\frac{1-\beta}{\beta}U_1^{\omega}}.$$

The first-order conditions are therefore:

$$e^{-kU_1^\omega}u'\left(c_1^\omega\right)=\mu$$
$$\beta e^{-\frac{k}{\beta}u\left(c_1^\omega\right)}e^{-ku\left(c_2^\omega\right)}u'\left(c_2^\omega\right)e^{k\frac{1-\beta}{\beta}U_1^\omega}=\mu\delta.$$

B First-order conditions in the deterministic case

Since c_0 is exogenous, maximization of U_0 is equivalent to that of

$$W_0\left(c_0,\left\{c_1^\omega\right\},\ldots,\left\{c_T^\omega\right\}\right)=-\frac{1}{k}\sum_{\omega\in\Omega}p_\omega\left[\pi_0^\omega e^{-kU_1^\omega}+\left(1-\pi_0^\omega\right)\right].$$

In the deterministic case, for any $\omega\in\Omega$, $c_t^\omega=c_t$. One may rewrite the objective function as follows:

$$\widetilde{W}_0\left(c_0,c_1,\ldots,c_T\right)=W_0\left(c_0,\left\{c_1^\omega\right\},\ldots,\left\{c_T^\omega\right\}\right).$$

Thus,

$$\frac{\partial\widetilde{W}_0}{\partial c_t}=\sum_{\omega\in\Omega}\frac{\partial W_0}{\partial c_t^\omega}$$

where

$$\frac{\partial W_0}{\partial c_t^\omega}=u'\left(c_t\right)\beta^{t-1}e^{-\frac{k}{\beta}\sum_{\tau=1}^{t-1}u(c_\tau)}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)e^{k\frac{1-\beta}{\beta}\sum_{\tau=1}^{t-1}U_\tau^\omega}e^{-kU_t^\omega}.$$

Then,

$$\frac{\partial\widetilde{W}_0}{\partial c_t}=u'\left(c_t\right)\beta^{t-1}e^{-\frac{k}{\beta}\sum_{\tau=1}^{t-1}u(c_\tau)}\sum_{\omega\in\Omega}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)e^{k\frac{1-\beta}{\beta}\sum_{\tau=1}^{t-1}U_\tau^\omega}e^{-kU_t^\omega}.$$

Let λ be a positive Lagrange parameter. For any $t\in\mathcal{T}$, the first-order condition with respect to c_t writes

$$u'\left(c_t\right)\beta^{t-1}e^{-\frac{k}{\beta}\sum_{\tau=1}^{t-1}u(c_\tau)}\sum_{\omega\in\Omega}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)e^{k\frac{1-\beta}{\beta}\sum_{\tau=1}^{t-1}U_\tau^\omega}e^{-kU_t^\omega}=\lambda\delta^{t-1}\sum_{\omega\in\Omega}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right).$$

Let

$$K_t=\frac{\sum_{\omega\in\Omega}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)e^{k\frac{1-\beta}{\beta}\sum_{\tau=1}^{t-1}U_\tau^\omega}e^{-kU_t^\omega}}{\sum_{\omega\in\Omega}p_\omega\left(\Pi_{\tau=0}^{t-1}\pi_\tau^\omega\right)}.$$

The Euler equation follows:

$$\frac{u'(c_t)}{u'(c_{t+1})} = \beta\delta^{-1}e^{-\frac{k}{\beta}u(c_t)}\frac{K_{t+1}}{K_t}.$$