SECURITIZING AND TRANCHING LONGEVITY EXPOSURES

ENRICO BIFFIS* AND DAVID BLAKE$^\circ$

Abstract. We consider the problem of optimally designing longevity risk transfers under asymmetric information. We focus on holders of longevity exposures that have superior knowledge of the underlying demographic risks, but are willing to take them off their balance sheets because of capital requirements. In equilibrium, they transfer longevity risk to uninformed agents at a cost, where the cost is represented by retention of part of the exposure and/or by a risk premium. We use a signalling model to quantify the effects of asymmetric information and emphasize how they compound with parameter uncertainty. We show how the cost of private information can be minimized by suitably tranching securitized cashflows, or, equivalently, by securitizing the exposure in exchange for an option on mortality rates. We also investigate the benefits of pooling several longevity exposures and the impact on tranching levels.

Keywords: longevity risk, asymmetric information, security design, pooling, tranching.

1. Introduction

Over the last decade, longevity risk has become increasingly capital intensive for pension funds and annuity providers to manage. The reason is that mortality improvements have been systematically underestimated, making balance sheets vulnerable to unexpected increases in liabilities. Blake and Burrows (2001) were the first to advocate the use of mortality-linked securities to transfer longevity risk to capital markets. Their proposal

---

Date: August 15, 2008. This version: July 10, 2009. This paper has been written for research purposes only. No party reading this paper is entitled to rely on its content in any way. If any party wishes to rely on advice in relation to the contents of this paper, they must consult directly their own advisors.
Previous versions were presented at Bachelier Finance Society 2008 World Congress (London), 4$^{\text{th}}$ International Longevity Risk and Capital Markets Solutions Conference (Amsterdam), 5$^{\text{th}}$ Conference in Actuarial Science and Finance (Samos), 2008 Workshop on Insurance-linked Securities (Imperial College's Centre for Hedge Fund Research), 2009 Risk and Stochastics Day (LSE), and at the University of Modena and Reggio Emilia. We are grateful to Anna Rita Bacinello, Tom Boardman, Gilles Chemla, Guy Coughlan, Xuewen Liu, Pietro Millossovich, Ragnar Norberg, Michael Sherris, as well as to an anonymous referee and to participants in conferences and seminars for helpful comments and suggestions. We are solely responsible for any errors.
*Biffis (E.Biffis@imperial.ac.uk) is at Imperial College Business School, Imperial College London, South Kensington Campus, SW7 2AZ United Kingdom. $^\circ$Blake (D.Blake@city.ac.uk) is at The Pensions Institute, Cass Business School, 106 Bunhill Row, EC1Y 8TZ United Kingdom.

has generated considerable attention in the last few years, and major investment banks and reinsurers are now actively innovating in this field (see Blake *et al.*, 2006a, 2008, for an overview).

Nevertheless, despite growing enthusiasm, longevity risk transfers have been materializing only slowly. One of the reasons is the huge imbalance in scale between existing exposures and willing hedge suppliers. The bulk of longevity exposures is represented by liabilities of defined benefit pension funds and annuity providers. In 2007, these institutions' exposure to improvements in life expectancy amounted to a staggering 400 billion USD in the UK and the US alone (see Loeys *et al.*, 2007). In the UK, a pension plan buyout market started in 2006 to transfer corporate pension plan assets and liabilities to insurance companies. This market has grown very rapidly with estimates of 11 billion USD of transfers during 2008. However, there are already signs of capacity constraints - some of the insurers have been unable to attract additional shareholders' funds to expand their businesses - and the buyout market transfers all risks (including interest rate and inflation risks) and not just longevity risk. It is therefore important to consider other vehicles for transferring longevity risk. One such vehicle is securitization, the packaging of illiquid assets and liabilities into securities that are sold into the capital markets. Five types of securitization have taken place involving longevity-related assets or liabilities: blocks of business, regulatory reserving (triple-X securitization), life settlements, annuity books and reverse mortgages. The new securities created are known as insurance-linked securities (ILSs) (for more details, see Krutov, 2006; Blake *et al.*, 2008).

Considerable progress has been made in understanding mortality dynamics (e.g., Cairns *et al.*, 2007, 2008; Dowd *et al.*, 2008a,b; Gourieroux and Monfort, 2008) and in designing mortality-linked securities (e.g., Dowd *et al.*, 2005; Blake *et al.*, 2006a,b; Dawson *et al.*, 2007; Denuit, 2007; Loeys *et al.*, 2007). Yet, the optimal pricing of the embedded risks remains elusive. The pricing exercises used so far by practitioners are typically based on partial equilibrium arguments and shed little light on how supply and demand might equilibrate when longevity exposures are exchanged. We address this problem by adopting a simple general equilibrium model that shows how longevity risk premia are determined by the joint effect of regulatory costs and uncertainty in longevity trends. The literature on optimal security design is vast (see Duffie and Rahi, 1995, for an overview pertaining to financial innovation) and spans research on corporate finance (see Tirole, 2006), deductibles in insurance policies (see: Eeckhoudt *et al.*, 2005; Dana and Scarsini, 2007) and minimization of risk measures (see Barrieu and El Karoui, 2005). We abstract from heterogeneity in risk preferences and spanning issues and focus on the role played by asymmetric information on longevity trends in optimal security design. We consider holders of longevity exposures that have superior knowledge of the underlying demographic risks, in the sense that they have access to better experience data or forecasting technologies developed by monitoring the exposures. This assumption is realistic for life insurers, reinsurers and other intermediaries (e.g., pension buyout firms and investment banks) that have developed considerable expertise in managing longevity-linked cashflows. Moreover, investors currently still seem to be uncomfortable enough with longevity risk to make this assumption a reasonable one, even for securities written on publicly available demographic indices. The incentive to enter a transaction is given by an exogenously specified retention cost resulting from capital requirements or alternative investment opportunities. Knowledge of this cost is available to all agents, as it can be quantified from international regulatory rules and accounting standards.

Our starting point is the securitization of a book of annuity-like cashflows and their backing assets. The presence of asymmetric information means that the holder or originator of the longevity exposure (also denoted as issuer or seller below) faces a 'lemons' problem (as in Akerlof, 1970), in the sense that investors' demand for the new security may be downward sloping and expose the issuer to a liquidity problem. As is common in annuity reinsurance and the securitization of insurance assets and liabilities (e.g., Cowley and Cummins, 2005), retention of part of the exposure can be used to 'prove' the quality of the cashflows to the market and alleviate the impact of asymmetric information. We use a signalling model of Walrasian market equilibrium, as in Gale (1992) and DeMarzo and Duffie (1999), to determine optimal retention levels and securitization payoffs. The resulting separating equilibrium allows us to determine the issuer's retention costs and to examine how risk premia would emerge if there were departures from the optimal securitization level. As a particular example, we consider the situation where the responsibility for meeting payments linked to the realized death rate in a reference population is transferred to a counterparty, exactly as happens with mortality swaps and forwards (see: Coughlan *et al.*, 2007; Loeys *et al.*, 2007).

We then allow the holder of the book of liabilities and backing assets to issue a security that is contingent on the net exposure to longevity (i.e., the surplus on the assets in excess of the longevity exposure) and examine conditions under which the optimal contract results in tranching of the net exposure. By tranching, we mean slicing the net exposure so that, in exchange for a lump sum paid to the originator, investors who buy the tranche are entitled to a specific portion of the net exposure's cashflows. The optimal tranching level minimizes the sensitivity of these cashflows to both asymmetric information and the impact of unsystematic risk, which is material to risk-neutral agents when payoffs are nonlinear. Since we wish to minimize the cost to the originator of issuing the tranche, the optimal tranche is the one that is least risky from the investors' viewpoint. In other words, the optimal tranche is equivalent to the senior debt tranche in a debt financing operation. Although we focus on the design of a single tranche, our analysis could be extended to multiple tranches, as in Plantin (2004) and DeMarzo (2005), for example. It could also be extended beyond the securitization of a specific book of business to encompass securities written on publicly available mortality indices (such as the LifeMetrics indices[1]). For example, we can obtain insights into how the strike levels of options on mortality indices could be chosen by originators willing to repackage and sell their longevity exposures.

We also address the issue of securitizing pools of exposures. As in DeMarzo (2005), diversification benefits can be traded off against the detrimental effect of information loss from pooling together low-longevity- and high-longevity-risk cashflows. We obtain the result that pooling and then tranching can reduce the negative impact of unsystematic risk at high ages and in small portfolios. Also, the benefits from pooling and tranching are magnified when private information is highly correlated across exposures, while residual risk is not. This occurs, for example, when issuers of securities pool different cohorts of individuals belonging to the same geographic area or social class, or pool several small portfolios with comparable demographic characteristics. When considering securities written on publicly available demographic indices, the model shows that 'age-bucketing' (i.e.,

[1]See `http://www.lifemetrics.com`.

writing derivatives on the mortality experience of an entire age range in a given population) can reduce asymmetric information costs, in addition to mitigating basis risk (see Coughlan *et al.*, 2007).

We note that although we concentrate on longevity risk, our analysis carries over to cashflows exposed to the risk of systematic mortality increases (i.e., brevity risk). See Cowley and Cummins (2005) and Bauer and Kramer (2007) for an overview of recent transactions involving mortality catastrophe bonds.

The paper is organized as follows. In Section 2, we provide a stylized description of longevity exposures, identifying a systematic and an unsystematic longevity risk component. In Section 3, we introduce asymmetry in the information available to agents concerning the systematic risk component. In Section 4, we provide an equilibrium model of the securitization of longevity-linked cashflows. In Section 5, we study the optimal design of a derivative written on a longevity exposure. In Section 6, we extend our analysis to pools of longevity exposures. In Section 7, we give a simple example of how the setup can be used with stochastic intensities of mortality. Finally, Section 8 offers some concluding remarks.

2. Longevity exposures

Denote the current time by 0 and let $\tau_1, \ldots, \tau_m$ be the residual lifetimes of the individuals in a given population at time $z \geq 0$, such as insureds in a book of annuity policies or members of a pension plan. For fixed time horizon $T > z$, denote by S the proportion of individuals surviving up to time T and by D the proportion of individuals dying by time T, i.e.

$$S = \frac{1}{m} \sum_{j=1}^{m} 1_{\tau_j > T-z} = 1 - D, \tag{2.1}$$

where 1_A denotes the indicator function of the event A. Expression (2.1) is a simple representation of the longevity exposures faced by annuity providers and pension funds. In a book of annuities, for example, the annuity payments over $[z, T]$ can be assumed to be proportional to the random variable S, while the reserves released over the same period can be assumed to be proportional to D. In practice, reversionary or indexed benefits might be provided, but we abstract from additional sources of uncertainty such as financial risk factors and minimum guarantees.

A longevity exposure is affected by two main sources of risk: the risk of unsystematic fluctuations around expected levels of mortality, and the risk of systematic departures from expected levels. To fix ideas, consider first the random variable D and the case of i.i.d. death times. Consider, for example, individuals of the same gender, the same age, and with similar health characteristics at time z. Then $\tau_1, \ldots, \tau_m$ may be assumed to be independent and to obey the same law. If the population is large enough, the Law of Large Numbers yields

$$\lim_{m \to \infty} D(m) = E[D] = q, \tag{2.2}$$

almost surely, where the notation $D(m)$ emphasizes the dependence of D on the population size, and q denotes the probability of dying by time T (we use the notation $p := 1 - q$ for the corresponding survival probability). Expression (2.2) lies at the heart of any life insurance business, based on pooling together a large number of risks to make outflows more predictable. From (2.2), we could use the representation

$$D = q + \eta, \tag{2.3}$$

with η a zero mean error term capturing unsystematic fluctuations around the estimate q.

When dealing with longevity exposures, however, the most challenging uncertainty is the one surrounding the trend in longevity improvements, i.e., systematic longevity risk. One way of formalizing this situation is to rewrite expression (2.3) as

$$D = q(Y) + \varepsilon, \tag{2.4}$$

with ε another error term. The estimate q in (2.2) is replaced by a random variable dependent on a vector $Y = (Y_1, \ldots, Y_l)'$ of relevant risk factors. The law for $\tau^1, \ldots, \tau^m$ depends now on the common factor Y influencing systematic longevity risk, and the predictability of the exposure is considerably reduced. For instance, when the individuals' death times are conditionally independent, given Y, the limit in (2.2) holds only conditionally on each outcome y of Y (e.g., Schervish, 1995, Lemma 1.61), reducing the scope for risk reduction by pooling together a large number of individuals.

3. Asymmetric information

We consider two types of agent interested in exposure (2.1): the holder of the exposure (e.g., an annuity provider, a pension buyout firm or an investment bank) and a multitude of agents that might seek an investment exposure to longevity risk (e.g., other institutional investors, endowments or hedge funds). We build on representation (2.4), and focus on exposures with the following structure:

$$D = q(Y) + \varepsilon. \tag{3.1}$$

Much of our analysis actually applies to the alternative representation

$$D = q(Y)\exp(\epsilon). \tag{3.2}$$

Equivalent expressions can be formulated for S, with the trend component denoted by $p(Y)$.

We assume that the holder has access to the entire information carried by Y, while investors have access to a limited number of signals on the trend component of (3.1) or (3.2). We can think of the holder possessing valuable experience data or having developed a superior forecasting technology by monitoring the book of liabilities over time. We denote the private components of Y by $X = (X_1, \ldots, X_d)$, for some $d \leq l$, and we will look at some specific examples.

A typical situation is when the private signal allows the holder to represent $q(Y)$ in terms of a mortality estimate relative to a different group of individuals (e.g., aggregate population or industry-wide experience), denoted by $\widetilde{q}$, which is publicly available. We can then set $Y = (\widetilde{q}, X)$ and write

$$q(Y) = \widetilde{q} + a(X), \tag{3.3}$$

for some adjustment factor $a(X)$ ensuring that $q(Y)$ takes values in $[0, 1]$. In other words, the private information allows the holder to express D in terms of a reference exposure $\widetilde{D}$ publicly monitored. Similarly, we could set

$$q(Y) = \widetilde{q}\; m(X), \tag{3.4}$$

with $m(X)$ a privately computable scaling factor.

Both examples (3.3) and (3.4) cover common ways of producing mortality estimates based on adjusting reference forecasts provided by industry bodies and government statistics offices. Approach (3.4) is common in the life settlements market,[2] where mortality estimates are made by adjusting a baseline mortality table with suitable impairment factors (see Modu, 2008).

In the following, we will be interested in the role played by the private signal X, rather than in the optimal structure of Y to produce estimates for D or S. We therefore focus on the special case in which the signal is entirely private, i.e. $X = Y$, and assume that $q(X)$ is an unbiased conditional estimate of D, given X, i.e.

$$q(X) = E\left[D|X\right].$$

This means that $q(X)$ is all that matters for the private valuation of exposure D.

So far, we have made no restrictions on the distribution of X. We now require the following:

Assumption 3.1. *The distribution of X has compact support and the conditional distribution of D, given X, is continuous.*

Under the above assumptions, it is easy to show that the support of the distribution of $q(X)$ is a compact interval, which we denote by $[q^{\min}, q^{\max}] \subseteq [0, 1]$. For example, we set $q^{\min} := \min_x E\left[D|X = x\right]$. The corresponding support for the distribution of $p(X)$ is denoted by $[p^{\min}, p^{\max}]$, with $p^{\min} = 1 - q^{\max}$ and $p^{\max} = 1 - q^{\min}$. The width of these intervals gives a simple characterization of the systematic uncertainty associated with the exposure. Since we focus on longevity risk, the bounds $q^{\min}$ and $p^{\max}$ will play a crucial role in determining how the effects of asymmetric information influence the design of longevity-linked securities.

As a practical example of the risks we are analyzing, Figure 1 plots one-year death rate estimates for ages $65, 70$ and 75 from year 2008 to 2030, based on the projections of the Continuous Mortality Investigation (CMI).[3] These projections provide a standard

[2]An index based on the mortality experience in this market was developed by Goldman Sachs in 2007: see `http://www.qxx-index.com`.
[3]Projections based on CMI improvement factors applied to the PMA92 mortality table: see `http://www.actuaries.org.uk/knowledge/cmi`.

benchmark for annuity and pension liabilities in the UK. Each death probability has three levels corresponding to so-called short-cohort, medium-cohort and long-cohort projections. The terms 'short', 'medium' and 'long' refer to the degree of persistence in mortality improvements assumed in the forecasting model.

Figure 1 gives an idea of the magnitude of trend uncertainty, particularly at higher ages. The differences between the long-cohort and short-cohort estimates for the three ages in 2013 are 0.013%, 0.083% and 0.343%. Bearing in mind that the historical annualized volatilities for death rates at ages 65, 70 and 75 are[4] 0.95%, 1.22% and 1.59%, we see that we need to deal with the uncertainty associated with both longevity trends and unsystematic longevity risk.

< Figure 1 about here. >

4. Securitizing longevity exposures

We now consider the problem of transferring longevity exposures to the capital markets. We assume that there is a single riskless asset yielding an interest rate normalized to zero and that there is a market populated by a large number of risk-neutral investors that have no access to the information carried by X. The holder of the longevity exposure is also risk-neutral, but discounts future cashflows at a positive rate, because of solvency requirements or to reflect the opportunity cost of alternative investment opportunities. This provides an incentive to securitize the exposure in exchange for cash.

Suppose that an infinitely divisible riskless asset $\alpha \geq 0$ backs a promised payment that depends on the proportion of survivors in a given population at some future date T. The final cashflow is given by the difference between the value of the asset and the value of the exposure (which we denote below as the net exposure),

$$\alpha - S = (\alpha - 1) + D. \tag{4.1}$$

Denoting by $\delta \in (0, 1)$ the issuer's discount factor over $[0, T]$, her valuation of cashflow (4.1) yields $\delta E[\alpha - S]$, which is lower than the value $E[\alpha - S]$ she would place on (4.1) in the absence of holding costs. On the other hand, the access to extra information allows

[4]Estimates for England and Wales, based on annual death probabilities over the period 1841-2002 from the Human Mortality Database: see `http://www.mortality.org`.

the issuer to formulate a private valuation of the net exposure equal to

$$\delta E\left[\alpha - S|X\right] = \delta(\alpha - p(X)) = \delta(\alpha - 1 + q(X)). \tag{4.2}$$

The incentive to securitize the net exposure by issuing a security with a terminal payoff proportional to $\alpha - S$ can then be quantified by the privately assessed holding cost $(1 - \delta)(\alpha - p(X))$. We assume $\alpha \geq p^{\max}$, so that we think of (4.2) as an expected surplus (the extension to the case of a deficit or a pure liability transfer is discussed after Proposition 4.2 below).

A securitization transaction involves the following steps: (i) once the private signal X is observed, the issuer computes (4.2); (ii) anticipating an arbitrary market demand P for the security, the issuer chooses a fraction $\gamma \in [0, 1]$ of the net exposure to supply to the market; (iii) at a later time T, the realized cashflows from (4.1) are revealed to both the issuer and investors.

By an arbitrary market demand, we mean a bounded measurable function P from $[0, 1]$ to $\mathbb{R}_+$. For a fixed securitization fraction $\gamma \in [0, 1]$, $P(\gamma)$ represents the market price of $\alpha - S$ perceived by investors when the issuer retains $(1 - \gamma)$ of the net exposure. For given outcome x of X and demand schedule P, the issuer chooses the optimal fraction to securitize by maximizing her private valuation of the final cashflows originating from the transaction. These are given by the discounted value of the net exposure plus any securitization payoff:

$$E\big[\delta(1-\gamma)(\alpha - S) + \gamma P(\gamma)|X = x\big] = \delta(\alpha - p(x)) + \gamma\left[P(\gamma) - \delta(\alpha - p(x))\right]. \tag{4.3}$$

The optimal fraction can then be found by maximizing the securitization payoff,

$$\Pi(e) = \sup_{\gamma \in [0,1]} \gamma\left[P(\gamma) - \delta e\right], \tag{4.4}$$

with $e = \alpha - p(x)$, so that the ex-ante expected payoff resulting from the securitization transaction is given by $E\left[\Pi(\alpha - p(X))\right]$.

The domain of the payoff, Π, is represented by the outcomes of the issuer's private valuation, e. We would like to express the optimal securitization fraction supplied to the market as a function of the private valuation, denoted by $\Gamma : [\alpha - p^{\max}, \alpha - p^{\min}] \to [0, 1]$.

The following standard proposition characterizes both Π and Γ for an arbitrary demand schedule P:

Proposition 4.1. *For any demand function P, assume that the issuer determines the optimal fraction of the net exposure to securitize by solving problem* (4.4). *If* $\Gamma(\alpha - p(x))$ *solves* (4.4) *uniquely for each outcome x of X, then*

(i) *the issuer's payoff,* Π, *is nonincreasing and convex in the private valuation of the net exposure;*

(ii) *the securitization fraction,* Γ, *is nonincreasing in the private valuation of the net exposure.*

The greater the severity of longevity risk indicated by the issuer's private valuation, the greater the fraction of the longevity exposure the issuer will wish to securitize. This is illustrated in Figure 2, where the slope of Π is increasing in the private valuation of the net exposure, $\alpha - p(X)$. From (4.4), we see that, if Π is differentiable, its differential is $-\delta\Gamma$, so that the securitization fraction must rise when there is a lower private valuation of the net exposure to ensure that Π is convex. The intuition is that a lower private valuation of cashflow (4.1) makes retention relatively less valuable and hence securitization relatively more valuable.

< FIGURE 2 ABOUT HERE. >

Although P has been taken to be arbitrarily given, one would expect the demand schedule to be downward sloping, because investors cannot observe X and face a 'lemons' problem. If investors rationally anticipate that the amount of longevity exposure put up for sale is increasing in the private valuation of longevity risk, their demand for the security would decrease in response to an increase in the securitization fraction. This is illustrated in Figure 3.

< FIGURE 3 ABOUT HERE. >

While the above considerations are intuitive, Proposition 4.1 does not provide insights into how P and Γ might arise endogenously in equilibrium. To explore this issue, we use a signalling model of Walrasian market equilibrium, as in Gale (1992) and DeMarzo and

Duffie (1999). Because the retention costs faced by the issuer are assumed to be publicly known, retention can be used as a credible signal to 'prove' the quality of cashflow (4.1) to the market. Conditional on X, the issuer computes $p(X)$ and puts up for sale a fraction $\Gamma(\alpha - p(X))$ of the net exposure. The uninformed agents then bid for the security using Γ to infer the level of longevity risk associated with the exposure. In a rational expectations equilibrium, their demand function satisfies

$$P(\Gamma(\alpha - p(X))) = E\left[\alpha - p(X) | \Gamma(\alpha - p(X))\right], \tag{4.5}$$

and the signalling game results in a unique equilibrium (P^*, Γ^*) characterized in the following proposition.

Proposition 4.2. *Consider the net exposure* (4.1)*, with* $\alpha > p^{\max}$*. The outcome of the signalling game is a unique separating equilibrium* (P^*, Γ^*)*, satisfying*

$$P^*(\Gamma^*(\alpha - p(X))) = \alpha - p(X), \tag{4.6}$$

and admitting the explicit representation

$$\begin{aligned} \Gamma^*(\alpha - p(X)) &= \left(\frac{\alpha - p(X)}{\alpha - p^{\max}}\right)^{\frac{1}{\delta - 1}} \\ P^*(\gamma) &= \gamma^{\delta - 1}\left(\alpha - p^{\max}\right), \qquad 0 \leq \gamma \leq 1. \end{aligned} \tag{4.7}$$

The equilibrium securitization payoff to the issuer is then given by

$$\Pi(\alpha - p(X)) = (1 - \delta)\ \left(\alpha - p^{\max}\right)\ \Gamma^*(\alpha - p(X))^{\delta}. \tag{4.8}$$

By letting $p^{\max}$ increase towards α, we can cover the case in which the private valuation implies that the issuer expects the asset α in (4.1) to be fully utilized in meeting the longevity exposure. The limiting optimal securitization fraction is 1 for $p(x) = \alpha$ and 0 for any $p(x) \in [p^{\min}, \alpha)$.

The above proposition provides a useful benchmark for understanding the effects of asymmetric information on the demand for and liquidity of longevity exposures. From (4.7), we see that a fundamental driver of investors' demand is the worst-case private valuation $\alpha - p^{\max}$, irrespective of the distribution of X. If $p(X) = p^{\max}$, it is optimal for the issuer to securitize the entire net exposure in order to obtain

$$\Pi(\alpha - p^{\max}) = (1-\delta)(\alpha - p^{\max}), \tag{4.9}$$

since $\Gamma^*(\alpha - p^{\max}) = 1$ in this case. Whenever the realized private valuation $p(x)$ is below the worst-case valuation, the issuer retains a fraction of the exposure to signal her private assessment of longevity risk, and the payoff from securitization is lower than the full holding cost which equals (4.9). The issuer could always obtain $(1-\delta)E[\alpha - S]$ in the case of symmetric information, but in the presence of asymmetric information, the issuer will have to incur a securitization cost that in expectation is at least as high as

$$(1-\delta)E[(\alpha - S) - (\alpha - p^{\max})] = (1-\delta)\left(p^{\max} - E[p(X)]\right). \tag{4.10}$$

Another way of viewing asymmetric information costs, is to examine off-equilibrium paths.[5] As a result of the existence of a separating equilibrium, we know that the optimal transaction always involves a partial exchange of the exposure at a fair price, irrespective of the realized private valuation $p(x)$. When $p(x) < p^{\max}$, the holder of the exposure incurs the cost associated with retaining a fraction $(1-\gamma^*)$ of the exposure. If she were to transfer a fraction γ higher than γ^*, investors would infer that the severity of longevity risk is higher, no matter what the issuer's private valuation is. As a result, they would demand a premium to take on the exposure given by $P^*(\gamma^*) - P^*(\gamma)$.

We note that in Γ^* the following ratio plays a fundamental role:

$$\frac{\alpha - p(X)}{\alpha - p^{\max}}. \tag{4.11}$$

Expression (4.11) provides a simple measure of the degree of uncertainty associated with the trend component in (3.1). When the fact that X is private is taken into account, the ratio also represents the sensitivity of the exposure to private information (see DeMarzo and Duffie, 1999, Proposition 4), since it compares the random outcome of the private valuation to the worst possible valuation. Hence, Proposition 4.2 shows how trend uncertainty combines with asymmetric information in determining either the equilibrium retention levels or the risk premia demanded by investors in a longevity risk transfer.

[5] We have not discussed the issue of out-of-equilibrium beliefs, which lead to multiple equilibria in signalling models. We note that Proposition 4.2 gives the unique equilibrium satisfying standard restrictions on off-equilibrium beliefs, as discussed in DeMarzo and Duffie (1999).

The above model can be extended to the situation in which $\alpha < p^{\max}$, i.e. the worst-case private valuation results in an expected deficit. To illustrate how this can be done, consider for example the case of $\alpha = 0$ in (4.1), implying that we are transferring only the liability S to the capital markets, not the assets. Retention costs are now captured by a discount factor $\delta > 1$, i.e. the issuer discounts liabilities at a rate lower than the risk-free rate (which is zero in our setting). We then allow the demand schedule P to take negative values, to reflect the fact that investors would be willing to take over a fraction γ of the liability in exchange for receiving a premium (instead of making a payment) equal to $\gamma P(\gamma)$ from the holder of the liability. Once these changes are enforced, expressions (4.7) and (4.8) in Proposition 4.2 can be directly applied to pure longevity risk transfers. The more general case when (4.2) could result in an expected deficit or expected surplus, depending on the private valuation, can be covered by introducing a discount factor that is function of $\alpha - p(X)$.

To give an idea of the practical implications of Proposition 4.2, we consider the special case of $\alpha = 1$, so that we are actually transferring the random death rate D to the capital markets, exactly as in mortality forwards and swaps[6] recently marketed by major investment banks (e.g., the q-forward contracts considered in Coughlan *et al.*, 2007). In this case, the ratio (4.11) reduces to $q(X)/q^{\min}$.

In Figure 4, we plot Γ^* for different values of the ratio and for δ equal to 90% and 95%. As expected, the securitization fraction is lower when δ is higher, since retention involves a higher cost and hence is a more credible signal. We consider one-year death rates for ages $65, 70$ and 75 in 2013 (five years from the assumed transaction year 2008, i.e. $z = 4$ and $T = 5$ in the notation of Section 2). The values $\delta = 0.90$ and $\delta = 0.95$ correspond in this case to annualized holding costs of approximately 2% and 1%, respectively. For the three ages, we assume that X makes the worst-case private valuation $q^{\min}$ coincide with the midpoint between the medium-cohort and the long-cohort forecast. In Figure 4, we plot the ratios $q(x)/q^{\min}$ corresponding to the event $\{X = x\}$ such that the private valuation coincides with the medium-cohort forecast (the values of the ratio are 101.02% for age 65, 102.96% for age 70, and 105.26% for age 75). We see that the higher parameter uncertainty associated with higher ages increases the potential cost of asymmetric information, simply

[6]Strictly speaking, our setup applies to fully-funded or fully-collateralized transactions in these instruments.

because the support of the distribution of the private valuations widens. The effective cost, however, depends on the discount factor δ, which is likely to vary depending on the relative contribution of the different ages to the capital provisions required by exposure (4.1).

To emphasize the role of the risk-free cashflows in the longevity risk transfer (i.e., those generated by the backing assets), we examine the equilibrium retention levels for ages 65 and 70 when the backing assets are increased by 2.04%. The case $\alpha = 1.0204$ corresponds to a level at which the retention for age 70 is equal to that for age 65 when $\alpha = 1$. Figure 5 plots the results, from which it is apparent that the increase in the backing assets generates a substantial increase in the optimal securitization fraction for both ages. The reason is that the ratio (4.11) affecting liquidity decreases quite rapidly as α increases. This means that the effect of asymmetric information is diluted when longevity exposures are bundled with other risks, as happens when entire books of liabilities and their backing assets are securitized. Recent evidence from the UK confirms this: bulk annuity buyouts (i.e., transfers of books of longevity-linked pension liabilities together with the pension fund assets) have been much more significant than pure longevity risk transfers, because of the huge risk premia demanded by investors in exchange for accepting these transfers.[7]

< Figure 4 about here. >

< Figure 5 about here. >

5. Tranching longevity exposures

While Proposition 4.2 is useful for understanding the effects of asymmetric information on the liquidity of longevity exposures, it does little to help improve the current design of longevity-linked securities. We now consider the possibility of manufacturing a derivative contract written on cashflow (4.1), rather than transferring the exposure directly to capital markets as in a securitization. Under the assumption of risk-neutrality, the residual risks ε and ϵ in (3.1) and (3.2) have so far played no role in our analysis, but will become crucial

[7]The general sentiment is that longevity solutions are currently "very expensive", if not "outrageously expensive". Mark Wood, CEO of Paternoster, the first of the new pension buyout firms established in the UK in 2006, says "We quote for mortality only buy-outs [insurance against people living longer than expected], but we tend to find that when people want a quote for a mortality buy-out, they end up comparing it to a bulk buy-out [a complete buyout of all pension liabilities] and go for that instead" ('Paying for a longer life', *Financial Times*, June 1, 2008).

in the present context. Throughout this section, we focus on the case $\alpha = 1$ and securities written on D. Equivalent results can be obtained for the more general setup of Section 3.

We consider limited liability contracts with monotone payoffs. A contract design is a nondecreasing function ϕ acting on the range of D such that $\phi(d) \leq d$ for every $d \in [0,1]$. The following steps are involved: (i) before observing the outcome of X, the holder of the exposure designs the contract; (ii) after X is observed, the contract is put up for sale and passed on to investors according to the signalling game described in Section 4; (iii) at maturity T, the final payoff is revealed to both the seller and the investors. By steps (i)-(ii), investors cannot make inference on the severity of longevity risk based on the design of the contract.

Setting $C := \phi(D)$, we wish to maximize the expected payoff to the seller over all possible contracts C. As in the previous section, given an arbitrary demand P^C for contract C, the private valuation of the final cashflows involves any proceeds from selling a fraction γ of contract C, plus the expected present value of the residual exposure $D - C$ and of the cashflows from the unsold fraction of C:

$$\begin{aligned} &E\big[\gamma P^C(\gamma) + \delta(D - C) + \delta(1-\gamma)C|X\big] \\ &\quad = \delta q(X) + \gamma\left[P^C(\gamma) - \delta c(X)\right]. \end{aligned} \tag{5.1}$$

Hence, the results of the previous section apply and the equilibrium payoff Π^C takes the same form as (4.8). To simplify notation, we drop superscript C from Π and P.

Since X is not known at the time when ϕ is chosen, we need to solve

$$V(D) = \max_C E\left[\Pi(c(X))\right], \tag{5.2}$$

where $c(x) := E[C|X = x]$ represents the conditional private valuation of contract C. What does it mean to optimize over all possible contract designs? From Proposition 4.2, the equilibrium price of the contract depends on ϕ only through the worst-case private valuation $c^{\min} := \min_x c(x)$, while the fraction of the longevity exposure hedged depends on the sensitivity of ϕ to the private information, as measured by the ratio $c(X)/c^{\min}$ (since $\delta \in (0,1)$). Hence, for each fixed $c^{\min}$, problem (5.2) identifies the contract that minimizes the sensitivity ratio among all those yielding a worst-case $c^{\min}$. If all possible contract designs were to achieve their worst-case at a common $x_0 \in \mathbb{R}^l$, the solution would drastically simplify. This is shown in the following proposition.

Proposition 5.1. *Assume that $q(X)$ is continuous and that ε (or ϵ) is log-concave. Then, for every contract C, we have $c^{\min} = c(x_0)$, for some $x_0 \in \mathbb{R}^l$, and the optimal contract design is given by*

$$C^* = \min(q^*, D) = q^* - \max(0, q^* - D), \tag{5.3}$$

for suitable tranching level q^*.

At maturity, the contract pays the amount q^* to the investor when the realized death rate exceeds the tranching level. However, the amount is reduced when the realized death rate falls below the tranching level, thereby capping the exposure of the contract seller to a survival rate of $p^* = 1 - q^*$. The seller therefore has an effective hedge against longevity risk. In turn, one can think of the investor as holding a riskless bond issued by the seller and paying q^* at maturity, but also short a put option which is held by the seller and which will expire with a maturity value of $q^* - D$ if the realized death rate falls below the tranching level.

The log-concavity assumption requires that the cumulative distribution function of ε (or ϵ) is log-concave, a property satisfied in a number of interesting cases[8] (see Karlin, 1968, for an overview). Continuity of $q(X)$ and log-concavity of ε (or ϵ) ensure that we can find an outcome x_0 of X that is a uniform worst case for D, and hence for C (see the appendix for details). The optimal design problem then simplifies because we can replace problem (5.2) with the simpler problem of optimizing over all possible tranching levels. At the optimum, given by

$$\{q^*\} = \arg\max_k E\left[\Pi\left(E\left[\min(k, D)|X\right]\right)\right],$$

the marginal benefit to the seller from making the contract more attractive to the investor by adding further positive cashflows (i.e., offering a higher q^*) is at most zero. When $q^* < q^{\max}$, it is possible to tranche the exposure to create a derivative contract written on D. Otherwise, the optimal contract transfers the entire exposure to the capital markets, and the signalling game results in the kind of arrangement that is common in annuity

[8]The property is satisfied, for example, by the Uniform, Normal, Exponential, and Extreme Value distributions and is preserved under truncation and convolution.

reinsurance, namely a 'quota share' (i.e., a contract whereby the liabilities arising from a portfolio are shared in fixed proportions between the annuity provider and the reinsurer).

Another way of understanding Proposition 5.1 is to rewrite expression (5.3) as

$$C^* = D - \max(0, D - q^*) = q^* - \max(0, S - p^*).$$

The first equality shows that the seller is transferring the random death rate D to investors in exchange for receiving a call option on the realized death rate, with strike level q^*, so that the call will expire in the money if at maturity $D > q^*$. The seller recognises that the exposure could turn out to be less risky in the end than originally anticipated (for example, if the outcome of D is some $d > E[D]$) and does not want to pay investors more than q^* to hedge her longevity risk. The second equality shows that the seller promises to pay investors q^* in exchange for a call option on survivorship, with strike price $p^* := 1 - q^*$. The intuition is that p^* measures the optimal level of protection against longevity risk that the seller is willing to obtain from the market, given the costs of capital requirements and asymmetric information. Hence, Proposition 5.1 provides a framework for determining the optimal strike levels for mortality options in the presence of differential information about longevity trends by holders and investors.

We note that an optimal 'debt contract' such as (5.3) was obtained by Innes (1990) in a 'hidden effort' (or moral hazard) model under the monotone likelihood ratio property[9] and by DeMarzo and Duffie (1999) in a 'hidden information' (or adverse selection) model under the more general condition of existence of a uniform worst case. Related issues are studied in the risk sharing literature (e.g. Eeckhoudt *et al.*, 2005; Dana and Scarsini, 2007), where the tranching level could be interpreted as a deductible. While we have concentrated on identifying a single tranche, multiple tranches could be obtained in equilibrium by allowing for heterogeneous investors (e.g., Plantin, 2004), or by considering the possibility of designing multiple tranches before X is revealed (DeMarzo, 2005, explores this situation under the monotone likelihood ratio property).

In analogy with expression (4.9) above, the maximum payoff the seller can expect to obtain from selling a contract C is $(1 - \delta)c^{\min}$, i.e. the cost from holding the contract in

[9]In our setting, this means imposing the extra restriction that $q(x)$ is monotone in x. This is usually a strong restriction for multifactor stochastic mortality models.

the worst case. Since $\phi(d) \leq d$ for all $d \in [0,1]$, we have

$$c(X) = E\left[\phi(D)|X\right] \leq E\left[D|X\right] = q(X),$$

which yields $c^{\min} \leq q^{\min}$. Using the fact that Π is decreasing in the private valuation, we can write

$$\Pi(c(X)) \leq \Pi(c^{\min}) = (1-\delta)c^{\min} \leq (1-\delta)q^{\min}, \tag{5.4}$$

showing that the maximum payoff that can be achieved is lower than the one given in (4.9). The reason is that by issuing contract C, we allow the residual risk to enter into the picture, and hence an additional risk premium payable to investors materializes in the transaction. Despite agents being risk neutral, a positive risk premium arises endogenously through the contract design on account of the nonlinear payoffs of the optimal contract.

If the seller observes $c(x) > c^{\min}$, she will retain a fraction of the contract, in order to signal the lower severity of longevity risk, and hence reduce her payoff from the transaction. Still, the ex-ante expected payoff associated with C might be higher than the one delivered by securitizing the exposure D directly. Contract (5.3) therefore optimally resolves (ex-ante) the trade-off between the costs associated with residual risk, asymmetric information and capital requirements.

As in the previous section, a useful benchmark is the case in which there is no asymmetric information and the seller can expect a transaction payoff equal to $(1-\delta)E[C]$. By analogy with expression (4.10), we can use (5.4) to compute

$$(1-\delta)E\left[C - c^{\min}\right] = (1-\delta)\left(E[c(X)] - c^{\min}\right), \tag{5.5}$$

the minimum cost the seller can expect to incur when selling C if only residual risk is taken into account.

To emphasize these two dimensions affecting the contract design, we depict in Figure 6 the densities for the one-year death rate $D = q(X) + \varepsilon$ of a 75-year-old individual in year 2009 and year 2013, as seen from year 2008. Clearly, the longer the time horizon, the greater the variance of the death rate. The solid lines on the plane (x, y) represent the CMI short-cohort and long-cohort forecasts over the period $2008 - 2013$. The worst-case private valuation $q^{\min}$, in this case, is assumed to coincide with the long-cohort forecast. In the setup in the previous section, the density curves were of no interest to either side of the transaction, and had no impact on the equilibrium results. When nonlinear contracts are written instead, the situation changes dramatically and the entire distribution of D (jointly determined by $q(X)$ and ε) enters the picture. This can be seen by comparing the two lines representing the strike levels for a suboptimal contract (strike q) and for the optimal contract (strike q^*). The optimal contract maximizes the payoff to the seller and optimally trades off the level of protection against longevity risk and the cost of doing so, as can be seen from Figure 7. The suboptimal contract (strike q) is cheaper, but provides too high a cap on survival rates (too low a floor on death rates).

< FIGURE 6 ABOUT HERE. >

< FIGURE 7 ABOUT HERE. >

To assess the practical implications of Proposition 5.1, we consider the case of an exposure $D = q(X) + \varepsilon$, where the one-year death rate refers to individuals aged 70 in 2013 (5 years from the assumed transaction year 2008), with X such that $q^{\min}$ corresponds to the long-cohort CMI estimate and $q(X)$ is uniformly distributed with lower bound $q^{\min}$ and mean coinciding with the medium-cohort CMI forecast. The residual risk component is distributed as a truncated Normal with mean 0, variance σ^2, and with the truncation ensuring that D takes values in $[0, 1]$.

Figure 8 reports the expected payoff for different strike levels and different volatility parameters σ. For low values of σ (cases a, b, and c), the relative weight of residual risk with respect to the cost of asymmetric information means that it is optimal for the seller to get protection against longevity risk at an optimal level $q^* < q^{\min}$. The lower the residual risk, the more q^* approaches $q^{\min}$, since it becomes easier to design a contract with the conditional private valuation close to the level yielding the maximum payoff computed in (5.4). In the case of no residual risk ($\varepsilon = 0$), the strike level coincides with $q^{\min}$. For high enough residual risk, the contract issued is so risky to investors (i.e., they are willing to accept only a very low tranching level q^* to assume the risk, one which is suboptimally low for the seller), that the seller's expected payoff might be higher by securitizing the entire net exposure rather than tranching it (case d).

< FIGURE 8 ABOUT HERE. >

6. Pooling multiple exposures

Extending the analysis of the previous sections to multiple exposures is relevant for two main reasons. First, if a market for longevity-linked securities were to take off, an important role would be played by institutions willing to take on (i.e., aggregate) longevity exposures for later repackaging and selling on to investors. Second, while a book of liabilities is typically regarded as a single liability, despite including several cohorts of policyholders, the holder might be better off securitizing only those exposures incurring high retention costs or likely to yield a higher securitization payoff. Similarly, some hedge suppliers find that aggregation of mortality rates by age ranges ('age bucketing') make mortality derivatives more attractive by reducing basis risk and enhancing the hedging potential of the instrument (see, e.g., Coughlan *et al.*, 2008). Here we show that age bucketing can also reduce the costs associated with asymmetric information.

We assume that the seller has n longevity exposures such as the one formalized in (3.1),[10] i.e.

$$D^i = q^i(X) + \varepsilon^i, \qquad (i = 1, \dots, n) \tag{6.1}$$

where $q^i(X) = E\left[D^i|X\right]$ represents a conditional unbiased estimate of the trend component of D^i, given the private signal X. We further require the following:

Assumption 6.1. *For each $i = 1, \dots, n$, the conditional distribution of $q^i(X)$, given the private valuations $(q^1(X), \dots, q^{i-1}(X), q^{i+1}(X), \dots, q^n(X))$, is continuous and has support $[q^{i,\min}, q^{i,\max}] \subseteq [0, 1]$.*

The above implies that there remains uncertainty with respect to (say) $q^i(X)$, i.e., $q^{i,\min} < q^{i,\max}$, no matter what information is available about any other $q^j(X)$, with $j \neq i$. Also, each lower bound $q^{i,\min}$ is independent of the available information on the other exposures.

If the holder were to securitize the exposures separately, she could use Proposition 4.2 to determine the equilibrium payoff (4.8) for each individual exposure, after observing the outcome of X. Note that Π depends on the distribution of the private valuation $q^i(X)$ only through the worst case $q^{i,\min}$. This means that even if investors were to learn something about the distribution of $q^i(X)$ from prior sales of D^j $(j \neq i)$, the equilibrium would be unchanged. Similarly, any information on X acquired by the seller before selling

[10]The results of this section do not apply to the alternative representation (3.2).

D^i would have no impact on the worst-case $q^{i,\min}$, which is constant by Assumption 6.1. In other words, we do not allow for learning.

Moving on to the aggregation of the exposures, the seller could determine the equilibrium payoff by treating $\sum_{i=1}^n D^i$ as a single exposure. Convexity of Π immediately reveals that selling the pool is suboptimal, since

$$\sum_{i=1}^{n} \Pi(q^i(X)) \geq \Pi\left(\sum_{i=1}^{n} q^i(X)\right).$$

The intuition is that the seller holds an option to choose the retention levels for the individual exposures, which means that her payoff is convex in the privately assessed severity of longevity risk (see Figure 2 for example). Hence combining low-longevity- and high-longevity-risk cashflows destroys the private information on the trend component.

On the other hand, we know that aggregation of exposures should allow the seller to benefit from risk pooling. We therefore examine whether the above result can be weakened by writing a contract on the pooled cashflows, as we did in the previous section, rather than passing them on directly to capital markets as in securitization. We denote the average exposure by

$$D_n = \frac{1}{n}\sum_{i=1}^{n} D^i, \tag{6.2}$$

and by $\overline{q}_n^{\min} := \frac{1}{n}\sum_{i=1}^n q^{i,\min}$ the average worst case. We do not consider the possibility of assigning different weights to the individual exposures, although the analysis could be extended to that case at the price of analytical tractability. For some nondecreasing function ψ, we define a contract on D_n by setting $C_n := \psi(D_n)$. As before, we denote its private valuation by $c_n(X) := E[C_n|X]$. The idea is to find conditions under which the expected payoff per exposure to the issuer is higher under pooling than under individual sale. Using the fact that Π is homogeneous of degree one, this happens when

$$\frac{1}{n} V\left(\sum_i D^i\right) = V(D_n) \geq \frac{1}{n}\sum_{i=1}^{n} V(D^i). \tag{6.3}$$

As in DeMarzo (2005), the trade-off between diversification benefits and the loss of private information from pooling can be optimally resolved as follows:

Proposition 6.2. *For $i = 1, \ldots, n$ assume that each exposure D^i has worst case $q^{i,\min} > 0$, and that the (ε^i) are log-concave i.i.d. and independent of $(q^i(X))$. Then, it is optimal* *to pool and tranche the exposures, and the optimal contract is given by*

$$C^* = \min(q_n^*, D_n) = q_n^* - \max(0, q_n^* - D_n).$$

Consider now a sequence of exposures $D^1, D^2, \ldots$ *satisfying the above assumptions and having bounded second moments. Assume that* $\lim_{n\to\infty} \overline{q}_n^{\min} = \overline{q}^{\min}$ *exists. Then, the limiting ex-ante securitization payoff is given by*

$$\lim_{n\to\infty} V\left(D_n\right) = (1-\delta)\overline{q}^{\min}, \tag{6.4}$$

with limiting tranching level $\lim_{n\to\infty} q_n^* = \overline{q}^{\min}$.

Interpretation of the optimal strike level q_n^* is as in the previous section. As expected, the benefit from pooling is increasing in the size of the pool. In the limit, the optimal contract is the one which is least insensitive to both private information and residual risk, as we now show. By following the same steps as in (5.4), we can write

$$V(D_n) = E\left[\Pi(c_n^*(X))\right] \le (1-\delta)c_n^{*,\min} \le (1-\delta)\overline{q}_n^{\min}, \tag{6.5}$$

where $c_n^{*,\min} := \min_x c_n^*(x)$. The term on the right-hand side represents the average holding cost of the n exposures. By taking limits on both sides of (6.5) and using (6.4), we then see that $\lim_{n\to\infty} c_n^{*,\min} = \overline{q}^{\min}$. This shows that pooling and tranching drives the securitization cost to the level computed in (4.10), the lowest possible cost associated with asymmetric information. Hence the additional cost associated with residual risk vanishes in a large enough pool, provided the assumptions of Proposition 6.2 are satisfied.

In the previous section, we saw (e.g., Figure 8) that tranching might be suboptimal when the residual risk is very high. The diversification benefits arising from pooling different exposures can be used to make tranching optimal. This is important when the seller writes contracts on volatile exposures, such as mortality rates relating to high ages or to small annuity books. For the opposite case when residual risk is undiversifiable, the information loss effect prevails and the following result holds:

Proposition 6.3. *If* $\varepsilon^i = \varepsilon$ *for all* i*, then it is optimal to securitize the exposures individually.*

The above proposition applies to cases where ε relates to events that affect all members of the pool equally, such as a catastrophic mortality event (e.g., an avian flu pandemic) or the development of a cure for cancer.

In Assumption 6.1 and Proposition 6.2, we have imposed no restriction on the dependence structure of the $q^i(X)$'s, apart from not being perfectly correlated (for otherwise the conditional support of $q^i(X)$, given the other private valuations, would collapse to a singleton). The following proposition helps us understand the role played by dependence in shaping the optimal contract (6.2):

Proposition 6.4. *Set $D'_n = \overline{q}'(X) + \frac{1}{n}\sum_{i=1}^n \varepsilon^i$, with $\overline{q}'(X) := \frac{1}{n}\sum_{i=1}^n q^{i\prime}(X)$. If $\overline{q}'(X)$ is a mean preserving spread of $\overline{q}(X) := \frac{1}{n}\sum_{i=1}^n q^i(X)$, then $V(D'_n) \geq V(D_n)$.*

Hence, pooling and tranching are more valuable when the private information on the severity of longevity risk affects the entire pool rather than being specific to each individual exposure. This is the case of death rates relating to individuals of different ages, but belonging to the same population, or relating to different cohorts of policyholders, but with similar underwriting characteristics. The above result may also justify the higher liquidity of mortality derivatives based on age bucketing.

As a practical example, consider several exposures $D^i = q^i(X) + \varepsilon^i$ $(i = 1, \dots, n)$ modelled as in the example at the end of the previous section. For simplicity, we assume that each $q^i(X)$ has the same lower bound $q^{i,\min} = \overline{q}^{\min}$. Figure 9 plots the expected payoff per exposure that can be obtained by writing contracts on pools of different sizes. As n increases, the payoff per exposure increases toward $(1-\delta)\overline{q}^{\min}$ and the optimal strike level q_n^* becomes closer to $\overline{q}^{\min}$.

To emphasize the role played by residual risk, we compare in Figure 10 the expected payoffs in the case of i.i.d. residual risks with those relative to the case of undiversifiable residual risk ($\varepsilon^i = \varepsilon$ for all i). If residual risk is diversifiable, the holder achieves not only a higher securitization payoff by pooling and tranching, but also a higher protection level against longevity risk. When exposures are tranched separately instead, the resulting payoff is unaffected by whether the residual risk is diversifiable or undiversifiable.

Finally, in Figure 11, we provide an illustration of Proposition 6.4, by considering a mean-preserving spread for the average trend component. When the private information on individual exposures is positively correlated, the payoff from pooling and tranching increases and the optimal contract results in a higher degree of protection against longevity risk for the seller.

< Figure 9 about here. >

< Figure 10 about here. >

< Figure 11 about here. >

7. Intensities of mortality

The framework of Section 3 can be easily recast in the familiar setting of stochastic intensities of mortality. Consider, for example, the survival rate S defined in (2.1). Set $z = 0$ and denote by $p(Y)$ the trend component of S. Assume that the random times $\tau^1, \ldots, \tau^n$ are conditionally Poisson over $[0, T]$, given the information generated by $(\mu_t)_{t\in[0,T]}$, a process representing the common intensity of mortality of the individuals in the population. We can then write (e.g., Brémaud, 1981)

$$E\left[S\right] = E\left[e^{-\int_0^T \mu_s \mathrm{d}s}\right]. \tag{7.1}$$

Consider now the case where μ can be expressed as $\mu = \widetilde{\mu} + Z$, where $(\widetilde{\mu}_t)_{t\in[0,T]}$ is the intensity of a reference population, while $(Z_t)_{t\in[0,T]}$ is an independent adjustment process. If the dynamics of $\widetilde{\mu}$ can be publicly estimated, while those of Z are available to the informed agent only, we can set $X = Z_0$, $Y = (\widetilde{\mu}_0, Z_0)$ and obtain an expression for the private valuation of S similar to (3.4):

$$\begin{aligned} p(Y) = E[S|Y] &= E\left[e^{-\int_0^T (\widetilde{\mu}_s + Z_s)\mathrm{d}s}\middle| Y\right] \\ &= E\left[e^{-\int_0^T \widetilde{\mu}_s \mathrm{d}s}\middle|\widetilde{\mu}_0\right] E\left[e^{-\int_0^T Z_s \mathrm{d}s}\middle| Z_0\right] = \widetilde{p}(\widetilde{\mu}_0)m(Z_0). \end{aligned} \tag{7.2}$$

In Section 4, we allowed the holder of the longevity exposure to transfer it to the capital markets. The following steps are involved: (i) immediately before time 0, the seller chooses the contract design; (ii) at time 0, the seller observes Z_0 and sells a fraction of the contract to the investors based on her private valuation of the exposure; (iii) at date T, the final payoff is revealed to both the seller and the investors. If an investor were to take over the exposure at time 0, she would be able to back out the dynamics of the signal Z by monitoring the exposure over $(0, T]$. Still, information on Z_0 is only available to the holder of the exposure before sale.

Proposition 5.1 can be used with representation $S = p(Y) + \varepsilon$ (or $S = p(Y)\exp(\epsilon)$) and $p(Y)$ given by (7.2), provided S admits a uniform worst case. This occurs, for example, when ε (or ϵ) and $\widetilde{p}(\widetilde{\mu}_0)$ are log-concave and $m(Z_0)$ is continuous. A special case is when $\widetilde{\mu}$ is deterministic, which means that the signal is entirely private.

8. Conclusions

In this paper, we examined the impact of asymmetric information and parameter uncertainty on the securitization and tranching of longevity exposures when agents are risk neutral. We first considered the securitization of annuity-like liabilities and their backing assets. The use of a signalling model has allowed us to understand the optimal retention levels or risk premia that would materialize in a transaction where investors have partial information on longevity risk and the holder faces regulatory retention costs. We then considered the optimal design of a derivative contract written on longevity-linked cash-flows, emphasizing the joint role played by both systematic and unsystematic longevity risk. In several interesting cases, it is optimal to tranche the exposure or, equivalently, to transfer the liabilities in exchange for a call option on mortality rates (or a put option on survival rates). We used the equilibrium model to understand how the strike levels on such options can be set to minimize the impact of asymmetric information and maximize the option seller's payoff. Finally, we examined the benefits arising from pooling and tranching several longevity exposures. The results shed light into the way regulatory costs and asymmetric information can shape the longevity risk premia originating from transactions in longevity-linked securities. Our analysis provides a useful theoretical basis for designing strategies for securitizing, tranching and pooling of longevity exposures as well as designing option contracts on these exposures.

References

Akerlof, D. (1970). The market for 'lemons': Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics*, vol. 89:488–500.

Barrieu, P. and N. El Karoui (2005). Pricing, hedging and optimally designing derivatives via minimization of risk measures. In R. Carmona (ed.), *Indifference Pricing*. Princeton University Press, *to appear*.

Bauer, D. and F. Kramer (2007). Risk and valuation of mortality contingent catastrophe bonds. Tech. rep., University of Ulm.

Blake, D. and W. Burrows (2001). Survivor bonds: Helping to hedge mortality risk. *Journal of Risk and Insurance*, vol. 68:339–348.

Blake, D., A. Cairns and K. Dowd (2006a). Living with mortality: Longevity bonds and other mortality-linked securities. *British Actuarial Journal*, vol. 12:153–197.

Blake, D., A. Cairns, K. Dowd and R. MacMinn (2006b). Longevity bonds: Financial engineering, valuation and hedging. *Journal of Risk and Insurance*, vol. 73:647–72.

Blake, D., A. Cairns and K. Dowd (2008). The birth of the Life Market. *Asia-Pacific Journal of Risk and Insurance*, vol. 3(1):6–36.

Brémaud, P. (1981). *Point Processes and Queues: Martingale Dynamics*. Springer.

Cairns, A., D. Blake, K. Dowd, G. Coughlan, D. Epstein, A. Ong and I. Balevich (2007). A quantitative comparison of stochastic mortality models using data from England & Wales and the United States. Tech. rep., Pensions Institute Discussion Paper PI-0701.

Cairns, A., D. Blake, K. Dowd, G. Coughlan, D. Epstein and M. Khalaf-Allah (2008). Mortality density forecasts: An analysis of six stochastic mortality models. Tech. rep., Pensions Institute Discussion Paper PI-0801.

Coughlan, G., D. Epstein, A. Sinha and P. Honig (2007). q-forwards: Derivatives for transferring longevity and mortality risks. Tech. rep., J.P. Morgan's Pension Advisory Group.

Coughlan, G., D. Epstein, C. Watts and M. Khalaf-Allah (2008). Hedging pension longevity risk: Practical capital markets solutions. *Asia-Pacific Journal of Risk and Insurance*, vol. 3(1):65–89.

Cowley, A. and J. Cummins (2005). Securitization of life insurance assets and liabilities. *Journal of Risk and Insurance*, vol. 72(2):193–226.

Dana, R. and M. Scarsini (2007). Optimal risk sharing with background risk. *Journal of Economic Theory*, vol. 133:152–176.

Dawson, P., K. Dowd, A. Cairns and D. Blake (2007). Completing the survivor derivatives market: A consistent pricing framework. Tech. rep., Pensions Institute Discussion Paper PI-0712.

DeMarzo, P. (2005). Pooling and tranching of securities: A model of informed intermediation. *Review of Financial Studies*, vol. 18(1):65–99.

DeMarzo, P. and D. Duffie (1999). A liquidity based model of security design. *Econometrica*, vol. 67(1):65–99.

Denuit, M. (2007). Indexed life annuities and longevity-based financial instruments. Tech. rep., Université Catholique de Louvain.

Dowd, K., D. Blake, A. Cairns and P. Dawson (2005). Survivor swaps. *Journal of Risk and Insurance*, vol. 73(1):1–17.

Dowd, K., A. Cairns, D. Blake, G. Coughlan, D. Epstein and M. Khalaf-Allah (2008a). Evaluating the goodness of fit of stochastic mortality models. Tech. rep., Pensions Institute Discussion Paper PI-0802.

Dowd, K., A. Cairns, D. Blake, G. Coughlan, D. Epstein and M. Khalaf-Allah (2008b). Backtesting stochastic mortality models: An ex-post evaluation of multi-period-ahead density forecasts. Tech. rep., Pensions Institute Discussion Paper PI-0803.

Duffie, D. and R. Rahi (1995). Financial innovation and security design. *Journal of Economic Theory*, vol. 65:1–42.

Eeckhoudt, L., C. Gollier and H. Schlesinger (2005). *Economic and Financial Decisions under Risk*. Princeton University Press.

Gale, D. (1992). A Walrasian theory of markets with adverse selection. *Review of Economic Studies*, vol. 59:229–255.

Gourieroux, C. and A. Monfort (2008). Quadratic stochastic intensity and prospective mortality tables. *Insurance: Mathematics and Economics*, vol. 43:174–184.

Innes, R. (1990). Limited liability and incentive contracting with ex-ante action choices. *Journal of Economic Theory*, vol. 52:45–67.

Karlin, S. (1968). *Total Positivity*. Stanford University Press.

Keilson, J. and U. Sumita (1982). Uniform stochastic ordering and related inequalities. *Canadian Journal of Statistics*, vol. 10(3):181–198.

Krutov, A. (2006). Insurance-linked securities: An emerging class of financial instruments. *Financial Engineering News*, (48 (March-April)):7–16.

Loeys, J., N. Panigirtzoglou and R. Ribeiro (2007). Longevity: A market in the making. Tech. rep., JP Morgan's Global Market Strategy.

Mailath, G. (1987). Incentive compatibility in signalling games with a continuum of types. *Econometrica*, vol. 55:1349–1366.

Modu, E. (2008). Life settlement securitization. Tech. rep., A.M. Best Structured Finance.

Plantin, G. (2004). Tranching. Tech. rep., Tepper School of Business.

Rothschild, M. and J. Stiglitz (1970). Increasing risk: I. A definition. *Journal of Economic Theory*, vol. 2:225–243.

Schervish, M. (1995). *Theory of Statistics*. Springer.

Tirole, J. (2006). *The Theory of Corporate Finance*. Princeton University Press.

Whitt, W. (1980). Uniform conditional stochastic order. *Journal of Applied Probability*, vol. 17:112–123.

Whitt, W. (1982). Multivariate monotone likelihood ratio and uniform conditional stochastic order. *Journal of Applied Probability*, vol. 19:675–701.

Appendix A. Proofs

Proof of Proposition 4.1. *(i)* For fixed δ and γ, the function $\pi(p) = \gamma(P(\gamma) - \delta(\alpha - p))$ is increasing affine. Hence Π is the upper envelope of a family of affine functions decreasing in the private valuation of (4.1). As a result, Π is nonincreasing and convex in the private valuation. *(ii)* The subdifferential of Π contains $-\delta\Gamma$. Hence, Γ is nonincreasing in the private valuation of the net exposure. □

Proof of Proposition 4.2. We can first verify that the separation property (4.6) holds for Γ^* defined by (4.7). Since the latter is decreasing in the private valuation, by (4.5) we have:

$$P^*(\Gamma^*(\alpha - p(X))) = E\left[\alpha - p(X) | \Gamma^*(\alpha - p(X))\right] = \alpha - E\left[p(X)|p(X)\right] = \alpha - p(X).$$

Denote by $U(e, P(\gamma), \gamma) := \gamma(P(\gamma) - \delta e)$ the securitization payoff to the holder when the private valuation of the net exposure is $e = \alpha - p(x)$. It can be verified that U satisfies the conditions of Mailath (1987, Theorem 3). In particular, the single crossing property is satisfied, in that the map $U_3(\cdot, P(\gamma), \gamma)/U_2(\cdot, P(\gamma), \gamma)$ is strictly monotone (with U_h denoting partial derivative with respect to the h-th argument). Hence there exists a unique separating equilibrium (P^*, Γ^*), with P^* differentiable. The latter can be found by using the first-order conditions for problem (4.4), yielding the differential equation

$$\begin{cases} \gamma P'(\gamma) + P(\gamma) - \delta e = 0 \\ P(1) = \alpha - p^{\max}. \end{cases}$$

By the separation property, we can replace e with $P(\gamma)$ and solve the above to obtain $P^*(\gamma) = \gamma^{\delta - 1}(\alpha - p^{\max})$. Finally, the optimal Γ^* can be found by inverting P^* and using once again the separation property:

$$\Gamma^*(\alpha - p(X)) = \left(\frac{P^*(\gamma)}{\alpha - p^{\max}}\right)^{\frac{1}{\delta - 1}} = \left(\frac{\alpha - p(X)}{\alpha - p^{\max}}\right)^{\frac{1}{\delta - 1}}.$$

□

Proof of Proposition 5.1. We begin by giving a definition of the uniform worst case property, introduced by DeMarzo and Duffie (1999) and related to the concept of uniform conditional stochastic order (Whitt, 1980, 1982). Letting λ_x denote the conditional law of D, given an outcome x of X, we define:

Definition A.1. *An outcome $x_0 \in \mathbb{R}^l$ of X is a 'uniform worst case' for D if, for every subinterval B of $[0,1]$ and every other outcome x of X, we have that $\lambda_x(B) > 0$ implies $\lambda_{x_0}(B) > 0$ and that, conditionally on B, λ_x stochastically dominates λ_{x_0}.*

Since $q(X)$ is continuous, and ε (or ϵ) is log-concave, x_0 is a uniform worst case for C (e.g., Keilson and Sumita, 1982). Denoting by l_{x,x_0} the Radon-Nykodim derivative of λ_x with respect to λ_{x_0}, by the definition of uniform worst case we have that l_{x,x_0} exists and can be chosen to be increasing. As a result, for any contract $C = \phi(D)$ (with ϕ nondecreasing and satisfying $\phi(d) \leq d$ on $[0,1]$), the following holds:

$$\begin{aligned}\min_x E\left[C|X=x\right] &= \min_x \int_0^1 \phi(u)\ \lambda_x(\mathrm{d}u) = \min_x \int_0^1 \phi(u)\ l_{x,x_0}(u)\ \lambda_{x_0}(\mathrm{d}u)\\ &= \int_0^1 \phi(u)\ \lambda_{x_0}(\mathrm{d}u) = E\left[C|X=x_0\right] =: c^{\min}.\end{aligned}$$

From here the proof is essentially as in DeMarzo and Duffie (1999, Proposition 10). We cover it for completeness. Consider the contract $C = \min(k, D)$. Since both $c(x) = E\left[C|X=x\right]$ and $c^{\min}$ are continuous in k, we can choose $k = k^*$ such that $c^{\min} = c(x)$. For some nondecreasing function ψ satisfying $\psi(d) \leq d$ on $[0,1]$, consider another contract $\hat{C} = \psi(D)$ and set $\hat{c}(x) = E[\hat{C}|X=x]$. Define now $\tilde{C} := \hat{C} - C$. Since x_0 is a uniform worst case for D, we have $\hat{c}^{\min} = c^{\min}$ and hence $\tilde{c}^{\min} = 0$. For each outcome x of X, we can write

$$\tilde{c}(x) = E[\tilde{C}|X=x] = \int_0^1 (\psi(u) - \phi(u))\ l_{x,x_0}(u)\ \lambda_{x_0}(\mathrm{d}u).$$

Since ϕ and ψ are both nondecreasing and bounded by d for all $d \in [0,1]$, we have that $\tilde{C} \geq 0$ if and only if $D \geq d^*$ for some d^*. From the above, we then have

$$\begin{aligned}\tilde{c}(x) &\geq \int_0^1 (\psi(u) - \phi(u))\ l_{x,x_0}(d^*)\ \lambda_{x_0}(\mathrm{d}u)\\ &= l_{x,x_0}(d^*) \int_0^1 (\psi(u) - \phi(u))\ \lambda_{x_0}(\mathrm{d}u) = l_{x,x_0}(d^*)\ \tilde{c}^{\min} = 0.\end{aligned}$$

As a consequence, $\hat{c}(X) \geq c(X)$ for any contract $\hat{C}$. Since Π is decreasing in the private valuation by Proposition 5.1(ii), we have that $\Pi(\hat{c}(X)) \leq \Pi(c(X))$ and hence

$$E\left[\Pi(\hat{c}(X))\right] \leq \max_C E\left[\Pi(c(X))\right] = V(D).$$

The optimal contract is then some $C^* = \min(k^*, D)$, with the optimal strike level k^* given by

$$\{k^*\} = \arg\max_k E\left[\Pi(E\left[\min(k, D)|X\right])\right].$$

$\square$

Proof of Proposition 6.2. See DeMarzo (2005, Theorem 2). □

Proof of Proposition 6.3. For fixed n, by Proposition 5.1, the optimal contract written on D_n is $C_n^* = \min(q_n^*, D_n)$. Since the private valuation $c_n^*(x) = E\left[C_n^*|X = x\right] = E\left[\min(q_n^*, \overline{q}(x) + \varepsilon)\right]$ is concave in $\overline{q}(x) = \frac{1}{n}\sum_i q^i(x)$, we have $c_n^*(x) \geq \frac{1}{n}\sum_i c^{i,*}(x)$ for every outcome x of X. Using the fact that the equilibrium payoff (4.8) is decreasing and convex, we then have

$$\Pi\left(c_n^*(X)\right) \leq \Pi\left(\frac{1}{n}\sum_i c^{i,*}(X)\right) \leq \frac{1}{n}\sum_i \Pi\left(c^{i,*}(X)\right).$$

Taking the expectation, we finally obtain $V(D_n) \leq \frac{1}{n}\sum_i V(D^i)$. □

Proof of Proposition 6.4. For each outcome x of X, we have that $c_n^*(x) = E[\min(q_n^*, \overline{q}(x) + \frac{1}{n}\sum_{i=1}^n \varepsilon^i)]$ is nondecreasing and concave in $\overline{q}(x)$. Similarly, $c_n'^*(x) = E\left[\min(q_n'^*, \overline{q}'(x) + \frac{1}{n}\sum_{i=1}^n \varepsilon^i)\right]$ is nondecreasing and concave in $\overline{q}'(x)$. Since Π is decreasing and convex, and $\overline{q}'(x)$ is a mean preserving spread of $\overline{q}(x)$, we have $E[\Pi(c'^*(X))] \geq E[\Pi(c^*(X))]$ (Rothschild and Stiglitz, 1970). □

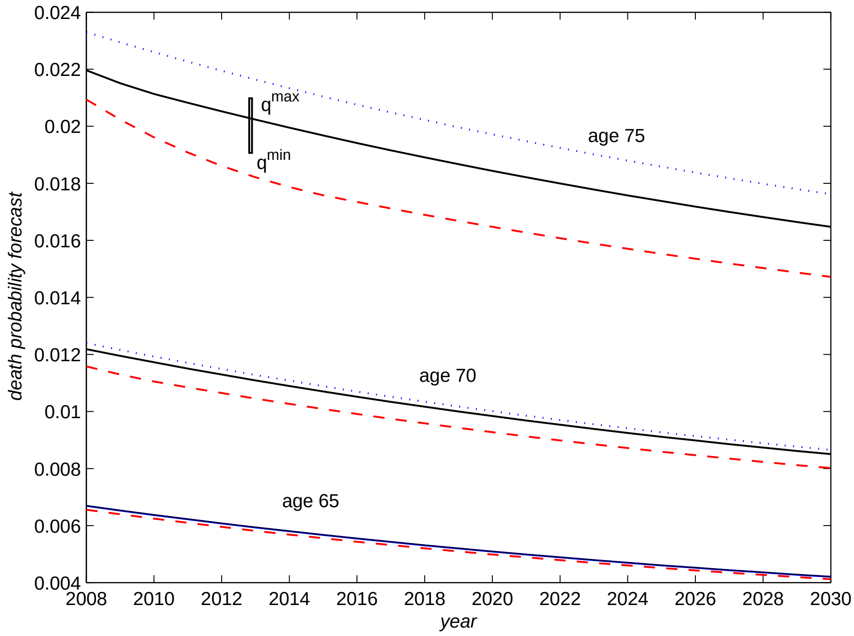

Figure 1: CMI forecasts for death probabilities at ages 65, 70 and 75 over 2008 − 2030: short-cohort (dotted), medium-cohort (solid) and long-cohort (dashed) forecast. The interval $[q^{\min}, q^{\max}]$ provides an example for the support of the private valuation of the death rate relative to age 75 in 2013.

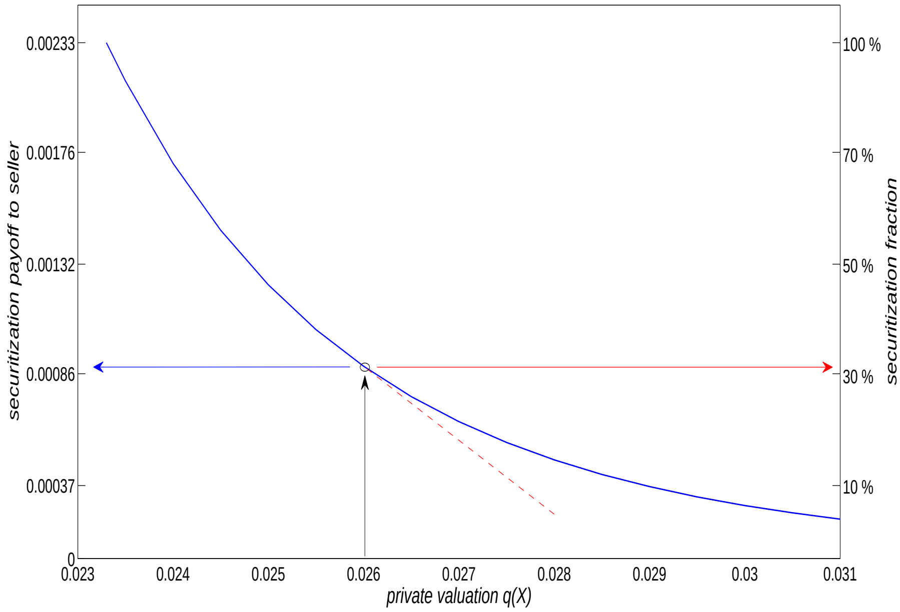

Figure 2: Securitization payoff (4.4) and securitization fraction as a function of the private valuation of the net exposure. The subdifferential of Π is $-\delta\Gamma$. The plot is based on the parameter values $\alpha = 1$ and $\delta = 0.9$, so that $\alpha - p(X) = q(X)$ represents the one-year death rate for age 75 in year 2009, with $q^{\min}$ coinciding with the CMI long-cohort forecast.

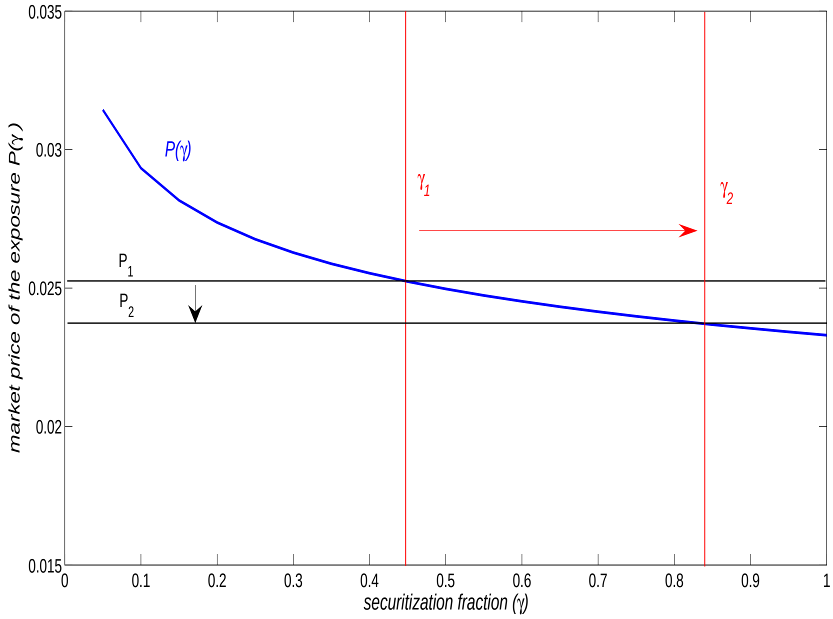

Figure 3: Investor demand for longevity exposure: If investors infer the severity of longevity risk from Γ, an increase from $\gamma_1 = \Gamma(\alpha - p(x_1))$ to $\gamma_2 = \Gamma(\alpha - p(x_2))$ induces a decrease in the market price for the net exposure from $P_1 = P(\gamma_1)$ to $P_2 = P(\gamma_2)$.

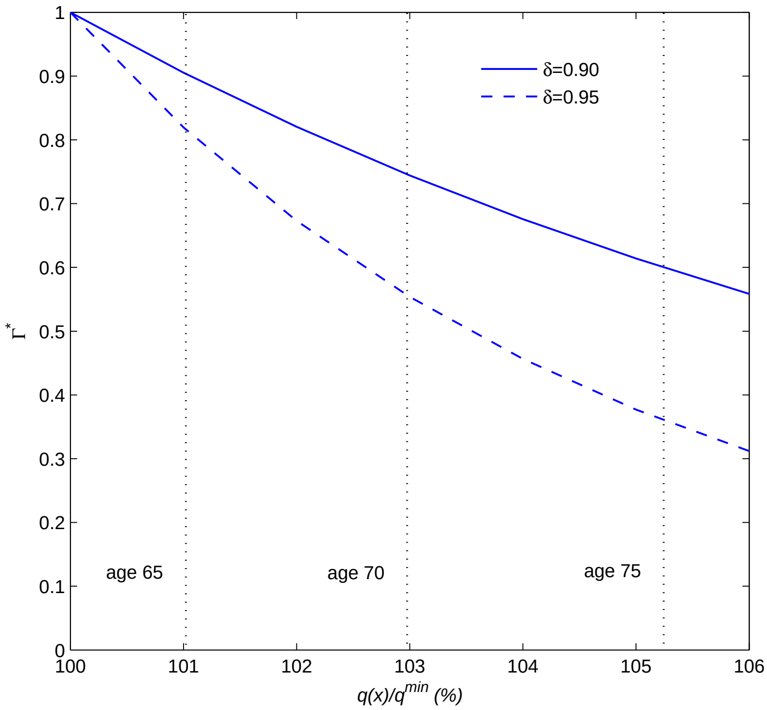

Figure 4: Securitization fraction Γ^* for $\delta = 0.90$ (solid line) and $\delta = 0.95$ (dashed line) as a function of $q(x)/q^{\min}$, and values of the ratio for ages 65, 70 and 75 (dotted lines).

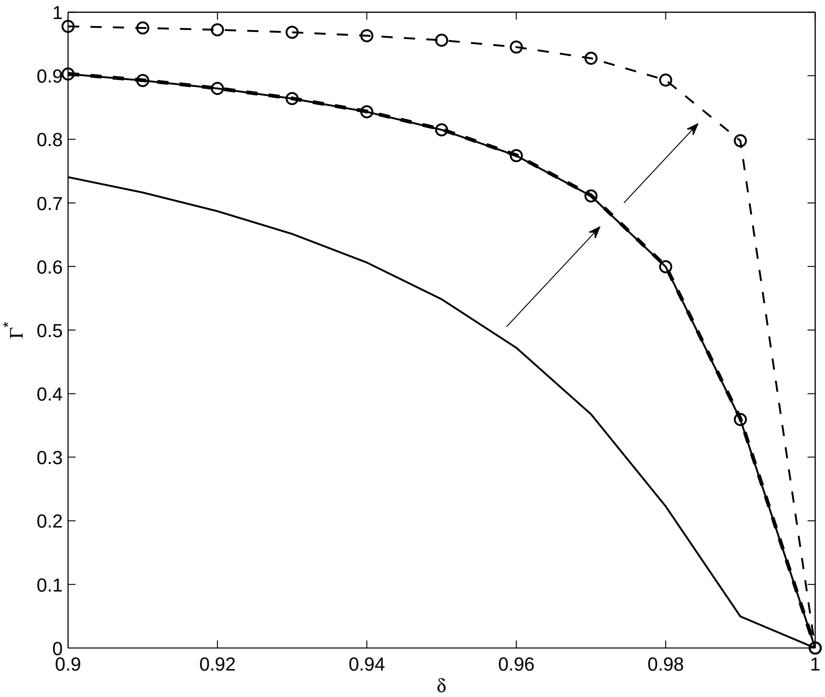

Figure 5: Securitization fraction Γ^* as a function of δ for ages 65 (dashed) and 70 (solid), when $\alpha = 1$ and $\alpha = 1.021$ (circles).

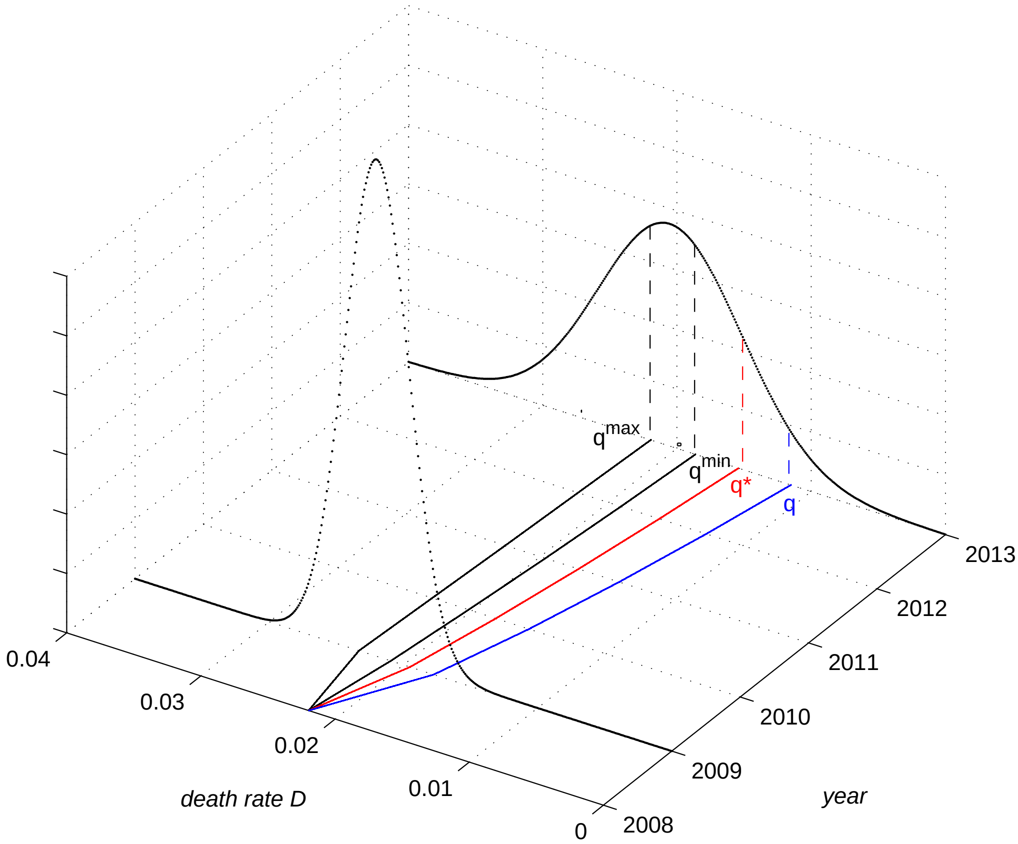

Figure 6: Densities of the death rate of a 75-year old in 2009 and 2013 viewed from 2008. The solid lines $q^{\max}$ and $q^{\min}$ represent, respectively, the CMI short-cohort and long-cohort forecasts for age 75 over $2008-2013$. The lines indicated by q^{*} and q represent, respectively, the optimal tranching level and a suboptimal tranching level over $2008-2013$.

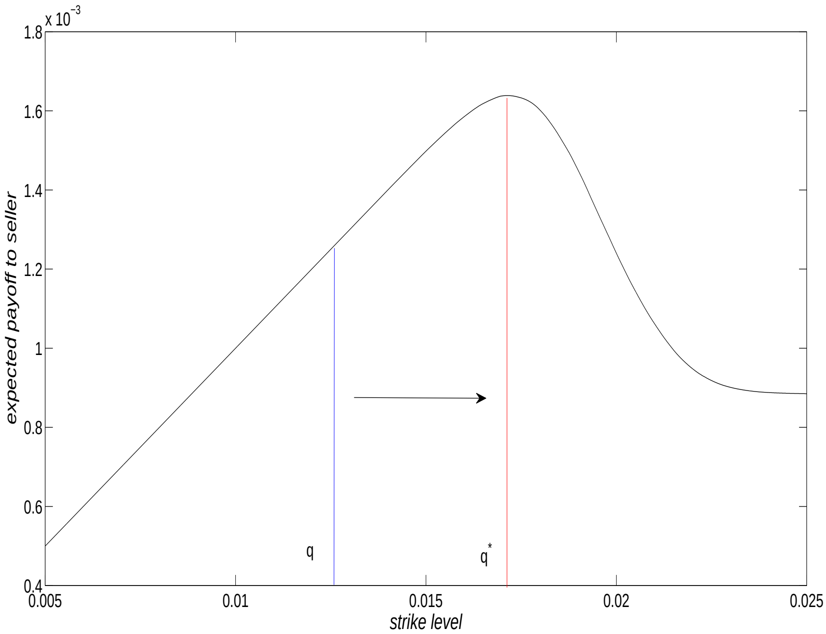

Figure 7: Optimality of the tranching/strike level q^* in Figure 6.

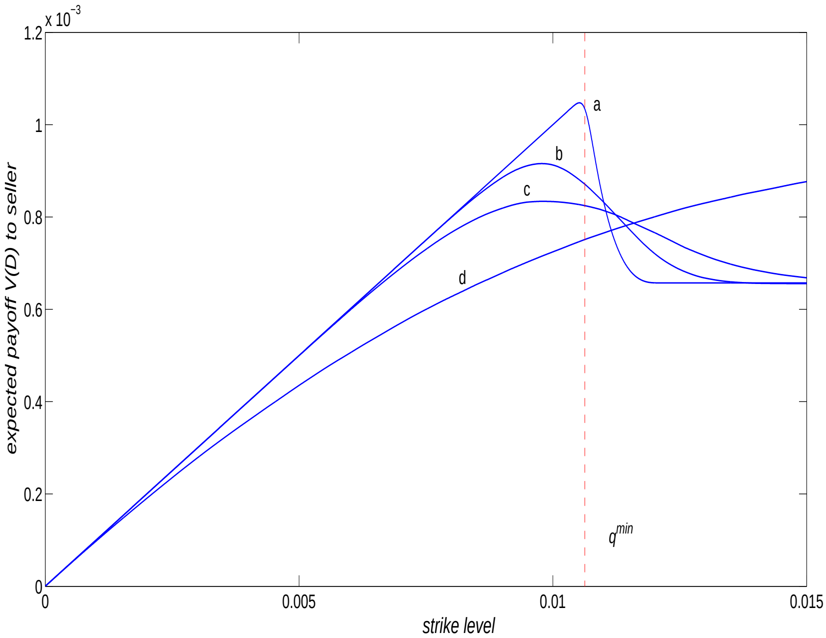

Figure 8: Ex-ante securitization payoff for $\delta = 0.90$ and different values for the standard deviation of ε: $\sigma = 0.0001$ (0.885% of $E[D]$) in case a, $\sigma = 0.0010$ (8.85% of $E[D]$) in case b, $\sigma = 0.0020$ (17.71% of $E[D]$) in case c, $\sigma = 0.0050$ (44.27% of $E[D]$) in case d.

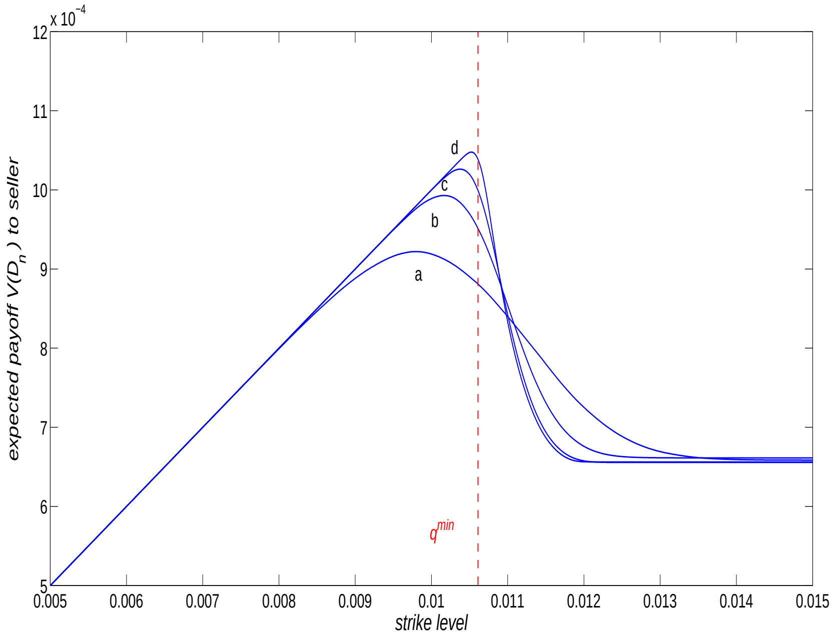

Figure 9: Expected payoff as a function of the tranching level for different pool sizes: $n = 1$ in case a, $n = 5$ in case b, $n = 10$ in case c and $n = 20$ in case d. As n grows larger, the optimal tranching level approaches $\overline{q}^{\min}$ (dashed).

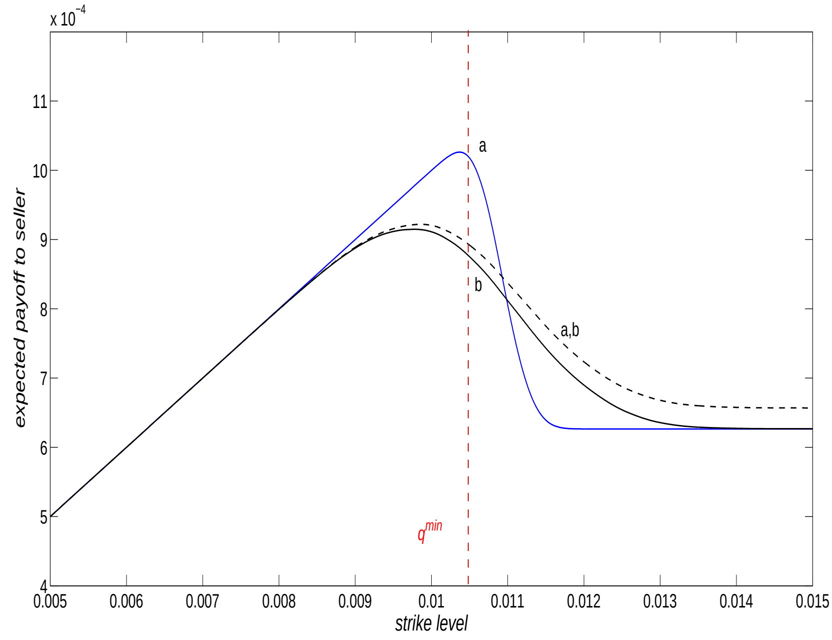

Figure 10: Expected payoff $V(D_{20})$ (solid lines) and $\frac{1}{20}\sum_{i=1}^{20} V(D^i)$ (dashed line) for net exposures of type $q^i(X)+\varepsilon^i$ (case a) and $q^i(X)+\varepsilon$ (case b), with $q^i(X)$ i.i.d. uniformly distributed as in Section 5, and ε^i and ε i.i.d. truncated Normal with mean zero and standard deviation 0.001.

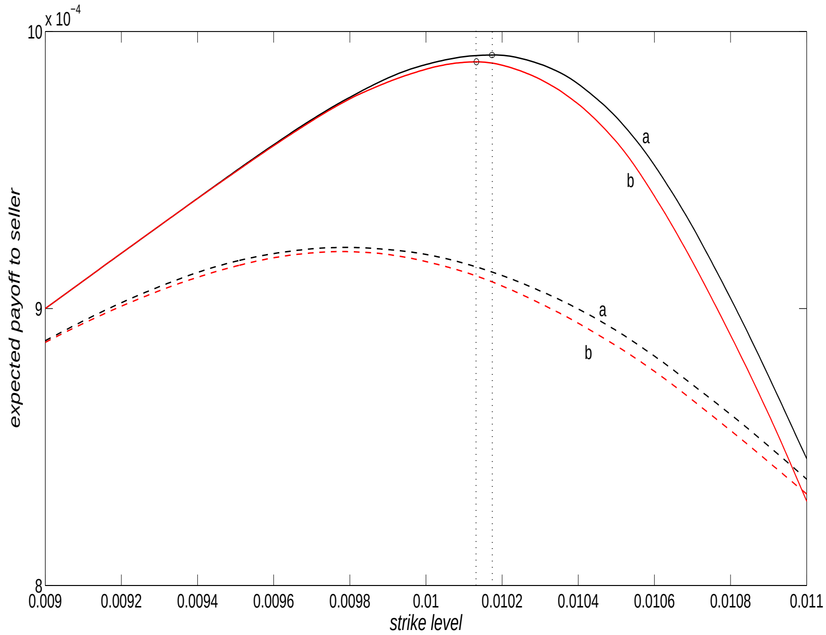

Figure 11: Expected payoff $V(D_1)$ (dashed) and $V(D_5)$ (solid) with $q^{i\prime}(X)+\varepsilon^i$ (case a) and $q^i(X)+\varepsilon^i$ (case b), with the $(q^i(X))$ independent and the $(q^{i\prime}(X))$ having 0.5 correlation. The (ε^i) are i.i.d. truncated Normal with mean zero and standard deviation 0.001.