City Research Online

City, University of London Institutional Repository

DISCUSSION PAPER PI-1207

Keeping Some Skin in the Game: How to Start a Capital Market in Longevity Risk Transfers

Enrico Biffis and David Blake

December 2013

Keeping Some Skin in the Game: How to Start a Capital Market in Longevity Risk Transfers

Enrico Biffis[*]   David Blake[†]

October 25, 2013

Abstract

The recent activity in pension buy-outs and bespoke longevity swaps suggests that a significant process of aggregation of longevity exposures is under way, led by major investment banks and buy-out firms with the support of leading reinsurers. As regulatory capital charges and limited reinsurance capacity constrain the scope for market growth, there is now an opportunity for institutions that are pooling longevity exposures to issue securities that appeal to capital market investors, thereby broadening the sharing of longevity risk and increasing market capacity. For this to happen, longevity exposures need to be suitably pooled and tranched to maximize diversification benefits offered to investors and to address asymmetric information issues. We argue that a natural way for longevity risk to be transferred is through suitably designed principal-at-risk bonds.

[*]Biffis (`E.Biffis@imperial.ac.uk`) is at Imperial College Business School, Imperial College London, South Kensington Campus, SW7 2AZ United Kingdom.

[†]Blake (`D.Blake@city.ac.uk`; corresponding author) is at the Pensions Institute, Cass Business School, City University London, 106 Bunhill Row, EC1Y 8TZ United Kingdom.

1 Introduction

The total amount of pension-related longevity risk exposure in private-sector corporations is globally estimated at $25 trillion.[1] In 2006, a market started in the UK to transfer longevity risk from the pension plans of UK corporations. From 2010, this market became international, with transfers taking place in Holland, the US and Canada. However, the total value of the transfers since the market started has been little more than $100 billion, a small fraction of the total exposure. In addition, most of the transfers have ended up with insurance and reinsurance companies.[2] A small proportion of the total has involved investment banks which have passed the risk exposure on to capital markets investors. This rate of transfer to the capital markets needs to grow significantly, since there is insufficient capital in the insurance and reinsurance industry to absorb the total exposure. Furthermore, long-term capital markets investors need to be persuaded to hold longevity risk exposure in a form in which they are comfortable and with a suitable longevity risk premium which reflects the true level of risk they are assuming. Long-term investors, such as sovereign wealth funds, endowments, family offices, should find a longevity-linked asset an attractive one to hold in a diversified portfolio, since its return will have low correlation with the returns on other asset classes, such as financial assets, real estate and commodities.

However, such investors are currently wary of longevity-linked assets. If experts in longevity risk, such as insurance companies, are trying to sell this risk on to them, is this just because they see this simply as an opportunity to gain some capital relief from their regulator, a perfectly legitimate business practice? Or is it because they have

---

[1]The Life and Longevity Markets Association.

[2]The market was started by newly established monoline insurers known as buy-out companies. Later, larger traditional insurance companies and reinsurance companies and investment banks entered the market. See LCP (2012) for an overview of market participants and recent transactions.

become aware that the risk is greater than they initially believed and are trying to offload it on to unsuspecting investors before the true extent of the exposure becomes more widely known? A relevant problem here is asymmetric information. One party to the transaction knows more about what is being offered for sale than the other party. In the extreme case, the potential buyer in the transaction might be so suspicious of the potential seller's intention that they will not transact at any price: in this case, there will be no market. For a market to exist in the presence of asymmetric information, the potential seller needs to find a way of communicating to the potential buyer the true extent of the information that they have acquired about the item up for sale. In other words, they have to provide an appropriate signal to the potential buyer that they are not being sold a 'lemon'.[3]

In this article, we outline how a market in longevity risk transfers may grow out of the recent activity in pension liability transfers. Figure 1 provides a stylized representation of the potential participants in such a market. There are some holders of longevity risk exposures, which we may regard as 'informed'. An example is a large pension plan which has undertaken a longevity study of its members and beneficiaries, and therefore has acquired private information about its own longevity exposure. There are then holders of longevity exposures which (oversimplifying) we may regard as 'uninformed'. An example is a small pension plan which does not have the resources to undertake a longevity study of its members. There are also informed intermediaries, such as insurance companies and investment banks, which have conducted their own longevity studies of the members of pension plans, but are considering selling on their acquired longevity risk exposure either directly to investors or to intermediaries in insurance-linked securities (ILSs).

_____

[3]A 'lemon' is something, such as a second-hand car, which is revealed to be faulty only after it has been purchased (e.g., Akerlof, 1970).

Finally, there are end investors (and other intermediaries), whom we may regard as 'uninformed', in the sense that they have not conducted any longevity studies of their own, possibly because they do not have access to the plan data needed to conduct such a study. They therefore have less information about the true extent of the longevity risk in which they are being invited to invest than those selling it, a classic case of asymmetric information. An effective way of overcoming the problem of asymmetric information for an informed seller is to disclose information about the relevant risk characteristics of the exposures being transferred (e.g., gender, age, medical status, similar details on potential beneficiaries of any survivor benefits), and to signal their quality by retaining part of the exposures. Another effective tool is to structure the risk transfer vehicle so as to minimize the sensitivity of the vehicle's value to the issuer's private information. This would typically entail pooling and tranching the exposure, as we outline below.

Since the bulk of global longevity risk exposures is represented by the liabilities of defined benefit pension plans, the market for pension buy-outs and longevity swaps is an important origination market that could support the development of a market in longevity-linked securities. In particular, buy-out firms act as aggregators of the pension liabilities of small companies into larger pools, and are *de facto* the natural candidates to intermediate the transfer of longevity exposures originating from the pension buy-out market to capital market investors. Because of capital requirements, and the returns that can be generated by deploying resources in the buy-out market, buy-out firms have an incentive to securitize their exposures to diversify risk and free up capital. There is an important role for regulators in this process, as they may provide incentives to disclose and use detailed information from the very same internal models used to demonstrate the capital resilience of (re)insurers and buy-out firms in the pension buy-out market.
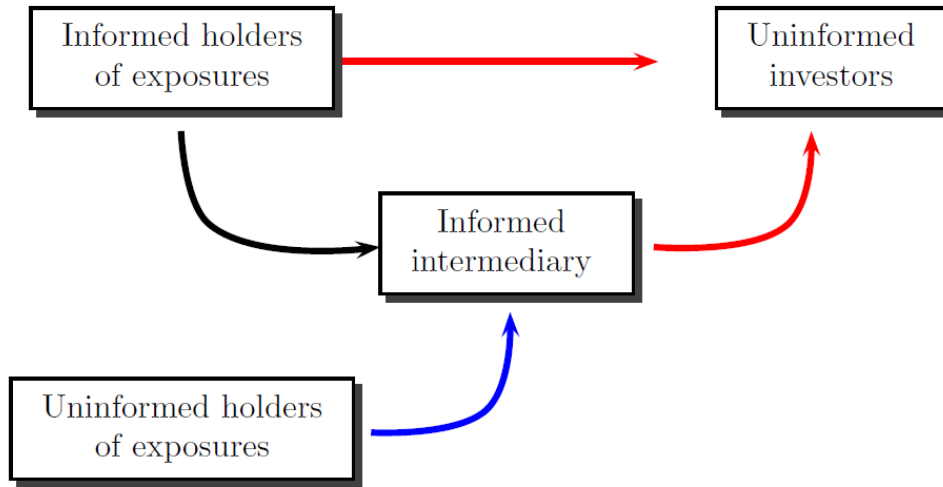
4

Figure 1: The participants in a market for longevity risk transfers.

2 Lessons learned so far, and the way forward

During the last decade, several attempts have been made to launch standardized longevity-linked securities. The longevity bond designed by the European Investment Bank (EIB) in 2004, and immediately withdrawn due to insufficient demand, is probably the most widely discussed example, due to the important lessons that can be learned from its failure to launch. The EIB instrument was a 25-year bond with an issue price of £540m and coupons linked to a cohort survivor index based on the realized mortality rates of English and Welsh males aged 65 in 2002. The initial coupon was set equal to £50m, while the subsequent coupons would have reduced in line with the realized mortality of the reference cohort (see Blake *et al.*, 2006, for additional details). Hence, the higher the number of survivors in the population each year, the higher the coupons paid to investors, meaning that the instrument should have been appealing to pension plans and annuity providers. There are two main reasons why the EIB bond did not launch

and we would like to discuss:[4]

- **Design issues.** As a hedging instrument, the EIB bond did not offer sufficient flexibility. The hedge was bundled up within a *bond*, and provided no leverage opportunities. The bond format meant that a considerable upfront payment was required to access the longevity hedge component of the instrument, represented by a longevity swap which paid the longevity-linked coupons. Furthermore, the *basis risk* in the bond was considered to be too great: the bond's mortality index covered just a single cohort of 65-year-old males from the national population of England and Wales, representing only a fraction of the average pension plan's and annuity provider's exposure to longevity risk.[5]

- **Transparency issues**. The EIB bond's projected cash flows depended on projections of the future mortality of 65-year-old males from England & Wales prepared by the UK Government Actuary's Department (GAD). The forecasting model used for the projections is not published, and the projections themselves are a result of adjustments made to the baseline forecasts in response to expert opinion. This represented a major barrier to investors and hedgers either not familiar with longevity risk or with strong views on specific mortality projection models. On the pricing side, the longevity risk premium built into the initial price of the EIB bond was set at 20 basis points (Cairns *et al.*, 2005): in the absence of agreement on a baseline best estimate, investors had no real feeling as to how appropriate this figure was.

---

[4]See Blake *et al.* (2006); Biffis and Blake (2010b); Blake *et al.* (2013) for further details.

[5]The typical longevity exposure of a pension plan spans different cohorts of active and retired members, not to mention the fact that a large portion of pensions paid by pension funds and annuity providers are indexed to inflation.

In the wake of the EIB bond withdrawal, there has been considerable innovation taking the form of longevity indices and mortality derivatives, for example. Despite several attempts, so far standardized solutions have not been as successful as those following the more traditional insurance paradigm, where the hedger is indemnified (fully hedged) against changes in longevity risk. Successful longevity risk transfers have taken the form of pension buy-outs, pension buy-ins, and longevity swaps. Pension buy-outs involve the transfer of a pension plan's assets and liabilities to a regulated life insurer. The transaction allows the employer to off-load the pension liabilities from its balance sheet, or to replace pension assets and liabilities with a regular loan in case buy-out costs are financed by borrowing.[6] In a pension buy-in, a bulk annuity contract is purchased as a plan asset to match some or all of the pension plan's liabilities, meaning that interest rate, inflation, and longevity risk are all fully hedged, while the liabilities remain in the pension plan. Longevity swaps allow a hedger to receive longevity-linked payments (the floating leg) in exchange for a stream of fixed payments (the fixed leg). The most popular transactions to date involve bespoke, or customized, swaps, where the floating leg is linked to the specific mortality experience of the hedger. These forms of longevity risk transfer have successfully addressed the design issues mentioned above, by removing the basis risk originating from differences in the mortality experience of the hedger and the evolution of a mortality index. Pension buy-outs and buy-ins still require substantial upfront costs, which longevity swaps avoid. The latter have limited upfront costs, mainly associated with set-up fees and initial margins in case of collateralization.[7]

---

[6]In the case of a plan deficit, a company borrows the amount necessary to pay an insurer to buy-out its pension liabilities in full.

[7]See Biffis, Blake, Pitotti and Sun (2012) for details on the structuring of longevity swaps, including collateral arrangements.

The limited take-up of standardized hedges such as indexed longevity swaps is probably due to (perceived) difficulties in managing longevity basis risk within broader liability driven investment (LDI) strategies.[8] There is no reason to believe, however, that the success of standardized instruments should depend exclusively on the demand of indexed hedges by pension plans and annuity providers. The huge capital inflows that have recently targeted the reinsurance and the ILSs space, as well as the natural process of aggregation of longevity exposures carried out by reinsurers and buy-out firms in the pension buy-out market, suggest that customized hedges may well co-exist with standardized instruments appealing to capital market investors. It is reinsurers and buy-out firms, rather than smaller pension plans and annuity providers, that are in a natural position to originate instruments written on large pools of exposures and to identify payoff structures that may appeal to ILS investors. The next sections aim at giving an idea of how the development of a 'life market' may come about through this channel.

3 The optimal level of securitization of longevity exposures

Securitization,[9] in general, involves the bundling together of a set of illiquid assets or liabilities (in our case mainly the latter) with a similar set of risk exposures, and then the sale of this bundled package to investors. The purpose of such a bundling exercise is to reduce the idiosyncratic (or diversifiable) risk which is contained in each of the individual assets or liabilities when they are considered in isolation, while leaving the systematic (or non-diversifiable) risk to be borne by the investor. In our case, the idiosyncratic

---

[8]Longevity risk management in LDI is discussed, for example, in Aro and Pennanen (2013).
[9]See Cowley and Cummins (2005), for an overview of insurance securitization.

risk relates to the uncertainty surrounding the length of life of individual plan members, while the systematic risk relates to the uncertainty surrounding the average length of life of the members of the group taken as a whole. We call these two types of risk idiosyncratic and systematic longevity risk, respectively. We would not expect investors to be interested in holding idiosyncratic longevity risk, any more than they would be interested in being exposed to the idiosyncratic risk from holding just a single company's shares or bonds in their investment portfolio. This is why longevity risk needs to be pooled before securitization. However, investors would be willing to invest in a particular asset class if they are appropriately rewarded for the true level of systematic longevity risk that they are assuming. So for investors to become interested in buying into this particular asset class, they need to have confidence that the premium they are being paid by buy-out firms, reinsurers, and pension plans reflects the true level of systematic longevity risk underlying the transaction. A natural way for informed sellers to persuade potential investors that this is the case is to retain some of the exposure themselves. If the seller keeps some 'skin in the game', this provides an important as well as a credible signal to investors, since longevity exposures are very capital intensive - in terms of both regulatory and economic capital requirements - for the insurance companies operating in longevity space. Investors will also realize that the larger the securitization fraction (γ) - the proportion of the total exposure that the seller seeks to offload - the greater the severity of the longevity risk exposure indicated by the seller's own assessment of longevity risk, and the lower the price the investor will offer. Since investors will try to neutralize the seller's informational advantage through a downward sloping demand curve for the longevity exposure transferred, the seller will try to securitize a fraction of exposure that is optimal, in the sense that it minimizes the cost to the seller from offloading the risk exposure. If the chosen fraction is too low, an insufficient amount

9

of the longevity risk will be transferred to investors and the seller will have to bear higher retention costs (in terms of higher capital requirements than desired). If the fraction is too high, investors will demand too high a risk premium, thus reducing the benefits of capital relief from securitization. The seller thus faces two costs, retention costs (from the capital requirements) and liquidity costs (from the downward sloping demand curve), and will wish to minimize the sum of these costs.[10]

4 The optimal tranching of longevity exposures

How should the securitized longevity exposure be transferred to investors? The seller of the exposure wants to transfer a liability, while the buyer of the exposure wants to hold an asset. How can the sale of a liability by an informed holder or intermediary become an asset to an uninformed investor? Further, we can assume that the investor will not want to be responsible for making the pension payments once the securitization transaction has taken place. To answer these questions, note that the holders of the longevity exposure are really looking for a longevity risk hedge for a fraction (γ) of their current holdings. The holders need to persuade the investors to provide that hedge in exchange for receiving an appropriate longevity risk premium. When sellers are more informed, a possible way to do this is for the holders of the longevity exposure to issue a principal-at-risk bond which the investors will buy. The return on the bond will incorporate a longevity risk premium[11] in addition to the usual credit risk premium associated with the issuer.[12] By issuing such a bond, it is the pure longevity risk that is

---

[10]This situation is formalized in the model of DeMarzo and Duffie (1999), for example.

[11]In the model of Biffis and Blake (2010a), this premium is zero as both supply and demand are risk-neutral.

[12]The bond would typically be collateralized and issued via a special purpose vehicle, a procedure that would influence the size of the credit risk premium.

transferred, not the exposure itself or the obligation to make the pension payments.[13] The securitization of longevity risk is therefore different from other forms of longevity risk transfers (such as pension buy-outs) in which assets and liabilities are bundled up and transferred in their entirety. Note that the bond-like structure discussed here is not to be confused with capital intensive hedging instruments such as the EIB longevity bond. What we have in mind is the format of catastrophe bonds,[14] where the premiums paid by investors to acquire longevity-linked notes are used to purchase securities that are held in a special purpose vehicle and used as collateral for the bond's payments.[15]

In particular, the principal-at-risk bond will repay the principal in full if the proportion of the members from the population underlying the risk transfer who have survived between the bond's issue and maturity dates is below a preset level. However, if the proportion who have survived (i.e., the survival rate) is higher than this preset level (known as the point of attachment), the amount of principal returned to the investor will be reduced according to some preset formula. If the survival rate exceeds another preset level (known as the point of exhaustion), the investor gets no principal back. The point of attachment will be determined to minimize the sensitivity of the bond's value to private information, and hence the cost to the holder of the exposure from issuing the bond.[16] Multiple attachment points can be used to meet the demand of investors with different risk preferences.

In financial engineering terms, the principal-at-risk bond is a combination of a stan-

---

[13]Should the holder of the exposure wish to do so, this responsibility could be transferred to a third party administrator who would make the pension payments.

[14]See Grace, Klein and Phillips (2001) and Lakdawalla and Zanjani (2012), for example.

[15]In the baseline model of Biffis and Blake (2010a), for example, one unit of collateral is used to collateralize a representative survival probability, S. The resulting payoff from the instrument is therefore a death rate $D = 1 - S$.

[16]It is a central tenet of corporate finance that debt-like instruments can minimize the adverse effects of information asymmetry when raising capital through the issuance of securities (e.g. Tirole, 2005).

dard bond paying (say) 1 at maturity and a call option which has been sold by the investors to the bond's issuer (in the case of a single point of attachment). The option will expire in-the-money if, on the bond's maturity date, the survival rate from the reference population (S) exceeds the preset level (S^*). The principal repayment to investors will be reduced in this case. The point of attachment is therefore the same as the strike price of the option. The option will expire worthless if S is below S^*. We can therefore write the formula for the repayment of principal as $B = 1 - \max(0, S - S^*)$, where $\max(0, S - S^*)$ means the larger of 0 and $(S - S^*)$ and is the formula for the value of a call option on its expiry date. The seller of the bond therefore has an effective hedge against systematic longevity risk for survival rates S which lie above S^*, but bears all the longevity risk below S^*.

Another way to think of this is that we have tranched the exposure at S^* with the current holder of the exposure bearing the longevity risk below S^* and investors bearing the longevity risk above S^*. The optimal tranching level will be the one that minimizes the cost to the holder of transferring the exposure, as shown in Figure 2. If S^* is set too low, the longevity risk premium demanded by investors will be excessive and the price that they will agree to pay for the bond will be too low. If S^* is set too high, the holder of the exposure bears too much of the longevity risk. A rough measure of the ex-post cost of the hedge is $B - P$, the difference between the principal actually repaid and the initial sale price of the bond. The bond issuer's objective is to choose the optimal tranching level (or, equivalently, the strike price or point of attachment) to minimize the expected cost of the longevity hedge to the issuer.

A key issue to address is the choice of reference population that determines the payments on the bond. So far, we have assumed that the bond issuer would prefer it if the reference population were the same as its own population of pension plan members

and annuitants. This would allow the bond issuer to establish a very effective hedge against its own exposure. But in a world of asymmetric information, investors who buy the bond might not feel comfortable about this, since the bond issuer will be determining the principal repayment in a manner which might be perceived to lack full transparency. Investors, however, might be more comfortable if the reference population were the same as the national population (or other publicly monitored population) since, in this case, information on survival and mortality rates are published by the country's independent official statistical agency. By doing this, asymmetric information issues are less of a concern, but a new problem emerges: using a national population index to hedge the longevity risk of a sub-group of the population introduces basis risk in the hedge. Basis risk arises because the mortality experience of the national population might differ from the sub-group being hedged. In particular, we know that members of pension plans tend to live longer than members of the population as a whole, so the national index might involve underestimating the longevity risk of the pensioner sub-group being hedged. While basis risk certainly exists, studies have shown that it can also be hedged quite effectively and so hedges and principal-at-risk bonds based on national indices can provide very good, if not 100%, hedge effectiveness (see, e.g., Coughlan *et al.*, 2011). Hedge effectiveness can be further increased if exposures are pooled or aggregated.

5 Pooling or aggregating exposures

If multiple small exposures, such as those associated with a number of small pension plans, are pooled or aggregated together,[17] there might be further diversification bene-

---

[17]buy-out and other insurance and reinsurance companies have already been acting as aggregators and may find themselves as naturally positioned to take on this intermediary role.

Figure 2: The optimal tranching level S^*.

fits,[18] although any informational advantage concerning the individual longevity trends within the separate exposures will be destroyed.[19] A principal-at-risk bond with n diversifiable exposures will have the following repayment value: $B = 1 - \max(0, S_n - S_n^*)$, where S_n^* is the optimal tranching level (or, equivalently, strike price or point of attachment) and S_n is the average exposure when there are a total of n exposures. Figure 3 shows the diversification benefits from pooling: the expected cost to the issuer of the bond and the strike price are both lower compared with the case where the exposures are tranched separately. 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 would be the case when the survival rates

[18]See Subrahmanyam (1991) and Gorton and Pennacchi (1993), for example.
[19]See DeMarzo (2005), for example.

relate to different cohorts of pension plan members, but the underlying characteristics are similar, for example, they come from the same geographical area, social class, gender or age grouping (e.g., men in their 70s). Pooling on the basis of such factors can reduce the costs associated with asymmetric information.


Figure 3: Tranching with and without pooling.

6 Conclusions

We argue that a natural format for transferring longevity risk to end investors is a principal-at-risk bond. The bond's payments need to be structured in a way that deals with the concerns that investors will have about the seller of the bond having more information about the true level of longevity risk underlying the exposure than the investor has. This problem of asymmetric information can be managed by suitably pooling and

securitizing the exposures in tranches. To give comfort to the investor that they are not being sold a 'lemon', the longevity risk hedger needs to keep some 'skin in the game', by retaining some of the exposure. To minimize the cost of the hedge, the hedger should only hedge the extreme longevity risk, that is, the longevity risk above a threshold level which we have called the optimal tranching level (or, equivalently, strike price or point of attachment). To further reduce the cost to the hedger, they should pool or aggregate their longevity exposures as much as possible. This reduces the information advantage of the seller about individual longevity trends, but it increases the degree of diversification and, in turn, might reduce the basis risk faced by hedger. Figure 4 shows the risk transfers that might characterize the stylized longevity risk market discussed in this article. Uninformed holders of longevity risk exposures, such as small pension plans, should pool and transfer their full exposures to informed intermediaries, since transferring exposures separately would expose them to 'cherry picking' from more informed counterparties.[20] Informed holders, such as large pension plans, will retain the fraction $(1 - \gamma_1^*)$ of their exposure and securitize the fraction γ_1^* by transferring it to (less) informed intermediaries. They could also deal directly with end investors by retaining a larger fraction $(1 - \gamma_2^*)$ of their exposure and securitizing a smaller fraction γ_2^*, possibly tranching it at an optimal level S^*. Finally, informed intermediaries would deal with end investors by pooling their (n say) exposures and tranching them at level S_n^*.[21] The issuance of instruments linked to public longevity indices would side step asymmetric information issues, but expose hedgers to basis risk. Even in this case, however, the principal-at-risk

---

[20]We note that Figures 1 and 4 encompass both cases of informed and uninformed sellers of longevity exposures. In the first case, the investors face the 'buyer's curse' (the risk of ending up with a 'lemon') when buying from more informed counterparties. In the second case, the uninformed holders of exposures face the 'seller's curse' (of being exposed to 'cherry picking') when selling to more informed counterparties.

[21]For a more formal analysis together with proofs of the propositions discussed in this paper, see Biffis and Blake (2010a, 2013).

format would appear a natural solution, given the familiarity of investors in ILSs with catastrophe bonds with payoffs linked to indices or parametric triggers.

It so happens that the principal-at-risk format was precisely the format chosen by Swiss Re when it issued its Vita bonds[22] to hedge the extreme mortality risk it was hedging. A similar format was adopted for Kortis by Swiss Re in 2010. Indexed on England & Wales and US population, Kortis was the first ever longevity trend bond: it would reduce payments to investors in the case of a large divergence between the mortality improvements experienced by male lives aged $75 - 85$ in England & Wales and by male lives aged $55 - 65$ in the US. So we have very successful precedents in the form of mortality/longevity trend bonds that should help promote an equally successful capital market in longevity risk transfers.


Figure 4: The market for longevity risk transfers.

[22]See "Swiss Re transfers 180 million of extreme mortality risk" (accessed 1 July 2012).

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, vol. 84(3):488–500.

Aro, H. and T. Pennanen (2013). Liability-driven investment in longevity risk management. Tech. rep., Aalto University and King's College London. Available at `http://math.tkk.fi/~haro/ldi.pdf`.

Biffis, E. and D. Blake (2010a). Securitizing and tranching longevity exposures. *Insurance: Mathematics and Economics*, vol. 46:186–197.

Biffis, E. and D. Blake (2010b). Mortality-linked securities and derivatives. In M. Bertocchi, S. Schwartz and W. Ziemba (eds.), *Optimizing the Aging, Retirement and Pensions Dilemma*. John Wiley & sons.

Biffis, E. and D. Blake (2013). Informed intermediation of longevity exposures. *Journal of Risk and Insurance*, vol. 80(3):559–584.

Biffis, E., D. Blake, L. Pitotti and A. Sun (2012). The cost of counterparty risk and collateralization in longevity swaps. Tech. rep., available at SSRN: `http://ssrn.com/abstract=1801826`.

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

Blake, D., A. Cairns, G. Coughlan, K. Dowd and R. MacMinn (2013). The new life market. *Journal of Risk and Insurance*, vol. 80(3):501–558.

Cairns, A., D. Blake, P. Dawson and K. Dowd (2005). Pricing the risk on longevity bonds. *Life and Pensions*, vol. 1(2):41–44.

Coughlan, G., M. Khalaf-Allah, Y. Ye, S. Kumar, A. Cairns, D. Blake and K. Dowd

(2011). Longevity hedging 101: A framework for longevity basis risk analysis and hedge effectiveness. *North American Actuarial Journal*, vol. 15(2):150–176.

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

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.

Gorton, G. B. and G. G. Pennacchi (1993). Security baskets and index-linked securities. *The Journal of Business*, vol. 66(1):1–27.

Grace, M. F., R. W. Klein and R. D. Phillips (2001). Regulating onshore special purpose reinsurance vehicles. *Journal of Insurance Regulation*, vol. 19(4):551–590.

Lakdawalla, D. and G. Zanjani (2012). Catastrophe bonds, reinsurance, and the optimal collateralization of risk transfer. *Journal of Risk and Insurance*, vol. 79(2):449–476.

LCP (2012). Pension buy-ins, buy-outs, and longevity swaps 2012. Tech. rep., Lane Clark & Picock LLP.

Subrahmanyam, A. (1991). A theory of trading in stock index futures. *Review of Financial Studies*, vol. 4(1):17–51.

Tirole, J. (2005). *The theory of corporate finance*. Princeton University Press.