

Pensions
Institute

DISCUSSION PAPER PI-0604

Survivor Swaps

Kevin Dowd, David Blake, Andrew G.J Cairns
and Paul Dawson

March 2006

ISSN 1367-580X

The Pensions Institute
Cass Business School
City University
106 Bunhill Row London
EC1Y 8TZ
UNITED KINGDOM

<http://www.pensions-institute.org/>

SURVIVOR SWAPS

Kevin Dowd
David Blake
Andrew J. G. Cairns
Paul Dawson

ABSTRACT

A survivor swap (SS) is an agreement to exchange cash flows in the future based on the outcome of at least one survivor index. This article discusses the possible uses of SSs as instruments for managing, hedging, and trading mortality-dependent risks. SSs are especially useful for insurance companies, but also offer other interested parties low beta avenues into the acquisition of mortality risk exposure. The article also investigates vanilla SSs in some detail, and suggests how their premiums and values might be determined in an incomplete market setting.

INTRODUCTION

Mortality or survivor risk has long been a major issue for insurance companies. If people live longer than anticipated, insurance companies make losses on their annuity books; and if people die sooner than projected, companies make losses on their life books. The need to be able to absorb these losses and remain solvent requires companies to maintain adequate levels of reserves or capital. However, capital is costly and, given the difficulties of forecasting mortality over even short horizons, it is hard to determine what adequate capital levels might be. To compound these problems, many insurance companies have suffered considerable losses in recent years, especially on stock market investments, so capital has also become more scarce. Insurance companies badly need additional, low-cost, tools to manage their mortality risk exposures.

Kevin Dowd, David Blake, Andrew J. G. Cairns, and Paul Dawson are at Nottingham University Business School, Cass Business School, Heriot-Watt University, and Kent State University, respectively. The corresponding author is Professor Kevin Dowd, Centre for Risk and Insurance Studies, Nottingham University Business School, Jubilee Campus, Nottingham NG8 1BB, UK. The author can be contacted via e-mail: kevin.dowd@nottingham.ac.uk. We thank Chris O'Brien, Sam Cox, Ralf Klett, Richard MacMinn, Krzysztof Ostraszewski, and two referees for very helpful contributions that have much improved the article. However, the usual academic caveat applies. The authors thank the Economic and Social Research Council for their financial support under grants RES-000-23-1036 and RES-000-27-0014.

An obvious answer is to use some form of mortality or survivor derivative (i.e., a financial derivative with a mortality-dependent payoff) and a number of such derivatives have recently been suggested. The first of these were “survivor bonds” (alternatively known as “mortality bonds,” or “longevity bonds”), which were proposed by Blake and Burrows (2001). The idea behind these bonds was to make their payments contingent on the proportion of a population cohort of a certain age who survive some further period, and the archetypal example was an annuity bond with coupon payments tied to a survivor index. Blake and Burrows suggested that survivor bonds have considerable potential as mortality-hedging instruments for insurance companies. However, they were also concerned that there might be insufficient supply of survivor bonds (on the grounds that there were insufficient private-sector parties willing to go short mortality improvement risk), and suggested as a remedy the controversial proposal that survivor bonds be issued by the state.¹ A more extensive discussion of these instruments was provided by Lin and Cox (2005), who set out a detailed and explicit treatment of some of the technical issues involved, especially those relating to their pricing and to the difficulties of forecasting mortality risk.

It was not long before financial institutions began to issue such securities. In December 2003, Swiss Re issued a bond whose principal payment was linked to adverse mortality risk scenarios. Under this transaction, Swiss Re set up a special purpose vehicle, which raised \$400 million from investors. The maturity of the issue was 4 years, and investors receive a floating coupon rate of U.S. LIBOR plus 135 bp. This coupon rate is generous by normal standards, but in this case, the principal payment was at risk if the weighted average of general population mortality across five reference countries (United States, UK, France, Italy, and Switzerland) should exceed 130 percent of the 2002 level. Since mortality is improving, the chances of such high mortality were judged to be very low, so investors obtained a high coupon rate in return for assuming some degree of exposure to extreme mortality risk, and Swiss Re was able to lay off some of its extreme mortality exposure.

A second bond followed in November 2004, when the European Investment Bank (EIB) issued an instrument worth £540 million. The bond issue was arranged by BNP Paribas, and involved time t -coupon payments that were tied to an initial annuity payment of £50 million indexed to the time t -survivor rates of English and Welsh males aged 65 in 2003. The EIB/BNP bond was closer in spirit to the bonds discussed by Blake, Lin, and Cox, because it tied coupon payments to a survivor index and dealt with more likely mortality risks rather than extreme ones.

However, there are other possible forms of survivor derivative, and a natural alternative is a survivor swap (SS). An SS is an agreement to exchange one or more cash flows in the future based on the outcome of at least one (presumably random) survivor index. SSSs have certain advantages over survivor bonds: they can be arranged at lower transactions cost than a bond issue and are more easily cancelled; they are more flexible and they can be tailor-made to suit diverse circumstances. They do not require the existence of a liquid market, just the willingness of counterparties to exploit their comparative advantages or trade views on the development of mortality over time.

¹ The issues involved were debated by Blake (2003) and Dowd (2003) and raised again by the Governor of the Bank of England in the 2004 British Academy Lecture (King, 2004).

SSs also have advantages over traditional insurance instruments: they involve lower transactions costs and are more flexible than reinsurance treaties, and so on. They are therefore a promising form of survivor derivative.²

SSs were discussed briefly by Blake (2003) and Dowd (2003), and in more detail by Dawson (2002) and Lin and Cox (2005). However, the most extensive analysis of them to date is provided by Lin and Cox (2005).³ In particular, they show how SSs can be priced, and provide a thorough analysis of how insurance companies might use them to exploit natural hedge opportunities across their annuity and life businesses.

This article further explores the mechanics and potential uses of SSs. It is organized as follows. “Mechanics of Survivor Swaps” defines SSs and provides an extended analysis of vanilla SSs, including how they might be “priced” in an incomplete market setting. “Uses of Survivor Swaps” discusses possible uses of SSs, including their uses as hedging instruments for insurance companies, their uses by capital market and other financial institutions to create mortality-dependent investments, and their uses as vehicles to speculate on mortality. “Warehousing and Market Making” discusses market making and warehousing issues associated with SSs, “Dealing With Counterparty Credit Risks” examines ways of managing the credit risks associated with SSs, followed by the “Conclusions.”

MECHANICS OF SURVIVOR SWAPS

Definition and Nature of Survivor Swaps

A swap is an agreement by which two parties agree to exchange one or more future cash flows, at least one of which is random. An SS can then be defined as a swap involving at least one random mortality-dependent payment.

In the most basic case, an SS would involve the exchange of a single preset payment for a single random mortality-dependent payment. More precisely, suppose that at time 0,

² Some insurance companies are already trading SSs, although the market is very much in its early stages. Concrete details are still hard to pin down, but off-the-record discussions with practitioners indicate that a number of reinsurers are trading over-the-counter (OTC) vanilla SSs in which the preset-rate leg is linked to a published mortality projection, and the floating leg is linked to realized mortality. There are also related derivatives being traded that involve the securitization of life offices’ annuity books. Typically, the reinsurers also act in syndicates to spread their exposures. As far as we can tell, the counterparties are life companies, but we know that at least one investment bank is also interested. The attractions of these arrangements are the obvious ones of risk mitigation and capital release for those laying off mortality risk, and low-beta risk exposures for those taking it on.

In addition, an article in the December 19, 2003 issue of *Grant’s* discussed curious financial arrangements known as Life Insurance and Life Assurance-backed Securities, or LILACs. These were issued by UBS and involved Warren Buffett, Hank Greenberg, and other rich, elderly, non-smoking Texans. Again, hard details are lacking, but these appear to be some kind of tontine or mortality swap promising beneficiary payments to be made to specified Texan charities when those involved die. The charities seem to have been brought into the arrangements to ensure that there is an insurance interest involved in the transaction, which is necessary to make the arrangements legal insurance contracts.

³ However, a rather different form of SS—a tax-arbitrage swap between annuities and life insurance policies—was also the subject of a recent article by Charupat and Milevsky (2001).

two firms enter into an agreement to swap a preset amount $K(t)$ for a random amount $S(t)$ at some future time t . As with a conventional forward rate agreement (FRA), $K(t)$ can be interpreted as a coupon associated with an implicit notional principal, and to keep mutual credit risks down, it makes sense for the agreement to specify that the two parties exchange only the net difference between the two payment amounts: so firm A pays firm B an amount $K(t) - S(t)$ if $K(t) > S(t)$ and B pays A an amount $S(t) - K(t)$ if $S(t) > K(t)$. $S(t)$ is related to the number of people from a specified reference population (e.g., the whole population or the number of annuity holders at time 0) who have actually survived to time t . *Ex post*, A benefits if $S(t)$ turns out to be high relative to $K(t)$ and loses if $S(t)$ turns out to be low: firm A has a long exposure to $S(t)$, while B has a short exposure to $S(t)$.

Vanilla Survivor Swaps

We can regard this basic one-payment swap as the core building block in a vanilla survivor swap (VSS), in which the parties agree to swap a series of payments periodically (i.e., for every $t = 1, 2, \dots, T$) until the swap matures in period T .

A VSS is reminiscent of a vanilla interest-rate swap (IRS), which involves one fixed leg and one floating leg typically related to a market rate such as LIBOR. However, there are several key differences. The fixed leg of the IRS specifies payments that are constant over time, whereas the corresponding leg of the VSS involves preset payments that decline over time in line with the survivor index anticipated at time 0. Also, the floating leg of the IRS is tied to a market interest rate, whereas the floating leg of the VSS depends on the realized value of the survivor index at time t . Finally, the IRS can be valued using a zero-arbitrage condition because of the existence of a liquid bond market. This is not the case with a VSS which must be valued in an incomplete markets setting.

We suggest making the preset payments equal to the product of the rates anticipated in the current mortality table for the reference population, $H(t)$, say, and a fixed proportional premium π . This suggests that a VSS would have the preset counterparty pay $(1 + \pi)H(t) - S(t)$ if this amount is positive, and receive $S(t) - (1 + \pi)H(t)$ otherwise. The premium π might be positive, zero, or negative. In line with vanilla IRSs, π would be set so that the initial value of the swap is zero to each party. Once the swap is entered into, its subsequent value would evolve over time depending on the realized survival experience of the reference population group.

A Concrete Illustration

To illustrate a VSS more concretely, suppose we have the following situation based loosely on Lin and Cox (2004, 2005). By time 0, an insurer has sold a large number (n , equal to 10,000) of annuities to a homogeneous group of 65-year-old individuals, each of which involves a commitment to pay \$1 a year to each annuitant for the rest of their lives. We now let $S(t)$ be the number of annuitants who survive to t , where $t = 0, 1, 2, \dots$. Hence, $S(0) = n = 10,000$, and $S(t)$ gradually falls over time as annuitants die off. $S(t)$ is also equal to the payments to be made in period t . The individual mortality process is binomial and each individual survives to t with probability $p(t)$, but for large n , $S(t)$ will be approximately normal with mean $np(t)$ and variance $np(t)(1 - p(t))$.

We now need a model for the mortality process $p(t)$ that takes account of aggregate mortality risk. Note that $p(t)$ is the actual survival probability *ex post*, and let $p^*(t)$ be the probability of survival to t as specified in the base mortality tables at the time the swap is arranged. To make the discussion precise, let $p(s, t, u)$ be the probability based on information available at s that an individual alive at t survives to time u . In the present context, we are particularly interested in $p(0, 0, u)$ and $p(u, 0, u)$, the former being our initial estimate of the survival probability to u and the latter being the true survival probability. We now model the progression to $p(u, 0, u)$ on a year-by-year basis by noting that this can be written as the product of 1-year survival probabilities, viz.:

$$p(u, 0, u) = p(1, 0, 1) \times p(2, 1, 2) \times \cdots \times p(u, u - 1, u) \quad (1)$$

For each t , we want the future-period 1-year survival probabilities to be subjected to i.i.d. random shocks $\varepsilon(1), \varepsilon(2), \dots, \varepsilon(t)$. We first take

$$p(0, t - 1, t) = \frac{p^*(t)}{p^*(t - 1)} \quad (2)$$

as the probability of survival from $t - 1$ to t assuming no change in the mortality table. Then for each $s = 1, \dots, t$ we let

$$p(s, t - 1, t) = p(s - 1, t - 1, t)^{1 - \varepsilon(s)} \quad (3)$$

Thus, the survival probability to time t is affected by each of the 1-year longevity shocks $\varepsilon(1), \varepsilon(2), \dots, \varepsilon(t)$.⁴

We now need to specify the distribution of the longevity shocks $\varepsilon(s)$ for $s = 1, \dots, t$. We want $\varepsilon(s)$ to fall in the range $[-1, +1]$, with $\varepsilon(s) > 0$ indicating that mortality unexpectedly improves in s , and $\varepsilon(s) < 0$ indicating that mortality unexpectedly deteriorates. More specifically, we assume that $\varepsilon(s)$ is a transformation of an underlying beta distribution. The beta is a very flexible distribution with two parameters, a and b , the values of which can be set to achieve the particular shape of distribution we desire. However, a beta-distributed random variable is constrained to the range $[0, 1]$, and this range is inappropriate for our purposes because we want to allow for the possibility of mortality deterioration ($\varepsilon(s) < 0$) as well as the possibility of mortality improvement ($\varepsilon(s) > 0$). To get around this constraint, we take a beta-distributed random variable, $y(s)$ say, and take $\varepsilon(s)$ to be the following transformation of $y(s)$:

$$\varepsilon(s) = 2y(s) - 1 \quad (4)$$

This transformation expands the range of the random shock to $[-1, +1]$ as desired. At the same time, the general shape of the $\varepsilon(s)$ distribution reflects that of the underlying

⁴ This model is similar to the forward mortality model suggested by Olivier and Jeffery (2004). However, they use a gamma random variable whereas we use the transformed beta.

beta distribution. Thus, the transformed beta has the desired range of $[-1, +1]$, but also retains the flexibility and tractability of the beta itself.

We want the underlying beta distribution to have a unimodal peak inside the range (i.e., not at its edges), and this is achieved if we set a and b bigger than 1. The mean and variance of the beta are $a/(a + b)$ and $ab/[(a + b)^2(a + b + 1)]$, respectively, so we can get $\varepsilon(s)$ to have a positive mean if we set $a > b$ and a negative mean if $a < b$. Other things being equal, we can also get the variance of $\varepsilon(s)$ to fall by increasing a and b together. In principle, the values of a and b would be determined by an appropriate stochastic mortality analysis. However, for our purposes we select three hypothetical sets of values, each reflecting an illustrative (and, we believe, not-too-implausible) stochastic scenario. The first is a scenario involving zero-aggregate mortality improvement where $\varepsilon(s)$ has a mean of zero and a standard deviation of 2.2 percent p.a.; this is obtained by setting $a = b = 1,000$. The second is a scenario of strong expected mortality improvement, where $\varepsilon(s)$ has a mean of 1 percent p.a. and a standard deviation of 2.2 percent p.a., and this is obtained by setting $a = 1,010$ and $b = 990$. The third scenario is one of strong expected mortality deterioration, where $\varepsilon(s)$ has a mean of -1 percent p.a. and a standard deviation of 2.2 percent p.a., obtained by setting $a = 990, b = 1,010$.⁵

This stochastic mortality model determines the aggregate mortality risk exposure faced by the insurer. The extent of this exposure is illustrated in Figure 1. This figure is based on the first mortality scenario (i.e., $a = b = 1,000$). This figure plots the means, and the bounds of the 90% confidence intervals of the insurer's prospective payments for each of the next 50 years. For example, if we take $t = 25$, we see that expected payments are around \$3,350, and that the 90 percent confidence interval is about [\$3,050, \$3,700]. The width of this range indicates the extent of the uncertainty faced by the insurer. In this particular instance, the width of the confidence interval is about 19 percent of the expected payment. (However, in general, we can also show that this width rises with the assumed volatility of $\varepsilon(s)$.) We can also see from the figure that the earlier and very distant confidence intervals tend to be narrow, but the intervals in the middle are wider.⁶ This tells us that most of the uncertainty attaches to payments in the middle period, and is exactly what we would expect bearing in mind that the earlier survival rates can be predicted reasonably confidently, while the most distant ones, although subject to the greatest uncertainty, will nevertheless be very low.

The insurer can hedge these payment risks almost perfectly if it now enters into a 50-year maturity VSS as the preset payer, provided that the survivor index ($S(t)$) on which the swap is based is equal to the survivor experience of the insurer's own annuity book. In this case, the SS provides an almost perfect hedge, because there is almost no basis risk in the selected survivor index.⁷

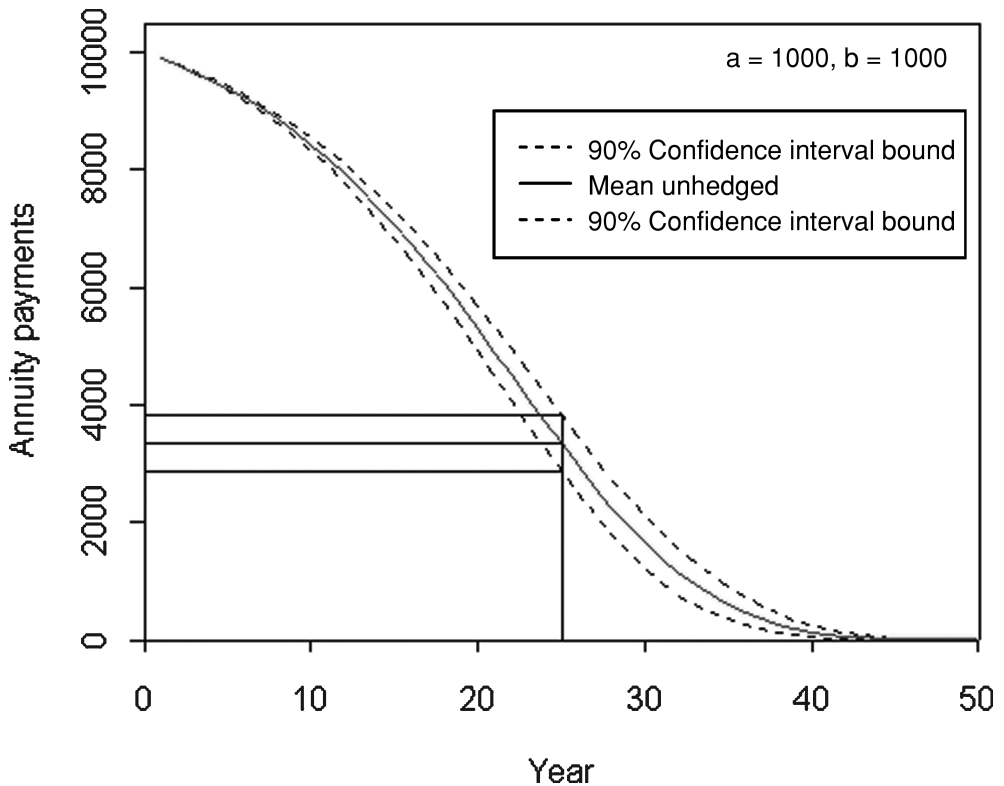
However, this hedge is only as good as it is because we have assumed that the reference index on which the swap is based is the insurer's own mortality experience. While

⁵ These and the other results reported in this article were obtained using specially written R functions, which are available on request.

⁶ Comparable figures for the other mortality scenarios are broadly similar.

⁷ There is some slight basis risk because even though $E[1 - \varepsilon(s)] = 1$, Jensen's inequality means that $E[p(s - 1, t - 1, t)^{1 - \varepsilon(s)}] \neq p(s - 1, t - 1, t)$. However, the low variance of the $\varepsilon(s)$ terms means that the degree of basis risk is very small.

FIGURE 1
Payments on an Annuity Book



Notes: Based on projected cumulative survival probabilities taken from the 1996 U.S. Annuity 2000 Basic Mortality Table, base age 65, and 10,000 Monte Carlo simulation trials. The assumptions about a and b imply that the estimated mean and standard deviation of $\varepsilon(s)$ are 0 percent p.a. and 2.2 percent p.a.

SSs could certainly be referenced on an insurer's own mortality experience, we would also expect that many swaps would involve more standardized reference indices (e.g., indices of national population mortality).⁸ In such cases, there would be some basis risk because the insurer's mortality experience would be imperfectly correlated with the reference mortality. Even so, we would usually expect reference indices and insurers' own mortality experiences to be highly correlated, and in such cases an SS can still hedge the insurer against a considerable amount of the aggregate mortality risk it faces.

⁸ One reason for the use of standardized reference indices is to ensure that reference indices are reliable, less prone to manipulation and asymmetric information problems (adverse selection, etc.), and as timely as possible. However, the development of such indices is more involved than the simple swap-rate curves that have evolved in the interest-rate swap market, but such problems are not insurmountable, and we already see parallel developments—the establishment of reliable default reference rates, and so on—in the credit derivatives markets. For more on these issues, see Lin and Cox (2005).

But how is the premium π determined? With a vanilla IRS, the terms of the swap are set by the zero-arbitrage condition implied by the spot-rate term structure: the spot-rate curve enables us to determine the values of both fixed and floating legs, and thus determine the fixed rate that makes the values of both legs equal. This fixed rate is the “price” of the swap. However, we cannot use zero-arbitrage conditions to determine the premium of the VSS because of market incompleteness. More particularly, although we can use the spot-rate curve to value the preset leg of the VSS, we cannot use it to value its floating leg.

To value the latter, we need an alternative valuation approach, and a good choice (and one that is already widely used in insurance applications) is a distortion valuation approach based on the Wang transform (Wang, 2000, 2002).⁹ If $\Phi(x)$ is the standard normal cdf, the Wang distortion operator is

$$g_\lambda(u) = \Phi[\Phi^{-1}(u) - \lambda] \quad (5)$$

where $0 < u < 1$ is a cumulative probability and the parameter λ is the market price of risk. If an instrument produces a random cash flow X with distribution function $F(x)$, then

$$F^*(x) = g_\lambda(F(x)) \quad (6)$$

can be interpreted as its “risk-adjusted” density function, and the “fair” value of the instrument is the mean of X under $F^*(\cdot)$ discounted at the risk-free rate.

We can now obtain the premium π as follows. Recall that the preset payments are $(1 + \pi)H(t)$ and the floating payments are $S(t)$. Let $V[(1 + \pi)H(t)]$ be the value of the former and $V[S(t)]$ be the value of the latter. The value of the swap to the preset-rate payer is then:

$$\text{Swap value} = V[S(t)] - V[(1 + \pi)H(t)] = V[S(t)] - (1 + \pi)V[H(t)]. \quad (7)$$

The premium π is set so that the swap value is 0. Thus, setting the left hand of (7) to 0 enables us to solve for the premium:

$$\pi = \frac{V[S(t)]}{V[H(t)]} - 1. \quad (8)$$

All that then remains is to obtain $V[S(t)]$ and $V[H(t)]$, and we obtain the former using the Wang transformation as described above¹⁰ and obtain the latter using standard present-value methods.

⁹ An alternative is to use one of the mortality securitisation frameworks suggested by Cairns, Bake, and Dowd (2005), but these are also *ad hoc* in that they too are conditional on assumed mortality processes that may or may not be empirically valid. We use the Wang transform here because it is simpler.

¹⁰ More precisely, $V[S(t)]$ is the expected value of the $S(t)$ cash flows under the Wang-transformed probability measure (5) discounted at the risk-free rate of interest.

TABLE 1
Illustrative Premiums π for Vanilla Survivor Swap

	$\lambda = 0\%$	$\lambda = 5\%$	$\lambda = 10\%$	$\lambda = 15\%$	$\lambda = 20\%$	$\lambda = 25\%$
Case 1: $a = b = 1,000$: $\varepsilon(s)$ has mean 0 and standard deviation 2.2%						
$r = 3\%$	-0.03%	-0.14%	-0.25%	-0.36%	-0.47%	-0.57%
$r = 6\%$	-0.03%	-0.11%	-0.19%	-0.26%	-0.34%	-0.41%
Case 2: $a = 1,010$, $b = 990$: $\varepsilon(s)$ has mean 1% and standard deviation 2.2%						
$r = 3\%$	5.08%	4.95%	4.83%	4.70%	4.57%	4.44%
$r = 6\%$	3.26%	3.18%	3.09%	3.01%	2.92%	2.84%
Case 3: $a = 990$, $b = 1,010$: $\varepsilon(s)$ has mean -1% and standard deviation 2.2%						
$r = 3\%$	-4.09%	-4.20%	-4.30%	-4.41%	-4.51%	-4.61%
$r = 6\%$	-2.78%	-2.85%	-2.93%	-3.00%	-3.08%	-3.15%

Some illustrative premiums are shown in Table 1. These are based on values of λ from 0 percent to 25 percent p.a.,¹¹ risk-free rates (r) equal to 3 percent p.a. and 6 percent p.a., and our three earlier sets of a and b parameter value (i.e., $a = b = 1,000$ corresponding to $\varepsilon(s)$ having a mean of 0 and a standard deviation of 2.2 percent p.a.; $a = 1,010$ and $b = 990$ corresponding to $\varepsilon(s)$ having a mean of 1 percent p.a. and a standard deviation of 2.2 percent p.a.; and $a = 990$ and $b = 1,010$ corresponding to $\varepsilon(s)$ having a mean of -1 percent p.a. and a standard deviation of 2.2 percent p.a.). Our results indicate that the premiums are relatively insensitive to the values of λ and r . However, they are very sensitive to the values of a and b (or to the $\varepsilon(s)$ parameters): with $\bar{\varepsilon}(s) \approx 0$ they are all fairly close to 0; with $\bar{\varepsilon}(s) > 0$ they are strongly positive and around 4 percent; and with $\bar{\varepsilon}(s) < 0$, they are strongly negative and around -3.5 percent. The key result is that premiums are very sensitive to the mean mortality improvement, and in a very intuitive way: if the parties to the swap operate with a mortality model that implies strong mortality improvement (deterioration), then we expect the insurer to pay (receive) a considerable premium, relative to the case where no mortality improvement or deterioration is expected.¹² In short, our results suggest that it is very important to get the “right” mean trajectory for future mortality improvements, and less important to get the “right” risk-free rate or λ .

USES OF SURVIVOR SWAPS

Uses by Insurance Companies

The previous section shows how an insurer could use a VSS to manage the risks of an annuity book. It is fairly obvious that a life insurer could also use a VSS to hedge

¹¹ This range encompasses reasonable estimates of λ . For example, using comparable U.S. data, Lin and Cox (2004a, Table 1) estimate λ to be 0.1792 for male annuitants and 0.2312 for female annuitants.

¹² The results in Table 1 indicate that the premium π falls as the price of risk λ increases. This follows because a higher λ reduces the risk-adjusted value of the floating payments receivable by the preset payer and hence lowers the premium the preset payer is willing to pay to have preset rather than floating payments. The results also indicate that π falls in absolute terms (i.e., gets closer to zero) as the risk-free interest rate increases. The explanation is that future cash flows are discounted more heavily as the interest rate increases and this reduces the premium the preset payer is willing to pay to hedge the volatility in these cash flows.

the mortality risks on its life book: the only differences between the two cases are that the life book has a long exposure to mortality improvement risk, whereas the annuity book has a short exposure, and a VSS hedge for the life book would require that the preset payments match anticipated policy payouts that would rise rather than fall over time. Both annuity providers and life companies could therefore use appropriately designed VSSs to reduce their exposures to aggregate mortality risk.

One could also imagine many other ways in which insurance companies might use SSs to manage their mortality risk exposures. An example is where annuity providers and life companies engage in SSs with each other: since the annuity providers and life companies have opposite exposures to mortality improvement risk, an SS can provide a partial mortality hedge to both parties simultaneously. As Lin and Cox (2005) explain at greater length, this implies that an SS (or in their terminology, mortality swap) can be used to help firms that run both annuity and life insurance lines of business manage the natural hedges implicit in their positions. An SS allows the firm to rebalance the exposures of the two lines of business, and so exploit the natural hedging opportunities created by the ways in which the exposures offset each other. This reduces a firm's overall mortality exposure which, in turn, allows them to reduce the risk premiums built into their product prices. SSs therefore allow firms to manage natural hedges and in so doing promote their comparative advantage.

We would also expect SSs to come in many different flavors besides plain vanilla. As with their IRS counterparts, some SSs could involve the exchange of one floating-rate payment for another. A floating-for-floating SS would be appropriate in the case just mentioned where an annuity provider swaps with a life company: one floating leg would be tied to the annuity provider's annuity payments, and the other to the life insurer's insurance payouts. We could also imagine more elaborate types of swap: swaps on mortality spreads, cross-currency SSs, SSs in which one or more floating payment depends on a nonmortality random variable (e.g., an interest rate, a stock index, etc.), and SSs with embedded features (e.g., options).

There are also other ways in which SSs can be used by insurance companies to manage their mortality exposures. Bearing in mind that swap payments would be conditioned on particular time periods and reference populations, firms can use SSs to manage their exposures across both reference populations and the "mortality term structure":

- *Managing risk exposures across reference populations.* Parties can trade risks across different reference populations. We might have a U.S. insurance company entirely exposed to U.S. mortality risk, and a UK company entirely exposed to UK mortality risk. Given that U.S. and UK mortality risks are not perfectly correlated, each company could reduce its mortality risk exposure by entering into an SS with the other. SSs therefore enable insurance companies to diversify their exposures across different reference populations, without having to sell products directly to all the populations concerned.
- *Managing risk exposures across the mortality term structure.* A company might use SSs to reduce its exposure to mortality at one or more particular horizons. For example, it can use SSs to transform a mortality term exposure that is concentrated on a particular horizon into a more balanced exposure across the mortality term structure.

SSs can also help insurance companies in other ways, and one of these is to give them greater opportunity to exploit their comparative advantages. Some firms will typically have a comparative advantage in originating business, and others in bearing the resulting exposures. For example, a firm might be particularly good at originating a specific line of business, but if it exploits this comparative advantage, it is likely to end up with more of that business than it is comfortable holding. On the other hand, if it restricts the amount of business it generates to suit its ability to bear the exposure, then it is failing to take full advantage of its ability to generate that business. A firm that is particularly well suited to bear the risks entailed by the same line of business may have the opposite problem, and be unable to generate the volume of business it desires. An SS would then provide an ideal way for both firms to alleviate these problems and play to their relative strengths.

There are some obvious parallels with the banking sector. Banks have been using securitization methods and swaps to exploit specialization economies since the early 1980s: one bank might originate a loan, another might service it, a third might warehouse it, and a fourth might fund it. Such methods have proven especially useful with standardized products (such as mortgage loans, credit card loans and auto loans), and there are clear parallels with standardized insurance products.¹³ There is also another analogy. A long-standing issue in banking is that banks frequently develop sector expertise that could (and sometimes does) lead to a dangerously large part of their loan book being concentrated in a small range of industrial sectors: specialization often leads to vulnerability. However, in recent years, credit default swaps have helped to ease this problem by providing banks with a low-cost and flexible means of exchanging sector-specific risk for risks outside their core territory. This has allowed banks to reap economies arising from their sector expertise while also achieving a greater degree of portfolio diversification than would otherwise be possible.¹⁴ There is every reason to think that insurance companies could benefit in similar ways.

Underlying these arguments is the more basic point that even where alternatives exist, SSs often offer insurance companies less costly and much more flexible ways of managing their mortality risks. Should circumstances change and the firm later wish to change its desired exposure, an SS also gives it a much more flexible means of altering its exposure.¹⁵

¹³ For more on the use of such methods in banking, see, e.g., Berlin (1992) or Brown (1992).

¹⁴ Hull (2003, p. 638 *et seq*) gives further details of such applications of credit derivatives

¹⁵ SSs also give insurance companies a useful means of avoiding legal, regulatory, and tax burdens. Imagine a U.S. insurance company that would like to take on Japanese business but operates under a constraint that prohibits it from operating abroad. Now imagine a Japanese insurance company that would like to take on U.S. business but is also prohibited from operating abroad. The two companies could use SSs to exchange some of their domestic exposures, so the U.S. company acquires exposure to Japanese mortality and the Japanese company acquires exposure to U.S. mortality. An SS would thus enable both companies to acquire surrogate business exposure overseas without breaching the legal or regulatory rules they were operating under. Similar arguments can be used to show that insurance companies can use SSs to take advantage of tax loopholes, and a good example is Charupat and Milevsky's recent analysis of the use of SSs to exploit arbitrage loopholes in the Canadian tax system (Charupat and Milevsky, 2001).

Uses by Other Investors

SSs also have major attractions to other financial institutions—such as capital market institutions, banks, and long-term investors—as means of acquiring risk exposures that would otherwise be much more difficult for them to achieve. Such institutions are always seeking to improve their risk-expected return trade-offs. Thinking in Capital Asset Pricing Model (CAPM) terms, they are looking for new investment assets with low (and preferably negative) betas measured relative to their existing portfolios. From this perspective, mortality-based assets are particularly attractive because mortality risks have low correlations with financial market risks—one imagines that these correlations must be fairly close to zero—and hence low betas relative to their existing portfolios. Hitherto, the problem for interested investors has been how to acquire such assets: mortality-dependent investment outlets are few and far between. Yet SSs fill this gap nicely. By allowing such investors to acquire mortality risk exposure, they enable them to acquire surrogate mortality assets: a bond-based investor might engage in one or more VSS, paying a preset amount and receiving a mortality-dependent floating payment, and thus convert a “straight” bond into a survivor bond; or an equity-based investor might swap some of the floating return on an equity portfolio for a floating mortality-dependent return. Since the mortality risks have a low or zero correlation with their other financial risks, such swaps would enable these firms to reduce their overall risk exposure for no little or no loss of expected return.¹⁶ Indeed, once such institutions start to see SSs as low-cost means of acquiring surrogate low beta investment assets, the market for SSs could expand to a considerable size.

Speculative Uses

SSs also have their uses as vehicles to speculate on mortality risk. Suppose a firm has a view that prospective future mortality is likely to be considerably less than is currently widely believed. The firm might be an insurance company whose in-house research indicated that mortality improvements would be greater than expected. The firm could exploit its beliefs about future mortality by entering into SSs as the preset-rate payer. If its views were correct, mortality would subsequently fall, and the payments it would receive on the floating-rate leg of the swap would rise. Furthermore, the firm would not necessarily have to wait until its prediction turned out to be correct to reap the benefit: as soon as market expectations of future mortality fell into line with its own view, the market values of SSs would be revised, and the firm would profit from the capital gain on its swap.

We can also imagine more elaborate “plays” on mortality risk. A firm might have a view that the market over-estimates mortality risk at one end of the mortality term structure relative to mortality risk at the other: for example, it might consider that the market over-estimates short-term mortality relative to long-term mortality. In this case, an appropriate trading strategy would be a mortality horizon spread: it enters into a short-horizon swap as the preset-rate receiver, and it enters into a long-horizon

¹⁶ A good example here is catastrophe reinsurance. The potential losses from catastrophes dwarf the capital of reinsurance companies, and their vulnerability to catastrophes was highlighted by events such as Hurricanes Fran in 1996 and Katrina in 2005. As with the present argument for SSs, the low correlations between catastrophe risks and other capital market risks are often presented as the principal attraction of these products for capital market participants.

swap as the preset-rate payer. Applying similar logic further, a firm might expect the shape of the mortality term structure to become more humped: if its expectations are correct, such a firm would profit from a butterfly spread in SSs.

A firm could also use SSs to speculate on differences between survivor rates among different populations. In the simplest case, a firm might expect that period- T mortality in population 1 might rise relative to period- T mortality in population 2. However, one can also imagine more elaborate speculations:

- A firm might take a spread on a period- T_1 U.S.-referenced SS and a period- T_2 UK-referenced SS: this would allow it to bet, say, that short-term U.S. mortality rises or falls relative to UK long-term mortality.
- A firm might use a spread between two butterfly positions to exploit anticipated changes in the relative shapes of two mortality term structures.
- A firm might speculate using SSs with payoffs contingent on the survivorships of three or more reference populations: for instance, a firm might take a view that one population will experience lower mortality than the average mortality of two other populations.

WAREHOUSING AND MARKET MAKING

We might expect the market for SSs to develop in much the same way as markets developed for interest-rate swaps and, more recently, for credit-default swaps. SSs would be OTC contracts with many trades going through market making intermediaries. Many early trades might be bilateral—such as U.S. and UK insurance companies swapping exposures to their respective reference populations—with no secondary markets as such. Such trades might, or might not, be assisted by third parties operating as brokers who bring the two counterparties together. However, as the market develops, brokering will increase and, as market liquidity and secondary trading grow, brokers will transform into market makers who would “stand between” trades in a true market making sense and act as intermediary counterparties themselves. The potential for market making arises from two interrelated sources.

- *The existence of warehousing economies.* These economies are much the same as those we see in interest-rate and credit-default swap markets, and arise because a market maker would be better able to economize on transactions costs and would typically have a large number of separate exposures that would offset each other or diversify. So a market maker would operate at lower costs for any given level of transactions; it would also have an overall net exposure that would be much smaller than the sum of its gross exposures, and so typically be better able to handle the exposure associated with any given trade. It could therefore offer better terms than most other potential counterparties. Hence, a firm wishing to engage in an SS is more likely to find a suitable trade with a market maker than with some other company.
- *The existence of specialist information economies.* These economies are the gains to be obtained from specialized knowledge of SS products and the SS market. A firm that engages only rarely in SSs has less incentive to acquire this information, but an intermediary that trades SSs on a daily basis is well placed to invest in such knowledge. Such a firm would benefit itself (e.g., through superior knowledge of how to structure swaps, how to manage the risks involved, how to exploit market

opportunities, etc.) and be in a position to provide better service to its clients (e.g., because more expert market makers would be able to offer better rates, etc.).¹⁷ So we would expect market makers to become specialist experts, and not just frequent traders, in the SS business.

Thus, the market for SSs would be much like that for specialist OTC derivatives today, and be dominated by market makers whose activities and willingness to trade would provide much of the market's liquidity.

DEALING WITH COUNTERPARTY CREDIT RISKS

We have hitherto ignored the issue of counterparty risk. This is a major problem with most swaps, because swaps can entail large exposures to counterparty credit risk. One very easy way to ameliorate these mutual counterparty risks is for the swap to specify that payments are to be made on a net rather than gross basis, which is why existing swaps routinely specify net payment. Achieving further reductions in counterparty exposure is more difficult, but such problems are common to OTC derivatives trades, and the standard methods of dealing with them in OTC derivatives markets could also be used to handle counterparty credit issues in SS markets as well:

- *Special purpose vehicles (SPVs)*. SS transactions can be channeled through SPVs or strongly capitalized wholly-owned subsidiaries.¹⁸
- *Credit insurance and credit derivatives*. SSs can be supplemented with a credit insurance arrangement such as a financial guaranty or surety bond: a firm might purchase such insurance to protect itself in the event that its counterparty defaults, or a firm might make itself a more attractive counterparty by purchasing insurance for its counterparty to protect the latter from loss in the event of its own default. SSs can also be supplemented by credit derivatives that promise payments if specified credit events occur (e.g., such as the default or downgrading of a counterparty). However, the reference credit events need to be chosen carefully and, as with SPVs, credit insurance and credit derivatives can be expensive.
- *Credit enhancement*. Counterparty credit exposures can also be managed using standard credit enhancement methods—collateral agreements, recollateralization with marking to market (i.e., so positions are periodically marked to market, and collateral reassessed accordingly in line with pre-agreed formulas), recouping (in which, cash is exchanged when exposures hit pre-agreed limits and payment schedules are reset to bring the swap value back to zero), credit triggers (in which a counterparty suffering a specified credit downgrade is obliged to close out its swap position and settle its outstanding debts), and mutual termination options (giving either party options to terminate a swap agreement). Each of these methods has proven to be useful in helping firms to manage the counterparty credit risks

¹⁷ The specialist expertise of market makers can also help the market in another way: market makers can use it—and it is in their own self-interest to do so—to help establish market conventions in SS markets: templates for standardized contracts, and so forth.

¹⁸ For more on the use of SPVs for such purposes, see, e.g., Culp (2002) or Kavanagh (2003).

of existing types of swap, and they are especially useful when the swaps are very long-dated ones—as would typically be the case with SSs.¹⁹

CONCLUSIONS

SSs are ideal instruments for managing, hedging, and trading mortality-dependent risks, and they offer many benefits to insurance companies that need to manage such risks.²⁰ Insurance companies can use them to hedge mortality risks, exploit natural hedges in their lines of business, diversify their mortality risk exposures, and rearrange their mortality exposures across the mortality term structure. SSs can also enhance their competitive advantage, enable them to better exploit the benefits of specialization, improve strategic flexibility, and avoid regulatory and tax burdens. They are easy to set up, but are also very flexible and can be tailor-made to suit individual users' circumstances. They also offer capital market institutions and other interested parties an easy avenue into the mortality risk business, allowing them to construct and trade synthetic mortality assets: from their perspective, SSs are very attractive because they can give them access to low beta investment assets, and so enable them to reduce their overall risk exposures for little or no loss of expected return. In short, there are good reasons to anticipate that the market for SSs—currently in its very early stages—may mature into a strong and healthy one.²¹

Of course, if SSs are so beneficial, one might ask why the SS market is still in its very early stages. This is a good question, and not a particularly easy one to answer. Part of the explanation would appear to be that the problems of survivorship risk have been largely concealed until relatively recently. For example, the long bull run in the equity markets engendered a sense of complacency in the management of long-term insurance and pension funds, and it has taken the subsequent stock market downturn to shake many institutional investors out of their earlier complacency. The need to take

¹⁹ Wakeman (1998) has more on credit enhancement methods, and on related issues such as the measurement of credit exposures.

²⁰ It is also reasonable to anticipate other developments in survivor derivatives markets, such as trading based on standardized parameters. For example, FX traders have long traded FX risk using standardized volatility measures, and then used models such as Garman-Kohlhagen to determine prices. In survivor derivatives markets, we might anticipate traders quoting π at each other, and then using, say, an International Swaps and Derivatives Association (ISDA) contract template to turn an agreed π value into a payment schedule. Indeed, if π is eventually seen as a standard reference point for market views on survivorship, then it also possible a market for π derivatives might emerge.

²¹ We might also expect to see derivatives based on SSs, such as survivor swaptions. A swaption gives the holder the right, but not the obligation, to buy or sell an underlying SS on specified terms. If the underlying swap is a VSS, the swaption might be a payer swaption, giving the holder the right to enter as the preset-rate payer, or a receiver swaption, giving the holder the right to enter as the preset-rate receiver. As with conventional swaptions, a payer swaption can be regarded as a put on survivor rates, because its value would go up when survivor rates fall, and a receiver swaption can be regarded as a call on survivor rates, because its value would increase when survivor rates rise. Survivor swaptions can be used for various risk management purposes, and an obvious use is to provide the option to lock-in future swap rates.

account of mortality risk was reinforced even more dramatically by recent failures in the insurance industry, most strikingly that of Equitable Life in the UK. These problems highlighted the hitherto hidden inadequacies of many actuarial mortality projections. At the same time, it has only relatively recently become *possible* to estimate mortality risks in a reasonably satisfactory manner: in other words, the *technology* to estimate mortality risks was not available until recently, and mortality risk markets of the type considered here are not feasible without it. So perhaps the "late" development of these markets is not so much of a puzzle after all. Again, there is a strong parallel here with credit derivatives: the credit derivatives market only took off once the technology to estimate credit risks had reached a level adequate to sustain a market, and we suspect that the same will be true of survivorship derivatives as well.

REFERENCES

- Berlin, M., 1992, Securitization, in: P. Newman, M. Milgate, and J. Eatwell eds., *The New Palgrave Dictionary of Money and Finance*, vol. 3 (London: Macmillan Press), 433-435.
- Blake, D., and W. Burrows, 2001, Survivor Bonds: Helping to Hedge Mortality Risk, *Journal of Risk and Insurance*, 68: 339-348.
- Blake, D., 2003, Reply to 'Survivor Bonds: A Comment on Blake and Burrows,' *Journal of Risk and Insurance*, 70: 349-351.
- Brown, B., 1992, Swap Markets, in: *The New Palgrave Dictionary of Money and Finance*, vol. 3 (London: Macmillan Press), 618-622.
- Cairns, A. J. G., D. Blake, and K. Dowd, 2005, Pricing Death: Frameworks for the Valuation and Securitization of Mortality Risk. *ASTIN Bulletin*, forthcoming.
- Charupat, N., and M. A. Milevsky, 2001, Mortality Swaps and Tax Arbitrage in the Canadian Insurance and Annuity Markets, *Journal of Risk and Insurance*, 68: 277-302.
- Culp, C. L., 2002, *The ART of Risk Management: Alternative Risk Transfer, Capital Structure, and the Convergence of Insurance and Capital Markets* (New York: John Wiley).
- Dawson, P., 2002, Mortality Swaps, Mimeo, Cass Business School.
- Dowd, K., 2003, Survivor Bonds: A Comment on Blake and Burrows, *Journal of Risk and Insurance*, 70: 339-348.
- Hull, J. C., 2003, *Options, Futures, and Other Derivatives* (Upper Saddle River, NJ: Prentice Hall).
- Kavanagh, B. T., 2003, An Introduction to the Business of Structured Finance, in: C. L. Culp and W. A. Niskanen, eds., *Corporate Aftershock: The Public Policy Lessons From the Collapse of Enron and Other Major Corporations* (Washington, D.C.: Cato Institute), 153-173.
- King, M., 2004, What Fates Impose: Facing up to Uncertainty, *The Eighth British Academy Annual Lecture*, December 1.
- Lin, Y., and S. H. Cox, 2004, Securitization of Mortality Risks in Life Annuities, *Journal of Risk and Insurance*, 72: 227-252.
- Lin, Y., and S. H. Cox, 2005, Natural Hedging of Life and Annuity Mortality Risks. Mimeo, Georgia State University.

- Olivier, P., and T. Jeffery , 2004, Stochastic Mortality Models. Presented to the Society of Actuaries of Ireland.
- Wakeman, L., 1998, Credit enhancement, in: C. Alexander, ed., *Risk Management and Analysis, vol 1: Measuring and Modelling Financial Risk* (Chichester: John Wiley), 255-275.
- Wang, S., 2000, A Class of Distortion Operations for Pricing Financial and Insurance Risks, *Journal of Risk and Insurance*, 67: 15-36.
- Wang, S., 2002, A Universal Framework for Pricing Financial and Insurance Risks, *ASTIN Bulletin*, 32: 213-234.

