IIJ Institutional Investor Journals

The Voices *of* Influence | iijournals.com

The Journal of Alternative Investments

Summer 2014 – Volume 17 Number 1 | www.iiJAI.com

Strategy for Increasing the Global Capacity for Longevity Risk Transfer: *Developing Transactions That Attract Capital Markets Investors*

AVERY MICHAELSON AND JEFF MULHOLLAND

Official Publication of the CAIA Association®

Strategy for Increasing the Global Capacity for Longevity Risk Transfer: *Developing Transactions That Attract Capital Markets Investors*

AVERY MICHAELSON AND JEFF MULHOLLAND

AVERY MICHAELSON
is a vice president of Insurance and Pension Solutions at SG Americas Securities, LLC, in New York, NY.
avery.michaelson@sgcib.com

JEFF MULHOLLAND
is the managing director and head of Insurance and Pension Solutions at SG Americas Securities, LLC, in New York, NY.
jeff.mulholland@sgcib.com

> *There's an iceberg dead ahead. It's called global aging, and it threatens to bankrupt the great powers. As the populations of the world's leading economies age and shrink, we will face unprecedented political, economic, and moral challenges. But we are woefully unprepared. Now is the time to ring the alarm bell…*
>
> —Peter G. Peterson [1999]

This call to action—penned by Pete Peterson, the founder of Blackstone, in his 1999 book *Gray Dawn*—refers, in part, to the vast distance between the mountains of accrued retirement liabilities in the world's developed economies and the ocean of capital that will be required to manage and ultimately fulfill these obligations. As Peterson [1999] identified, this risk threatens the stability of the global financial system and therefore demands the collaboration of all the numerous stakeholders—individuals who seek to ensure the safety of their future retirement income; employers who sponsor defined benefit pension schemes for employees; insurers who offer products that protect against longevity risk; governments who are bound to pay state and civil service pensions; and taxpayers who will ultimately be left with the tab if systems were to fail.

At the heart of this issue is longevity risk. Increases in population longevity—whether from medical advancements, healthier lifestyles, or even emerging technologies such as regenerative medicine—will dramatically increase the life expectancy, therefore the retirement needs, of the world's elderly. While having longer, healthier lives is undoubtedly a positive development for humanity, the burden of supporting the elderly in their later years will increasingly fall on the shoulders of a shrinking working-age population. In order to avoid intergenerational equity transfer (wealth transfer from younger to older generations)—which many would claim is a moral imperative—the full cost of tomorrow's retirement benefits needs to be fully capitalized and properly secured in the present (Society of Actuaries [2014]).

Achieving this imperative will require action on a number of fronts—increased contributions from retirement plan sponsors, consideration of proper retirement age and plan structure, and, perhaps most important, the use of accurate assumptions as to current rates of mortality as well as the range of possible outcomes for future longevity improvement. This largely actuarial exercise is crucial to assessing the present value of retirement obligations as well as the sensitivity of those obligations to increases in population longevity. Adequacy of these assumptions is paramount to ensuring that the financial resources put in place today—current assets and reserves against the variability of outcomes—are sufficient to meet tomorrow's obligations with a high degree of confidence.

Problematically, preparations for this impending global concern have been based on mortality rate forecasts that have consistently underestimated longevity. In fact, underestimation has been widespread and persistent: 20-year forecasts of longevity made in recent decades in Australia, Canada, Japan, New Zealand, and the United States have been too low by an average of three years (Bongaarts and Bulatao [2000]). This systematic error appears to have arisen from the assumption that currently observed rates of longevity improvement would slow down in the future. Regulatory frameworks are partly to blame, as many do not mandate the use of the most recent actual longevity data and, what's more, do not require future expected improvements in longevity to be included in calculations of pension liabilities (IMF [2012]).

QUANTIFYING GLOBAL LONGEVITY RISK

To begin quantifying global longevity risk, in order to get a sense of the order of magnitude of this issue, consider first the aggregate accrued liabilities of the developed world's retirement systems:

1. At year-end 2012, the accumulated assets of private pension systems in the OECD were $32.1 trillion, comprising: pension funds (67.9%), banks and investment companies (18.5%), insurance companies (12.8%), and employers' book reserves (0.8%) (OECD [2013]).
2. The U.S. Social Security Administration estimated the current unfunded obligation for past and current participants to be $24.3 trillion, as of the end of 2013 (SSA [2013]).
3. The aggregate liability of U.S. State Retirement Systems was an additional $3 trillion, as of the end of 2012 (Morningstar [2013]), which does not capture the liabilities of countless U.S. local and municipal pension systems.
4. Further, there are public social security systems in 170 countries (excluding the United States) that provide old-age benefits of some sort for which reliable size estimates are not readily available but which certainly amount to a substantial amount.[1]

Totaling these global obligations, we see that the accrued retirement obligations of the world's developed economies may be a staggering $60 trillion to $80 trillion!

What is more, few stakeholders adequately recognize the magnitude of the longevity risk inherent in these obligations. Each additional year of unanticipated life expectancy—roughly equivalent to a 1% increase in mortality improvement—can increase pension liabilities by as much as 4%–5% (Swiss Re [2012]). A fair estimate of the standard deviation of a sustained shock to annual mortality improvement (lasting 10 years or more) relative to expectations would be ~0.80%, and thus a longevity tail event (i.e., a 2.5 standard deviation event) would be akin to a 2.00% change in trend ($0.80\% \times 2.5 = 2.00\%$)(Risk Management Solutions [2014]). It follows then that longevity-related liabilities could grow by as much as 8%–10% as a result of unforeseen longevity improvement. Taken in the context of the $60–80 trillion aggregate retirement obligations (using current estimates of mortality), we see that liabilities could balloon by a further $5–8 trillion.

To put this into context, the total assets held by the global insurance industry (for all classes of insurance business) were estimated to be $3.66 trillion as of the end of 2013 (AON Benfield [2014]). Add to that an estimated $540 billion of capital in the reinsurance and insurance-linked securities industries, and we see that the combined capital of the insurance and reinsurance industries is barely 80% of the existing global potential for longevity risk, at the low end of the range. This drastic shortfall presents a substantial problem, as the longevity risk inherent in the world's aggregate retirement obligations is far in excess of the amount of risk capital the global insurance industry could realistically bring to bear against this risk.

Seen in this light, it becomes painfully obvious that vast sums of additional risk capital must be dedicated to adequately managing longevity risk. It is similarly evident that the only source capable of providing such quantities of capital, and thus assuming a meaningful amount of the world's longevity risk, are the global capital markets. As of the end of 2012, the amount of combined assets managed by institutional investors, globally, was $97.6 trillion dollars, roughly the same amount as the combined global bond markets, at $97.5 trillion (Accenture [2012]).

The mission is clear—longevity risk must be successfully turned into an asset class capable of attracting these vast pools of capital, or else the world's retirement systems will struggle to significantly reduce their longevity exposures in an efficient manner. However, developing capital markets solutions that are readily acceptable by a wide spectrum of institutional investors—given the complexity and uncertainty in modeling this long-term risk—requires innovative solutions from dedicated and experienced financial institutions.

THE EMERGING VALUE CHAIN FOR TRANSFERRING LONGEVITY RISK

Fortunately, a more developed global value chain is already emerging for transferring longevity risk from traditional holders of such risk—public and private pension funds—to a broader set of risk takers, including the capital markets.

At the inception of this value chain are the public and private retirement systems (detailed in the previous section) that are presently responsible for meeting the vast majority of retirement obligations. Public and private pension plan sponsors are not compensated for holding longevity risk and, in some cases, are not particularly well suited to manage it, either. Therefore, they are increasingly taking advantage of opportunities to shift liabilities off their balance sheets using a variety of transactions, called pension buy-ins, pension buy-outs and longevity swaps (collectively, "pension risk transfer contracts"). Perhaps the most highly publicized of these was the transaction between General Motors and Prudential Financial, completed in November 2012, which saw GM transfer $29 billion in assets from its pension plan to Prudential in exchange for Prudential assuming the $26 billion pension obligation attributable to the 110,000 U.S. salaried retirees covered by the transaction (Terlep [2012]). This provided a myriad of benefits, as it (a) helped to significantly reduce the economic volatility caused by the pension plan; (b) improved GM's valuation transparency; and (c) enabled GM to focus more on making cars instead of managing a pension fund (Morgan Stanley [2012]).

While the features of each type of pension risk transfer contract differ—for example who manages the dedicated pool of assets, who bears the investment risks, and who will administer the pension payments—all three transaction types are equivalent in transferring away the longevity risk inherent in retirement obligations. All of these transactions shift the longevity risk related to the specific lives in the associated pension system; as such, they are said to be "indemnity-based." Further, they are typically designed to cover any increases in the liability on a dollar-for-dollar basis above baseline assumptions; as such, they can be described as "at-the-money."

These pension risk transfer contracts are typically transacted with large life insurance companies and reinsurers that have the size, scope, and financial stability to assume such large and long-dated liabilities. However, these insurers and reinsurers ("(re)insurers") are required to hold capital (under their own internal economic capital models or regulatory capital models such as Solvency II) sufficient to cover longevity liabilities, typically at the 99% confidence level relative to potential longevity outcomes.

These capital requirements, while essential to ensuring that retirement obligations will be met, necessitate that the (re)insurers assuming the liabilities hold onerous levels of capital. This is especially true given the nature of longevity risk—a long-term trend risk that may have a relatively low probability of changing dramatically over a short period of time. Thus, the majority of capital set aside today against extreme longevity scenarios is not assumed to be needed to meet obligations in the short-to-medium term, but rather is required to be held for decades to ensure the solvency of the (re)insurers over the duration of the liability in the event of remote longevity outcomes.

Longevity risk is, thus, very capital-intensive, and therefore the amount of longevity-related liabilities that can be transferred is restricted by the quantity of capital currently dedicated by the existing risk-taker universe. To date, the binding constraint has been the (re) insurer's own economic capital models; however, over the coming years, regulatory and rating agency capital models are likely to continue moving asymptotically towards economic capital calculations.

It is in this regard—providing additional risk capital—that capital markets can be put to work. In general, capital markets will be most effective in providing capital against the most remote pieces of longevity risk, called tail risk. This can be accomplished by creating "out-of-the-money" hedges against extreme longevity outcomes featuring option-like payouts that will occur if certain predefined thresholds are breached. These hedges would be capable of alleviating certain capital requirements to which the (re)insurers are subject, thereby enabling additional risk assumption.

However, a well-constructed hedge program must perform a delicate balancing act to be effective. On the one hand, it must provide an exposure that sufficiently mimics the performance of the underlying portfolio so as not to introduce unacceptable amounts of basis risk; while, on the other hand, it must simplify the modeling and underwriting process to a level that is manageable by a broad base of investors. Further, the hedge transaction must compress the 60+ year duration of the underlying retirement obligations to an investment horizon that is appealing to institutional investors.

One strategy for accomplishing these diverse objectives is described in the following sections.

DEVELOPING AN INDEX-BASED HEDGE WITH MINIMAL BASIS RISK

As we described in the previous section, retirement systems, and the (re)insurers that assume their liabilities through pension risk transfer contracts, are exposed to the performance of a specific set of lives. While these pools may contain exposure to tens or even hundreds of thousands of lives, the underlying lives are typically from a distinct subset of the entire population. Taking the aforementioned GM transaction as an example, it would be safe to assume the lives are geographically concentrated in areas where GM has based its operations, such as the U.S. Midwest.

Any distinct population subset will have a very specific mortality profile that requires access to plan-specific mortality data and deep underwriting expertise to properly assess. This exercise is imminently achievable by a large life insurance company such as Prudential, which views this as a core competency; however, for an institutional investor, such as a university endowment, it may be a prohibitively daunting exercise. In the first place, sufficiently robust datasets about historical plan mortality data often do not exist; further, even if they did, ensuring that significant information asymmetries do not exist would be difficult (Cairns et al. [2010], Lu et al. [2013]). As such, pricing and assuming an "indemnity-based" exposure are likely beyond the abilities of even quite sophisticated capital markets investors, at present.

To address this issue, hedge programs can be designed using customized index-based exposures that reference mortality data covering the population of an entire country. National statistical reporting agencies, such as the Centers for Disease Control (CDC) in the United States or the Office of National Statistics (ONS) in the United Kingdom, publicly report the mortality rates of their populations on an annual basis by age and gender. For many developed economies, the historical mortality dataset is quite robust, often containing 40+ years of well-reported mortality statistics. Basing a transaction's payout on such data switches the pricing exercise from one that is heavily based on the underwriting of a specific population toward a statistical analysis of broad-based population mortality risk. It also ensures that all parties can access the same underlying mortality data, historically and prospectively, to reduce potential information asymmetries.

While using an index-based exposure simplifies the analysis for investors, it also introduces basis risk between the hedger's underlying book of business and the population referenced by the hedge transaction. This basis risk will reduce the hedge's effectiveness, and thus diminish the capital relief achievable (which is the hedger's ultimate goal). As such, the hedge exposure should be tailored in a way that attempts to minimize basis risk while still referencing the publicly available, national mortality rates. This is largely achievable by allowing the hedger to customize three elements of the hedge exposure:

1. The hedger selects the "cohorts," combinations of age and gender (e.g., 65-year-old males) they want in the reference exposure. Hedgers will thus limit exposure to the cohorts that are present in their underlying book of business, or a specific subset thereof that they desire to hedge. For example, a hedger may define an exposure of 70 cohorts, males and females aged 65–99, to broadly cover the lives already in retirement.
2. The hedger then stipulates the "relative weighting" of each cohort over time (an "exposure vector"). For each cohort, in each year of the risk period, the hedger will indicate the sum of the anticipated annuity payments across the underlying book to define the related exposure vector. For younger cohorts, the weighting may be zero, until the individuals reach retirement age and begin receiving payments (e.g., a 50-year-old cohort may have a zero weighting for 15 years, until they reach retirement age of 65) (see Exhibit 1).
3. The hedger then produces an "experience ratio matrix" based on an experience study of their

EXHIBIT 1
Relative Weighting of Cohorts Over Time (Exposure Vector)

Cohort	Year 1	Year 2	Year 3	....	Year 15	Year 16	Year 17	....	Year 54	Year 55
Male 65	1000	995	985	....	590	565	535	....	65	55
Male 66	980	975	960	....	505	485	450	....	45	40
....	....	....	....	....	....	....	....	....	....	....
Female 99	125	120	115	....	20	10	5	....	0	0

EXHIBIT 2
Experience Ratio Matrix

Cohort	Year 1	Year 2	Year 3	....	Year 15	Year 16	Year 17	....	Year 54	Year 55
Male 65	90%	89%	88%	....	81%	80%	80%	....	75%	75%
Male 66	89%	88%	87%	....	80%	79%	79%	....	75%	75%
....	....	....	....	....	....	....	....	....	....	....
Female 99	77%	77%	76%	....	75%	75%	75%	....	75%	75%

underlying book of business. For each cohort, in each year of the risk period, a fixed adjustment is applied to the published general-population mortality rate to adjust for anticipated differences between the mortality profile of the hedger's book of business and the corresponding reference population. For example, if the hedger's underlying lives are healthier than the general population, they may assign experience ratios of less than 100% to "scale down" the mortality rate applied in the payout (see Exhibit 2).

Customizing the exposure using these three elements should greatly reduce basis risk (a) by referencing the relevant subset of the whole population, (b) in proportion to the exposure in the underlying book of business, and (c) by further adjusting for anticipated differences in the relative mortality profiles of the two populations. This customization should reduce basis risk to tolerable levels for the hedger, while allowing the transaction payout to be determined purely in reference to publicly available general population mortality rates.

COMPRESSING THE RISK PERIOD USING A RE-PARAMETERIZED COMMUTATION FUNCTION

Most retirement systems begin making payments to retirees at around age 65 and will continue making payments until the death of the retiree (or any surviving spouse). As such, retirement obligations can be assumed to extend for 60+ years, given that a nontrivial portion of the population may live well into their 100s.

An investment horizon anywhere near this length will be categorically unappealing to the vast majority of investors, especially considering the limited liquidity that is available for longevity transactions at present. Even investors with long-term investment horizons will typically desire maturities of 15 years or shorter. Further, as the risk period lengthens, the range of outcomes for longevity experience widens due to greater opportunity for technological change (or other advances in mortality improvement).

In order to bridge the gap between investors desiring shorter maturities and hedgers looking for protection against long-dated liabilities, the hedge program can use a "commutation function." To do so, the hedger defines an exposure that covers each cohort until "maturity"—typically assumed to be age 120, or some other sufficiently remote possibility. Thus, the "exposure period" may be 60+ years; however, the "risk period" (transaction length) will be much shorter, say, 15 years.

This is accomplished by basing the final index calculations on the combination of two elements: 1) the actual mortality experience, as published by the national statistical reporting agency, applied to the exposure defined for the risk period; and 2) the present value of the remaining exposure at the end of the risk period calculated using a "re-parameterized" longevity model that takes into account the realized mortality experience over the life of the transaction.

This re-parameterization process involves:

1. Selecting an appropriate longevity risk model and establishing the initial parameterization of the model using publicly available historical mortality data that exist as of the trade date. For a basic longevity model, the parameters that may be established, on a cohort-by-cohort basis, are (a) the current rate of mortality; (b) the expected path of mortality improvement; and (c) the variability in the expected path of mortality improvement.
2. "Freezing" the longevity risk model, with regard to the related structure; but also defining, in advance, an objective process for updating the model's parameters based on the additional mortality experience that will be reported over the risk period. A determination needs to be made as to which parameters are subject to updating, as well as the relative importance that will be placed on the historical data versus the data received during the risk period.
3. Re-parameterizing the longevity model by incorporating the additional mortality data reported over the life of the trade. This occurs at the end of the transaction risk period, once the mortality data for the final year in the risk period have been received.
4. Calculating the present value of the remaining exposure using the re-parameterized version of the initial longevity model. This is done by projecting future mortality rates, either stochastically or deterministically, and then discounting the cash flows using forward rates determined at the inception of the transaction.

This approach enables a comparatively much more accurate calculation of the expected net present value of the remaining exposure at the end of the risk period. For the hedger, this should improve the hedge's effectiveness by greatly reducing the "roll risk." However, given that the structure of the longevity model is not subject to change, the payout, as a function of general population experience during the risk period, is known by investors in advance. As such, the only driver of cash flow uncertainty is the realization of general-population mortality for the relevant subset over the risk period. (See Exhibit 3).

Using a re-parameterized commutation calculation, in combination with the elements described in the previous section, should achieve the hedger's objectives of reducing tail-longevity risk, subject to tolerable amounts of "basis" and "roll" risk, while also encouraging broad-based capital market participation.

Having defined the transaction exposure and commutation function, it is now important to understand

EXHIBIT 3

Illustrative Use of a Re-Parameterized Commutation Function

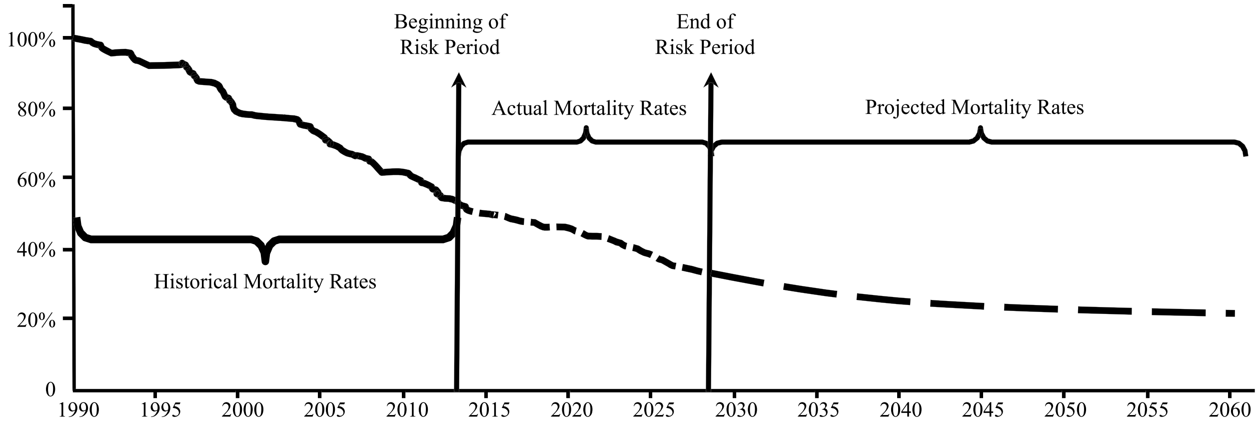

Note: Projected mortality rates are calculated using experience data available at end of the risk period.

how the transaction payout will be defined and, therefore, how capital relief is generated.

DEFINING THE TRANSACTION PAYOUT

As mentioned previously, these hedge transactions can be thought of as out-of-the-money options on future longevity outcomes. More specifically, they are "spread options" that feature two strike points—or, using reinsurance terminology, a layer of protection with an attachment point and an exhaustion point. These strikes are defined in reference to the distribution of "final index values" as calculated using a contractually agreed-upon longevity model.

The final index value will be a combination of two elements:

1. The "actual" mortality experience throughout the risk period calculated by applying the mortality rates reported by the statistical reporting agency to the predefined "exposure vector" and "experience ratio matrix," for each cohort in each year of the risk period, and accumulating with interest, using forward interest rates defined on the trade date.
2. The "commutation calculation," which captures the expected net present value of the remaining exposure at the end of the risk period, calculated using the re-parameterized version of the initial longevity model.

Once the distribution of final index values has been computed (using the agreed upon longevity model), the attachment point and the exhaustion point are selected in order to minimize the cost of capital for the capital relief taking into account market dynamics and investor preferences. The premium required by risk takers to assume a given amount of expected loss does not follow a linear pattern. Further, a "minimum premium" may be required by investors to cover the cost of committing resources to the transaction. As such, it is incumbent on those responsible for placing the risk with investors to understand the optimal layer of risk to target for the transaction relative to the hedger's strategic objectives.

Prior to implementing the hedge transaction, the hedger will be required to hold a quantifiable amount of capital to cover longevity outcomes. For example, if the "best estimate" of the longevity liability is $1 billion dollars, the (re)insurer may actually be required to hold $1.2 billion dollars, $200 million of which is reserve capital to cover the potential increase in liability due to unanticipated longevity improvement.

In this case, the (re)insurer may decide to implement a hedge transaction with a notional (maximum payout) of $100 million. This transaction would begin making a payment to the hedger in the event the attachment point is breached, paying out linearly between zero and $100 million from the attachment to the exhaustion points, and paying out the full $100 million in the event the longevity outcome meets or surpasses

EXHIBIT 4

Distribution of Final Index Values

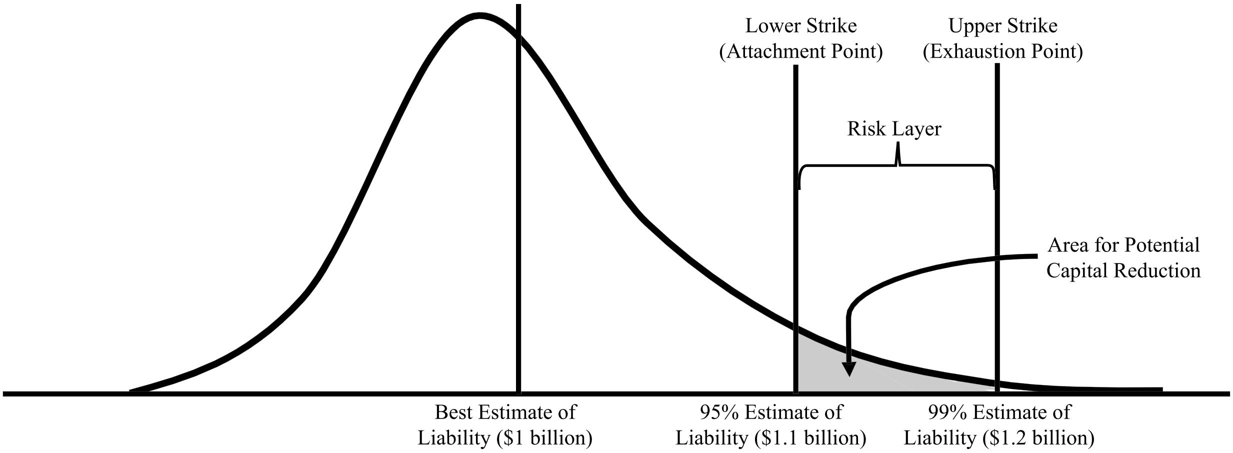

the exhaustion point. A hedge transaction designed in this manner should provide "contingent capital," thereby enabling the hedger to reduce the amount of capital it must hold (Exhibit 4).

BUILDING AN EFFICIENT MARKET FOR LONGEVITY RISK TRANSFER

To paraphrase a memorable line from the 1989 movie *Field of Dreams*: If you build it, they will come. A well-designed structure that satisfies the objectives of both hedgers and investors could lead to the development of a well-functioning market for longevity risk transfer. Specifically, longevity risk transfer instruments must offer investors an attractive expected return on capital, taking into account the correlation benefit, which will entice additional sources of capital to the asset class.

Many life insurers with limited or no exposure to longevity risk focus instead on life insurance products (which have mortality risk) or investment-oriented products with minimal longevity risk exposure. These insurers have the expertise and infrastructure to understand and manage life risks, so longevity risk is a natural extension for them. Given that these insurers have significant mortality risk, adding longevity risk will provide an offsetting exposure for some part of the distribution of mortality outcomes, with some level of correlation, and therefore can offer a high potential expected return on capital. Further, these longevity risk transfer transactions enable them to enter the longevity risk market in a scalable fashion, without incurring the costs of building the infrastructure to originate annuities or manage pension exposures directly.

More importantly, the index-based transactions we've discussed here invite participation from a broad base of capital markets investors because they package longevity risk into a format that can be understood and analyzed using the existing competencies of most institutional investors. These investors are drawn primarily because longevity risk is not correlated with the returns of other financial markets; therefore these investments may provide a great diversification benefit. When adding a small percentage of an independent (or weakly correlated) asset to a portfolio, investors should be concerned mostly with the expected return of the new asset, since the impact of its variance and its higher moments is highly suppressed. If the new asset has an expected return larger than the risk-free rate, investors can increase the Sharpe ratio of their portfolio by moving a small amount into the new asset (Goldman Sachs [1998]).

In this way, the development of an efficient market for longevity risk transfer will increase global capacity by enabling capital-constrained (re)insurers to reduce their capital requirements by securing contingent capital from a diverse set of institutional investors. This facilitates the transfer of the most capital-intensive portion of the longevity risk distribution by transferring it in a standardized format. The benefits derived from deepening the longevity risk value chain are twofold—the global capacity for longevity risk transfer is increased, and spreading this systemic risk strengthens the stability of retirement systems.

CONCLUSION

Using the Global Value Chain to Alleviate Systemic Longevity Risk

Longevity risk is a major, global, systemic risk capable of bankrupting the various institutions that are exposed to it. Unfortunately, low interest rates and recent market drawdowns have put retirement systems in a dire state. This has been exacerbated by the persistent underestimation of population longevity, resulting in existing assets and reserves that are inadequate to ensure that liabilities are met.

Relative to the size of global longevity risk, the present capital base of the insurance and reinsurance industry is insufficient to bear a meaningful portion of this significant systemic risk. The only sources of capital capable of assuming a risk of this magnitude are the global capital markets—and, as such, they must be mobilized to do so.

This "call to action" requires the creation of innovative risk transfer contracts to standardize longevity risk by packaging it into a format that can be digested by a wide base of investors. Fortunately, such solutions have already been implemented by certain pioneering investment banks and reinsurers.

For these institutions to succeed in that mission, there must be collaboration from the various stakeholders; hedgers, investors, and regulators must be educated on the risks and benefits; contracts and models must be standardized; transparency must be provided; and liquidity must be fostered. If these objectives are met, the global capacity for longevity risk transfer will increase and a critical systemic risk can be reduced substantially. This will ensure greater satisfaction of retirement obligations and the stability of financial systems in general.

APPENDIX

ADDRESSING POTENTIAL SYSTEMIC RISKS HIGHLIGHTED BY THE BASEL COMMITTEE ON BANKING SUPERVISION JOINT FORUM

In 2013, the Basel Committee on Banking Supervision published "Joint Forum—Longevity Risk Transfer Markets: Market Structure, Growth Drivers and Impediments, and Potential Risks" (BIS [2013]). The aim of this report was threefold:

1. It provided a comprehensive picture of the market for longevity risk transfer;
2. It investigated the incentives driving insurers, pension funds, banks, reinsurers, and other parties to participate in longevity risk transfer markets (or not); and,
3. It assessed the potential risks and cross-sector issues arising from longevity risk transfer, for pensioners, market participants, policymakers, and supervisors.

To this end, the report attempted to illuminate the linkages between and across firms that are created through longevity risk transfer and to analyze the potential breakdown of the risk transfer chain in stressed longevity scenarios. Recommendations were then made to promote an orderly functioning of longevity risk transfer markets, now and in the future.

According to the report, the longevity risk transfer market is not large enough to present systemic concerns yet, but its significant potential size and the growing interest from investment banks in transferring this risk make it important to ensure that these markets are safe, both on a prudential and a systemic level. In that regard, the Joint Forum highlighted four potential key systemic risks that any longevity risk transfer solution should address. Two of these have been covered extensively already—basis risk and rollover risk—while the other two are addressed in the following.

1. *Basis Risk*. The longevity risk transfer approach described in this article attempts to mitigate basis risk through the customization of the general population longevity index, using the three-step process detailed in the section "Developing an Index-Based Hedge with Minimal Basis Risk."
2. *Rollover Risk*. The approach described in this article proposes to minimize rollover risk by using a commutation calculation based on a re-parameterized version of an agreed-upon longevity risk model as detailed in the section "Compressing the Risk Period Using a Re-Parameterized Commutation Function."
3. "*Opacity*" *Risk*. Longevity risk transfer may lead to opacity risk arising due to differences in the knowledge, skills, and expertise of the buyer and seller of longevity risk. According to the Joint Forum, opacity risk likely increases with the number of links in the risk transfer chain, as the original seller of longevity risk and the ultimate buyer become more and more widely separated.
   One approach to mitigating this concern is to develop transparency and standardization in the modeling and documentation of longevity risk. As detailed previously, a transaction-specific longevity risk model should be provided to all parties to a transaction such that everyone is able to understand the risk metrics using a common framework. Each user will then be able to input their own assumptions as to the model's parameters—including projections for mortality rates and volatility around those projections.
4. *Risk Concentration/Credit Risk*. In the case of longevity risk transfer, as in the credit risk transfer markets prior to the financial crisis of 2008, risk concentration seems especially likely, given the complexity and specialized nature of these transactions. In fact, currently only a handful of (re)insurers and investment banks are active in the longevity risk transfer market.

This risk can be partially mitigated by attracting participation from a diverse set of credit worthy counterparties, as further detailed in "Building an Efficient Market for Longevity Risk Transfer." By expanding the universe of market participants, systemic concentration risk can be diffused in the initial distribution process.

Concentration risk goes hand-in-hand with credit risk, as each deals with the reliability of hedgers receiving payment when required by the hedge transaction. These risks can be further defrayed through the implementation of a dynamic collateralization approach. One such approach would be to bifurcate collateral into two categories, each of which serves a specific purpose:

1. Maintenance collateral should also be exchanged, on a two-way basis, as frequently as is operationally possible. This form of collateral is used to cover changes in the mark-to-model value of the contract.
2. An up-front collateral amount (independent amount) should be posted in relation to the credit worthiness of

the investor. It can be used in the replacement process, should an investor default, to cover price moves during the replacement period. In addition to reflecting the credit risk of the investor, it should also be scaled in proportion to the potential volatility of longevity reinsurance prices during the replacement period.

In conclusion, the longevity risk transfer methodology described in this article should materially address each of the primary concerns outlined by the Joint Forum report.

ENDNOTE

[1] Social Security Administration: http://www.ssa.gov/policy/docs/progdesc/ssptw/.

REFERENCES

Accenture Capital Markets. "Key Facts 2012: Know Your Numbers." Accenture, 2012.

AON Benfield. "Reinsurance Market Outlook—April 1, 2014 Update." April 1, 2014.

Bank for International Settlements (BIS). "Joint Forum—Longevity Risk Transfer Markets: Market Structure, Growth Drivers and Impediments, and Potential Risks." Basel Committee on Banking Supervision, 2013. http://www.bis.org.

Bongaarts, J., and R.A. Bulatao, eds. *Beyond Six Billion: Forecasting the World's Population*. Panel on Population Projections, Committee on Population, Commission on Behavioral and Social Sciences and Education. Washington: National Academy Press, 2000. www.nap.edu/catalog.php?record_id=9828.

Cairns, A.J.G., D. Blake, K. Dowd, G.D. Coughlan, and M. Khalaf-Allah. "Bayesian Stochastic Mortality Modeling for Two Populations." Pensions Institute Discussion Paper PI-1001, January 2010.

Goldman Sachs. "Analyzing Insurance-Linked Securities." Fixed Income Research October 1998.

International Monetary Fund (IMF). "The Financial Impact of Longevity Risk." In *Global Financial Stability Report—The Quest for Lasting Stability*, Chapter 4, April 2012.

Lu, J., W. Wong, and M. Bajekal. "Mortality Improvements by Socio-Economic Circumstances in England (1982 to 2006)." The Institute and Faculty of Actuaries Discussion Paper, September 2012.

Morgan Stanley. "General Motors: Pension Down, Credibility Up." June 1, 2012.

Morningstar. "The State of State Pension Plans 2013—A Deep Dive Into Shortfalls and Surpluses." 2013.

Organisation for Economic Co-Operation and Development (OECD). "Pension Markets in Focus." 2013. www.oecd.org.

Peterson, P.G. *Gray Dawn—How the Coming Age Wave Will Transform America—and the World*. New York: Random House, 1999.

Risk Management Solutions. "RMS Internal Analysis & Estimate." Privately communicated, 2014.

Social Security Administration (SSA). "2013 OASDI Trustees Report." 2013. http://www.ssa.gov/OACT/tr/2013/IV_B_LRest.html#267528.

Society of Actuaries (SOA). "Report of the Blue Ribbon Panel on Public Pension Plan Funding." Schaumburg, Illinois Society of Actuaries, 2014.

Swiss Re Europe (U.K.) "A Mature Market: Building a Capital Market for Longevity Risk 2012." Swiss Re Europe research, 2012.

Terlep, S. "GM Acts to Pare Pension Liability." *The Wall Street Journal*, June 2012.

Disclaimer

The view and opinions expressed in this article are those of the authors and do not necessarily reflect those of Société Générale or any of its affiliates. This article was authored by Avery Michaelson and Jeff Mulholland in their personal capacities and not in their respective capacities at SG Americas Securities, LLC.