A Multivariate Model of Strategic Asset Allocation with Longevity Risk

Emilio Bisetti, Carlo A. Favero, Giacomo Nocera, Claudio Tebaldi

► To cite this version:

Emilio Bisetti, Carlo A. Favero, Giacomo Nocera, Claudio Tebaldi. A Multivariate Model of Strategic Asset Allocation with Longevity Risk. *Journal of Financial and Quantitative Analysis*, 2017, 52 (05), pp.2251 - 2275. 10.1017/S0022109017000692 . hal-01633544

HAL Id: hal-01633544

<https://hal-audencia.archives-ouvertes.fr/hal-01633544>

Submitted on 13 Nov 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Multivariate Model of Strategic Asset Allocation with Longevity Risk

Emilio Bisetti, Carlo A. Favero, Giacomo Nocera, Claudio Tebaldi[☆]

Abstract

Population-wide increase in life expectancy is a source of aggregate risk. Longevity-linked securities are a natural instrument to reallocate it. This paper extends the standard Campbell and Viceira (2005) strategic asset allocation model by including a longevity-linked investment possibility. Model estimation, based on prices for standardized annuities publicly offered by United States insurance companies, shows that aggregate shocks to survival probabilities are predictors for long-term returns of the longevity-linked securities, and reveals an unexpected predictability pattern. Valuation of longevity risk premium confirms that longevity-linked securities offer inexpensive funding opportunities to asset managers.

Keywords: Longevity Risk, Strategic Asset Allocation.

JEL Classification: [G11, G12, G22]

[☆]We would like to thank B. Hollifield, E. Biffis, S. Paci for valuable comments, the participants to the WU Gutmann Center Symposium 2015: Retirement and Asset Management and the discussant M. Fischer, the participants to the International Pension Workshop Netspar 2015 Amsterdam and in particular B. Melenberg (the discussant) for useful insights, the usual disclaimer applies. Emilio Bisetti is affiliated with the Carnegie Mellon University, Tepper School of Business, Pittsburgh, PA 15213. Carlo A. Favero is Deutsche Bank Chair Università L. Bocconi and IGIER, Milano, Italy. Giacomo Nocera is affiliated with Audencia - PRES LUNAM, Nantes France. Claudio Tebaldi is affiliated with the Department of Finance at Bocconi University, IGIER and CAREFIN Milano, Italy. Corresponding author: Carlo A. Favero, carlo.favero@unibocconi.it. Department of Finance, Bocconi University, Via G. Röntgen n. 1 - 20136 - Milano.

I. Introduction

Four hours a day. This is by how much the expected lifetime for an individual aged sixty-five has been increasing over the last fifty years.¹ This phenomenon highlights an important source of risk: longevity risk. Broadly speaking, longevity risk is any potential risk attached to the increasing life expectancy of pensioners and policyholders, which may eventually translate into higher than anticipated pay-out ratios for many pension funds and insurance companies. Longevity-linked securities are instruments designed to reduce the impact of undiversifiable longevity risk on public and private balance sheets. While over the counter transactions of longevity-linked assets and liabilities are nowadays a consolidated industry practice,² a fully-fledged liquid market for longevity risk transfer is still missing and little is known about the impact of longevity risk on the financial investors' risk-return tradeoff.

This paper extends the multivariate strategic asset allocation framework of Campbell and Viceira (2005) to produce a quantitative assessment of the impact of longevity risk on the term structure of the Markowitz (1952) risk-return trade-off and on optimal investment allocations. More precisely, Campbell and Viceira (2005) estimate a vector autoregressive (VAR) model including the returns on U.S. stocks, treasury bonds and bills and a set of associated predictors; namely, the dividend-price ratio, the spread between long-term and short-term bonds and the nominal T-bill yield. In this paper, the VAR is extended to include the excess return of a synthetic financial security indexed to longevity risk annuity-

¹Life expectancy at age 65 in OECD (Organization for Economic Cooperation and Development) countries has increased from about 13 years to about 20 years in the last 42 years.

²For a recent review on the market practice see Blake, Cairns, Coughlan and Dowd (2013).

linked security (ALS) and an aggregate longevity predictor. Although the return of the longevity-linked security is not observable, it can be estimated using the observations of log-price changes for standardized annuities offered and publicly reported by North American insurance companies. The aggregate longevity shock is obtained from the Lee and Carter (1992) stochastic mortality model.

Our model estimation shows that the aggregate longevity shock is not spanned by the other variables in the VAR and that realized aggregate longevity is mean-reverting, so that it can be used to forecast longevity growth. Predictability impacts the term structure of risk that describes the dependence of the risk-return tradeoff on the investor's holding period.

Estimates from the VAR model are then used to compute optimal portfolios for investors adopting a buy-and-hold strategy with a holding period between one and forty years. The model predicted asset allocations show that investors with a long horizon (longer than 12 years) want to take a long position in the ALS security, since it has low risk and diversification properties. On the other hand, investors with a shorter horizon want to take a short position in the ALS security, using the proceeds to fund equity and bond investments. Finally, we provide a quantitative assessment of the longevity risk compensation. Its computation shows that for short- and medium-term investment horizons, the optimal (zero investment) longevity risk hedging portfolio is formed by a long position in traded financial assets financed by a short position in the synthetic longevity risk security. Evaluation of the Sharpe ratios for portfolios corresponding to different hedging policies produce reasonable risk compensation estimates. When hedging instruments are restricted to conventional fixed income securities, then risk compensation is in line with those used in the

insurance industry practice; when allocations are left unconstrained and hedges may include equity, then large Sharpe ratio estimates confirm that longevity-linked securities offer a good source of cheap leverage to risk-seeking asset managers. Quite remarkably the optimal portfolios finance equity investment shorting longevity-linked liabilities. Interestingly Frazzini, Kabiller and Pedersen (2013) estimate that 36% of Berkshire Hathaway’s liabilities consist of insurance float, thus suggesting that strategies consistent with the indication of our empirical model are already being implemented.

Our precise description of the term structure of longevity risk-return trade-offs illustrates that the creation of liquid longevity-linked securities with a stable demand critically hinges on an efficient maturity transformation activity. In fact, only in this case it is possible to diversify longevity risks among investors with shorter holding periods, who are averse to liquidity and credit risk of long-duration bonds.

The rest of the paper is organized as follows: Section II places the contribution in the context of existing actuarial and financial literature. Section III describes the construction of the aggregate longevity risk state variable and the estimation of an extended VAR that includes the aggregate longevity risk shock and annuity price changes. Section IV describes the optimal allocation for investors who have the opportunity to invest in a synthetic longevity-linked security with short duration and discusses the normative implications for the design of an efficient market for longevity risk transfer. Section V defines an hedging portfolio for aggregate longevity risk and quantifies longevity risk compensation as measured by the Sharpe ratio of this hedging portfolio. Section VI concludes. Details about the estimation of the Lee-Carter model and about the derivation of the VAR specification are reported in the Online Appendix.

II. Related Literature

Long horizon mean-variance allocations share many properties with the strategic asset allocations chosen by intertemporal utility maximizing investors, but they are easier to compute (Campbell and Viceira (2005)). In this respect, our VAR is very close in spirit to the Campbell, Chan and Viceira (2003) strategic asset allocation model. Our extension builds on Cocco and Gomes (2012), who analyze the portfolio choice problem of an agent investing in financial assets whose returns are correlated with the shocks to survival probabilities and can in turn be used to buy insurance against aggregate longevity risk. In particular, the authors study both the portfolio allocation between these bonds and risk-free assets and how their demand changes over the life-cycle depending on individual characteristics. Differently from these authors, we assume a point of view closer to the one of a reinsurance (representative) investor assuming that the idiosyncratic life cycle component is optimally managed within the insurance sector and does not play any role in the analysis. Allocation in longevity-linked securities is freely determined by an unconstrained reinsurance investor who optimizes his allocation conditionally on the risk-return performance of tradable investment opportunities.

In this respect, it is important to remark that standardized retail insurance annuity contracts that are used to estimate the VAR differ significantly from tradable financial securities, in that they are individual-specific and their purchase is irreversible. The approach pursued in this paper overcomes this problem, proving (see Section 3) that annuity log-price changes can be considered as the returns of a tradable synthetic security with a payoff that depends on aggregate longevity risk - as measured by the (ideally publicly available) index of our construction - and thus offers a stylized example of a longevity-linked security.

The distinction between annuities and these synthetic longevity-linked securities is particularly important as the valuations of financial and actuarial contracts differ significantly, as previously highlighted in the actuarial literature. From an empirical point of view Mitchell, Poterba, Warshawsky and Brown (1999), Poterba (2001), Horneff, Mitchell and Stamos (2009) among others, study the welfare benefits from purchasing annuities and discuss the well-known under-annuitization puzzle. On the theoretical side, on the other hand, the actuarially fair pricing of annuities is a well-known, investigated problem (see e.g. Pitacco, Denuit, Haberman and Olivieri (2009) and Milevsky (2006) and references therein).

Longevity risk can be decomposed in two underlying components: an idiosyncratic random variation risk and a common trend risk. Random variation risk is the risk that mortality rates differ from their expected outcome as a result of chance or individual-specific characteristics. Trend risk, on the other hand, is the risk that unanticipated changes in lifestyle behavior or medical advances significantly improve longevity for the population as a whole. Idiosyncratic risk is dealt with by pooling a large number of different individuals. Trend risk, similarly to any macroeconomic risk, is, on the other hand, an “aggregate risk” that cannot be diversified away by pooling. One path toward the reduction of the impact of longevity risk on the balance sheets of public and private insurance providers passes through the creation of a market for longevity-linked securities to enhance risk-sharing among different categories of financial investors and insurance sellers and to produce an efficient valuation of the cost of longevity risk³. In addition to idiosyncratic random

³An early proposal regarding design of financial instruments for hedging longevity risk is done in Blake and Burrows (2001).

variation risk and a common trend risk, Biffis, Denuit and Devolder (2010) point out a third source of risk affecting insurance securities: basis risk, i.e. the risk that the population from which the survival probabilities were estimated differs from the insurer's cohort. This paper analyzes only the aggregate mortality risk component; the one that is not removable through conventional actuarial policy actions (e.g. pooling, screening of subscribers) and is priced by rational agents.

The point of contact between the valuation of an annuity and that of the longevity-linked security employed in this paper is a “fair pricing” argument of widespread use in actuarial science (Milevsky (2006)): rational agents decide whether to annuitize or to defer the purchase of the annuity for a given period of time by comparing the return offered by the annuity with the potential return from self-annuitization during the deferral period. As the efficiency of self-annuitization depends on the alternative financial investments available in the market, it is reasonable to expect that an annuity contract will offer a return that is both competitive as compared to that of similar financial securities and is inclusive of a specific mortality credit component. In fact, the main goal for this modelling procedure is a precise quantification of the potential benefits that an integrated market for (aggregate) longevity risk-sharing would offer to investors and insurance providers.

The present paper assumes that individual specific risk is optimally diversified within the insurance sector by risk-pooling and that reinsurance business is carried out by financial investors who are available to trade aggregate longevity risk in order to optimize the risk-return tradeoff of their portfolio. In this respect, the present approach is complementary to the analysis of Koijen and Yogo (2013) who focus on the valuation of risks and benefits arising from concentration of shadow, off balance-sheet, reinsurance activities carried out

by insurance companies.

An approach similar to ours for the pricing of longevity risk is followed by Lin and Cox (2005) and Lin and Cox (2008), who apply the 1-factor and 2-factor Wang transform to estimate longevity premia from annuity prices. Milevsky, Promislow and Young (2006) and Bayraktar, Milevsky, Promislow and Young (2009) develop a theory for pricing undiversifiable mortality risk in an incomplete market. They postulate that an issuer of a life contingency requires compensation for this risk according to a pre-specified instantaneous Sharpe ratio. Within the model proposed in the paper, the incompleteness generated by demographic uncertainty is accounted for by including an additional state variable which is extracted from the Lee and Carter (1992) model for stochastic mortality. Previous attempts to quantify the impact of longevity risk on market prices, like Friedberg and Webb (2007), who apply the Capital Asset Pricing Model (CAPM) and the Consumption Capital Asset Pricing Model (CCAPM) to quantify risk premia for potential investors in longevity bonds, produce very low estimates of such a premium. The authors acknowledge that there is likely to exist a “mortality premium puzzle” similar to the well-known “equity premium puzzle” (Mehra and Prescott (1985)) driving higher mortality risk premia in the data than those economic models would suggest.

III. Risk and Returns in a VAR Model for Financial Securities and Annuity Prices

Our empirical strategy follows the approach to the optimal portfolio choice problem under return predictability proposed by Campbell et al. (2003) and Campbell and Viceira (2005). As a longevity-linked security is included in the investment opportunity set, an

appropriate associated predictor is built from the estimation of a stochastic mortality model and included in the VAR. We therefore first illustrate how annuity valuation implies that the unexpected generalized mortality innovation as from the popular Lee and Carter (1992) model can be used as a predictor of the return of a longevity-linked security. We then propose a VAR model of the joint dynamics of the returns on stocks, bills, bonds, longevity-linked securities and their predictors.

A. A Reduced-Form Model for Annuity Valuation

Our estimate of the longevity risk-return trade-off is based on the historical time series of observed prices for standardized annuity contracts offered by insurance companies to voluntary individual annuitants. By “standardized” annuities we mean single premium (involving a one-time investment), immediate (commencing regular income payments one period after the premium has been paid), single life (guaranteeing to make payments only to a single beneficiary until her death) and fixed (providing fixed payments) annuities.

It is important to remark that annuity contracts significantly differ from tradable financial contracts, as they are individual-specific and their purchase is irreversible. Moreover, insurance companies cannot liquidate the subscribers and annuities cannot be replicated or sold short.⁴ Finally, informational asymmetry between the subscriber and the insurance company is known to affect annuity pricing: it is documented that voluntary subscribers of life annuities live longer than average population.⁵

⁴An important element distinguishing insurance companies from several other financial intermediaries is the lack of a secondary market where the contracts written by insurance companies may be traded. The holder cannot sell her insurance policy to a third final investor, albeit, in recent years, secondary markets for some insurance contracts have developed.

⁵The adverse selection problem in annuity pricing has been discussed, among others, by Mitchell et al.

Despite these differences, it is possible to set bounds on the returns from annuities using the information on the returns offered from alternative financial investment opportunities using a simple argument. Assume that a rational agent of age x at time t faces the alternative between immediate annuitization at price $P_{x,t}$ or a deferral of the entry decision by one year, buying the annuity at time $t+1$ and age $x+1$ at a price $P_{x+1,t+1}$. In order for the agent to opt for immediate annuitization, the expected return provided by the annuity must be at least as large as the one she would expect from a financial investment with a similar risk-return profile *plus a mortality credit*, the extra return required by the annuitant as a compensation for the exposure to the (actuarial) risk of a premature death between age x and age $x+1$. If this is not the case, the agent would prefer self-annuitization, i.e. deferral of the annuitization anticipated by a short-term financial investment of the pension savings. To precisely quantify such risk, let $q_{x,t}$ denote the mortality rate for individuals of age x in year t , i.e. the probability that a person aged x and alive at the beginning of year t dies before the end of the year. We define by $s_{\bar{x}+i,t+i}$ the probability to be alive in year $t+i$, of age $\bar{x}+i$, conditional on being alive at age $\bar{x}$ in year t , so that:⁶

$$\begin{aligned} s_{\bar{x},t} &= 1, \\ s_{\bar{x}+i,t+i} &= s_{\bar{x}+i-1,t+i}[1 - q_{\bar{x}+i-1,t+i}], \quad i = 1, \dots, \infty. \end{aligned}$$

Life expectancy for a person aged x at time t is defined as $le_{x,t} = \sum_{i=1}^{\infty} s_{x+i,t+i}$. Survival

(1999) and Finkelstein and Poterba (2004). More recently, the role of informational asymmetries in longevity markets has been analyzed in Biffis and Blake (2010).

⁶The common actuarial notation for the survival probability $s_{\bar{x}+i,t+i}$ would be ${}_ip_{\bar{x},t}$. It is modified in order to keep using the common financial convention where p indicates the logarithmic price of a risky security, e.g. $p-d$ will indicate the price-dividend ratio.

probabilities tend to zero as time increases, given that mortality rates $q_{x,t}$ increase with age x , and the probability of death between x and age $x+1$ is then quantified by $1 - s_{x+1,t+1}/s_{x,t}$.

Assuming that annuities are offered to rational agents in a competitive market, prices set by insurance companies should correspond to the lowest return making the investor indifferent between immediate annuitization or deferral. Given this premise and the previous definitions, we derive an approximate accounting identity providing an explicit expression of these contributions.

By definition the one-period holding return for an annuity paying a coupon C in each period (year) to a person aged x at time t is given by:

$$(1) \quad R_{t+1}^A = \frac{(P_{x+1,t+1} + C)s_{x+1,t+1}}{P_{x,t}s_{x,t}} - 1.$$

Dividing both sides of (1) by $(1 + R_{t,t+1}^A)$ and multiplying both sides by $P_{x,t}/C$, we have:

$$\frac{P_{x,t}}{C} = \frac{1}{(1 + R_{t,t+1}^A)} \left(\frac{s_{x+1,t+1}}{s_{x,t}} \right) \left(1 + \frac{P_{x+1,t+1}}{C} \right).$$

Denoting with lowercase letters the natural logarithms of uppercase letters we have:

$$p_{x,t} - c = -r_{t+1}^A + \ln \left(\frac{s_{x+1,t+1}}{s_{x,t}} \right) + \ln (1 + e^{p_{x+1,t+1} - c}).$$

Finally, taking a Taylor expansion of the last term about the average log price-coupon

ratio, $\frac{P}{C} = e^{\overline{p-c}}$, we have:

$$p_{x,t} - c \simeq -r_{t+1}^A + \ln \left(\frac{s_{x+1,t+1}}{s_{x,t}} \right) + k^s + \rho (p_{x+1,t+1} - c),$$

where $\rho \equiv \frac{e^{\overline{p-c}}}{1+e^{\overline{p-c}}}$. Therefore, recalling that:

$$\ln \left(\frac{s_{x+1,t+1}}{s_{x,t}} \right) = \ln (1 - q_{x,t}),$$

we derive by rearranging that annuity prices can be written as:

$$p_{x,t} = k + (1 - \rho) c + \rho (p_{x+1,t+1}) - r_{t+1}^A + \ln (1 - q_{x,t}).$$

Taking the term-by-term difference of the valuation equation between time $t + 1$ and t , the value of the coupon (nominal payment is fixed) disappears from the valuation equation, so that:

$$(2) \quad \Delta p_{x,t} = \rho (p_{x+1,t+1} - p_{x+1,t}) - \Delta r_{t+1}^A + \Delta \ln (1 - q_{x,t}),$$

with $\Delta p_{x,t} \equiv p_{x,t} - p_{x,t-1}$.

Consider now the popular Lee and Carter (1992) model for stochastic mortality. This model has both strong within-sample fitting properties and remarkable out-of-sample predictive power. Together with the relative ease of its computation, these characteristics have made it the standard mortality forecasting model among practitioners and academics. The

model consists of a system of equations for the logarithm of the central death rate $m(x, t)$ of each age cohort x at time t , and a time-series equation for an unobservable time-varying mortality index k_t , common among all age cohorts. In particular, in terms of the mortality rate $q_{x,t}$ we have

$$(3) \quad \ln(q_{x,t}) \simeq -(a_x + b_x k_t + \epsilon_{x,t}),$$

$$(4) \quad k_t = c_0 + c_1 k_{t-1} + e_t,$$

$$\epsilon_{x,t} \sim \text{NID}(0, \sigma_\epsilon^2), \quad E[e_t] = 0,$$

where a_x and b_x are age-specific constants.⁷ The error term $\epsilon_{x,t}$ captures cross-sectional errors in the model-based prediction of mortality for different cohorts, while the error term e_t captures random fluctuations in the time series of the common factor k_t driving mortality at all ages. This common factor k_t evolves over time as an auto-regressive process and the favorite Lee-Carter specification makes it a unit-root process by setting $c_1 = 1$ while e_t is assumed to be a zero mean stationary process. Identification is achieved by imposing the restrictions $\sum_t k_t = 0$ and $\sum_x b_x = 1$, so that the unobserved mortality index k_t is estimated through Singular Value Decomposition.⁸

Under the Lee-Carter specification the revision of the mortality credit contribution is

⁷In the original formulation Lee and Carter assume an exponential affine expression of the central death rates, while in our framework the exponential affine expression is applied directly to the mortality rates. The two coincide within the log-linear approximation of the mortality rates evolution that is considered in the VAR specification. Unreported estimation tests do show that results are essentially unaffected if the original specification is adopted.

⁸See Appendix A for a full description of the adopted identification and estimation strategy.

linear in the innovation to the unobserved common factor component k_t . In fact,

$$1 - q_{x,t} = 1 - \exp(-[a_x + b_x k_t + \varepsilon_{x,t}]), \quad k_t = c_0 + k_{t-1} + e_t,$$

so that for small variations in mortality rates:

$$\ln(1 - q_{x,t}) - \ln(1 - q_{x,t-1}) \simeq \Delta q_{x,t} = -[b_x e_t + \varepsilon_{x,t} - \varepsilon_{x,t-1}].$$

A similar measure has already been introduced and discussed in Friedberg and Webb (2007).

Assuming that agents will compensate only aggregate risk, the priced contribution to mortality credit is given by e_t , and we can re-write (2) as

$$(5) \quad \Delta p_{x,t} = \rho(\Delta p_{x+1,t+1}) - \Delta r_{t+1}^A + \Delta q_{x,t}.$$

Solving this relation forward up to period $t + m$ and taking expectations given the information set available at time t , we have:

$$\Delta p_{x,t} \simeq - \sum_{j=0}^m \rho^j E_t \Delta r_{t+1+j}^A + \sum_{j=0}^m \rho^j E_t \Delta q_{x+j,t+j},$$

which shows that the annuity log price change is determined by future expectations on changes Δr_{t+k}^A in the annuitant holding period returns and on revisions in trend longevity factor $\Delta q_{x,t}$.

It becomes useful at this point to observe that the annuity log price change can be

interpreted as the return on a longevity-linked security whose payoff is indexed to the log price variation $\Delta p_{x,t}$ of annuity prices sold each year t to annuitants in the cohort x . By convention we define Annuity-Linked Security (ALS hereafter) as a reference synthetic security indexed on annuity prices based on the cohort of 65-year-old US annuitants. A long position in the ALS corresponds to an (nontraded) tontine insurance in which contracts are terminated and then possibly renegotiated every year.⁹ A short position in the ALS allows the investor to sell protection on the longevity risk of the cohort of 65-year-old US annuitants. Then the last equation identifies a relation between the return on ALS, the innovations in the common mortality factor in the Lee-Carter model and the annuitant holding period returns.

B. A Model for Stochastic Mortality and its Performance on the U.S. Data

We apply the Lee and Carter (1992) model to estimate shocks to mortality for cohorts in the age interval between 65 and 110. We restrict the estimation to the cohorts of the retired population for several reasons. First, the active rebalancing of the contributions is not feasible for these cohorts, as they typically consist of people who have left the accumulation phase and entered the decumulation phase. Hence, reallocation via securitization or reinsurance is the only viable strategy that insurance companies can pursue to hedge the associated longevity risk. Second, the largest publicly available empirical data sets on annuity prices apply to annuitants belonging to these cohorts. Third, limiting the specification to the retired cohorts alleviates some well-known limitations of the Lee-Carter model when applied to all cohorts (see Lee (2000)). Finally, the Lee-Carter model, being a single-factor model of mortality, minimizes the number of parameters to be estimated and produces a

⁹For an actuarial discussion of such synthetic contracts see Milevsky (2006), p. 224.

parsimonious set of mortality predictors. In fact, the Lee-Carter stochastic mortality model produces a mortality predictor that fits nicely the log-normal approximation conventionally applied to model prices and predictors within the VAR approach. Note, however, that the approach proposed is fully flexible and we do not see specific obstacles to extend it to any other (log-linearized) version of stochastic mortality models as, for example, the one proposed by Blake, Cairns and Dowd (2006). In Appendix A we report evidence on the performance of the Lee-Carter model in fitting US mortality rates. To derive an observable counterpart of the priced contribution to mortality credit, we consider an estimation of the Lee-Carter model restricted to the post-retirement cohorts of individuals aged between 65 and 110. This measure, called qk_t , is similar to the one discussed in Friedberg and Webb (2007), and coincides with the unexpected variation in the survival rate pooled over all retired cohorts:

$$qk_t \simeq \frac{1}{45} \sum_{x=65}^{110} \{\ln(1 - q_{x,t}) - E_{t-1}[\ln(1 - q_{x,t})]\}.$$

Using the Lee-Carter specification, the observable index qk_{t+1} can be approximated by:

$$qk_t \simeq -\frac{1}{45} \sum_{x=65}^{110} (\alpha_x + b_x k_t + \varepsilon_{x,t}) - (\alpha_x + b_x E_{t-1}[k_t]) = -\frac{1}{45} \sum_{x=65}^{110} (b_x e_t + \varepsilon_{x,t})$$

and, taking into account the normalization condition $\sum_{x=65}^{110} b_x = 1$ we have:

$$(6) \quad qk_t \simeq -\frac{1}{45} \sum_{x=65}^{110} (b_x e_t + \varepsilon_{x,t}) = -\frac{e_t}{45} - \frac{1}{45} \sum_{x=65}^{110} \varepsilon_{x,t}.$$

The longevity shock qk_t describes the time evolution of the unexpected variation in mortality rates which has a uniform impact across cohorts, and is estimated by applying the Lee and Carter model only to the retired cohorts. Notice that this shock includes two contributions; the first contribution is proportional to the opposite of the Lee Carter aggregate mortality shock e_t , the second one is the sum of the cohort-specific innovations. The filtered innovation qk_t is included as a predictor in the vector of autoregressive variables to account for unexpected trend variation in mortality rates. This variable offers a publicly available, cohort-independent information which investors can observe and use to quantify the impact of the variability of aggregate longevity on prices.

Notice that the second contribution accounts for longevity fluctuations that are non-uniform across cohorts, e.g. an increase of mortality rates in older cohorts compensated by a reduction of mortality at younger ages. It is interesting to analyze the dynamic properties of the second contribution. By construction it has zero expectation but, as originally highlighted by Lee and Carter, $\varepsilon_{x,t}$ are generally correlated across cohorts and across time. Importantly, we retain only the testable assumption that this second contribution is stationary: the only non-stationary component in the evolution of mortality is the common factor k_t .

C. A Reduced VAR Dynamic Model for the Annuity-Linked Security Returns

In this subsection we show that, using the above reduced valuation approach, it is possible to model the stochastic evolution of ALS returns (log annuity price changes) $\Delta p_{x,t}$ using a VAR specification which extends that of Campbell and Viceira (2005) (hereinafter CV-VAR(1)). Following Barberis (2000), Campbell and Viceira (2002), we describe dynamics of asset returns and relevant predictors using a VAR(1) model:

$$(7) \quad \mathbf{z}_t^{\text{Mkt}} = \Phi_0^{\text{Mkt}} + \Phi_1^{\text{Mkt}} \mathbf{z}_{t-1}^{\text{Mkt}} + \nu_t^{\text{Mkt}},$$

where

$$\mathbf{z}_t^{\text{Mkt}} = \begin{bmatrix} r_0 & x_t^{\text{Mkt}} & s_t^{\text{Mkt}} \end{bmatrix}^T$$

is a $m \times 1$ vector, with r_{0t} being the log real return on the asset used as a benchmark to compute excess returns on all other asset classes, $\mathbf{x}_t$ being the $n \times 1$ vector of log excess returns on all other asset classes with respect to the benchmark, and $\mathbf{s}_t$ being the $(m - n - 1) \times 1$ vector of returns predictors. The exact specification and its estimation results are reviewed in online Appendix B.

Although an annuity is not a financial security and cannot be priced accordingly, rationality of the annuitant forces the (log) holding period return r_t^A to be comparable (but for the mortality credit) to the compensation one would receive by investing in a portfolio of traded financial securities with similar risk and return characteristics while deferring by one year the annuitization. Hence we claim, and later show empirically, that the financial component of the return r_t^A can be replicated using a portfolio of securities whose evolution is described by the CV-VAR(1) model. Moreover, assuming a stationary evolution for r_t^A , the VAR(1) specification implies that also $-\sum_{j=0}^m \rho^j E_t \Delta r_{t+1+j}^A$ is a linear function of the state variables $\mathbf{z}_t^{\text{Mkt}}$. This is equivalent to the assumption:

$$-\sum_{j=0}^m \rho^j E_t \Delta r_{t+1+j}^A = \phi_0^A + \phi_1^{A, \text{Mkt}} \mathbf{z}_t^{\text{Mkt}}.$$

Similarly, the mortality credit component affecting the evolution of ALS price is proxied by the observed innovations in the longevity trend. For this reason an additional predictor state variable qk_t is introduced and it is assumed that:

$$\sum_{j=0}^m \rho^j E_t \Delta q_{x+j,t+j} = \phi_3 qk_t.$$

In conclusion, the standard CV-VAR(1) model can be augmented to include the evolution of the (65 cohort) annuities' (log) price growth in excess to the return of the safe asset, $x\Delta p_t := \Delta p_{65,t} - r_{f,t}$, following the specification:

$$x\Delta p_{t+1} = \phi_0^A + \phi_1^{A,\text{Mkt}} \mathbf{z}_t^{\text{Mkt}} + \phi_2 x\Delta p_t + \phi_3 qk_t + \nu_{t+1}^A,$$

where ν_{t+1}^A is the combination of all shocks in the state variables and idiosyncratic mortality shocks. Excess returns of the ALS are determined by a combination of market returns, market return predictors and the aggregate longevity predictor. We now analyze the effect of extending the traditional portfolio to include excess annuity prices by considering the following augmented VAR specification:

$$(8) \quad \mathbf{z}_t = \Phi_0 + \Phi_1 \mathbf{z}_{t-1} + \nu_t, \quad \nu_t \sim \mathcal{N}(0, \Sigma_\nu),$$

where

$$\mathbf{z}_t = \begin{bmatrix} r_0 & x_t^{\text{Mkt}} & x\Delta p_t & s_t^{\text{Mkt}} & qk_t \end{bmatrix}^T$$

and Σ_ν is the $(m + 2) \times (m + 2)$ variance-covariance matrix of the returns on financial assets, the annuity prices and their associated predictors.

D. The Dynamics of Returns of U.S. Bonds, Bills, Stocks and Annuities

To evaluate how the inclusion of the ALS excess returns and of a predictor for the change in their prices modifies the optimal portfolio allocation at different horizons, we compare the results obtained from the CV-VAR(1) estimation over the yearly sample 1953-2010 (the most recent update of the human mortality database) to those obtained from our extended VAR. The first model includes six variables: the ex-post real T-bill rate, the annual excess returns on stocks, the annual excess returns on long-term (20-year) bonds, the log yield on a 90-day T-bill, the log dividend-price ratio and the yield spread (defined as the difference between the 20-year zero-coupon bond yield from the CRSP Fama-Bliss data file - the longest maturity yield available in the file - and the T-bill rate). The second model is an eight-variable VAR obtained by adding to the standard CV-VAR(1) the log difference in the annuity premium minus the risk-free rate (which extends the set of excess returns) and the aggregate longevity shock (which extends the set of predictors). Table 1 shows sample statistics for all variables. Our sample, which includes observations up to the most recent update of mortality data, compares well with the annual sample used in previous studies. Only the statistics on long-term bond indicate a lower expected return, a result which is clearly driven by recent trends in monetary policy.

Insert Table 1 about here

Both estimated VARs include constants in each equation. Table 2 shows the results for the extended VAR including longevity predictor and ALS returns, while the results of the

original CV-VAR(1) estimation over the yearly sample 1953-2010 and detailed information on the data used in the estimation are reported in Appendix C. The results from the standard model are well in line with those reported in Campbell et al. (2003). When the extended VAR is estimated, no major changes take place in the coefficients attached to the six financial variables in the original model. For those variables whose explanatory power is significantly different from zero, the impulse response coefficients are qualitatively similar and confirm all the stylized properties found in the original estimation: real T-bill, stock and bond returns are predicted by nominal short rate, dividend-price and term spread. The longevity shock is persistent and helps to predict the change in annuity prices (a positive shock to longevity increases the price of annuities). It also has some significance in predicting excess return on bonds. Although it is not possible to identify the structural origin of this predictive relation within our reduced partial equilibrium model, it is in line with some macroeconomic arguments supporting a relation between population age structure and price fluctuations of long duration safe assets.

Insert Table 2 about here

The two new equations included in the extended VAR describe the evolution of the aggregate longevity shock and the logarithmic yearly change of the annuity price in excess to the nominal T-bill. The estimated aggregate longevity shock dynamics qk_t is substantially a univariate mean-reverting process with a root of -0.75 and zero expectation (estimation provides a value $E_{t-1}[qk_t] = 0.0005$), indicating that the information conveyed by the aggregate longevity shock is not spanned by other variables and that the spread between expected and realized aggregate longevity is mean-reverting and thus may be used

to forecast aggregate longevity growth.

Notice that the persistence of this shock does not contradict the standard Lee-Carter specification; in fact the usual specification assumed for the trend component k_t is a unit root evolution process. This requires only the first difference of the process to be a stationary process, which is consistent with our findings. In addition, as shown in Figure 1, a closer analysis of the contributions to qk_t highlights that mean reversion is almost entirely driven by the second contribution to qk_t in eq.(6), i.e. by the sum of cohort specific shocks. These fluctuations measure the time t deviation from zero of the cohort-specific shocks sample mean; the longer the persistence of this process, the more slowly they return to the mean. Notice that from a demographic point of view these fluctuations measure also the deviations from the original expected distribution of the cohort specific mortality rates which, especially at older ages, are driven by health and life quality improvements. As should be expected, we verify in sub-samples that this mean reversion is slowly increasing as the time span of the sample increases.

Insert Figure 1 about here

The aggregate longevity shock is a significant predictor for $x\Delta p_{t+1}$ (the ALS return in excess to the T-Bill rate), which is also significantly predicted by past real and nominal T-bill rates and the excess returns on long-term bonds. Annuity price growth in excess of the T-bill rate have a positive loading on the real rate and a negative loading on the nominal T-bill rate, a negative loading on long-term excess bond rate returns and a positive dependence on the aggregate longevity shock. In Figure 2 the time series of historical (real) logarithmic price changes is compared to the replication as operated by the VAR dynamic

model aggregating the information of financial securities returns and of the forecasting variables, including the aggregate longevity shock.

Insert Figure 2 about here

The good fit indicates that the VAR estimation produces a realistic “reduced-form” pricing model for the annuity contract offered by insurance companies to annuitants. Note that our approach concentrates only on aggregate longevity risk and does not account for the actuarial components of insurance premia. These have to be included to hedge basis risks or adverse selection effects, which would in turn require a discussion of the specific characteristics of annuitants.

IV. The Impact of Longevity Securitization on Optimal Allocations

The possibility to trade longevity-linked securities extends the set of investment opportunities and offers a new diversification dimension. As in Campbell and Viceira (2005), optimal portfolios are computed for investors adopting a buy-and-hold strategy with holding period between 1 and 40 years. The set of investment opportunities is composed by T-bills, equity, a rolling strategy in a long-term bond and the ALS. This security grants to its holder a yearly return equal to $\Delta pr_{65,t}$ the variation of the mean (logarithmic) price observed on the US insurance market for a standardized annuity contract. The estimation indicates that, as expected, the return from this contract will rise if aggregate longevity is rising and will decrease if aggregate mortality increases. As in Campbell and Viceira (2005), the VAR is used to derive the term structure of risk while expected returns are estimated unconditionally: the VAR framework is not exploited to derive a time-varying

asset allocation that depends on conditional first and second moments. From this perspective, all that is required in order to derive the same term structure of risk over time is the structural stability in the relation between predictors and returns. In practice, we have checked that the term structure of risk does not vary by comparing our full sample results with those obtained by reducing the sample via the omission of the first twenty years¹⁰. Note however, that optimal allocations do vary across the full and the limited sample as a consequence of the fact that unconditional mean of returns are different in the two periods. The VAR is never used to forecast first moments; therefore, the difference between within-sample performance and out-of-sample performance will not have a significant impact on our VAR-based portfolio allocation. The results of a more extended robustness check are reported in Appendix C.

A term structure of conditional volatilities at different horizons can be naturally derived from the estimation of our VAR process for returns and predictors.

Insert Figure 3 about here

In Figure 3, we compare the term structures of the standard deviation of the ALS and of traditional financial securities. Notice that the annualized volatility increases with the holding period, the clear sign of the long-term nature of the risks underlying annuity prices. For holding periods shorter than 10 years, ALS is less risky than equity and (rolling) bond investments, while on longer horizons its risk exceeds that of other securities: the irreversible nature of annuitization implies that, from a pure financial point of view, this contract has a risk-return profile similar to the one by a buy-and-hold strategy on a long-term bond. ALS price fluctuations reflect changes in the long-run expectations regarding

¹⁰Results of this robustness exercise are available upon request.

inflation and longevity trends. A small persistent change to future expectations can have a relevant impact on the current evaluation of the annuity contract. Note that an ALS is not the contract that an insurance company would like to use to reinsure aggregate longevity risk. In fact, an insurance company with a portfolio of annuities under management would be willing to reinsure only aggregate longevity risk but would prefer to retain the remunerative core business of the insurance industry through the diversification of the residual risk by pooling together annuitants of all ages and different lifestyles.

Figure 4 plots the correlations among the set of financial securities and ALS as a function of the holding period.

Insert Figure 4 about here

These correlations have a sharp decline with the holding period. These correlations, and more in general the term structure of assets' risk and returns, determine the weight that each asset receives in the portfolio allocation of an investor with mean-variance preferences for any given horizon. To have a sense of how this portfolio allocation changes by including the possibility to invest in the ALS, we again follow Campbell and Viceira (2005) and first consider the generalized absolute minimum variance portfolio (henceforth GMV), the portfolio with the lowest variance on the mean-variance efficient frontier. For each holding period this portfolio is described in Figure 5.

Insert Figure 5 about here

Figure 5 shows that the investor overweights the allocation in the T-bill to buy a combination of ALS and long-term bond independently of the investment horizon. This combination is a long position in the bond and a short one in the ALS when the holding

period is below 10 years, while for longer horizons the two positions are switched. Hence, over periods of time less than a decade, risk exposure is minimized by selling protection to longevity while the same investment becomes speculative over longer holding periods. As a consequence, it is expected that demand for longevity exposure and the liquidity of longevity-linked securities can be considerably increased by offering products with short durations.

Figure 6 compares the term structure of risk of the extended GMV with that of the “Campbell-Viceira” GMV and that of a T-bill. Inclusion of the ALS reduces risk for all holding periods with the exception of the interval between 10 and 15 years, where the optimal allocation in the ALS shifts from negative to positive values and the rolling position in long-term bonds shifts from positive to negative.

Insert Figure 6 about here

Roughly speaking, the allocation and the risk profile of the GMV confirm that a long position in aggregate longevity is financially appealing, with low risk and good diversification properties, only for an investor with a horizon longer than 12-13 years. This result is consistent with the hypothesis that annuity prices offer a return that is competitive with the alternative investments available to the investor: annuities for a 65-year-old investor have an effective duration of around 12 years (see, for example, Loeys, Panigirtzoglou and Ribeiro (2007)). Note that the 12-year minimum variance portfolio corresponds to an allocation in the ALS that is essentially zero.

Based on these observations, it is possible to conclude that the creation of short-duration (less than 10 years) longevity-linked securities is the key step for an efficient securitization of longevity risk. These securities offer a stochastic liability that can efficiently be used

to finance investments with good diversification properties. The risk-return analysis of the ALS shows that securitization of longevity, which is a long run risk in the sense of carrying a small but persistent component, does not necessarily require the use of long-duration securities. On the contrary, upon a precise quantification of the term structure of longevity risk exposures, a more efficient management of maturity transformation can be realized using structured securities, as, for example, swap contracts. These findings suggest that the problems that affected early longevity-indexed security issuances were determined by long durations. For these securities liquidity and credit risk components were so large as to overwhelm the effect of longevity risk both for pricing and hedging.

As a second illustration of the optimal mean-variance allocations including a position in the short-term ALS, in Figure 7 we plot the optimal allocations for a portfolio with an expected return of 10% as a function of the holding period returns. As expected, the ALS short position is used to leverage a portfolio of T-bills, equity and long-term bonds.

Insert Figure 7 about here

V. Longevity Securitization and Intertemporal Hedging of the Aggregate Longevity Risk

The results from the estimation of the extended VAR provide evidence of a significant response of annuity prices to variations in aggregate longevity rates. Quoted annuity prices are expected to include a compensation for the insurance company to bear a risk exposure for the unexpected rise of the undiversifiable longevity risk component. Following the conventional intertemporal CAPM (ICAPM) interpretation (Merton (1973)), expected utility maximizers attempt to hedge the stochastic changes of their investment opportunity

created by unexpected aggregate longevity shocks. The hedging portfolio is determined by an allocation in traded securities whose return is maximally correlated with the longevity shock qk_t . This portfolio is determined by the constrained minimization problem:

$$\begin{aligned} \min_{\mathbf{w}} \quad & Var_{t-1} \left[R_t^{qk}(\mathbf{w}_{t-1}) - qk_t \right] \\ s.t. \quad & R_t^{qk}(\mathbf{w}_{t-1}) = \mathbf{w}_{t-1} \cdot \mathbf{x}_t + W_{0,t-1}rtb_t \\ \mathbf{w}_{t-1} \quad &= [w_{xr,t-1}, w_{xb,t-1}, w_{x\Delta p,t-1}], \quad \mathbf{x}_t = [xr_t, xb_t, x\Delta p_t], \end{aligned}$$

where $\mathbf{x}_t$ includes the log excess returns of market securities plus the annuity log price growth in excess to the T-bill rate, R_t^{qk} is the return on the replication portfolio and $W_{0,t}$ is the investor's wealth at time t .

Recall that the purchase of an annuity is irreversible and payments are done until the death of a single beneficiary, while $x\Delta p_t$ is the ALS return, the annuity price variation between time $t-1$ and t for the 65-year-old male cohort in excess of the T-bill rate. Hence the Aggregate Longevity Hedging Portfolio (hereinafter ALHP) is not tradable unless an Annuity-linked Security paying off the return $\Delta p_{65,t}$ on a yearly basis is made available to investors. This is a benchmark example of the theoretical motivations underpinning the necessity of longevity securitization.

By construction, the ALHP tracks the aggregate longevity shock qk_t , and is therefore the best available product to reinsure aggregate longevity risk. Note that the efficiency of the replication increases with the number of investment opportunities exposed to aggregate longevity risk. In practice, the hedging portfolio is determined by performing the minimization over the set of unconstrained allocations $\mathbf{w}_t = [w_{xr,t}, w_{xb,t}, w_{x\Delta p,t}]$, while the

position in the short rate is set equal to $w_{rtb,t} = W_0 - \mathbf{w}_t \cdot \mathbf{1}$. Since our VAR model generates a stationary dynamics, the minimum variance replication portfolio corresponds to a time-independent allocation $\mathbf{w}_{t-1}^* = \mathbf{w}$. The first order condition, therefore, implies the solution:

$$\mathbf{w}^T = \text{Var} [\mathbf{x}_t]^{-1} \{ \text{Cov} [\mathbf{x}_t, qk_t] \}, \quad w_{rtb} = W_0 - \mathbf{w} \cdot \mathbf{1}.$$

We consider three alternative replication portfolios corresponding to zero initial investment ($W_0 = 0$) with an increasing set of restrictions on the allocations. Table 3 reports three hedging portfolios. The first portfolio refers to an unrestricted allocation, while the second reports an allocation in which investment in equity is not allowed ($w_{xr} = 0$). In both cases the allocation strategy is a short position in T-bill and equity and a long position in long-term bond and in the ALS.

Insert Table 3 about here

The third hedging portfolio is further restricted by forcing a zero allocation in the T-bill, $w_{rtb} = 0$, thus making the long-term bond the only available financial security available to finance the annuity liability. In all three cases the volatility induced by the aggregate longevity shock, as measured by the volatility of the aggregate longevity replication portfolio, is close to 60 basis points in annual terms. This value remains almost constant for any holding period with an essentially flat term structure of volatility. This value is slightly higher but comparable to the 50 basis points which are usually considered as the market standard for longevity risk (see Loeys et al. (2007)).

The above measures of aggregate longevity risk and the values of its replication portfolios can be revised on a yearly basis, at the highest revision frequency of the mortality rates. By making ALHP tradable on a yearly basis, one could sell aggregate longevity protection without incurring the liquidity problem of long duration securities.

According to Campbell (1996), an intertemporal utility-maximizing agent will optimally demand to invest or sell the hedging portfolio for aggregate longevity, if the state variable qk_t forecasts changes in financial or human capital. The empirical estimation of the extended VAR shows that the aggregate longevity shock indeed predicts price changes in annuity prices and in long-term bonds, and supports the hypothesis of existence of non-zero potential demand for ALS. While a complete discussion of the demand for longevity-linked securities requires a structural equilibrium framework as in, for example, the model of Cocco and Gomes (2012), in the next section we estimate the size of the compensation for bearing longevity risk assuming that the set of investment opportunities also includes the ALS.

VI. Pricing Longevity Risk

Milevsky, Promislow and Young (2005) propose to use the notion of Sharpe Ratio as an actuarial measure of aggregate longevity risk compensation. While the Sharpe Ratio of an investment is determined by the ratio between the expected return from the investment in excess to a benchmark security (usually the T-bill) and the expected volatility, in actuarial science the Sharpe Ratio determines the excess markup per unit of volatility that an aggregate longevity protection seller would charge to the protection buyer.

The discussion of the previous subsections suggests the possibility of using the informa-

tion conveyed by the VAR dynamic model in order to estimate a longevity risk compensation. It is easy to understand that within our framework the Sharpe Ratio of the ALHP is a reliable measure of such compensation. Note that by definition the ALHP is a zero investment portfolio, as:

$$\mathbf{w}^{\text{ALHP}} = [w_{xrtb,t}^{\text{ALHP}}, w_{xr,t}^{\text{ALHP}}, w_{xb,t}^{\text{ALHP}}, w_{x\Delta p,t}^{\text{ALHP}}], \quad \mathbf{w}^{\text{ALHP}} \cdot \mathbf{1} = 0,$$

and that the corresponding return can be split as the differential between the return of a long position in financial securities and a short position in the ALS liability. Hence, the ratio between the expected differential return and its risk provides a properly-defined Sharpe Ratio. Moreover, different from previous approaches, our estimation procedure identifies a dynamic compensation component whose evolution is maximally correlated with longevity shocks as opposed to other intermediation margins, which are not.

Conditional expectations of risks and returns may substantially differ from unconditional ones, and therefore the conditional Sharpe Ratios depend on the holding period τ and generate a term structure. This term structure depends on the initial level of the VAR state variables corresponding to the current level of financial returns and the current level of predictors. Over a horizon τ the Sharpe Ratio is the ratio between the τ -period expected excess simple return and the τ -period standard deviation. Hence, recalling that the extended VAR(1) models logarithmic returns we have:

$$\text{SR}_{\tau, \mathbf{r}_t} = \frac{\text{Ex}_{\tau, \mathbf{r}_t}}{\text{Std}_{\tau, \mathbf{r}_t}},$$

The relative compensation of different securities as measured by the Conditional Sharpe

Ratios depends on the horizon and on the state of the economy. When the holding period goes to infinity ($\tau \rightarrow \infty$) the Sharpe Ratio $SR_{\tau, \mathbf{r}_t}$ for a generic portfolio w converges to a limit SR_∞ .¹¹

$$\begin{aligned}
(9) \quad SR_\infty &= \frac{Ex_\infty}{Std_\infty}, \\
Ex_\infty &= \lim_{\tau \rightarrow +\infty} \left\{ E_{t, \mathbf{z}_t} \left[\frac{1}{\tau} \sum_{k=1}^{\tau} \mathbf{w} \cdot \mathbf{r}_{t+k} \right] + \frac{Std_{\tau, \mathbf{r}_t}^2(\mathbf{w})}{2\tau} \right\}, \\
Std_\infty &= \lim_{\tau \rightarrow +\infty} \sqrt{\text{Var}_{t, \mathbf{r}_t} \left[\sum_{k=1}^{\tau} \mathbf{w} \cdot (\mathbf{r}_{t+k} + r_{0,t+k} \mathbf{1}) \right] / \tau},
\end{aligned}$$

As shown in the Appendix, convergence of the $SR_{\tau, \mathbf{r}_t}$ to SR_∞ does occur at very long horizons, but the limiting procedure is necessary in order to produce a bona fide unconditional measure of expected performance consistent with the predictability patterns we have documented. Table 4 reports the estimation of long-term Sharpe Ratios for all the financial securities included in the extended VAR(1): an equity index, a rolling position in long-term bonds and the ALS. As expected, the Sharpe Ratio of the ALS security is negative, as its performance is lower than that of benchmark risk free security, the T-Bill. On the other hand, the low level of the period variance implies that the unconditional level of the ALS Sharpe ratio is as high as $SR_\infty^{\text{ALS}} = (-) 0.43$, indicating the potential usefulness of this synthetic security as a “liability” offering a good potential reward to investors seeking new

¹¹These quantities can be easily computed using the following expressions of the long run mean and covariance:

$$\begin{aligned}
\mu_\infty &= (\mathbb{I} - \Phi_1)^{-1} \Phi_0 \\
\Sigma_\infty &= (\mathbb{I} - \Phi_1)^{-1} \Phi_1 (\mathbb{I} - \Phi_1)^{-T}
\end{aligned}$$

diversification strategies. However, as this liability is financed by a short-term T-Bill, and the maturity mismatch is known to increase interest rate risk variations, the high value of the Sharpe Ratio can be misleading.

Insert Table 4 about here

A similar problem arises when measuring the Sharpe Ratio of the ALHP, the measure of aggregate longevity risk premium. The ALHP can be split in a short position in a ALS security and a long position in a portfolio of traded financial securities, thus making the Sharpe Ratio dependent on the composition of the portfolio used to finance the short position.

Table 5 reports the Sharpe measure of aggregate longevity risk compensation for the three alternative allocations defined in Table 3: the unrestricted one $\mathbf{w}^{\text{ALHP-Unr}}$, the one excluding allocation to equity, $\mathbf{w}^{\text{ALHP-1}}$, and the one where the ALS stochastic liability can be hedged using only long-term bonds $\mathbf{w}^{\text{ALHP-2}}$. As the Table 3 shows, Sharpe Ratios decrease with increasing restrictions. Notice that the opposite of the unrestricted allocation can well be interpreted as a feasible investment policy for an asset management company that is using a leverage policy similar to the one adopted by Berkshire Hathaway. In this case the optimal allocation corresponds to assume an insurance liability that is selling longevity risk protection to finance a long position in equity and bond markets.¹² The Sharpe Ratio value of 0.50 for this strategy witnesses the benefits that can be achieved using such a form of optimal asset and liability management policy and confirms that this

¹²As of february 5 2013, Berkshire Hathaway announced a deal to reinsure the risks associated with \$4 billion of future claims for two of health insurer Cigna Corp's run-off variable annuity businesses.

allocation strategy is one of the ingredients that contributed to the increase in Berkshire performance up to the 0.66 Sharpe ratio computed in Frazzini et al. (2013).

Insert Table 5 about here

In fact, this anomalous performance sharply disappears when the allocation used to hedge the stochastic liability induced by the short position in the ALS are constrained. The Sharpe Ratio is the lowest if the investor is allowed to hedge using long-term bonds, the security which is typically held in the reserves of insurance companies. In fact, the hedging portfolio $\mathbf{w}^{\text{ALHP-2}}$ using only long-term bonds is safer (but less profitable) than $\mathbf{w}^{\text{ALHP-1}}$ that also uses T-bills, as its duration matches that of the stochastic liability and thus has an identical response to (small) interest rate fluctuations.

In conclusion, the actuarial longevity premium estimate consistent with a prudent hedging policy is given by $\text{SR}_{\infty}^{\text{ALHP-2}} = 0.33$. Its value is not far from the conventional level 0.25 used in the actuarial pricing of longevity products as discussed in Loeys et al. (2007). This estimation is expected to overestimate the potential Sharpe Ratio from longevity liability, as shorting costs are not explicitly accounted for in this analysis. In addition the adverse selection effect is also expected to play a role here: the mortality rates of annuitants are known to be significantly smaller from those of average population (see Poterba (2001) and Mitchell et al. (1999)). Given the scarcity of data on prices of traded longevity-linked securities, the same problems affect virtually any empirical measure of longevity risk compensation.

It is possible to draw some conclusion on the robustness of the performance of a portfolio like $\mathbf{w}^{\text{ALHP-Unr}}$ with respect to exogenous liquidity shocks, thanks to the analysis of Koijen

and Yogo (2014) on the dynamics of annuity prices during the recent financial crisis period 2008-2009. They prove that during the financial crisis, life insurers sold long-term policies at deep discounts relative to their actuarial value and that this discount was supply driven. The average markup was as low as -19 percent for annuities despite the fact that interest rates were falling. In fact, insurance companies were trying to sell more annuities in the attempt to increase statutory capital i.e., assets relative to accounting liabilities. By selling new policies, life insurers were able to raise capital because their reserve valuation was more aggressive than mark to market during the financial crisis. In a way, life insurance policies sold directly to households were used by insurance companies as a source of short term funding when traded assets like equity and bonds were hit by financial markets' illiquidity. On the other hand, longevity risk is poorly correlated with market risk and short-term volatility of longevity shocks is extremely low, thus liquidity was provided by annuities at the expense of a rise in longevity risk exposure, having a negligible impact over the short term. In other terms, diversification benefits created by the possibility to use longevity liabilities to finance financial market investment seem to be robust to exogenous liquidity shocks.

From a normative point of view these considerations indicate a further indirect motivation to promote the integration between financial and actuarial markets: their development would drive a more transparent and precise assessment of the price for aggregate longevity risk. Among other benefits, it is worth mentioning that this assessment can certainly reduce the dangerous lack of awareness regarding the public and private costs deriving from generalized longevity increase.

VII. Conclusions

Our analysis shows that integration between insurance and financial markets is a promising direction to improve the efficiency of longevity risk-sharing. We believe that our results uncover some critical issues to improve longevity risk securitization. First, the long-term nature of the longevity risk requires an accurate analysis of the term structure of the risk return trade-offs generated by including a longevity-linked security in the set of investments. Second, a potentially large number of short-term investors would be willing to increase their exposure to longevity risk without increasing their investment horizon. This requires the organization of a maturity transformation activity by financial intermediaries that seems to be a crucial step in increasing the interest of the market for longevity-linked securities, as well as their liquidity. Finally, an integrated market for insurance and financial contracts with a publicly-traded longevity index would also imply a more transparent and efficient pricing of life annuities with a direct benefit to annuity subscribers.

References

- Barberis, N., “Investing for the Long Run when Returns are Predictable.” *The Journal of Finance*, 55 (2000), 225–264.
- Bayraktar, E., M. A. Milevsky, S. D. Promislow, and V. R. Young, “Valuation of Mortality Risk via the Instantaneous Sharpe Ratio: Applications to Life Annuities.” *Journal of Economic Dynamics and Control*, 33 (2009), 676–691.
- Biffis, E., and D. Blake, “Securitizing and Tranching Longevity Exposures.” *Insurance: Mathematics and Economics*, 46 (2010), 186–197.
- Biffis, E., M. Denuit, and P. Devolder, “Stochastic Mortality Under Measure Changes.” *Scandinavian Actuarial Journal*, 2010 (2010), 284–311.
- Blake, D., and W. Burrows, “Survivor Bonds: Helping to Hedge Mortality Risk.” *Journal of Risk and Insurance*, 68 (2001), 339–348.
- Blake, D., A. Cairns, G. Coughlan, and K. Dowd, “The New Life Market.” *Journal of Risk and Insurance*, 80 (2013), 501–557.
- Blake, D., A. Cairns, and K. Dowd, “A Two Factor Model for Stochastic Mortality with Parameter Uncertainty: Theory and Calibration.” *Journal of Risk and Insurance*, 73 (2006), 683–718.
- Campbell, J. Y., “Understanding Risk and Return.” *Journal of Political Economy*, 104 (1996), 298–345.
- Campbell, J. Y., Y. L. Chan, and L. M. Viceira, “A Multivariate Model of Strategic Asset Allocation.” *Journal of Financial Economics*, 67 (2003), 41–80.
- Campbell, J. Y., and L. M. Viceira, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*, Oxford University Press, New York, NY (2002).
- Campbell, J. Y., and L. M. Viceira, “The Term Structure of the Risk-Return Tradeoff.” *Financial Analysts Journal*, 61 (2005), 34–44.
- Cocco, J. F., and F. J. Gomes, “Longevity Risk, Retirement Savings, and Financial Innovation.” *Journal of Financial Economics*, 1034 (2012), 507–529.
- Finkelstein, A., and J. M. Poterba, “Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market.” *Journal of Political Economy*, 112 (2004), 183–208.

- Frazzini, A., D. Kabiller, and L. H. Pedersen, “Buffetts Alpha.”, Technical report, National Bureau of Economic Research (2013).
- Friedberg, L., and A. Webb, “Life is Cheap: Using Mortality Bonds to Hedge Aggregate Mortality Risk.” *The B.E. Journal of Economic Analysis & Policy*, 7 (2007), 31.
- Horneff, R. H., Wolfram J. and Maurer, O. S. Mitchell, and M. Z. Stamos, “Asset Allocation and Location over the Life Cycle with Survival-Contingent Payouts.” *Journal of Banking and Finance*, 33 (2009), 1688–1699.
- Koijen, R. S., and M. Yogo, “Shadow Insurance.”, Technical report, National Bureau of Economic Research (2013).
- Lee, R., “The Lee-Carter Method for Forecasting Mortality, with Various Extensions and Applications.” *North American Actuarial Journal*, 4 (2000), 80–93.
- Lee, R. D., and L. R. Carter, “Modeling and Forecasting U.S. Mortality.” *Journal of the American Statistical Association*, 87 (1992), 659–671.
- Lin, Y., and S. H. Cox, “Securitization of Mortality Risks in Life Annuities.” *The Journal of Risk and Insurance*, 72 (2005), 227–252.
- Lin, Y., and S. H. Cox, “Securitization of Catastrophe Mortality Risks.” *Insurance: Mathematics and Economics*, 42 (2008), 628–637.
- Loeys, J., N. Panigirtzoglou, and R. M. Ribeiro, “Longevity: a Market in the Making.”, Technical report, J.P. Morgan Global Market Strategy (2007).
- Markowitz, H., “Portfolio Selection.” *The Journal of Finance*, 7 (1952), 77–91.
- Mehra, R., and E. C. Prescott, “The Equity Premium: a Puzzle.” *Journal of monetary Economics*, 15 (1985), 145–161.
- Merton, R. C., “An Intertemporal Capital Asset Pricing Model.” *Econometrica*, 41 (1973), 867–887.
- Milevsky, M. A., *The Calculus of Retirement Income. Financial Models for Pension Annuities and Life Insurance*, Cambridge University Press Cambridge, UK (2006).
- Milevsky, M. A., S. D. Promislow, and V. R. Young, “Financial Valuation of Mortality Risk via the Instantaneous Sharpe Ratio.” *IFID Center, Toronto, Canada*, (2005).
- Milevsky, M. A., S. D. Promislow, and V. R. Young, “Killing the Law of Large Numbers: Mortality Risk

Premiums and the Sharpe Ratio.” *Journal of Risk and Insurance*, 73 (2006), 673–686.

Mitchell, O. S., J. M. Poterba, M. J. Warshawsky, and J. R. Brown, “New Evidence on the Money’s Worth of Individual Annuities.” *The American Economic Review*, 89 (1999), 1299–1318.

Pitacco, E., M. Denuit, S. Haberman, and A. Olivieri, *Modelling Longevity Dynamics for Pensions and Annuity Business*, Oxford University Press (2009).

Poterba, J. M., “The History of Annuities in the United States.”, in “The Role of Annuity Markets in Financing Retirement,” The MIT Press, Cambridge, MA (2001).

Figures and Tables

	rtb	xr	xb	$x\Delta pr$	y	$d - p$	spr	qk
Mean	0.027	0.056	0.0046	−0.027 5	0.056	−3.493	0.004	0.00(4)
Std Dev	0.052	0.160	0.108	0.057	0.031	0.411	0.014	0.041
Sharpe Ratio	-	0.353	0.04 259 3	−0.482 46	-	-	-	-

Summary Statistics.

Table 1: Mean returns is computed including the Jensen correction term, thus are computed as $\mu + 0.5\sigma^2$. Sharpe Ratio is computed as the ratio between Mean and Std Dev. **Note:** rtb = ex post real T-bill rate, xr = excess stock return, xb = excess bond return, $(d - p)$ = log dividend-price ratio, y = nominal T-bill yield, spr = yield spread.

	rtb_t (t)	xr_t (t)	xb_t (t)	$x\Delta pr_t$ (t)	y_t (t)	$(d-p)_t$ (t)	spr_t (t)	$(qk)_t$ (t)	R^2	$adjR^2$
rtb_{t+1}	0.477 (3.265)	0.007 (0.203)	-0.106 (-1.753)	0.074 (0.640)	0.590 (2.534)	-0.010 (-0.568)	-0.680 (-1.452)	-0.016 (-0.116)	0.475	0.389
xr_{t+1}	0.190 (0.388)	0.048 (0.363)	0.131 (0.679)	0.929 (1.660)	-0.215 (-0.225)	0.154 (2.423)	-0.892 (-0.462)	-0.236 (-0.372)	0.167	0.031
xb_{t+1}	0.917 (4.121)	-0.116 (-1.972)	-0.502 (-4.311)	0.116 (0.370)	0.267 (0.550)	-0.054 (-1.921)	4.362 (4.203)	0.588 (2.721)	0.562	0.491
$x\Delta pr_{t+1}$	0.531 (4.294)	-0.099 (-2.228)	0.057 (0.970)	-0.188 (-0.968)	-1.146 (-3.950)	-0.034 (-1.871)	-0.446 (-0.675)	0.271 (2.199)	0.516	0.437
y_{t+1}	-0.204 (-3.702)	0.025 (2.477)	0.038 (1.702)	-0.046 (-0.841)	0.940 (9.253)	0.007 (1.896)	0.286 (1.315)	-0.064 (-1.728)	0.794	0.761
$(d-p)_{t+1}$	-0.812 (-1.392)	0.149 (0.808)	-0.081 (-0.363)	-0.898 (-1.504)	-0.038 (-0.035)	0.885 (13.163)	2.475 (1.141)	0.148 (0.221)	0.845	0.819
spr_{t+1}	0.104 (2.452)	-0.017 (-1.775)	0.014 (1.040)	0.048 (1.069)	0.036 (0.466)	-0.002 (-0.585)	0.339 (2.076)	0.011 (0.404)	0.456	0.367
$(qk)_{t+1}$	-0.113 (-1.341)	-0.051 (-1.671)	0.000 (0.002)	-0.059 (-0.617)	-0.008 (-0.034)	-0.009 (-1.172)	0.373 (1.292)	0.747 (6.625)	0.575	0.506

Cross-Correlations of Residuals

	rtb	xr	xb	Δpr	y	$(d-p)$	spr	(qk)
rtb	3.763	0.206	0.117	0.192	-0.298	-0.234	0.318	-0.104
xr	-	14.584	0.040	-0.063	-0.248	-0.970	0.295	-0.118
xb	-	-	7.184	0.176	-0.620	-0.089	0.132	0.014
Δpr	-	-	-	3.927	-0.547	0.022	0.583	0.044
y	-	-	-	-	1.368	0.262	-0.843	0.041
$(d-p)$	-	-	-	-	-	15.768	-0.277	0.098
spr	-	-	-	-	-	-	1.047	-0.047
(qk)	-	-	-	-	-	-	-	2.617

VAR(1) - Matrix Φ_1 - Yearly Sample 1953-2010. Annuities

Table 2: VAR(1) coefficients with relative t -statistics and Cross-Correlations of Residuals. **Note:** rtb_t = ex post real T-Bill rate, xr_t = excess stock return, xb_t = excess bond return, $x\Delta pr_t$ = log difference on annuities premium ($d-p$) $_t$ = log dividend-price ratio, y_t = nominal T-bill yield, spr_t = yield spread, qk = aggregate longevity shock.

<i>HedgingPortfolio</i>	w_{xr}	w_{rtb}	w_{xb}	$w_{x\Delta pr}$
<i>Unconstrained</i>	-0.042	-0.014	0.009	0.047
<i>Constrained.1</i> $w_{xr} = 0$	0	-0.062	0.004	0.058
<i>Constrained.2</i> $w_{rtb} = w_{xr} = 0$	0	0	-0.059	0.059

Table 3: Optimal Allocation for Aggregate Longevity Hedging Portfolio under different constraints on the set of financial securities included in the hedging portfolio.

SR_{∞}^{Equity}	$SR_{\infty}^{LongBond}$	SR_{∞}^{ALS}
0.429	0.0323	(-) 0.437

Table 4: Long term Sharpe Ratios for securities included in the Extended VAR. Values are computed as illustrated in eq.9.

SR_{∞}^{Unc}	$SR_{\infty}^{Cons.1}$	$Cons.2_{\infty}$
(-) 0.504	(-) 0.421	(-) 0.333

Table 5: Long term Sharpe Ratios for Aggregate Longevity Hedging Portfolios. Values are computed as illustrated in eq.9.

Figure 1: Contributions to the qk_t shock: time series of aggregate vs sum of cohort-specific innovations as defined in eq.(6.)

Figure 2: Time series of historical (real) logarithmic price changes (dashed line) vs the replication as operated by the VAR dynamic model.

Figure 3: Term structure of risks for the securities included in the Extended VAR model.

Figure 4: Term Structure of correlations between financial securities and the Annuity-Linked Security.

Figure 5: Term structure of allocations forming the GMV portfolio at different horizons.

Figure 6: Term structure of risks for an allocation in T-Bill (continuous line), in the GMV portfolio restricted to financial securities (dashed line), in the GMV portfolio including also the Annuity-Linked Security.

Figure 7: Term structure of the efficient allocation with target expected return of 10%.