Intergenerational Risk Sharing in Life Insurance: Evidence from France*

Johan Hombert† Victor Lyonnet‡

October 11, 2017

Abstract

We study intergenerational risk sharing in Euro-denominated life insurance contracts. These savings products represent 80% of the life insurance market in Europe. Using regulatory and survey data for the French market, which is €1.3 trillion large, we analyze the patterns of intergenerational redistribution implemented by these products. We show that contract returns are an order of magnitude less volatile than the return of assets underlying these contracts. Contract return smoothing is achieved using reserves that absorb fluctuations in asset returns and that generate intertemporal transfers across generations of investors. We estimate the average annual amount of intergenerational transfer at 1.4% of contract value, i.e., €17 billion or 0.8% of GDP. Finally, we provide evidence that smoothing makes contract returns predictable, but inflows react only weakly to these predictable returns.

*We thank Autorité de Contrôle Prudentiel et de Résolution for access and assistance in putting together the data as well as comments. In particular we thank Frédéric Ahado, Anne-Lise Bontemps-Chanel, Fabrice Borel-Mathurin, Charles-Henri Carlier, Edouard Chrétien, Pierre-Emmanuel Darpeix, Olivier DeBandt, Dominique Durant, Henri Fraisse, Samuel Slama. The views expressed here are those of the authors and do not necessarily reflect those of the ACPR or the Banque de France.

†HEC Paris and CEPR, hombert@hec.fr

‡HEC Paris and ACPR, victor.lyonnet@hec.edu

1 Introduction

With the secular decline of defined benefit pension plans and Social Security around the world, life insurers are increasingly taking on the role of providing insurance against market risk by selling savings products with embedded insurance against aggregate market risk (Koijen and Yogo (2017)). Insurance against market risk can be provided to households through two different mechanisms (Allen and Gale (1997)). The first one is cross-sectional risk sharing between the insurer and households. The insurer bears part of the risk in order to provide households with a return partially hedged against market risk. Cross-sectional risk sharing is implemented by guaranteed return products like the ones studied by Koijen and Yogo (2017) in the US.

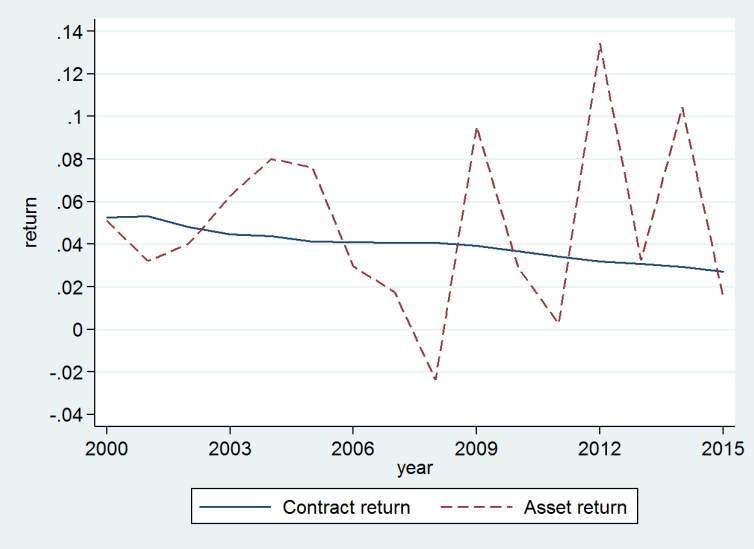
The other way to provide households with insurance against aggregate risk is to share risk intertemporally across generations of households, as was traditionally done in defined benefits plans and by Social Security. Intergenerational risk sharing is implemented by the most common type of life insurance contracts in Europe. Although the details of these products vary across countries (as well as their name: "Euro-denominated", "participating", "with profits", etc.), their common key feature is that they are backed by reserves that absorb fluctuations in market return over time and that are passed on between successive generations of investors. These products account for 80% of aggregate life insurance premiums in Europe in 2014 (Insurance Europe (2016)).

We focus in this paper on the French life insurance market, which is a large and mature market. At the end of 2015, the value of Euro-denominated life insurance contracts in France is €1,300 billion, which represents 80% of aggregate life insurance provisions and 40% of aggregate household financial wealth. These savings products work as follows. When a retail investor buys a life insurance contract, an account is created on which she can deposit and withdraw cash. Cash in all outstanding accounts associated to insurance contracts sold by the insurer is pooled into a single fund invested in asset markets and managed by the insurer. At the end of each year, account values are credited at a rate that is different from, and typically smoother than the rate of return of the fund's assets. It implies that account values diverge from the fund asset value. The difference between the two make up fund reserves. The role of these reserves is to cushion fluctuations in asset returns in order to smooth contract returns over time. Crucially, fund reserves are owed to (current or future) investors and are passed on between successive generations of investors as investors redeeming their contract give up their share of reserves while new investors share in existing reserves. Thus, fluctuations in the value of reserves over time generates redistribution across generations of investors.

We use regulatory and survey data from the national insurance supervisor to study intergenerational risk sharing in life insurance. Our objective is twofold. The first one is to study the smoothing of contract returns through fund reserves and to quantify the transfers across generations of investors it induces. The second objective is to analyze the implications of smoothing for the predictability of contract returns and for household flows into life insurance contracts.

Regarding the first objective, we have four main findings. The first finding is that contract returns are significantly smoother than funds' asset returns. This finding is illustrated in Figure 1, which plots value-weighted average asset returns and contract returns over 2000–2015. While the annual volatility of asset returns is 4.1%, the volatility of contract returns is only 0.8%. The smoothing of contract returns reflects insurance against market risk. This insurance may arise because fluctuations in asset returns are absorbed by insurer equity or because they are absorbed by fund reserves, i.e., by past and future generations of investors.

Figure 1: Asset Return vs. Contract Return



Source: Dossiers Annuels 1999–2015. Average contract returns (solid blue) and asset return (dashed red) weighted by account value.

The second finding is that fluctuations in asset returns are absorbed by fund reserves. Fluctuations in fund reserves explain almost entirely the wedge between asset returns and contract returns. In contrast, cross-sectional transfers between investor accounts and insurers are an order of magnitude smaller and less volatile than fluctuations in fund reserves. Fund reserves, which represent on average 12% of account values, fluctuate with realized asset returns: they increase when asset returns are high and they decrease when asset returns are low.

The third finding is that the amount of investor account value transferred across time through fluctuations in fund reserves is large. The average annual transfer between current investors and fund reserves is equal to 3.7% of account values. Evaluated at the 2015 level of aggregate account value of €1,200 billion, it amounts to an annual €44 billion that shifts from year to year on average, or 2% of GDP.

The fourth finding is that the amount of account value transferred across generations of investors is large. Transfers across years overstate transfers across generations of investors, because investors

hold their contracts for longer than one year. Investors receive positive transfers in some years (when contract return is above asset return) and negative transfers in other years (when contract return is below asset return) that partially net out. For instance, if annual transfers from fund reserves are i.i.d. and normally distributed, then the expected lifetime net transfer for an investor with a holding period of T year is equal to $1/\sqrt{T}$ of the expected annual transfer. Given an average holding period of 12 years, this back-of-envelop calculation yields an average transfer across generations of investors of 1.1% per year.

To obtain a more precise estimate of intergenerational transfers, we use information on inflows and outflows at the insurer-year level to estimate the share of each generation (characterized by their entry date and exit date in the contract) in total account value. Combining this estimate with annual transfers between investor accounts and fund reserves, we calculate that the amount of annual transfers across generations of investors is 1.4% of total account value on average. Evaluated at the 2015 level of aggregate account value of €1,200 billion, it amounts to an annual €17 billion that shifts across generations of investors on average, or 0.8% of GDP.

The second objective of the paper is to analyze insurer and investor behavior. We study how contract returns are related to the accumulated stock of reserves; how this creates predictability in contract returns; and how investors react to this predictability. These issues are crucial, because they determine the possibility (or impossibility) to implement intergenerational risk sharing, thereby diversifying risks that cannot be diversified at a given point in time. It is well-known that, although intergenerational risk sharing improves over cross-sectional risk sharing from a welfare perspective, competitive markets cannot implement intergenerational risk sharing because it requires future generations to share risk occurring before they start participating in the market (Stiglitz (1983), Gordon and Varian (1988)). Even if an intermediary implements a mechanism that achieves intergenerational risk sharing, it can be undone by competition from financial markets (Allen and Gale (1997)). The implementation of intergenerational risk sharing requires some market friction. An extreme form of friction is to force households to participate in the risk sharing scheme, which solves the problem that future generations would not agree to share in past losses in a competitive market. In this case, the first best allocation can be achieved (Ball and Mankiw (2007), Gollier (2008)).

A less extreme form of friction is imperfect competition between life insurers. Generally speaking, imperfect competition refers to situations where demand is imperfectly elastic to the price, leading to price dispersion in equilibrium. In markets for savings products, the notion of price corresponds to expected risk-adjusted net-of-fees contract return. However, in Euro-denominated contracts, insurers do not commit to a contract return but choose it ex post. To determine whether there is dispersion in expected contract returns, we test whether expected contract returns are predictable. Theory guides our choice of the predictive variable for expected contract returns. In Allen and Gale (1997) and Gollier (2008), the socially optimal contract return is an increasing function of the level

of fund reserves. The intuition is rooted in the role of reserves as a buffer against fluctuations in asset returns. When the asset return is high, part of it is hoarded as reserves for future distribution, leading to a high level of fund reserves that predicts large future contract returns. Conversely, when the asset return is low, insurers draw on reserves, leading to a low level of fund reserves that predicts low future contract returns.

We find robust empirical support for this prediction. A higher level of fund reserves predicts higher contract returns after controlling for insurer and year fixed effects. Our point estimate implies that a one standard deviation increase in fund reserves (8.8% of account value) is associated with an average increase in subsequent contract returns of 25 basis points per year or one-fourth of its standard deviation. In euro terms, a one euro increase in fund reserves is associated with an additional 2.9 cents of contract return per year. At this rate, if flows into life insurance contracts were zero or independent of fund reserves, it would take around $1/0.29 = 35$ years to consume the marginal euro of reserves.

Then, we test whether investor flows react to these predictable returns. If inflows were highly correlated with the level of reserves, this would quickly dilute reserves and undo the predictive power of reserves for future contract returns. We find that inflows are responsive to the level of reserves, but only to a limited extend. One additional euro of reserves leads on average to 8 cents of additional inflows per year. Given that reserves represent on average 12% of account value, the estimated inflow–reserves sensitivity implies that reserves are diluted by endogenous inflows at a rate of $0.08 \times 0.12 \approx 1\%$ per year. In other words, it would take a 100 times larger inflow–reserves sensitivity than the observed one to fully dilute reserves at a one year horizon. These results shed light on why life insurers are able to achieve large amounts of intergenerational redistribution in a decentralized market. Imperfectly elastic demand allows insurers to smooth contract returns through reserves management even though it creates return predictability, because investor flows react only weakly to this predictability.

The rest of the paper is organized as follows. The institutional framework is described in words in Section 2 and in equations in Section 3. Section 4 presents the data and summary statistics. Section 5 presents evidence on intergenerational risk sharing. Section 6 presents evidence on the relation between reserves, returns, and flows. Section 7 offers concluding remarks.

2 Institutional Framework

We describe how Euro-denominated life insurance contracts work. We then explain how fund reserves are used to smooth contract returns and to implement transfers across generations of investors.

2.1 Contract Returns

Life insurers sell savings products called Euro-denominated life insurance contracts. When an investor buys such a contract, she opens an account with the insurer on which she can deposit cash and withdraw cash at any time. The cash deposited in investor accounts is invested in asset markets through a fund managed by the insurer.

At the end of each calendar year, the account is credited by an amount equal to the account value times an interest rate (*taux de revalorisation*) that we will refer to in this paper as the contract return.[1] Most contracts have management fees proportional to the account value that are debited from the account at the end of the year. Some contracts also have entry and exit fees.

The contract return is chosen at the full discretion of the insurer at the end of each year, subject to two constraints. The first constraint is that the insurer commits to a minimum guaranteed return fixed at the subscription of the contract. Against a backdrop of decreasing interest rates, French life insurers have strongly reduced guaranteed rates close to zero since the 1990s (Darpeix (2016)). In 2015, the average guaranteed return of in-force contracts is 0.4% (see Panel A of Figure 2, and Appendix A for a description of the data used to produce this graph) while 77% of contracts have a zero guaranteed return. For the subset of contracts still open to new subscriptions, guaranteed returns are even lower: 0.1% on average and equal to zero for 87% of contracts. As a matter of comparison, the average contract return in 2015 is 2.6%. It implies that the minimum guaranteed rate is typically not binding. The contract return is strictly larger than the guaranteed return for 94% of contracts over 2011–2015 (see Panel B of Figure 2).

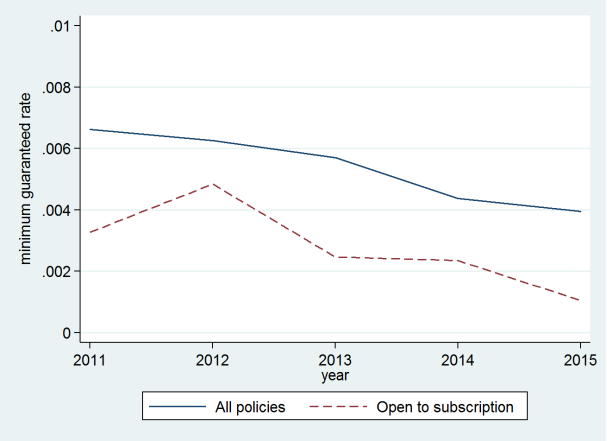
The insurer must credit the same rate of return to all investors holding the exact same product, that is, ruled by the same contract. However, the insurer is allowed to pay different returns to investors holding products ruled by different contracts. Insurers can have several in-force contracts at a given point time for two reasons. First, insurers can close a product to new subscriptions and create a new one for new clients. They may do so when they want to change the characteristics of the contract, like fees or the guaranteed rate, or for marketing purposes by changing the name of the product. Second, some insurers sell products targeted to specific clienteles. For instance, some insurers sell contracts with minimum investment amounts targeted to a wealthier clientele that may carry lower fees than the regular contract. Some insurers also sell group contracts for instance to corporations for their pension plans.

In practice, there is a little bit of dispersion in net-of-fees contract returns across the products sold by a given insurer, part of it because of different before-fees returns, part of it because of different fees. Using exhaustive data on net-of-fees contract returns at the product level collected by the national insurance supervisor for the 2011–2015 period, we calculate in Appendix A that the within-insurer-year standard deviation of annual net-of-fees return is 0.3 percentage points on

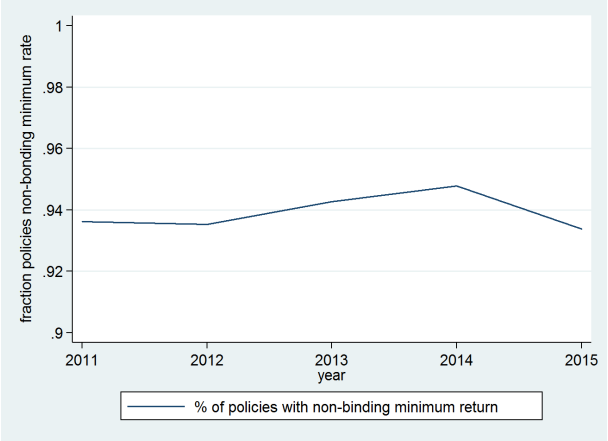
[1] Investors benefit from a tax relief on gains corresponding to cash withdrawn at least four years after the subscription of the contract and a larger tax relief after eight years.

Figure 2: Guaranteed Returns

Panel A: Average guaranteed return



Panel B: Fraction of contracts with non-binding minimum return



Source: Enquête Revalo 2011–2015. Panel A plots the weighted average minimum guaranteed return for all in-force contracts (solid blue) and for the subset of contracts still open to new subscriptions (dashed red). Panel B plots the fraction of contracts with a return strictly above the minimum guaranteed return.

average during this period. As a matter of comparison, the average net-of-fees return over the same period is 2.7% and the average time-series standard deviation over the longer period 2000–2015 is 1.0 percentage point. Thus, cross-contract dispersion in returns is small relative to the time-series dispersion.

More importantly, this cross-contract return dispersion is mostly explained by product fixed effects, which means that the returns of different contracts sold by a given insurer all follow the same time-series variation. When we include a product fixed effect, the residual cross-contract standard deviation of annual contract returns drops to 0.1 percentage points. This feature of the market will be important for our empirical analysis, because it implies that the pattern of intertemporal smoothing of contract returns is the same across all contracts sold by a given insurer, even if the absolute level of contract returns may differ.

The second constraint on contract returns is that life insurers are required by regulation to credit at least 85% of the intertemporal returns of underlying investments plus 90% of the fund net operating income (fees minus operating costs) to investors. Crucially, the regulatory constraint applies intertemporally and imposes no constraint on the timing of contract returns.[2] In the following section, we describe this regulation in more detail and we explain how insurers use fund reserves to choose the timing of contract returns.

[2]Another regulation imposes a constraint on the timing of contract returns, but it is typically not binding. See Footnote 4.

2.2 Fund Reserves

Each insurer has a unique fund holding the assets backing all the contracts. The economic balance sheet of the fund is as follows.[3] Assets are equal to the market value of the asset portfolio. Liabilities are equal to total account values plus fund reserves, where fund reserves are by definition equal to the difference between total asset value and total account value.

Fund reserves have two key features that are at the root of the intergenerational risk sharing mechanism. First, they are due to investors. Second, they are pooled across all investors. In particular, new investors share in reserves accumulated by previous investors while leaving investors give up their share of reserves.

In the French regulatory framework, fund reserves are made of three components associated to three components of asset returns.

Profit-sharing reserves At the end of each calendar year, fund income is calculated as the sum of financial income and technical income. Financial income is equal to asset yield (dividends on non-fixed income securities plus yield on fixed income securities) plus realized gains and losses on non-fixed income securities. Technical income is equal to fees paid by investors minus operating costs. The insurer decides on the sharing of the fund income between itself and investors. By law, the investors' share must be equal to at least 85% of financial income plus 90% of technical income.

The share of fund income attributed to investors is further split into a part credited immediately to investor accounts and a part credited to a reserve account called a profit-sharing reserve (*provision pour participation aux bénéfices*). The profit-sharing reserve account can only be used for future distribution to investor accounts. It implies that profit-sharing reserves effectively belong to (current and future) investors.

Crucially, profit-sharing reserves are not attached to the contracts already in force at the time of the creation of the reserve account. Instead, profit-sharing reserves are pooled across all contracts. When a investor redeems her contract or withdraw cash from her account, she gives up the associated right to future distribution of the profit-sharing reserves. Conversely, when a new investor buys a contract or an exiting investor deposits more cash on her account, she shares in the existing profit-sharing reserves.[4]

Capitalization reserves Realized gains and losses on fixed income securities have to be credited to or debited from a reserve account called the capitalization reserve account (*réserve de capitalisation*). The capitalization reserve account can only be used to offset future losses on fixed income

[3] The economic balance sheet differs from the accounting balance sheet because the former is marked-to-market while life insurance accounting principles are mostly based on historical cost accounting.

[4] The insurer must credit the funds in the profit-sharing reserve account to investor accounts within eight years. It can therefore hoard up to eight years of contract returns in profit-sharing reserve accounts. In practice, the eight-year constraint is never binding. Profit-sharing reserves represent less than one year of contract returns on average, and two years and a half at the 99th percentile.

securities and cannot be credited to investor accounts or to insurer income. Thus, capitalization reserves represent deferred financial income for the fund. Since at least 85% of the fund financial income must be distributed to investors, it implies that at least 85% of capitalization reserves effectively belong to (current and future) investors.[5]

Unrealized gains Unrealized capital gains are not booked as fund income. Thus, accumulated capital gains and losses create a wedge between the market value and the book value of the fund assets, which makes up the third component of reserves.[6] Since unrealized gains represent deferred fund financial income, at least 85% of their value effectively belong to (current and future) investors.

3 The Accounting of Intergenerational Risk Sharing

In this section, we present a simple framework of the accounting of intergenerational risk sharing in a life insurance fund.

$V_{i,t}$ denotes investor i's account value at the end of year t after payment of the annual net-of-fees return $y_{i,t}$. It evolves according to

$$V_{i,t} = (1 + y_{i,t})V_{i,t-1} + Flow_{i,t}, \tag{1}$$

where $Flow_{i,t}$ is net flows of investor i in year t.[7] In line with the institutional framework described in Section 2 and the evidence presented in Appendix A, we assume that the insurer pays the same return to all investors in a given year up to an investor fixed effect:

$$y_{i,t} = \eta_t + \xi_i. \tag{2}$$

The value-weighted average contract return is

$$y_t = \sum_i \frac{V_{i,t-1}}{V_{t-1}} y_{i,t}. \tag{3}$$

The balance sheet of the fund at the end of year t is

$$A_t = V_t + R_t, \tag{4}$$

[5]Technically, capitalization reserves are booked as insurer equity and not as liabilities. However, from an economic point of view, at least 85% of capitalization reserves represent fund reserves because at least 85% of their value are due to investors.

[6]While unrealized capital gains are never booked as fund profit, there are deviations from historical cost accounting principles that force insurers to recognize "long and lasting" unrealized losses to prevent them from running too large unrealized losses. See Appendix B for details.

[7]We assume here that flows take place at the end of the year after payment of the annual return to simplify the exposition. We relax this assumption in the empirical analysis.

where A_t is the market value of fund assets, $V_t = \sum_i V_{i,t}$ is total account value, and R_t is fund reserves. Fund assets evolve according to

$$A_t = (1 + x_t)A_{t-1} + Flow_t - \Pi_t, \tag{5}$$

where x_t is the return of fund assets net of the fund operating costs, $Flow_t = \sum_i Flow_{i,t}$ is total net flows, and Π_t is the payoff of the insurer.

Combining (1), (3), (4), and (5), we obtain

$$x_t A_{t-1} = y_t V_{t-1} + \Pi_t + \Delta R_t. \tag{6}$$

Equation (6) describes how asset income (on the left-hand side) is shared across three groups of agents (on the right-hand side): (i) current investors who receive $y_t V_{t-1}$; (ii) the insurer who gets Π_t; and (iii) fund reserves whose amount varies by $\Delta R_t = R_t - R_{t-1}$. Since beginning-of-year reserves have been accumulated by past investors and end-of-year reserves are available for distribution to future investors, the change in fund reserves represents a payoff to past and future investors. The last term of (6) is thus at the root of intergenerational risk sharing.

Equation (6) highlights that both *cross-sectional* risk sharing and *intergenerational* risk sharing are potentially at play. Investors may be insured against market risk because fluctuations in asset returns are absorbed by the insurer. This mechanism reflects static (cross-sectional) risk sharing between the insurer and investors. The covariation between Π_t and x_t determines how much insurance is provided by the insurer to investors. Alternatively, investors may be insured against market risk because fluctuations in asset returns are absorbed by fund reserves. This mechanism reflects intertemporal (intergenerational) risk sharing between current insurers and past and future investors. The covariation between ΔR_t and x_t determines how much insurance is provided to investors through intergenerational risk sharing. Both risk sharing mechanisms may also be at play simultaneously.

Equation (6) can be compared to analogous decompositions of investment returns for other savings products. In traditional savings products like mutual funds or unit-linked life insurance products, the first term on the right-hand side moves one-for-one with the fund asset return while the second and third terms are zero. In structured savings products like variable annuities with guaranteed returns, only the first two terms on the right-hand side arise as risk is shared between investors and the insurer while the third term is zero.

To quantify the amount of intergenerational risk sharing, we compare investors' payoff to what they would have received in a counterfactual situation with a constant level of reserves, same asset return, and same insurer payoff. Relative to this benchmark, the payoff of investors holding a contract in year t is higher by $-\Delta R_t$. $-\Delta R_t$ represents a transfer from investors holding a contract in year $t' \neq t$ to investors holding a contract in year t. We denote the amount of *InterTemporal*

Transfer between investor accounts in year t and investor accounts in other years by

$$\mathcal{ITT}_t = |-\Delta R_t|. \tag{7}$$

The amount transferred across investors is less than the amount transfer across years, because investors hold their contracts for more than one year. Investors may thus receive positive transfers in some years and negative transfers in other years that partially net out. Given the institutional feature that all investors receive the same return up to an investor fixed effect as specified in equation (2), each investor i receives in year t a transfer proportional to her account value and equal to $\frac{-\Delta R_t}{V_{t-1}} V_{i,t-1}$. Thus, investor i buying a contract at the beginning of year t_i^0 and redeeming it at the end of year t_i^1 receives a lifetime net transfer throughout her holding period $\mathcal{H}_i = [t_i^0, t_i^1]$ equal to[8]

$$\sum_{s \in \mathcal{H}_i} \frac{-\Delta R_s}{V_{s-1}} V_{i,s-1}.$$

Finally, the *Annualized Lifetime Transfer* to investor i in year t, calculated as a constant fraction of the beginning-of-year account value $V_{i,t-1}$, is equal to

$$\mathcal{ALT}_{i,t} = \frac{V_{i,t-1}}{\sum_{s \in \mathcal{H}_i} V_{i,s-1}} \sum_{s \in \mathcal{H}_i} \frac{-\Delta R_s}{V_{s-1}} V_{i,s-1}. \tag{8}$$

$\mathcal{ALT}_{i,t}$ is the net transfer in year t to investor i from other generations of investors. We say that two investors $i \neq i'$ belong to the same generation if they have the same holding period and same timing of investment: $V_{i,t} = \kappa V_{i',t}$ for all t. Two investors of the same generation are always on the same side of the redistribution because $\mathcal{ALT}_{i,t}$ and $\mathcal{ALT}_{i',t}$ have the same sign. In contrast, two investors belonging to different generations can be on opposite sides of the redistribution. Therefore, $\mathcal{ALT}_{i,t}$ represents intergenerational transfer. We denote the total amount of *InterGenerational Transfer* in year t by

$$\mathcal{IGT}_t = \sum_i |\mathcal{ALT}_{i,t}|. \tag{9}$$

[8] Transfers taking place in different years are not discounted differently. How they should be discounted depends on the assumption regarding whether current investors are entitled to the return on fund reserves. We adopt the convention that they are. Thus, they receive zero transfer when $R_t = R_{t-1}$ and the annual transfer amount is proportional to $(-\Delta R_t)$. The discount factor in the calculation of lifetime net transfer must be chosen such that aggregate lifetime net transfers is equal to zero (or equal to the difference between initial and final level of fund reserves if these are not equal to zero). Denoting the discount factor by ρ_t and summing over all investors: $\sum_i \sum_{t \in \mathcal{H}_i} \rho_t \frac{-\Delta R_t}{V_{t-1}} V_{i,t-1} = \sum_t \rho_t(-\Delta R_t)$. It follows immediately that $\rho_t = 1$ for all t.

The argument can be illustrated with a numerical example. Suppose there is only one investor starting with $V_0 = 100$ and investing for two years. Asset return net of insurer payoff is 10% in both years. The initial level of fund reserves is $R_0 = 0$. The first year return $0.1 \times 100 = 10$ is entirely hoarded as fund reserves: $A_1 = 110$, $V_1 = 100$, and $R_1 = 10$. In the second year, asset return $0.1 \times 110 = 11$ and fund reserves are credited to the investor account: $A_2 = V_2 = 121$ and $R_2 = 0$. Annual transfer to the investor is $T_1 = -10$ in the first year and $T_2 = 10$ in the second year. Since there is only one investor, her lifetime net transfer should be zero. Since $T_1 = -T_2$, it implies that the discount factor is constant.

In Section 5, we shall quantify intertemporal transfer $\mathcal{ITT}_t$, investor lifetime transfer $\mathcal{ALT}_{i,t}$, and intergenerational transfer $\mathcal{IGT}_t$ in the French life insurance market.

4 Data and Summary Statistics

We use regulatory data from the national insurance supervisor *Autorité de Contrôle Prudentiel et de Résolution*. The data covers all stock and mutual companies with life insurance operations in France from 1999 to 2015. It contains information on all the variables that appear in the analytical framework of the previous section: account value, inflows, outflows, contract return, profit-sharing and capitalization reserves, book value and market value of assets, asset return. The information is reported by type of contracts, which allows us to focus on Euro-denominated contracts. We restrict the analysis to stock insurance companies, which represent 97% of aggregate life insurance provisions. We drop insurers with less than €10 million of life insurance provision. Because we need lagged values to calculate the change in fund reserves, the sample period of our analysis is 2000–2015. The final sample has 76 insurers and 978 insurer-year observations.

Table 1: Summary Statistics

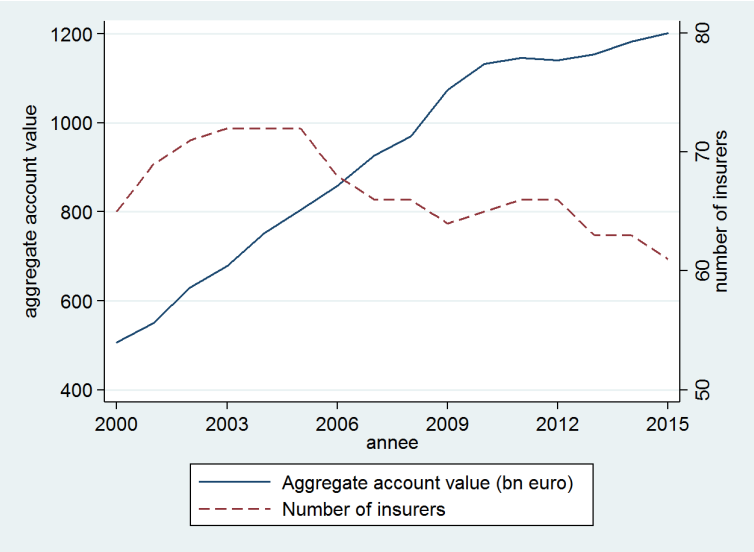
	Mean	Median	Std.Dev.	N
Account value (bn euro)	13.9	3.1	30.1	978
Inflows (% account value)	10.8	9.9	6.5	978
Outflows (% account value)	8.5	8.0	3.6	978
Total fund reserves (% account value)	11.9	11.0	8.8	978
Profit-sharing reserves (% account value)	2.5	1.9	2.1	978
Capitalization reserves (% account value)	1.9	1.3	2.1	978
Unrealized gains (% account value)	7.5	6.4	7.6	978
Share bonds (%)	81.2	83.3	11.6	978
Share stocks (%)	12.4	11.0	8.4	978
Asset return (%)	5.0	4.4	5.4	978
Contract return (%)	4.0	4.0	1.0	978

Source: Dossiers Annuels 1999–2015.

Table 1 reports summary statistics for the main variables associated to Euro-denominated contracts (see Appendix C for details about variable construction). The average (median) insurer has €13.9 billion (€3.1 billion) of account value. Inflows (premiums), which include cash deposited in newly opened contracts and in existing contracts, represent on average 10.8% of account value per year. Outflows, which include partial and full redemptions, either voluntarily or at expiration of the contract (investor death), represent on average 8.5% of account value per year.

The combination of positive net flows and compounded contract returns generates an increasing trend in aggregate account value plotted in Figure 3. Aggregate account value grows from €505

Figure 3: Aggregate Account Value (billion €)



Source: Dossiers Annuels 1999–2015. Aggregate account value of Euro-denominated contracts in billion 2015 euros (solid blue) and of insurers in the sample (dashed red).

billion in 2000 to €1,200 billion in 2015 (all amounts are in constant 2015 euros). Aggregate growth reflects internal growth of existing life insurers rather than entry of new insurers. The number of insurers in the sample is 65 at the beginning of the period and 61 at the end.

Fund reserves represent on average 11.9% of account value, of which about two-third are unrealized gains (7.5%) and one-third are accounting reserves (2.5% of profit-sharing reserves and 1.9% of capitalization reserves). Figure 4 plots the time-series of aggregate reserves as a fraction of account value over time. Reserves fluctuate widely over time, reflecting active reserve management in order to smooth contract return.

On the asset side, 81.2% of the fund portfolio is invested in government and corporate bonds, 12.4% in stocks, the rest in real estate, loans, and cash. Average asset return net of operating costs is 5.0% per year.

Average contract return before fees is 4.0% per year. The wedge between average asset return and average contract return is due to two factors. First, the insurer can keep up to 15% of asset return, which represents about 75 basis points on average. However, competition between insurers usually keeps the insurer's share of fund income below this level. Second, part of asset returns have been retained to offset the dilution of fund reserves caused by positive net flows over the sample period. Because new investors share in accumulated reserves while leaving investors give up their share of reserves, reserves are mechanically diluted at a rate equal to the net flow rate. Given the average net flow rate of 2.3% per year during the sample period and the average level of fund reserves of 12% of account value, reserves have been mechanically diluted at an annual rate of about 0.3%

Figure 4: Fund Reserves (% account value)



Source: Dossiers Annuels 1999–2015. Aggregate fund reserves as a fraction of account value (blue) and breakdown into the three components of reserves: unrealized gains (long dashed red), profit-sharing reserves (dashed green), and capitalization reserves (short dashed orange).

of account value. It implies that insurers would have had to retain about 30 basis points of asset returns per year in order to keep the ratio of reserves to account value constant. In fact, the average ratio of reserves is higher at the end of the period than at the beginning by about 4 percentage points. It implies that insurers have retained over this 15-year period an additional $0.04/15 \approx 25$ basis points per year on average.

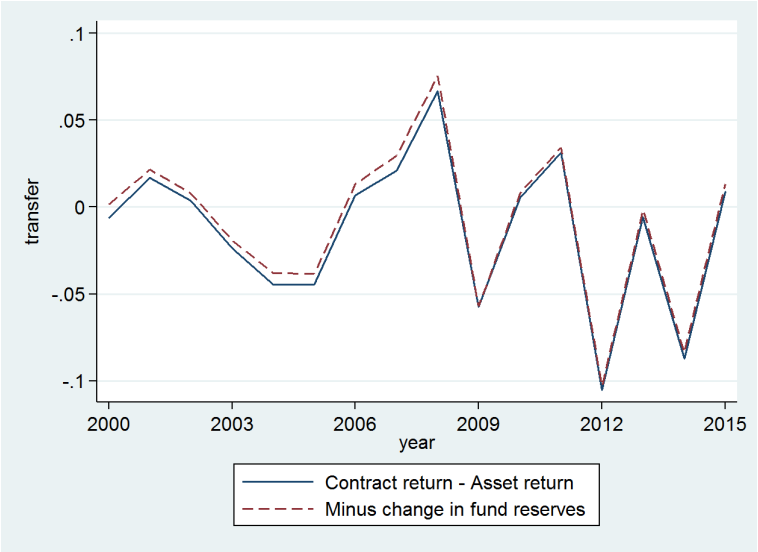
5 Intergenerational Risk Sharing: Evidence

5.1 Intertemporal Smoothing: The Role of Fund Reserves

Figure 1 plots weighted-average asset return and contract return from 2000 to 2015. The key pattern is that the contract return is much less volatile than the return of underlying assets. The annual volatility of asset return is 4.1% vs. 0.8% for contract return.

Contract return smoothing can be implemented by two mechanisms reflected in equation (6). Contract return ($y_t V_{t-1}$) can be hedged against variation is asset return ($x_t A_{t-1}$) by offsetting transfers from the insurer (Π_t) or by offsetting transfers from fund reserves (ΔR_t). To visualize graphically the contribution of fund reserves to contract return smoothing, Figure 5 plots two series: the difference between contract return and asset return ($y_t V_{t-1} - x_t A_{t-1}$) and transfers from reserves ($-\Delta R_t$), both aggregated across all insurers and normalized by aggregate beginning-of-year account value (V_{t-1}).

Figure 5: Fund Reserves Absorb Asset Return Fluctuations



Source: Dossiers Annuels 1999–2015. Difference between aggregate contract return and asset return normalized by account value $(y_t V_{t-1} - x_t A_{t-1})/V_{t-1}$ (solid blue) and aggregate transfers from reserves normalized by account value $-\Delta R_t/V_{t-1}$ (dashed red).

Figure 5 shows that almost all of the difference between asset and contract return is explained by variation in fund reserves. It means that the smoothing of contract return is generated by intertemporal transfer across investor accounts. The next section analyzes the extent to which these intertemporal transfers translate into intergenerational transfers.

5.2 Intertemporal Transfers and Intergenerational Transfers

We now quantify intertemporal transfer $\mathcal{ITT}_t$, investor annualized lifetime transfer $\mathcal{ALT}_{i,t}$, and intergenerational transfer $\mathcal{IGT}_t$ at the insurer level. We focus on the sample of insurers for which we have data throughout 1999–2015, which leads us to make two adjustments to the sample. First, when an insurer is acquired by another one, their reserves are pooled together. In this case, we consolidate both entities into a single one before the acquisition date such that we have a single insurer with a constant perimeter throughout the sample period.[9] Second, we drop a few insurers that enter or exit during the sample period. The final sample has 50 insurers that we observe continuously from 1999 to 2015 and that account for 94% of the aggregate account value in the initial sample.

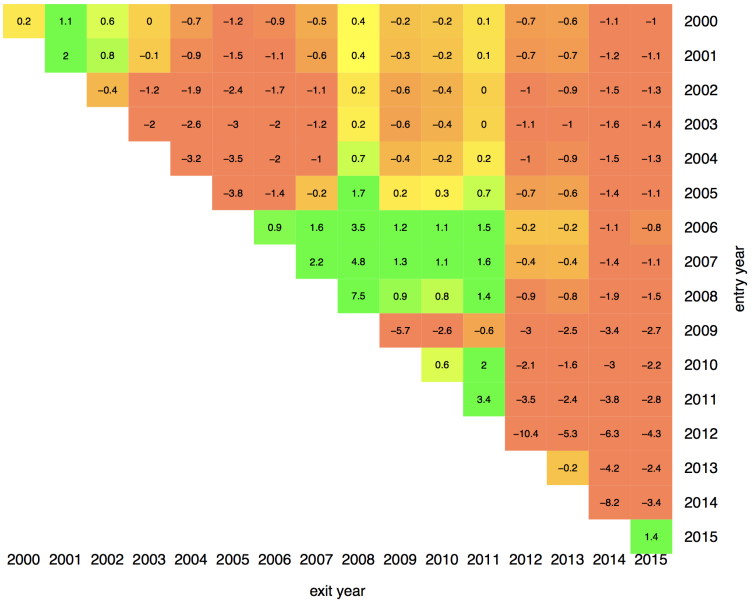
Variation in fund reserves generates transfer of investor account value across years. $-\Delta R_t$ is the transfer to investor accounts in year t from fund reserves, that is, from investor accounts

[9]We apply this procedure for six acquisitions: Natio by Cardif in 2005, Generali Assurances, GPA, and Guardian by Generali in 2006, Ecureuil by CNP in 2007, and Assurances Banque Populaire by Natixis Assurances in 2013.

in years $t' \neq t$. The value-weighted (equal-weighted) average amount of intertemporal transfer $\mathcal{ITT}_t = |-\Delta R_t|$ is equal to 3.7% (4.2%) of total account value. Evaluated at the 2015 level of aggregate account value of €1,200 billion, it amounts to an annual €44 billion that shifts from year to year on average, or 2% of GDP.

Investor annualized lifetime transfer is smaller than intertemporal transfer, because investors hold their contract for several years and thus receive positive transfers in some years and negative transfers in other years that partially net out over the holding period. To visualize how the annualized lifetime transfer $\mathcal{ALT}_{i,t}$ given by equation (8) depends on the holding period, Table 2 displays its value for the average contract and for every possible holding period $[t^0, t^1]$ during our sample period, $2000 \leq t^0 \leq t^1 \leq 2015$. We do the calculation for an investor who holds the representative contract with underlying fund reserves equal to the value-weighted average level of fund reserves, and invests a constant amount $V_{i,t} = 100$ by withdrawing the interests paid at the end of each year.

Table 2: Annualized Lifetime Transfer as Function of Holding Period



Source: Dossiers Annuels 1999–2015. The table reports $\mathcal{ALT}_{i,t}$ given by equation (8) for an investor buying the representative contract at the beginning of year t^0 (rows) and redeeming it at the end of year t^1 (columns). Reading: An investor buying a contract at the beginning of 2006 and redeeming it at the end of 2011 has received an additional 1.5 percentage points per year relative to an investment in a fund or a unit-linked contract that would be invested in the same portfolio of assets but would not smooth investor returns with reserves.

Table 2 is displayed as a heatmap. The numbers in the heatmap should be interpreted as the additional annual returns of having invested in the representative contract relative to having invested in a fund that (a) is invested in the same portfolio of assets but (b) do not smooth investor returns with reserves. For instance, the counterfactual can be an investment in a unit-linked contract that

would be invested in the same asset portfolio as the Euro-denominated contract.

Transfers for holding periods spanning the 2008 stock market crash and the 2011 sovereign debt crisis tend to be positive. The poor asset returns in these years were offset by large transfers from fund reserves. For instance, an investor buying a contract at the beginning of 2006 and redeeming it at the end of 2011 received an extra return of 1.5 percentage points per year on her contract relative a counterfactual with no reserve management. Taking an ex-post perspective, these investors have been net beneficiaries of the intergenerational redistribution implemented by life insurance.

Transfers for holding periods spanning the recent period of low interest rates are negative. Drops in interest rates led to high bond returns that were not credited to investor accounts, but hoarded as reserves. This hoarding behavior can be seen on Figure 4, which shows that fund reserves increased from 2% of account value in 2011 to 18% of account value in 2015. This makes exit from life insurance contracts in the post-2011 low interest rate environment detrimental compared to the counterfactual with constant reserves. Investors who have redeemed their contracts during this period have been ex-post net contributors of the intergenerational redistribution mechanism.

The high level of fund reserves in 2015 also suggests that insurers might draw on their reserves in the coming years to pay investors returns above market returns, which might be low given the currently low level of interest rates and the potential for an increase in interest rates. In this scenario, the heatmap for holding periods spanning the years to come might turn green again.

The last step is to aggregate individual annualized lifetime transfers $\mathcal{ALT}_{i,t}$ at the insurer level to calculate the amount of intergenerational transfer $\mathcal{IGT}_t$ given by equation (9). A back-of-the-envelop calculation illustrates the relation between intertemporal transfer $\mathcal{ITT}_t$ and intergenerational transfer $\mathcal{IGT}_t$. Suppose intertemporal transfer $\mathcal{ITT}_t$ is i.i.d. and normally distributed with zero mean and all investors have T year-holding periods. Then, expected intergenerational transfer is equal to $1/\sqrt{T}$ of expected intertemporal transfer. During the sample period, the average outflow rate is 8.5%, which implies an average duration of 12 years and an average amount of intergenerational transfer of $3.7\%/\sqrt{12} = 1.1\%$ per year. Accounting for holding period heterogeneity would lead to a higher estimate because of the convexity of $1/\sqrt{T}$.

To have an exact measure of intergenerational transfer, we would need to observe the entire investment history $V_{i,t}$ of all investors. This is not possible since the investment histories of current investors still holding a contract are not over. There are two other data limitations. First, the regulatory data starts in 1999, so we do not observe the entire investment history of investors who entered their contract before 1999 even if they no longer hold a contract in 2015. It implies that we can calculate annualized lifetime transfer $\mathcal{ALT}_{i,t}$ for investors with holding periods within 2000–2015 (we need one lagged year to calculate asset returns). Second, we observe inflows and outflows at the insurer level but not at the investor level. It implies that we know the average holding period but not its entire distribution. To calculate $\mathcal{IGT}_{i,t}$, we assume that the outflow rate is constant across contract age at the insurer-year level and that investors only make one-off investments.

Formally, denoting by $V_t(t^0)$ the total account value of contracts subscribed in year t^0, we assume $V_t(t^0) = (1-\phi_t)(1+y_t)V_{t-1}(t^0)$ for all $t^0 < t$, where the outflow rate ϕ_t is calculated to match observed outflows at the insurer level: $\sum_{t^0<t} \phi_t(1+y_t)V_{t-1}(t^0) = Outflow_t$; and account value of new contracts is calculated to match observed inflows at the insurer level: $V_t(t) = Inflow_t$.[10] Under this assumption, we can reconstruct the investment history of all generations of investors and calculate the intergenerational transfer amount $\mathcal{IGT}_t$.

The value-weighted (equal-weighted) amount of intergenerational transfer is equal to 1.4% (1.5%) of account value. Evaluated at the 2015 level of aggregate account value of €1,200 billion, it amounts to an annual €17 billion that shifts across generations of investors on average, or 0.8% of GDP.

The assumption of outflow rates independent of contract age is likely to under-estimate the amount of intergenerational transfer. Actual outflow rates are decreasing in contract age (FFSA-GEMA (2013)). It implies that the actual dispersion of holding periods is higher than the dispersion obtained under the assumption of age-independent outflow rate, for a given average holding period that we match to the observed one of 12 years. Since expected annualized life transfer is convex in the holding period, under-estimating the dispersion of holding periods leads to under-estimating intergenerational transfer.

6 Sensitivity of Returns and Flows to Reserves

As argued by Stiglitz (1983), competitive markets cannot implement intergenerational risk sharing because it requires future generations to share risk occurring before they start participating in the market. Gordon and Varian (1988) show that generations borne after periods of low returns may be worse off in the intergenerational risk sharing scheme, because they expect to be net contributors to the scheme to rebuild depleted reserves. Even if the government or a monopolistic intermediary implements a mechanism that achieves intergenerational risk sharing, Allen and Gale (1997) show that it may be undone by competition from financial markets.

The implication of these results is that the implementation of intergenerational risk sharing requires some market friction. An extreme form of friction is to force households to participate in the risk sharing scheme, which solves the problem that later generations would not agree to share in past losses in a competitive market. Ball and Mankiw (2007) and Gollier (2008) characterize the socially optimal allocation in such environments. The model of Gollier (2008) is particularly relevant, because it fully characterizes the socially optimal reserves management of a fund that closely resembles life insurance funds, except that flows are exogenous.

A less extreme form of friction is imperfect competition due to finitely elastic flows. It implies that even if the level of fund reserves predict future contract returns, as it does in the social optimum

[10] These formulas assume that inflows and outflows take place at year-end whereas in practice they occur throughout the year. To take this into account, we use modified formulas that account for interests accrued during the year. See Appendix D for details.

studied by Gollier (2008), investor flows would respond only weakly to these predictable expected returns, thereby sustaining return predictability in equilibrium and enabling intergenerational redistribution.

To explore this hypothesis, we test whether the level of fund reserves predict future contract returns in Section 6.1, and flows respond to the level of fund reserves in Section 6.2.

6.1 Contract Return Predictability

Our approach to test for contract return predictability is guided by the model of Gollier (2008), who solves for the optimal reserves management policy that achieves first best intergenerational risk sharing. At the first best, investor return depends on the level of fund reserves, but not on current asset return beyond its effect on reserves. The intuition behind this result is similar to the one for permanent income hypothesis according to which the optimal level of consumption depends only on wealth and not on current income beyond its effect on wealth. Conversely, in a perfectly competitive market, funds would not provide any insurance and would pay the asset return as in a traditional mutual fund.

We regress the contract return chosen at the end of the year by the insurer on the end-of-year level of reserves just before investor accounts are being credited this annual return, normalized by total account value. The regressor is thus equal to the beginning-of-year level of reserves plus the current year asset return (see equation (6)). Using the notations of Section 3, we regress y_t on $\frac{1}{V_{t-1}}(R_{t-1} + x_t A_{t-1})$. We include insurer fixed effects and year fixed effects and double-cluster standard errors by insurer and year. We report the results for regressions weighted by the insurer share of account value in aggregate account value in the current year in columns (1) and (2) of Table 3, and equal-weighted regressions in columns (3) and (4).

The coefficient on end-of-year reserves in column (1) is positive and statistically significant at the 1% level. The point estimate indicates that a 1 percentage point increase in fund reserves in associated with a 2.9 basis point increase in annual contract return. In other words, out of each additional euro of fund reserves, 2.9 cents per year are credited to investor accounts. At this rate, if flows did not react to the level of reserves, the distribution of the marginal euro of fund reserves to investors would be spread over the next $1/0.029 = 35$ years. Another way to gauge the economic significance of the point estimate is to multiply it by the standard deviation of fund reserves (8.8% of account value). It implies that a one standard deviation increase in fund reserves is associated with an average increase in subsequent contract returns of 25 basis points per year or one-fourth of its standard deviation (100 basis points).

Under first best intergenerational risk sharing, the contract return should only depend on the current asset return through its effect on the level of reserves. If we decompose the year-end level of reserves into beginning-of-year reserves plus current asset return, the contract return should depend on both terms equally. To see whether this holds, we include separately both terms in the

Table 3: Contract returns

	Contract return			
	(1)	(2)	(3)	(4)
Lagged reserves + Asset return	.029***		.029**	
	(.008)		(.013)	
Lagged reserves		.035***		.031**
		(.0078)		(.012)
Asset return		.017		.025
		(.011)		(.017)
Insurer FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Weights	Value	Value	Equal	Equal
Adjusted-R2	.8	.81	.53	.53
Observations	978	978	978	978

Source: Dossiers Annuels 1999–2015. Panel regressions of contract return on beginning-of-year fund reserves, annual asset return, insurer fixed effects, and year fixed effects. Regressions in columns (1) and (2) are weighted by insurer share in aggregate account value and regressions in columns (3) and (4) are equal-weighted. Standard errors are double clustered by insurer and year. ***, **, and * mean statistically significant at the 1%, 5%, and 10% levels, respectively.

regression. The coefficient on lagged reserves in column (2) is positive and significant while the coefficient on current asset return is positive but insignificant. The evidence is thus consistent with optimal intergenerational risk sharing. The results of equal-weighted regressions in columns (3) and (4) are similar.

6.2 Inflow Elasticity

Are investor inflows sensitive to the level of fund reserves? To study this question, we regress inflows (premiums) on beginning-of-year reserves, both normalized by beginning-of-year account value. We exclude insurers that enter during the sample period because the dynamics of inflows is likely to be different out of steady state. Results are reported in Table 4 for value-weighted and equal-weighted regressions.

The positive coefficient on beginning-of-year reserves in column (1) is positive and statistically significant at the 10% level. The point estimate indicates that an additional €1 in fund reserves leads on average to 8 cents of additional inflows per year. Given that the average level of reserves is 12% of account value, the estimated inflow–reserves sensitivity implies that reserves are diluted by endogenous inflows at a rate of $0.08 \times 0.12 \approx 1\%$ per year. In other words, it would take a 100 times larger inflow–reserves sensitivity than the observed one to fully dilute reserves at a one year

Table 4: Inflows

	Inflows	Inflows unit-linked	Inflows	Inflows unit-linked
	(1)	(2)	(3)	(4)
Lagged reserves	.08*	-.28**	.069*	-.029
	(.042)	(.097)	(.037)	(.18)
Insurer FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Weights	Value	Value	Equal	Equal
Adjusted-R2	.73	.5	.62	.31
Observations	735	735	735	735

Source: Dossiers Annuels 1999–2015. Panel regressions of inflows into Euro-denominated contracts and unit-linked contracts on beginning-of-year fund reserves, with insurer fixed effects and year fixed effects. Regressions in columns (1) and (2) are weighted by account value and regressions in columns (3) and (4) are equal weighted. Standard errors are double clustered by insurer and year. ***, **, and * mean statistically significant at the 1%, 5%, and 10% levels, respectively.

horizon.

There are two potential mechanisms by investor flows react to the level of fund reserves. First, investors may select insurers based on reasons unrelated to the level of reserves, but once they have selected an insurer, their choice of life insurance product may depend on the level of reserves. Life insurers offer two types of products: Euro-denominated contracts, which are the contracts we focus on in this paper and represent 80% of the market, and unit-linked contracts, which work like traditional mutual funds, benefit from the same tax treatment as Euro-denominated contracts, and represent the remaining 20% of the market. Anecdotal evidence suggests that insurers try to nudge investors into buying unit-linked contracts when they want to reduce inflows into Euro-denominated contracts.[11] In this case, inflows into unit-linked contracts should be negatively related to the level of fund reserves.

Second, investors may understand that the level of reserves predicts future contract returns and select insurers based on this. In this case, inflows into unit-linked contracts should be unrelated to the level of fund reserves. They may even be positively related to the level of fund reserves if Euro-denominated contracts and unit-linked contracts are complements. It might be so for instance if some investors want to buy both types of contracts and find it easier to buy both from the same insurer. In practice, insurers frequently advise investors to buy both types of contracts for diversification purposes. For investors following this advice, there is a tax advantage to buy both

[11] See for instance EIOPA (2016) and *L'Express: Votre Argent*, March 2016, "Assurance vie: pourquoi et comment les assureurs vous poussent à prendre des risques".

contracts from the same insurer, because transferring funds across contracts with the same insurer does not lead to the realization of gains for tax purposes.

To disentangle between these two mechanisms, we analyze how inflows into unit-linked contracts depends on the level of fund reserves. The coefficient on reserves in column (2) is negative and statistically significant at the 5% level, consistent with the first mechanism. The size of the coefficient is not directly comparable to the one in column (1), because inflows into each type of contract are normalized by account value for the same type of contract. Since Euro-denominated contracts represent 80% of total account value and unit-linked contracts 20%, a euro-for-euro comparison of the effect of fund reserves on inflows into each type of contract requires to divide by four the coefficient for unit-linked inflows. Thus, the effect of reserves on inflows has a similar magnitude in euro terms for each type of contract, and opposite signs. This is consistent with the interpretation that insurers are able to channel inflows from Euro-denominated contracts into unit-linked contracts when they have low reserves.

Column (3) and (4) report the results of equal weighted regressions. The positive relation between fund reserves and subsequent inflows into Euro-denominated contracts is robust and the magnitude is similar to the value-weighted regression. However, there is no longer a significant relation between fund reserves and subsequent inflows into unit-linked contracts. It suggests that only large insurers are able to channel inflows into Euro-denominated or unit-linked contracts as a function of their level of reserves.

These results shed light on why life insurers are able to achieve large amounts of intergenerational redistribution in a decentralized market, whereas theory predicts that this is not possible in a perfectly competitive market. Imperfectly elastic demand allows insurers to smooth contract returns through reserves management even though doing so create contract return predictability, because investor flows react only weakly to predictable returns.

7 Concluding Remarks

The main finding of this paper is that Euro-denominated life insurance contracts, which are the most widespread type of life insurance products in Europe, achieve large amount of intergenerational redistribution. Using regulatory and survey data from the national supervisor for the French market, we estimate that the average transfer across generations of investors is 1.4% of total account value per year, or 0.8% of GDP. Given Allen and Gale (1997) show that intergenerational risk sharing does not survive perfectly competitive markets, our evidence of significant intergenerational transfers suggests a form of market imperfection, that makes intergenerational risk possible. This notion is confirmed by our evidence of a limited sensitivity of inflows to predictable returns.

Intergenerational risk sharing is unique to Euro-denominated contracts, which use reserves to smooth returns across generations of investors. This is in contrast to traditional savings products

like mutual funds and unit-linked life insurance contracts, which deliver the asset return to investors with no insurance against market risk. Euro-denominated contracts also stand in contrast to variable annuities with guaranteed returns sold by life insurers (Koijen and Yogo (2017)) in the US and structured savings products sold by banks in Europe and the US (Célérier and Vallée (2017)), because they generate insurance against aggregate risk through intergenerational risk sharing whereas the latter types of products implement purely cross-sectional risk sharing. This specificity of Euro-denominated contracts suggests that an aggregate shift from these contracts to unit-linked contracts would change the pattern of aggregate risk sharing in the economy, possibly to a sizeable extent given the large aggregate amount of intergenerational transfers we estimate. An additional potential benefit of intergenerational risk sharing is that it may allow insurers to hold more risky assets, as risk is better shared (Gollier (2008)).

References

Allen, Franklin and Douglas Gale. 1997. "Financial Markets, Intermediaries, and Intertemporal Smoothing." *Journal of Political Economy* 105 (3):523–546.

Ball, Laurence and N Gregory Mankiw. 2007. "Intergenerational risk sharing in the spirit of Arrow, Debreu, and Rawls, with applications to social security design." *Journal of Political Economy* 115 (4):523–547.

Célérier, Claire and Boris Vallée. 2017. "Catering to Investors Through Security Design: Headline Rate and Complexity." *Quarterly Journal of Economics* (Forthcoming).

Darpeix, Pierre-Emmanuel. 2016. "Le taux technique en assurance vie (Code des Assurances)." *Analyses et Synthèses (Banque de France)* .

EIOPA. 2016. "EIOPA Fifth Consumer Trends Report." *EIOPA Annual Reports* .

FFSA-GEMA. 2013. "Les rachats des contrats d'assurance vie après 60 ans." *Document de travail* Conseil d'Orientation des Retraites (8).

Gollier, Christian. 2008. "Intergenerational risk-sharing and risk-taking of a pension fund." *Journal of Public Economics* 92 (5):1463–1485.

Gordon, Roger H and Hal R Varian. 1988. "Intergenerational risk sharing." *Journal of Public economics* 37 (2):185–202.

Insurance Europe. 2016. "European Insurance – Key Facts."

Koijen, Ralph SJ and Motohiro Yogo. 2017. "The Fragility of Market Risk Insurance." .

Stiglitz, J. 1983. "On the Relevance of Public Financial Policy: Indexation, Price Rigidities and Optimal Monetary Policies."

A Return Dispersion Across Contracts

In this appendix, we study the dispersion of contract return across contracts sold by the same insurer. In particular, we want to assess whether one of the following describes the data accurately: (a) contract return is constant across all contracts sold by the same insurer; (b) contract return is constant up to a contract-specific fixed effect. The goal of this test is to determine whether assuming a constant return across contracts (up to a constant in case (b)) is a good approximation of the extent of intertemporal return smoothing experienced by investors. The answer is yes if fluctuations in contract return is the same across all contracts of the same insurer (i.e., (a) or (b) hold) whereas the answer is no if there is significant idiosyncratic fluctuations in contract returns. To test (a), we study whether the dispersion of contract returns in the cross-section of contracts is small relative to the time-series dispersion of contract returns. To test (b), we calculate the cross-sectional dispersion after controlling for a product-specific fixed effect.

A.1 Data

We use data from a survey (*Enquête Revalo*) run by the insurance supervisor that collects information at the product level. Insurers are required to report for every life insurance product they have sold in the past and that still has in-force contracts the following information: product name, contract category, date of first commercialization, minimum guaranteed rate, a dummy variable indicating whether the product is open to new subscriptions; and the value at the end of the survey year and of the previous year of the following variables: number of in-force contracts, total account value, and contract return.

The survey has been run for all the years from 2011 to 2015. It does not include a unique product identifier to track a given product across survey years, which we will need to test whether contract return is well explained by a product fixed effect. We construct a unique product identifier using the following procedure:

- We match contracts on name, contract category, date of first commercialization, and minimum guaranteed rate across the five waves of the survey. We refer to the resulting level of observation as a product.
- If a product has several observations in a given year, we check whether they all have the same contract return. If not, we drop all the observations associated with this product-year. If yes, we collapse them into a single product-year observation by summing account values across observations.
- We drop products with gaps in their time-series.
- For each pair of subsequent years of a product, we check whether the lagged values of account value and contract return reported in the survey for year t match the current values reported

in the survey for year $t-1$. When the discrepancy is larger than 1% for at least one of the two variables, we drop all the years associated with the product.

This procedure allows us to assign a unique product identifier for 71% of account value.

A.2 Return Dispersion

To compute the time-series dispersion, we use the longer time-series of regulatory data from 1999 to 2015 and we restrict the sample to insurers that we observe throughout the sample period. First, we calculate the average contract return at the insurer-year level. Then, we compute the time-series standard deviation of this variable at the insurer level. Finally, we take the average across all insurers. The time-series standard deviation of annual contract returns is 0.97 percentage points.

To compute the cross-contract dispersion, we use the survey data at the product level. First, we calculate the value-weighted standard deviation of contract return across products at the insurer-year level. Then, we take the average over all insurer-year observations. The cross-contract standard deviation of annual contract returns is 0.30 percentage points. The dispersion of returns across contracts is thus three times smaller than the time-series dispersion. Therefore, although condition (a) of no return dispersion across contracts does not perfectly hold, it would still be a reasonable approximation of the data relative to the large time-series variations.

In fact, we only need condition (b) that the dispersion in contract returns reflects a contract fixed effects. First, we restrict the sample to products that we observe in every survey year and we compute the residual of a panel regression of annual product return on product fixed effect. Then, we calculate the value-weighted standard deviation of the residual across products at the insurer-year level. Finally, we take the average over all insurer-year observations. The cross-sectional standard deviation after absorbing product fixed effects drops to 0.10 percentage points, that is, 10 times less than the time-series dispersion. We conclude that condition (b) that contract return is constant across all contracts up to a product fixed effect is a good approximation of the data.

B Deviations From Historical Cost Accounting

There are two deviations from historical cost accounting principles to force insurers to recognize large unrealized losses. First, when an asset has "lasting and significant" unrealized capital losses, its book value is partially adjusted downwards through the creation of a provision on the asset side of the balance sheet (*provision pour dépréciation durable*) to reflect the paper loss. This adjustment is booked as a realized loss. It thus increases unrealized gains (makes them less negative). If the return credited to contracts and to the insurer are held constant, this realized loss reduces the profit-sharing reserve account, and total fund reserves are not affected. However, the goal of this provision is to induce the insurer to reduce the return credited to investor accounts and thus reduce profit-sharing reserves by less than the realized loss, which increases total fund reserves.

The second deviation from historical cost accounting is that, when the market value of the fund portfolio of non-fixed income securities is less than the book value, the overall paper loss is recognized through a provision on the liability side of the balance sheet (*provision pour risque d'exigibilité*). This is booked as a loss. Therefore, if the return credited to contracts and to the insurer are held constant, this reduces the profit-sharing reserve account and thus total fund reserves. However, the goal of this provision is to induce the insurer to reduce the return credited to investor accounts and thus offset the reduction in the amount of fund reserves.

C Variables Construction

The raw data are from *Dossiers Annuels* from 1999 to 2015.

Account value `Provisions d'assurance vie à l'ouverture` (beginning-of-year account value) and `Provisions d'assurance vie à la clôture` (end-of-year account value) in C1V1–C1V3 statements summed over contract categories 1, 2, 4, 5, and 7, which is the set of contracts backed by the same pool of underlying assets and associated to the same pool of reserves. The main excluded contract categories are 8 and 9, which are unit-linked contracts.

Profit-sharing reserves `Provisions pour participations aux bénéfices et ristournes` in BILPV statement.

Capitalization reserves `Réserve de capitalisation` in C5P1 statement.

Unrealized gains Book value (`Valeur nette`) minus market value (`Valeur de réalisation`) of assets underlying life insurance contracts measured as `Placements représentatifs des provisions techniques` minus `Actifs représentatifs des unités de compte` in N3BJ statement.

Inflows `Sous-total primes nettes` in C1V1–C1V3 statements summed over contract categories 1, 2, 4, 5, and 7. It includes initial cash deposits at subscription and subsequent cash deposits in existing contracts. The inflow rate is calculated as inflow amount divided by beginning-of-year account value plus one half of net flows.

Outflows `Sinistres et capitaux payés` plus `Rachats payés` in C1V1–C1V3 statements summed over contract categories 1, 2, 4, 5, and 7. It includes partial and full redemptions, either voluntary or at death of investor. The outflow rate is calculated as outflow amount divided by beginning-of-year account value plus one half of net flows.

Contract return We calculate the value-weighted average contract return as the amount credited to investor accounts divided by beginning-of-year account value plus one half of net flows (i.e., we assume that flows are uniformly distributed throughout the year and thus receive on average one half of the annual contract return). The amount credited to investor accounts is measured as `Intérêts techniques incorporés aux provisions d'assurance vie` plus `Participations aux bénéfices` plus `Intérêts techniques inclus dans exercice prestations` plus `Participations aux bénéfices incorporées dans exercice prestations` in C1V1–C1V3 statements summed over contract categories 1, 2, 4, 5, and 7.

Asset return We sum the three components of asset returns, which are reported separately in insurers' financial statement. First, `Produits des placements nets de charges` in C1V1–C1V3 statements summed over contract categories 1, 2, 4, 5, and 7, measures asset yield (dividends on non-fixed income securities plus yield on fixed income securities) and realized gains and losses on non-fixed income securities, net of operating costs. Second, the change in capitalization reserves account value reflects realized gains and losses on fixed income securities. Third, the change in unrealized gains captures measures unrealized gains. We calculate asset return as the sum of these three components divided by account value plus fund reserves.

D Account Value by Generation

We describe in this appendix how we estimate account value by generation from insurer-level account value, inflows, and outflows, under parametric assumptions on the inflow rate and the outflow rate.

Regarding inflows, we assume that investors only make one-off investments. They make an initial deposit when they buy a contract and never add more cash at a subsequent date. Regarding outflows, we assume that investors only proceed to full redemptions and that the redemption rate does not depend on contract age at the insurer-year level.

We call generation (t^0, t^1) the set of investors who buy their contract in year t^0 and redeem it in year t^1, for $t^0 < t^1$. We denote $V_t(t^0, t^1)$ the account value of generation (t^0, t^1) at the end of year t and by $V_t^+(t^0, t^1)$ and $V_t^-(t^0, t^1)$ their inflows and outflows, respectively, during year t. Under the maintained assumption that inflows and outflows are uniformly distributed throughout the year and are entitled to one half of the annual contract return, account value of generation (t^0, t^1) evolves according to

$$V_{t^0-1}(t^0, t^1) = 0, \tag{10}$$

$$V_t(t^0, t^1) = (1 + y_t)V_{t-1}(t^0, t^1) + (1 + y_t/2)(V_t^+(t^0, t^1) - V_t^-(t^0, t^1)), \quad t = t^0, ..., t^1 - 1, \tag{11}$$

$$V_{t^1}(t^0, t^1) = 0, \tag{12}$$

where y_t is the net-of-fees contract return. The data only reports gross-of-fees contract return. Since we observe beginning-of-year account value, inflows, and outflows, and end-of-year account value, we back out the net-of-fees contract return from the law of motion of total account value

$$V_t = (1 + y_t)V_{t-1}(1 + y_t/2)(V_t^+ - V_t^-). \tag{13}$$

Finally, the assumption of no inflow after initial subscription implies

$$V_t^+(t^0, t^1) = 0, \qquad t > t^0. \tag{14}$$

The assumption of no partial redemption before exit implies

$$V_t^-(t^0, t^1) = 0, \qquad t < t^1. \tag{15}$$

The assumption of outflow rate independent of contract age at the insurer-year level implies

$$\frac{V_t^-(t^0, t)}{V_{t-1}^-(t^0, t)} = \frac{V_t^-}{V_{t-1}}, \qquad t > t^0. \tag{16}$$