The Review of Financial Studies

Can Risk Be Shared across Investor Cohorts? Evidence from a Popular Savings Product

Johan Hombert
HEC Paris and CEPR, France

Victor Lyonnet
The Ohio State University, USA

We study how retail savings products can share market risk across investor cohorts, thereby completing financial markets. Financial intermediaries smooth returns by varying reserves, which are passed on between successive investor cohorts, thereby redistributing wealth across cohorts. Using data on euro contracts sold by life insurers in France, we estimate this redistribution to be large: 0.8% of GDP. We develop and provide evidence for a model in which low investor sophistication, while leading to individually suboptimal decisions, improves risk sharing by allowing intercohort risk sharing. (*JEL* G22, G52)

Received July 31, 2020; editorial decision July 21, 2022 by Editor Ralph Koijen. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Even in well-developed financial markets, aggregate risk can be shared only among investors participating in the market when this risk is realized. This limit to risk sharing sometimes results in significant losses: in 2008, a perfectly diversified portfolio of stocks lost 40% of its value. Superior risk sharing can be

---

We thank Joseph Briggs, Anne-Laure Delatte, Christian Gollier, Valentin Haddad, Anastasia Kartasheva, Robert Novy-Marx, Ishita Sen, Boris Vallée, and Moto Yogo; conference participants at the 2017 NBER Insurance Project Workshop, European Winter Finance Conference, FIRS, CSEF-IGIER Conference, SFS Cavalcade, Paul Woolley Annual Conference, Finance Meets Insurance Conference, EFA, and AFA; seminar participants at Autorité de Contrôle Prudentiel et de Résolution, Chinese University of Hong Kong, ESCP Europe, Hong Kong University, Hong Kong University of Science and Technology, Singapore Management University, Nanyang Technological University, Deutsche Bundesbank, University of Toronto, and MIT-Sloan for helpful comments. We thank Autorité de Contrôle Prudentiel et de Résolution, and 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, for access and assistance in putting together the data as well as comments. Hombert acknowledges financial support from the European Research Council under the EU's Horizon 2020 research and innovation program [grant no. 101002355] LIFE, and from Investissements d'Avenir Labex [ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047]. A preliminary version of this paper circulated under the title "Intergenerational Risk Sharing in Life Insurance: Evidence from France." All errors remain our own. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Johan Hombert, hombert@hec.fr.

The Review of Financial Studies 35 (2022) 5387–5437

https://doi.org/10.1093/rfs/hhac054 Advance Access publication August 24, 2022

achieved by diversifying risk intertemporally across investor cohorts (Gordon and Varian 1988), but financial markets do not allow current and future investor cohorts to trade with each other (i.e., financial markets are incomplete).[1] In principle, long-lived financial intermediaries can complete the market by transferring risk between successive cohorts. However, Allen and Gale (1997) show that competition in the savings market unravels intercohort risk sharing implemented by an intermediary under the assumption that investors always identify and pick the best investment opportunities. This paper theoretically and empirically shows how intercohort risk sharing can be achieved when one relaxes that assumption.

Our first contribution is to show how one of the most popular savings products in Europe shares market risk across investor cohorts. These products are sold by life insurers to retail investors. Their name varies by country: "euro contracts," "participating contracts," and so on. In Europe, as of 2017, these products represent 15% of households' financial wealth, and 60% of life insurers' provisions (statistics from EIOPA and Eurostat). We focus on the 1.4-trillion-euro French market, where they are called *euro contracts* and are pure savings products (i.e., they are not traditional life insurance products). Euro contracts work as follows. When a retail investor buys a contract, an account is created, on which she can invest and withdraw cash at any time. In turn, each insurer pools the cash deposited by all its investors into a single fund invested in a portfolio of assets.

The fund holds reserves that vary to offset shocks to asset returns. Reserves increase when asset returns are high and decrease when asset returns are low, so that contract returns are an order of magnitude less volatile than funds asset returns. Reserves are pooled across investors and passed on between investors who hold contracts over different periods. Throughout the paper, we define the set of investors who hold a contract over a given period as an investor cohort. The fact that reserves are passed on between successive investor cohorts causes redistribution across cohorts. Investors receive a transfer from reserves when asset returns are low, and contribute to reserves when asset returns are high. Part of these transfers net out within investors' holding period. Only the *net* transfer received from or contributed to reserves over investors' holding period represents redistribution across investor cohorts. Using regulatory and survey data from France, we estimate that intercohort redistribution amounts to 1.4% of total account value per year, which represents 17 billion euros redistributed across investor cohorts every year, or 0.8% of gross domestic product (GDP).

This finding challenges the notion that intercohort risk sharing cannot be achieved in competitive markets. Allen and Gale (1997) study savings contracts that share market risk across investor cohorts through a reserve mechanism

[1] Gollier (2008) estimates that intercohort risk sharing increases the certainty equivalent of capital income by 25% relative to an economy without intercohort risk sharing.

similar to that of euro contracts. They study two polar cases, showing that (a) a financial intermediary can implement intercohort risk sharing if it is protected from competition, that is, if investors must invest with the intermediary regardless of the reserves level, and (b) competition unravels intercohort risk sharing if investors are fully strategic and hence invest in contracts only when reserves are high and opt out when reserves are low, that is, if demand for contracts is infinitely elastic to reserves. In practice, (a) does not apply to euro contracts, because insurers compete with each other as well as with alternative investment options. Moreover, the large amount of intercohort redistribution we observe in the data rules out the assumption of infinitely elastic demand in (b). To our knowledge, no theoretical framework exists to analyze intercohort risk sharing in real-world euro contracts. Our second contribution is to study, theoretically and empirically, the conditions enabling intercohort risk sharing.

We develop a model in which long-lived intermediaries compete in selling savings products to successive cohorts of investors. We characterize how the amount of intercohort risk sharing depends on the elasticity of demand for contracts with respect to the expected contract return conditional on reserves. The model nests the two polar cases of perfectly inelastic demand and perfectly elastic demand that have been studied in the literature. In line with this literature, we show asset risk can perfectly be shared across investor cohorts when demand is inelastic. Instead, when demand is elastic, investors behave opportunistically and exploit the predictability of contract returns: they flow into (out of) contracts when reserves are high (low), partially unravelling risk sharing across cohorts. In the limit when demand is perfectly elastic, intercohort risk sharing fully unravels so that the savings products are akin to pass-through mutual funds, in line with Allen and Gale (1997). In a nutshell, the equilibrium level of intercohort risk sharing crucially (and monotonically) depends on demand elasticity.

The intercohort risk sharing achieved by the contracts cannot be replicated using market instruments. In this sense, the contracts complete markets. The reason being that contracts exploit a dimension of risk sharing—cross-cohort risk sharing—that cannot be achieved in financial markets, implying the contract has a lower risk exposure than the underlying insurers asset portfolio.

Our model shows we can estimate demand elasticity and the corresponding amount of intercohort risk sharing from two moments in our data. The first moment is the regression coefficient of contract return on contemporaneous asset return, conditional on the level of reserves. When demand is inelastic, contracts share risk across investor cohorts. In this case, the contract return depends on the level of reserves, but not on contemporaneous asset return beyond its effect on reserves. The intuition behind this idea is similar to that behind the permanent income hypothesis, whereby optimal consumption does not depend on current income beyond its effect on permanent income. By contrast, when demand is elastic, little intercohort risk sharing occurs, and the contract return strongly depends on the contemporaneous asset return. We estimate panel regressions and, controlling for the level of reserves, we show the contract return does not depend on the asset return in the current year. Therefore, the evidence is consistent with low demand elasticity and, correspondingly, sizeable intercohort risk sharing.

The second moment that is informative about demand elasticity is the regression coefficient of investor flows on reserves. A high level of reserves predicts high expected contract returns, so that the sensitivity of flows to reserves is directly related to demand elasticity. We run panel regressions and find that the sensitivity of flows to reserves is close to zero, again consistent with low demand elasticity. One issue when regressing flows on reserves is that reserves are potentially endogenous to unobserved demand shocks, which is a standard issue one encounters when estimating demand functions by regressing quantity on price. Our model shows the past asset return is a valid instrument for reserves to estimate the sensitivity of flow to reserves. Instrumenting reserves, the sensitivity of flows to reserves remains close to zero.

Why is demand inelastic to reserves, allowing for intercohort risk sharing? We first rule out explanations based on switching costs related to taxes or fees. In particular, we study investors buying a new contract. These new investors do not face switching costs. Despite no switching costs, and although high reserves predict high contract returns, we find new investors' flows do not react to reserves. An alternative hypothesis is that investors are inelastic to expected contract return because they have strong preferences for other contract characteristics, such as risk exposures. Inconsistent with this hypothesis, we show that high- and low-reserves contracts are similar in their exposure to both systematic and idiosyncratic risks.

We hypothesize that demand is inelastic to reserves because investors lack the knowledge to predict contract returns using reserves. In line with this hypothesis, we show that contracts held by investors with a small investment amount (less than 250,000 euros) have a flow-reserves sensitivity indistinguishable from zero, whereas contracts with a large investment amount (greater than 250,000 euros) exhibit a positive and significant flow-reserves sensitivity. This result helps us further rule out the hypothesis that investors sort into contracts based on other contract characteristics, such as risk exposures: rationalizing this result under this alternative hypothesis would require to explain why the heterogeneity in risk preferences is significantly smaller among wealthier investors than among less wealthy investors. Instead, this result is consistent with interpreting the investment amount as a proxy for wealth and financial sophistication, whereby less sophisticated investors fail to predict contract returns using reserves.[2] Differences in demand elasticity

[2] Using data from a French life insurer, Bianchi (2018) studies households' portfolio allocation between mutual funds and euro contracts. He constructs a survey-based measure of financial literacy and shows this measure is highly correlated with household wealth.

across investors with different wealth levels can arise if investors must incur a fixed cost to acquire the knowledge or information necessary to understand the sources of contract return predictability (Lusardi and Mitchell 2014).

Perhaps paradoxically, the lack of household financial sophistication enables more risk sharing than would be possible if households were perfectly informed and acted accordingly. The idea that ignoring privately valuable information can be socially beneficial because it improves risk sharing goes back to Hirshleifer (1971). Our results are an illustration of this principle in the context of aggregate risk sharing: investor inertia, while individually suboptimal, improves (intercohort) risk sharing.[3]

To quantity the welfare benefit of intercohort risk sharing, we consider a simple portfolio choice model in which the investor can invest in the risk-free asset, the risky asset, and evaluate the welfare gains from adding the euro contract to the set of investable assets. We define the welfare benefit of intercohort risk sharing as the certainty equivalent gain of adding the contract to the set of investable assets. We estimate it to be 90 basis points (bps) per year for an investment horizon equal to the sample average of 12 years. The welfare benefit decreases with the investment horizon. The reason being that return smoothing is less effective at a long horizon because shocks to asset return are absorbed by reserves in the short run but are eventually passed through to investors as reserves are eventually distributed though contract returns. We also show that, perhaps surprisingly, the welfare benefit does not depend much on risk aversion. The reason being that the contract is quite safe, so the investor optimally uses it as a substitute for the risk-free asset. The magnitude of the welfare benefit is therefore determined by the premium earned by the contract over the risk-free asset and not so much by the reduction in risk relative to the risky asset.

We quantity the welfare benefit of investor stickiness. We run counterfactual analyses for different values of the elasticity of demand. We find that, if the elasticity increases from zero to the level of elasticity of deposits estimated by Drechsler, Savov, and Schnabl (2017), the contract return becomes more volatile and thus a less good substitute for the risk-free asset, reducing its welfare benefit by a factor of two.

Our paper contributes to the recent literature that studies savings products implementing cross-sectional sharing of aggregate risk between investors and the financial intermediary. Examples include variable annuities with return guarantees sold by U.S. life insurers (Koijen and Yogo 2022; Ellul et al. 2021) and structured products sold by European banks (Célérier and Vallée 2017; Calvet et al. forthcoming). In these products, the intermediary bears part of

[3] In the context of health insurance, Handel (2013) shows that consumer health plan choice inertia reduces adverse selection and hence improves (cross-sectional) risk sharing. Hortaçsu and Syverson (2004), Drechsler, Savov, and Schnabl (2017), and Koijen and Yogo (2022) study the implications of investor inertia on competition between financial intermediaries.

the risk by hedging investors' returns from market risk. Such cross-sectional risk sharing is also at play in euro contracts, but we show it is an order of magnitude smaller than intertemporal risk sharing across investor cohorts. Crucially, cross-sectional risk sharing hinges on intermediaries' risk-bearing capacity.[4] By contrast, euro contracts shift most of the risk to households and share it across cohorts.

Similar to euro contracts, defined benefits pension plans contain an element of intercohort risk sharing, because defined benefits sponsors can spread shocks across cohorts by adjusting the contributions of futures employees, and they also can be bailed out by future taxpayers (Novy-Marx and Rauh 2011; Novy-Marx and Rauh 2014). However, the market structure and thus the determinants of demand elasticity are different for euro contracts and defined benefits plans. Greenwood and Vissing-Jorgensen (2019) study the implications of pension funds and insurance companies' behavior for asset prices, and Scharfstein (2018) examines their role in shaping the financial system.

We also contribute to the theoretical literature on the private implementation of intercohort risk sharing. The notion that financial markets cannot implement intercohort risk sharing because they do not allow current and future investor cohorts to trade with each other goes back at least to Stiglitz (1983) and Gordon and Varian (1988), whereas Ball and Mankiw (2007) study intercohort risk sharing in a hypothetical economy in which current investors can trade with future investors. Allen and Gale (1997) and Gollier (2008) study how intercohort risk sharing can be implemented by an intermediary having monopoly power over households' savings. Allen and Gale (1997) show intercohort risk sharing unravels if investor demand is infinitely elastic to reserves. We extend this literature by considering the case of finite elasticity, showing the equilibrium level of intercohort risk sharing decreases monotonically from perfect to nonexistent as the elasticity increases from zero to infinity.

1. Euro Contracts

1.1 Institutional framework

European life insurers sell savings contracts designed to implement intercohort risk sharing. We study the market for these contracts in France, where they are called euro contracts. Despite being sold by life insurers, euro contracts are pure savings products which do not entail insurance against longevity or mortality risk. Life insurers selling euro contracts can be subsidiaries of insurance

[4] Life insurers' product supply shifts inward when their capital position weakens (Koijen and Yogo 2015; Koijen and Yogo 2022; Ge forthcoming; Sen and Humphry 2020). Their capital position also affects their asset portfolio choices (Ellul et al. 2015; Becker and Ivashina 2015; Ge and Weisbach 2021).

holding companies, subsidiaries of bank holding companies, or stand-alone life insurance companies.[5]

Investments in euro contract amount to 1.4 trillion euros as of 2015, which represent one-third of French household financial wealth (Insee 2016). Another one-third of household financial wealth is invested in risky securities and investment funds, held directly or through special vehicles. The last third is invested in short-term instruments including checking accounts, savings accounts, and regulated savings products, such as the ones analyzed at the end of Section 4.3.

When an investor buys a euro contract, the insurer creates an account on which the investor can deposit and withdraw cash at any time. We refer to the cash balance on the investor's account as the account value. The insurer pools the cash deposited by all investors in a single fund called the euro contract fund, which is invested in a portfolio of assets. Regulation defines the set of assets euro contract funds can invest in. This set includes most assets issued in Organisation for Economic Co-operation and Development (OECD) countries, such as sovereign and corporate bonds, loans, public and private equities, real estate, and shares in investment funds holding such assets.[6] Table 1 reports the summary statistics on the asset composition of French euro contract funds.

At the end of each calendar year, the insurer chooses how to split the asset return into three parts: the contract return paid to current investors, the change in reserves, and the insurer's profit. This flexibility allows insurers to provide investors with insurance against market risk. The contract return chosen by the insurer can differ from the return on the assets held by the insurer if the insurer chooses to use reserves or its profit as a buffer.

Reserves have two key features that jointly give rise to intercohort risk sharing. First, reserves are owed to investors. While the insurer can choose how to split the asset return between current contract return, change in reserves, and insurer profit, this choice is subject to the regulatory constraint that the sum of the first two components must be at least 85% of asset income. Therefore, the insurer cannot transfer funds from reserves to its equity, because its profit must be at most 15% of the asset income.[7] However, the insurer can freely choose how much funds to transfer between reserves and investors' accounts to smooth contract returns over time, that is, how much market risk insurance it provides to its investors.

Second, reserves are pooled across all investors and passed on between successive investor cohorts. In particular, new investors share in reserves

5 Mutual insurance companies, pension institutions, and reinsurance companies can also sell euro contracts. These institutions are subject to a different regulation and account for only 4% of aggregate provisions (ACPR 2016). We abstract from them in the empirical analysis.

6 See article R.332 of *Code des Assurances*. Asset regulation also includes diversification requirements, such as preventing insurers from investing more than 65% of their asset portfolio in a given asset class. These diversification requirements are not binding in our sample.

7 See Internet Appendix B.2 for a detailed description of the regulatory framework.

Table 1
Summary statistics

	Mean	SD	P25	P50	P75	N
A. Regulatory filings						
Account value (bn euro)	13.9	30.1	0.9	3.1	11.9	978
Inflows (% account value)	10.5	3.8	7.8	10.5	12.3	978
Outflows (% account value)	8.1	2.0	7.1	7.9	8.8	978
Reserves (% account value)	10.9	6.8	6.7	10.5	14.3	978
Portfolio share: bonds (%)	80.4	8.0	75.4	81.5	85.6	978
Portfolio share: stocks (%)	13.5	6.3	10.0	12.5	15.7	978
Asset return (%)	4.9	4.4	2.1	4.4	7.5	978
Contract return (%)	4.0	0.9	3.3	4.0	4.5	978
B. Prospectus data						
Management fee (%)	.7	.13	.64	.73	.77	578
Entry fee (%)	3.3	.87	3	3.5	3.8	578
C. Survey data						
Net-of-fees return (%)	2.7	.45	2.4	2.8	3	13,672
Minimum guaranteed return (%)	.35	.73	0	0	0	13,672

Panel A presents regulatory filings data at the insurer-year level for 76 insurers over 2000–2015. All statistics (except for account value) are weighted by the insurer share in aggregate account value in the current year. *Account value* is total account value at year-end in constant 2015 billion euros. *Inflows* are inflows (premiums) divided by beginning-of-year account value plus one-half of net flows. *Outflows* are outflows (redemptions plus payment at contract termination) divided by beginning-of-year account value plus one-half of net flows. *Reserves* is total reserves divided by year-end account value. *Portfolio share: bonds* is the share of (corporate and sovereign) bonds, held either directly or through funds, in the asset portfolio. *Portfolio share: stocks* is the share of stocks, held either directly or through funds, in the asset portfolio. *Asset return* is the asset return. *Contract return* is the average before-fees contract return. Panel B presents prospectus data on fees at the insurer-year level for 48 insurers over 2000–2015. *Management fees* is the average management fees across contracts offered by the insurer and open to new subscriptions in the current year. *Entry fees* is the average entry fees across contracts offered by the insurer and open to new subscriptions in the current year. Panel C presents survey data at the contract-year level for about 2,700 outstanding contracts per year from 56 insurers over 2011–2015. *Net-of-fees return* is the contract net-of-fees return. *Minimum guaranteed return* is the before-fees minimum return guaranteed by the insurer.

accumulated by previous investors, and investors redeeming their contracts give up their share of reserves. The pooling of reserves across investor cohorts happens because the contract return on a given contract is the same for all investors who hold this contract, regardless of when they bought the contract.

Insurers offer a range of contracts, for instance, a basic contract and a premium contract with a minimum investment amount and a lower fee. Since they are allowed to pay different returns on different contracts, insurers could close existing contracts to new subscriptions when reserves are high, and create a new vintage of contracts on which they will pay higher returns. Doing so would undo reserve pooling across investor cohorts. Using data at the contract level, we show in Section 4.2 that insurers do not follow this strategy; therefore, reserves are effectively pooled across investor cohorts.

The rest of this section summarizes the other main features of the institutional framework.

Minimum return guarantees. Euro contracts include a minimum guaranteed return for the annual contract return. This minimum guaranteed return is fixed at the subscription of the contract. Against a backdrop of decreasing interest rates, French insurers have strongly reduced guaranteed rates since the 1990s. The average guaranteed rate is now close to zero (Darpeix 2016), such that minimum return guarantees are not binding for the vast majority of contracts during the sample period (see Section 1.2).

Fees. Insurers usually charge entry fees when investors deposit cash on their account (front-end loads), and annual management fees. They are not allowed to charge exit fees (back-end loads). By regulation, at least 90% of insurers' annual technical income, equal to fees minus the insurer's operating costs, must be paid to investors. This fraction goes up to 100% if technical income is negative. Insurers can choose to pay investors immediately via the contract return, or later by increasing reserves. The implication of this regulation is that insurers cannot extract money from reserves by raising fees on new investors, because 90% (or 100%) of these fees must eventually be returned to investors. Neither can they extract money from reserves by increasing management fees, which are contractually set for each contract, nor by manipulating their operating costs, which accrue to their employees or brokers.

Taxes. Contract returns are taxed upon withdrawal at a rate that depends on the age of the contract at the time of withdrawal. The tax rate is decreasing in contract age for the first 8 years of the contract. This creates a potential switching cost that we analyze in Section 5.1. The tax treatment of euro contracts is the same as that of unit-linked contracts, which are investment vehicles also sold by life insurers through which households can hold mutual funds. Therefore, households can invest in mutual funds at the same fiscal cost as in euro contracts.

Solvency regulation. Insurers are subject to the Solvency I regulation during the sample period. This regulation imposes that insurers hold a minimum amount of capital such that the sum of capital and unrealized capital gains is at least equal to 4% of total account value.[8] These capital requirements are independent of the portfolio asset composition or the minimum return guarantees. Solvency II came into effect in 2016, which is after the end of the sample. Although insurers anticipated the new regulation, Solvency II did not change the regulatory framework at the basis of intercohort risk sharing, such as the constraint that at least 85% of asset returns must eventually be paid out to investors and how reserves are created and can be used. Under Solvency II, capital requirements depend on the portfolio asset composition and the return guarantees issued. Despite the new regulation, the insurer supervisor found no significant change in asset riskiness in response to Solvency II (Baddou et al. 2016). There was no effect on minimum guaranteed returns, which were already close to zero during the sample period (Darpeix 2016).

8 See article L.334 of *Code des Assurances.* We show in Internet Appendix B.3 that insurers' equity alone is never below 4% of total account value throughout the sample period, including during the 2008 Global Financial Crisis and the 2011–2012 European sovereign debt crisis.

1.2 Data and summary statistics

Our main source of data comprises regulatory filings obtained from the national insurance supervisor (*Autorité de Contrôle Prudentiel et de Résolution*) for the years 1999 to 2015. The data cover all companies with life insurance operations in France and contain detailed financial statements.[9] We focus on insurance companies with more than 10 million euros of euro contracts total account value. Because we need lagged values to calculate the change in reserves, the sample period of our analysis is 2000–2015. The final sample contains 76 insurers and 978 insurer-year observations.

Panel A of Table 1 reports summary statistics from the regulatory filings. Statistics on ratios are value-weighted by the insurer's share in aggregate account value in the current year. The average (median) insurer has 13.9 (3.1) billion euros of account value. Inflows (premiums), which include cash deposited in newly opened contracts and in existing contracts, represent, on average, 10.5% of account value per year. Outflows, which include partial and full redemptions, either voluntarily or at contract termination (investor death), represent on average 8.1% of account value per year. The combination of positive net flows and compounded contract returns generates an increasing trend in aggregate account value as plotted in Internet Appendix Figure B.1. Aggregate account value grows from 500 billion euros in 2000 to 1,200 billion euros in 2015 (all amounts are in constant 2015 euros). Market concentration is relatively low, with a Herfindahl-Hirschman Index around 800 and total market shares of the top-five insurers slightly below 50%.

The average reserve ratio is 10.9%. On the asset side, 80.4% of funds' portfolios are invested in sovereign and corporate bonds, 13.5% in stocks, and the rest in real estate, loans, and cash. The average asset return is 4.9% per year and the average contract return before fees is 4.0% per year, both in nominal terms.

Three factors can explain the wedge between the average asset return and the average contract return. First, as described in Section 1.1, the insurer can keep up to 15% of the asset return as profit, which represents about 75 bps on average. Second, part of the asset return has been retained to offset the dilution of reserves induced by positive net flows over the sample period. Given the average net flow rate of 2.4% per year and the average reserve ratio of 10.9%, insurers would have had to retain $0.024 \times 0.109 \simeq 25$ bps of asset returns per year to maintain the reserve ratio constant. Third, the average reserve ratio is actually about 3.5 percentage points higher at the end of the period than at the beginning (see Internet Appendix Figure B.2), which implies insurers have

9 See Internet Appendix C.1 for details about the data used in the paper and variable construction. The data are available through the Banque de France's open data room (https://www.banque-france.fr/en/statistics/access-granular-data/open-data-room).

retained over this 15-year period an additional $0.035/15 \simeq 25$ bps per year on average.

We complement the regulatory data with contract-level information from two sources. First, we retrieve information on fees from the data provider *Profideo*, which collects information on contract characteristics from contract prospectuses. The data consist in a snapshot of contracts with positive outstanding account value in 2017, even if the contract is closed to new subscriptions at that date. The fee structure is fixed at the subscription of the contract and written in the contract prospectus. Given that for every contract, some investors hold their contract for many years, it is sufficient to have a snapshot of outstanding contracts in 2017 to retrieve a complete picture of the fee structure of all contracts sold throughout the sample period 2000–2015. The data also include information on the time period during which contracts were open to new subscriptions. We keep contracts for which this period overlaps with the sample period. Fifty-seven percent of insurers, representing 68% of account value in the regulatory filings, can be matched with this data set. Panel B of Table 1 shows summary statistics on fees aggregated at the level of insurer-years in which the contract is open to new subscriptions, which is the level at which we run regressions using these data. Management fees are, on average, 70 bps of account value. Entry fees are, on average, 3.3%.[10]

Our second source of contract-level information is a survey (*Enquête Revalo*) conducted by the insurance supervisor every year from 2011 to 2015 among all the main insurers. The data cover 81% of aggregate account value in the regulatory filings. We retrieve information on net-of-fees contract returns, minimum guaranteed return, total account value, and number of investors, which allows us to calculate the average invested amount for every contract. Panel C of Table 1 presents summary statistics from this data set at the contract-survey year level. The average net-of-fees contract return is 2.7%.[11] The average (75th percentile) minimum guaranteed return is 35 bps (0), which is well below the average contract return of 2.7 percentage points over the same period. Thus, the minimum guaranteed rate is not binding for the vast majority of contracts: the net-of-fees contract return is strictly larger than the guaranteed return for 98% of contracts. This figure actually overstates the extent to which the minimum guaranteed return is binding, because the guaranteed return is before-fees. Assuming uniform management fees at the sample average of 70 bps, over 99% of contracts have a nonbinding minimum guaranteed return.

10 Recall that fees do not map one-to-one into insurer profit because insurers must return at least 90% of fees to investors (see Section 1.1).

11 It is lower than the average before-fees contract return in regulatory filings (4% in panel A) minus average management fees (0.7% in panel B) because the sample period is 2011–2015 for the survey data, whereas it is 2000–2015 for the regulatory filings, and contract returns are lower toward the end of the sample period (see Figure 1).

2. The Accounting of Intercohort Risk Sharing

In this section, we quantify intercohort risk sharing in euro contracts based on an accounting framework which formalizes the institutional framework presented in Section 1.1.

2.1 Accounting framework

$V_{j,t}$ denotes the total account value with insurer j at the end of year t after payment of the annual net-of-fees return $y_{j,t}$. It evolves according to

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

where $Flow_{j,t}$ is net flow to insurer j in year t.[12] The balance sheet of the fund at the end of year t is

$$A_{j,t}=V_{j,t}+R_{j,t}, \tag{2}$$

where $A_{j,t}$ is asset value and $R_{j,t}$ is reserves at the end of year t. Assets evolve according to

$$A_{j,t}=(1+x_{j,t})A_{j,t-1}+Flow_{j,t}-\Pi_{j,t}, \tag{3}$$

where $x_{j,t}$ is asset return and $\Pi_{j,t}$ is insurer j's profit in year t. Combining (1), (2), and (3), we obtain

$$x_{j,t}A_{j,t-1}=y_{j,t}V_{j,t-1}+\Pi_{j,t}+\Delta R_{j,t}, \tag{4}$$

where $\Delta R_{j,t}=R_{j,t}-R_{j,t-1}$. Equation (4) describes how asset income (on the left-hand side) is split into three parts: the amount credited to current investors (first term on the right-hand side), insurer profit (second term), and change in reserves (third term). Returns paid to investors can therefore be hedged against market risk if fluctuations in asset returns are absorbed by the insurer and/or by reserves. 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 reserves represents a payoff to past and future investors and is at the root of intercohort risk sharing.

2.2 Contract return smoothing

Figure 1 compares the time series of asset return $x_{j,t}$ and contract return $y_{j,t}$, averaged across insurers. The key pattern is that the contract return is an order of magnitude less volatile than the return on underlying assets. Thus, euro contracts provide investors with insurance against market risk.

Equation (4) shows two potential sources of market risk sharing exist. Static (cross-sectional) risk sharing between investors and the insurer arises if the

[12] To simplify the exposition, we write Equation (1) assuming that flows take place at the end of the year after payment of the annual return. In the empirical analysis, we assume that flows are uniformly spread throughout the year and therefore earn half the annual contract return.

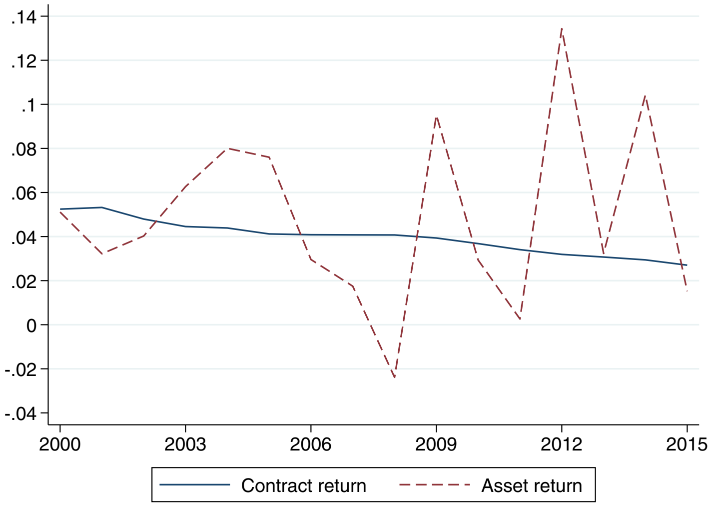

Figure 1
Asset return versus contract return
The figure graphs the value-weighted average contract return (solid blue), which is an order of magnitude less volatile than the value-weighted average asset return (dashed red).

difference between asset income and the amount credited to current investors is absorbed by the insurer. Intertemporal risk sharing between successive cohorts of investors arises if reserves absorb this difference.

To assess the contribution of reserves to the provision of insurance, Figure 2 compares two series. The solid blue line represents the difference between the amount credited to current investors and asset income: $y_{j,t}V_{j,t-1} - x_{j,t}A_{j,t-1}$. It represents the total transfer to current investors, that is, the transfer from the insurer plus the transfer from reserves. The dashed red line represents minus the change in reserves: $-\Delta R_{j,t}$. It represents the transfer from reserves. Both series are summed across insurers and normalized by aggregate account value. The figure shows how the two series track each other very closely; that is, variation in reserves absorbs almost all of the difference between asset return and contract return. Therefore, virtually all insurance against market risk is provided to investors through reserves.

2.3 Intercohort redistribution

Contract return smoothing using reserves implies that wealth is transferred across time, but it does not necessarily imply that wealth is transferred across investor cohorts because part of intertemporal transfers net out within investors' holding period. To illustrate this point with a stylized example, consider an investor holding a contract for 2 years during which the asset return and contract return are as follows:

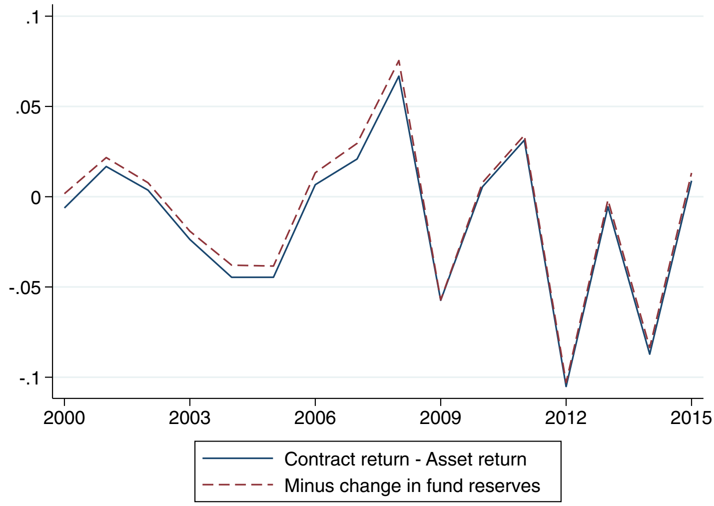

Figure 2
Reserves absorb asset return fluctuations
The figure visualizes the time-series variation in the difference between aggregate contract return and asset return normalized by account value $(y_t V_{t-1} - x_t A_{t-1})/V_{t-1}$ (solid blue) explains almost all of the time-series variation in aggregate transfer from reserves normalized by account value $-\Delta R_t / V_{t-1}$ (dashed red).

	Year 1	Year 2
Asset return	0	6
Contract return	4	4

Reserves absorb the difference between the asset return and contract return. In year 1, the investor receives a positive transfer from reserves equal to four. In year 2, the investor makes a transfer to reserves equal to two. Therefore, part of the year-on-year transfers net out over the investor's holding period. The net transfer to the investor is then $4-2=2$ over 2 years, or 1 per year. Our methodology to quantify intercohort redistribution follows the same logic as in this example, netting out transfers within investors' holding period in order to isolate the intercohort component.

To quantify intercohort redistribution induced by changes in reserves, we compare the actual contract return paid out to investors with the return they would obtain in a counterfactual with constant reserves, the same asset return, and the same insurer profit as in the data. Relative to the counterfactual, investors holding a contract with insurer j in year t receive a transfer from reserves equal to $-\Delta R_{j,t}$. Consider investor i holding a contract from beginning of year t_0 to end of year t_1, and $V_{i,j,\tau-1}$ denotes her account value at the beginning of year τ. She receives in year τ a transfer proportional to her weight in the insurer's total account value, equal to $\frac{V_{i,j,\tau-1}}{V_{j,\tau-1}}(-\Delta R_{j,\tau})$. Summing over her holding period as in the two-period example, we obtain investor i's holding period net transfer, which we apportion to each year in proportion to the beginning-of-year account value:[13]

$$NetTransfer_{i,j,t} = \frac{V_{i,j,t-1}}{\sum_{\tau=t_0}^{t_1} V_{i,j,\tau-1}} \sum_{\tau=t_0}^{t_1} \frac{V_{i,j,\tau-1}}{V_{j,\tau-1}} (-\Delta R_{j,\tau}). \tag{5}$$

The net transfer received by an investor depends on her holding period, that is, the year in which she starts investing and the year she redeems (and on the time profile of her investment within the holding period). Because investor cohorts are defined by the contract holding period of the cohort's investors, transfers across investors are transfers across investor cohorts. Therefore, investors from the same cohort are on the same side of redistribution. By contrast, investors from different cohorts may be on opposite sides of the redistribution.

Panel A of Table 2 shows the net transfer (5) received by an investor as a function of her holding period, for every possible holding period within the sample period. We calculate the net transfer for an investor who holds the value-weighted average contract and keeps a constant investment amount of 100 by withdrawing interests paid at the end of each year.[14] The numbers in the table represent the additional annual returns of the representative euro contract relative to a counterfactual with constant reserves. For instance, an investor buying a euro contract at the beginning of 2006 and redeeming it at the end of 2011 earned an additional 1.5 percentage points per year relative to a counterfactual with no smoothing, because insurers tapped reserves during the 2008 stock market crash and the 2011 sovereign debt crisis. Conversely, transfers turn negative for holding periods spanning the end of the sample period characterized by decreasing interest rates, because insurers hoarded the high bond returns as reserves during this period. Finally, shorter holding periods are associated with larger intercohort transfers in absolute value, because a smaller share of the intertemporal transfer takes place within the investor's holding period when the holding period is shorter. Hence, intercohort risk sharing plays out not only at long investment horizons but also at shorter ones.

2.4 Aggregate intercohort redistribution

The aggregate amount transferred across cohorts each year t is obtained by summing up the net transfers (5) across all investors in the set $\mathcal{I}^{j,t}$ of investors

[13] Transfers taking place in different years are not discounted differently because (85% of) asset returns are due to investors regardless of the level of reserves; that is, investors are entitled to the same share of asset returns whether assets are credited to the reserves or to their accounts. Therefore, only the total amount of reserve distribution matters, but not its timing within an investor's holding period.

[14] We estimate the amount of intercohort transfer on the sample of insurers for which we have data throughout 1999–2015. Doing so leads us to make two adjustments to the sample. First, when an insurer acquires another insurer, 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 scope throughout the sample period. Second, we drop a few insurers that enter or exit during the sample period or have missing data in some years. 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.

Table 2
Intercohort redistribution

A. Net transfer by investor cohort

entry year \ exit year	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015
2000	0.2	1.1	0.6	0	−0.7	−1.2	−0.9	−0.5	0.4	−0.2	−0.2	0.1	−0.7	−0.6	−1.1	−1
2001		2	0.8	−0.1	−0.9	−1.5	−1.1	−0.6	0.4	−0.3	−0.2	0.1	−0.7	−0.7	−1.2	−1.1
2002			−0.4	−1.2	−1.9	−2.4	−1.7	−1.1	0.2	−0.6	−0.4	0	−1	−0.9	−1.5	−1.3
2003				−2	−2.6	−3	−2	−1.2	0.2	−0.6	−0.4	0	−1.1	−1	−1.6	−1.4
2004					−3.2	−3.5	−2	−1	0.7	−0.4	−0.2	0.2	−1	−0.9	−1.5	−1.3
2005						−3.8	−1.4	−0.2	1.7	0.2	0.3	0.7	−0.7	−0.6	−1.4	−1.1
2006							0.9	1.6	3.5	1.2	1.1	1.5	−0.2	−0.2	−1.1	−0.8
2007								2.2	4.8	1.3	1.1	1.6	−0.4	−0.4	−1.4	−1.1
2008									7.5	0.9	0.8	1.4	−0.9	−0.8	−1.9	−1.5
2009										−5.7	−2.6	−0.6	−3	−2.5	−3.4	−2.7
2010											0.6	2	−2.1	−1.6	−3	−2.2
2011												3.4	−3.5	−2.4	−3.8	−2.8
2012													−10.4	−5.3	−6.3	−4.3
2013														−0.2	−4.2	−2.4
2014															−8.2	−3.4
2015																1.4

B. Aggregate intercohort redistribution

Intercohort transfer	
in % account value/year	1.4
in 2015 euros/year	17 billion
in % GDP	0.8

In panel A, *Net transfer* is defined in (5) for an investor buying a contract at the beginning of year t_0 (rows) and redeeming it at the end of year t_1 (columns). An investor buying a contract at the beginning of 2006 and redeeming it at the end of 2011 received an additional 1.5 percentage points per year relative to a counterfactual with constant reserves. In panel B, *Intercohort transfer* is defined in (6) and equal to the sum of lifetime net transfer across investors divided by total account value.

who hold a contract with insurer j in year t:

$$InterCohort\ Transfer_{j,t} = \sum_{i \in \mathcal{I}^{j,t}} |NetTransfer_{i,j,t}|. \tag{6}$$

Before calculating the aggregate intercohort transfer using (6), we show how the relation between variation in reserves and intercohort redistribution can already be quantitatively approximated using a back-of-the-envelope calculation. Suppose all investors have T-year holding periods and the annual transfer from reserves $-\Delta R_{j,t}$ is i.i.d. across time and normally distributed with zero mean. Then, the expected annualized net transfer over T years (i.e., expected $\left|\sum_{t=1}^{T} -\Delta R_{j,t}/T\right|$) is equal to $1/\sqrt{T}$ times the expected yearly transfer from reserves (i.e., expected $|-\Delta R_{j,t}|$). Intuitively, a longer holding period reduces intercohort transfers because a larger fraction of transfers from reserves net out over investors' holding period. The average outflow rate is 8.1% per year, which implies an average holding period of 12 years. The average absolute value of the yearly transfer from reserves is 3.7% of account value, implying an average intercohort transfer of the order of $3.7\%/\sqrt{12} \simeq 1.1\%$ of account value per year. Accounting for holding period heterogeneity across investors would lead to a larger intercohort transfer because of the convexity of $1/\sqrt{T}$.

To have an exact measure of aggregate intercohort transfer (6), we would need to observe the entire investment history of all investors, which is not possible, because the investment history of investors still holding a contract at the end of the sample period is not over. Two data limitations also exist. First, regulatory data start in 1999; therefore, we do not observe the entire investment history of investors who entered their contract before 1999. We can calculate the net transfer 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, which implies that we know the average holding period, but not its entire distribution. To calculate intercohort transfers, we assume the outflow rate is constant across cohorts at the insurer-year level and that investors only make one-off investments.[15] Under this assumption, we can reconstruct the investment history of all cohorts of investors and calculate the total intercohort transfer.

The value-weighted average amount of intercohort transfer is 1.4% of account value per year (panel B of Table 2). Evaluated at the 2015 level of aggregate account value of 1,200 billion euros, it amounts to an annual 17 billion euros that shift across investor cohorts on average, or 0.8% of GDP.[16]

3. Model

The large-scale intercohort redistribution we document in the previous section challenges the notion that intercohort risk sharing cannot be achieved in competitive markets. This notion follows from the assumption that investors are fully strategic and hence invest in contracts only when reserves are high and opt out when reserves are low. When the elasticity of the demand for contracts

15 Formally, $V_{j,t}(t_0)$ denotes the year t-total account value of contracts sold by insurer j in year t_0, we assume $V_{j,t}(t_0)=(1-\theta_{j,t})(1+y_{j,t})V_{j,t-1}(t_0)$ for all $t_0<t$, where the outflow rate $\theta_{j,t}$ is calculated to match observed outflows for insurer j in year t, that is, $\sum_{t_0<t}\theta_{j,t}(1+y_{j,t})V_{j,t-1}(t_0)=\mathit{Outflow}_{j,t}$; and account value of new contracts is calculated to match observed inflows to insurer j in year t, that is, $V_{j,t}(t)=\mathit{Inflow}_{j,t}$. See Internet Appendix C.2 for details.

16 The assumption of outflow rates independent of contract age is likely to underestimate the amount of intercohort transfer. Actual outflow rates are decreasing in contract age (FFSA-GEMA 2013), implying the true dispersion of holding periods is higher than the dispersion obtained under the assumption of the age-independent outflow rate. Because expected annualized life transfer is convex in the holding period, underestimating the dispersion of holding periods leads to underestimating intercohort transfer.

to reserves is infinitely large, intercohort risk sharing unravels (Allen and Gale 1997). In this section, we relax this assumption and instead assume that the elasticity of demand is finite. The model allows us to characterize how the equilibrium amount of intercohort risk sharing depends on the elasticity of demand, and to derive econometric specifications to estimate this elasticity.

3.1 Setup

The backbone of the model is the accounting framework presented in Section 2.1. Every period $t = 1, 2, \ldots, +\infty$, $J \geq 1$ long-lived intermediaries, indexed by $j = 1, \ldots, J$, sell one-period saving contracts. The contract offered by intermediary j in period t promises a return $y_{j,t}$ contingent on all information observable at the end of period t. At the beginning of period t, intermediary j has reserves $R_{j,t-1}$ and collects $V_{j,t-1}$ from investors. The intermediary has total assets $V_{j,t-1} + R_{j,t-1}$, which generate an exogenous return $x_{j,t}$ with $E_{t-1}[x_{j,t}] = r$, where E_{t-1} denotes expectation conditional on information at the beginning of period t. Asset risk may include a systematic component and an idiosyncratic component determined by the covariance structure of $x_t \equiv (x_{0,t}, \ldots, x_{J,t})$, where $j = 0$ defines investors' outside option described below.

As described in Section 1.1, the insurer profit is pinned down by the regulation ruling profit sharing between investors and the insurer. Regulation requires that the amount credited to investor accounts in any given year must be at least 85% of financial income in this year, where financial income is equal to asset return excluding unrealized capital gains. Empirically, unrealized capital gains are the main component of reserves and by far its main source of variation.[17] Accordingly, we write the regulatory constraint as

$$y_{j,t} V_{j,t-1} = (1-\phi)\big(x_{j,t} A_{j,t} - \Delta R_{j,t}\big),$$

where $\phi \in (0,1)$ equals 15% in the French regulatory framework. Using the budget constraint (4), this implies

$$\Pi_{j,t} = \frac{\phi}{1-\phi} y_{j,t} V_{j,t-1}. \tag{7}$$

The regulation imposes a cap on profits whereas for simplicity, we write the regulatory constraint with an equality. In Internet Appendix A.2, we derive a sufficient condition for the regulatory constraint to be binding. Intuitively, the constraint is binding if demand is sufficiently inelastic, in which case the equilibrium profit is large absent the regulatory constraint. We show this condition is empirically validated in Section 4.

---

17 See Internet Appendix B.2 for a detailed description of the regulatory framework and an empirical decomposition of reserves. The fact that the variation in reserves is mostly driven by unrealized capital gains can be visualized in Internet Appendix Figure B.2. In the insurer-year panel, the change in unrealized capital gains explains 98% of the variation in total reserves.

By constraining the intermediary's profit to be proportional to the contract return, regulation exogenously pins down the share of asset risk borne by the intermediary, and prevents the intermediary from using its equity to provide additional insurance to investors. Even if the regulatory constraint is written with "$\leq$" instead of "=", transferring wealth between the intermediary and investors across states of nature is still not feasible, because doing so involves lowering the intermediary's profit below the regulatory cap in some states (which is feasible) and increasing it above the cap in other states (which would violate the regulatory constraint). Therefore, insurance is provided to investors only through intercohort risk sharing. This property of the model is in line with the evidence in Section 2.2 that insurers' equity does not absorb losses on the asset portfolio, even during the 2008 Global Financial Crisis and the 2011–2012 European sovereign debt crisis. Since the model rules out that intermediaries absorb losses, we can assume for simplicity that intermediaries are risk neutral without creating the possibility of risk sharing between intermediaries and investors. Given the evidence in Appendix Figure B.3 that capital requirements never bind over the sample period, including during two crises, we also abstract from modeling solvency requirements. Intermediaries maximize expected profit discounted at the expected rate of asset return

$$E_0\left[\sum_{t=1}^{+\infty}\frac{\Pi_{j,t}}{(1+r)^t}\right]. \tag{8}$$

Intermediaries face the sequential budget constraint (4) for all $t \geq 1$. We normalize initial reserves $R_{j,0}$ to zero. To rule out Ponzi schemes, reserves must satisfy the transversality condition

$$\lim_{t\to+\infty}\frac{R_{j,t}}{(1+r)^t} \geq 0. \tag{9}$$

We model investor demand for contracts using a multinomial logit model. Every period a mass one of investors have one unit of wealth to invest. Each investor buys the contract that provides her with the highest expected utility. Investor i's expected utility from investing with intermediary j in period t is

$$\alpha\, E_{t-1}[u(y_{j,t})]+\xi_{j,t-1}+\psi_{i,j,t-1}. \tag{10}$$

The term $\alpha\, E_{t-1}[u(y_{j,t})]$ represents the expected indirect utility provided by contract return $y_{j,t}$, where $\alpha > 0$, $u' > 0$, $u'' < 0$, and without loss of generality we normalize $u'(r(1-\phi)) = 1$. $\xi_{j,t-1}$ is nonreturn preference for intermediary j in period t shared across all investors and $\psi_{i,j,t-1}$ is investor i's idiosyncratic preference. $\xi_{j,t-1}$ and $\psi_{i,j,t-1}$ are indexed by $t-1$ because they are realized at the end of period $t-1$. The vector of demand shocks $\xi_t \equiv (\xi_{1,t},\ldots,\xi_{J,t})$ follows a random walk that is uncorrelated with asset returns, $E_{t-1}[\xi_t \,|\, x_t] = \xi_{t-1}$, such that in equilibrium, the asset return will affect investor demand through its effect on the contract return and through this effect only. $\psi_{i,j,t-1}$ is distributed i.i.d.

extreme value across investors in each cohort, yielding the usual logit demand function (11).

α is the key parameter of the model. It parameterizes the elasticity of demand to the expected utility of contract return, capturing in reduced form several, nonmutually exclusive, mechanisms leading to imperfectly elastic demand, such as nonreturn product differentiation, switching costs, and information frictions (Hortaçsu and Syverson 2004). Investors might be able to anticipate future contract returns yet do not necessarily buy the contract with the highest expected utility of return, because they trade off returns against other contract attributes or because portfolio rebalancing is costly. Alternatively, information frictions might prevent investors from figuring whether certain contracts have higher expected returns than others, and which ones. We provide evidence for and against these mechanisms in Section 5.

α should be interpreted as the elasticity of demand to the expected utility of contract return *conditional on reserves*. Indeed, in the model, variation in reserves is the only source of contract return predictability. In practice, other sources of return predictability might exist, such as heterogeneity in asset risk and management skills across intermediaries, which is absent from the model. Demand may also react to changes in expected returns on outside investment opportunities. The elasticity of demand to these other sources of return predictability may differ from α, for example, if these other factors are more salient or easier to apprehend than reserves. We provide evidence supporting this interpretation in Section 4.3.

Investors have access to an outside investment opportunity indexed by $j=0$, which yields expected utility given by (10) with $\xi_{0,t-1}$ normalized to zero, $\psi_{i,0,t-1}$ distributed i.i.d. extreme value, and $y_{0,t}=(1-\phi_0)x_{0,t}$. The parameter $\phi_0>0$ captures the fees and other costs of investing in the outside investment opportunity. The outside investment opportunity can be thought of as other liquid saving instruments, for example, mutual funds or direct investment in financial markets. We assume the cost of investing through the outside option is the same as the cost of investing through intermediaries, that is, $\phi_0=\phi$.[18]

Contract return in period t is contingent on all observable information at the end of period t, which includes the history of asset returns and demand shocks. Thus, $y_{j,t}$ is a function of (x^t,ξ^t), where the t exponent denotes history up to end-of-period t. The demand for intermediary j's contract in period t is

$$V_{j,t-1}=\frac{\exp\{\alpha E_{t-1}[u(y_{j,t})]+\xi_{j,t-1}\}}{\sum_{k=0}^{J}\exp\{\alpha E_{t-1}[u(y_{k,t})]+\xi_{k,t-1}\}}. \tag{11}$$

The problem of an intermediary is to maximize profit (8) by choosing a contract return policy subject to the budget constraint (4), the profit

[18] In France, as in several other European countries, life insurers sell mutual funds through unit-linked contracts that are subject to the same fee structure and tax treatment as euro contracts. In such cases, $\phi_0=\phi$ by design.

function (7), the transversality condition (9), and the demand function (11). Each intermediary takes other intermediaries' contract return policies as given. An equilibrium is defined as a fixed point of this problem.[19]

Finding a general analytical solution to this problem is difficult. To simplify the problem and obtain an explicit solution, we solve the model using a first-order approximation.[20] We assume shocks have bounded support, that is, there exists $\sigma > 0$ such that deviations of $x_{j,t}$ and $\xi_{j,t}$ from their $t-1$-conditional expectations lie in $[-\sigma,\sigma]$ for all j and t, and that, for some period T, these deviations are zero for $t > T$. The value of T can be any positive integer, so that our analysis covers any finite number of shocks, however large. We calculate an explicit solution that is valid as long as σ is small, that is, fluctuations in asset return and demand are not too large.

3.2 Equilibrium

Our first result shows how the equilibrium contract return depends on the history of past and current asset returns.

Proposition 1. Contract return of intermediary j in period t is

$$y_{j,t} \simeq (1-\phi)\left[r+\sum_{s=1}^{t}\beta_{j,t}(s)(x_{j,s}-r)\right]+f_{j,t}(\overline{x}^t-r,\xi^t), \tag{12}$$

where

$$\beta_{j,t}(s) = \frac{\rho_j}{\alpha+\frac{1+r}{r}\rho_j} \qquad \text{for } s<t, \tag{13}$$

$$\beta_{j,t}(t) = \frac{\alpha+\rho_j}{\alpha+\frac{1+r}{r}\rho_j}. \tag{14}$$

$\rho_j > 0$ is an intermediary-specific constant that is independent of α and proportional to investors' coefficient of relative-risk aversion. $f_{j,t}(.)$ is a function of the history of weighted-average asset return shocks $\overline{x}^t - r$ and the history of demand shocks ξ^t. Closed-form expressions for these variables are in Internet Appendix A.1.

Equation (12) shows that the contract return is equal to the expected asset return, r, plus a function of the history of shocks to the intermediary's asset

---

[19] Because we do not clear the capital market, our model is in partial equilibrium. The model is equivalent to a general equilibrium model with constant returns to capital as in Allen and Gale (1997) and Ball and Mankiw (2007). Suppose each intermediary j can lend capital to competitive firms using a linear production function $Y_{j,t}=(1+x_{j,t})K_{j,t-1}$, where $Y_{j,t}$ is output and $K_{j,t-1}$ is capital. In such an economy, an increase in reserves leads to an increase in the aggregate capital stock. An alternative interpretation of our model is that of a small open economy, in which case an increase in reserves leads to a capital account deficit.

[20] The advantage of using a first-order approximation is that it eliminates any possible interaction between the shocks occurring in different periods. Ball and Mankiw (2007) use a similar method to solve the complete-market equilibrium in which investors are allowed to trade with future investor cohorts.

return, $x_{j,s}-r$ for $1\leq s\leq t$, minus the compensation of the intermediary which consists in a fraction ϕ of those asset returns, and a function of the history of average asset returns $\overline{x}^t$ and demand shocks ξ^t. The key coefficients in (12) are the $\beta_{j,t}(s)$, which pin down the extent of risk sharing across investor cohorts. $\beta_{j,t}(s)$ measures the sensitivity of period-t contract return, $y_{j,t}$, to period-s asset return, $x_{j,s}$. When $\beta_{j,t}(s)>0$, the period-t investor cohort bears some of period-s asset risk.

The contract return policy (12) implies period-s asset risk is shared between the current (period-s) cohort (because $\beta_{j,s}(s)>0$) and all future cohorts (because $\beta_{j,t}(s)>0$ for $t>s$). In turn, Equations (13) and (14) show the extent of intercohort risk sharing depends on investor risk aversion through the parameter ρ_j, and on the elasticity of demand to expected utility of contract return, α. Intuitively, there is more intercohort risk sharing when investors are more risk averse: when ρ_j is higher, $\beta_{j,s}(s)$ is lower and $\beta_{j,t}(s)$ is higher for $t>s$. In words, shocks to current asset returns are better shared across investor cohorts when investors are more risk averse. More importantly, there is more intercohort risk sharing when investor demand is less elastic. When demand is inelastic ($\alpha\simeq 0$), asset risk is perfectly shared between the current and future cohorts: $\beta_{j,s}(s)=\beta_{j,t}(s)$ for all $t>s$. When demand is elastic ($\alpha>0$), more asset risk is borne by the contemporaneous cohort: $\beta_{j,s}(s)>\beta_{j,t}(s)$ for $t>s$. In this case, the asset risk is imperfectly shared across investor cohorts.

The intuition is that when demand is elastic, future investor cohorts behave opportunistically by investing more (less) when reserves are higher (lower). For instance, when the asset return is high, the intermediary would like to share in gains with future investor cohorts by hoarding part of the return as reserves. When demand is elastic, however, future investors flow in, diluting reserves and undoing the sharing of gains. Conversely, when the asset return is low, the intermediary would like to share in losses with future investor cohorts by tapping reserves and replenishing them in future periods. In this case, future investors flow out, preventing the intermediary from replenishing reserves and undoing the sharing of losses. In the limit when demand is perfectly elastic ($\alpha\simeq\infty$), intercohort risk sharing unravels completely: $\beta_{j,s}(s)=1$ and $\beta_{j,t}(s)=0$ for $t>s$.

Our next result shows that the contract return depends on past asset returns, $x_{j,s}$ for $s<t$, *only through* the impact of past asset returns on the reserve ratio. R_{j,t^-} denotes the amount of reserves at the end of period t just before distribution to investors. This amount is equal to beginning-of-period reserves plus asset income:

$$R_{j,t^-}=R_{j,t-1}+x_{j,t}(V_{j,t-1}+R_{j,t-1}). \tag{15}$$

The next proposition states that the reserve ratio to total account value, $\mathcal{R}_{j,t^-}\equiv R_{j,t^-}/V_{j,t-1}$, is a sufficient statistics for how the history of past asset returns determines the current contract return.

Proposition 2. Contract return of intermediary j in period t is

$$y_{j,t} \simeq (1-\phi)r + \frac{1-\phi}{1+r}\,\frac{\alpha}{\alpha+\frac{1+r}{r}\rho_j}(x_{j,t}-r) + \frac{(1-\phi)r}{1+r}(\mathcal{R}_{j,t^-}-r) + \mu_j(\overline{x}_t - r) + \nu_j\,\Delta\xi_{j,t}, \tag{16}$$

where $\rho_j > 0$ is a constant independent of α, $\mu_j < 0$ goes to zero when α goes to zero or infinity, $\overline{x}_t$ is a weighted average of $x_{k,t}$ over $k = 1,\ldots,J$, $\nu_j > 0$ goes to zero when α goes to infinity, and $\Delta\xi_{j,t}$ is a demand shock. Closed-form expressions for these variables are in Internet Appendix A.1.

Proposition 2 shows how the share of asset risk borne by current investors depends on the elasticity of demand. When demand is inelastic ($\alpha \simeq 0$), the coefficient in front of $x_{j,t}$ in (16) is equal to zero. The contract return then does not depend on the current asset return beyond its effect on the reserve ratio; that is, asset risk is perfectly shared across investor cohorts. When demand is elastic ($\alpha > 0$), the coefficient in front of $x_{j,t}$ is strictly positive. The intermediary then shifts more asset risk to the current cohort. In this case, the contract return depends on the current asset return above and beyond its effect on the end-of-period reserve ratio; that is, asset risk is imperfectly shared across investor cohorts.

An implication of Proposition 2 is that the reserve ratio $\mathcal{R}_{j,t^-}$ is a sufficient statistic for the history of shocks. All that matters for setting the contract return is the intermediary's current reserve ratio, not the path leading to that ratio. Indeed, in (16), the contract return does not depend on past shocks beyond their effect on the reserve ratio. The sensitivity of the contract return to the reserve ratio results from the following tradeoff faced by the intermediary. On the one hand, paying out a larger fraction of reserves to current investors leads to higher demand and thus higher profit in the current period. On the other hand, tapping reserves today implies paying lower returns to future investors, lowering future demand and hence future profit. The optimal choice is to pay out in the current period a fraction of reserves equal to the weight of current investors in intertemporal profit, equal to $\frac{r}{1+r}$, a fraction $1-\phi$ of which accrues to investors.

Proposition 2 also implies an intermediary's contract return depends negatively on other intermediaries' asset returns, because $\mu_j < 0$. Intuitively, when other intermediaries have high asset returns they increase contract returns both in the current period and in future periods, which reduces intermediary j's future demand, but not its current-period demand, which is realized before asset returns. Intermediary j's optimal response is then to increase future contract returns by lowering the current contract return so as to avoid losing too large future market shares.[21] This effect vanishes when demand is perfectly inelastic ($\alpha \simeq 0$), because intermediary j then has no incentive to react; and it vanishes when demand is infinitely elastic ($\alpha \simeq \infty$), because other intermediaries then do not change the future contract return in response to asset return shocks. Finally, the contract return depends on the demand shock $\Delta\xi_{j,t}$ realized at the end of the period. Intuitively, the intermediary has incentives to lean against a negative shock to future demand, by lowering the contract return in the current period and increasing reserves to promise higher returns in the future.

3.3 Contracts complete markets

Proposition 2 implies that contracts complete markets in the sense that the contract return cannot be replicated using existing market instruments. More precisely, the risk profile of the contract can be replicated but only at a higher price than with other market instruments. A first, mundane reason the risk profile of the contract may not be replicated is because the contract return (16) includes an error term that depends on the realized demand shock, which may not be tradable. Let us exclude this reason by focusing on the case without demand shocks, that is, $\Delta\xi_{j,t} = 0$. In this case, the following proposition shows that the risk profile of contract j can be replicated by positions in (i) the assets held by insurer j, generating the return $x_{j,t}$; (ii) a weighted portfolio of assets held by all insurers, generating the return $\overline{x}_t$; and (iii) the risk-free asset, generating the return $r_f < r$. This portfolio replicates the risk profile of the contract in the sense that its return is equal to the contract return up to a constant. This constant is, however, nonzero.

Proposition 3. In the absence of demand shocks, the contract return is equal to the return on a portfolio containing the assets held by insurers and the risk-free asset plus a constant equal to

$$\left[1-(1-\phi)\frac{\alpha+\rho_j}{\alpha+\frac{1+r}{r}\rho_j}-\mu_j\right](r-r_f)+(1-\phi)r\mathcal{R}_{j,t-1}-\phi r, \tag{17}$$

where $\mathcal{R}_{j,t-1} \equiv R_{j,t-1}/V_{j,t-1}$ is the beginning-of-period reserve ratio.

The first term in (17) is positive and proportional to the risk premium on the assets held by insurers, $r - r_f > 0$. It reflects the fact that the contract earns the risk premium without bearing all the associated risk (and almost none of it when $\alpha \simeq 0$), because some of this risk is diversified across investor cohorts. Therefore, the portfolio with the same risk exposure as the contract has a lower risk exposure than the insurer's asset portfolio and thus earns a lower risk premium than the contract. Correspondingly, the term in large brackets which

[21] This best-response reflects the strategic complementarity property of logit demand, that is, the property whereby contract return best-response functions are increasing in other intermediaries' contract returns.

multiplies the risk premium in (17) is positive and equal to the difference between the risk exposure of the insurer's portfolio and that of the contract. The second term is proportional to the reserve ratio and arises from the predictable distribution of reserves in the contract return. The third term is negative and equal to the fees.

When reserves are equal to their unconditional mean (normalized to zero) and fees are not too high, Proposition 3 implies that the contract strictly dominates the portfolio replicating the risk profile of the contract. The contract reaches a point beyond the efficient frontier based on market instruments, because the contract shares asset risk with future cohorts of investors who do not yet participate in the market. By contrast, intercohort risk sharing cannot be achieved using market instruments. In this sense, the contract completes financial markets.

Proposition 3 could imply that there exists an arbitrage opportunity consisting in buying the contract and shorting the replicating portfolio. We analyze this possibility in more depth in Section 5.4 and show that, in practice, arbitrage is made unprofitable by the nondeductibility of interest expenses on levered financial investments by households.

3.4 Empirical implications

We now derive two relations that can be estimated in the data to back out the elasticity of demand, which is the key determinant of equilibrium risk sharing.

The first relation is the contract return policy. The coefficients ρ_j, μ_j, and ν_j from Proposition 2 are intermediary specific because the optimal contract return depends on the elasticity of demand, itself a function of the intermediary's market share due to logit demand. A closer inspection of these coefficients (reported in Internet Appendix A.1) reveals they only depend on market shares up to second-order terms. When market shares are not too large, equilibrium contract returns can be approximated as follows:

Relation 1 (contract return policy). For small market shares, the period-t contract return of intermediary j is

$$y_{j,t} \simeq cste_t + \frac{1-\phi}{1+r}\frac{\alpha}{\alpha+\frac{1+r}{r}\rho}x_{j,t} + \frac{(1-\phi)r}{1+r}\mathcal{R}_{j,t^-} + \varepsilon_{j,t}, \qquad (18)$$

where $cste_t$ is a period-specific constant, ρ is proportional to the coefficient of relative-risk aversion, and $Cov\big((x_{j,t},\mathcal{R}_{j,t^-}),\varepsilon_{j,t}\big)\simeq 0$.

Relation 1 implies the coefficients in the equilibrium contract return policy (18) can be estimated by running a linear regression with time fixed effects in a panel of intermediaries. Relation 1 also implies the model can be easily rejected, because it predicts the coefficient in front of the reserve ratio should be commensurate with the expected asset return. The coefficient in front of the current asset return is informative about the elasticity of demand: it varies monotonically from zero to one as α varies from zero to infinity.

The second relation is that between flows and reserves:

Relation 2 (flow-reserves relation). Net flows to intermediary j in period t are given by

$$\log(V_{j,t-1}) \simeq cste_j + cste_{t-1} + \alpha(1-\phi)r\mathcal{R}_{j,t-1} + \xi_{j,t-1}, \tag{19}$$

where $cste_j$ and $cste_{t-1}$ are intermediary-specific and period-specific constants, respectively, $\mathcal{R}_{j,t-1}$ is the beginning-of-period reserve ratio, and $Cov(\mathcal{R}_{j,t-1},\xi_{j,t-1})<0$. Moreover, lagged asset return $x_{j,t-1}$ is a valid instrument for $\mathcal{R}_{j,t-1}$.

We know from Proposition 2 that the contract return paid at the end of period t depends on the end-of-period reserve ratio, itself determined by the beginning-of-period reserve ratio. Correspondingly, Relation 2 states that investor demand depends on the reserve ratio at the beginning of the period. The sensitivity of investor demand to the reserve ratio is equal to the product of the sensitivity of investor demand to expected contract return, α, and the sensitivity of expected contract return to the beginning-of-period reserve ratio, $(1-\phi)r$.

The coefficient $\alpha(1-\phi)r$ in the flow-reserves relation can be estimated by running a linear regression with time and intermediary fixed effects. The ordinary least squares (OLS) estimate is unbiased if intermediary-specific demand shocks are zero, or if they are observable to the econometrician and can be controlled for. In the presence of unobservable demand shocks, however, the error term is negatively correlated with reserves. Intuitively, when the intermediary anticipates a negative demand shock, it optimally increases reserves to increase future contract returns and lean against the demand shock, generating a spurious negative correlation between reserves and demand. This correlation creates a downward bias in the OLS estimate. The bias can be corrected by instrumenting reserves using lagged asset returns. Indeed, lagged asset returns affect reserves because a fraction of asset returns are hoarded as reserves (relevance condition), but they are not directly correlated with demand shocks beyond their effect on reserves (exclusion restriction).

The model considers one-period contracts such that the total account value is equal to inflows in the current period. In practice, although euro contracts are demandable thus conceptually similar to one-period contracts, they are automatically renewed until investors redeem shares. Given the evidence in Section 4 that investments in euro contracts are sticky, with an average holding period of 12 years, an alternative specification of the model could allow investors to hold contracts for several periods. We propose an extension of our model in Internet Appendix D where we assume that inflows come from new investors choosing between the euro contracts or the outside investment, and outflows come from existing investors choosing between staying with their contract or leaving for the outside investment. We show that in that extension, the equilibrium contract return can be written as a function of the current asset return and the reserve ratio, as in (19), in which the coefficient in front of the current asset return goes to zero when the elasticities of both inflows and outflows to the expected utility of contract return go to zero, and the coefficient in front of the reserve ratio is the same as in (19).

4. Demand (In)elasticity to Reserves

The key insight from the model is that the equilibrium level of intercohort risk sharing depends on the elasticity of demand to expected returns conditional on reserves, that is, on the value of α. Specifically, a demand that is inelastic to reserves ($\alpha \simeq 0$) allows for perfect intercohort risk sharing, whereas a demand that is perfectly elastic to reserves ($\alpha \simeq \infty$) unravels intercohort risk sharing. The model shows that α can be identified from two moments given by Relations 1 and 2. Guided by the results in Section 3.4, we estimate each of these moments in turn, by running panel regressions.

4.1 Relation 1: Contract return policy

The first implication of the model is that after controlling for the current reserve ratio, equilibrium contract returns depend positively on current asset returns if $\alpha > 0$, and do not depend on current asset returns if $\alpha \simeq 0$. We estimate the contract return policy given by Relation 1 by running a panel regression with year fixed effects. According to our model, insurer fixed effects are not necessary, because the model assumes no heterogeneity in expected asset return across insurers. Yet, our preferred specification does include insurer fixed effects to account for such heterogeneity in the data.[22] We estimate weighted regressions using the insurer share of account value in aggregate account value as weights.[23] We calculate standard errors two-way clustered by insurer and by year.

Table 3 reports the results. In line with the model, the coefficient for the reserve ratio is positive and statistically significant at the 1% level, in both specifications. In our preferred specification with insurer fixed effects (Column 2), the point estimate implies a one-percentage-point increase in the reserve ratio is associated with a 3.5-basis-point increase in the annual contract return. That is, out of each additional euro of reserves, 3.5 cents are credited to investor accounts per year. The model predicts a regression coefficient equal

[22] In the model, including insurer fixed effects does not lead to a misspecified regression, because it only adds regressors that are uncorrelated with the dependent variable and with the other explanatory variables. In the data, insurer fixed effects in contract return regressions are always jointly significant at statistical levels below 1%.

[23] We obtain similar results when we estimate nonweighted regressions (untabulated).

Table 3
Contract return

	Contract return ($y_{j,t}$)					
	(1)	(2)	(3)	(4)	(5)	(6)
Reserve ratio ($R_{j,t-}$)	.026***	.035***	.039***	.03***	.032***	.034***
	(.0078)	(.0082)	(.01)	(.0095)	(.0068)	(.0079)
Asset return ($x_{j,t}$)	−.017	−.018**	−.019**	−.015*	−.016**	−.017*
	(.011)	(.0079)	(.0085)	(.0075)	(.0066)	(.0097)
Reserve ratio × High insurer capital			−.0045			
			(.008)			
Reserve ratio × Large insurer				.011		
				(.0075)		
Reserve ratio × Low reserves					.017	
					(.017)	
Reserve ratio × Crisis						.015***
						(.0047)
Asset return × High insurer capital			.0025			
			(.0068)			
Asset return × Large insurer				−.011		
				(.0076)		
Asset return × Low reserves					−.011	
					(.0098)	
Asset return × Crisis						−.011
						(.013)
High insurer capital			−.00049			
			(.0016)			
Large insurer				−.0025		
				(.0016)		
Low reserves					−.0019	
					(.0019)	
Year FE	✓	✓	✓	✓	✓	✓
Insurer FE		✓	✓	✓	✓	✓
R^2	.69	.81	.81	.81	.81	.81
Observations	978	978	978	978	978	978

Panel regressions at the insurer-year level for 76 insurers over 2000–2015. *Contract return* is the annual before-fees contract return paid at the end of year t. *Reserve ratio* is total reserves at the end of year t just before annual distribution normalized by total account value. *Asset return* is asset return in year t. *High insurer capital* is dummy equal to one if the insurer's capital ratio is above the sample median. *Large insurer* is a dummy equal to one if the insurer is in the top quartile of insurer size as measured by total account value. *Low reserves* is a dummy equal to one if the insurer has a reserve ratio in the bottom quartile of the sample distribution. *Crisis* is a dummy equal to one in crisis years 2008 (global financial crisis) and 2011–2012 (European sovereign debt crisis). All regressions are weighted by the insurer share in aggregate account value in the current year. Standard errors two-way clustered by insurer and year are reported in parentheses. $^*p < .1$; $^{**}p < .05$; $^{***}p < .01$.

to $(1-\phi)r/(1+r)$. Thus, the estimate of 0.035 and $\phi=0.15$ imply $r=4.3\%$, consistent with the average asset return observed during the period.[24]

The key result is that the coefficient for the asset return is a precisely estimated zero. The coefficient is not statistically different from zero with a 95% confidence interval between negative 0.04 and positive 0.01 when insurer fixed effects are not included (Column 1). In other words, the contract return does not depend on the contemporaneous asset return beyond its effect on the

[24] The sample average asset return is 4.9% (Table 1), perhaps because realized asset returns have been above expected returns during the sample period. As discussed in Section 1.2, the reserve ratio rose by 25 bps per year, while positive net flows should have diluted reserves at a rate of 25 bps per year. Therefore, insurers have retained approximately 50 bps of realized asset returns in reserves, consistent with an ex ante expected return of approximately 4.4%.

reserve ratio, which is consistent with $\alpha \simeq 0$. The coefficient is also a precisely estimated zero when insurer fixed effects are included (Column 2). Note that although insurer fixed effects reduce the standard deviation of the point estimate and make it statistically significant, the economic magnitude is essentially zero: the coefficient is negative 0.018, which is an order of magnitude smaller than the value of the coefficient implied by the model when demand is infinitely elastic, equal to $\frac{1-\phi}{1+r} \simeq 0.81$ using $\phi = 0.15$ and $r = 4.3\%$. Note also that the negative point estimate does not imply that the pass-through of asset return to contract return is negative. Indeed, the asset return enters positively into the reserve ratio (see Equation (15)), so the pass-through of asset return to contract return is given by the sum of the coefficient for the reserve ratio and the coefficient for the asset return. This pass-through is positive (equal to 0.017 with p-value of 0.15). The negative coefficient for asset return means that the contract return in year t is less sensitive to the asset return in year t than to asset return in years $s < t$. This can be explained by two institutional factors. First, when investors withdraw money, say in July of year t, they are usually credited a pro rata contract return over January to June based on the contract return in year $t-1$. In this case, the contract return in year t does not depend on the asset return in year t. Second, insurers sometimes offer new clients a guaranteed return in the first year of the contract for marketing purposes. In this case, the first-year contract return does not depend on the contemporaneous asset return. These two factors imply that contract returns associated with inflows and outflows do not depend on the current year asset return, which weakens the relation between the contract return and the contemporaneous asset return.

In conclusion, the empirical contract return policy exhibits sizeable intercohort risk sharing, consistent with $\alpha \simeq 0$.

In columns 3 to 6 of Table 3, we analyze whether the contract return policy varies across insurers and according to market conditions. We regress the contract return on the reserve ratio and the contemporaneous asset return, which we both also interact with several insurer characteristics. In column 3, we interact with a dummy equal to one if the insurer's capital ratio is above the sample median. In column 4, we interact with a dummy equal to one if the insurer is in the top quartile of insurer size as measured by total account value. We find that neither the capital ratio nor the insurer size is associated with significantly different contract return policies. In column 5, we interact with a dummy equal to one if the insurer has a reserve ratio in the bottom quartile of the sample distribution, and in column 6, we interact with a dummy equal to one in crisis years 2008 (Global Financial Crisis) and 2011–2012 (European sovereign debt crisis). We find that the sensitivity of contract return to reserves is higher in crisis years, that is, insurers with low reserves cut contract return relatively more during crises than outside crises. However, the extent of return smoothing, captured by the coefficient for the current asset return, does not increase in crisis time. The pass-through from asset return to contract return remains precisely estimated and very close to zero in all cases.

4.2 Are reserves really pooled across investor cohorts?

Intercohort risk sharing arises to the extent that reserves are pooled across investor cohorts. As discussed in Section 1.1, in principle, insurers could undo reserves pooling by closing existing contracts to new subscriptions when reserves are high, and creating a new vintage of contracts to price in the high level of reserves for new investors. Pricing of reserves could be done by creating new contracts with (a) higher entry fees, (b) higher management fees, (c) lower before-fees contracts return, or any combination of (a), (b), and (c), when reserves are higher.

We use two different sources of contract-level information to test whether insurers follow any of the (a), (b), or (c) strategies. First, we use data on fees to test for (a) and (b). The data provide a snapshot of contracts with positive outstanding account value in 2017 (even if the contract is no longer commercialized in 2017). The fee structure is fixed at the subscription of the contract and written in the contract prospectus. Because, for a given contract, some investors hold their contract for many years, it is sufficient to have a snapshot of outstanding contracts in 2017 to retrieve the fee structure of all contracts sold throughout the sample period 2000–2015. The data also report the time period during which each contract was open to new subscriptions. For each insurer j and each year t over 2000–2015, we calculate the average entry fee and average management fee across all contracts offered by insurer j and open to new subscriptions in year t. We regress the average fee on the insurer's beginning-of-year reserve ratio. If insurers price reserves into fees, the coefficient for the reserve ratio would be positive. Results in Table 4 show that insurers neither follow strategy (a) by adjusting entry fees (Column 1) to the level of reserves nor strategy (b) by adjusting management fees (Column 2) to the level of reserves.

Second, we use data on net-of-(management-)fees returns at the contract level to test whether insurers do a combination of (b) and (c). The data are from a survey conducted by the insurance supervisor since 2011. Each survey is a snapshot of contracts with positive outstanding contract value (even if the contract is no longer commercialized in the survey year) with information on the net-of-fees contract return. The data report the first year in which the contract was commercialized. For each contract c of vintage s, we retrieve the insurer's reserve ratio at the beginning of year s from the regulatory filings. We obtain a panel at the contract (c) $\times$ vintage year (s) $\times$ return year (t) level, where vintage years run throughout the sample period 2000–2015 and return years are from 2011 to 2015. We regress the net-of-fees return (of contract c in year t) on the reserve ratio in the contracts' vintage year (at beginning of year s) with insurer and vintage year fixed effects.[25] If insurers followed strategies

[25] Because we stack the five snapshots of return data, we interact the fixed effects with return-year dummies.

Table 4
Fees

	Entry fee (1)	Management fee (2)	Net-of-fee contract return (3)
Lagged reserves	−.016	.000054	−.005
	(.011)	(.0011)	(.0056)
Year FE	✓	✓	✓
Insurer FE	✓	✓	✓
R^2	.92	.95	.72
Observations	578	578	13,659

Columns 1 and 2 present panel regressions at the insurer-year level for 48 insurers over 2000–2015. The dependent variable in column 1 is *Entry fee* constructed as the average entry fee (frond-end load) of contracts sold by the insurer j in year t. The dependent variable in column 2 is *Management fee* constructed as the average management fee of contracts sold by insurer j in year t. The independent variable in columns 1 and 2 is *Lagged reserves* constructed as insurer j's reserves at beginning-of-year t normalized by total account value. The regressions in columns 1 and 2 include insurer and year fixed effects and are weighted by the insurer share in aggregate account value in the current year. Column 3 presents a panel regression at the contract-vintage year-return year level for about 2,700 outstanding contracts per year from 56 insurers over 2011–2015. The dependent variable in column 3 is contract return in year t of contract c of vintage year s offered by insurer j. The independent variable in column 3 is *Lagged reserves* constructed as insurer j's reserves at beginning-of-year s normalized by total account value. The regression in column 3 includes insurer-return year and vintage year-return year fixed effects and are weighted by the contract share in aggregate account value in the current return year. Standard errors two-way clustered by insurer and year (return year for column 3) are reported in parentheses. *$p<.1$; **$p<.05$; ***$p<.01$.

(b) or (c) of pricing reserves by adjusting future net-of-fees contract returns, the coefficient for the reserve ratio in the contract's vintage year would be negative. Column 3 of Table 4 shows this is not the case. Insurers do not discriminate across investor cohorts based on the level of reserves when investors enter into the contract.

In conclusion, reserves are indeed pooled across investor cohorts.

4.3 Relation 2: Flow-reserves relation

The second implication of the model is that investor flows depend positively on the reserve ratio if $\alpha > 0$, whereas flows are insensitive to reserves if $\alpha = 0$. We estimate the flow-reserves relations by running panel regressions with insurer and year fixed effects. We use several specifications, reported in panel A of Table 5. In column 1, the dependent variable is the net flow in rate (inflow minus outflow as a fraction of total account value). Columns 2 to 4 decompose the three components of net flows: (plus) inflows, that is, premiums, which come either from investors already holding a contract and adding money to their account or from new investors; (minus) redemptions, which are voluntary outflows; and (minus) payments at contract termination, which are involuntary outflows (because of investor death).

The dependent variable in the flow-reserves relation given by Equation (19) is the log level of the invested amount, rather than the investment flow as in column 1. The reason being that investments are assumed to be one-period in the model, so the stock and flow of investment are identical. In Appendix D, we work out an extension of the model in which contracts are held for several periods, leading to two flow-reserves relations, one for inflows and one for

Table 5
Investor flows

	A. OLS regressions				
	Net flows (1)	Inflows (2)	Redemptions (3)	Termination (4)	log(Inflows) (5)
Lagged reserves	.035	.031	−.013	.012	−1.2
	(.039)	(.038)	(.019)	(.0098)	(1.1)
Year FE	✓	✓	✓	✓	✓
Insurer FE	✓	✓	✓	✓	✓
R^2	.63	.73	.74	.79	.91
Observations	978	978	978	978	978

	B. IV regressions				
	Net flows (1)	Inflows (2)	Redemptions (3)	Termination (4)	Log(Inflows) (5)
Lagged reserves	.12	.069	−.041	−.008	.31
	(.086)	(.074)	(.037)	(.013)	(1)
Year FE	✓	✓	✓	✓	✓
Insurer FE	✓	✓	✓	✓	✓
R^2	.65	.76	.77	.8	.92
Observations	910	910	910	910	910

Panel regressions at the insurer-year level for 76 insurers over 2000–2015. In column 1, the dependent variable is net flow (total premiums minus voluntary redemptions minus involuntary redemptions at contract termination, i.e. the investor's death) normalized by total account value. In columns 2 to 4, the dependent variable is each of these three components of the net flow rate. In column 5, the dependent variable is the log inflow amount. *Lagged reserves* is the beginning-of-year level of reserves normalized by total account value. Panel A shows OLS regressions. Panel B shows IV regressions in which the insurer's beginning-of-year reserve ratio is instrumented using the insurer's asset return in the previous year. IV regressions include as control variables the shares of the insurer's asset portfolio in five broad asset classes at the beginning of the previous year. Each insurer's first year of observation is dropped from the IV regressions because these regressions use lagged variables. All regressions are weighted by the insurer share in aggregate account value in the current year. Standard errors two-way clustered by insurer and year are reported in parentheses. $*p < .1$; $**p < .05$; $***p < .01$.

outflows. The inflow-reserves relation has the log inflow amount on the left-hand side and is governed by a parameter controlling the elasticity of inflows. In accordance with this model extension, the specification in column 5 uses the log inflow amount as the dependent variable. The outflow-reserves relation has the outflow rate on the left-hand side and is governed by the elasticity of outflows. Therefore, the specification in column 3 is already consistent with the model extension.

In all specifications, the sensitivity of flows to the beginning-of-year reserve ratio is not significantly different from zero. More importantly, it is economically small. In column 1, we can reject at the 5% level that the regression coefficient of net flow on the reserve ratio is larger than 0.12. Combined with our estimate of the predictive power of reserves for future contract returns, a coefficient of 0.12 implies a semielasticity of the net flow rate to expected returns conditional on reserves of 4.6; that is, a change in reserves implying a one-percentage-point increase in future contract returns increases net flow by 4.6 percentage points.[26] In

[26] The semielasticity is computed as the regression coefficient of net flows on reserves rejected at 5% (0.12 in column 1 of panel A of Table 5) divided by the regression coefficient of the contract return on reserves (0.026 in column 1 of Table 3).

comparison, Drechsler, Savov, and Schnabl (2017) estimate the semielasticity of bank deposits net flows to the spread between interest rates on deposits and the Fed Fund rate to be 5.3. In the specification in log inflow in column 5, we reject at 5% a semielasticity of inflow to expected returns conditional on reserves of 0.4; that is, a change in reserves implying a one-percentage-point increase in future contract returns increases inflows by 40%.[27] This estimate can be compared to estimates of log specifications in the literature. Hortaçsu and Syverson (2004) estimate a semielasticity of the purchase of S&P 500 index funds to the fee of seven, and Koijen and Yogo (2022) estimate a semielasticity of the purchase of variable annuities to the fee of 16. We conjecture that demand for euro contracts is inelastic to reserves whereas the demand for deposits is elastic to deposit spreads and the demand for other investment products is elastic to fees, because the predictive power of reserves for contract returns is not easily comprehensible whereas deposit spreads and fees are readily observable and easy to understand. We provide evidence consistent with this interpretation at the end of this section and in Section 5.3.

Relation 2 implies that the OLS estimate of the flow-reserves relation has a downward bias if insurers face flow shocks anticipated by insurers but unobservable to the econometrician.[28] The bias caused by unobservable demand shocks can be corrected by instrumenting reserves using past asset returns. The instrument satisfies the exclusion restriction under the model assumption that past asset return is uncorrelated with current unobservable demand shocks. In practice, anticipation of demand shocks can lead insurers to adjust their asset allocation so that past risk premiums, and hence past asset returns, can be correlated with current demand shocks. To control for this potential threat to the validity of the instrument, we control for insurers' asset portfolio weights reported in the regulatory data (bonds, equities, real estate, loans, and other assets).[29,30] The first stage, reported in Internet Appendix Table E.1, is strongly significant: the t-stat is equal to 5.8 with standard errors two-way clustered by insurer and year. Panel B of Table 5 presents the second-stage regressions. The IV estimates of the flow-reserves sensitivity are slightly larger than the OLS estimates, but they remain statistically insignificant, and

---

27 The semielasticity is computed as the regression coefficient of log inflows on reserves rejected at 5% (1.1 in column 5 of panel A of Table 5) divided by the regression coefficient of the contract return on reserves (0.026 in column 1 of Table 3, multiplied by 100 to convert the return into percentage points).

28 To account for observable demand shocks, in Internet Appendix Table E.2, we estimate the flow regressions controlling for potential determinants of demand, and find that the coefficient for reserves remains small and insignificant.

29 We can further lag the asset return to further mitigate this identification concern. Internet Appendix Table E.3 shows that results are similar when we use the 2-year-lagged asset return as the instrument.

30 Another threat to the exclusion restriction could be that past asset returns are correlated with fees, because fees likely affect the demand for new contracts. In Internet Appendix Table E.1, we show that this is not the case: neither entry fees nor management fees is predicted by past asset return.

the economic magnitudes are small. To conclude, the empirical flow-reserves relationship rejects $\alpha > 0$ and is instead consistent with $\alpha \simeq 0$.[31]

Are there insurer characteristics that matter for investor demand? In Internet Appendix Table E.2, we regress the flow measures on lagged reserves and other insurer-level characteristics. We find that bank-insurance conglomerates enjoy a level of inflow 4.5 times larger than stand-alone insurers as well as a lower outflow rate. This suggests that the distribution channel is a key determinant of investor demand, with bank-insurance conglomerates benefiting from the advantage of tunneling depositors toward the contracts of their subsidiaries. Interestingly, we don't find net flows to be correlated with insurers' capital ratio.

The fact that flows are inelastic to reserves even though reserves predict returns does not imply that investors do not care about returns. Investors may fail to predict returns using reserves (see Section 5.3) but react to more transparent sources of predictability. In Internet E.4, we study how flows to euro contracts react to changes in the interest rate on a regulated savings product which competes with euro contracts. This product's regulated interest rate is readily observable and regularly discussed in the press, so that even unsophisticated investors are aware of its evolution. Consistent with flows reacting to more salient sources of predictability, we find that flows to euro contracts decrease when the regulated interest rate on the competing savings product increases. Also consistent with flows reacting to more salient events, we show in Internet Appendix Table E.5 that the flow-reserves relation becomes statistically significant during periods of financial crisis (Global Financial Crisis of 2008 and European sovereign debt crisis over 2011–2012). The economic magnitude remains small, however.

4.4 Reserves predict contract returns

Because reserves are not diluted by investor flows (as shown in the previous section) and are owed to investors (by regulation), the reserve ratio should predict future contract returns. We verify this prediction in Table 6. In column 1, we regress the contract return paid at the end of year t on the reserve ratio at the beginning of year t in the insurer-year panel with year fixed effects.[32] The coefficient for the beginning-of-year reserve ratio is positive and statistically significant at the 1% level. Therefore, the reserve ratio predicts the expected

[31] The IV analysis also allows us to distinguish investor strategies based on reserves from contract return chasing. In the OLS flow-reserves regressions, current reserves are correlated with past reserves because reserves are sticky, and past reserves are correlated with past contract return (see Equation (18) and Table 3). Therefore, the impact of current reserves on flows could be confounded by the impact of the past contract return. The IV flow-reserves regressions overcome this issue because the instrument (past asset return) is correlated with reserves (because shocks to asset returns are absorbed by reserves) but not with past contract return when $\alpha \simeq 0$ (see Equation (18) and Table 3). We thank an anonymous referee for raising this issue.

[32] We do not include insurer fixed effects because we are running a predictive regression, which would estimate insurer fixed effects on the entire sample period. In the (untabulated) regression with insurer fixed effects, the coefficient for the lagged reserve ratio is 0.03 and significant at the 1% level.

Table 6
Contract return predictability

	Contract return in year				
	t	$t+1$	$t+2$	$t+3$	$t+4$
	(1)	(2)	(3)	(4)	(5)
Reserves at beginning of year t	.025***	.024***	.023**	.019**	.019*
	(.0074)	(.0073)	(.0078)	(.0087)	(.0088)
Year FE	✓	✓	✓	✓	✓
R^2	.69	.71	.62	.61	.57
Observations	978	859	783	717	645

Panel regressions at the insurer-year level for 76 insurers over 2000–2015. *Contract return* is the annual before-fees contract return the end of years t (column 1), $t+1$ (column 2), …, $t+4$ (column 5). *Reserves at beginning of year t* is total reserves at the beginning-of-year t normalized by total account value. All regressions include year fixed effects and are weighted by the insurer share in aggregate account value in the current year. Standard errors two-way clustered by insurer and year are reported in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

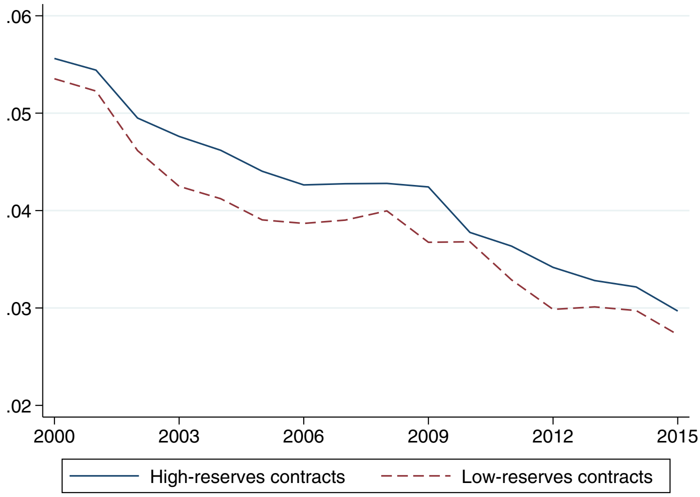

Figure 3
High-reserves contracts outperform low-reserves contracts
The figure illustrates the return of the portfolio of contracts with the above-median reserve ratio (solid blue) is always higher than the return of the portfolio of contracts with the below-median reserve ratio (dashed red).

contract return at a 1-year horizon: contracts with higher reserves have higher expected returns.

Higher reserves predict a higher expected contract return because reserves are eventually distributed to investors, not because higher reserves are associated with higher risk. To show this, we consider a zero-cost portfolio that is invested long in contracts with high reserves and short in contracts with low reserves. At the beginning of each year, we rank insurers on the $[0,1]$ interval based on the beginning-of-year reserve ratio, and use portfolio weights proportional to insurers' rank minus one-half. Figure 3 shows that the long leg of the portfolio outperforms the short leg in every year of the sample period. Table 7

Table 7
High-reserves contracts are not riskier

	Low-reserves contracts (1)	High-reserves contracts (2)	High minus Low portfolio (3)
Mean return	.039	.042	.0034***
	(.0029)	(.0030)	(.00035)
Alpha	.039	.042	.0034***
	(.0027)	(.0029)	(.00032)
Beta	−.012	−.0092	.0027*
	(.0070)	(.0075)	(.00092)
SD return	.0051	.0042	−.00091
	(.00034)	(.00049)	(.00069)

Performance of a portfolio long contracts with beginning-of-year reserves above median and short contracts with beginning-of-year reserves below median with portfolio weights proportional to the contract rank rescaled between minus one and one times the contract's total account value. Column 1 shows the performance of the short leg, column 2 of the long leg, and column 3 the performance of the long-short portfolio. *Mean return* is the average return of the leg/portfolio. *Alpha* and *Beta* are the intercept and loading on the market in the market model. *SD return* is the time-series average of the cross-sectional standard deviation of contract return within the leg in columns 1 and 2, and it is the difference between that of the long leg and that of the short leg in column 3. Newey-West standard errors with two lags are reported in parentheses. In column 3, $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

shows performance statistics for each leg of the portfolio and for the long-short portfolio. The first row confirms that higher reserves predict higher expected returns: Average returns are 34 bps per year higher for high-reserves contracts than for low-reserves contracts.

The second and third rows of Table 7 report the estimates of a market model. The difference in market beta between high-reserves and low-reserves contracts is a precisely estimated zero (a difference in beta larger than 0.01 is rejected at the 1% level), implying alpha is 34-bps higher for high-reserves contracts than for low-reserves contracts, on average.[33] Therefore, the predictability of expected contract returns does not reflect a compensation for market risk.

The fourth row reports the cross-sectional standard deviations of high- and low-reserves contracts returns, averaged over time. We find the difference between the two groups is a precisely estimated zero. Therefore, the predictability of expected contract returns does not reflect a compensation for idiosyncratic risk either.

Reserves should predict contract returns not only at 1 year but also at longer horizons because reserves are only progressively distributed to investors. We show in Internet Appendix F that the predictive power of reserves for future contract returns should decay at the same rate as the one at which the reserve ratio mean reverts. The reserve ratio mean reverts for two reasons. First, reserves are progressively credited to investors' accounts (at a rate of 3% per year in columns 1 and 2 of Table 3). Second, inflows dilute reserves at a rate equal

[33] The large t-stat of the alpha estimate reflects the fact that the predictive power of reserves is almost mechanical. Because reserves must eventually be distributed to investors, they must predict future contract returns. The null hypothesis rejected by the nonzero alpha is merely that insurers do not divert reserves.

to the unconditional net flow rate (2.4% per year in Table 1) plus a term that depends on the sensitivity of flows to reserves (equal to zero in Table 5). Thus, the reserve ratio should mean revert at a rate of 5.4% per year. The predictive power of reserves for future contract returns should also decay at a rate of 5.4% per year.

In columns 2 to 5 of Table 6, we check that the data are consistent with the above calculation. We regress contract return in years t, $t+1$, ..., $t+4$, on the reserve ratio at the beginning of year t. The regression coefficient for the initial reserve ratio decays at a rate of about 7%, which is close to the predicted rate of 5.4%. In conclusion, reserves predict future contract returns over many years.

4.5 Where does cross-sectional variation in reserves come from?

Our analysis of the contract return policy and flow-reserves relation relies on variation in reserves across insurers. Where does cross-sectional variation in reserves come from in the first place? To answer this question, we start from the accounting identity (4) that determines the evolution of reserves over time, and we estimate the contribution of each variable to the cross-sectional variance in the reserve ratio. Three terms must be added to the model equation (4) for this accounting identity to hold exactly in practice. First, when insurer j reinsures euro contract liabilities, net income from reinsurance $Reins_{j,t}$ is added to the asset return in year t and at least 85% of it must be paid to investors either immediately as contract return or later via reserves.[34] Second, technical income, which is equal to fees collected from investors minus operating costs, must be shared with investors either immediately as contract return or later via reserves.[35] Therefore, net income from reinsurance and fees must be added to, and operating costs must be subtracted from, the left-hand side of (4) for this accounting identity to hold exactly empirically. Rearranging terms, the reserve ratio evolves according to

$$\frac{R_{j,t}}{V_{j,t}}=\left[\frac{R_{j,t-1}}{V_{j,t-1}}+x_{j,t}\left(1+\frac{R_{j,t-1}}{V_{j,t-1}}\right)-y_{j,t}\right.$$
$$\left.-\left(\frac{\Pi_{j,t}}{V_{j,t-1}}+\frac{Costs_{j,t}}{V_{j,t-1}}-\frac{Fees_{j,t}}{V_{j,t-1}}\right)+\frac{Reins_{j,t}}{V_{j,t-1}}\right]\frac{V_{j,t-1}}{V_{j,t}}. \tag{20}$$

The regulatory filings do not allow us to decompose the insurer profit plus operating costs minus fees (the second big parenthesis in (20)) into its three components, so we use a variable that encompasses the three components. Iterating equation (20), the reserve ratio in year t writes as a function of the reserve ratio in year $t-h$, with $h\geq 1$, and the sequence from year $t-h+1$ to

34 Insurers rarely reinsure their liabilities. The absolute value of net income from reinsurance represents 5 bps of account value on average.

35 At least 90% of the technical income must be paid to investors when the technical income is positive, and 100% when it is negative. See Internet Appendix B.2.

Table 8
Asset return explains most of the cross-sectional variation in reserves

	Share of variance of h-year change in reserve ratio explained by:					
	Asset return (1)	Contract return (2)	Insurer profit, costs, fees (3)	Reinsurance (4)	Account value growth (5)	Covariance terms (6)
$h=1$ year	0.88	<0.01	0.07	<0.01	0.02	0.02
$h=2$ years	0.81	0.01	0.07	0.01	0.03	0.07
$h=3$ years	0.79	0.01	0.08	0.02	0.04	0.06
$h=4$ years	0.77	0.01	0.12	0.01	0.04	0.05
$h=5$ years	0.72	<0.01	0.14	0.01	0.04	0.08
$h=6$ years	0.69	<0.01	0.17	<0.01	0.04	0.09
$h=7$ years	0.64	<0.01	0.21	<0.01	0.05	0.1

Each cell in columns 1 to 5 corresponds to a different regression at the insurer-year level. The statistics reported in each cell is the R^2 of the regression of the h-year change in the reserve ratio on the variable indicated in the column heading, run after partialling out the initial reserve ratio, year fixed effects and the initial ratio interacted with year fixed effects, divided by the R^2 of the same regression but including the five explanatory variables. The statistics in column 6 are calculated as one minus the sum of the other columns.

year t of five variables: asset return; contract return; insurer profit, costs, and fees; reinsurance; and asset growth. To estimate the contribution of each of these five variables to the cross-sectional variance of reserves, we first partial out initial reserves from the reserve ratio and from the six explanatory variables. Specifically, we compute the residuals of insurer-year-level panel regressions of the reserve ratio on the h-year lagged reserve ratio, year fixed effects, and the interaction between the lagged reserve ratio and year fixed effects. We do the same for each one of the explanatory variables. Then, for each explanatory variable, we regress the residual of the current reserve ratio on the residual of the explanatory variable and its lags.

The R^2s of these regressions measure the contribution of each explanatory variable to the cross-sectional variance of the reserve ratio. To adjust for the nonlinearity of the accounting identity and the small data discrepancies, we report normalized R^2s in columns 1 to 5 of Table 8.[36] One minus the sum of the normalized R^2s reflects the contribution of covariance terms to the total variance. This contribution is reported in the rightmost column of Table 8.

At horizon $h=1$ year, the asset return explains 88% of the cross-sectional variance in reserves while the other variables explain little. This result complements the pattern in Figure 2 showing that the asset return explains virtually all of the time-series variation in reserves. At longer horizon, the asset return still explains most of the cross-sectional variation in reserves. The contribution of insurer profit, costs, and fees increases at longer horizons

[36] Each normalized R^2 is equal to the original R^2 of the regressions using each one of the explanatory variables divided by the R^2 of the regression using all the explanatory variables. The normalized R^2s measure the fraction of the variance of the reserve ratio explained by each explanatory variable. The R^2 from the regression including all five explanatory variables is less than one because the regression specification is additive whereas account value growth enters multiplicatively in the accounting identity (20), and because of small discrepancies in the data during acquisitions and divestures. The five normalized R^2s would sum to one only if the explanatory variables were uncorrelated with each other.

because these variables are persistent and hence their effect builds up over time. The contribution of the contract return is small at every horizon, accounting for less than 1% of the cross-sectional variance in reserves, which reflects that there is little variation in the contract return policy across insurers.

In Internet Appendix Table G.2, we decompose the asset return $x_{j,t}$ into a systematic risk component and an idiosyncratic risk component. This decomposition brings the additional result that heterogeneity in systematic risk exposure only explains a small part of cross-sectional variation in reserves, so that the heterogeneity in asset return is mainly because of idiosyncratic risk. This finding is consistent with anecdotal evidence on the heterogeneity in asset holdings across insurers. For instance, the European sovereign debt crisis of 2011–2012 has shed light on the differential exposure to sovereign bond issuing countries across insurers.

5. Why Is Demand Inelastic to Reserves?

5.1 Switching costs

We test whether the low elasticity of flows to reserves is explained by switching costs created by the tax treatment of euro contracts. As described in Section 1.1, contract returns are taxed upon withdrawal at a rate that depends on the age of the contract at the time of withdrawal: the tax rate is 35% if contract age is less than 4 years, 15% between 4 and 8 years, and 7.5% after 8 years.[37] Therefore, an investor owning a contract and willing to increase her investment in euro contracts faces a tax incentive to add cash to her existing contract rather than buying a new contract.

In contrast to other investors, new investors are not subject to the tax-induced switching cost. Therefore, if the low elasticity of flows to reserves is explained by switching costs, purchases of new contracts should react to reserves. Instead, if the low elasticity is explained by something else, then purchases of new contracts should be as inelastic to reserves as total flows are. We test whether purchases of new contracts react to reserves using information on the number of new contracts purchased from each insurer in each year. Insurers have been required to report this information since 2006, therefore, the sample period for this test is restricted to 2006–2015. We regress the number of new contracts purchased divided by the number of outstanding contracts on the beginning-of-year reserve ratio. Table 9 shows that both in our OLS and IV estimations, new investors' inflows are not sensitive to the level of reserves. We conclude that switching costs induced by taxes cannot explain the low elasticity of inflows to reserves.

Another switching cost stems from entry fees that investors incur when they add cash to their contract, creating a disincentive to move cash from one contract

[37] See Internet Appendix H for an estimate of the tax-induced switching cost.

Table 9
Inflows from new investors

	Purchases of new contracts	
	OLS	IV
	(1)	(1)
Lagged reserves	.029	.092
	(.05)	(.15)
Year FE	✓	✓
Insurer FE	✓	✓
R^2	.49	.49
Observations	581	577

Panel regressions at the insurer-year level for 67 insurers over 2006–2015 *Purchases of new contracts* is the number of new contracts purchased in the current year divided by the beginning-of-year outstanding number of contracts. *Lagged reserves* is the beginning-of-year level of reserves normalized by total account value. Column 1 shows the OLS regression. Column 2 shows the IV regression in which the insurer's beginning-of-year reserve ratio is instrumented using the insurer's asset return in the previous year. All regressions include insurer and year fixed effects and are weighted by the insurer share in aggregate account value in the current year. Standard errors two-way clustered by insurer and year are reported in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

to another. However, entry fees do not distort the choice of contract for newly invested money, because entry fees are incurred regardless of the contract chosen, and we have shown insurers do not adjust entry fees to the level of reserves (Table 4). Therefore, entry fees cannot explain the low elasticity of total inflow to reserves (Column 2 of Table 5) nor the low elasticity of new investors' inflow (Table 9).

5.2 Preferences for nonprice contract characteristics

An alternative hypothesis is that investors are inelastic to reserves because they have heterogeneous and strong preferences for other contract characteristics, such as risk exposures. Inconsistent with this hypothesis, we showed in Table 7 that high- and low-reserves contracts are similar in their exposure to both systematic and idiosyncratic risks. In addition, as we show in the next subsection, the elasticity of demand to reserves varies systematically with investor wealth, a proxy for financial sophistication. The hypothesis that investors are inelastic to reserves because of preferences for certain risk exposures would require to explain why the heterogeneity in such risk preferences is significantly smaller among more sophisticated investors.

5.3 Investor sophistication

We test whether flows are inelastic to reserves because investors lack the knowledge to predict contract returns using reserves. This lack of knowledge could be due to investors simply not understanding that reserves predict returns, or perhaps investors not being able to obtain information on the level of reserves.[38] To test that hypothesis, we study whether the flow-reserves sensitivity varies across investors with different levels of financial

[38] Although insurers' annual reports contain information on the level of reserves, it often is incomplete or consolidated at the group level.

sophistication. We proxy for investor sophistication using the investment amount, the idea being that financial sophistication is correlated with wealth, for instance, if investors must incur a fixed cost to acquire the knowledge necessary to predict returns (Lusardi and Mitchell 2014).

We construct the proxy for investor sophistication using contract-level data collected by the insurance supervisor for the years 2011 to 2015. The data contain information on the number of investors, the total account value, and the net-of-fees return for every contract. We calculate the average individual account value as the total account value divided by the number of investors. We classify contracts into three bins according to the average account value: less than 50,000 euros, 50,000–250,000 euros, and more than 250,000 euros. We also construct net flows at the contract level.

We exploit cross-sectional variation in investor sophistication along two dimensions. First, we exploit variation across insurers. Some insurers cater to wealthier—hence more sophisticated—clienteles. Second, we exploit variation across contracts within a given insurer. As described in Section 1.1, insurers often offer different contracts with different minimum investment amounts that target different clienteles. A crucial feature of the institutional framework is that reserves are pooled across all contracts of a given insurer, so that reserves predict returns for all contracts. Therefore, we can exploit cross-contract variation in investor sophistication to test whether the flow-reserves sensitivity varies within a given insurer-year. We regress net flows at the contract level on the beginning-of-year reserve ratio interacted with dummy variables for each bin of average account value (and on the noninteracted dummy variables).

The first specification (column 1 of Table 10) does not include insurer-year fixed effects and thus exploits cross-insurer variation in investor sophistication. The flow-reserves sensitivity is small and statistically insignificant for contracts with both small and intermediate average account values (below 250,000 euros per investor). By contrast, the flow-reserves sensitivity is positive and statistically significant at the 10% level for contracts with larger average account value (above 250,000 euros per investor). Combined with our estimate of the predictive power of reserves for future contract returns, the point estimate implies a semielasticity of sophisticated net flows to expected returns conditional on reserves of 14; that is, a change in reserves implying a one-percentage-point increase in future contract returns increases sophisticated net flow by 14 percentage points.[39]

The second specification (column 2 of Table 10) includes insurer-year fixed effects and thus isolates cross-contract variation in investor sophistication within insurer-years. In that case, the absolute level of the flow-reserves sensitivity is no longer identified, because it is defined at the insurer-year

[39] The semielasticity is computed as the regression coefficient of net flows on reserves in contracts with average account value above 250,000 euros (0.36 in column 1 of Table 10) divided by the regression coefficient of contract return on reserves (0.026 in column 1 of Table 3).

Table 10
Financial sophistication

	Contract-level net flows			
	OLS (1)	OLS (2)	IV (3)	IV (4)
Lagged reserves × (Avg account value 0–50 k€)	−.059		−.19	
	(.17)		(.26)	
Lagged reserves × (Avg account value 50–250 k€)	.014	.13	.0028	.32
	(.17)	(.076)	(.22)	(.21)
Lagged reserves × (Avg account value 250+ k€)	.36*	.41***	.54*	.84**
	(.13)	(.0031)	(.21)	(.27)
Avg account value bin FE	✓	✓	✓	✓
Year FE	✓		✓	
Insurer FE	✓		✓	
Insurer-year FE		✓		✓
R^2	.13	.16	.15	.18
Observations	7,272	7,272	7,268	7,268

Panel regressions at the contract-year level *Contract-level net flows* is contract net flows normalized by contract total account value. *Lagged reserves* is insurer beginning-of-year level of reserves normalized by insurer total account value. *Avg account value RANGE* is a dummy variable equal to one if the contract average account value (calculated as contract total account value divided by number of investors) lies in *RANGE*. All regressions include these noninteracted dummy variables in addition to their interaction with lagged reserves. Columns 1 and 3 include insurer and year fixed effects. Columns 2 and 4 include insurer-year fixed effects. Columns 1 and 2 show OLS regressions. Columns 3 and 4 show IV regressions in which the insurer's beginning-of-year reserve ratio is instrumented using the insurer's asset return in the previous year. All regressions are weighted by the contract share in aggregate account value in the current year. Standard errors two-way clustered by insurer and year are reported in parentheses. $*p < .1$; $**p < .05$; $***p < .01$.

level. We use the small-average-account-value category as the reference group. The results are consistent with those obtained in the first specification: the flow-reserves sensitivity is larger for contracts with large account values than for contracts with smaller account values. The difference is significant at the 1% level. The IV estimates yield similar results (Columns 3 and 4).

These results are consistent with low financial sophistication explaining the low elasticity of flows to reserves. Only investors with large investments time reserves, because they are more likely sophisticated and have incentives to understand the mechanics of intercohort redistribution through reserves. Therefore, perhaps surprisingly, the lack of household financial sophistication enables more risk sharing than would be possible if households were perfectly informed and acted accordingly.

5.4 Do arbitrage opportunities exist?

Does the predictability of contract returns generate an arbitrage opportunity that an investor who perfectly understands euro contracts could trade on? If this was the case, a single arbitrageur would unravel the intercohort risk sharing equilibrium. In this section, we show the role of the capital income tax is crucial to prevent this from happening.

Proposition 3 implies contract returns can be replicated up to a constant by a portfolio composed of the assets held by insurers and the risk-free asset. Because euro contracts cannot be sold short, if an arbitrage strategy exists, it consists in long positions in euro contracts and short positions in the assets held by insurers and the risk-free asset. In France, as in many other countries, households interest expenses in levered financial investments are not tax deductible.[40] Therefore, the return on the long leg of the arbitrage strategy is taxable, whereas the return on the short leg is not tax deductible. We write the capital income tax rate as τ and the risk-free rate as r_f. We show in Internet Appendix A.7 that one euro invested long in contract j hedged with short positions in the replicating portfolio generates a risk-free profit:

$$\pi_{j,t}^{arb} \simeq \left[1-\frac{(1-\tau)(1-\phi)r}{1+r}\right](r-r_f) \\ +(1-\tau)(1-\phi)r\mathcal{R}_{j,t-1}-\tau r-(1-\tau)\phi r \tag{21}$$

when $\alpha \simeq 0$.

Equation (21) highlights the two sources of arbitrage profits and the two arbitrage costs. First, contract returns are hedged against asset risk, yet they earn the risk premium on the risky assets held by insurers. An arbitrageur going long the contract and short the underlying assets earns the risk premium without bearing the associated risk. This source of arbitrage profits is reflected in the first term of $\pi_{j,t}^{arb}$: $r-r_f>0$ is the risk premium, and the term in brackets is equal to one minus the exposure of the after-tax contract return to asset risk. This term is close to one because contract returns are almost perfectly hedged against asset risk. The second source of arbitrage profits comes from the predictable distribution of reserves to contract holders. It is reflected in the second term of $\pi_{j,t}^{arb}$, which is proportional to the reserve ratio. The costs of the arbitrage strategy are the tax on the expected asset return (third term of $\pi_{j,t}^{arb}$) and the insurer compensation (fourth term).

The key insight from (21) is that a capital income tax is sufficient to eliminate arbitrage opportunities; that is, $\pi_{j,t}^{arb}<0$ if τ is large enough. This result does not rely on euro contracts benefiting from a tax advantage, because it assumes the returns on all long positions are taxed at a uniform rate τ. Neither does this result rely on euro contracts being expensive, because it holds even when ϕ is arbitrarily close to zero.

We calibrate the terms in (21) in Internet Appendix A.7 and show that arbitrage opportunities are eliminated if the capital income tax rate is greater than 26%. In reality, the applicable tax rate depends on the contract holding period. At the end of the sample period, the lowest possible tax rate is 23% (15.5% of social security contributions plus 7.5% of capital income tax). Hence, the actual minimum tax rate is close to our estimate of the minimum tax rate necessary to eliminate arbitrage opportunities.

[40] In some countries, including France and the United States, interest paid on mortgages, student loans and business loans often are tax deductible, but interest expenses in levered financial investments typically are not. Since euro contracts can only be purchased by households, the relevant tax regime is that of households.

Note the absence of arbitrage opportunities is not contradictory with our finding in Section 5.3 that the flow-reserves relation is statistically significant among investors with large invested amounts. Indeed, conditionally on saving a positive amount, sophisticated households always should buy those contracts with high reserves rather than those with low reserves, even in the presence of a capital income tax. Yet, an arbitrage strategy that consists in buying euro contracts and shorting the underlying assets and the risk-free asset is not profitable in the presence of a large enough capital income tax.

6. Welfare Analysis

We now analyze the welfare implications of intercohort risk sharing. Consider an investor with initial wealth W_t, who invests for h years and derives utility over her final wealth W_{t+h} with constant relative-risk aversion (CRRA) γ. To measure the impact of intercohort risk sharing on welfare, we calculate the change in the investor's expected utility induced by adding euro contracts to the investment opportunity set. To make this calculation, we compare expected utility in two cases. In the first case, the investor has access to two assets: the portfolio of assets held by insurer j with return $x_{j,t}$; and the risk-free asset with return r_f. CRRA preferences and i.i.d. returns imply that the optimal portfolio allocation does not depend on wealth and is constant through time.

In the second case, the euro contract sold by insurer j with return $y_{j,t}$ is added to the investment opportunity set in addition to the insurer asset portfolio and the risk-free asset. Contract return smoothing implies that the contract return is serially correlated, so the optimal portfolio allocation may not be constant through time. However, in Section 4.3, we show that investor flow fails to respond to investment opportunities and, in Section 5.3, we show this lack of responsiveness is due to investors' lack of sophistication. In this context, assuming that investors do not rebalance their portfolio may be a more realistic assumption than frictionless rebalancing. Accordingly, we assume that the investor chooses the optimal portfolio allocation in the class of time-invariant allocations.

In both cases, we allow the portfolio weight on the insurer asset portfolio to exceed one. In practice, this does not necessarily require the investor to leverage because the insurer asset portfolio includes the risk-free asset. Insurers hold approximately 80% of corporate and sovereign bonds, 14% of equities, and some real estate and loans. Regulatory filings do not report more granular information on insurers' asset allocation, unfortunately, so for the sake of illustration, suppose that two-third is in risky assets and one-third is risk-free. An individual investing her entire portfolio in risky assets would therefore have a portfolio weight of $1/(2/3)=1.5$ on the insurer asset portfolio and a weight of negative 0.5 on the risk-free asset, even though the investor does not have a short position in the risk-free asset. In the case in which the euro contract is in the investment opportunity set, we impose that the weight on the contract cannot exceed one; for example, we do not allow the investor to leverage to buy euro contracts. This assumption is justified by the fact that, in practice, leveraging to buy euro contracts is made unprofitable by the capital income tax, as shown in Section 5.4.

The contract return is given by Equation (18) as a function of the insurer asset return and reserves. In this equation, we calibrate the elasticity of demand using the empirical estimate in Table 3 and set $\alpha=0$. The expression for ρ is given in Internet Appendix A.8 and is equal to relative-risk aversion times the share of the investor's total wealth held in euro contracts. We calibrate this share using that euro contracts represent one-third of aggregate household financial wealth and that financial wealth represents 40% of total household wealth. Therefore, $\rho=\gamma\times 0.33\times 0.4$. Regulation imposes $\phi=0.15$. We calibrate the moments of the insurer asset return distribution to the data: mean $r=4.4\%$ per year (see footnote 24) and standard deviation 4.4% per year (see Table 1), and we assume it is lognormal and i.i.d. The risk-free rate is 3% per year such that the risk premium on the insurer asset portfolio is 1.4% per year. We assume that initial reserves are equal to their unconditional model mean, which is normalized to zero. This implies that the unconditional expected contract return is equal to the expected asset return on a before-fee basis, for example, $r=4.4\%$. To isolate the impact of intercohort risk sharing, we further assume that fees on direct investment in the insurer asset portfolio are equal to fees on euro contracts. This assumption implies that the expected contract return is equal to the expected asset return on an after-fee basis, $(1-\phi)r$. Using these parameters, we simulate the contract return based on the empirical moments of the asset return distribution. The moments of the contract return distribution are implied by the model and depend on the model parameters, in particular the elasticity of demand α, which is estimated in the data in Section 4. We solve for the optimal portfolio in the two cases described above (analytically when the contract is not in the investment opportunity set, and numerically when it is), and we calculate expected utility in each case.

Panel A of Table 11 contains the results for a level of relative-risk aversion $\gamma=2$ and investment horizon $h=12$ years. Before analyzing optimal portfolios, we present summary statistics for a portfolio fully invested in the insurer asset portfolio (first row) and a portfolio fully invested in the euro contract (second row) to illustrate the impact of return smoothing. By assumption (see above), both have the same net-of-fees expected return, equal to $(1-\phi)r$ per year, yielding a risk premium over the risk-free rate equal to 12.8% over 12 years.[41] Due to return smoothing, the standard deviation of the euro contract is about four times lower than the standard deviation of the insurer asset portfolio. As a

[41] The euro contract has a slightly higher 12-year risk premium than the insurer assets even though both have the same annual expected return as a result of Jensen's inequality when compounding returns and the fact that the contract return has lower variance.

Table 11
Welfare analysis

	Portfolio weights			h-year return (%)			
	Insurer assets (1)	Risk-free asset (2)	Euro contract (3)	Risk premium (4)	SD (5)	Sharpe ratio (6)	Certainty equivalent gain (7)
				A. $\gamma=2$, $\alpha=0$, $h=12$ years			
Insurer assets	1			12.8	19.5	0.66	
Contract			1	12.9	6.1	2.13	2.2
Optimal w/o contract	2.76	−1.76		37.9	62.9	0.6	
Optimal w/ contract	2.62	−2.62	1	53.4	71.8	0.74	10.8
				B. $\gamma=5$, $\alpha=0$, $h=12$ years			
Insurer assets	1			12.8	19.5	0.66	
Contract			1	12.9	6.1	2.13	5.2
Optimal w/o contract	1.11	−0.11		14.3	21.8	0.65	
Optimal w/ contract	0.9	−0.9	1	25.9	25	1.04	9.7
				C. $\gamma=2$, $\alpha=5$, $h=12$ years			
Insurer assets	1			12.8	19.5	0.66	
Contract			1	12.9	11.7	1.11	1.6
Optimal w/o contract	2.76	−1.76		37.9	62.9	0.6	
Optimal w/ contract	2.24	−2.24	1	46.3	67.7	0.68	5.9

Each row is a portfolio. Columns 1 through 3 show the portfolio weights on the insurer's asset portfolio with return $x_{j,t}$, insurer's contract with return $y_{j,t}$, and risk-free asset with return r_f. Columns 4 through 6 show properties the portfolio return over the $h=12$ year holding period: risk premium over the risk-free asset, standard deviation, and Sharpe ratio. Column 7 shows the certainty equivalent gain (again over the $h=12$ year holding period) of the portfolio relative to the portfolio on the previous row.

In each panel, the first row is a portfolio fully invested in the insurer's asset portfolio. The second row is a portfolio fully invested in the insurer's contract. The third row is the optimal portfolio for a CRRA investor with $h=12$ year investment horizon when the investment opportunity set is the insurer's asset portfolio and the risk-free asset. The fourth row is the optimal portfolio when the insurer's contract is added to the investment opportunity set.

Each panel considers a different set of parameter values for the investor's relative-risk aversion is γ, elasticity of demand to the expected contract return condition on reserves α, and investor's investment horizon h.

result, the Sharpe ratio of the euro contract is approximately four times larger than the insurer asset portfolio.[42] The certainty equivalent gain from investing in the euro contract relative to investing in the insurer asset portfolio is 2.2 percentage points over 12 years.[43]

The third row shows the optimal portfolio when the investor can invest in the insurer asset portfolio and the risk-free asset, but not in the contract. The fourth row is the optimal portfolio when the investor can also invest in the euro contract. The optimal weight on the euro contract is one, which means that the investor is up against the constraint that she cannot use leverage to buy additional euro contracts. Comparing the optimal portfolio with the euro contract to the case

42 The high Sharpe ratio of the euro contract follows mechanically from the fact that the contract earns the risk premium associated to the insurer assets while bearing only a fraction of the risk, because risk is shared across cohorts. This high Sharpe ratio cannot be leveraged to produce highly profitable trading strategies because of the tax treatment of short positions for retail investors (see Section 5.4).

43 The certainty equivalent CE is calculated as $\left(\prod_{\tau=t+1}^{t+h}(1+r_{1,\tau})\right)^{1-\gamma}/(1-\gamma)=\left(CE+\prod_{\tau=t+1}^{t+h}(1+r_{2,\tau})\right)^{1-\gamma}/(1-\gamma)$ where $r_{1,\tau}$ is the contract return and $r_{2,\tau}$ is the insurer asset return.

without the euro contract, we observe that the investment in the euro contract substitutes for the investment in the risk-free asset (whose weight decreases by 0.86) and barely changes the investment in insurer assets (whose weight decreases by 0.14). The reason being that the contract is quite safe owing to intercohort risk sharing. It is therefore a good (and superior) substitute to the risk-free asset, which leads the investor to substitute the risk-free asset with the euro contract. This finding is consistent with the analysis in Section 3.3, where we show that contract returns are hedged against asset risk yet earn the risk premium on the risky assets held by insurers. The weight on the insurer assets decreases slightly because the euro contract is not completely safe and is correlated with the insurer assets, especially at long horizon because asset risk is absorbed by reserves and reserves are eventually distributed to contracts.

The certainty equivalent gain from adding the euro contract to the investment opportunity set is 10.8 percentage points over the average 12-year holding period. It is substantially higher than the certainty equivalent gain from investing fully in the contract relative to investing fully in the insurer assets reported above. The reason being that the euro contract is optimally used as a substitute to the risk-free asset, such that adding the euro contract to an optimal portfolio yields higher welfare gains than the relative gain from investing in the euro contract only versus the insurer assets only. In line with survey evidence showing that euro contract investors value the safety of these products (Darmon and Pagenelle 2005; Bianchi 2018), the key advantage of holding a euro contract is to earn a long-run average return similar to that of a portfolio of long-term bonds and stocks while facing low volatility.

Panel B considers a higher level of relative-risk aversion: $\gamma = 5$. The first two rows show that the certainty equivalent gain from investing fully in the contract relative to investing fully in the insurer assets increases. The reason is intuitive: the reduction is risk yields a larger welfare gain when the investor is more risk averse. However, these portfolios are not optimal portfolios. When we consider optimal portfolios in the bottom two rows, the certainty equivalent gain from adding the euro contract to the investment opportunity set hardly depends on the level of risk aversion. The reason being that the investor optimally uses the euro contract as a substitute to the risk-free asset, so the welfare gain stems from the return spread between the contract and the risk-free asset and barely varies with the investor's risk aversion.

The estimated welfare gain from intercohort risk sharing can be interpreted as the maximum difference in fees that investors are willing to pay to hold a euro contract instead of the insurer's assets. Future empirical work using data on fees can use these estimates to tease out the split of the welfare gain from euro contracts between insurers and investors.

In panel C, we perform a counterfactual analysis of the impact of demand elasticity on intercohort risk sharing. We consider a higher elasticity of demand to expected returns conditional on reserves and calibrate it to the elasticity of deposits to the deposit spread estimated by Drechsler, Savov, and Schnabl

Table 12
Welfare analysis and investment horizon

		Certainty equivalent gain (% per year)	
Investment horizon	SD contract return / SD insurer asset return (1)	Contract relative to insurer assets (2)	Optimal portfolio with contract relative to without contract (3)
1 year	0.04	0.15	0.71
5 years	0.14	0.13	0.68
12 years	0.31	0.13	0.62
20 years	0.51	0.12	0.56
30 years	0.78	0.09	0.47
Geometric with mean 12 years		0.12	0.57

Column 1 shows the ratio of the standard deviation of the h-year contract return divided by the standard deviation of the h-year insurer asset return. Column 2 shows the certainty equivalent gain from investing in the euro contract relative to investing in the insurer asset portfolio. Column 3 shows the certainty equivalent gain from investing in the optimal portfolio when the investment opportunity set is the insurer asset portfolio, the euro contract, and the risk-free asset, relative to investing in the optimal portfolio when the euro contract is excluded from the investment opportunity set.

Each row corresponds to an investor with a different holding period h. h is deterministic, except in the last row where h has a geometric distribution truncated at 30 years with mean 12 years.

(2017): $\alpha = 5$. Compared to the case of inelastic demand, a higher elasticity reduces the ability of insurers to share asset risk across investor cohorts, increasing the riskiness of contracts. The standard deviation of contract return is roughly twice as large when $\alpha = 5$ than when $\alpha = 0$ (second row of panel C vs. panel A). Because the contract becomes riskier and correlated with insurer assets, it is less good of a substitute for the risk-free asset. The optimal portfolio weight on insurer assets decreases, reducing the expected return of the optimal portfolio (fourth row of panel C vs. panel A). As a result, the certainty equivalent gain from adding the contract to the investment opportunity set is approximately halved compared to the case of inelastic demand.

We repeat the welfare analysis by varying the investment horizon in Table 12. Column 1 shows the ratio of the standard deviation of the contract return to that of the insurer asset return over the corresponding holding period. The ratio is below one, reflecting return smoothing, but increases with horizon, meaning that the impact of return smoothing decreases over time; for example, the contract is relatively riskier at longer horizon. The reason being that return smoothing generates positive autocorrelation in the contract return, so that increasing the investment horizon increases the variance of the compounded contract return by more than that of the underlying asset return. Therefore, the annualized certainty equivalent gain from investing in the portfolio fully invested in the euro contract relative to the portfolio fully invested in the insurer assets decreases with the investment horizon (Column 2). The annualized certainty equivalent gain from adding the euro contract to the investment opportunity set of the optimal portfolio also decreases with horizon but remains substantial even at a 30-year horizon: 47 bps per year (Column 3).

We consider an uncertain investment horizon in the bottom row of Table 12. Each year, the investor is hit by a liquidity shock with probability λ, in which case she liquidates her portfolio. Specifically, the holding period h has a geometric distribution truncated at 30 years with a mean of 12 years. The optimal portfolio is determined under rational expectations about the distribution of h. The welfare gain brought about by the euro contract is similar to that when the investment horizon is deterministic and equal to 12 years. The reason being that the welfare benefits of intercohort risk sharing decrease approximately linearly with the horizon. Therefore, uncertainty regarding the investment horizon does not have a significant effect on these welfare benefits.

7. Conclusion

We provide the first evidence of a large scale, and private, implementation of intercohort risk sharing. The evidence implies that financial intermediaries can complete markets, by allowing different investor cohorts to share risk, which they cannot achieve even in fully developed financial markets. Such intercohort risk sharing is desirable from an ex ante welfare perspective, that is, under the Rawlsian veil of ignorance (Gordon and Varian 1988; Ball and Mankiw 2007).

Private implementation of intercohort risk sharing requires a two-sided commitment problem be overcome (Allen and Gale 1997). First, investors must remain invested in contracts even when reserves are low. We show that investment flows are inelastic to reserves, and do not tumble when reserves are low. The demand elasticity is lower among investors who are expected to have lower financial sophistication. Therefore, perhaps paradoxically, lower investor sophistication enables a better sharing of risk—across investor cohorts—than what would be possible if investors were perfectly informed.

Second, insurers must credibly commit not to run away with reserves, which they might be tempted to do when reserves are high. Regulation solves this side of the commitment problem, ensuring reserves are eventually returned to investors. This suggests a reason intercohort risk-sharing savings products exist in several European countries, where such regulation exists, but not in the United States, where it does not.

These results have implications for real investment, which we leave for future research. First, spreading aggregate risk across cohorts implies that aggregate consumption is smoothed over time, which requires the capital stock to increase in good time and to decrease in bad time. Hence, intercohort risk sharing has implications for the cyclicality of aggregate investment. Second, as Gollier (2008) theoretically shows, intermediaries can invest in more risky assets when risk is shared across cohorts. Therefore, intercohort risk sharing has implications for the composition of aggregate investment.

References

ACPR. 2016. Les chiffres du marché français de la banque et de l’assurance.

Allen, F., and D. Gale. 1997. Financial markets, intermediaries, and intertemporal smoothing. *Journal of Political Economy* 105:523–46.

Baddou, S., C.-H. Huynh, F. Malgras, D. Marionnet, S. Slama, and Others. 2016. La situation des assureurs en France au regard des premières remises Solvabilité II en 2016.

Ball, L., and N. G. Mankiw. 2007. Intergenerational risk sharing in the spirit of arrow, debreu, and rawls, with applications to social security design. *Journal of Political Economy* 115:523–47.

Becker, B., and V. Ivashina. 2015. Reaching for yield in the bond market. *Journal of finance* 70:1863–902.

Bianchi, M. 2018. Financial literacy and portfolio dynamics. *Journal of Finance* 73:831–59.

Calvet, L. E., C. Célérier, P. Sodini, and B. Vallee. forthcoming. Can Security Design Foster Household Risk-Taking? *Journal of Finance*.

Célérier, C., and B. Vallée. 2017. Catering to investors through security design: Headline rate and complexity. *Quarterly Journal of Economics* 132:1469–508.

Darmon, C., and H. Pagenelle. 2005. Patrimoine: quand les ménages prennent de l'assurance.

Darpeix, P.-E. 2016. Le taux technique en assurance vie (code des assurances). *Analyses et Synthèses (Banque de France).*

Drechsler, I., A. Savov, and P. Schnabl. 2017. The deposits channel of monetary policy. *Quarterly Journal of Economics* 132:1819–76.

Ellul, A., C. Jotikasthira, A. V. Kartasheva, C. T. Lundblad, and W. Wagner. 2021. Insurers as asset managers and systemic risk. *Kelley School of Business Research Paper* 18.

Ellul, A., C. Jotikasthira, C. T. Lundblad, and Y. Wang. 2015. Is historical cost accounting a panacea? Market stress, incentive distortions, and gains trading. *Journal of Finance* 70:2489–538.

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

Ge, S. forthcoming. How Do Financial Constraints Affect Product Pricing? Evidence from Weather and Life Insurance Premiums. *Journal of Finance*.

Ge, S., and M. S. Weisbach. 2021. The role of financial conditions in portfolio choices: The case of insurers. *Journal of Financial Economics* 142:803–30.

Gollier, C. 2008. Intergenerational risk-sharing and risk-taking of a pension fund. *Journal of Public Economics* 92:1463–85.

Gordon, R. H., and H. R. Varian. 1988. Intergenerational risk sharing. *Journal of Public Economics* 37:185–202.

Greenwood, R. M., and A. Vissing-Jorgensen. 2019. The Impact of Pensions and Insurance on Global Yield Curves. *Harvard Business School Finance Working Paper No. 18-109*.

Handel, B. R. 2013. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103:2643–82.

Hirshleifer, J. 1971. The private and social value of information and the reward to inventive activity. *American Economic Review* 61:561–74.

Hortaçsu, A., and C. Syverson. 2004. Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds. *Quarterly journal of economics* 119:403–56.

Insee. 2016. Les revenus et le patrimoine des ménages. *Insee Références* édition 2016.

Koijen, R., and M. Yogo. 2015. The cost of financial frictions for life insurers. *American Economic Review* 105:445–75.

Koijen, R. S., and M. Yogo. 2022. The fragility of market risk insurance. *The Journal of Finance* 77:815–62.

Lusardi, A., and O. S. Mitchell. 2014. The economic importance of financial literacy: Theory and evidence. *Journal of economic literature* 52:5–44.

Novy-Marx, R., and J. Rauh. 2011. Public pension promises: how big are they and what are they worth? *Journal of Finance* 66:1211–49.

———. 2014. Linking benefits to investment performance in US public pension systems. *Journal of Public Economics* 116:47–61.

Scharfstein, D. S. 2018. Presidential address: Pension policy and the financial system. *Journal of Finance* 73:1463–512.

Sen, I., and D. Humphry. 2020. Capital regulation and product market outcomes.

Stiglitz, J. 1983. On the relevance of public financial policy: indexation, price rigidities and optimal monetary policies.