ISSN 0956-8549

Regulator Model-Implied Beliefs

By

Manav Chaudhary

FINANCIAL MARKETS GROUP DISCUSSION PAPER NO. 960

June 2026

Regulator Model-Implied Beliefs*

Manav Chaudhary†

October 13, 2024

Click here for the latest draft

Abstract

Financial regulations rely on regulator-controlled models to generate probabilistic forecasts which determine firm constraints. I refer to these forecasts as *regulator model-implied beliefs*. Using the U.S. life insurance sector as a laboratory, I measure both regulator model-implied and insurer expectations. I show that the regulator model disagrees with insurers yet quantitatively mirrors the systematic belief patterns observed in human forecasters, embedding these patterns into regulatory constraints. Insurers, in turn, pass through these belief dynamics into their portfolio decisions, effectively coordinating behavior at the sector level. This mechanism reveals a new channel linking well-documented belief dynamics to financial intermediary decisions.

*I am grateful to my advisors Ralph Koijen, Stefan Nagel, Niels Gormsen, and Lars Hansen, for formative discussions and continuous guidance. I am also very grateful to Francesca Bastianello, Aditya Chaudhry, Douglas Diamond, Zhiyu Fu, Xavier Gabaix, John Heaton, Joseph Kupferberg, Jingoo Kwon, Jane Li, Yueran Ma, Federico Mainardi, Ben Marrow, Pascal Noel, Simon Oh, Amir Sufi, Quentin Vandeweyer, Constantine Yannelis, Anthony Zhang, Eric Zwick, and participants of Boston College, Colorado Finance Summit, Cornell, HBS Finance, HBS BGIE, Imperial College, LBS, LSE, OSU, University of Chicago, UC Boulder, USC, and University of Toronto. I thank staff at regulators and industry practitioners for many helpful conversations. All errors are my own. A previous version of this paper was under the title "Regulator Beliefs."

†London School of Economics, Department of Finance; email: m.chaudhary6@lse.ac.uk.

1 Introduction

Modern financial regulations increasingly rely on models to quantify future risks. These models generate probabilistic forecasts that are then embedded into regulatory requirements. For example, bank risk weights depend on model-estimated default probabilities, while insurer liabilities are based on the model-simulated distribution of future losses. Because regulators prescribe or approve these models, entire industries are governed by the same outlooks (or a narrow range of approved ones), binding firms to a common view of the future. I refer to these regulator model-generated probability distributions, which directly determine firm constraints, as **regulator model-implied beliefs**.[1]

These beliefs are everywhere, yet remain poorly understood. What do they look like? How do they differ from those of regulated firms? And how much do they shape firm behavior?

The U.S. life insurance sector provides a unique laboratory for answering these questions. A single regulator prescribed interest rate model determines statutory requirements. Because the model is public, we can directly recover its implied probability distribution and compare them to the yield assumptions insurers disclose in investor calls. In addition, regulatory filings provide detailed data on portfolios, letting us trace how these different outlooks shape firm actions. The sector is not only economically large—managing over $10 trillion in liabilities and serving as the biggest intermediary of corporate bonds—but is also representative of model-based regulation in banking and pensions.[2]

The setting reveals three main findings. First, the regulator's model forecasts at least as well as professional surveys, but exhibits the same systematic patterns: persistent negative errors and underreaction to news. Second, the model's yield expectations often disagree with insurers' assumptions by more than 200 basis points at long horizons. Third, because the model-implied beliefs are directly embedded into regulation, the disagreement pressures insurers to act in ways they otherwise would not: they rebalance portfolios and adjust their own assumptions to align more closely with the regulator's model.

Together, the evidence shows that the dynamics of regulator model-implied beliefs drive the dynamics of intermediary balance sheets. This opens a new channel through which systematic belief patterns, long documented in surveys, become consequential for the economy. Once em-

[1]The models represent beliefs in the sense that they imply a coherent set of probabilities over future outcomes; these beliefs may, however, differ from officials' own subjective beliefs due to incentives or institutional frictions.

[2]e.g.,Banks: Basel III's market, credit and operational risk components; Insurer: U.S. valuation laws (AG43, VM20, VM21, VM22) and risk-based capital (C3 P1/P2), and EU regulations (Solvency II); Pensions: EU fund risk (IORP II)

bedded in regulation through models, they transmit directly into the actions of regulated firms. More broadly, because entire sectors are coordinated by the same set of regulator model-implied beliefs, they deserve a central place in the study of capital allocation, asset prices, and systemic risk.

I begin by measuring the model-implied beliefs that determine insurer constraints. Life insurers calculate their risk based capital and policy reserves from the tail of simulated losses. Traditional life reserves, for example, use the average of the worst 30% of simulated outcomes. To prevent firms from gaming these simulations, the regulator prescribes its own model economic model. I extract the model's implied expectations of future yields and refer to them as regulator model-implied expectations. Because these simulations extend over decades, outcomes are largely driven by the drift of the interest rate process, making expectations the key determinant of regulatory constraints.[3]

In the analysis below I takes these expectations as given, since they shape insurer constraints regardless of how they arise. Even so, it is useful to understand the regulator's model design philosophy. The regulator's stated objective in developing the model was to replicate the observed distribution of yields rather than to embed additional prudence through pessimistic assumptions. Consistent with this objective, the model's forecasts perform as well as those of professional surveys, indicating that its probabilities are not systematically tilted toward adverse outcomes. In practice, prudence in regulation arises from the use of tail-loss measures in capital and reserve requirements rather than from distortions in the model's probabilities. Some aspects of the model, however, reflect institutional frictions—for instance, a built-in parameter update rule that avoids the need for legislative discretion when parameters are changed—though these features do not appear to impair forecast accuracy.

Insurers, meanwhile, form their own interest rate assumptions to guide planning, pricing, and allocation. These assumptions are stated in investor calls, providing a direct window into their expectations. Using recent advances in Natural Language Processing (NLP), I build a Retrieval Augmented Generator (RAG) that parses transcripts of publicly traded firms and extracts the stated numbers. Here too, I take these assumptions as given, though I show later that they are partly shaped by firms' incentives to align with the regulator.

Using these data, I begin by examining the systematic properties of the regulator's model-implied expectations and compare them to those of professional forecasters in the Blue Chip Financial Forecasts (BCFF). The model's forecasts are at least as accurate as those of profes-

[3]See Appendix A for a discussion and calibration exercise quantifying the relative impact of the drift.

sional surveys in predicting future yields, yet they exhibit the same well-documented patterns as professionals: persistent negative forecast errors and underreaction in forecast revisions. Quantitatively, these patterns are nearly identical. In effect, the model reproduces the same belief dynamics observed in human forecasters. But with one crucial difference: the model's dynamics are embedded in the evolution of regulatory constraints.

Next, I compare the model's expectations with those of the insurers it constrains. Figure 1 plots regulator model-implied and average insurer expectations for ten-year yields against realized outcomes. The model typically expects yields to stay lower for longer than insurers. This difference does not reflect disagreement about the long-run mean—on that, they largely agree—but stems from the model's assumption of greater persistence in the interest rate process. The model implies a half-life of more than ten years, compared with three to seven years insurers typically assume. In the low-rate environment, this persistence assumption produced disagreements as large as 200 basis points for ten-year-ahead yields. More recently, as rates have risen, the same mechanism has generated disagreement in the opposite direction, with the model now expecting yields to stay higher for longer.

Figure 1: 10yr yield expectations: regulator $\mathbb{E}_t^R[y_{10,t+h}]$ vs average insurer $\mathbb{E}_t^I[y_{10,t+h}]$

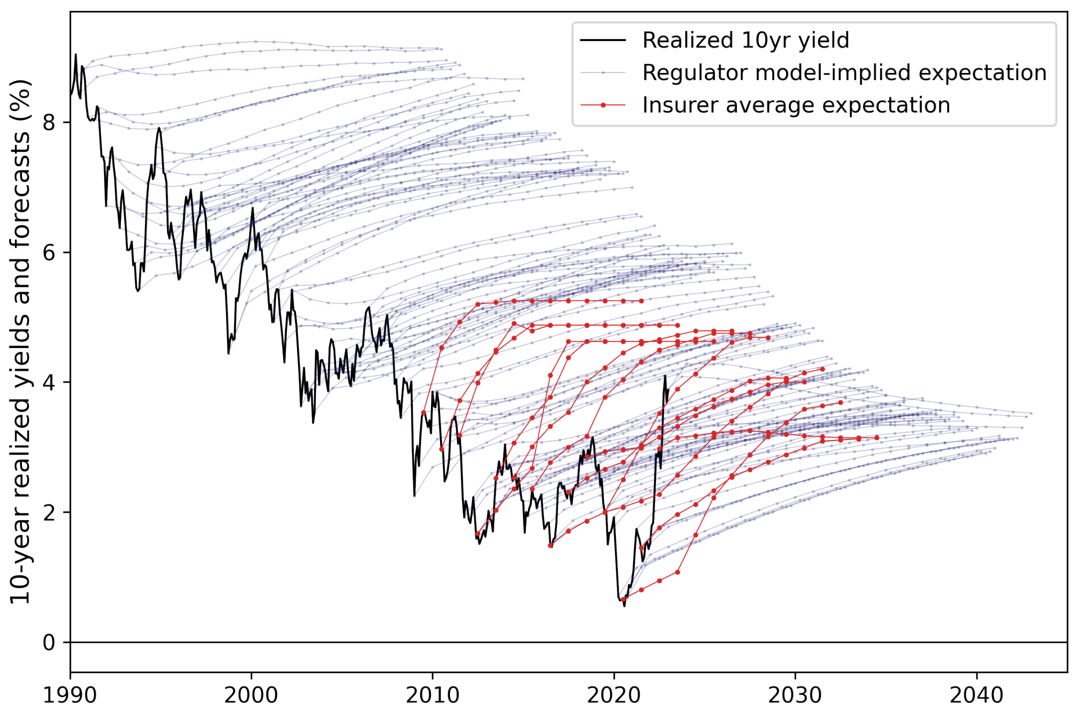

The figure plots the realized path of 10-year yields against vintages of regulator model-implied expectations $\mathbb{E}_t^R[y_{10,t+h}]$, and insurer's average expectations $\mathbb{E}_t^I[y_{10,t+h}]$. Due to data availability limitations insurer expectation are only plotted starting 2009.

Because these model-implied expectations are embedded in regulatory constraints, insurers cannot simply ignore them. The large disagreement pressures firms to act in ways they otherwise would not. Next, I therefore look at how the regulator model-implied expectations impact firm behavior, focusing on two dimensions: bond allocations and reported interest rate assumptions.

On the allocation side, disagreement with the regulator creates a trade-off. Insurers can raise additional capital and continue investing based on their own views, insulating their portfolios from the model-implied beliefs they disagree with. Or, if capital constrained, they can satisfy regulatory constraints by aligning their portfolios with the model—passing the model-implied beliefs through to their allocations.[4] If the latter effect is strong, then the systematic patterns in model-implied beliefs will be relevant for broader phenomena like overall capital allocation and asset prices.[5] Hence, I next assess the portfolio passthrough using two quasi-experiments.

The first experiment exploits differences in insurers' exposure to model-based regulation. Traditional life and fixed annuity policies carry similar interest rate risk and, before 2017, followed similar reserving rules. The introduction of Valuation Manual 20 (VM20) in 2017 changed this: traditional life reserves had to be calculated using simulated losses based on the regulator's model, while fixed annuities remained exempt. Firms with a larger share of traditional life business thus became more exposed to model-implied beliefs. A difference-in-differences analysis shows that these firms increased their holdings of long-term bonds, with portfolio duration rising from 6.7 to 7.8 years (a one standard deviation shift) consistent with rebalancing toward assets favored under the model's lower-for-longer outlook.

The second experiment exploits arbitrary variation in the regulator's model-implied expectations that stems from an ad hoc rule for updating the model's long-run mean. Each year, the regulator drops old yield observations from the input sample and rounds the parameter to the nearest 25 basis points. This procedure generates discrete, quasi-random shifts in the model's outlook that are unrelated to fundamentals. Using these shifts as an instrument, I find that insurers significantly rebalance their portfolios in response, adjusting portfolio duration in line with the direction of the model's belief changes. The results indicate that insurers respond mechanically to fluctuations in the regulatory model-implied expectations, even when those fluctuations are economically uninformative. Highlighting that temporal variation in regulator model-implied beliefs passes through to intermediaries' allocation, shaping the dynamics of their balance sheets.

[4]e.g., if the model forecasts rising yields, holding shorter-maturity bonds reduces projected losses compared to holding long-term bonds. Helping satisfy regulatory requirements without raising capital or hedging.

[5]For example, life insurer portfolio's duration is known to drive credit spread dynamics (Li, 2024).

Turning to insurers' reported expectations, I find similar evidence of passthrough. After the 2017 introduction of VM20, insurers significantly revised their interest rate assumptions, moving them closer to those implied by the model. Long-horizon disagreements fell by over 100 basis points. Because the model's design was public several years before VM20, this convergence likely reflects strategic alignment rather than learning. Given that management uses these assumptions to guide risk management, pricing, and planning, the response of reported expectations suggests that regulator belief dynamics influence multiple dimensions of insurer decision-making beyond just portfolio allocations.

Related literature: The findings broadly speak to the literatures on subjective beliefs and and risks in the insurance sectors.

This paper relates to the literature linking systematic belief patterns to asset prices, macro fluctuations, and financial stability (Barberis et al., 2015; Adam et al., 2017; Greenwood et al., 2019; Bordalo et al., 2021; d'Arienzo, 2020; Nagel and Xu, 2023; Krishnamurthy and Li, 2020; Maxted, 2024), by identifying a new mechanism through which these belief dynamics become economically consequential. A central challenge in this literature is the *belief passthrough puzzle*: survey expectations display systematic patterns, yet the agents surveyed (analysts and households) often do not make portfolio decisions, and those that do, the link between beliefs and actions is weak (Merkle and Weber, 2014; Giglio et al., 2021; Charles et al., 2024). In contrast, regulator model-implied beliefs naturally have strong passthrough. They exhibit the same systematic patterns as surveyed expectations but are embedded directly in financial constraints, creating a regulatory channel through which belief dynamics coordinate sector-wide capital allocation.

This paper also contributes to the subjective expectations literature by measuring regulator model-implied and insurer interest rate expectations. Prior work has largely focused on households and professional forecasters (see Adam and Nagel, 2023 for a comprehensive review). A growing strand now studies institutional investors, whose expectations are naturally consequential given the scale of assets they manage (Andonov and Rauh, 2022; Dahlquist and Ibert, 2024; Couts et al., 2024). I extend this work by introducing two new sets of interest rate expectations: (i) those embedded in regulator models that define the constraints faced by institutional investors, and (ii) those disclosed by life insurers in investor communications. Relevant to this strand of the literature, I show that insurer assumptions respond to regulations. This suggests that institutions' assumptions should not be interpreted as purely reflecting subjective expectations, but rather as outcomes of broader optimization decisions shaped by strategic considerations.

This paper relates to the literature on insurance and financial regulation by identifying a new driver of insurer behavior and highlighting the limits of model-based regulation. As life insurers have become increasingly complex and systemically important (Koijen and Yogo, 2016; Acharya et al., 2017; Koijen and Yogo, 2017; Ellul et al., 2022; Koijen and Yogo, 2022), regulators have turned to model-based frameworks. To prevent firms from gaming these models (Behn et al., 2022; Plosser and Santos, 2014; Sen and Sharma, 2020), regulators now prescribe their own models. Sen (2023) shows that when such regulations incorporate new risks, insurers respond by hedging them. I complement this finding by showing a speculation channel, distinct from the hedging channel: insurers reallocate portfolios to align with the model's projections—reducing losses when the view proves correct, but amplifying them when it does not.[6] Moreover, because statutory liabilities are now re-evaluated each period using market sensitive models, insurer balance sheets are becoming less historical cost based and more market sensitive. This shift potentially weakens the asset insulator effect documented by Chodorow-Reich et al. (2021)—reducing the sector's role in dampening systemic risk.[7]

Outline: Section 2 describes the measurement of regulator and insurer interest rate expectations. Section 3 documents the behavior of these beliefs, focusing on the disagreement between regulators and insurers and systematic patterns in their expectations. Section 4 examines the impact of regulator beliefs on insurer bond allocation decisions, and explores the passthrough of regulator model-implied expectations on insurer assumptions. Finally, Section 5 concludes and discusses the broader implications of these findings.

2 Measurement and Data

This section describes how I measure expectations relevant to insurers' regulatory constraints and operating decisions. First, I use the interest rate model prescribed by regulators to construct model-implied expectations. Second, I collect insurers' assumptions about future interest rates—assumptions that insurers say they use for internal risk management and portfolio allocation decisions—which they disclose in conference calls.

[6] The finding relates more broadly to work on how regulatory requirements and financial constraints shape insurer behavior (Becker and Ivashina, 2015; Koijen and Yogo, 2015; Ellul et al., 2015; Ge and Weisbach, 2021; Egan et al., 2022; Becker et al., 2022; Tang, 2022; Ge, 2022).

[7] More generally, coordinated insurer rebalancing can generate persistent price impact on corporate bonds (Ellul et al., 2011; Nanda et al., 2019; Murray and Nikolova, 2022; Chaudhary et al., 2023).

2.1 Regulator model and measuring model-implied beliefs

U.S. life insurance companies are regulated at the state level, with oversight harmonized at the federal level by the National Association of Insurance Commissioners (NAIC). Catalyzed by the financial crisis, the NAIC has shifted toward forward-looking and risk-based regulation. In particular, they now rely on model-based approaches to ensure that insurers manage their asset-liability positions in a way that is resilient to interest rate fluctuations. Under this framework, insurers calculate the expected shortfall of their assets and liabilities. The results of which then determines statutory risk-based capital and policy reserve requirements.

The expected short-fall calculation proceeds in two steps. First, based on assumptions about how interest rates will evolve, insurers estimate a multi-year loss distribution of assets and liabilities. Second, using this simulated distribution, they then calculate the expected shortfall i.e., the average of the worst tail-losses. The regulator sets the tail percentile for the amount of risk they are willing to allow insurers to have on their balance sheets (e.g., 30% for statutory reserves, 5% for risk-based capital).

For the first step, earlier regulations allowed insurers to make their own assumptions about future interest rates and how they depend on current yields. However, to ensure these assumptions accurately reflected interest rate dynamics, starting in 2017 the regulator began prescribing a pre-calibrated model of yields. As a result, all insurers now face regulatory constraints based on the same projected path of interest rates. I refer to these projections as the regulator model-implied beliefs. Model-implied beliefs of this kind are widely used across financial regulation and are a common feature of how regulatory constraints are set, including in banking and pension regulation.[8] In this insurance setting, however, the prescribed model is publicly available, allowing us to directly measure these model-implied beliefs and study their effects on insurer behavior.

2.1.1 Model

The regulator models long term interest rates as a mean reverting process with stochastic volatility. The model has three endogenous variables: the log 20-year nominal yield ($\ln y_{20,t}$), the slope of the yield curve (measured as the spread between the nominal 20-year and 1-year yields, $s_t := y_{20,t} - y_{1,t}$), and log volatility of the log long-rate ($\ln \sigma_{20,t}$). Specifically, these

[8]In some regulatory settings, rather than prescribing a single model, regulators control model quality by approving firms' internal models. This leads to a similar convergence of models, but toward the set of admissible ones that minimize regulatory requirements (Behn et al., 2022).

variables follow the mean-reverting processes described by the monthly frequency equations,[9]

$$\begin{aligned} \Delta \ln y_{20,t} &= \beta_1(\ln \bar{y}_{20}^{(h)} - \ln y_{20,t-1}) + \psi(\bar{s} - s_{t-1}) + \sigma_{20,t}u_{1t} && \text{(log long rate)} \\ \Delta s_t &= \beta_2(\bar{s} - s_{t-1}) + (y_{20,t-1})\sigma_2 u_{2t} && \text{(slope)} \\ \Delta \ln \sigma_{20,t} &= \beta_3(\ln \bar{\sigma} - \ln \sigma_{20,t-1}) + \sigma_3 u_{3t} && \text{(log volatility)} \end{aligned}$$

where the shocks $u_t := [u_{1t}, u_{2t}, u_{3t}]'$ are jointly normally distributed and correlated,

$$u_t = \Sigma^{1/2}\varepsilon_t \quad \text{s.t.} \quad \Sigma = \begin{pmatrix} 1 & \rho_{12} & 0 \\ \rho_{12} & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad \text{and} \quad \varepsilon_t \sim N(0, I)$$

Additionally, the 20-year yield and 1-year yields are bounded between $y_{20,t} \in [1.15\%, 18\%]$ and $y_{1,t} \in [1\%, 40\%]$, respectively.[10]

For any simulations made in year h, insurers have to use the 20-year yield mean reverting point $\bar{y}_{20}^{(h)}$. To capture secular changes in nominal rates, at the start of a given year h, the parameter is calculated using a linear combination of various backward looking medians and means, which is then rounded to the nearest 25 basis points,

$$\bar{y}_{20}^{(h)} = \textbf{Round}_{25\text{bps}}\Big\{0.2 \times \text{Med}_{600m,h}(y_{20}) + 0.3 \times \text{Mean}_{120m,h}(y_{20}) + 0.5 \times \text{Mean}_{36m,h}(y_{20})\Big\}$$

where $\text{Med}_{600m,h}(y_{20}) := \text{Median}\{y_{h-600}, \ldots, y_{h-1}\}$ is a 600-month moving median, $\text{Mean}_{120,h}(y_{20}) := \frac{1}{120}\sum_{l=1}^{120} y_{20,h-l}$ is a 120-month moving average, and $\text{Mean}_{36m,h}(y_{20}) := \frac{1}{36}\sum_{l=1}^{36} y_{20,h-l}$ is a 36-month moving average.

The model parameters were calibrated in December 2008. Table 1 presents these parameters.[11] Long-term rates converge particularly slowly, with an implied half-life of around 11.3 years. Therefore, the transition path from current yields to the long-term mean critically drives long-horizon projected losses.

To generate interest rate projections: current yields (the model's state variables) are initialized using the prevailing 20-year and 1-year interest rates, and stochastic volatility is initialized using its long-run mean calibration. The model then draws shocks $\varepsilon_t \sim N(0, I)$ and uses these

[9]The slope equation has an additional mean-reverting term on long-yields, but its contribution is effectively zero, so is omitted here for brevity.

[10]The short-rate 1% floor is a "soft bound"; if it goes below 1% $y_{1,t}$ is set to $0.25y_{20,t}$. Hence, in theory $y_{1,t}$ can go as low as 0.2875%.

[11]ψ is an error correction rate and hence the implied half-life is omitted.

Table 1: Regulator model parameters

	Reversion/Correction Rates				Long-Run Means		Shock Covariances		
	β_1	β_2	β_3	ψ	$\bar{s}$	$\bar{\sigma}_{20,t}$	σ_2	σ_3	ρ_{12}
Parameter (3 d.p)	0.005	0.027	0.040	0.252	0.010	0.029	0.041	0.115	-0.192
Half-life (years)	11.3	2.2	1.4	—	—	—	—	—	—

with the initialized values to project the model forward. Using the resulting projected path of 20-year and 1-year yields, the model then calculates other yields by fitting a Nelson-Siegal curve.[12] This process then is repeated many times to generate a distribution of interest rate projections, which are then used to determine the loss distributions.

2.1.2 Measuring model-implied expectations

Using the regulator's model, we can in principle construct any moment of the regulator model-implied belief distribution. I focus on expectations, as the multi-year horizon of the expected shortfall calculations for regulatory requirements makes them the primary driver of insurers' constraints (see Appendix A for a formal quantification). Over such horizons, the drift/expectation of the loss process dominates other moments—much like how, for equities, the equity premium drives long-run returns while volatility matters mainly over short horizons.

To calculate the model-implied expectations, the bounds in the model prevent a closed-form solution, so the expectations must be evaluated numerically. For any given period, this is done by taking the relevant initialization conditions, simulating a large number of interest rate paths, and averaging the outcomes. Figure 1 plots vintages of the regulator's model-implied expected path of 10-year yields, $\mathbb{E}_t^R[y_{10,t+h}]$, alongside the realized path of 10-year yields, $y_{10,t}$.

2.1.3 Interpreting regulator model outputs

The regulator's model represents a coherent set of beliefs in the sense that it implies a probability distribution over future outcomes. These beliefs may not align with the regulator's own subjective views, whether due to regulatory incentives to emphasize certain adverse outcomes or institutional frictions that limit model updating. Ultimately, however, they are the beliefs that determine the constraints faced by all insurers. Hence, understanding their properties is required to understand how insurer constraints evolve and drive their decision-making.

[12]The model uses the yield curve function, $y_{n,t+h} = \beta_{0,t+h} + \beta_{1,t+h}\frac{\lambda}{n}\left(1 - e^{-n/\lambda}\right)$, where $\lambda = 5/2$ is the preset shape parameter. The $\beta_{0,t+h}$ and $\beta_{1,t+h}$ coefficients are determined using the simulated $y_{1,t+h}$ and $y_{20,t+h}$.

My main analysis does not hinge on why these beliefs take their particular form—it takes them as given and studies their properties and implications for firm actions. Nonetheless, understanding the regulator's model design philosophy, and whether it reflects motivations other than correctly forecasting interest rates, is of interest in its own right.

The regulator states that the model is designed to match the observed distribution of interest rates (P-measure distribution), rather than to distort probabilities for pricing or prudential purposes. As outlined in their December 2008 report:

> *"The scenarios created by the SLV [(stochastic local volatility)] model are real world, P-measure, scenarios...it does not produce scenarios that are arbitrage free in respect of the starting yield curve. As such, the SLV scenarios are not appropriate for market pricing of financial instruments (assets or liabilities); they are intended for real world cashflow projections."*

Empirical properties of the model-implied beliefs are consistent with this stated design goal (see Section 3). First, the model's forecast errors are similar to or smaller than those of professional forecasters, suggesting it is not systematically skewed toward adverse scenarios and that there is limited scope for the regulator to materially improve forecast accuracy. Second, the model's long-run mean is in line with that of other market participants, rather than being set unusually low to build in conservatism. Third, while the model tended to expect "lower for longer" yields than other agents prior to 2022, it has since shifted to expecting "higher for longer" yields—implying a less prudent view than other agents in the current higher yield environment.[13] Overall, the evidence suggests that prudence in the regulation is expressed through the choice of tail percentile in expected shortfall calculations, rather than by distorting the model's underlying probabilities.

While a prudence motive does not appear to drive the model-implied beliefs, institutional frictions may still influence their form. For example, the consensus-based decision-making process within the NAIC, combined with the decentralized nature of state-level legislation, makes it difficult to implement discretionary parameter revisions. These fictions help explain why the long-run mean is updated using a mechanical rule, and why there have not been changes to other parameters since 2008. With that said, the lack of discretionary changes to the model does not seem to have impacted its accuracy, which is comparable to that of professionals. This may be because estimating the parameters of an interest rate process is challenging even with multiple decades of data (Farmer et al., 2021), so the scope for substantial revisions based on a few additional years of data is limited. Moreover, the parameter that has varied most over

[13]Higher for longer view is less prudent because insurers' assets tend to have a shorter duration than liabilities, and hence, in a higher rate environment they are more likely to have a surplus.

time—the long-run mean—is updated every year using a rule that approximates the Kalman filter a Bayesian learner would apply.[14]

2.2 Insurer interest rate assumptions

Insurance companies make assumptions about the path of interest rates to assess risks to their business and to guide their policy issuance and portfolio allocation decisions. Standard industry practice is to assume that interest rates follow the forward curve for three years, and then converge linearly to a long-run mean over a specified convergence period. Firms review and update their assumptions about the long-run mean and convergence period on an annual basis.

Given the importance of these assumptions for insurers' operations, they are frequently discussed and scrutinized in calls between firm management and equity analysts.[15] To comply with disclosure regulations, publicly traded insurers publish transcripts of these meetings.

I use conference call transcript data from FactSet, covering 3,094 transcripts from 2002 to 2023 (see Table 3 in Appendix B.1)—coverage is less comprehensive in the early years but becomes more complete over time. The key challenge in using these transcripts is the high volume of text and the unstructured way in which assumptions are often described. Making extracting insurer assumptions a needle-in-a-haystack problem. To address this, I use recent advances in large language models (LLMs) to build a retrieval-augmented generator (RAG) that parses the transcripts and produces a firm-year dataset of insurers' interest rate assumptions.[16]

2.2.1 Retrieval Augmented Generator

A RAG combines a retrieval system with a generative LLM. In this setting, the retrieval system searches through transcripts to identify passages that are likely to contain interest rate assumptions, while the LLM then interprets these passages and extracts the relevant numbers. For each firm-year, I first subset the transcripts to statements made by company management, then use embedding-based semantic search to retrieve the most relevant passages (e.g., discussions of long-term rate assumptions or convergence horizons). These retrieved excerpts are then passed to the LLM, which returns the identified assumptions together with the supporting text. This

[14]Appendix A.3 shows that the update rule closely matches the form of a Kalman filter for learning the long-run mean of a standard interest rate process.

[15]These assumptions have also been subject to legal scrutiny, with lawsuits brought by shareholders and policyholders alleging inappropriate interest rate assumptions. For example, the MetLife class action lawsuit in 2012, the AXA Equitable Life Insurance Co. litigation in 2016, and the Transamerica Life Insurance Co. lawsuit in 2018.

[16]Similarly, Gormsen and Huber (2023) use conference call transcripts to extract firms' discount rates. The RAG methodology here can be viewed as an automated version of the data collection procedure in their paper.

design reduces the chance of the model hallucinating and ensures that each extracted value can be directly verified against the transcript. For more technical details on the retrieval system, see Appendix B.2.

As an illustration, the excerpt below shows the RAG applied to MetLife's 2019 transcripts, where it extracts the company's long-run mean assumption (3.75%) and convergence horizon (8 years) together with the supporting quote. The first box shows the prompt sent to the LLM, which contains the transcript excerpts retrieved by the semantic search system as likely to include insurer assumptions. The second box shows the LLM's output: the extracted assumptions along with the supporting text.

MetLife 2019 example prompt

Answer the question based only on the following context:

- (10/31/2019) The most significant driver was the reduction of our long-term U.S. 10-year treasury interest rate assumption from 4.25% to 3.75% with the rate mean reverting over the next 8 years.
- ...

Based on the above context, answer the following question:

- What is the long-term U.S. Treasury interest rate assumption? Give answer as %.
- Over what time period are rates expected to revert this assumption? Give answer in years or date e.g. 10 years or end of 2016.

Output formatting instructions: ...

MetLife 2019 example response

Long-term mean assumption: 3.75%
Mean-reversion time assumption: 8 years
Relevant statement: "(10/31/2019) The most significant driver was the reduction of our long-term U.S. 10-year treasury interest rate assumption from 4.25% to 3.75% with the rate mean reverting over the next 8 years."

Using this approach, I construct an unbalanced panel of 12 publicly traded insurers from

2009-2022, yielding 175 collected assumptions. I then manually verify their validity comparing the extracted assumptions with the associated statement in the transcript.

Following industry convention, I use these to build insurer assumption-implied expectation paths $\mathbb{E}_t^I[y_{10,t+h}]$: the forward curve is assumed for the first three years, after which yields converge linearly to the long-run mean at the pace implied by the convergence assumption. Figure 1 plots average insurer expectations alongside regulator model-implied expectations and the realized path of 10-year yields.

2.2.2 Interpreting insurer assumption-implied beliefs

The collected assumptions add to a growing body of work that seeks to measure institutional investors' expectations directly, such as asset managers' equity return assumptions (Dahlquist and Ibert, 2024; Couts et al., 2024) and public pension funds' return expectations (Andonov and Rauh, 2022). As is the case with such financial market assumptions, however, interpretation and coverage require care.

First, while the extracted assumptions provide a unique window into insurers' views, they should not necessarily be interpreted as firms' unbiased beliefs. Despite scrutiny from equity analysts, investors, and policyholders, management may have strategic reasons to deviate from their true expectations. For example, because managers benefit from upside gains but do not fully bear the downside costs of failure, they may have incentives to emphasize higher interest rate assumptions, which imply stronger profitability—I find evidence suggestive of such upward tilting in the stated assumptions (see Section 4.4).

Second, there are limitations in coverage. The data are drawn from publicly traded insurers, since transcripts are only available for firms subject to SEC disclosure rules. A large share of the U.S. life insurance market, however, consists of life insurers organized as mutuals, for which no comparable data are available. I view this paper as a first step in collecting and analyzing insurer interest rate assumptions, leaving room for future research to expand coverage to a broader set of institutions.

3 Regulator and insurer expectations properties

Section 3.1 studies systematic patterns in regulator model-implied and insurer assumption-implied expectations, and them with each other and patterns observed in the expectations of professional forecasters. Section 3.2 shows systematic disagreement between the regulator and insurer, and discusses how these differences are generated by disagreement about the persistence of the interest rate process.

3.1 Systematic patterns in expectations

I compare regulator model-implied and insurer assumption-implied expectations to those of professional forecasters. Professional forecasts provide a natural benchmark in two ways. First, they help assess whether insurers or regulators perform systematically worse, which could reflect strategic incentives or limited forecasting ability. Second, the large subjective-expectations asset pricing literature has documented systematic patterns in professional forecasts and argued that these may underlie anomalous bond yield dynamics. Comparing insurers and regulators to this benchmark allows me to see whether their expectations exhibit similar patterns. This matters because, unlike professionals who do not directly hold assets, these beliefs shape regulatory constraints and guide the business planning of insurers.

I focus on two patterns commonly documented in the literature on professional expectations. First, expectations tend to be too high on average, with realized yields being less than those forecasted, particularly at longer horizons. This provides a direct test of overall forecast accuracy and speaks to whether insurers or regulators perform systematically worse than professional forecasters. Second, expectations underreact to news, so forecast revisions are too muted. These belief dynamics are particularly relevant for asset pricing, as sluggish adjustment of expectations has been argued to underlie phenomena such as excess return predictability, overshooting of long-term yields, and excess volatility in bond returns (d'Arienzo, 2020).

Formally, let $\mathbb{E}_t^k[y_{n,t+h}]$ denote agent k's period t expectation for the realization of maturity n yield y in period $t+h$. I run the following regressions,

1. **Average expectation error**:

$$y_{10,t+h} - \mathbb{E}_t^k[y_{10,t+h}] = \alpha_h^k + \varepsilon_{t+h}^k$$

   where $H_{FIRE} : \alpha_h^k = 0$.

2. **Underreaction**: Coibion-Gorodnichenko (CG) regression,

$$y_{n,t+1Q} - \mathbb{E}_t^k[y_{n,t+1Q}] = \alpha_n^k + \beta_n^k(\mathbb{E}_t^k[y_{n,t+1Q}] - \mathbb{E}_{t-1Q}^k[y_{n,t+1Q}]) + \varepsilon_{n,t+1Q}^k$$

   where $H_{FIRE} : \alpha_n^k = 0$ and $\beta_n^k = 0$, and $\beta_n^k > 0$ implies underreaction (hence positive correlation in forecast errors).

To make these comparisons, I compile a quarterly panel from 1988 to 2021 covering regulator model-implied expectations, insurer assumption-implied expectations, and survey forecasts from Blue Chip Financial Forecasts (BCFF).[17] The survey data provide only short-horizon forecasts (up to one year ahead). The results are robust to restricting the sample to the post-2008 period—the period after the regulator model's calibration—and using longer lags in HAC standard errors (see Appendix D.1).

It is worth noting that, at the horizons studied in this subsection, insurer expectations are mechanically tied to the forward curve.[18] So their expectations simply reflect the path implied by the expectation hypothesis—implicitly forward rates are taken as reflecting only the path of expected rates, ignoring term premia in the forward curve—which is strongly rejected in the data.

Figure 2 plots the average expectation errors for 10-year yields across different expectation horizons, comparing the regulator model, insurer assumptions, and professionals. Each point represents the average forecast error for a specific horizon, with error bars indicating the 95% confidence intervals calculated using Newey-West errors with a 1-lag. On average, all expectation errors are negative, indicating that forecasters missed the secular decline in interest rates during the 1988 to 2021 sample period. This observed pattern in expectation errors aligns with the difficulty of accurately forecasting the slow-moving latent mean of interest rates (Farmer et al., 2021).

Notably, the regulator's model produces expectation errors that are similar to, or smaller than, those of professional forecasters, with the relative performance improving at longer horizons. This benchmark comparison highlights two points. First, the regulator's model is not systematically underperforming compared to sophisticated market participants, consistent with the model's objective to mimic the data-generating process of interest rates. Second, because even professionals exhibit systematic errors, it suggests it is difficult to avoid such forecast errors.

Figure 3 plots the CG coefficients measuring expectation underreaction to news across maturities. A positive coefficient indicates underreaction, as a positive correlation between expectation errors and expectation revisions implies that expectations failed to fully adjust to new information. The figure shows that insurer expectations—based on forward curves—do not ex-

[17] To ensure comparability, I measure insurer and regulator expectations using information available at the end of the first month of each quarter, which matches the timing of BCFF submissions. I also adjust the target to be the average yield over the relevant quarter at the forecast horizon, consistent with the survey definitions.

[18] The insurer expectation sample is too short to study the time-series properties of long-horizon expectations that depend on their internal interest rate assumptions.

Figure 2: 10-year yield average expectation error ($:= y_{10,t+h} - \mathbb{E}_t^k[y_{10,t+h}]$)

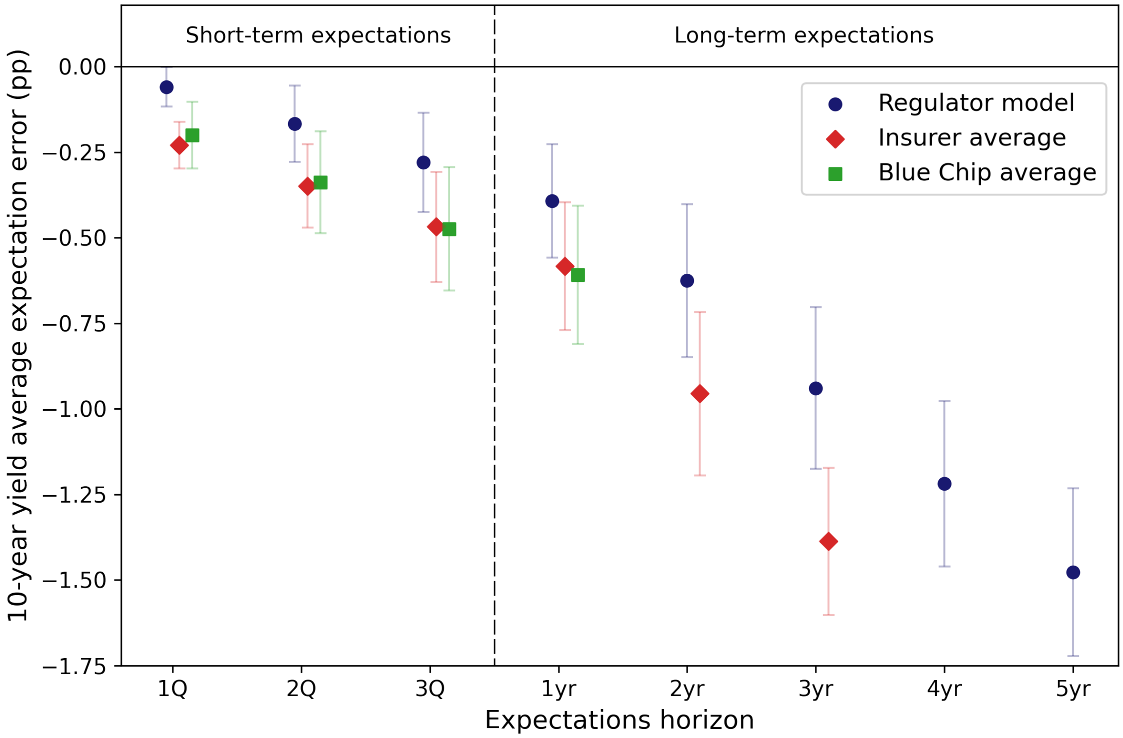

The figure plots the 10yr yield average expectation error for $h = 1Q, 2Q, 3Q, 1yr, 2yr, 3yr, 4yr, 5yr$ horizons for regulator model, average of insurer assumptions, and Blue Chip forecasters. The expectations are a panel of quarterly frequency expectations from 1988 to 2021. However, expectations for all horizons are not available for all four forecasters. The 95% confidence intervals are calculated using Newey-West errors with 1-lag. Results are robust to using a post-2008 sample and HAC standard errors with 4 lags (see Appendix D.1).

hibit significant under- or overreaction. By contrast, regulator expectations display pronounced underreaction, quantitatively similar to that observed for professional forecasters. Consistent with earlier work (Wang, 2021; d'Arienzo, 2020), the underreaction is strongest at short maturities and diminishes at longer horizons.[19]

Overall, the results highlight that the regulator's model quantitatively mirrors the expectation patterns observed in human forecasters. Crucially, however, unlike survey forecasts, these expectations directly determine insurer constraints. Moreover, the model's structure allows us to assess the sources of these patterns. Appendix D.1.1 shows that the persistent negative forecast errors stem primarily from the model's long-run mean parameters, while the underreaction result arises mainly from its persistence parameters.

[19]Subjective expectations theories suggest that such sluggish belief dynamics could help explain anomalies in bond markets, including excess return predictability, overshooting of long yields, and excess volatility (d'Arienzo, 2020).

Figure 3: Coibion-Gorodnichenko coefficient

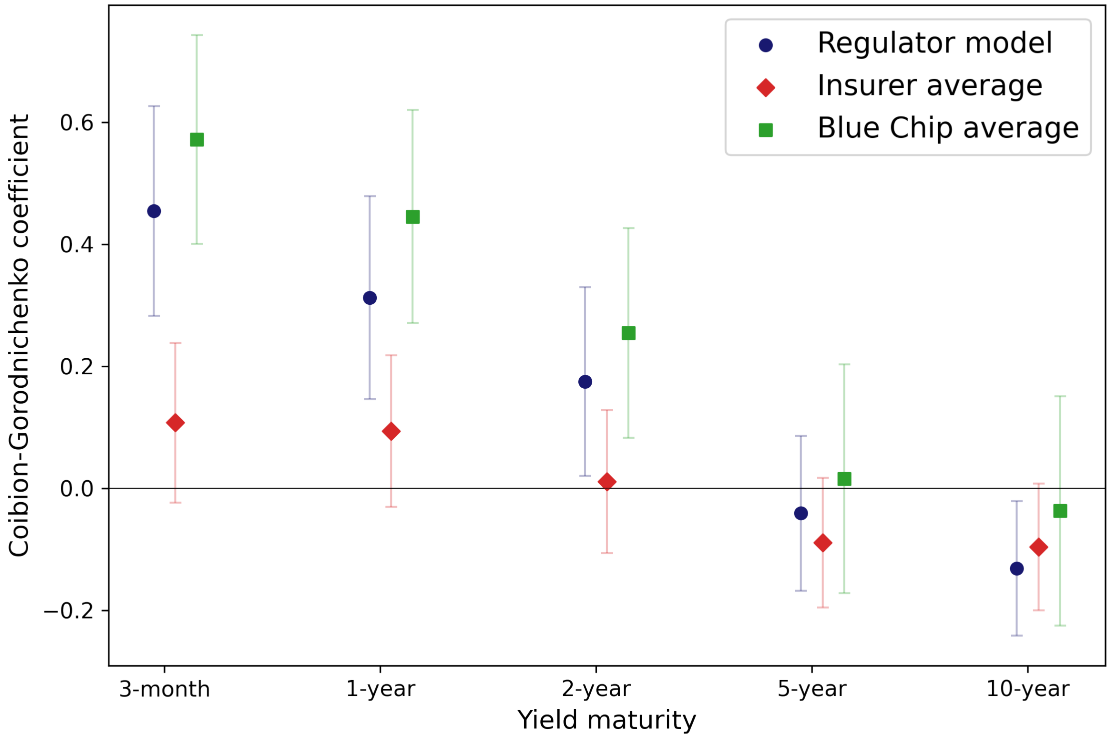

The figure plots the Coibion-Gorodnichenko (CG) coefficient for different maturities, comparing the regulator model, average insurer assumptions, and Blue Chip survey. The expectations are a panel of quarterly frequency expectations from 1988 to 2021. However, expectations for all horizons are not available for all four forecasters. The 95% confidence intervals are calculated using Newey-West errors with 1-lag. Results are robust to using a post-2008 sample and HAC standard errors with 4 lags (see Appendix D.1).

3.2 Disagreement

Figure 4 plots the difference between the regulator's model-implied and insurer's assumption-implied 10 year yield expectations for different horizons. The plots shows that there is significant disagreement between regulators and insurers, with the regulator model's paths being lower for longer than insurers assume. For example, for 6-year ahead expectations, the regulator model on average expected 10-year yields to be 75 basis points lower than the average insurer's assumptions. Therefore, under the regulator model's expectations long-term bonds are relatively more attractive investment than under the insurer assumptions.

To better understand the driver of this disagreement, figure 5 plots the regulator model's and insurers' assumptions about the long-run mean of yields (left) and the persistence of yields (right). Note that regulator model and insurers generally agree about the long-run mean of yields. Both have revised down their estimates of the long-run mean, while there are periods in the middle of the sample where some disagreement opens up, it closes by the end of the sample period. In contrast, they significantly disagree about the persistence of the interest rate process.

Figure 4: Average disagreement between the regulator and insurers ($\mathbb{E}_t^R[y_{10,t+h}] - \mathbb{E}_t^I[y_{10,t+h}]$)

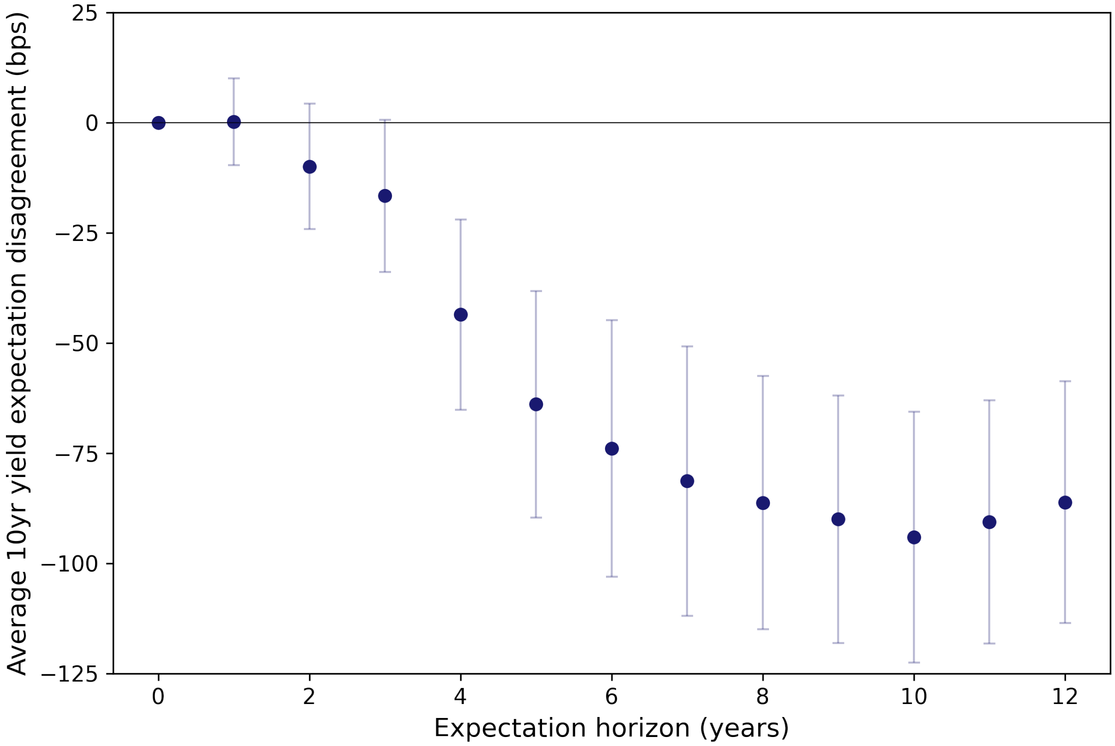

The figure shows the average difference between regulator and insurer expectations of 10yr yields for different expectation horizons, using annual data from 2009 to 2022, and 95% confidence intervals based on SE clustered at the insurer level.

While insurers have increased their persistence assumption from reversion taking around 5 years to 11 years. This is still substantially lower than the regulator model, which assumes that yields will take over 25 years to converge just 80%. Given that low interest rates prevailed during the majority of this sample period, the regulator model's slow reversion assumption naturally implied a low for long view on yields.

Since the disagreement in expectations stems from differing views on the persistence of the interest rate process rather than its mean, the direction of disagreement depends on prevailing interest rates. In a higher interest rate environment, the direction of disagreement will reverse, as the regulator model will assume that higher rates will persist for longer. Consistent with this prediction, in 2023 when yields rose above the long-run mean assumed by both insurers and the regulator, the disagreement was minor and flipped sign. Furthermore, since the sign of disagreement is state dependent, it suggests that the regulator is not actively distorting their model to emphasize scenarios where insurers are more likely to incur losses—in line with the regulator's objective of designing their model to mimic the observed behavior of yields.

Figure 5: Insurer long-term 10-year US Treasury yield assumptions

(a) Mean-reversion point assumption (b) Mean-reversion time assumption

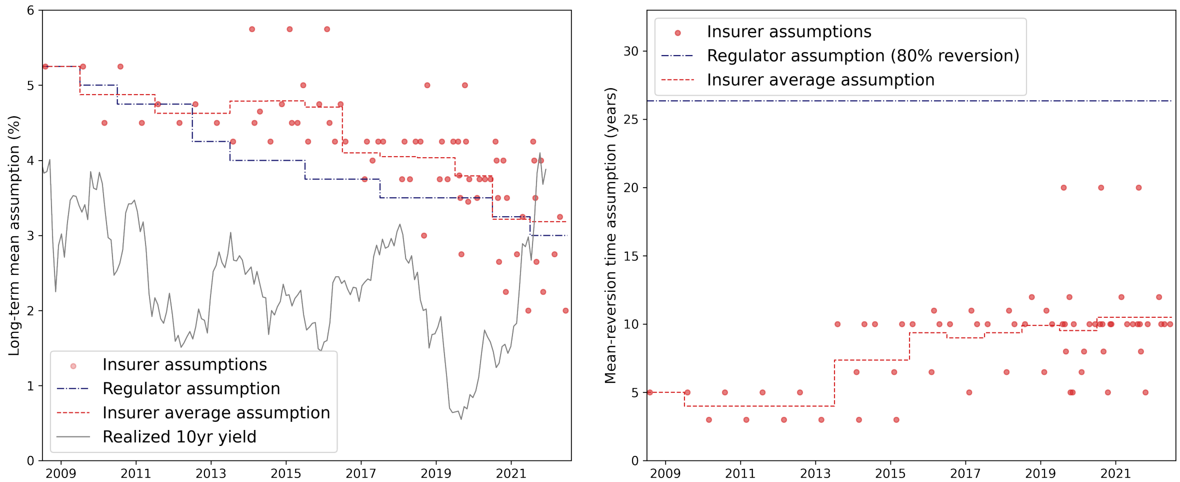

The figures plot insurers' long-term assumptions for the 10 year US Treasury yields: mean-reversion point (left), and the time-taken to reach the mean-reversion point (right). The red scatters correspond to insurer assumptions, the dashed red line are the average of insurer assumptions, the dashed blue lines are the regulator's corresponding model assumption, and the solid grey in the left plot is the realized path of 10-year Treasury yields. Since, the regulator model's interest rates as an AR processes, the complete reversion takes ∞ periods, hence, the right plots the 80% reversion time.

4 Model-implied expectations and insurer actions

In this section, I study the impact regulator model-implied expectations have on insurer behavior. I focus on their impact on insurer's bond allocation decisions, and also provide evidence of broader effects on insurer's own economic assumptions.

4.1 Motivating model

To illustrate the key trade-off faced by insurers, I present a stylized mean-variance bond allocation problem in which the insurer maximizes firm value based on both its own and the regulator's model-implied beliefs. Because firm value depends jointly on two distributions: future losses evaluated under the insurer's own beliefs and statutory reserves determined under the regulator model-implied beliefs. The insurer's optimal portfolio choice must therefore balance the two, allocating bonds in a way that reflects both private expectations and the model's expectation embedded in statutory constraints.

Balance sheet: In period t, firms compete in a monopolistic insurance market and when they sell Q_t policies at price P_t it generates: (i) revenue $A_t := P_tQ_t$, (ii) economic liabilities $L_t := Q_tV_t$ where V_t is the fair value per policy, and (iii) additional statutory liabilities $L_t^{statutory}$ that are evaluated using losses calculated based on regulator model-implied beliefs. The revenue from selling these policies is invested in short-term bonds M_t and long-term bonds B_t.

$$\underbrace{M_t + B_t}_{:=A_t=P_tQ_t} = Q_tV_t + L_t^{statutory} + K_t$$

where K_t is the period t capital of the insurance company.

Short-term bonds pay risk-free return $1 + r_t^f$, while the future value long-term bonds and liabilities is unknown and exposed to interest rate risk. I capture this risk, by modeling long-term bonds and liability return as a function of change in long-term yields $\Delta y_{l,t+1}$,

$$\frac{B_{t+1}}{B_t} = -D_B \Delta y_{l,t+1} \quad \text{and} \quad \frac{Q_tV_{t+1}}{Q_tV_t} = -D_L \Delta y_{l,t+1}$$

$\Delta y_{l,t+1}$ is not known known at time t, and D_B and D_L are the duration of long-term bonds and liabilities, respectively. The period $t+1$ value of capital (net of statutory reserves) is given by,

$$K_{t+1} = (K_t - L_t^{statutory})(1 + r_{t+1}^f) + B_t \underbrace{(-D_B\Delta y_{l,t+1} - r_{t+1}^f)}_{:=rx_{t+1}^B} - L_t \underbrace{(-D_L\Delta y_{l,t+1} - r_{t+1}^f)}_{:=rx_{t+1}^L}$$

where rx_{t+1}^B and rx_{t+1}^L are the excess returns on bonds and liabilities respectively.

Insurer and regulator model-implied beliefs: Assuming the insurer and the regulator's model agree about the variance of long-term yields is σ_y^2, but disagree about the mean, $\mathbb{E}^R[\Delta y_{l,t}] \neq \mathbb{E}^I[\Delta y_{l,t}]$, where R and I superscript denote the regulator and insurer respectively. As a result of this disagreement, the regulator's model and insurer also disagree about the expected excess return on bonds and liabilities.

Portfolio problem: When external financing is costly, insurers aim to hedge potential losses and minimize statutory reserve requirements to maximize firm value (Froot et al., 1994; Koijen and Yogo, 2015; Sen, 2023). The insurer evaluates potential losses using their own beliefs, while regulation requires statutory reserves to be calculated using potential losses based on regulator model-implied beliefs. Appendix C shows that, to a second-order approximation, the bond allocation decision of such an insurer can be represented by the following mean-variance problem:

$$\max_{B_t} \alpha_t \mathbb{E}_t^R[K_{t+1}] + (1-\alpha_t)\mathbb{E}_t^I[K_{t+1}] - \frac{\gamma}{2}\mathbb{V}_t[K_{t+1}]$$

where $\alpha_t \in [0,1]$ is the cost of regulation coming from statutory liabilities, and γ is the cost of variance. If the cost of regulation is higher, then it is more costly for the insurer to realize losses under the regulator model-implied beliefs compared to their own beliefs, and hence they put more weight on regulator model-implied beliefs in their bond allocation decisions.

<u>Proposition</u>: *Optimal holding of long-term bonds B_t is,*

$$B_t = \underbrace{\frac{\mathbb{E}_t^I[rx_{t+1}^B] + \overbrace{\alpha_t}^{\text{Cost of regulation}} \times \overbrace{(\mathbb{E}_t^R[rx_{t+1}^B] - \mathbb{E}_t^I[rx_{t+1}^B])}^{\text{Disagreement}}}{\gamma\sigma_B^2}}_{\text{MV optimal bond demand}} + \underbrace{\frac{D_L}{D_B}L_t}_{\text{Duration hedging demand}}$$

The optimal allocation to long-term bonds is relatively standard. There is a mean-variance component based on expected excess returns and a hedging demand component that aims to match the duration of assets and liabilities. However, what is non-standard here is that the expected excess returns is a linear combination of the insurer's expectations and the regulator model-implied expectations. Specifically, the allocation tilts more towards the regulator's model if: (i) the relative cost of the regulation (α_t) is high, and (ii) there is more disagreement between the regulator's model and the insurer ($\mathbb{E}_t^R[rx_{t+1}^B] - \mathbb{E}_t^I[rx_{t+1}^B]$).

In terms of disagreement, section 3.2 shows that the regulator's model expects yields to rise less relative to insurers, hence, the regulator expects long-term bonds to earn higher excess returns, i.e., $\mathbb{E}_t^R[rx_{t+1}^B] > \mathbb{E}_t^I[rx_{t+1}^B]$. Therefore, the model predicts that if the cost of regulation (α_t) or the regulator model's expected excess returns ($\mathbb{E}_t^R[rx_{t+1}^B]$) increases, we should observe an increase in long-term bond holdings.

The model points to two complementary approaches for identifying the impact of regulator model-implied expectations on bond allocations:

1. Exploiting variation in the relative cost (α_t) of the model-dependent loss constraints, and
2. Isolating quasi-random shifts in the regulator model-implied expectations ($\mathbb{E}_t^R[rx_{t+1}^B]$)

I begin with the first, which helps quantify the overall impact of the model-dependent constraint on insurers holdings. I then proceed to the second, which provides a cleaner test of the mechanism that dynamics in regulator model-implied expectations directly drive the dynamics in insurer allocation decisions.

4.2 Cost exposure experiment: enactment of Valuation Manual 20

In 2017, insurance regulators enacted Valuation Manual 20 (VM20) which introduced a new methodology for calculating statutory reserves for traditional life insurance products. Under

VM20, for each policy sold, insurers need to use asset/liability loss simulations to determine statutory reserves, which they must recalculate yearly as long as the policy is in effect.[20] To ensure the efficacy of financial projections underlying these simulations, the regulator prescribes the pre-calibrated interest rate model outlined in section 2.1.

In terms of the roll-out, the regulation was announced at the end of 2012 and enacted in 2017. VM20 allowed insurers a three-year opt-out transition period, with full compliance required starting in 2020. It also grandfathered old policies, so only newly issued policies were required to calculate VM20 reserves. Both these aspects of the roll out mean that a larger share of policies are covered by VM20 over time.

In terms of product coverage, the regulation impacted only traditional life products and excluded fixed annuities, marking a departure from previous statutory reserving rules. Previously, regulation required a similar method to calculate the reserves for both product types, consistent with their comparable exposure to interest rate risk. Traditional life policies, once pooled, have their mortality risk nearly fully diversified, resulting in a cash flow structure akin to that of fixed annuities. Consequently, the payoffs of both can be replicated by zero-coupon bonds.

The staggered roll-out and product-specific implementation of VM20 generate both cross-sectional and temporal variation in the cost of disagreeing with the regulator's model. Temporally, because the regulation applies only to new policies, its influence strengthens over time as a growing share of insurer liabilities becomes subject to regulator's model—that is, α_t increases over time. Cross-sectionally, traditional life insurers must use the regulator's projections to determine reserves, while fixed annuity writers remain exempt, creating heterogeneity in exposure to the model despite similar underlying interest rate risk. In this setting, $\alpha_t = 0$ for fixed annuity providers, and an increasing $\alpha_t > 0$ for traditional life insurers. I exploit both dimensions of variation in a difference-in-differences design to estimate the effect of model-based regulation on insurers' bond allocations.

Data: Using S&P Capital IQ's database of NAIC annual fillings, I get insurer bond holdings (schedule D) and other balance sheet items. I measure insurer's long-term bond holdings using the Macaulay duration of their portfolios. To avoid any price effects driving the results, I con-

[20] $\text{Reserves} = \max\{\text{Net Premium Reserve (NR)}, \text{Deterministic Reserve (DR)}, \text{Stochastic Reserve (SR)}\}$. The NPR is the present value calculated using a prescribed discount rate that remains unchanged for the policy's life. The DR and SR require simulating losses, which are recalculated yearly. The SR is the 70% expected shortfall of the assets invested to meet policy payouts. The DR is the maximum value of liability simulations from sixteen prescribed scenarios; each scenario involves a prescribed path of shocks input into the regulator's model. The resulting financial projections determine asset returns, which they use to discount liabilities.

struct the duration measure using the bonds par value.

For the treatment group, I use insurance companies specializing in selling traditional life products, i.e., those subject to VM20 regulation. I define these firms as those for whom traditional life insurance comprised at least 90% of their life and annuity reserves in 2013, four years before the VM20 was enacted, to help avoid contamination from VM20 in the definition of the treatment group. The results are robust to using earlier and later life shares to define the treatment group, and using different thresholds (see table 8).

For the control group, I use fixed annuity providers, defined as all insurance companies not included in the treatment group. Fixed annuity providers serve as a good control group because, as discussed earlier, life and fixed annuity products have similar underlying cash flow structures and used the same statutory reserving methodology prior to 2017. After restricting to the treatment and control groups, the sample period spans from 2011 to 2022, with a total of 3,692 insurer-year observations 30% of which are for the treatment group. In 2022, the treatment group held \$324 billion in assets, and the control group held \$4.7 trillion in assets.

Estimation: I run the following insurer-year difference-in-differences (DiD) regression, with all estimates relative to the year 2016, the last year before VM20 was enacted:

$$B_{i,t} = \alpha + \beta_1 \delta_{i \in LI} + \sum_{h \notin 2016} \beta_{3,h} \delta_{t \in h} + \sum_{h \notin 2016} \beta_h^{VM20} (\delta_{t \in h} \times \delta_{i \in LI}) + \gamma \text{VA Share}_{i,t} + \varepsilon_{i,t} \tag{1}$$

where $B_{i,t}$ is measured using the portfolio's duration in years, $\delta_{i \in LI} := \mathbf{1}\{i \in LI\}$ is a dummy for life insurers (treatment group), and $\delta_{t \in h} := \mathbf{1}\{t = h\}$ is a year dummy. β_h^{VM20} is the estimated treatment effect of interest in year h, i.e., the change in long-term bond holdings due to the VM20 regulation in year h.[21]

VA $\text{Share}_{i,t}$ controls for the variable annuity (VA) share of a firm. It is important to control for VA share because VAs have a very different cash flow profile compared to life insurance and fixed annuity policies. For example, they are exposed to equity returns and often contain put option-like payoffs that give them highly non-linear exposure to interest rates. Additionally, VAs transitioned to model-based reserving approaches much earlier.[22] Both the distinct cash flow profile VAs and their earlier transition to model-based reserving require controlling for VA

[21]Due to the opt-out period, for 2017-2019, the estimate captures the average Intent-to-Treat (ITT) effect, while for the post-2019 period, it captures the Average Treatment-on-the-Treated (ATT).

[22]Enacted in 2008, Actuarial Guideline 43 (AG43) requires insurers to use 70% expected shortfall for VA reserve calculations. Unlike VM20, AG43 allows for considerable freedom in how interest rates are simulated (e.g. no model prescription). In 2020, Valuation Manual 22 for VAs went into effect (with a three-year opt-out period), which now effectively mandates the use of the regulator's pre-calibrated interest rate model.

share to ensure parallel trends holds—Appendix D.3 shows the results are robust to dropping firms that have a VA business.

Assuming parallel trends hold between the treatment and control groups, the β_t^{VM20} estimand captures the effect (see Appendix D.2 for derivation),

$$\beta_t^{VM20} = \mathbb{E}^{(i)}\left[\alpha_{i,t}^{VM20}\frac{\mathbb{E}_t^R[rx_{t+1}^B] - \mathbb{E}_{i,t}^I[rx_{t+1}^B]}{\gamma_i\sigma_B^2}\middle| VM20 = 1\right]$$

where $\mathbb{E}^{(i)}[\cdot]$ represents the cross-sectional average across insurers. In other words, according to the model, the DiD estimand captures the total effect from (i) the VM20-induced increase in the relative cost of regulation and (ii) the disagreement in expectations.

Regarding the relative cost of regulation, the rollout timeline of VM20 implies that $\alpha_{i,t}^{VM20}$ was zero before 2017 and then gradually increased as fewer insurers opted out and newly issued policies became subject to the regulation. Regarding the disagreement, as shown in Section 2, on average, the regulator's model expects higher excess returns than insurers, i.e., $\mathbb{E}_t^R[rx_{t+1}^B] > \mathbb{E}_{i,t}^I[rx_{t+1}^B]$. Therefore, we should see no treatment effect before 2017, with a steady increase over time after that.

Figure 6a plots the average portfolio duration (adjusted for the group fixed effect) for the treatment and control groups, and figure 6b plots the estimated treatment effect $\hat{\beta}_t^{\text{VM20}}$.[23] Overall, the results align with the predictions. Prior to the enactment of VM20, the portfolio duration of both groups is similar, suggesting the parallel trends assumption holds. Once the legislation is enacted, traditional life providers' portfolio duration begins to steadily increase, and by the end of the sample period, they had risen by over a one year relative to the control group; a rise from around 6.7 year portfolio duration to a 7.8 year portfolio duration. This effect will likely continue grow as a larger portion of outstanding policies become subject to VM20.

A potential concern is that insurers may have been differentially affected by the Tax Cuts and Jobs Act (TCJA), which went into effect in 2018. The Society of Actuaries found that the reforms generally had a modest impact on the profitability of product lines, with the exception of products reinsured through captives, which became less profitable (Clingerman et al., 2018). Traditional life insurers are generally less likely to use captive reinsurance than other insurers. With that said, Appendix D.3 shows that the estimates are robust to controlling for year fixed effects interacted with insurers' exposure to the TCJA.

[23]See table 8 and figure 16 in the Appendix for the DiD estimates and permutation inference test, respectively.

Figure 6: Difference-in-Difference: VM20 impact on long-term bond holdings

(a) Treatment vs control group (b) Treatment effect estimate, $\hat{\beta}_t^{VM20}$

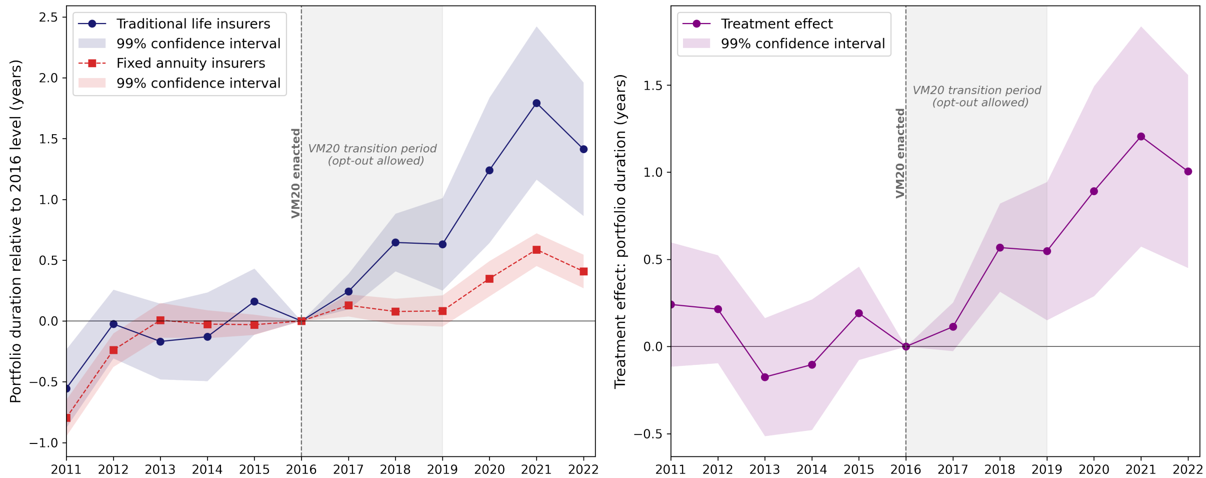

The figure on the left plots the treatment group's (life insurers) and control group's (non-variable annuity providers) portfolio duration (years) relative to their 2016 level. The figure on the right plots the difference between the treatment group's and control group's portfolio duration. The error bars are the 99% confidence intervals, which are calculated using standard errors that are two way clustered at the insurer and year level.

4.3 Expectation shock experiment: exploiting ad hoc re-calibration

A more direct test of the mechanism that dynamics regulator model-implied expectations directly driver insurer allocation is to assess how insurers' duration changes in response to changes in the model's expectations. A naive method would be to simply run OLS. However, the issue with this approach is that changes in regulator model-implied expectations may be correlated with other confounders, such as changes in insurers' expectations or other belief moments.[24] To overcome this identification challenge, I isolate plausibly exogenous changes in regulator model-implied expectations by exploiting the ad-hoc rule the regulator uses to update the mean reversion point parameter in their model.

As outlined in Section 2.1, the regulator models long-term interest rates as a mean-reverting process. Hence, for a bond with duration D_B, the regulator model's one-period ahead expected excess returns are approximately given by:

$$\mathbb{E}_t^R[rx_{t+1}^B] \approx -D_B\mathbb{E}_t^R[\Delta y_{l,t+1}] - r_{t+1}^f = D_B\psi \underbrace{\left(y_{20,t} - \bar{y}_{20}^{(t)}\right)}_{:=y^R_{gap,t}} - r_{t+1}^f$$

[24]We cannot use insurers' stated expectations as a control because we do not observe the beliefs of all life insurers. Furthermore, as I will show, insurers' stated beliefs seem to be an equilibrium outcome of optimization that depends directly on regulator model-implied beliefs. Hence, using these stated beliefs as a control would introduce a bad control bias.

where $\Delta y_{l,t+1}$ is the change in the long-run yield (as proxied in the model by the 20 year yield), ψ is the mean-reverting rate, $y_{20,t}$ is the current 20-year yield, and $\bar{y}_{20}^{(t)}$ is period t parameter of the mean-reversion point. When $y_{gap,t}^{R}$ is positive, yields are above their long-run mean, implying that yields are expected to decline and bond prices are expected to rise. Consequently, any changes in the value of the mean reversion parameter will induce changes in the regulator model-implied expected excess returns.

Specifically, at the start of every year, the regulator updates the mean-reversion point parameter using a linear combination of three different moving averages/medians of 20-year Treasury yields (a 600-month moving median, a 120-month moving average, and a 36-month moving average), which they then round to the nearest 25 basis points:

$$\bar{y}_{20}^{(t)} = \mathbf{Round}_{25\text{bps}}\Big\{0.2 \times \text{Med}_{600m,t}\left(y_{20}\right) + 0.3 \times \text{Mean}_{120m,t}\left(y_{20}\right) + 0.5 \times \text{Mean}_{36m,t}\left(y_{20}\right)\Big\}$$

Hence, changes in the mean-reversion point parameter, $\Delta\bar{y}_{20}^{(t)} := y_{20}^{(t)} - y_{20}^{(t-1)}$, can be decomposed into three sources: (i) contemporaneous yield information entering the moving window, (ii) old information exiting the window at the 36-month, 120-month, and 600-month points, and (iii) changes in the extent of rounding:

$$\Delta\bar{y}_{20}^{(t)} = \Delta\bar{y}_{20,\text{ new data entering windows}}^{(t)} + \underbrace{\Delta\bar{y}_{20,\text{ old data exiting windows}}^{(t)} + \Delta\bar{y}_{20,\text{ rounding}}^{(t)}}_{:=\Delta\bar{y}_{20,\text{ad hoc}}^{(t)}}$$

see Appendix D.4 for the derivation of this decomposition identity.

Figure 7a plots contribution to changes in mean reversion parameter coming from the three terms of the decomposition identity. The contributions from dropping observations at the arbitrary window cut-off points and the rounding rule (the ad hoc aspects of the update rule, $\Delta\bar{y}_{20,\text{ad hoc}}^{(t)}$) are major contributors of the changes in the parameter.

As figure 7b shows, the weight assigned to various lagged yields has a constant contribution that changes discretely once they age from 36 to 37 months and from 120 to 121 months, experiencing 70% and 100% declines at those arbitrary thresholds, respectively. As long as insurers' latent expectations do not exhibit these exact same discrete weighting changes, the variation induced from observations exiting the window should be exogenous. Similarly, if insurers do not use the same 25 basis point rounding, the variation induced by rounding should also be exogenous.

Estimation: To estimate the effect of changes in regulator model-implied excess returns, proxied

Figure 7: Drivers of regulator's long-run mean, $\bar{y}_{20}^{(t)}$, update rule

(a) Decomposing changes in long-run mean, $\Delta\bar{y}_{20}^{(t)}$ (b) Update rule weight assigned to lagged yields

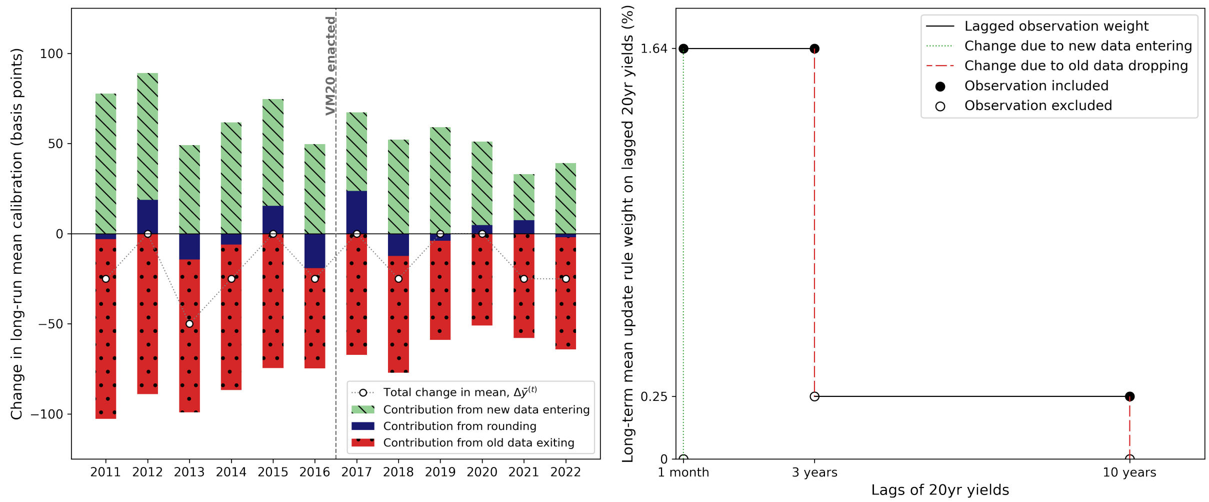

Figure on the left decomposes changes in the regulator's long-run mean calibration, $\Delta\bar{y}_{20}^{(t)}$, into three components: (i) contemporaneous yield information entering the moving windows (grey), (ii) old information exiting the windows at the 36 month, 120 month, and 600 month points (red), and (iii) changes due to rounding to the nearest 25bps (blue). Figure on the right, plots the weights in $\bar{y}_{20}^{(t)}$ assigned to yields of different lags by the 36 month and 120 month moving averages.

using $\Delta y^R_{gap,t}$, I run the following regression, leveraging both cross-sectional exposure to VM20 and time variation in regulator model-implied expectations:

$$\Delta B_{i,t} = \alpha_i + \delta_t + \beta(\mathbf{1}_i\{VM20\} \times \Delta y^R_{gap,t}) + \varepsilon_{i,t}$$

where $\Delta B_{i,t}$ is the change in long-term bond holdings, measured as the change in the insurer's portfolio duration in years, α_i and δ_t are insurer and year fixed effects, respectively, and $\mathbf{1}_i\{VM20\}$ is an indicator for whether insurer i is subject to VM20 (i.e., a traditional life insurance provider that did not opt out of the regulation).[25]

To address the concern of potential confounders, I construct an instrument using the ad hoc variation in $\Delta\bar{y}_{20}^{(t)}$:

$$z_{i,t}^{\text{ad hoc}} = \mathbf{1}_i\{VM20\} \times \Delta\bar{y}^{(t)}_{20,\text{ad hoc}}$$

The key identification assumption is that insurers' latent expectations do not use the exact same rules for rounding and window observations. The relevance of this instrument is some-

[25] I identify life insurers that opted out using the supplementary filings for firms using VM20's small insurer exemption. This may underestimate the number of insurers that opted out, as it misses those life insurers that did not apply for the small insurer exemption but took advantage of the 2017 to 2019 transition opt-out. As a result, the estimates in the regressions below may underestimate the effect of changes in regulator model-implied expectations.

what ensured since $\Delta\bar{y}^{(t)}_{20,\text{ad hoc}}$ is a substantial contributor to $\Delta\bar{y}^{(t)}_{20}$, which directly determines the regressor $\Delta y^R_{gap,t} := y_{20,t} - \Delta\bar{y}^{(t)}_{20}$. For robustness, I also construct instruments that isolate the variation induced by rounding and dropping observations: $z^{\text{round}}_{i,t} = \mathbf{1}_i\{VM20\} \times \Delta\bar{y}^{(t)}_{20,\text{rounding}}$ and $z^{\text{drop}}_{i,t} = \mathbf{1}_i\{VM20\} \times \Delta\bar{y}^{(t)}_{20,\text{old data exiting windows}}$, respectively.

Table 2 presents the regression results. I run these regressions for the period when VM20 was in effect, using the same sample of insurers as in the difference-in-differences exercise. Column 1 presents the OLS regressions and finds a 1pp increase (approximately a standard deviation increase) in Δy^{gap}_t is associated with a portfolio duration increase of 0.27 years (compared to a typical portfolio duration of 6.8 years). The positive estimate is consistent with the mechanism, however, this estimate may suffer from bias. To address this, column 3 instruments $(\mathbf{1}_i\{VM20\} \times \Delta y^R_{gap,t})$ with $z^{\text{ad hoc}}_{i,t}$. Column 2 shows that the first stage of this instrument is strong, with an r-squared of 48% (with-in r-squared of 34%). The IV estimate in column 3 shows that a 1 pp increase in Δy^{gap}_t causes portfolio duration to significantly increase by 0.29 years. Suggesting that bias is likely not a major concern in this setting; hence giving more confidence in the DiD estimates too.

In terms of robustness checks, columns 4 and 5 repeat the IV specification but use instruments that only leverage variation induced by rounding ($z^{\text{round}}_{i,t}$) and by dropping observations from the moving window ($z^{\text{drop}}_{i,t}$). Both estimates are significant and statistically similar to the instrument that uses both sources of variation. Finally, I conduct a placebo test by re-running the specification for the pre-VM20 sample when model-based beliefs should not affect insurer allocation. I find the point estimate has the incorrect sign and are statistically insignificant. Overall, these results further support the outlined mechanism of regulator model-implied expectations significantly effecting insurers' portfolio allocation.

Overall, the findings show that even arbitrary changes in regulator model-implied expectations passthrough into insurer allocation decisions. This suggests that the belief patterns in the model's expectations—that mimic those in surveyed expectations—passthrough into insurer actions.

4.4 Regulator model-implied beliefs and insurer reported beliefs

I next assess the passthrough from regulator model-implied expectations to insurers' own reported interest rate assumptions. As with institutional disclosures more generally, insurers' assumptions may reflect their internal subjective beliefs but also strategic considerations shaped by regulatory and market pressures.

Table 2: Effect of changes in regulator expectations on insurer portfolio duration

	OLS	IV				Placebos
	Baseline	First stage	Baseline	$z_{i,t}^{\text{round}}$	$z_{i,t}^{\text{drop}}$	PreVM20
$\mathbf{1}_i\{\text{VM20}\} \times \Delta y_{gap,t}^R$	0.27***		0.29***	0.33***	0.23***	-0.35
	(0.05)		(0.09)	(0.11)	(0.08)	(0.27)
$z_{i,t}^{\text{ad hoc}}$		-2.85***				
		(0.13)				
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	981	981	981	981	981	1,596
R-squared	0.17	0.49	—	—	—	—
First stage F-stat	—	—	459	378	317	166

Driscoll-Kraay standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table presents results from regressing ΔB_{it} on $\Delta y_{gap,t}^R := \Delta y_{20,t} - \Delta \bar{y}_{20}^{(t)}$. Columns 1 is an OLS regressions to provide a comparison for the baseline IV specification in presented in columns 3 and its corresponding first stage is in column 2. The IV specifications instrument $\Delta y_{gap,t}^R$ with $z_{i,t}^{\text{ad hoc}} := \mathbf{1}_i\{VM20\} \times \Delta \bar{y}_{20,\text{ad hoc}}^{(t)}$. Column 4 and 5, repeat the baseline IV specification but use $z_{i,t}^{\text{round}}$ and $z_{i,t}^{\text{drop}}$ as instruments, respectively. Columns 6 does a placebo test, rerunning the baseline specification for the same treatment assignment but for the pre-VM20 period. The standard errors are based on Driscoll-Kraay errors.

The enactment of VM20 represents a shift in these strategic incentives. The regulation required insurers to calculate statutory reserves using the regulator's prescribed interest rate model, pushing them to take actions consistent with the regulator model's outlook even when it diverged from their own assumptions. Creating an incentive to align their disclosed view closer to the regulators, so that their actions were more in line with their own assumptions.

To quantify this passthrough, I conduct an event study around the implementation of VM20. Figure 8a plots the average difference between insurers' and the regulator's six-year-ahead interest rate expectations by year. Following VM20's introduction in 2017, insurers' assumptions converged sharply toward the regulator's model: the average disagreement declined from about 140 basis points before 2017 to roughly 40 basis points afterward. Figure 8b shows that this convergence towards the model is evident not just for six-year-ahead expectations, but across the term structure of expectations.

The sharp convergence in insurer assumptions around VM20 is most consistent with strategic behavior rather than learning. Investors, analysts, and rating agencies closely scrutinize these

Figure 8: Pre/post VM20 impact on regulator vs. insurer disagreement

(a) Disagreement for 6-year horizon by year (b) Disagreement by forecast horizons

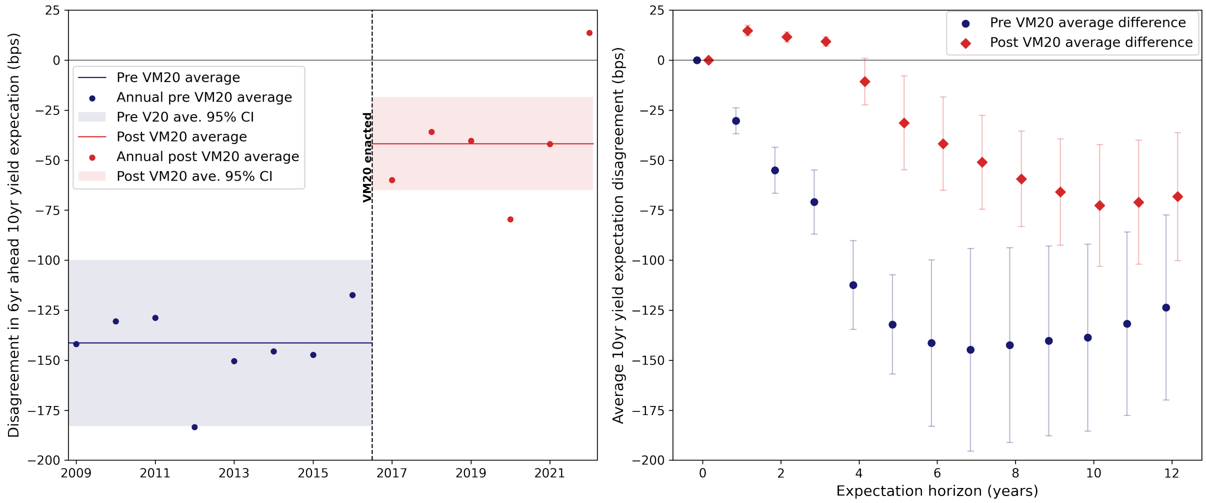

The left figure shows disagreement for different years for the 6yr projection horizon. The 95% confidence band is based on SE clustered at the insurer level. The right figure shows the average difference between regulator model-implied and insurer expectations of 10yr yields for different forecast horizons, using annual data for pre-VM20 (2009-2016) and pos-VM20 (2017-2022), and 95% confidence intervals based on SE clustered at the insurer level.

assumptions. Once VM20 was enacted, insurers had incentives to align their reported outlooks with the regulator's model to avoid discrepancies between their stated views and regulator-driven actions that could attract scrutiny. A learning-based explanation is less plausible: the regulator's interest rate model had been publicly available since 2009, yet expectations shifted only when VM20 became binding in 2017. If learning were the driver, one would expect a gradual convergence or an earlier adjustment following the model's release, neither of which is observed. Taken together, these findings suggest that institutionally disclosed assumptions should not be interpreted as purely subjective beliefs but rather as shaped by strategic and regulatory considerations.

Because insurers use these assumptions to guide risk management, pricing, and planning decisions, the convergence in reported expectations highlights that regulator model dynamics likely influence multiple dimensions of insurer behavior beyond the portfolio allocation channel directly tested.

5 Conclusion

Intermediaries' regulatory constraints are increasingly shaped by regulator model-implied beliefs. Using the U.S. insurance sector as a laboratory, I show that the regulator model quantitatively mirrors the systematic belief patterns observed in human forecasters, embedding these dynamics into the evolution of regulatory constraints. As a result, intermediaries directly pass through these model-implied beliefs into their portfolio decisions, effectively coordinating their actions at the sector level. Together, these findings reveal a new channel through which systematic belief patterns become consequential for the real economy: once embedded in regulation through models, they transmit and synchronize the behavior of regulated firms.

These findings have implications for both asset pricing and financial stability. Regulator model-implied beliefs coordinate and drive intermediary demand, linking the dynamics of these beliefs to the dynamics of intermediary balance sheets. This connection generates a new source of common shocks that move asset prices and amplify systemic risk. It also raises normative questions about regulatory design. First, how should regulators design and update models when faced with ambiguity about the underlying data generating process? Second, how should policymakers balance the benefits of prescriptive forward-looking models—which help anticipate risks and limit the scope for firms to game the rules—against their costs, as such models synchronize allocation decisions across institutions? Coarser, less prescriptive constraints grant firms more discretion in their portfolio choices and may therefore reduce correlated exposures, but at the expense of weaker control over risk.

Beyond regulation, the same mechanism likely operates more broadly. As model-based risk management and algorithmic trading become widespread, their implied beliefs may also exhibit systematic patterns that pass through to firm actions. Similar coordination can arise whenever institutions rely on common models, whether through shared risk systems, third party software, or learning algorithms. For example, intermediaries' internal risk teams often converge on similar platforms, such as BlackRock's Aladdin system, which is used to manage risk of over $21 trillion assets as of 2020. Understanding how such model-driven convergence shapes capital allocation and market dynamics remains an open question.

References

Acharya, Viral V., Lasse H. Pedersen, Thomas Philippon, and Matthew Richardson, "Measuring Systemic Risk," *Review of Financial Studies*, 2017, *30* (1), 2–47.

Adam, Klaus, Albert Marcet, and Johannes Beutel, "Stock price booms and expected capital gains," *American Economic Review*, 2017, *107* (8), 2352–2408.

— **and Stefan Nagel**, "Expectations data in asset pricing," in "Handbook of economic expectations," Elsevier, 2023, pp. 477–506.

Andonov, Aleksandar and Joshua D Rauh, "The return expectations of public pension funds," *The Review of Financial Studies*, 2022, *35* (8), 3777–3822.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, "X-CAPM: An extrapolative capital asset pricing model," *Journal of Financial Economics*, 2015, *115* (1), 1–24.

Becker, Bo and Victoria Ivashina, "Reaching for yield in the bond market," *The Journal of Finance*, 2015, *70* (5), 1863–1902.

— **, Marcus M Opp, and Farzad Saidi**, "Regulatory forbearance in the US insurance industry: The effects of removing capital requirements for an asset class," *The Review of Financial Studies*, 2022, *35* (12), 5438–5482.

Behn, Markus, Rainer Haselmann, and Vikrant Vig, "The limits of model-based regulation," *The Journal of Finance*, 2022, *77* (3), 1635–1684. Publisher: Wiley Online Library.

Bordalo, Pedro, Nicola Gennaioli, Andrei Shleifer, and Stephen J Terry, "Real credit cycles," Technical Report, National Bureau of Economic Research 2021.

Carbonell, Jaime and Jade Goldstein, "The use of MMR, diversity-based reranking for reordering documents and producing summaries," in "Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval" 1998, pp. 335–336.

Charles, Constantin, Cary Frydman, and Mete Kilic, "Insensitive investors," *The Journal of Finance*, 2024, *79* (4), 2473–2503.

Chaudhary, Manav, Zhiyu Fu, and Jian Li, "Corporate bond multipliers: Substitutes matter," *Available at SSRN*, 2023.

Chodorow-Reich, Gabriel, Andra Ghent, and Valentin Haddad, "Asset insulators," *The Review of Financial Studies*, 2021, *34* (3), 1509–1539.

Clingerman, Curt, Paul Fedchak, Casey Malone, and Craig Reynolds, "Tax Reform Impacts onLife Insurance Pricingand Profitability," Technical Report 110, Society of Actuaries June 2018.

Cormack, Gordon V., Charles L A Clarke, and Stefan Buettcher, "Reciprocal rank fusion outperforms condorcet and individual rank learning methods," in "Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval" ACM 2009, pp. 758–759.

Couts, Spencer J., Andrei S. Goncalves, and Johnathan Loudis, "The Subjective Risk and Return Expectations of Institutional Investors," 2024.

Dahlquist, Magnus and Markus Ibert, "Equity return expectations and portfolios: Evidence from large asset managers," *The Review of Financial Studies*, 2024, *37* (6), 1887–1928.

d'Arienzo, Daniele, "Maturity increasing overreaction and bond market puzzles," *Available at SSRN 3733056*, 2020.

Egan, Mark, Shan Ge, and Johnny Tang, "Conflicting interests and the effect of fiduciary duty: Evidence from variable annuities," *The Review of Financial Studies*, 2022, *35* (12), 5334–5386.

Ellul, Andrew, Chotibhak Jotikasthira, Anastasia Kartasheva, Christian T Lundblad, and Wolf Wagner, "Insurers as asset managers and systemic risk," *The Review of Financial Studies*, 2022, *35* (12), 5483–5534.

_ , _ , and Christian T Lundblad, "Regulatory pressure and fire sales in the corporate bond market," *Journal of Financial Economics*, 2011, *101* (3), 596–620.

_ , _ , _ , and Yihui Wang, "Is historical cost accounting a panacea? Market stress, incentive distortions, and gains trading," *The Journal of Finance*, 2015, *70* (6), 2489–2538.

Farmer, Leland, Emi Nakamura, and Jon Steinsson, "Learning about the long run," Technical Report, National Bureau of Economic Research, Inc 2021.

Froot, Kenneth A, David S Scharfstein, and Jeremy C Stein, "Risk management: Coordinating corporate investment and financing policies," *the Journal of Finance*, 1993, *48* (5), 1629–1658.

Froot, Kenneth A., David S. Scharfstein, and Jeremy C. Stein, "A Framework for Risk Management," *Harvard Business Review*, November 1994. Section: Financial analysis.

Ge, Shan, "How do financial constraints affect product pricing? evidence from weather and life insurance premiums," *The Journal of Finance*, 2022, *77* (1), 449–503.

_ and Michael S Weisbach, "The role of financial conditions in portfolio choices: The case of insurers," *Journal of Financial Economics*, 2021, *142* (2), 803–830.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus, "Five facts about beliefs and portfolios," *American Economic Review*, 2021, *111* (5), 1481–1522.

Gormsen, Niels Joachim and Kilian Huber, "Corporate discount rates," Technical Report, National Bureau of Economic Research 2023.

Greenwood, Robin, Samuel G Hanson, and Lawrence J Jin, "Reflexivity in credit markets," Technical Report, National Bureau of Economic Research 2019.

Koijen, Ralph SJ and Motohiro Yogo, "The Cost of Financial Frictions for Life Insurers," *American Economic Review*, 2015, *105* (1), 445–475.

_ and _ , "Shadow insurance," *Econometrica : journal of the Econometric Society*, 2016, *84* (3), 1265–1287.

_ and _ , "Risk of life insurers: Recent trends and transmission mechanisms," 2017. Publisher: National Bureau of Economic Research.

_ and _ , "Global life insurers during a low interest rate environment," in "AEA papers and proceedings," Vol. 112 2022, pp. 503–508.

Krishnamurthy, Arvind and Wenhao Li, "Dissecting mechanisms of financial crises: Intermediation and sentiment," Technical Report, National Bureau of Economic Research 2020.

Li, Ziang, "Long rates, life insurers, and credit spreads," 2024.

Maxted, Peter, "A macro-finance model with sentiment," *Review of Economic Studies*, 2024, *91* (1), 438–475.

Merkle, Christoph and Martin Weber, "Do investors put their money where their mouth is? Stock market expectations and investing behavior," *Journal of Banking & Finance*, 2014, *46*, 372–386.

Murray, Scott and Stanislava Nikolova, "The bond-pricing implications of rating-based capital requirements," *Journal of Financial and Quantitative Analysis*, 2022, *57* (6), 2177–2207.

Nagel, Stefan and Zhengyang Xu, "Dynamics of subjective risk premia," *Journal of Financial Economics*, 2023, *150* (2), 103713.

Nanda, Vikram, Wei Wu, and Xing Alex Zhou, "Investment commonality across insurance companies: Fire sale risk and corporate yield spreads," *Journal of Financial and Quantitative Analysis*, 2019, *54* (6), 2543–2574.

Plosser, Matthew C. and Joao A. C. Santos, "Banks' Incentives and the Quality of Internal Risk Models," December 2014.

Sen, Ishita, "Regulatory limits to risk management," *The Review of Financial Studies*, 2023, *36* (6), 2175–2223.

_ **and Varun Sharma**, "Internal Models, Make Believe Prices, and Bond Market Cornering," *SSRN Electronic Journal*, 2020.

Tang, Johnny, "Regulatory competition in the us life insurance industry," Technical Report, working Paper 2022.

Verani, Stephane and Pei Cheng Yu, "What's Wrong with Annuity Markets?," July 2021.

Wang, Chen, "Under-and overreaction in yield curve expectations," *Available at SSRN 3487602*, 2021.

A Regulator expectations and loss-based constraints

A.1 Investment horizon and loss distributions

Typically means are ignored when working with short-horizon loss distributions based constraints e.g. Value-at-Risk (VaR), and Conditional Tail Expectations (CTEs). In this section, I show that means are of first order importance when considering loss distribution based constraints at longer horizons. Intuitively, the return on buying and holding stocks for just a few days depends critically on volatility, but over longer horizons, the expected return becomes more significant, and over sufficiently long periods, positive returns are typically realized despite recessions, wars, and pandemics.

As a warm-up example, suppose daily returns are iid and follow $r_{t+1} \sim N(\mu, \sigma^2)$. Then the daily Value-at-Risk per dollar is given by,[26]

$$VaR(p) = -\mu - \sigma\Phi^{-1}(p)$$

Daily expected return are typically several orders of magnitudes smaller than daily volatility $\sigma \gg \mu$. Hence, VaR essentially is a scaling of σ; so much so that in practice typically μ is assumed to be zero.

However, consider an h-period buy and hold return. The h-horizon return is approximately distributed as $r_{t:t+h} \sim N(h\mu, h\sigma^2)$, where both means and variances scale by T. In this case the h-period Value-at-Risk per dollar is given by,

$$Var_h(p) = -h\mu - \sqrt{h}\sigma\Phi^{-1}(p)$$

where $p < 0.5$ and $\Phi^{-1}(p) < 0$. At longer horizons μ's contribution is non-neglegible. In fact in this iid returns case, the contribution share of the mean goes to 1,

$$\lim_{h\to\infty} \frac{-h\mu}{Var_h(p)} = 1$$

This poses a challenge for managing long-horizons risks as first moments are much harder to correctly estimate than second moments.

[26] Given the normality assumed here, other commonly used risk measures (e.g. CTE) simply scale up the variance term relative to VaR. Hence, the qualitative implications are unchanged by working with VaR.

A.1.1 Long-horizon VaR and calibrated interest rate process

With mean-reverting interest rate process, the ultimate dominance of the mean is not gauranteed, since as $h \to \infty$, yields return to their long-run mean and do not generate capital gains/losses. In this subsection, I calibrate a mean-reverting model using common parameters and explore the relative importance of the mean for VaR.

Suppose long-term interest rates follow a mean reverting process with long term mean $\bar{y}$,

$$\Delta y_{t+1} = -\psi(y_t - \bar{y}) + \varepsilon_{t+1} \quad \text{where} \quad \varepsilon_{t+1} \sim^{iid} N(0, \sigma^2)$$

We can iterate this process forward to find the h-horizon change in yields $\Delta y_{t:t+h} := y_{t+h} - y_t$,

$$\Delta y_{t:t+h} = \psi\bar{y}\sum_{l=1}^{h}(1-\psi)^{l-1} + [(1-\psi)^h - 1]y_t + \sum_{l=1}^{h}(1-\psi)^{h-l}\varepsilon_{t+l}$$

This implies that the (approximate) h-period return of a perpetual bond with duration D is distributed,

$$r_{t:t+h} \sim N\left(-D[(1-\psi)^h - 1](y_t - \bar{y}), (D\sigma)^2\frac{1-(1-\psi)^{2h}}{1-(1-\psi)^2}\right)$$

Hence, the h period Value-at-Risk per dollar is given by,

$$VaR_h(p) = D[(1-\psi)^h - 1](y_t - \bar{y}) - D\sigma\sqrt{\frac{1-(1-\psi)^{2h}}{1-(1-\psi)^2}}\Phi^{-1}(p)$$

For $h = 1$ and $h = \infty$ the $VaR_h(p)$ are given by,

$$\begin{aligned} VaR_1(p) &= -D\psi(y_t - \bar{y}) - D\sigma\Phi^{-1}(p) \quad \text{and,} \\ \lim_{h\to\infty} VaR_h(p) &= -D\left[(y_t - \bar{y}) + \sigma\frac{\Phi^{-1}(p)}{\sqrt{1-(1-\psi)^2}}\right] \end{aligned}$$

Typical estimates of ψ are around 0.1 per year ($\approx$ 7 year half life). Hence, mean have a negligible contribution to daily VaR. But in the long-run, ψ doesn't attenuate the contribution of the mean. Unlike, the iid return model in the previous subsection, where the mean's relative contribution tended to 1. In this case, the mean's relative contribution tends to a constant share.

To assess whether the contribution of the mean is quantitatively important, I calibrate the model and quantify the contribution of the mean to $VaR_h(p)$ for different horizons.[27] Figure

[27] $\sigma = 4.3bps$ based on standard deviation of daily changes in 1 year ZCB US government bond; $\psi = 0.1313$ per year (convert this to daily rate for the simulations) from Wang et al. (2013); $D = 10$ years

9 plots the contribution of the mean and variance components of $VaR_h(5\%)$ for different horizons. Figure 9b assumes that yields are initially 200bps below their long-run mean, and hence on average, the bond will experience capital losses as yields rise towards their long-run mean. These expected capital losses for short horizon holding periods have a negligible impact on VaR calculations. However, they meaningfully increase VaR for longer holding periods, accounting for over 40% of total VaR at the 20-year horizon. Figure 9a assumes that yields are initially 200 bps above their long-run mean, and hence on average, the bond will experience capital gains as yields decline towards their long-run mean. For shorter holding periods, the contribution of the mean is negligible, but increases sizably over longer holding periods, and begins to offset the contribution from the variance.

For insurance companies, whose equity typically has negative duration, the response to a positive vs. negative $y_t - \bar{y}$ gap will be flipped. When y_t is above $\bar{y}$, the expected return helps offset the contribution from variance, as equity increases in value as yields rise. The opposite is true for when y_t is below $\bar{y}$.

Figure 9: 95% Value-at-Risk per dollar for different holding horizon

(a) y_t 200bps below $\bar{y}$ (b) y_t 200bps above $\bar{y}$

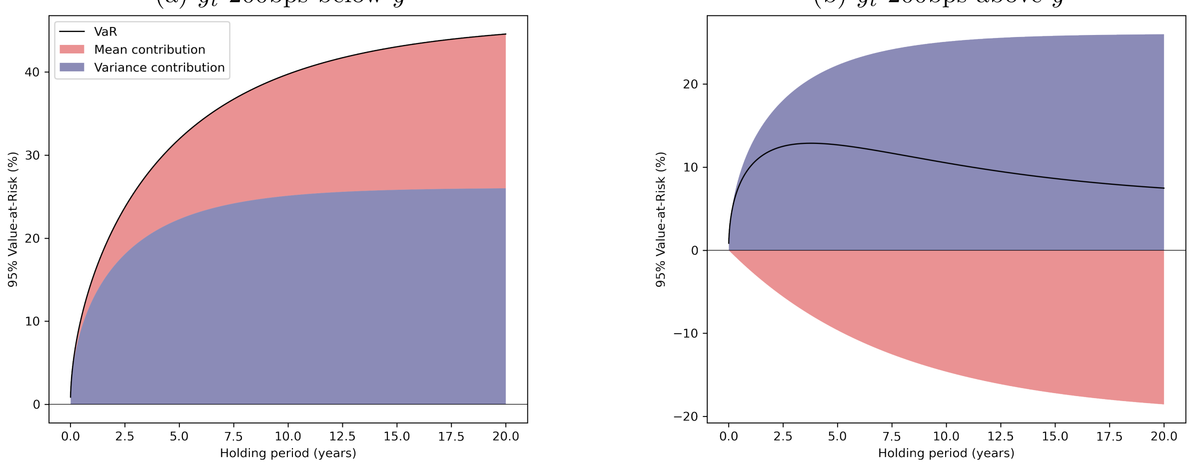

A.2 Sensitivity of constraints to projection horizons and changes in $\bar{y}_{20}$

As discussed, various risk-based capital and statutory reserving calculations require the expected shortfall as an input. Here I simulate loss distributions and 95% expected shortfall for a hypothetical insurance companies surplus, and assess its sensitivity to the regulators assumption about the long-run mean reversion point.

For these simulations, I assume the insurer has an initial surplus of \$100 million, and assets and liabilities with constant duration of 10 and 15 years, respectively. Specifically, for the loss distribution construction, I use the prescribed interest rate model and follow the NAIC's prescribed loss distribution construction methodology. Namely, for a T horizon projection, the value used in the loss distribution is the minimum point on the path, rather than the terminal point. For these simulations, I initialize the model with the prevailing 20-year and 1-year yields on 18th November, 2023, and simulated losses for 1 to 40 years out. The minimum projection horizon is the life of the policy in the case of statutory reserving, and a minimum of 20 years for risk-based capital calculations. While there is no maximum projection horizon.

Figure 10: Share of simulations with negative surplus (left) and ES95% (right)

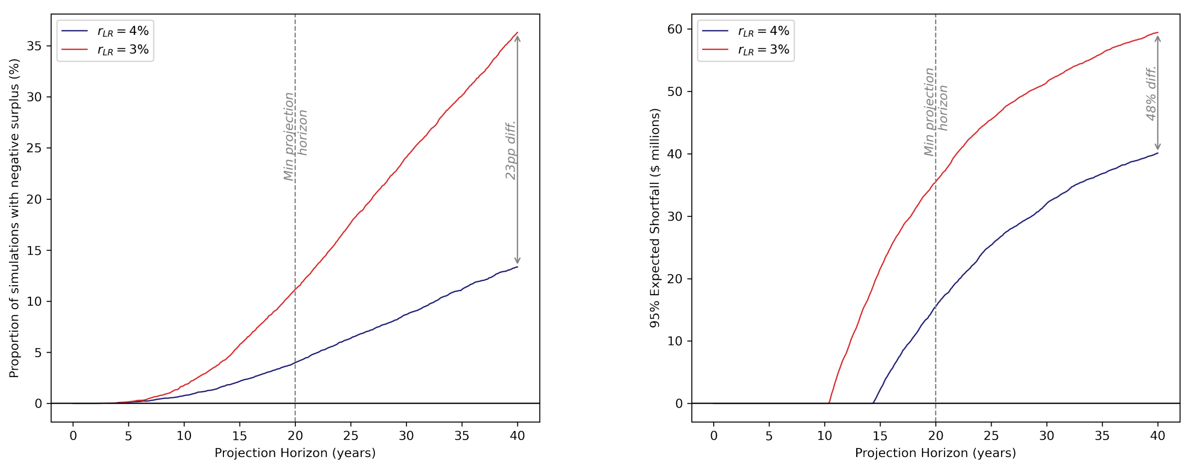

Figure 10 plots the results of these simulations for $\bar{y}_{20} = 4\%$ (blue) and $\bar{y}_{20} = 3\%$ (red) for different projection horizons. The figure on the left plots the share of simulations in which there was negative surplus. The share share of negative simulations is monotonically increasing in projection horizon, and lower $\bar{y}_{20}$ assumptions result in higher proportion of negative simulations. Notably at the 40 year projection horizon, the model assuming 3% long-term mean has a three times higher share of negative simulations than the model assuming 4% long-term mean. The plot on the right shows the 95% expected short fall calculation (this is required level in the RBC legislation). The results are mirrored here, the risk-based capital requirement is almost 50% higher in the model with the more conservative long-term mean assumption. Overall, capital requirements are highly sensitive to assumptions about the long-run mean, and their sensitivity is increasing for longer projection horizons.

A.3 Optimality of regulator learning rule

In this section, I explore whether the regulator's seemingly ad-hoc updates to the long-run mean can be reconciled with any optimal learning rules. As a reminder, the regulator uses a constant-gain learning algorithm, which is a linear combination of various backward-looking medians and means, and then rounds the result to the nearest 25 basis points:

$$\textbf{Round}_{25\text{bps}}\Big\{0.2 \times \text{Med}_{600m,h}\left(y_{20}\right) + 0.3 \times \text{Mean}_{120m,h}\left(y_{20}\right) + 0.5 \times \text{Mean}_{36m,h}\left(y_{20}\right)\Big\}$$

where $\text{Med}_{600m,h}\left(y_{20}\right)$ is a 600-month moving median, $\text{Mean}_{120,h}\left(y_{20}\right)$ is a 120-month moving average, and $\text{Mean}_{36m,h}\left(y_{20}\right)$ is a 36-month moving average.

To assess which rules are best suited to learn the long-run mean, we must take a stance on the data-generating process (DGP) of interest rates. A commonly assumed process is one where yields exhibit mean reversion, and the mean-reverting point evolves randomly over time:

$$dy_t = \theta(\mu_t - y_t)dt + \sigma_y dZ_t^y \qquad \text{(Observation yield)}$$

$$d\mu_t = \sigma_\mu dZ_t^w \qquad \text{(Unobserved mean)}$$

where dZ_t^y and dZ_t^w are standard Brownian motions that are mutually uncorrelated. I assume all parameters in the model are known and that y_t is observed, but the time-varying mean μ_t is unobservable.

The optimal update rule implied by the (steady-state) Kalman filter is a linear combination of the exponentially weighted moving average (EWMA) and the current yield:

$$\hat{\mu}_t = Ky_t + (1-K)\left[\nu \int_{-\infty}^{t} e^{-\nu(t-s)} y_s ds\right]$$

where $\nu := \theta K$, and $K = \sigma_\mu^2/\sigma_y^2$. In practice, the variance of changes in $\hat{\mu}_t$ is orders of magnitude smaller than the variance of changes in y_t, meaning $K \approx 0$. Hence, the optimal update rule of $\hat{\mu}_t$ is well-approximated by the EWMA; a constants gains learning rule.

Next, I assess whether the regulator's rule is well-approximated by an EWMA process. Figure 11 shows the realized path of 20-year yields (grey), the regulator's rule-implied long-run mean (blue), and the EWMA rule-implied mean (dashed red). The EWMA parameter is fitted to minimize the mean squared error between the EWMA-implied path and the regulator's rule-implied path. The close fit of the EWMA rule to the regulator's path suggests that, while the regulator's rule appears ad hoc, it is in practice incorporating past information in a somewhat optimal way.[28]

Figure 11: Regulator- vs EWMA- rule implied long-run mean

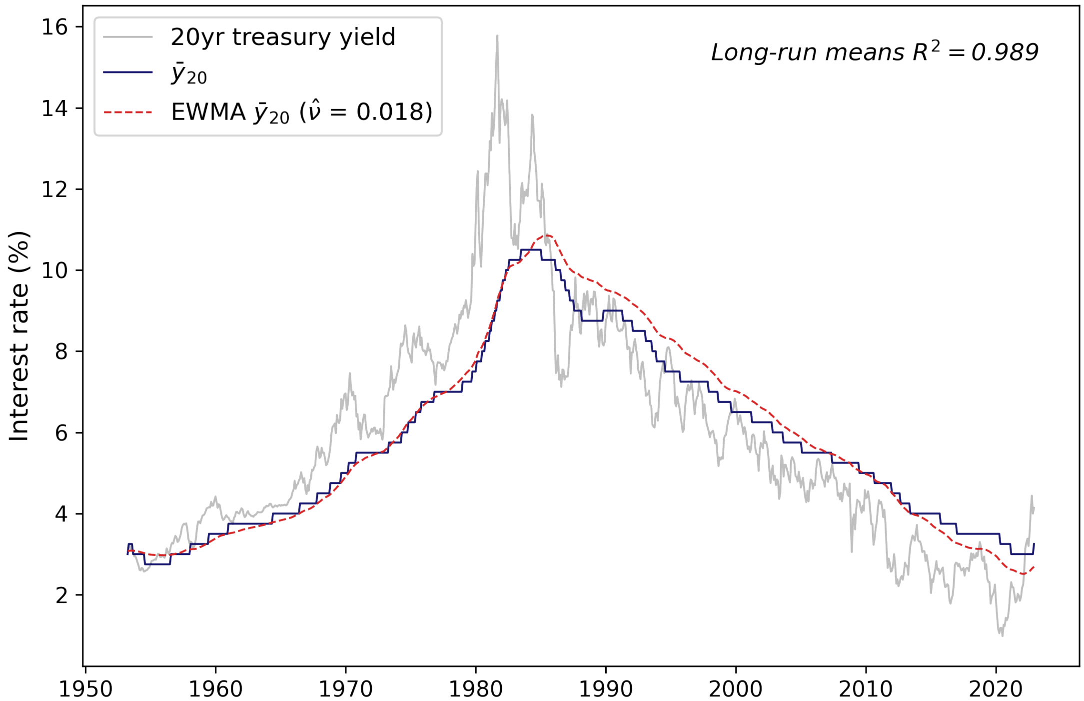

The figure shows the realized path of 20-year yields (grey), the regulator's rule-implied long-run mean (blue), and the EWMA rule-implied mean (dashed red), where the EWMA parameter is fitted to minimize the mean squared error between the EWMA-implied path and the regulator's rule-implied path.

Proof. Here I derive the steady state Kalman filter for estimating μ_t. The derivation finds the discrete-time Kalman filter for a small Δt approximation of the model, then takes the limit of $\Delta t \to 0$, and then finds the steady state of the updating process. With that in mind, for a small time interval Δt the model can be approximated as,

$$y_t = (1-\theta\Delta t)y_{t-1} + \theta\Delta t\mu_t + \varepsilon_t \qquad \text{(Observation equation)}$$

$$\mu_{t+1} = v_{t+1} \qquad \text{(State equation)}$$

where $\varepsilon_t \sim N(0, \sigma_y^2\Delta t)$ and $v_{t+1} \sim N(0, \sigma_\mu^2\Delta t)$. In this case the $\Delta t \to 0$ limit of the discrete time Kalman gain is given by,

$$\lim_{\Delta t\to 0} K_t = \lim_{\Delta t\to 0} \frac{P_{t|t-\Delta t}\theta}{\theta^2 P_{t|t-\Delta t}\Delta t + \sigma_y^2} = \frac{\theta}{\sigma_y^2}P_t$$

[28]Ideally, we would estimate the true parameters of the process described above to further assess whether the regulator has, in fact, adopted the optimal rule. However, it is notoriously difficult to accurately estimate the parameters of such processes in practice (Farmer et al., 2021).

To find the evolution of P_t note that,

$$\begin{aligned}
P_{t+\Delta t|t} =& P_{t|t-\Delta t} - K_t\theta\Delta t P_{t|t-\Delta t} + \sigma_\mu^2\Delta t \\
\frac{P_{t+\Delta t|t} - P_{t|t-\Delta t}}{\Delta t} =& - K_t\theta P_{t|t-\Delta t} + \sigma_\mu^2 + O(\Delta t) \\
\implies dP_t =& [-K_t\theta P_t + \sigma_\mu^2]dt && (\text{as } \Delta t \to 0)
\end{aligned}$$

In terms of the update rule we have,

$$\begin{aligned}
\hat{\mu}_{t+\Delta t|t} =& \hat{\mu}_{t|t-\Delta t} + K_t[y_t - (1-\theta\Delta t)y_{t-\Delta t} - \theta\Delta t\hat{\mu}_{t|t-\Delta t}] \\
\frac{\hat{\mu}_{t+\Delta t|t} - \mu_{t|t-\Delta t}}{\Delta t} =& K_t\frac{y_t - y_{t-\Delta t}}{\Delta t} + K_t\theta(y_{t-\Delta t} - \hat{\mu}_{t|t-\Delta t}) \\
\implies d\hat{\mu}_t =& K_t\theta(y_t - \hat{\mu}_t)dt + K_t dy && (\text{as } \Delta t \to 0)
\end{aligned}$$

In steady state (i.e. assuming $T_0 \to -\infty$) $dP_t = 0 \implies P = \frac{\sigma_\mu^2}{K\theta}$, Plugging this into K we get,

$$K = \frac{\theta}{\sigma_y^2}\frac{\sigma_\mu^2}{K\theta}$$

We can express the update rule $d\hat{\mu}_t = K_t\theta(y_t - \hat{\mu}_t)dt + K_t dy$ as:

$$\hat{\mu}_t = Ky_t + K\theta(1-K)\int_{-\infty}^{t} e^{-\theta K(t-s)}y_s ds$$

To see this note,

$$\begin{aligned}
d\hat{\mu}_t =& K\theta(y_t - \hat{\mu}_t)dt + Kdy \\
\iff d\hat{\mu}_t + \theta K\hat{\mu}_t dt =& \theta K y_t dt + Kdy \\
e^{\theta Kt}d\hat{\mu}_t + e^{\theta Kt}\theta K\hat{\mu}_t dt =& e^{\theta Kt}\theta K y_t dt + e^{\theta Kt}Kdy \\
d\left[e^{\theta Kt}\hat{\mu}_t\right] =& e^{\theta Kt}\theta K y_t dt + e^{\theta Kt}Kdy \\
\left|e^{\theta Ks}\hat{\mu}_s\right|_{t_0}^{t} =& \theta K\int_{t_0}^{t} e^{\theta Ks}y_s ds + \left|e^{\theta Kt}Ky_t\right|_{t_0}^{t} - \theta K^2\int_{t_0}^{t} e^{\theta Ks}y_s ds \\
e^{\theta Kt}\hat{\mu}_t =& e^{\theta Kt}Ky_t + K[\theta(1-K)]\int_{-\infty}^{t} e^{\theta Ks}y_s ds && (t_0 \to -\infty) \\
\hat{\mu}_t =& Ky_t + (1-K)\left[\nu\int_{-\infty}^{t} e^{-\nu(t-s)}y_s ds\right] && (\nu := K\theta)
\end{aligned}$$

□

B Insurer expectations data collection details

B.1 Insurer transcript data

Table 3 lists the publicly traded insurers whose transcripts I use to construct the insurer interest rate assumption dataset.

Table 3: Life insurer transcript sample

Firm	**Ticker**	**Number of transcripts**
Aegon	AEG	130
American Equity	AEL	103
American International Group	AIG	179
Athene Holding	ATH	39
Aviva	AVVIY	113
CNO Financial Group	CNO	125
China Life Insurance	CILJF	34
Citizens	CIA	25
Corebridge Financial	CRBG	9
Delphi Financial Group	DFG	49
Equitable Holdings	EQH	59
F&G Annuities & Life	FG	24
FBL Financial Group	FBL	90
Genworth Financial	GNW	135
Globe Life	GL	121
HRG Group	HRG	5
Horace Mann Educators Corp	HMN	106
ING N V	ING	156
Jackson Financial	JXN	14
Lincoln National	LNC	234
Manulife Financial	MFC	189
Mercury General New	MCY	74
Met Life	MET	223
Patriot National	PNBK	7
Primerica	PRI	75
Protective Life	PL	86
Prudential Financial	PRU	175
Prudential	PUK	111
Reinsurance Group Of America	RGA	108
Sun Life Financial	SLF	160
Symetra Financial	SYA	37
Voya Financial	VOYA	99
Total	—	3094

B.2 Additional information on Retrieval Augmented Generator

Figures 12 outlines the RAG process for constructing the vector database. This involves taking the raw transcript data, chunking statements, mapping them to embedding vectors and creating associated metadata. In terms of chunking, the transcripts are chunked to ensure that two different speakers don't overlap. Since embedding search works best when chunks are not too large, I use a chunking length of 300 characters (one-to-two sentence length). Chunking is done in a manner to ensure that sentences are not split in the middle, hence in the case of long sentences, chunks can be larger than 300 characters. In term of metadata I record: the conference call date, firm, speaker of the chunk, and role of the speaker e.g. CEO, CFO, equity analyst etc.

In terms of the embedding model, I use the Beijing Academy of Artificial Intelligence's M3 model. This is an open source model that is small enough to be run locally on CPU. And at the time of writing, this model outperforms OpenAI's embedding models, as well as various other popular models, on the Massive Text Embedding Benchmark (MTEB); the standard measure of embedding model performance. For each chunk, the model creates a 1024-dimensional embedding vector, which is then stored in a ChromaDB vector database for fast retrieval and comparison. With the vector database and embedding model in hand, semantic search can be implemented.

Figure 12: Vector database construction

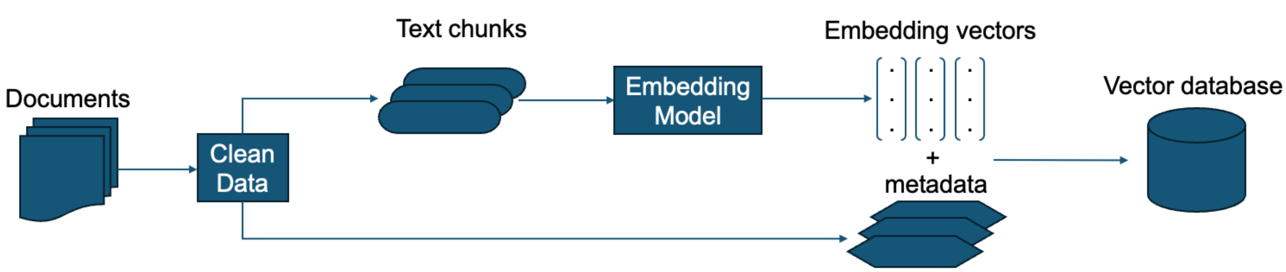

Figure 13 outlines the RAG process for answering queries. For retrieving statements, I loop over each firm-year. First, I subset to the firm-year statements made by firm employees. Then I do embedding search to retrieve the top five relevant contexts potentially containing information about interest rate assumptions. There are embedding search optimization techniques used to identify these five contexts, see below for details. The retrieved contexts are then combined into a prompt. The prompt contains: (i) top 5 retrieved context, (ii) query about mean-reversion point and time assumptions, and (iii) additional return instructions (see example boxes in the main text). The prompt is then sent to a generative LLM. For this purpose, I use GPT4, namely the "gpt-4-0125-preview" variant of the model, which was the highest performing publicly available OpenAI model available at the time of writing. The responses are then automatically converted into a dataframe, which I then manually verify.

Figure 13: Querying vector database

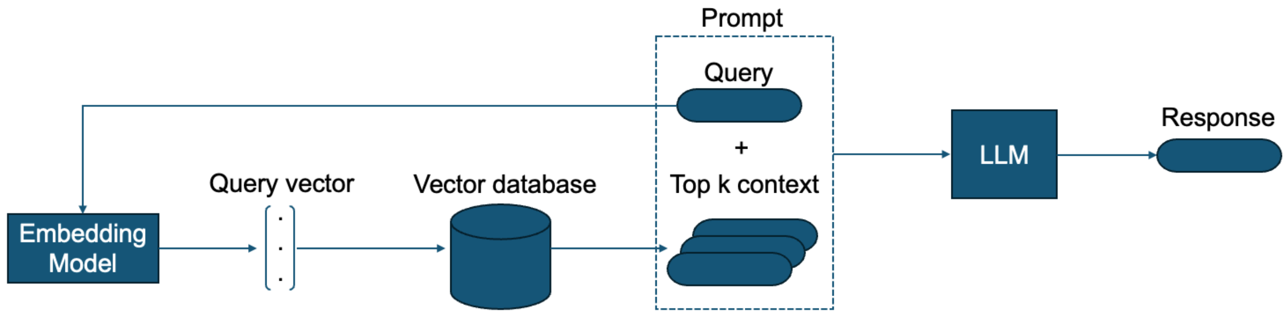

A plain vanilla retriever would simply embed the query and return the k nearest transcript statements in the embedding space. This simple approach has several shortcomings, such as, returning too many redundant statements, not finding all relevant statements in the embedding space, and returning statements without appropriate context to fully understand them. I implement several retrieval optimization techniques to alleviate these shortcomings,

- **Maximum marginal relevance (MMR):** For a given query embedding, if we simply return the nearest k neighbors around a embedded query, the retrieved statements will often contain the same information. Hence, providing more statements in this case will not necessarily provide more information to the LLM for answering the question. MMR is an algorithm that improves of the nearest neighbor approach, by additionally, penalizing the retrieval process if a statement is similar to the already retrieved statements (Carbonell and Goldstein, 1998). Specifically, MMR iteratively solves,

$$\arg \max_{s_i \in S \setminus R} \left[\lambda \cdot d\left(s_i, q\right) - (1-\lambda) \cdot \max_{s_j \in R} d\left(s_i, s_j\right) \right]$$

  where S is the set of all candidate statements, R is the set of already retrieved statements, s_i and s_j are embedding vectors of a statements in S and R, d measures the cosine similarity between two vectors, and λ is hyperparameter trading off the relevance of the statement to the query and its similarity with already retrieved statements. For each query, I first find the 60 nearest neighbors, $|S| = 60$, and then apply MMR to retrieve 20 of these statements.

- **Multi-query:** assumptions can be stated in many different ways, and hence, statements containing assumptions can end up in different parts of the embedding space. Multi-query is a technique to ensure we broadly search through the embedding space for relevant statements. This involves phrasing the same query in a different ways, as these queries may then get embedded onto different points in the embedding space, hopefully ensuring we capture as many possible statement clusters containing interest rate assumption information. Specifically, I search through the following five queries,

1. What is your long-term interest rate assumption?
2. How long will it take interest rates to revert to your long-term assumption?
3. What do you think is the long-term yield on 10-year US Treasuries?
4. What is your mean reversion assumption for long-term rates?
5. What rate does the forward curve eventually converge to?

Applying MMR to each of these queries gives us 100 statements (with possible repetition).

- **Reciprocal rank fusion (RRF):** This is an algorithm to rank the results from multi-query, and in turn select the most relevant ones to feed to the LLM (Cormack et al., 2009). Each of the five queries from the multi-query step gives 20 statements and associated MMR scores. We can use these MMR scores to create five sets of rankings; if a statement was not returned for one the queries then it is given a rank of ∞ for that query. RRF calculates the following RRFscore for each retrieved statement s_i,

$$\text{RRFscore}(s_i) = \sum_{q=1}^{5} \frac{1}{k + r_q(s_i)}$$

where $r_q(s_i)$ is the query specific ranking based on the MMR score, and k is a hyperparameter than ensures that low ranked statements get a quantitatively higher score than statements not returned for a given query.[29] After applying RRF to multi-query results, the retriever returns the top 5 ranked statements.

- **Parent-child retrieval:** Embedding search works best when the statements that are being embedded are not too long. This effectively requires chunking the transcript into statements that are no longer than a couple of sentences. Naturally, even though though small chunks sizes improve embedding search, they result in losing relevant context. Parent-child retrieval helps alleviate this context loss by additionally returning the sentences surrounding the identified statements.

[29] Since, a query can return at max 20 statements, for a returned statement the lowest possible contribution to RRFscore is $\frac{1}{20+k}$, whereas, if a statement is not returned as part of a query it is assigned $r_q = \infty \implies \frac{1}{\infty+k} = 0$.

C Model details

C.1 Set-up

Consider a two period economy in which the insurance company competes in a monopolistic insurance market. In period t, they set prices P_t and face a downward sloping demand curve $Q(P_t)$ with constant elasticity of demand $\varepsilon := \partial \log Q_t / \partial \log P_t$. When they sell Q_t policies it creates economic liabilities $L_t = Q_t V_t$, they need to also recognize statutory reserves $L_t^{statutory}$, and selling policies creates assets $A_t = P_t Q_t$ equal to the revenue from selling policies. The revenue is invested in short-term and long-term bonds with market values M_t and B_t, respectively. Accordingly, the period t balance sheet is,

$$\underbrace{M_t + B_t}_{:=A_t = P_t Q_t} = Q_t V_t + L_t^{statutory} + K_t$$

where K_t is the period t capital of the insurance company.

At period $t+1$, short-term bonds pay risk-free return $1+r_t^f$, there are no additional statutory reserves required as the policies payout, and the long-term bonds and the return on the economic values of liabilities is unknown and exposed to interest rate risk. I capture this risk, by modeling long-term bonds and liability return as a function of change in long-term yields $\Delta y_{l,t+1}$,

$$\frac{B_{t+1}}{B_t} = -D_B \Delta y_{l,t+1} \quad \text{and} \quad \frac{Q_t V_{t+1}}{Q_t V_t} = -D_L \Delta y_{l,t+1}$$

$\Delta y_{l,t+1}$ is not known known at time t, and D_B and D_L are the duration of long-term bonds and liabilities, respectively. The period $t+1$ value of capital is given by,

$$K_{t+1} = (K_t - L_t^{statutory})(1 + r_{t+1}^f) + B_t \underbrace{(-D_B \Delta y_{l,t+1} - r_{t+1}^f)}_{:=rx_{t+1}^B} - L_t \underbrace{(-D_L \Delta y_{l,t+1} - r_{t+1}^f)}_{:=rx_{t+1}^L}$$

where rx_{t+1}^B and rx_{t+1}^L are the excess returns on bonds and liabilities respectively.

Insurer and regulator model-implied beliefs: The insurer and regulators agree the variance of long-term yields is σ_y^2, but disagree about the mean, $\mathbb{E}^R[\Delta y_{l,t}] \neq \mathbb{E}^{*I}[\Delta y_{l,t}]$, where R and I superscript denote the regulator and insurer respectively. And the * superscript is added to distinguish the insurers latent expectations, from those effective beliefs discussed earlier. As a result of the disagreement about the change in yields, the regulator and insurer also disagree about the expected excess return on bonds and liabilities.

Statutory liabilities: When the insurer sells a policy, the regulator requires them to recognize an additional statutory liability $L_t^{statutory}$. Reflecting the regulator prescribed methodology, statutory liabilities are determined using the left tail of the capital loss distribution distribution, where the distribution is calculated using interest rate projections from the regulator's pre-calibrated interest rate model. In other words, the statutory liability is larger, if the insurer is expected to make large losses under regulator model-implied beliefs. I capture this feature of the of the regulation by modeling statutory liabilities as,

$$L_t^{statutory} = \frac{1}{1+r_{t+1}^f}\mathbb{E}^R[C^R(K_{t+1})]$$

where $C^R(\cdot) > 0$ is a strictly decreasing and convex cost function that puts more weight on large losses of K_{t+1}, and $\mathbb{E}^R[\cdot]$ evaluates the expectation under regulators model-implied beliefs about future yields.

Financial frictions: Financial frictions in raising external capital means that firm value depends on the distribution of future losses. Insurance policies promise payouts far into the future, which exposes internal finances to interest rate risk. Financial frictions make external financing costly, motivating insurers to mitigate the potential capital losses from interest rate fluctuations (Froot et al., 1993). Beyond external financing frictions, financial distress incurs additional costs, such as reputational damage and legal and administrative fees, further reducing the current value of the firm and encouraging the hedging of interest rate risk. These costs, stemming from financial frictions, are represented through a strictly decreasing and convex cost function $C^I(K_{t+1})$.

Firm value maximization: Given frictions, the period $t+1$ value of the firm is $J_{t+1} = K_{t+1} - C^I(K_{t+1})$. The insurer wants to choose policy prices and bond allocation in order to maximize the present value of the firm, $\mathbb{E}^I[\Gamma_{t+1}J_{t+1}]$ where Γ_{t+1} is the market SDF. Assuming the insurance company has no special advantage over the market in generating returns on bonds and their portfolio of policies, and assuming $L_t^{statutory}$ and $C^I(K_{t+1})$ are independent of Γ_{t+1}. The insurer's maximization problem is given by,

$$\mathbb{E}^I[\Gamma_{t+1}J_{t+1}(B_t^*, P_t^*)] = \max_{B_t, P_t} K_t - \frac{1}{1+r_{t+1}^f}\mathbb{E}^I[C^I(K_{t+1})] - \frac{1}{1+r_{t+1}^f}\mathbb{E}^R[C^R(K_{t+1})]$$

To elucidate the mechanisms, I work with a second-order approximation of this problem which allows for closed form solutions,[30]

$$\arg\max_{B_t,P_t} K_t - \frac{1}{1+r^f_{t+1}} c^I \left(\frac{\gamma^I}{2}\sigma^2_K - \mu^I_K\right) - \frac{1}{1+r^f_{t+1}} c^R \left(\frac{\gamma^R}{2}\sigma^2_K - \mu^R_K\right)$$

where $\mu^I_K = \mathbb{E}^I[K_{t+1}]$ and $\mu^I_R = \mathbb{E}^R[K_{t+1}]$ are the insurer and regulator's expectation of next period capital, and σ^2_K is their shared belief about capital volatility. The expected value and volatility of capital impact the firm value in opposite directions. On the one hand, higher variance makes large losses of K_{t+1} more likely, reducing firm value. On the other hand, higher expected future capital reduces the likelihood of extreme losses, in turn, offsetting the decline in firm value coming from higher volatility.

Given, the multi-year projection horizon, expectations play more prominent role than other moments (see appendix A). Hence, for expositional simplicity, in the main text I assume $\gamma := \gamma^R \approx \gamma^I$.

C.2 Solving the model

We want to solve for the optimal long-term bond allocation, B_t, and policy price, P_t,

$$\arg\max_{B_t,P_t} K_t - \frac{1}{1+r^f_{t+1}} c^I \left(\frac{\gamma^I}{2}\sigma^2_K - \mu^I_K\right) - \frac{1}{1+r^f_{t+1}} c^R \left(\frac{\gamma^R}{2}\sigma^2_K - \mu^R_K\right)$$

where,

- $K_t = (P_t - V_t)Q_t$
- $K_{t+1} = K_t(1 + r^f_{t+1}) + B_t(-D_B\Delta y_{t+1} - r^f_{t+1}) - Q_tV_t(-D_L\Delta y_{t+1} - r^f_{t+1})$
- $\mu^\tau_K = K_t(1 + r^f_{t+1}) + B_t\mathbb{E}^\tau[rx^B_{t+1}] - Q_tV_t\mathbb{E}^\tau[rx^V_{t+1}]$ for $\tau \in \{*I, R\}$
- $\sigma^2_K = (B_tD_B - D_LQ_tV_t)^2\sigma^2_y$

[30] See appendix section C.3 for details on the approximation.

Solving for B_t: *Take derivatives of the objective function with respect to B_t,*

$$
\begin{aligned}
0 =& \frac{1}{1+r^f_{t+1}} \sum_{\tau\in\{I,R\}} c^\tau \left(\frac{\partial \mu^\tau_K}{\partial B_t} - \frac{\gamma^\tau}{2} \frac{\partial \sigma^2_K}{\partial B_t} \right) \\
0 =& \alpha \frac{\partial \mu^R_K}{\partial B_t} + (1-\alpha) \frac{\partial \mu^I_K}{\partial B_t} - \frac{\tilde{\gamma}}{2} \frac{\partial \sigma^2_K}{\partial B_t} && (\alpha := \tfrac{c^R}{c^R+c^I};\ \tilde{\gamma} := \alpha\gamma^R + (1-\alpha)\gamma^I) \\
0 =& \mathbb{E}^I[rx^B_{t+1}] + \alpha(\mathbb{E}^R[rx^B_{t+1}] - \mathbb{E}^I[rx^B_{t+1}]) - \tilde{\gamma}(B_t D_B^2 - D_B D_L Q_t V_t)\sigma^2_y \\
B_t =& \frac{\mathbb{E}^I[rx^B_{t+1}] + \alpha(\mathbb{E}^R[rx^B_{t+1}] - \mathbb{E}^I[rx^B_{t+1}])}{\tilde{\gamma}\sigma^2_B} + \frac{D_L}{D_B} \underbrace{Q_t V_t}_{=L_t} && (\sigma^2_B := D^2_B \sigma^2_y)
\end{aligned}
$$

where $\lambda = K_t/A_t$ is the insurer's leverage.

Solving for P_t: *Since this is a bit algebraically involved, it is helpful to write-out and simplify some derivatives before hand,*

- $P_t Q'_t/Q_t = -\varepsilon \implies Q'_t = -\varepsilon Q_t/P_t$
- $\frac{\partial K_t}{\partial P_t} = Q'_t P_t + Q_t - Q'_t V_t = \frac{\varepsilon Q_t}{P_t}\left[V_t - \left(1 - \frac{1}{\varepsilon}\right) P_t\right]$
- $\frac{\partial \mu^\tau_K}{\partial P_t} = (1+r^f_{t+1})\frac{\partial K_t}{\partial P_t} + V_t \mathbb{E}^\tau[rx^L_{t+1}]\varepsilon Q_t/P_t$
- $\frac{\partial \sigma^2_K}{\partial P_t} = 2[D_B B_t - D_L V_t Q_t] V_t D_L \sigma^2_y \varepsilon Q_t/P_t$ *which when evaluated for equilibrium B_t is* $\frac{\partial \sigma^2_K}{\partial P_t} = 2 D_B B_{MV,t} V_t D_L \sigma^2_y \varepsilon Q_t/P_t$ *where* $B_{MV,t} := \frac{\mathbb{E}^I[rx^B_{t+1}]}{\tilde{\gamma}\sigma^2_B}$

Taking derivatives of the objective with respect to price,

$$
\begin{aligned}
0 =& \frac{\partial K_t}{\partial P_t} + \frac{1}{1+r^f_{t+1}} \sum_{\tau\in\{I,R\}} c^\tau \left(\frac{\partial \mu^\tau_K}{\partial P_t} - \frac{\gamma^\tau}{2} \frac{\partial \sigma^2_K}{\partial P_t} \right) \\
0 =& \frac{1+C}{C}(1+r^f_{t+1}) \frac{\partial K_t}{\partial P_t} + V_t \mathbb{E}^I[rx^L_{t+1}]\varepsilon Q_t/P_t - \frac{\tilde{\gamma}}{2}\frac{\partial \sigma^2_K}{\partial P_t} \\
& (C := c^R + c^I \text{ and } \tilde{\gamma} := \alpha\gamma^R + (1-\alpha)\gamma^I)
\end{aligned}
$$

We can simplify,

$$
V_t \mathbb{E}^I \left[rx^L_{t+1} - rx^B_{t+1} \frac{D_L}{D_B} \right] = V_t r^f_{t+1} \left(\frac{D_L}{D_B} - 1 \right)
$$

Note that because of mean-variance optimization the contribution from the mean exactly cancels out in equilibrium. This result echos Verani and Yu (2021) finding that without frictions restricting optimal investment in long-term bonds, fluctuations in interest rates do not impact *insurance policy prices. Substituting this back in,*

$$\implies P_t = \left(1 - \frac{1}{\varepsilon}\right)^{-1} \left(1 + \frac{Cr^f_{t+1}}{(1+C)(1+r^f_{t+1})}\left[\frac{D_L}{D_B} - 1\right]\right) V_t$$

Prices equal the marginal cost of issuing policies times a mark-up. In equilibrium, there are two contributors to the marginal cost,

1. The economic liability V_t created by selling the policy.
2. The net cost of hedging the economic liability. For each dollar of liability created the insurer needs to buy $\frac{D_L}{D_B}$ bonds, if $D_L > D_B$ then creating the hedge is costly. The overall contribution of the hedge cost to the price depends on the cost of frictions, which incentivize hedging in the first place. In the absence of frictions $C = 0$, and this second marginal cost disappears.

C.3 Approximating loss distribution based constraints

Suppose an investor chooses a control θ, which impacts their stochastic dollar payoff $X(\theta)$. Denote the mean and variance of $X(\theta)$ as $\mu_x(\theta)$ and $\sigma^2_x(\theta)$ respectively. Now, consider a loss distribution based statistic,

$$S(\theta) = \mathbb{E}[f(X(\theta))]$$

where $f(\cdot)$ is strictly positive, strictly decreasing, and strictly convex. In other words, $f(\cdot)$ puts increasingly more weight on losses of X_{t+1}; it can be viewed as the "disuitily" of losses. The second order approximation of the loss statistic around $\mu_x(\theta)$ is given by,

$$S(\theta) \approx f(\mu_x(\theta)) + \frac{f''(\mu_x(\theta))}{2}\sigma^2_x(\theta)$$

Indeed commonly used loss distribution statistics, such as Value-at-Risk and Expected Shortfall, for a large class of distributions—e.g. normal, log-normal, student-t and exponential—are fully characterized as functions of the mean and variance.

Suppose an investor wants to maximize utility subject to a loss based constraint,

$$\theta^* = \arg\max_{\theta} U(\theta) - \lambda\left[\nu - S(\theta)\right]$$

The approximate solution to this program is implicitly characterized by,

$$U'(\theta^*) + \lambda \underbrace{[-f'(\mu_x(\theta^*))]}_{:=c^S>0} \left[\underbrace{\frac{f''(\mu_x(\theta^*))}{-f'(\mu_x(\theta^*))}}_{:=\gamma^S} \sigma_x(\theta^*)\sigma_x'(\theta^*) - \mu_x'(\theta^*) \right] = 0$$

where, c^S is the "marginal cost" of higher variance increasing the loss based statistic, and γ is the relative offsetting effect of a higher expected mean. Notice, that if in the optimization we replace $S(\theta)$ with the function $c^S \left[\frac{\gamma^S}{2}\sigma_x^2(\theta) - \mu_x(\theta)\right]$, then it will admit the same θ^* (treating c^S and γ^S as constants in the optimization step).

D Additional empirical results

D.1 Systematic patterns in regulator beliefs additional results

Table 4: 10-year yield average expectation error for different samples

		Horizon							
Expectation	Sample	1Q	2Q	3Q	1yr	2yr	3yr	4yr	5yr
Regulator model	1988-2021	-0.06**	-0.17***	-0.28***	-0.39***	-0.63***	-0.94***	-1.22***	-1.48***
		(0.03)	(0.06)	(0.07)	(0.08)	(0.11)	(0.12)	(0.12)	(0.12)
	2009-2021	-0.07	-0.16*	-0.26**	-0.35***	-0.45**	-0.69***	-0.78***	-0.88***
		(0.05)	(0.09)	(0.11)	(0.13)	(0.2)	(0.18)	(0.13)	(0.16)
Insurer	1988-2021	-0.23***	-0.35***	-0.47***	-0.58***	-0.96***	-1.39***	—	—
		(0.03)	(0.06)	(0.08)	(0.1)	(0.12)	(0.11)	—	—
	2009-2021	-0.19***	-0.31***	-0.42***	-0.53***	-0.81***	-1.25***	—	—
		(0.06)	(0.1)	(0.13)	(0.15)	(0.26)	(0.24)	—	—
Blue Chip average	1988-2021	-0.2***	-0.34***	-0.47***	-0.61***	—	—	—	—
		(0.05)	(0.08)	(0.09)	(0.1)	—	—	—	—
	2009-2021	-0.21***	-0.38***	-0.57***	-0.76***	—	—	—	—
		(0.07)	(0.11)	(0.13)	(0.14)	—	—	—	—

Standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table plots the 10yr yield average expectation error for $h = 1Q, 2Q, 3Q, 1yr, 2yr, 3yr, 4yr, 5yr$ horizons for regulators, insurers, Blue Chip forecasters, and SPF forecasters—insurer forecasts are based on the forward curve. The expectations are a panel of quarterly frequency expectations for sample periods: (i) 1988 to 2021, and (ii) 2009 to 2021. The standard errors in the parenthesis are calculated using Newey-West errors with 1-lag.

Table 5: 10-year yield average expectation error for different samples (HAC-4)

Expectation	Sample	Horizon							
		1Q	2Q	3Q	1yr	2yr	3yr	4yr	5yr
Regulator model	1988-2021	-0.06*	-0.17***	-0.28***	-0.39***	-0.63***	-0.94***	-1.22***	-1.48***
		(0.04)	(0.07)	(0.08)	(0.09)	(0.13)	(0.14)	(0.14)	(0.14)
	2009-2021	-0.07	-0.16*	-0.26**	-0.35**	-0.45*	-0.69***	-0.78***	-0.88***
		(0.06)	(0.10)	(0.12)	(0.15)	(0.22)	(0.20)	(0.15)	(0.18)
Insurer	1988-2021	-0.23***	-0.35***	-0.47***	-0.58***	-0.96***	-1.39***	—	—
		(0.04)	(0.07)	(0.09)	(0.11)	(0.14)	(0.13)	—	—
	2009-2021	-0.19***	-0.31***	-0.42***	-0.53***	-0.81***	-1.25***	—	—
		(0.07)	(0.11)	(0.15)	(0.17)	(0.29)	(0.27)	—	—
Blue Chip average	1988-2021	-0.2***	-0.34***	-0.47***	-0.61***	—	—	—	—
		(0.06)	(0.09)	(0.10)	(0.12)	—	—	—	—
	2009-2021	-0.21***	-0.38***	-0.57***	-0.76***	—	—	—	—
		(0.08)	(0.12)	(0.15)	(0.16)	—	—	—	—

HAC standard errors with 4 lags in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table plots the 10yr yield average expectation error for $h = 1Q, 2Q, 3Q, 1yr, 2yr, 3yr, 4yr, 5yr$ horizons for regulators, insurers, and Blue Chip forecasters. The expectations are a panel of quarterly frequency expectations for sample periods: (i) 1988 to 2021, and (ii) 2009 to 2021. The standard errors in the parenthesis are calculated using Newey-West HAC standard errors with 4 lags for robustness to serial correlation and heteroskedasticity.

Table 6: Coibion-Gorodnichenko coefficient

Expectation	Sample	Maturity 3M	1Y	2Y	5Y	10Y
Regulator model	1988-2021	0.45***	0.31***	0.17**	-0.04	-0.13**
		(0.09)	(0.09)	(0.08)	(0.06)	(0.06)
	2009-2021	0.39**	0.47**	0.36**	0.08	-0.03
		(0.18)	(0.2)	(0.15)	(0.09)	(0.07)
Insurer	1988-2021	0.11	0.09	0.01	-0.09*	-0.1*
		(0.07)	(0.06)	(0.06)	(0.05)	(0.05)
	2009-2021	0.1	0.12	0.03	-0.05	-0.11
		(0.07)	(0.09)	(0.09)	(0.08)	(0.07)
Blue Chip average	1988-2021	0.57***	0.45***	0.25***	0.02	-0.04
		(0.09)	(0.09)	(0.09)	(0.1)	(0.1)
	2009-2021	0.29*	0.32**	0.17	-0.06	-0.13
		(0.17)	(0.15)	(0.12)	(0.11)	(0.12)

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

The table shows the Coibion-Gorodnichenko (CG) coefficient for different maturities, comparing expectations of the regulator, insurers, BCFF, and SPF. The CG regression is given by: $e^n_{t+1Q|t} = \alpha^n + \beta^n(\mathbb{E}_t[y^n_{t+1Q}] - \mathbb{E}_{t-1Q}[y^n_{t+1Q}]) + \varepsilon^n_{t+1Q}$ where $e^n_{t+1Q|t}$ is the expectation error, and β^n is the CG coefficient where $\beta^n > 0$ implies underreaction. The expectations are a panel of quarterly frequency expectations for sample periods: (i) 1988 to 2021, and (ii) 2009 to 2021. The standard errors are calculated using Newey-West errors with 1-lag.

Table 7: Coibion-Gorodnichenko coefficient (HAC-4)

Expectation	Sample	Maturity 3M	1Y	2Y	5Y	10Y
Regulator model	1988-2021	0.45***	0.31***	0.17*	-0.04	-0.13*
		(0.10)	(0.10)	(0.09)	(0.07)	(0.07)
	2009-2021	0.39*	0.47**	0.36**	0.08	-0.03
		(0.20)	(0.22)	(0.17)	(0.10)	(0.08)
Insurer	1988-2021	0.11	0.09	0.01	-0.09	-0.1
		(0.08)	(0.07)	(0.07)	(0.06)	(0.06)
	2009-2021	0.1	0.12	0.03	-0.05	-0.11
		(0.08)	(0.10)	(0.10)	(0.09)	(0.08)
Blue Chip average	1988-2021	0.57***	0.45***	0.25**	0.02	-0.04
		(0.10)	(0.10)	(0.10)	(0.11)	(0.11)
	2009-2021	0.29	0.32*	0.17	-0.06	-0.13
		(0.19)	(0.17)	(0.14)	(0.12)	(0.13)

HAC standard errors with 4 lags in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table shows the Coibion-Gorodnichenko (CG) coefficient for different maturities, comparing expectations of the regulator, insurers, and Blue Chip forecasters. The CG regression is given by: $e^n_{t+1Q|t} = \alpha^n + \beta^n(\mathbb{E}_t[y^n_{t+1Q}] - \mathbb{E}_{t-1Q}[y^n_{t+1Q}]) + \varepsilon^n_{t+1Q}$ where $e^n_{t+1Q|t}$ is the expectation error, and β^n is the CG coefficient where $\beta^n > 0$ implies underreaction. The expectations are a panel of quarterly frequency expectations for sample periods: (i) 1988 to 2021, and (ii) 2009 to 2021. The standard errors are calculated using Newey-West HAC standard errors with 4 lags for robustness to serial correlation and heteroskedasticity.

D.1.1 Drivers of patterns in regulator expectations

A key benefit of using model-elicited beliefs is the ability to gain insight into the drivers of the observed patterns in regulator beliefs. To precisely answer this question, we would need access to the true data-generating process (DGP) for interest rates in order to see how the regulator's DGP deviates from it. Since we do not have access to the true DGP, I adopt a DGP-agnostic approach. In this section, I assess the sensitivity of the observed patterns in regulator beliefs to changes in the parameters of the regulator's model.

I conduct a sensitivity analysis on observed patterns to several key parameters: the long-run mean of log long-term yields ($\bar{y}^{(h)}_{20}$), the long-run mean of the slope ($\bar{s}$), the mean reversion rates of long-term yields (β_1), and the mean reversion rates of the slope in both the long-rate (ψ) and slope equations (β_2). For each parameter, I increase its value by 10%, reconstruct the regulator's beliefs, and then re-estimate the three patterns: (i) forecast errors, and (ii) over/underreaction.

Figure 14 shows that expectation errors are most sensitive to changes in the long-run mean of interest rates, suggesting that these errors are primarily driven by the regulator's failure to anticipate the decline in yields. Figure 15 presents the sensitivity of the Coibion and Gorodnichenko (CG) regression coefficient for 3-month yields, showing that CG behavior is predominantly driven by the model's persistence parameters.

Figure 14: Regulator sensitivity of 5-year ahead expectation error to parameters

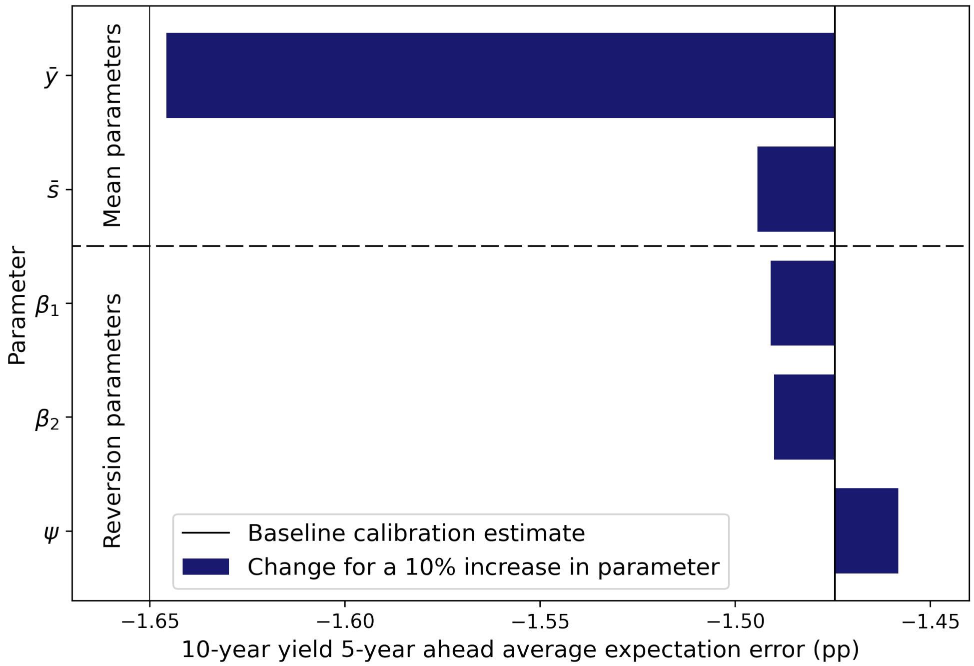

Figure 15: Regulator sensitivity of 1Q ahead CG coefficient to parameters

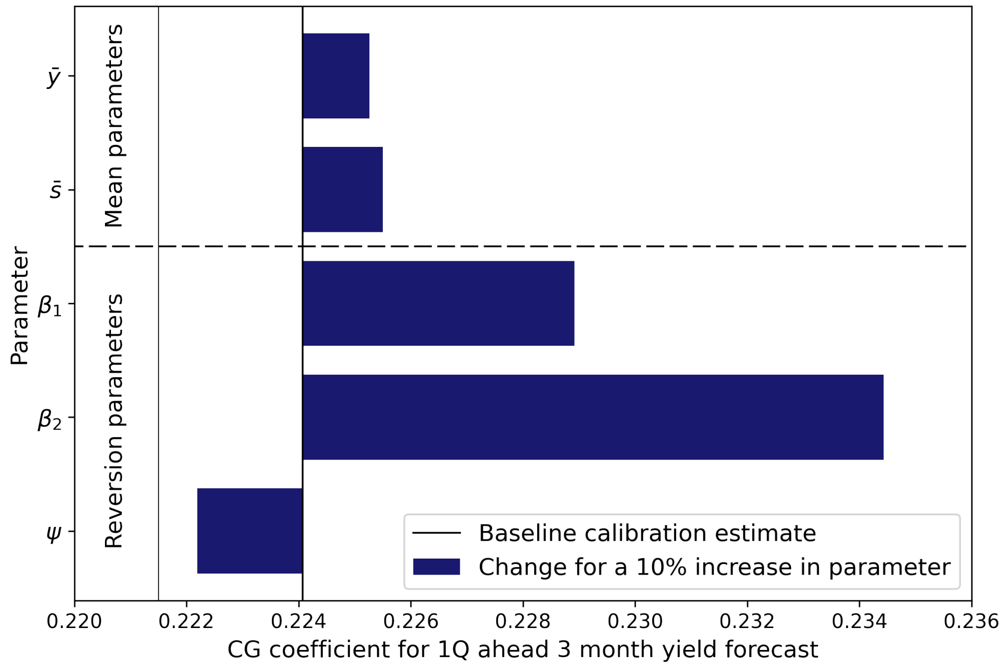

D.2 Difference-in-difference estimand derivation

Here I derive the model-based DD estimand β_t^{VM20}. In terms of notation, $\mathbb{E}^{(i)}[\cdot]$ denotes the cross-sectional average across insurers. Additionally, I follow the potential outcomes framework in denoting the actual treatment status with a superscript 1 for being subject to VM20, and 0 for not being subject to to VM20.

$$
\begin{aligned}
\beta_t^{VM20} =& \mathbb{E}^{(i)}[B_{i,t}^1 - B_{i,2016}^0 | VM20 = 1] - \mathbb{E}^{(i)}[B_{i,t}^0 - B_{i,2016}^0 | VM20 = 0] \\
=& \mathbb{E}^{(i)}[B_{i,t}^1 - B_{i,2016}^0 | VM20 = 1] - \mathbb{E}^{(i)}[B_{i,t}^0 - B_{i,2016}^0 | VM20 = 1] \quad \text{(common trends ass.)} \\
=& \mathbb{E}^{(i)}[B_{i,t}^1 - B_{i,t}^0 | VM20 = 1] \\
=& \mathbb{E}^{(i)}\Bigg[\frac{\mathbb{E}_{i,t}^I[rx_{t+1}^B] + \alpha_{i,t}^{VM20}(\mathbb{E}_t^R[rx_{t+1}^B] - \mathbb{E}_{i,t}^I[rx_{t+1}^B])}{\gamma_{i,t}\sigma_B^2} + \frac{D_L}{D_B}L_{i,t} \\
& - \left(\frac{\mathbb{E}_{i,t}^I[rx_{t+1}^B]}{\gamma_{i,t}\sigma_B^2} + \frac{D_L}{D_B}L_{i,t}\right)\Bigg| VM20 = 1\Bigg] \quad \text{(Substitute in model)} \\
=& \mathbb{E}^{(i)}\left[\alpha_{i,t}^{VM20}\frac{\mathbb{E}_t^R[rx_{t+1}^B] - \mathbb{E}_{i,t}^I[rx_{t+1}^B]}{\gamma_i\sigma_B^2}\Bigg| VM20 = 1\right]
\end{aligned}
$$

D.3 VM20 difference-in-difference estimates

This subsection assesses the robustness of the difference-in-differences results to alternative definitions of the treatment group, potential confounding factors, and inference procedures. Table

8 reports the baseline estimates in column (1), and various robustness checks in columns (2) through (7).

Table 8: Difference-in-Difference: VM20 impact on long-term bond holdings

	Baseline	Robustness					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
α	8.40***	8.43***	8.46***	8.04***	8.37***	8.55***	8.68***
	(0.18)	(0.17)	(0.17)	(0.20)	(0.17)	(0.20)	(0.19)
$\delta_{i\in LI}$	-1.67***	-1.76***	-1.96***	-1.18***	-1.88***	-1.30***	-1.48***
	(0.39)	(0.38)	(0.39)	(0.40)	(0.43)	(0.31)	(0.43)
VA Share$_{i,t}$	1.18*	1.06	1.16*	—	1.23*	0.98	0.02
	(0.70)	(0.65)	(0.68)	—	(0.69)	(0.73)	(0.74)
δ_{2011}	-0.80***	-0.86***	-0.96***	-0.53***	-0.79***	-0.83***	-0.84***
	(0.06)	(0.06)	(0.06)	(0.08)	(0.06)	(0.08)	(0.06)
δ_{2012}	-0.24***	-0.30***	-0.32***	-0.08	-0.26***	-0.27***	-0.33***
	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)	(0.07)	(0.05)
δ_{2013}	0.01	0.02	0.00	0.14*	-0.03	0.04	0.00
	(0.06)	(0.05)	(0.05)	(0.08)	(0.06)	(0.07)	(0.05)
δ_{2014}	-0.03	-0.02	-0.04	0.07	-0.06	-0.02	-0.10**
	(0.04)	(0.04)	(0.04)	(0.06)	(0.04)	(0.05)	(0.04)
δ_{2015}	-0.03	-0.00	-0.02	0.02	-0.03	-0.04	-0.13***
	(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	(0.03)	(0.02)
δ_{2016}	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	—	—	—	—	—	—	—
δ_{2017}	0.13***	0.14***	0.16***	0.16***	0.18***	0.13***	-0.13***
	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.03)	(0.02)
δ_{2018}	0.08**	0.11***	0.09**	0.04	0.10***	0.10***	-0.18***
	(0.03)	(0.03)	(0.03)	(0.04)	(0.03)	(0.04)	(0.03)
δ_{2019}	0.08*	0.11**	0.07	-0.05	0.13***	0.11**	-0.17***
	(0.04)	(0.04)	(0.04)	(0.06)	(0.04)	(0.05)	(0.04)
δ_{2020}	0.35***	0.39***	0.34***	0.20***	0.38***	0.38***	0.04
	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)	(0.06)	(0.05)
δ_{2021}	0.59***	0.63***	0.59***	0.42***	0.59***	0.66***	0.33***
	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)	(0.06)	(0.05)
δ_{2022}	0.41***	0.38***	0.34***	0.15**	0.42***	0.45***	0.16***
	(0.05)	(0.05)	(0.05)	(0.07)	(0.05)	(0.06)	(0.05)
$\delta_{2011} \times \delta_{i\in LI}$	0.24*	0.28*	0.41**	-0.09	0.26*	0.22**	0.42***
	(0.14)	(0.15)	(0.16)	(0.15)	(0.15)	(0.10)	(0.12)
$\delta_{2012} \times \delta_{i\in LI}$	0.21*	0.23*	0.26*	0.04	0.37***	0.21**	0.20*
	(0.12)	(0.14)	(0.15)	(0.14)	(0.13)	(0.09)	(0.11)
$\delta_{2013} \times \delta_{i\in LI}$	-0.17	-0.20	-0.15	-0.32**	-0.03	-0.18*	-0.01
	(0.13)	(0.14)	(0.15)	(0.15)	(0.15)	(0.10)	(0.11)
$\delta_{2014} \times \delta_{i\in LI}$	-0.10	-0.16	-0.05	-0.22	0.06	-0.06	0.18
	(0.15)	(0.13)	(0.14)	(0.17)	(0.17)	(0.09)	(0.13)
$\delta_{2015} \times \delta_{i\in LI}$	0.19*	0.20**	0.21**	0.16	0.24*	0.13**	0.23**
	(0.10)	(0.10)	(0.10)	(0.12)	(0.13)	(0.06)	(0.09)
$\delta_{2016} \times \delta_{i\in LI}$	0.00	0.00	0.00	0.00	0.00	0.00	0.00
	—	—	—	—	—	—	—
$\delta_{2017} \times \delta_{i\in LI}$	0.11**	0.02	-0.12**	0.11*	-0.12*	0.05	0.16***
	(0.05)	(0.05)	(0.06)	(0.06)	(0.06)	(0.04)	(0.06)
$\delta_{2018} \times \delta_{i\in LI}$	0.57***	0.33***	0.40***	0.68***	0.53***	0.26***	0.24***
	(0.10)	(0.09)	(0.09)	(0.11)	(0.11)	(0.07)	(0.07)
$\delta_{2019} \times \delta_{i\in LI}$	0.55***	0.25*	0.39***	0.79***	0.33*	0.29***	0.24**
	(0.15)	(0.15)	(0.15)	(0.18)	(0.19)	(0.10)	(0.11)
$\delta_{2020} \times \delta_{i\in LI}$	0.89***	0.45**	0.70***	1.17***	0.84***	0.44***	0.60***
	(0.23)	(0.22)	(0.23)	(0.27)	(0.29)	(0.13)	(0.20)
$\delta_{2021} \times \delta_{i\in LI}$	1.21***	0.77***	0.97***	1.53***	1.43***	0.52***	0.96***
	(0.24)	(0.23)	(0.25)	(0.28)	(0.32)	(0.13)	(0.22)
$\delta_{2022} \times \delta_{i\in LI}$	1.00***	0.80***	1.03***	1.21***	1.15***	0.46***	0.80***
	(0.21)	(0.21)	(0.22)	(0.25)	(0.27)	(0.13)	(0.20)
Observations	3,692	3,721	3,709	2,617	3,692	3,692	3,215
VA dropped	No	No	No	Yes	No	No	No
$\delta_t \times$TCJA exposure	No	No	No	No	No	No	Yes
$\delta_{i\in LI}$ definition year	2013	2012	2011	2013	2013	2013	2013
$\delta_{i\in LI}$ threshold	90%	90%	90%	90%	95%	70%	90%

Clustered standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table shows estimates from the difference-in-difference regression, all standard errors cluster at the insurer and year. Column 1 presents the baseline estimates, where treatment group are insurers with more 90% traditional life reserves in 2013, and the control group are all other insurers. Column 2 and 3 use 2012 and 2011 to define the treatment group. Column 4 drops insurance companies with variable annuity business. Column 5 and 6, define the treatment group as having more than 95% and 70% traditional life reserves in 2013, respectively. Column 7, controls for insurers' exposure to Tax Cuts and Jobs Act (TCJA) tax reforms.

Year used to identify traditional life insurers: In the baseline specification, I define traditional life insurers as those with more than 90% of reserves in traditional life policies in 2013. Columns (2) and (3) show that the results are essentially unchanged when reserves in 2012 or 2011 are used instead. This is consistent with product specialization being a relatively persistent characteristic of these firms.

Controlling for variable annuity exposure: The baseline specification controls for the share of variable annuity reserves to ensure comparability between treatment and control groups. Column (4) shows that the results are robust to an alternative specification that simply excludes insurers offering variable annuity products.

Threshold for defining traditional life insurers: The baseline uses a 90% threshold of traditional life reserves to identify firms specializing in traditional life insurance. Columns (5) and (6) use alternative thresholds of 95% and 70%, respectively. The overall pattern in duration remains similar to the baseline, with slightly larger estimates at the 95% threshold and slightly smaller ones at 70%, as expected given the degree of exposure to VM-20 regulation is positively related to the threshold.

Exposure to Tax Cuts and Jobs Act (TCJA) reforms: A potential concern is that the results may be confounded by the Tax Cuts and Jobs Act (TCJA), which went into effect in 2018—the year following the first phase of VM20 implementation. The Society of Actuaries (SOA) assessed the impact of the TCJA on life insurance product pricing and profitability, finding that most product types were only modestly affected, with the exception of term life products using captive reinsurance vehicles, which experienced a substantial decline in profit margins under the new tax regime (Clingerman et al., 2018).

If traditional life insurers were more likely to use captive reinsurance, and if such products tend to have lower duration, the tax reform could have induced these insurers to shift toward longer-duration products (and assets), potentially confounding the baseline results. To address this concern, I measure each insurer's exposure to the TCJA using its use of captive reinsurance vehicles and control for this exposure interacted with year dummies. Specifically, I follow Koijen and Yogo (2016) in constructing a shadow insurance measure that captures captive use. Column (7) includes these interactions, and the estimates remain quantitatively similar to the baseline specification. Consistent with this, I find that traditional life insurers are, on average, *less* likely to use captive reinsurance than non-traditional life insurers.

Placebo inference: Finally, to assess whether the observed treatment effects could arise by chance, I perform a placebo inference exercise following the spirit of a Fisherian sharp null. I randomly assign treatment status across insurers 10,000 times and re-estimate the difference-in-differences specification for each permutation. Figure 16 plots the distribution of placebo treatment effects and the actual 2022 estimate. The observed treatment effect lies in the right tail of the placebo distribution, corresponding to a p-value of 0.2%.

Figure 16: Plot of β_t^{VM20} against 10,000 $\beta_t^{VM20,placebo}$ which randomly assigning treatment

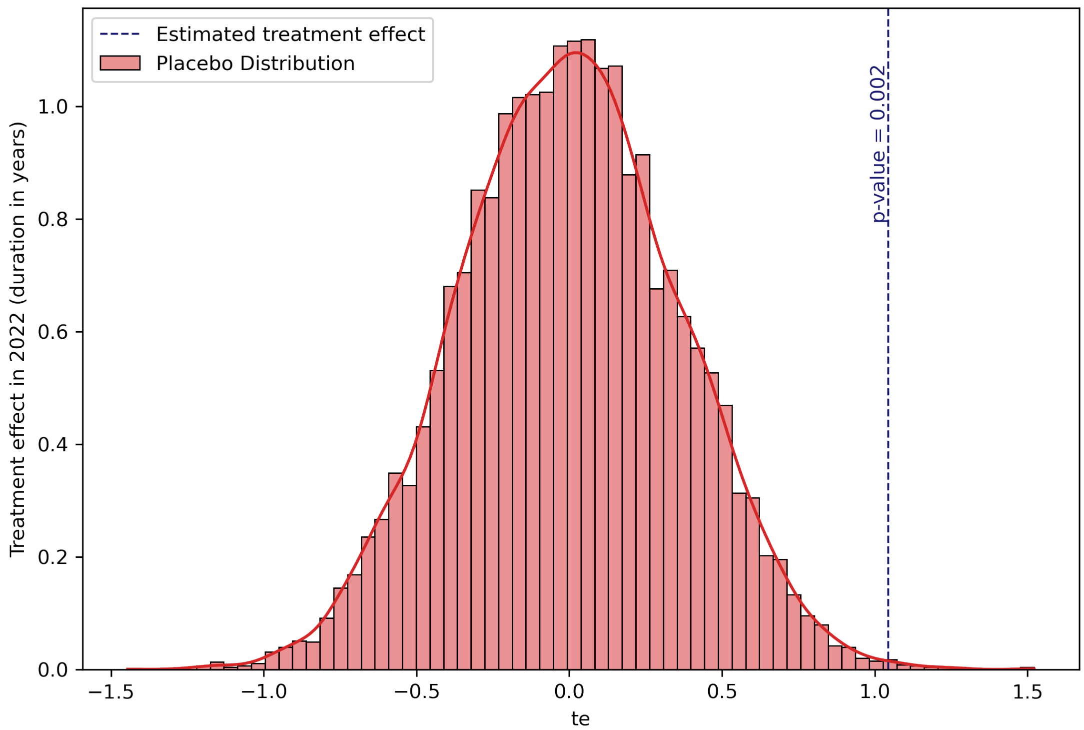

This figure plots the treatment effect (β_t^{VM20}) as outlined in the baseline DD regression. The dashed blue line represents the actual treatment effect, while the red bars show the placebo distribution.

D.4 Mean reversion point decomposition

The regulator uses the following rule to update the mean-reversion point parameter in January of each year using monthly frequency yield curve data,

$$\bar{y}_{20}^{(t)} = \text{Round}_{25\text{bps}} \left\{ 0.2 \times \text{Med}_{600m,t-1}\left(y_{20}\right) + 0.3 \times \text{Mean}_{120m,t}\left(y_{20}\right) + 0.5 \times \text{Mean}_{36m,t}\left(y_{20}\right) \right\}$$

The unrounded changes in the moving averages and medians from January to January can be decomposed into the parts coming from either: (i) new interest rate data released over the previous twelve months, or (ii) the parts coming from old observations dropping out/their weights changing,

1. $\Delta \,\text{Mean}_{36m,t}\,(y_{20}) = \frac{1}{36}\sum_{l=1}^{36} y_{20,t-l} - \frac{1}{36}\sum_{l=13}^{48} y_{20,t-l} = \underbrace{\frac{1}{36}\sum_{t=1}^{12} y_{20,t-l}}_{\text{new information}} - \underbrace{\frac{1}{36}\sum_{l=37}^{48} y_{20,t-l}}_{\text{old information}}$

2. $\Delta \,\text{Mean}_{120m,t}\,(y_{20}) = \frac{1}{120}\sum_{l=1}^{120} y_{20,t-l} - \frac{1}{120}\sum_{l=13}^{132} y_{20,t-l} = \underbrace{\frac{1}{120}\sum_{t=1}^{12} y_{20,t-l}}_{\text{new information}} - \underbrace{\frac{1}{120}\sum_{l=121}^{132} y_{20,t-l}}_{\text{old information}}$

3. $\Delta \,\text{Med}_{600m,t-1}\,(y_{20}) = \begin{cases} \underbrace{\text{Med}_{600m,t-1}\,(y_{20})}_{\text{new information}} - \underbrace{\text{Med}_{600m,t-12}\,(y_{20})}_{\text{old information}} & \text{if } \text{Med}_{600m,t-1} \in \{y_{t-1}, \ldots, y_{t-12}\} \\ \underbrace{\Delta \,\text{Med}_{600m,t-1}\,(y_{20})}_{\text{old information}} & \text{Otherwise} \end{cases}$

Collecting these terms we can decompose the changes in the parameter into,

$$
\begin{aligned}
\Delta\bar{y}^{(t)}_{20,\text{new data entering windows}} =& 0.2 \times \mathbf{1}\{\text{Med}_{600m,t-1} \in \{y_{t-1}, \ldots, y_{t-12}\}\} \times \text{Med}_{600m,t-1}\,(y_{20}) \\
&+ \left(0.5 \times \frac{1}{36} + 0.3 \times \frac{1}{120}\right) \sum_{t=1}^{12} y_{20,t-1-l} \\
\Delta\bar{y}^{(t)}_{20,\text{old data exiting windows}} =& 0.2 \times \mathbf{1}\{\text{Med}_{600m,t-1} \notin \{y_{t-1}, \ldots, y_{t-12}\}\} \times \Delta \,\text{Med}_{600m,t-1}\,(y_{20}) \\
&- 0.2 \times \mathbf{1}\{\text{Med}_{600m,t-1} \in \{y_{t-1}, \ldots, y_{t-12}\}\} \times \text{Med}_{600m,t-1}\,(y_{20}) \\
&- 0.5 \times \frac{1}{36} \sum_{l=37}^{48} y_{20,t-l} - 0.3 \times \frac{1}{120} \sum_{l=121}^{132} y_{20,t-l} \\
\Delta\bar{y}^{(t)}_{20,\text{ rounding}} =& \Delta\bar{y}^{(t)}_{20} - \Delta\bar{y}^{(t)}_{20,\text{unrounded}}
\end{aligned}
$$

Note that changes in the unrounded part can be fully decomposed into,

$$
\Delta\bar{y}_{20,\text{unrounded}} \equiv \Delta\bar{y}^{(t)}_{20,\text{new data entering windows}} + \Delta\bar{y}^{(t)}_{20,\text{old data exiting windows}}
$$

Hence, this gives the decomposition identity I use to construct the instrument,

$$
\Delta\bar{y}^{(t)}_{20} = \Delta\bar{y}^{(t)}_{20,\text{new data entering windows}} + \Delta\bar{y}^{(t)}_{20,\text{old data exiting windows}} + \Delta\bar{y}^{(t)}_{20,\text{ rounding}}
$$

D.5 Higher cost of regulation and probability of opting-out

Insurers for whom the regulation would have been more costly (high α_t^{VM20}) were more likely to opt out when given the option. The cost of regulation is higher for life insurers with lower RBC ratios. In this section, I test whether more financially constrained insurers, for whom VM20 was more costly, were more likely to opt out when given the option.

VM20 was enacted in 2017, but it allowed insurers to opt out until the end of 2019 to help them transition to the new reserving methodology. After 2020, a few in-state insurers and small insurers, for whom it would be too costly to follow VM20, were also allowed to apply for exemption from the regulation. Unfortunately, there isn't data available on who chose to opt out during the transition period; however, we can observe insurance companies opting out post-2020, as they need to reveal this in supplementary VM20 regulatory filings.

Table 9 estimates a linear probability model of the decision to opt out of VM20 in 2022—the results are similar for other years as there is little time-variation in the decision to opt-out. The estimates show that smaller firms and those that are more financially constrained are more likely to apply to opt out of VM20. The size effect is somewhat mechanical, as only sufficiently small firms are able to opt out of the regulation. However, the RBC effect is quantitatively large and significant. Conditional on size, having a one-standard deviation lower RBC ratio increases the probability of opting out by 8 percentage points.

Table 9: Linear probability model of VM20 opt out decision on insurer characteristics

	$\mathbf{1}\{\text{VM20 Opt-Out}_{2022}\}$
z-score($\log \text{RBC ratio}_{t-1}$)	-0.08***
	(0.03)
$\log(\text{Total Investments}_{t-1})$	-0.14***
	(0.05)
Observations	44
R-squared	0.20

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table shows the estimates from regression the binary variable of if an insurer chose to opt out of VM20 in 2022 on the laggged size and the lagged z-score of the risk based capital ratio.

D.6 Supporting evidence for regulatory constraint mechanism

In the paper, I argue that regulator model-implied beliefs influence insurer actions primarily through their effect on statutory liabilities—rather than through belief updating or learning. As supporting evidence, Figure 17 plots the relationship between annual changes in statutory liabilities and the corresponding changes in statutory equity and total assets, using the same insurer-year sample as in the baseline difference-in-differences analysis (2011-2022). Consistent with insurers managing assets to maintain roughly constant statutory equity, changes in statutory liabilities map almost one-for-one into changes in assets, with an R^2 of 0.99. By contrast, the relationship with statutory equity is much weaker ($R^2 = 0.15$), indicating that most liability movements are offset by asset adjustments rather than by changes in equity.

Figure 17: Relationship between changes in statutory liabilities, assets, and equity

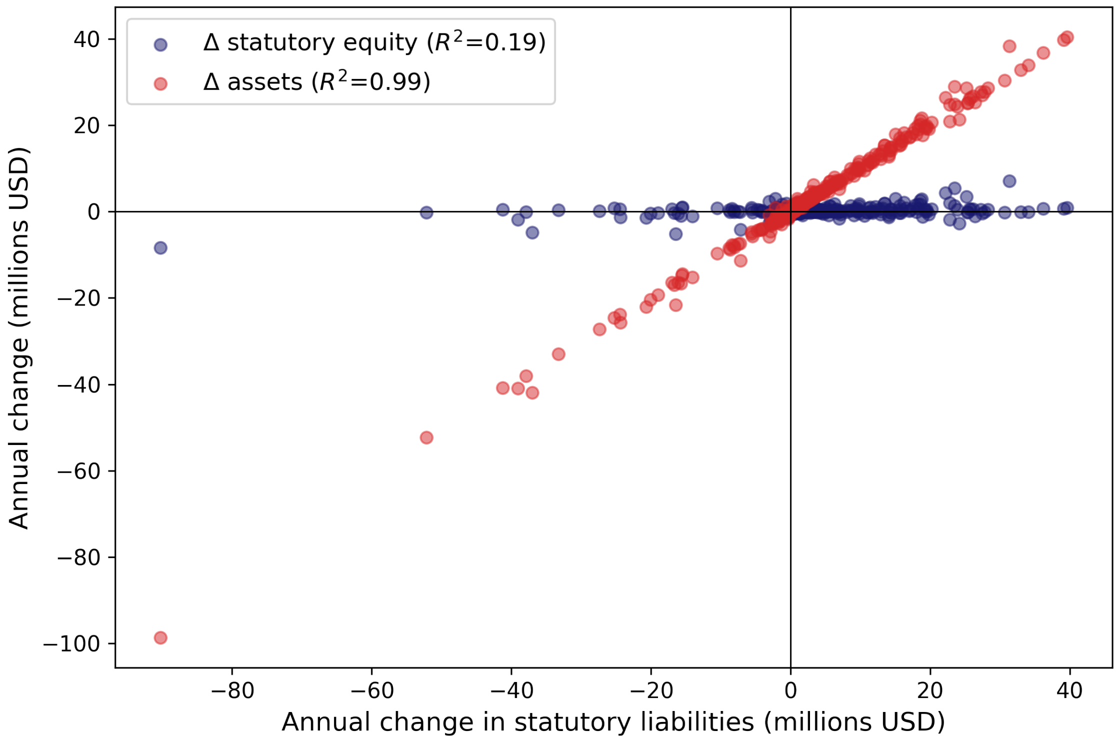

The figure plots the annual change in statutory equity (blue) and total assets (red) against the annual change in statutory liabilities (in millions of USD). Each point corresponds to an insurer-year observation over the period 2011-2022.

D.7 Disagreement estimates

Table 10: Average disagreement between regulator and insurer ($\mathbb{E}_t^R[y_{10,t+h}] - \mathbb{E}_t^I[y_{10,t+h}]$)

Horizon (years)	10yr yield expectation disagreement (bps)			
	Full Sample	Pre VM20	Post VM20	Difference
1	0.21	-30.36***	14.80***	45.16***
	(5.03)	(3.31)	(1.36)	(3.58)
2	-9.94	-55.07***	11.60***	66.67***
	(7.27)	(5.89)	(1.28)	(6.03)
3	-16.59*	-70.95***	9.35***	80.30***
	(8.80)	(8.18)	(1.20)	(8.27)
4	-43.55***	-112.43***	-10.68*	101.75***
	(11.03)	(11.32)	(5.94)	(12.78)
5	-63.90***	-132.11***	-31.34***	100.77***
	(13.10)	(12.66)	(11.98)	(17.42)
6	-73.94***	-141.41***	-41.74***	99.66***
	(14.86)	(21.21)	(11.91)	(24.32)
7	-81.32***	-144.77***	-51.04***	93.73***
	(15.60)	(25.85)	(11.97)	(28.49)
8	-86.20***	-142.44***	-59.36***	83.08***
	(14.65)	(24.86)	(12.19)	(27.69)
9	-89.97***	-140.35***	-65.92***	74.43***
	(14.35)	(24.16)	(13.55)	(27.70)
10	-94.00***	-138.69***	-72.67***	66.02**
	(14.52)	(23.83)	(15.53)	(28.45)
11	-90.60***	-131.73***	-70.98***	60.75**
	(14.07)	(23.41)	(15.81)	(28.24)
12	-86.10***	-123.60***	-68.20***	55.40*
	(14.00)	(23.60)	(16.33)	(28.70)

Clustered standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

The table shows the average difference between regulator model-based and insurer effective expectations of 10yr yields for different expectation horizons, using annual data from 2009 to 2022, and the standard errors are clustered at the insurer level. The The first sample shows the full sampe estimate, the second and the third columns are the pre VM20 (2009-2016) and post VM20 (2017-2022 respectively, and the fourth column is the difference between the estimates between the pre and post VM20 columns.