Risk Management, Product Offerings, and Consumer Surplus: Evidence from the Insurance Industry *

Cameron M. Ellis† Andrew Ellul‡ Chotibhak Jotikasthira§ Jianren Xu¶

March 26, 2026

Abstract

We study the causal impact of enterprise-wide risk management (ERM) — designed to move firms away from a "siloed" structure — on product decisions and consumer surplus. Exploiting the staggered rollout of an industry-wide ERM mandate in the insurance sector, we analyze life insurers' offerings of annuities, which now account for nearly 70% of their premium revenues. We find that insurers respond by reducing risky guarantees, raising fees on the riskiest products, and shifting from traditional variable annuities toward index-linked products that provide natural hedges. To examine mechanisms and welfare outcomes, we develop a structural model that links consumer demand with multi-product supply. ERM corrects firms' misperceptions about guarantee risk and cross-product risk interactions, effectively raising the marginal costs of offering risky guarantees. As a result, equilibrium prices increase and guarantee offerings contract, leading to significant losses in consumer surplus. Overall, ERM reshapes insurers' product strategies and risk exposures, enhancing financial stability but at a cost to consumers.

Keywords: Enterprise Risk Management · Variable Annuities · Index-Linked Annuities · Consumer Surplus · Insurance Companies · Financial Intermediaries · Stress Testing · Regulation
JEL classifications: G22 · G28 · G32 · G11

*The authors would like to thank Alexandru Barbu, Sharjil Haque (discussant), Anastasia Kartasheva, Steve Miller (discussant), Greg Nini (discussant), Joe Hong Zou, seminar or conference participants at California State University-Fullerton, Central Bank of Italy, Ludwig Maximilian University of Munich, Northeastern University, University of California Davis, University of Calgary, University of Hamburg, University of Iowa, University of North Texas, the 2026 Florida State University Risk Management and Insurance Research Symposium, 2026 Midwest Finance Association Annual Meeting, 2025 Modern Risk Society seminar, 2025 National Association of Insurance Commissioners (NAIC) Post-Summit Research Workshop, 2025 NTxEC Conference, 2025 Society for the Advancement of Economic Theory (SAET) Conference, 2025 World Risk and Insurance Economics Congress (WRIEC) for helpful comments and feedback. The authors have no conflicts of interest to report.

†Tippie College of Business, University of Iowa, cameron-ellis@uiowa.edu
‡Kelley School of Business, Indiana University, CEPR, CSEF, and ECGI, anellul@iu.edu
§Cox School of Business, Southern Methodist University, cjotikasthira@smu.edu
¶G. Brint Ryan College of Business, University of North Texas, jianren.xu@unt.edu

In imperfect capital markets, risk management affects the fragility of financial institutions' balance sheets by shaping how they respond to market frictions. Prior research finds that risk management constrains risk-taking activities or improves hedging practices (Berger et al., 2016; Ellul and Yerramilli, 2013; Giambona et al., 2025; Landier et al., 2009; Laeven and Levine, 2009; Stulz, 2008), with a primary focus on the asset side of the balance sheet. This emphasis on isolated balance-sheet components or specific types of risk mirrors firms' historical tendency to manage risks in a "siloed" manner—at the level of individual products or business units—potentially obscuring how risks interact and offset one another across the organization. Recent shifts toward enterprise-wide risk management (ERM) (SEC, 2007; Li et al., 2025) aim to overcome these limitations. Because ERM can reshape how a firm perceives risks and understands their propagation across its balance sheets, moving from a siloed to an enterprise-wide approach may influence not only the firm's asset portfolio but also its product offerings and, in turn, consumer surplus.

This paper examines the effects of ERM on product offerings within the life insurance industry. Life insurers supply a variety of retirement and savings products that constitute a significant share of their liabilities and expose them to differing types and levels of risk, with direct implications for their financial resilience. Providing empirical evidence on this previously unexplored mechanism contributes to a deeper understanding of how risk management practices shape insurers' product mix and, ultimately, balance sheet fragility. Furthermore, because consumers have heterogeneous preferences over product features, changes in the product mix induced by ERM can affect consumer surplus, highlighting a potential tradeoff between financial stability and consumer welfare.

Because underlying organizational frictions and managerial traits jointly influence both ERM adoption and product design, regressions linking risk management intensity to product portfolio changes are unlikely to recover causal effects.[1] To address this endogeneity, we leverage the staggered implementation of the Own Risk and Solvency Assessment (ORSA) Model Law, a regulatory mandate for insurers grounded in ERM principles, across U.S. states.[2] A key feature of our identification strategy is that, unlike many regulatory changes studied in prior work, ORSA (i) targets risk management practices rather than other bundled regulatory issues, limiting concerns about confounding channels, and (ii) was adopted state-by-state based on the domicile of the insurer's lead company. Our preliminary evidence suggests that staggered adoption was driven

[1] In theory, firms should assume risk in a value-maximizing manner, with risk management policies guiding behavior toward optimal choices or, at a minimum, constraining value-destroying actions. In practice, however, frictions such as high-powered executive incentives, leverage, and managerial short-termism often undermine effective risk governance (Kashyap et al., 2008; Stulz, 2008). In complex organizations, these frictions, together with corporate culture and executives' risk attitudes (Bodnar et al., 2019), can impede the recognition of risk silos and the adoption of ERM, while simultaneously shaping broader strategic choices, including product design.

[2] In the financial sector, firm- and system-level optimal risk taking can diverge due to interconnectedness and externalities. Firm-level risk taking and associated losses may generate negative spillovers through fire sales, contagion, and broader financial instability, motivating regulatory interventions such as the Dodd–Frank Act in banking and Solvency II in insurance.

mainly by differences in regulatory sophistication and enforcement priorities, rather than systematic insurer characteristics. Although adopting firms vary in size and product mix, they exhibit similar pre-adoption ERM ratings, balance sheet characteristics, and trends in product portfolios. Moreover, while insurers are domiciled in different states, they operate nationwide and face common demand conditions. Taken together, these facts support our identifying assumption that ORSA timing is plausibly unrelated to unobserved firm-level trends—such as shifts in consumer demand—that could otherwise influence product offerings.

Why should ERM, as represented by ORSA in our setting, affect risk taking and product choices? ORSA requires insurers to assess current and prospective risks on an enterprise-wide basis, particularly under adverse scenarios, and to report these assessments to state regulators. In doing so, it embeds forward-looking stress testing and capital adequacy analysis into strategic decision-making. By forcing firms to quantify and aggregate risks across business lines, ORSA can influence product strategy in two main ways.[3] First, it can clarify the capital intensity and tail-risk implications of different activities, moving firms toward a more optimal risk profile by reducing previously underappreciated exposures and expanding activities better understood to be manageable and capital-efficient. Second, enterprise-wide evaluation can reveal natural hedges across product lines, allowing firms to redesign products or adjust hedging strategies to reduce net risk and lessen capital pressures.[4]

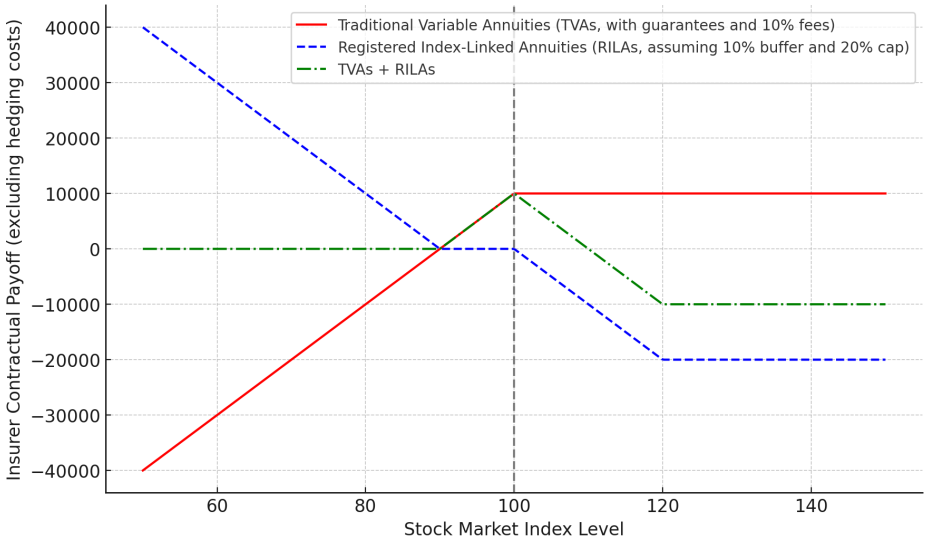
We use life insurers' annuity product offerings as a natural testing ground. With the decline of employer-sponsored defined benefit plans in the U.S., annuities have increasingly served as a substitute mechanism for retirement security, representing almost 70% of the direct premiums as well as the liabilities of life insurers in 2024 (National Association of Insurance Commissioners, 2025; American Council of Life Insurers, 2024, respectively). Historically, the main annuity products are traditional variable annuities (TVAs), which over time have become a key liability category for life insurers. While these contracts appeal to policyholders by offering equity market participation, many have embedded guarantees, which can expose insurers to sizable aggregate market risks (Ellul et al., 2022). Specifically, the guarantee feature of these annuity products are similar to long-dated out-of-the-money put options, exposing insurers to undiversifiable tail risk. Periods of financial distress magnify these exposures, making TVAs a prime laboratory for studying how risk management mandates alter product decisions (Koijen and Yogo, 2015).

[3]These effects require that ORSA induces real changes in risk management practices, rather than merely symbolic compliance. A key alternative hypothesis is that ERM may be implemented superficially—for example, by appointing risk managers without granting them meaningful authority—allowing insurers to satisfy reporting requirements without altering product strategies or risk exposures. Our empirical analysis is designed to distinguish between substantive and purely formal responses to the mandate.

[4]The importance of this approach for the financial industry was explicitly stated by the Basel Committee on Banking Supervision (2010).

A separate product line that has recently emerged and grown rapidly is registered index-linked annuities (RILAs). In 2023, RILA sales surpassed TVA sales for the first time (Almazora, 2024). Like TVAs, RILAs offer equity-based returns, but compared to TVAs, are less "guarantee-loaded": they provide partial downside protection through buffers or floors. In the more common buffer design, the policyholder is protected from the initial portion of market decline (e.g., the first 10–15% decline in the reference index) but is responsible for losses beyond this threshold. Economically, this structure (with a deposit contract as the base) is equivalent to the policyholder writing an out-of-the-money put option to the insurer, thereby offsetting the payoff profile of TVA guarantees. More generally, both buffer and floor designs limit the insurer's downside exposure (in exchange for some upside exposure) and thus serve as a natural hedge against the risks inherent in TVA guarantees (Barbu and Sen, 2024).[5]

Figure 1: Insurer Contractual Payoffs Under TVAs, RILAs, and Combined Portfolio



In our setting, where insurers hold substantial legacy liabilities from guaranteed TVAs, if ORSA induces meaningful changes in risk management and product strategy, we would expect a shift in the product mix away from TVAs, which heighten volatility and capital strain, and toward RILAs, which naturally offset these exposures. For the TVAs that continue to be offered, we

[5] Although insurers can, in principle, hedge the downside risk embedded in TVA guarantees using derivatives or reinsurance, practical and regulatory constraints often make such hedging costly or incomplete (Ellul et al., 2022; Sen, 2023). By contrast, hedging risks within the product portfolio—such as through the introduction of RILAs—is generally more flexible and is viewed more favorably by relevant regulations.

also expect less generous guarantees and higher fees. Using an imputation staggered difference-in-differences design following Borusyak et al. (2024),[6] we find evidence consistent with these predictions.

Specifically, we show that firms newly subject to ORSA reduce their offering of TVAs and expand their offering of RILAs, while leaving their offering of traditional fixed annuities unchanged, consistent with the fact that fixed annuities do not provide comparable hedging benefits. In other words, ORSA leads to a substantive reallocation within the product portfolio. Moreover, for the TVAs that remain, affected insurers are less likely to offer the most generous guarantees and significantly increase fees. We check the robustness of our estimates using the specification curve methodology of Simonsohn et al. (2020), which helps us to systematically examine how treatment effect estimates change across different combinations of specifications and methodologies.

We interpret these patterns as the causal effects of ERM, as represented by ORSA. Because ORSA was adopted at different times across states, it is unlikely that our results are driven by shifts in consumer demand or preferences. Any demand-based explanation would require that changes in consumer preferences for specific annuity features occurred precisely and exclusively in the states where insurers happened to be domiciled at the moment of ORSA adoption. This is an implausible scenario, especially given that insurers, regardless of their state of domicile, operate nationally and face largely uniform demand conditions.

One potential alternative explanation is that ORSA's stress-testing requirements lead insurers to hold higher capital buffers and, as consequence, change their product mix as one route to achieve the regulatory goal. However, we find no evidence that RBC ratios increase after ORSA; if anything, they decline, particularly for firms with ample regulatory capital in the pre-ORSA period, consistent with insurers adopting product strategies that improve their risk profile but are indirectly penalized by the regulatory capital formula that does not internalizes natural hedges. These considerations suggest that our results are unlikely to be driven by binding RBC constraints.

One related argument that links the demand channel with the regulatory channel discussed above is made by Barbu et al. (2025) who show that risk-based capital (RBC) regulation affects insurance products through the supply channel, i.e. adjustments made by insurers to respond to tighter regulation, and the demand channel, i.e. households shift their demand towards low risk insurers. To the extent that ORSA should lead to tighter risk management, insurers should experience higher demand for all their products from households, including TVAs, eventually impacting prices as well. Instead we find that there is a clear compositional effect and not a widespread shift in demand: insurers offer more RILAs and reduce TVAs.

[6]This approach avoids using early-treated firms as controls for later-treated ones and allows us to estimate causal effects on multiple product attributes jointly by predicting counterfactual outcomes using untreated insurer-years.

Motivated by our DiD results, we develop a structural model of annuity supply to explore the underlying mechanisms through which ERM causes insurers to alter their product mix. In contrast to traditional supply-side models, we explicitly allow for cross-product hedging to lower marginal costs. However, prior to the adoption of ERM, firms may not fully recognize the hedging effects and instead set prices based on the "siloed" first-order conditions. This structure and the staggered roll-out of ORSA allow us to identify how the correction of cost misperceptions and the recognition of natural, product-led hedging affect pricing. We combine our supply-side model with a demand model, similar to Koijen and Yogo (2022), to estimate the hidden costs of the changed product design caused by risk management. TVA guarantees are risky to insurers because there are no financial derivatives that insurers can use to perfectly hedge the risk. This also means there are no direct substitutes for the long-term market risk insurance that TVA guarantees provide that consumers can turn to when the product menu changes, and our demand model allows us to estimate this lost consumer surplus directly.

In our structural analysis, we estimate a Berry et al. (1995, BLP) demand system and pair it with a multi-product Bertrand supply model. Marginal costs include a portfolio "natural hedge": a larger non-TVA franchise lowers the marginal cost of guaranteeing TVA payoffs and, through multi-product pricing incentives, encourages insurers to steer marginal demand toward products that offset guarantee risk. We find that, pre-ORSA, insurers materially under-recognized these portfolio-level hedge benefits and, to a less extent, underpriced product-level guarantee costs. ORSA forces enterprise-wide risk aggregation and improves cost recognition. In equilibrium, guarantee-intensive TVAs become more expensive, demand shifts toward RILAs, and consumer surplus falls as households lose access to long-dated downside insurance.

Specifically, our structural estimates indicate that pre-ORSA, firm's substantially under-recognized portfolio-level natural-hedge benefits, about 53% accuracy for firms with weak pre-ORSA ERM ratings and 64% for strong-ERM firms, while recognizing guarantee costs at roughly 88–93%. We allow for time and firm fixed effects; we do not separately identify a post-ORSA compliance-cost shifter in the current specification. The welfare implications of the shifts in costs and prices induced by ORSA are substantial and adverse. Consumer surplus peaks at $24.0 billion in 2011 and falls to $8.2–9.1 billion by 2016–2017 as ORSA adoption reaches near-full coverage; it partially rebounds to $13.6 billion in 2021 but declines again to $7.5 billion by 2023. This implies annual welfare losses on the order of $6–17 billion under our logit measure. These losses reflect a fundamental tension: while ORSA corrects cost misperceptions and enhances risk management, the induced price increases outweigh any consumer benefits from improved pricing accuracy. The 24.9 percentage-point decline in the probability of offering TVAs suggests reduced systemic risk, achieving ORSA's macroprudential objectives. However, this stability imposes significant costs on retirement savers who depend on guaranteed annuities, transferring investment risk to consumers at the life-cycle stage when they are least capable of bearing it.

Our paper contributes first to the literature on enterprise-wide risk management and its impact on firm risk-taking. While this literature highlights ERM's importance, it has faced persistent challenges in establishing causality because ERM adoption and intensity are endogenous (Ellul and Yerramilli, 2013; Chen et al., 2020; Nocco and Stulz, 2022; Amiraslani et al., 2025; Li et al., 2025; Neel and Xu, 2025). The staggered implementation of ORSA provides a plausibly exogenous shock to insurers' organizational risk management structures, allowing us to isolate how strengthened ERM affects the design and selection of products that determine insurers' liabilities and risk exposures—an important but previously underexplored channel.

Our findings also speak to the broader literature on financial regulation, including stress testing, and its economic effects. Prior work shows that regulation can shape market outcomes (Barbu et al., 2025; Koijen and Yogo, 2016), but credible identification is difficult when reforms are broad-based and treated and untreated firms operate in overlapping regulatory environments (Houston et al., 2012; Ongena et al., 2013; Aiyar et al., 2014; Frame et al., 2020). By exploiting within-industry variation in the timing of ORSA adoption, we provide causal evidence that risk management mandates can reshape firms' product strategies and risk exposures, clarifying a mechanism through which financial regulation can influence both firm behavior and consumer welfare. Analogous to stress testing in banking, ORSA requires insurers to internalize systemic externalities by incorporating the costs caused by potential system-wide distress into risk management. Empirical evidence shows that stress-tested banks reduce lending, particularly to riskier borrowers (Acharya et al., 2018; Eber and Minoiu, 2016) and small businesses (Cortés et al., 2020). While this literature focuses on banks' asset-side adjustments, we instead examine insurers' funding behavior through their product offerings, complementing existing evidence by showing that firm-wide stress testing can reshape decisions across the entire balance sheet rather than within a single functional area.

We also contribute to the extensive literature on TVAs, where numerous studies examine their guarantee structures, risk exposures and hedging, and capital implications (e.g., Ellul et al., 2022; Koijen and Yogo, 2015; Hufeld et al., 2017; Koijen and Yogo, 2022; Li, 2024; Barbu, 2023; Rogers and Tonetti, 2023). Our findings highlight how an enterprise-wide risk management approach can reshape product portfolios, adding a new dimension to the analysis of TVA pricing, market dynamics, and risk management. In addition, we complement an emerging literature on RILAs (e.g., Friedberg and Webb, 2021; Barbu and Sen, 2024; Mezger, 2024; Moenig and Xu, 2023; Ellis et al., 2025) by showing that insurers respond to enterprise-wide oversight not simply by retreating from risky contracts but by adopting products like RILAs that permit strategic risk sharing with policyholders. This result is reminiscent of the supply effects in Barbu et al. (2025). Unlike their capital-regulation channel, we show that ORSA reshapes annuity supply through a different mechanism: it improves firms' recognition of cross-product natural hedges, leading to a redesign of product portfolios. Our results suggest that risk management can be a driver of product innovation.

Our findings have implications beyond insurance and financial institutions. Life insurers' liabilities arise directly from the contractual promises embedded in the products they sell, so shifts in product mix reflect changes in risk management on the liability side of the balance sheet. A similar dynamic applies in many industries where firms offer products or services with deferred or contingent obligations, such as warranties, guarantees, advance payments, or long-duration service commitments. In these settings (e.g., airlines, healthcare, energy, and consumer goods), our findings suggest that organizational risk management can directly shape the design of contractual promises and the capital required to support them.

1 Institutional Background and Hypotheses

1.1 Enterprise Risk Management

Enterprise Risk Management (ERM) is a firm-wide approach to organizational risk management that evaluates risks and their interconnectedness across the entire firm, rather than within isolated business units. Unlike traditional silo-based methods—where individual departments manage risks independently, ERM integrates risk identification, assessment, and control into a unified framework that links risk management with strategic and regulatory objectives. If done effectively, ERM should allow the firm to improve capital allocation, and strengthen organizational resilience and decision-making (COSO, 2004; SEC, 2007; Hoyt and Liebenberg, 2011).

The strategic, enterprise-wide view provided by ERM should allow top management to identify interdependencies across risks or projects and recognize how they interact and influence one another. Such an integrated understanding allows firms to exploit **natural hedges**—instances where the dynamics of one risk or project offset the effects of another.[7]

The importance of ERM's comprehensive approach was highlighted in a 2011 speech by the U.S. Securities and Exchange Commission (SEC) Chairman Mary L. Schapiro, who emphasized that a commitment to ERM is essential to prevent the formation of silos and the overlooking of interdependencies between risk categories and across the product and asset classes. Real-world anecdotes also illustrate the challenges of silo-based approaches and the advantages ERM provides in addressing them. For example, in 1984, Lufthansa, Germany's largest airline, signed a 3 billion U.S. dollar contract to purchase aircraft from Boeing and entered into a 1.5 billion U.S. dollar forward contract to hedge exchange rate risk, fearing continued appreciation of the U.S. dollar. However, a significant portion of Lufthansa's cash flow was U.S. dollar-denominated, naturally

[7] A key feature of ERM is its emphasis on a "tone from the top" framework, where the board of directors plays a key role in guiding enterprise-wide risk management practices (Braumann et al., 2020; Neel and Xu, 2025). This leadership ensures that risk management strategies are consistent across the organization and aligned with the firm's mission, vision, and core values (COSO, 2017; Li et al., 2025).

offsetting the exchange rate exposure. By overlooking this natural hedge and managing the risk in isolation, Lufthansa incurred a major and unnecessary foreign exchange loss when the U.S. dollar depreciated against the mark in 1985 (The Economist, February 10, 1996, p. 16-17).

Recent capital adequacy frameworks—Basel III for banks and Solvency II for insurers—represent a clear shift toward enterprise-wide, risk-sensitive supervision, replacing the compartmentalized approaches of earlier regimes. Under prior systems such as Basel I, Basel II, and Solvency I, capital standards relied on standardized formulas and separate treatment of credit, market, and operational risks. By contrast, Basel III and Solvency II adopt a holistic view of firm-wide risk, emphasizing diversification and the aggregation of exposures across business lines. They also introduce forward-looking stress testing, internal models, and governance requirements that link capital to the institution's overall risk profile rather than to individual assets or liabilities. Together, these changes align regulatory capital more closely with economic risk and ERM as a central element of prudential oversight.

The insurance industry offers a natural laboratory for studying ERM. Because insurers' business models center on risk transfer and pooling, they face a broad range of exposures—from policy performance and asset-liability mismatches to demographic, operational, and strategic risks. Effective risk management must therefore operate not only at the product or subsidiary level but also enterprise-wide, where interdependencies among risks are recognized and managed collectively. Reflecting this view, regulators have embedded the ERM concept in recent initiatives such as the Own Risk and Solvency Assessment (ORSA) in the United States and Solvency II in Europe.

1.2 Own Risk and Solvency Assessment Requirements

ORSA is a U.S. regulatory framework designed to enhance insurers' financial stability by embedding ERM into a verifiable and recurring process. Introduced by the NAIC following the 2008 financial crisis, ORSA moves beyond principles-based guidance toward a structured, evidence-based regime that integrates risk management, capital planning, and board oversight at both the legal-entity and group-wide levels. It requires insurers not only to identify and assess their material risks but also to demonstrate—through documentation and analysis—their ability to remain solvent under adverse scenarios. The assessment is documented in an ORSA Summary Report, submitted confidentially to the insurer's lead state commissioner (NAIC, 2012).[8]

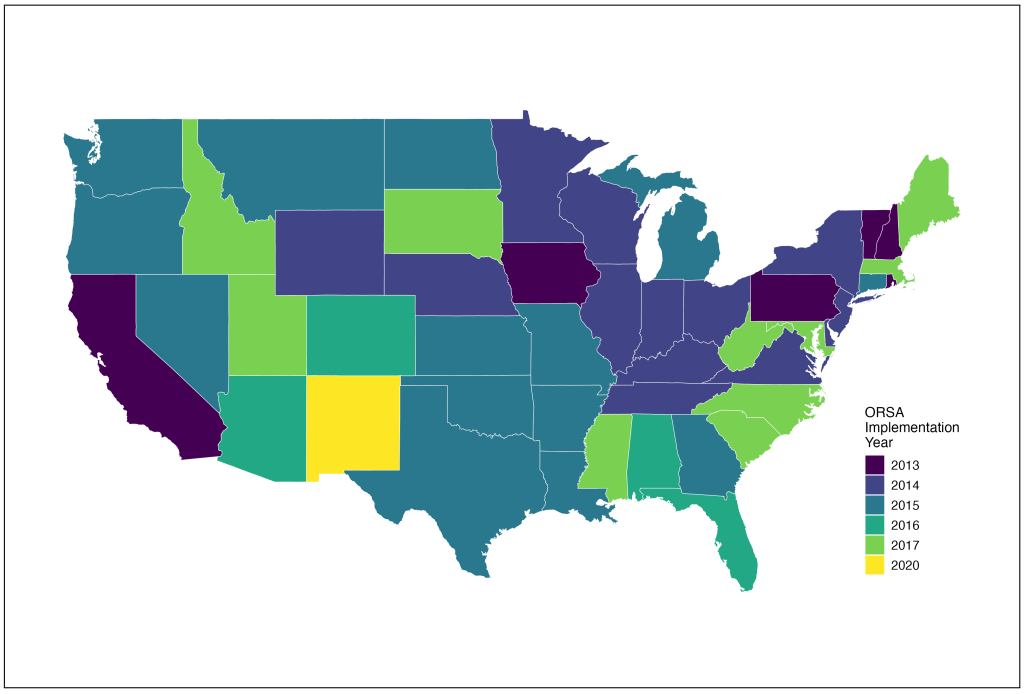
ORSA makes enterprise-wide risk management and stress testing verifiable through enforceable reporting and governance requirements. Each ORSA Summary Report must be signed by the Chief Risk Officer (or equivalent) and delivered to the board, creating a documented trail that dis-

[8]ORSA applies to U.S. insurers with more than $500 million in annual direct written and assumed premiums, or to insurance groups with over $1 billion. These insurers are required to conduct and document an ORSA on a regular basis. Because annuity writers are typically large insurers, all firms in our sample fall within the ORSA mandate.

tinguishes genuine implementation from symbolic compliance (Liu and Xu, 2024; Li et al., 2025). Conducted at least annually, ORSA turns ERM into a continuous, enterprise-level assessment of solvency under stress, rather than a one-off disclosure, with penalties for noncompliance. Regulators further shape the process by prescribing stress-test parameters and economic scenario assumptions, constraining insurers' discretion to produce overly optimistic projections. The ORSA Manual also specifies key report elements—accounting basis, data scope, solvency projections, and management actions under capital shortfalls—allowing supervisors to evaluate the firm's overall resilience and planned responses to systemwide stress.

The ORSA reporting requirement is outlined in the NAIC's Risk Management and Own Risk and Solvency Assessment Model Act (#505) (NAIC, 2012). U.S. states and jurisdictions adopted the Model Act between 2013 and 2017 in a staggered manner, as represented in Figure 2. As of 2024, 53 out of 56 U.S. jurisdictions have formally enacted the Model Act.

Figure 2: Map of ORSA Roll-out



Note: Figure shows the roll-out of ORSA regulations for insurer lead states. Most states adopted between 2013 and 2017; New Mexico adopted later in 2020. No insurers in our sample have New Mexico as their lead state, so this late adoption does not affect our analysis.

The above description indicates that ORSA for insurers is economically and financially analogous to bank stress testing because both frameworks are forward-looking, internal, risk-based capital adequacy frameworks designed to assess the resilience of financial intermediaries under

adverse but plausible states of the world. Both mechanisms address the same underlying market failures arising from leverage, maturity transformation, and risk opacity. Although ORSA is typically firm-focused and bank stress tests often have explicit system-wide components, the underlying objectives share the same economic principles: (a) individual firm failure creates negative externalities, and (b) market discipline alone is insufficient due to complexity and opacity. Both frameworks could function as macroprudential tools, even when applied at the micro level, by discouraging correlated risk-taking and procyclical behavior.

We want to emphasize from the outset that ORSA differs from other regulatory mandates in one important aspect, crucial for identifying the effects of ERM. Government mandates such as Dodd-Frank Act, Solvency II or SOX typically take effect across entire sectors (or even countries) at the same time. It is thus difficult to untangle their true causal impact on firm risk-taking because of the multiple effects on various dimensions of the firms' operations. Prior studies have attempted various identification strategies. However, since uniform regulatory roll-outs rarely generate clean control groups, they typically rely on either extremely strong assumptions or yield only "local" evidence.[9]

Eastman and Xu (2021) examine the market reactions of insurers with and without voluntary ERM adoption on key dates leading up to the passage of ORSA. Krupa (2024) finds that, due to ORSA's integrated approach in considering all firm-wide risks, insurers' managers are required to assess both tax and non-tax risks when making tax strategies. This leads to increased awareness of previously underestimated non-tax risks and a significant reduction in tax aggressiveness. The study also shows that this effect holds regardless of whether insurers had voluntarily adopted ERM prior to ORSA, suggesting that ORSA elevates ERM practices to a higher level, likely due to its mandatory nature and the detailed self-assessment reports required by regulators.

1.3 Product Lines

Life insurers' two primary product lines are annuities and life insurance (both individual and group). Over the past two decades, annuities have become the dominant source of premium income for U.S. life insurers, representing roughly 65–75% in recent years, depending on the accounting basis. This paper focuses on annuities because (a) they represent the principal driver

[9]Identification strategies include: Exemptions for very small companies (e.g., Balasubramanyan et al., 2024; Jiang and Ji, 2024); Cross-country comparisons (e.g., Houston et al., 2012; Ongena et al., 2013; Frame et al., 2020); Assumed variation in treatment intensity (e.g., Aiyar et al., 2014); and the inclusion of fixed effects (e.g., Amiraslani et al., 2025). A large separate literature also examines firms that voluntarily adopt ERM. Within this literature, earlier studies mainly rely on treatment effects models (e.g., Hoyt and Liebenberg, 2011; Eckles et al., 2014; Berry-Stölzle and Xu, 2018), while more recent research often employs difference-in-differences strategies with endogenous treatment, typically in matched samples (e.g., Eastman et al., 2024; Li et al., 2025; Neel and Xu, 2025).

of insurers' risk exposures, particularly systematic risk (Koijen and Yogo, 2022), and (b) the state-contingent payoffs across different annuity types exhibit rich cross-sectional heterogeneity and potential natural hedging opportunities within insurers' product portfolios.

Traditional Variable Annuities: Traditional Variable Annuities (TVAs) are hybrid financial instruments that combine investment products, typically mutual funds, with insurance features that provide income stability and downside protection. These products often offer an assortment of guaranteed benefits (many of which can be combined) in exchange for fees, with each benefit type shifting varying degrees of risk to the insurer. For detailed descriptions of specific annuity benefit types—including Guaranteed Minimum Death Benefits (GMDB), Guaranteed Minimum Accumulation Benefits (GMAB), Guaranteed Minimum Income Benefits (GMIB), Guaranteed Minimum Withdrawal Benefits (GMWB), and Guaranteed Lifetime Withdrawal Benefits (GLWB)—and their associated risk profiles, see Appendix A.

Among guaranteed benefits, lifetime living benefits—notably GMWB and GLWB—pose the greatest complexity and risk for insurers, as they depend on both market performance and policyholder longevity.[10] Financially, these guarantees resemble put options written on mutual fund portfolios, often combining index-based and actively managed funds. Hedging these exposures is costly and imperfect, despite affording partial relief under RBC requirements.[11] In practice, insurers often under-hedge due to regulatory frictions (Sen, 2023), and the marginal effectiveness of hedging declines as coverage expands (Ellul et al., 2022). Consequently, substantial basis risk remains in most portfolios (Ankirchner et al., 2014; Bauer, 2020).

Registered Index-linked Annuities: Registered Index-linked Annuities (RILAs) are a new form of variable annuity. RILAs have drastically gained in popularity in recent years, with annual sales exceeding $40 billion in 2022, after being virtually non-existent a decade earlier (Ellis et al., 2025). A RILA entails an investment into a separate account managed by an insurance company. The account evolution is linked to the performance of a popular market index—such as the S&P 500—over a crediting term of usually one year (but potentially up to six years). At the end of the year, the insurer credits the investor's account with the index return, subject to downside protection in the form of a Floor or Buffer. For instance, with a 10% Floor the investor can lose at most 10% of her funds over the crediting term, and with a 10% Buffer the loss from the index decline would be reduced by up to 10%. In exchange, the investor accepts an upside limit, usually in the form of a Cap, that the credited return cannot exceed. When combined with a target-date glide path

[10] By contrast, GMDBs are less complex, with more predictable triggers and shorter horizons. We focus on the lifetime benefits of GMWB and GLWB and classify the rest as non-lifetime benefits. Results are robust to including GMAB and GMIB as lifetime benefits, though magnitudes weaken somewhat.

[11] Insurers typically hedge through futures, ETF short positions, or index put options to offset guarantee deltas.

that gradually reduces equity exposure, these target-date RILAs (TD-RILAs) offer a passive, risk-adjusting investment option well-suited for qualified retirement plans (QRPs), such as 401(k) and 403(b) plans.

For insurers, RILAs offer several advantages. The embedded equity risk can be hedged efficiently with exchange-traded European options on major indices, avoiding the complex and dynamic hedging required for traditional variable annuity (TVA) guarantees. RILAs also provide pricing flexibility, as carriers can reset the Cap each term to reflect current option market conditions—an adjustment used effectively during the 2020 volatility spike (Moenig, 2022). Finally, RILAs are generally lower-fee products (Moenig, 2022; Moenig and Samuelson, 2024), reflecting their efficient risk structure and regulatory treatment: since investors have no direct exposure to the index, insurers invest premiums at their discretion (typically in corporate bonds) and earn a credit spread that helps offset costs and sustain profitability.

Synergies Between TVAs and RILAs: RILAs provide a natural hedge against TVA guarantees by effectively transferring part of the downside risk back to investors (Barbu and Sen, 2024). In TVAs, insurers provide downside protection to policyholders, whereas in RILAs, the policyholders accept some downside risk in exchange for some upside potential, effectively "selling" a put option to the insurer. This offsetting exposure allows carriers to hedge large realizations of equity (and interest rate) risk within their product portfolio. Unlike exchange-traded options, RILAs involve cash premiums held by insurers, which can be invested in safe assets—eliminating collateral, counterparty, and liquidity risks. Because these products are typically held in tax-advantaged retirement accounts, insurers naturally intermediate market risk, matching households seeking protection (typically older and more risk-averse) with those willing to bear limited downside exposure (yournger and more risk-tolerant). The result is lower capital volatility, a reduced need for external hedging, and a more balanced allocation of equity market risk across households with different risk profiles.

Figure 1 builds on the explanations above and illustrates insurer contractual payoffs under three product structures: TVAs, RILAs, and a portfolio combining one unit of each, assuming a 10% TVA fee and ignoring hedging costs. TVAs expose insurers to significant downside risk due to embedded guarantees, while passing through gains to policyholders when equity markets rise. The illustrative RILA example in the figure assumes a 10% buffer and a 20% cap (actual product designs vary). When combined, the TVA and RILA payoffs offset each other in both tails of the return distribution, flattening the insurer's aggregate exposure and creating a natural hedge.

1.4 Hypotheses

Taken together, ORSA's enterprise-wide requirements can influence insurers' product strategies through two mutually reinforcing channels. First, by compelling insurers to quantify and aggregate risks across business lines—especially under stress scenarios—ORSA improves forward-looking risk assessment and may correct prior misperceptions about the full risk distributions of different products, particularly their tail-risk exposures. This sharper understanding can move firms toward their optimal level of risk, leading them to scale back exposures that were previously underestimated or, in some cases, expand those now viewed as more manageable. Second, ORSA's emphasis on enterprise-wide evaluation can uncover natural hedging opportunities across product lines—such as the offsetting risk profiles of TVAs and RILAs—encouraging insurers to redesign product offerings or adjust hedging strategies to reduce net risk and relieve capital pressures.

We begin by assuming that ORSA's ERM requirements place greater weight on risk reduction than insurers would under their own objective functions (see, e.g., Kim and Santomero, 1988). This implies that prior to ORSA, insurers may have taken on more risk than regulators preferred—either to boost revenues or profits.[12] In this context—where insurers already hold substantial legacy liabilities from guaranteed TVAs—ORSA should push firms to rebalance their product portfolios: away from TVAs, whose embedded guarantees heighten volatility and capital strain, and toward RILAs, which are less guarantee-loaded and naturally offset TVA exposures (as illustrated in Figure 1). Among TVAs that remain in the market, we expect insurers to shift toward less generous guarantees and higher-fee designs that better compensate for their risk and, perhaps, incentivize customers to switch to RILAs.

As noted above, some insurers had already implemented ERM prior to ORSA. To the extent that an insurer's pre-ORSA risk management practices were already comprehensive and stringent, the introduction of ORSA would be expected to have little incremental effect on its product offerings or guarantee pricing. Conversely, for insurers whose earlier ERM frameworks were less comprehensive or insufficiently robust, ORSA should prompt meaningful adjustments in product design, product mix, and fee structures, consistent with the mechanisms discussed above.

The welfare implications for households arising from ORSA-driven shifts in insurers' annuity product design and pricing also merit consideration. Many households value TVAs with lifetime guarantees because these contracts embed a long-term put option that has no close substitute in broader financial markets. If so, any reduction in the availability of guaranteed TVAs—or any

[12] Such frictions as high-powered compensation and limited internalization of failure externalities may lead insurers to underweight tail risks or the broader systemic consequences of distress.

increase in the cost of these guarantees—is likely to produce a meaningful decline in consumer surplus. This loss cannot be offset by greater availability of RILAs, which offer attractive investment features but do not replicate the insurance protection inherent in lifetime guarantees.

2 Data

2.1 Data Sources

We combine data from multiple sources to construct firm-level, contract-level, and benefit-level panels for our analysis. Morningstar Annuity Intelligence is our primary source for product-level data. From Morningstar, we obtain comprehensive annuity contract characteristics including product type classifications (TVA, fixed, and index-linked), benefit rider descriptions and features, fee schedules, contract inception and close dates, and sales data for TVA contracts.[13] Morningstar's coverage is comprehensive for TVA products but more limited for non-TVA products (fixed and index-linked annuities). Consequently, sales data are available primarily for TVA contracts, though we observe contract offerings (inception and close dates) for all annuity types.

NAIC filings provide firm-level financial and operational data. From insurers' annual statements reported to the NAIC, obtained through S&P Global Market Intelligence, we obtain insurer characteristics including net total assets, surplus, direct and net premiums written, asset allocations, and capitalization measures such as RBC ratios. The NAIC data also include total annuity account values and gross reserves associated with annuities with guarantees, reported in the General Interrogatories (2003-2016) and the Variable Annuities Supplement, Parts 1 and 2 (2017 onward).[14] Separate NAIC data provide each firm's lead state—typically the state in which the firm (if unaffiliated) or its affiliated group (if part of a group) is domiciled or headquartered—which determines ORSA treatment timing in our staggered difference-in-differences design.

We merge Morningstar product-level data with NAIC firm-level data by company name and year, using a combination of automated string matching and manual validation to establish firm linkages. This merged dataset allows us to analyze how firm characteristics and regulatory treatment affect product-level decisions. For sales analyses, we use Morningstar's product-level TVA sales data aggregated to the firm level, supplemented by NAIC firm-level total annuity sales where appropriate to capture non-TVA products.

[13]The Morningstar annuity data were collected, with permission, through web scraping from Morningstar document archives and AnnuityIntel.com.

[14]These aggregate reserve and account value data from NAIC filings represent totals across all of a firm's annuity products with guarantees, complementing the product-level detail from Morningstar.

To explore heterogeneity based on prior usage of ERM, we obtain S&P ERM ratings from S&P Global. These ratings evaluate insurers' risk management practices across five dimensions: risk management culture, risk controls, emerging risk management, risk models, and strategic risk management. The ratings are scored on a 5-point scale: (1) Weak, (2) Adequate, (3) Adequate with strong risk controls, (4) Strong, (5) Very Strong/Excellent. The ratings assess whether insurers execute risk management practices in a systematic, consistent, and strategic manner across the enterprise. We use pre-ORSA ERM ratings (from 2010) as a proxy for firms' prior level of risk management sophistication.[15]

We calculate shadow insurance usage from Schedule S reinsurance filings obtained through S&P Global Market Intelligence. Shadow insurance is identified using three criteria: (1) the reinsurer is unauthorized (not licensed in the ceding company's state), (2) the reinsurer is affiliated with the ceding company, and (3) the reinsurer is not rated by A.M. Best. We measure shadow insurance usage as the sum of life reserves credit taken and modified coinsurance ceded to these shadow reinsurers. We match reinsurer names from Schedule S to A.M. Best rating data to identify unrated reinsurers, using automated string matching with manual review. We use pre-ORSA shadow insurance usage (from 2010) as a proxy for firms' prior level of risk management sophistication.

2.2 Sample Construction

Our full sample covers the period from 2005 to 2023. The reduced-form analyses use data from 2008 to 2016, which provides a balanced pre- and post-treatment window around the staggered ORSA adoption period (2013-2017) while maintaining a meaningful control group throughout. We end in 2016 because firms in the final adopting states became subject to ORSA in 2017, eliminating the control group necessary for our identification strategy. In contrast, the structural model is estimated using data from 2005 through 2023 to capture the longer-run evolution of annuity sales and consumer surplus after ORSA implementation, allowing us to examine whether the observed post-ORSA patterns persist over time.

We construct three nested panels from our merged data sources. First, for firm-level sample, we start with the universe of insurance companies that filed annual statements with the NAIC and offered at least one annuity contract in Morningstar during 2008-2016. This yields 941 firm-year observations at annual frequency across 120 unique firms. Second, in the contract-level sample, we observe all annuity contracts from Morningstar at quarterly frequency to capture the timing of contract introductions and removals precisely. This yields 173,744 contract-quarter observations across 5,198 unique contracts. Third, for benefit-level hazard and fee analyses, we focus on guar-

[15] We assign a value of 0 to firms that are not rated by S&P.

anteed benefit riders offered within TVA contracts. Benefits are observed at quarterly frequency. This hierarchical structure—where each contract may offer multiple benefit riders—yields 947,900 benefit-quarter observations across 27,525 unique benefit-contract combinations and 2,596 contracts. [16]

Summary statistics in Table 1 are reported at the observation-unit frequency used in the corresponding regression analyses: annual for firm-level variables (Panel A), quarterly for contract-level variables (Panel B), and quarterly for benefit-level variables (Panel C). The observation counts in Table 1 match exactly the samples used in our reduced-form analyses.

2.3 Variable Construction

Our firm-level variables, including assets, surplus, premiums, and RBC ratios, are taken directly from the NAIC data and are standard. Morningstar classifies contracts into four types of annuities: (1) fixed, (2) fixed-index, (3) index-linked, and (4) traditional variable annuities (TVAs).[17] We track whether each contract is being offered in a given year-quarter using its inception and close dates. Other variables, including mortality and expense risk fee and sales numbers, are taken directly from Morningstar. Contract sales are available only for TVA contracts. See Appendix B for the full definition list and unit conventions.

The annuity contracts in our data, particularly traditional variable annuities (TVAs), offer various guaranteed living benefit riders that shift different types of risk to the insurer.[18] These life guarantees include Guaranteed Minimum Withdrawal Benefits (GMWB), and Guaranteed Lifetime Withdrawal Benefits (GLWB). Each benefit type protects policyholders against different risks across the accumulation, annuitization, and post-retirement phases of the contract. The benefits—particularly GMWB and GLWB—expose insurers to both market risk and longevity risk, making them especially challenging to hedge and manage. We classify the benefits broadly into life guaranteed benefits and non-life guaranteed benefits. Specifically, the benefit names that contain the following words—"GMWB", "GLWB", and "Lifetime"—are considered life guaranteed benefits. Finally, we track whether each benefit rider is active in a given year-quarter using its inception and close dates.

[16] The benefit-level sample exceeds the contract-level sample because TVA contracts typically offer multiple benefit riders, reflecting the common practice of bundling different types of guarantees (e.g., GMWB combined with GMDB) within a single annuity product.

[17] Fixed annuities contain no equity exposure, while fixed-index and index-linked annuities are linked to an index and have downside protection. In our analysis, we group fixed-index and index-linked annuities together as index-linked annuities.

[18] See Appendix A for detailed descriptions of each benefit type.

2.4 Summary Statistics

Table 1: Summary Statistics

Panel A: Firm-level variables (annual)

Variable	Mean	SD	p25	Median	p75	N
Total Assets ($M)	39736.2	65540.8	4432.4	11591.7	37684.0	941
Surplus ($M)	2196.8	3961.1	281.9	760.5	2140.7	941
RBC Ratio (%)	1017.160	2800.515	407.174	489.433	647.487	941
Capital/Assets Ratio (%)	18.562	20.217	7.978	11.135	18.023	941
Has Shadow Insurance	0.312	0.464	0.000	0.000	1.000	941
ERM Rating (1-5)	2.190	0.392	2.000	2.000	2.000	522

Panel B: Contract-level variables (quarterly)

Variable	Mean	SD	p25	Median	p75	N
Traditional Variable Annuity	0.847	0.360	1.000	1.000	1.000	173,744
Fixed Annuity	0.053	0.223	0.000	0.000	0.000	173,744
Index-Linked Annuity	0.100	0.300	0.000	0.000	0.000	173,744
Contract Offered	0.331	0.471	0.000	0.000	1.000	173,744

Panel C: Benefit-level variables (TVA contracts only, quarterly)

Variable	Mean	SD	p25	Median	p75	N
Benefit Offered	0.3882	0.4873	0.0000	0.0000	1.0000	947,900
Life Guarantee	0.5566	0.4968	0.0000	1.0000	1.0000	947,900
Current Fee (bps)	81.99	65.39	30.00	85.00	125.00	947,900
Rollup Rate (bps)	564.62	149.02	500.00	500.00	600.00	359,796

Notes: This table presents summary statistics at the observation-unit frequency used in our reduced-form analyses, covering 2008-2016 for all panels. Panel A reports firm-year statistics (annual frequency) for 941 observations across 120 unique firms, combining NAIC financial data with Morningstar products. Firm-level variables including assets, surplus, and RBC ratios are from NAIC filings. Panel B reports contract-quarter statistics (quarterly frequency) for 173,744 observations across 5,198 unique contracts, representing all annuity contracts from Morningstar. Panel C reports benefit-quarter statistics (quarterly frequency) for 947,900 observations across 27,525 unique benefit-contract combinations and 2,596 contracts, representing benefit riders within TVA contracts from Morningstar. The benefit-level sample exceeds the contract-level sample because TVA contracts typically offer multiple benefit riders.

Table 1 presents summary statistics for our reduced-form analysis sample. Panel A shows firm-level statistics across 941 firm-year observations from 120 unique firms. Firms in our sample exhibit substantial heterogeneity in size and financial strength. Average total assets are $39.7 billion (standard deviation of $65.5 billion), with the median firm holding $11.6 billion in assets, reflecting the presence of both small specialized insurers and large diversified financial institu-

tions in the annuity market. Average surplus is $2.2 billion (standard deviation of $4.0 billion). Among the 522 firm-year observations with ERM ratings, the average rating is 2.19 on the 1-5 scale (standard deviation of 0.39), with the median firm rated 2.0 (Adequate). Shadow insurance usage is present in 31% of firm-year observations, indicating substantial variation in firms' use of affiliated reinsurance structures.

Panel B shows contract-level statistics across 173,744 contract-quarter observations from 5,198 unique contracts. In our sample, 84.7% of contract-quarter observations are traditional variable annuities (TVAs), 5.3% are fixed annuities, and 10.0% are index-linked annuities, reflecting the dominance of variable annuities with life guarantees in the annuity market during this period. On average, 33% of contract-quarter observations show the contract as actively offered in a given quarter, reflecting the panel structure where contracts may exist in firms' product lineups but not be actively marketed in all periods. The ORSA treatment indicator equals one for 31% of observations, reflecting the staggered adoption across states. Strong ERM rating (indicating strong pre-ORSA risk management practices) characterizes 27% of contract-quarter observations, showing substantial variation in firms' ERM sophistication.

Panel C shows benefit-level statistics for benefit riders across 947,900 benefit-quarter observations from 27,525 unique benefit-contract combinations spanning 2,596 contracts. The benefit-level sample is larger than the contract-level sample because each TVA contract typically offers multiple benefit riders, with many contracts bundling different guarantee types (e.g., GMWB combined with GMDB) within a single product. Across all benefit-quarter observations, 55.7% are classified as life guaranteed benefits (GMWB and GLWB), while 44.3% are non-life guaranteed benefits. Among observations with fee data, the average current fee is 0.82% (standard deviation of 0.65%), with the median fee at 0.85%. Among the 359,796 observations with rollup rate data, the average rollup rate is 5.65% (standard deviation of 1.49%), with most benefits offering a 5.0% rollup rate. Contract sales average $250.9 million (standard deviation of $925.3 million) across the 433,384 observations with sales data, though the median is just $7.6 million, reflecting the highly skewed distribution of annuity sales concentrated among a few large products.

3 Reduced-Form Methodology: Imputed Difference-in-Differences

We exploit the staggered adoption of ORSA regulations across states to investigate the impact of ERM on annuity companies using a DiD approach. Specifically, we examine whether the roll-out of ORSA caused insurers to (1) offer fewer TVAs and substitute towards more index-linked annuities; (2) offer fewer of the riskiest benefit types on those TVAs; and (3) increase the fee charged for those risky benefits.

Our DiD analysis replies on two primary assumptions. The first is that in the absence of ORSA, the treatment and control groups would have had parallel trends over time. Any deviation in trends between the two groups after the policy change is then attributed as the causal effect of the policy. While we cannot directly check the parallel post-trends assumptions, we do check for parallel pre-trends using a staggered event study analysis. The second primary assumption is known as the "stable unit treatment value" assumption (or SUTVA). This assumption requires that the response of a particular insurer depends only on the treatment value of that insurer, and not on the treatment value of other insurers. For our analysis, this assumption likely holds in the first order, however there may be second-order feedback mechanisms such as competitive effects. For example, treated firms raising their prices may also lead to non-treated firms raising their prices in response. As long as the "non-treated" impact is in the same direction the true treatment, this will attenuate our results.

In our setting, where states adopt ORSA at different times, the traditional two-way fixed effects (TWFE) DiD approach can yield biased results. This occurs because, in addition to comparing treated insurers with those not yet treated, TWFE also compares insurers incorporated in newly adopting states to those in states that adopted ORSA earlier. If the treatment effects are not homogeneous across time, then TWFE will be biased.[19] Once a state has adopted ORSA, insurers incorporated in that state can no longer serve as valid controls for insurers in states that adopt ORSA later.

Recent developments in the literature have established methods to address this limitation.[20] All of these new methods, some more explicitly than others, rely on a common framework of "imputing" the missing, untreated potential outcomes for treated units by assuming that the observed trend from an appropriate comparison group also happens in the treated group and then averaging, in some manner, the difference between the actual and predicted outcomes for the dependent variable of interest (Borusyak et al., 2024). Where the methods differ is in how the common trends are estimated and how the resulting treatment effects are weighted to get the average treatment effect on the treated (ATT).

Borusyak et al. (2024) formalizes this imputation idea.[21] For all of our reduced form models, we first fit a fixed effects model with OLS using data from insurers in the untreated state-years and then use that model to predict individual counterfactual outcomes for insurers in the treated state-years. We then regress the difference between the actual and predicted outcomes for the dependent variable of interest on the treatment variables. By not including the treated units in

[19] See Goodman-Bacon (2021) for discussions of the drawbacks of two-way fixed effects with differential timing.

[20] For example, Callaway and Sant'Anna (2020), Athey et al. (2021), Borusyak et al. (2024), Gardner (2021), and Jakiela (2021).

[21] For robustness, we also report estimates using two-way fixed effects, the group-time ATT estimator of Callaway and Sant'Anna (2020), and a stacked DiD approach following Cengiz et al. (2019).

the first-stage fixed effects model, we avoid the bias that would arise from the treated units being used as their own control. To account for correlation in treatment assignment, we cluster standard errors at the lead state × annuity type (or guarantee type) level (Abadie et al., 2023).

Contract Offerings: We first examine the opening and closing of individual contracts. We transform our contract-level data into a "long"-balanced panel with one observation per contract-quarter and a binary indicator for whether the contract is offered.[22] We then estimate an imputation-based specification on this contract-level panel.

Formally, we index time at the year-quarter level and write "qt" to denote a specific quarter q in year t (e.g., 2020Q4 and 2021Q4 are distinct). For annuity contract k, (potentially) offered by insurance company j in quarter q of year t, we estimate the following model using the imputation approach:

$$P(y^0_{kjqt} = 1) = \alpha_j + \delta_{a(k)qt} + \epsilon_{kjqt} \tag{1}$$

$$P(y^1_{kjqt} = 1) - \widehat{y^0_{kjqt}} = \sum_{m \in \{TVA, Index, Fixed\}} \beta_m \times ORSA_{jqt} \times \mathbb{1}[Annuity_Type_k = m] + \nu_{kjqt} \tag{2}$$

where y_{kjqt} is a binary indicator for if annuity contract k is offered in quarter q of year t; α_j are firm fixed effects; $\delta_{a(k)qt}$ are annuity-type-by-year-quarter fixed effects, with $a(k) \in \{TVA, Index, Fixed\}$ denoting the type of contract k and qt denoting a specific year-quarter; $Annuity_Type_k$ is the associated one-hot (3-dimensional) type indicator so that $\mathbb{1}[Annuity_Type_k = m] = \mathbb{1}[a(k) = m]$; and β_m captures the differential effect of ORSA on each annuity type.

Equation (1) uses untreated contract-quarters to predict potential counterfactual outcomes for the treated contract-quarters, which are then used in Equation (2) to estimate the differential impact of ORSA adoption on contract offerings across product types.

Benefit Offerings: We next examine the offering of different types of benefits for TVAs. For annuity contract k and benefit i, (potentially) offered by insurance company j in quarter q of year t, we estimate the following model using the imputation approach:

$$P(y^0_{ikjqt} = 1) = \alpha_k + \delta_{g(i)qt} + \epsilon_{ikjqt} \tag{3}$$

$$P(y^1_{ikjqt} = 1) - \widehat{y^0_{ikjqt}} = \beta_G \times ORSA_{jqt} \times LifeGuarantee_i \tag{4}$$

$$+ \beta_N \times ORSA_{jqt} \times (1 - LifeGuarantee_i) + \nu_{ikjqt} \tag{5}$$

[22] By "long" we mean one row per contract-quarter over the sample window, with $y_{kjqt} = 1$ if contract k offered by firm j is available for new sales in quarter q of year t and 0 otherwise. This yields a balanced panel across quarters for each contract so that entries and exits are captured by changes in the offered indicator rather than changes in the set of rows. For example, a contract that is offered in the first 3 periods and then closed in the fourth period will have $y_{kjt} = 1$ for $t = 1, 2, 3$ and $y_{kjt} = 0$ for $t = 4$. This is also known as a "discrete" hazard model.

where y_{ikjqt} is a binary indicator for if benefit i is offered on contract k in quarter q of year t; α_k are contract fixed effects; $\delta_{g(i)qt}$ are life-guarantee-group-by-year-quarter fixed effects, with $g(i) \in \{G, N\}$ indicating whether benefit i is a life guarantee (G) or non-life guarantee (N); $LifeGuarantee_i$ is a binary indicator for whether the benefit is a life guarantee benefit; and $ORSA_{jqt}$ is a binary indicator for insurance firm j being exposed to ORSA regulation in quarter q of year t.

Benefit Fees: Finally, we examine TVAs benefit fees, measured in basis points (bps), conditional on the benefit and contract being actively offered. Using the same benefit-contract-quarter panel, we estimate:

$$Fee^0_{ikjqt} = \alpha_k + \delta_{g(i)qt} + \epsilon_{ikjqt} \tag{6}$$

$$Fee^1_{ikjqt} - \widehat{Fee^0_{ikjqt}} = \gamma_G \times ORSA_{jqt} \times LifeGuarantee_i \tag{7}$$

$$+ \gamma_N \times ORSA_{jqt} \times (1 - LifeGuarantee_i) + \nu_{ikjqt} \tag{8}$$

where Fee_{ikjqt} is the total fee (mortality and expense risk (M&E) plus benefit rider fees) charged for benefit i on contract k in quarter q of year t, measured as an annualized rate in basis points. The specification mirrors the benefit hazard regressions, with contract fixed effects and life-guarantee-group-by-year-quarter fixed effects. The coefficients γ_G and γ_N capture the differential impact of ORSA on fees for life guarantee versus non-life guarantee benefits.

4 Reduced-Form DiD Results

4.1 Baseline Analysis

We begin by estimating average effects of ORSA adoption on insurers' product design and pricing decisions. Our analysis focuses on three margins of adjustment: contract-level offering decisions (whether to offer a particular TVA, RILA, or fixed annuity product), benefit-level design choices (whether to offer a particular guarantee benefit rider in offered TVA products), and pricing adjustments (fees charged on offered guarantee benefit riders in offered TVA products). Table 2 presents our main DiD estimates, while Figures 3–5 provide event study evidence supporting the parallel trends assumption. These event-study figures are constructed by re-estimating the corresponding specifications from Table 2 after replacing the post-adoption indicator with event-time indicators relative to adoption, with $t = -1$ omitted as the reference period. Note that our fee analysis focuses on TVA products because RILAs typically do not charge explicit fees; instead, they compensate for downside protection through participation caps that limit upside returns.

Table 2: Impact of ORSA on Annuity Contract Offering

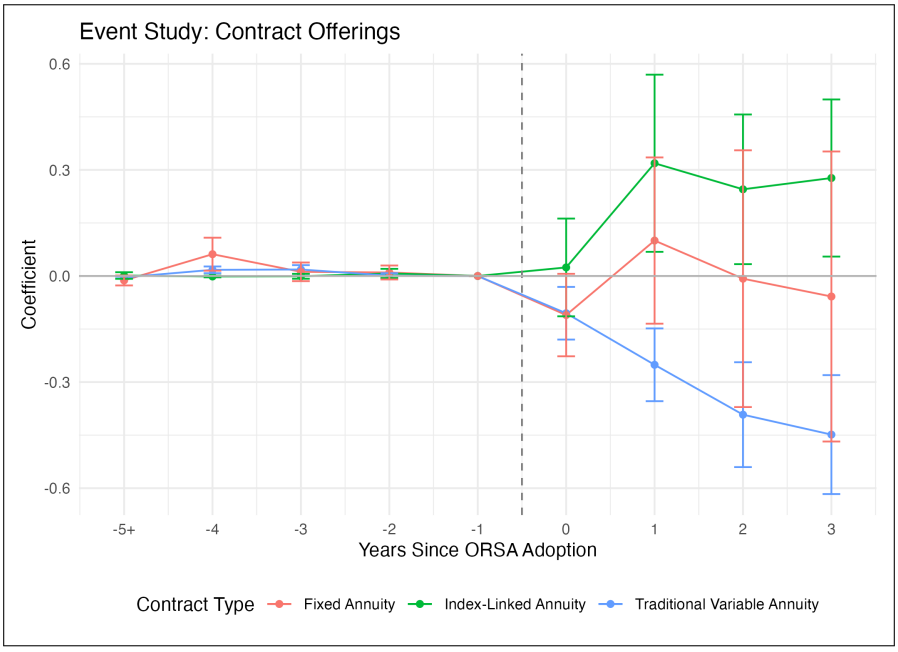
	P(Contract Offered)	P(Benefit Active)	Fee (bps)
Dependent Variables:			
ORSA x Traditional Variable Annuity	-0.249^{***}		
	(0.061)		
ORSA x Fixed Annuity	-0.008		
	(0.131)		
ORSA x Index-Linked Annuity	0.199^{**}		
	(0.098)		
ORSA x Life Guarantee		-0.200^{***}	8.437^{*}
		(0.063)	(4.319)
ORSA x Non-life Guarantee		0.021	1.029
		(0.052)	(4.823)
Firm FE	Yes		
Annuity Type × Year × Quarter FE	Yes		
Contract FE		Yes	Yes
Guarantee Type × Year × Quarter FE		Yes	Yes
Observations	170,296	947,900	178,436

Notes: This table reports reduced-form regression results of the impact of ORSA adoption on annuity product offerings, benefit activation, and benefit fees. Column (1) reports the probability a contract type is offered. Column (2) reports the probability a benefit rider is active (benefit-level) for TVA contracts. Column (3) reports benefit rider fees (in basis points) among active benefits for TVA contracts. The interaction terms capture how the effect of ORSA varies across product and guarantee types. Standard errors are computed via Dirichlet bootstrap (100 replications) clustered at the lead state x annuity type level for Column (1) and the lead state x guarantee type level for Columns (2) and (3). * p < 0.1, ** p < 0.05, and *** p < 0.01.

The contract-type coefficients in column (1) of Table 2 reveal a pronounced portfolio reallocation following ORSA adoption. Insurers reduce their probability of offering TVAs by about 24.9 percentage points and simultaneously increase their probability of offering RILAs by 19.9 percentage points. Both effects are statistically significant. Relative to the pre-ORSA baseline, this represents a significant reduction in the number of TVA offerings and a substantial expansion in RILAs. This pattern suggests direct product substitution, with insurers shifting from mutual fund-like products with guaranteed benefits toward structured products that offer downside protection through indexed participation rather than explicit guarantees. The coefficient on Fixed Annuities is small (-0.8 percentage points) and statistically insignificant, indicating that the primary substitution occurs within equity-linked products rather than toward simpler fixed instruments. The specification has both firm fixed effects and annuity type × year-quarter fixed effects. This means that any time invariant firm-level characteristic, such as the insurer's preferences vis-a-vis the type

of annuity it wants to offer, is absorbed by the fixed effects. Similarly, the interactive fixed effects absorb any other changes in the insurers' or consumers' preferences that vary with time outside of the adoption of ORSA.

Figure 3: Event Study Estimates: Contract Type



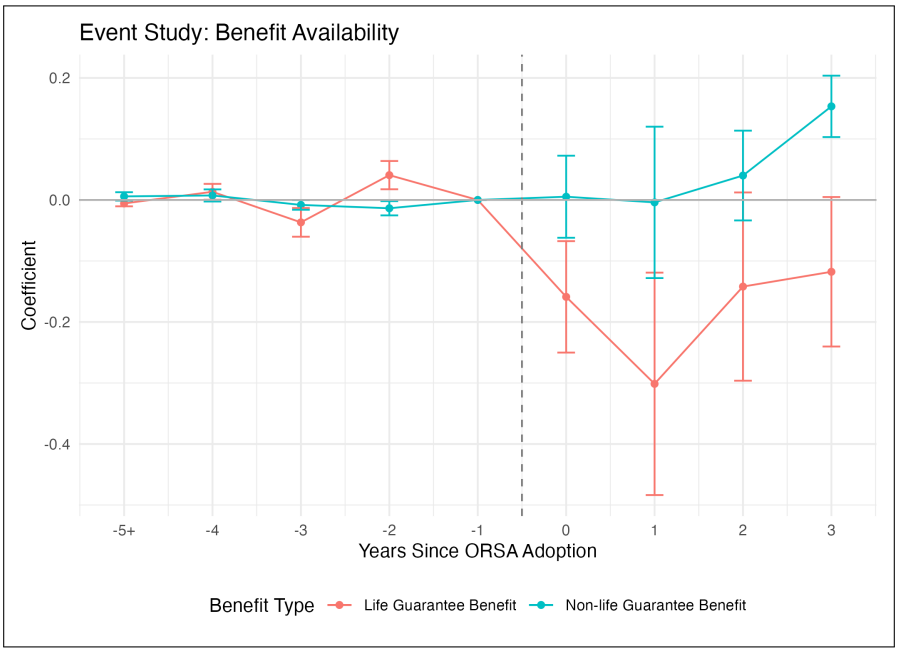
Note: Event-study coefficients are shown for contract types, corresponding to column (1) of Table 2. 90% confidence intervals are shown with standard errors clustered at the lead state x annuity type level.

Figure 3 provides dynamic evidence supporting a causal interpretation of these contract-level shifts. The event study estimates show no differential pre-trends in the five years preceding ORSA adoption, with pre-treatment coefficients centered near zero and statistically indistinguishable from the reference period (t - 1). Following adoption, TVA offerings decline sharply and persistently, while RILA offerings increase in a mirror pattern. The temporal alignment of these effects with ORSA implementation, combined with flat pre-trends, strengthens the inference that regulatory mandates drove portfolio restructuring rather than pre-existing firm trajectories.

Beyond these annuity-type (contract-level) decisions, we also investigate what happens within the TVA product line to study whether insurers start offering less of the TVAs with generous guarantees and offer more of the TVAs with less generous guarantees. To do so, we estimate a specification shown in column (2) of Table 2 that examines benefit-level adjustments within the TVA products that insurers continue to offer. The coefficient on Life Guarantee benefits is -0.200, statistically significant at the 1% level. This estimate implies that ORSA adoption reduces the

probability of offering life guarantee riders by 20.0 percentage points, a decline of 58% relative to the pre-ORSA baseline of 34.3%. In contrast, the coefficient on Non-life Guarantee benefits is small (2.1 percentage points) and statistically insignificant, suggesting that simpler riders like GMDBs and GMABs remain largely unaffected. This differential response indicates that insurers respond to enterprise risk management mandates by selectively removing the most complex and tail-risk-sensitive guarantees while preserving features with more tractable risk profiles. The specification has contract fixed effects, absorbing any contract-specific characteristic that may affect the decision outside of ORSA, and also guarantee type $\times$ year-quarter fixed effects.

Figure 4: Event Study Estimates: Benefit Type

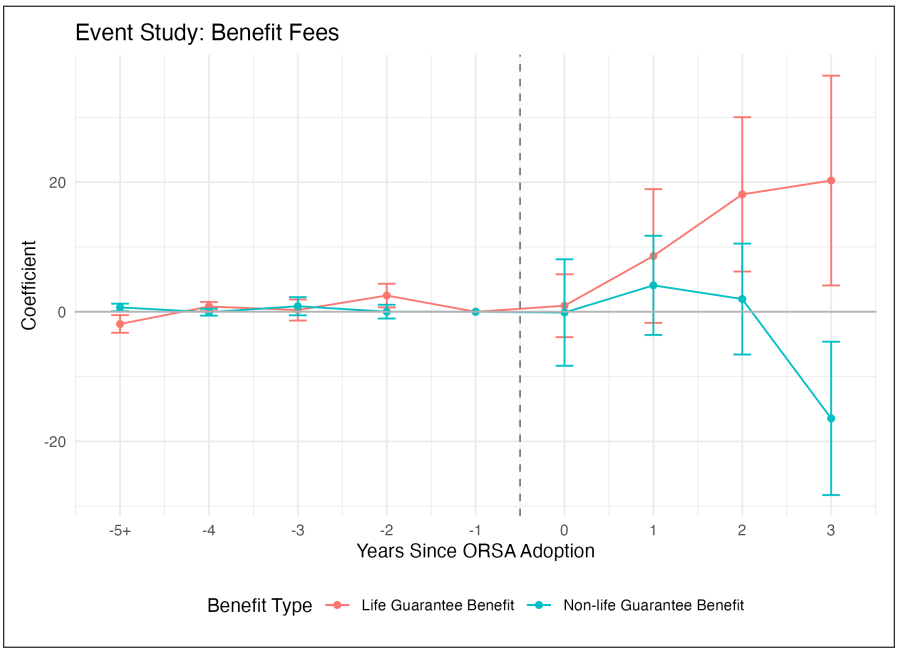


Note: Event-study coefficients are shown for benefit types, corresponding to column (2) of Table 2. 90% confidence intervals are shown with standard errors clustered at the lead state x guarantee type level.

Figure 4 reinforces these findings with event study evidence. Life guarantee TVA offerings exhibit flat pre-trends followed by a sharp, sustained decline beginning at adoption. The absence of differential pre-treatment dynamics supports the parallel trends assumption and suggests that the observed benefit removals reflect regulatory responses rather than secular trends. The event study pattern for non-life guarantees remains relatively flat throughout, consistent with the small and insignificant DiD coefficient.

Finally we investigate the impact on fees charged by the insurers for the TVA offerings. Column (3) examines pricing adjustments on benefits that remain actively offered following ORSA adoption. The coefficient on Life Guarantee fees is 8.4 basis points, statistically significant at the 10% level. Relative to the pre-ORSA baseline of 120.2 basis points, this represents a 7.0% fee increase on the riskiest riders that insurers choose to retain. This intensive-margin adjustment complements the extensive-margin removals documented in column (2), suggesting that insurers employ both quantity (removing products) and price (raising fees) mechanisms to manage tail risk. The coefficient on Non-life Guarantee fees is small (1.029 basis points) and statistically insignificant, indicating that fee adjustments concentrate on the products with the most substantial guarantee obligations. The specification has contract fixed effects, absorbing any contract-specific characteristic that may affect the decision outside of ORSA, and also guarantee type $\times$ year-quarter fixed effects.

Figure 5: Event Study Estimates: Fees



Note: Event-study coefficients are shown for fees, corresponding to column (3) of Table 2. 90% confidence intervals are shown with standard errors clustered at the lead state x guarantee type level.

Figure 5 provides temporal evidence on TVA offerings fee dynamics. Life guarantee fees exhibit no differential pre-trends but increase sharply following ORSA adoption, with effects that persist and potentially grow over subsequent years. This pattern suggests that fee adjustments

respond to regulatory implementation rather than anticipatory pricing changes. Non-life guarantee fees remain relatively flat throughout the event window, consistent with the insignificant DiD estimate.

Collectively, these reduced-form results demonstrate that ORSA induces systematic product redesign across multiple margins. At the portfolio level, insurers substitute from TVAs toward RILAs, reallocating capital toward products with natural hedging properties. Within retained TVA product lines, they remove life guarantee riders that expose them to longevity and market risk simultaneously. On the TVA products they continue offering, they raise fees to better align pricing with the economic costs of embedded guarantees. This comprehensive adjustment pattern is consistent with the hypothesis that ORSA impacts the organizational risk management of insurers: this enterprise risk management mandate leads to a change of insurers' product lines consistent with our maintained hypothesis that insurers start internalizing cross-product risk dependencies rather than evaluating each product line in isolation. By forcing a holistic assessment of solvency risk, ORSA appears to reveal cost complementarities and hedging opportunities that insurers previously overlooked under siloed risk management frameworks.

Demand-based explanations for these three findings, i.e., that consumer preferences rather than ORSA drove the outcomes, are highly implausible. For such an explanation to hold, one would need to argue that consumer preferences for specific annuity features shifted exactly with the timing of ORSA adoption, and only in the states where insurers happened to be incorporated. This scenario is unlikely, given that insurers operate nationally and face largely uniform demand conditions (Wenning, 2024).

One could argue that changes in insurers' annuity product lines reflect their desire to maintain higher RBC ratios under the ORSA Model Law, rather than the ERM principles underlying the law. However, ORSA is designed to alter how insurers assess and manage risk, not simply how much capital they are required to hold. Although its stress-testing requirements could, in principle, affect optimal capital decisions, Appendix E shows that RBC ratios do not increase following ORSA and, if anything, decline. This is particularly true for insurers with ample capital buffers. Our results are consistent with insurers pursuing product strategies that improve their true risk profile (e.g., shifting toward RILAs) but that the regulatory capital formula penalizes because it does not recognize natural hedging benefits. The evidence points to the ERM framework embedded in ORSA, which changes firms' economic assessment of risk, rather than binding regulatory capital constraints as the primary driver of the behavioral responses we document.

Barbu et al. (2025) connect the demand and regulatory channels by demonstrating that RBC rules influence insurance offerings both via insurers' reactions to stricter regulation (supply-side adjustments) and via consumers preferring lower-risk companies (demand-side shifts). If ORSA results in stronger risk management, one would expect insurers to see increased consumer interest

across their product lines, including TVAs, which could drive price changes as well. Our results, however, point to a different pattern: rather than a broad uptick in demand, there is a compositional shift in product mix because insurers expand their issuance of RILAs while cutting back on TVAs.

4.2 Robustness: Specification Curve Analysis

We assess the robustness of our reduced-form estimates using the specification curve methodology of Simonsohn et al. (2020). This approach systematically examines how treatment effect estimates vary across the space of defensible analytical choices, providing a more comprehensive assessment of robustness than traditional sensitivity tables that report only a handful of alternative specifications.

Table 3 summarizes the results across 320 specifications for each of our three main outcomes. We form this grid by crossing (i) eight control sets from toggling three financial controls (log assets, RBC ratio, and surplus ratio), (ii) four DiD estimators (imputation Borusyak et al., 2024, TWFE, Callaway-Sant'Anna, and stacked DiD), (iii) five cohort samples (all cohorts or leaving out one adoption cohort at a time: 2013–2016), and (iv) two samples that either include or exclude firms with New York as lead state, yielding $2^3 \times 4 \times 5 \times 2 = 320$ specifications per outcome. The table reports the median estimate, the 5th–95th percentile range, the percentage of specifications that are statistically significant at the 5% level, and the percentage with the same sign as the median.

Table 3: Specification Curve Analysis: Summary of Robustness

Outcome	N Specs	Median Estimate	[5th, 95th] Percentile	Sig. at 5%	Same Sign	Preferred Spec.
Contract Offered (TVA)	320	-0.209	[-0.272, -0.121]	87.2%	100.0%	-0.249 (0.054)
Benefit Active (Guaranteed)	320	-0.118	[-0.336, -0.065]	94.1%	99.7%	-0.200 (0.016)
Fee (Guaranteed, bps)	320	5.51	[0.15, 12.39]	59.4%	95.6%	8.44 (2.81)

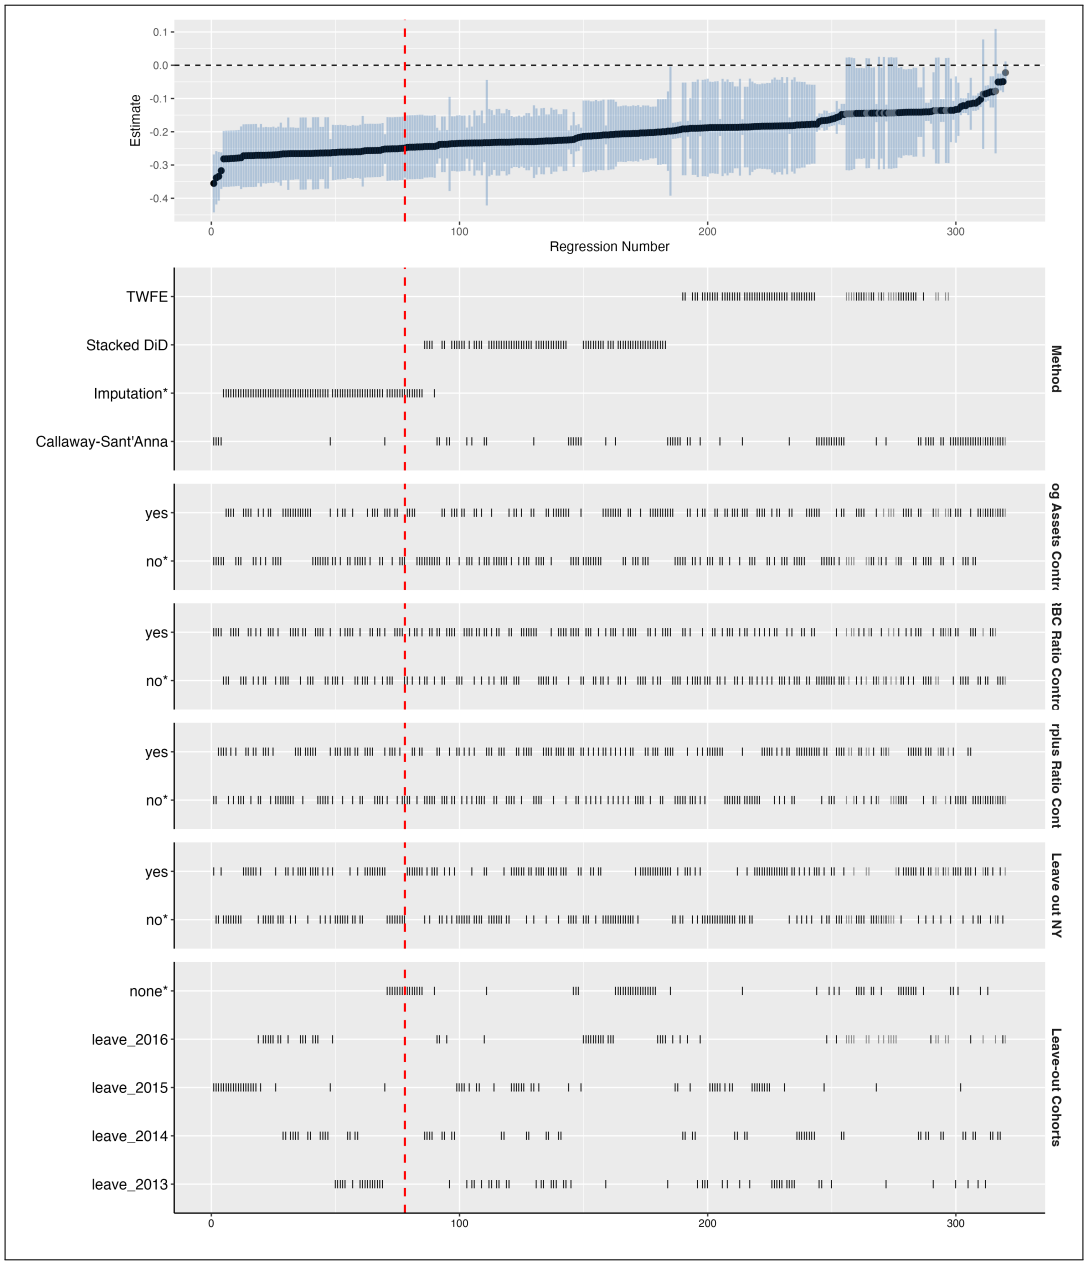
Notes: Specifications vary by control variables (log assets, RBC ratio, surplus ratio), estimation method (imputation, TWFE, Callaway-Sant'Anna, stacked DiD), cohort exclusions (leave-one-out by adoption year), and NY exclusion. "Sig. at 5%" reports the percentage of specifications with $p < 0.05$. "Same Sign" reports the percentage with the same sign as the median. "Preferred Spec." shows the estimate (SE) from the main specification: imputation method, no additional controls, all cohorts included. Contract Offered is a binary indicator for whether the firm offers the product type. Benefit Active is a binary indicator for whether the benefit feature is available. Fees are in basis points (bps) conditional on benefit being active. See Appendix for full specification curve figures.

The specification curve results demonstrate robustness for our key finding that ORSA reduces TVA contract offerings. Across 320 specifications, 100% of estimates are negative and 87.2% are statistically significant at the 5% level. The median estimate of -0.209 lies close to our preferred specification, and the 5th–95th percentile range remains negative and bounded away from zero ($[-0.272, -0.121]$). This consistency indicates that the TVA reduction is not an artifact of particular specification choices.

Figure 6 displays the specification curve for TVA contract offerings. The top panel plots 320 coefficient estimates (dots) sorted from most negative to least negative, with 95% confidence intervals, and the dashed vertical line marks our preferred specification from Table 2. The bottom panels provide a "specification map": each column aligns with an estimate above and marks which analytical choices generated it—the estimator (imputation, TWFE, Callaway-Sant'Anna, or stacked DiD), whether each control is included (Yes/No for log assets, RBC ratio, and surplus ratio), whether firms with New York as lead state are excluded (Yes/No), and which adoption cohort (if any) is left out (none or excluding 2013, 2014, 2015, or 2016). Asterisks indicate the preferred choice within each dimension. Every estimate is negative, and the preferred specification falls near the median of the distribution.

Table 3 also shows that the benefit hazard results (guaranteed benefit availability) are highly robust: 94.1% of specifications are significant at 5%, and 99.7% share the same negative sign. Fee effects are moderately robust, with 59.4% significant and 95.6% same-sign. For the contract hazard and fee specifications, we use Dirichlet bootstrap standard errors for the imputation estimator to account for imputation uncertainty, which widens confidence intervals relative to analytical standard errors. For brevity, we present the TVA offering curve in the main text; Appendix C reports the full set of specification curve figures for the remaining outcomes and product types (including fixed and index-linked contracts, benefit hazards, and fees).

Figure 6: Specification Curve: Effect of ORSA on TVA Contract Offerings



Note: Specification curve for the effect of ORSA adoption on traditional variable annuity (TVA) contract offerings. The top panel displays coefficient estimates from 320 specifications, sorted from most negative to least negative with 95% confidence intervals; the dashed vertical line marks the preferred specification from Table 2. The bottom panels map each estimate to its specification choices: estimator, control inclusion (Yes/No for log assets, RBC ratio, and surplus ratio), New York lead state exclusion (Yes/No), and leave-one-out adoption cohort (none or excluding 2013–2016). Asterisks indicate the preferred choice within each dimension. All estimates are negative, and 87.2% are statistically significant at the 5% level.

4.3 Heterogeneity Analysis

To better understand which firms respond most strongly to ORSA mandates, we examine heterogeneous treatment effects across two dimensions that capture firms' pre-existing risk management sophistication and regulatory arbitrage behavior. Tables 4 and 5 present interaction effects for contract hazard, benefit hazard, and fee outcomes, along the same specifications shown in Table 2 but with the difference that we additionally investigate the differential effect of ORSA on insurers that had high risk management capabilities in the pre-ORSA period. These heterogeneity analyses shed light on the mechanisms through which ORSA influences product design and pricing decisions.

4.3.1 ERM Rating Heterogeneity

We first examine whether firms with stronger pre-ORSA risk management systems respond differently to the regulatory mandate. Insurers that already maintained strong risk-management frameworks were more likely to implement an ERM system following their decision. If their existing ERM practices had already met the stringent ORSA requirements, we would expect ORSA adoption to produce only a modest effect—or none at all.

Table 4 interacts ORSA treatment with an indicator for high ERM rating (measured as "strong" or better on S&P's 5-point scale, measured in 2012 thus before the beginning of the ORSA adoption). We find that the baseline treatment effects demonstrate that insurers with weaker ERM ratings (a) reduce TVA contract offerings by 25.9 percentage points while raising the RILAs offerings by 22.2 percentage points (column (1)), (b) reduce life guarantee benefit offerings within the TVA product line by 21.0 percentage points (column (2)), and (c) raise fees on remaining life guarantees by 8.16 basis points.

The interaction coefficients for contract-level outcomes in column (1) reveal limited differential responses for firms with high ERM ratings. The TVA interaction coefficient is 0.044 and statistically insignificant, suggesting that high-ERM and low-ERM firms reduce TVA offerings by similar magnitudes. For RILAs, the interaction coefficient is -0.111, statistically significant at the 1% level, indicating that high-ERM firms increase RILA offerings by 11.1 percentage points less than low-ERM firms. This yields a total RILA effect of $0.222 - 0.111 = 0.111$ for high-ERM firms, roughly half the expansion observed among low-ERM firms.

In column (2) we present the specification for the benefit hazard within the TVA product line. The coefficient of the interaction variable is 0.030 but statistically insignificant, suggesting that high-ERM firms do not significantly differ in their propensity to remove life guarantee benefits relative to low-ERM firms. In column (3) we present the results for the fees specification. We find

Table 4: Heterogeneous Treatment Effects by ERM Rating (2012)

	Contract Hazard	Benefits Hazard	Fees (bps)
ORSA × Traditional Variable Annuity	-0.259^{***}		
	(0.060)		
ORSA × Traditional Variable Annuity × High ERM	0.044		
	(0.033)		
ORSA × Fixed Annuity	-0.010		
	(0.135)		
ORSA × Fixed Annuity × High ERM	0.010		
	(0.117)		
ORSA × Index-Linked Annuity	0.222^{**}		
	(0.099)		
ORSA × Index-Linked Annuity × High ERM	-0.111^{***}		
	(0.038)		
ORSA × Life Guarantee		-0.210^{***}	8.163^{*}
		(0.065)	(4.555)
ORSA × Life Guarantee × High ERM		0.030	0.643
		(0.029)	(5.655)
ORSA × Non-life Guarantee		0.001	-8.380^{**}
		(0.058)	(3.305)
ORSA × Non-life Guarantee × High ERM		0.092^{*}	26.826^{***}
		(0.047)	(6.240)
Firm FE	Yes		
Annuity Type × Year × Quarter FE	Yes		
Contract FE		Yes	Yes
Guarantee Type × Year × Quarter FE		Yes	Yes
Observations	170,296	947,900	178,436
Firms	120	100	84

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the lead-state level. Contract hazard specifications report coefficients for all annuity types (Traditional VA, Fixed, Index-Linked). Benefits and fees columns report effects for both life and non-life guarantees in the corresponding subsamples. Fixed effects are the same as in the combined hazard table and are not repeated here.

that the coefficient of the interaction term for the life guarantee products is small and statistically insignificant (0.643 basis points), whereas that for the non-life guarantee products is large (26.83, and statistically significant) undoing the decrease that we find for the firms with low-ERM scores (a coefficient of -8.38 and statistically significant at the 5% level).

The absence of meaningful ERM heterogeneity in TVA contract reductions is itself informative: it suggests that the corrective force of ORSA operates similarly across the ERM distribution. Even firms with strong pre-existing risk management frameworks adjust their product portfolios in comparable ways, consistent with the interpretation that ORSA reveals cost interdependencies that were not previously incorporated into product design regardless of ERM sophistication.

The RILA heterogeneity, by contrast, suggests that high-ERM firms expand into RILAs more cautiously, perhaps because their existing risk frameworks already accounted for some of the hedging benefits that RILAs offer.

4.3.2 Shadow Insurance Heterogeneity

Table 5: Heterogeneous Treatment Effects by Shadow Insurance (2010)

	Contract Hazard	Benefits Hazard	Fees (bps)
ORSA × Traditional Variable Annuity	-0.277^{***}		
	(0.063)		
ORSA × Traditional Variable Annuity × Shadow Ins.	0.046		
	(0.040)		
ORSA × Fixed Annuity	0.055		
	(0.154)		
ORSA × Fixed Annuity × Shadow Ins.	-0.112		
	(0.084)		
ORSA × Index-Linked Annuity	0.182^{*}		
	(0.095)		
ORSA × Index-Linked Annuity × Shadow Ins.	0.018		
	(0.057)		
ORSA × Life Guarantee		-0.175^{**}	15.624^{***}
		(0.069)	(4.331)
ORSA × Life Guarantee × Shadow Ins.		-0.032	-9.910^{**}
		(0.032)	(4.301)
ORSA × Non-life Guarantee		-0.023	1.539
		(0.063)	(2.544)
ORSA × Non-life Guarantee × Shadow Ins.		0.070	-0.732
		(0.109)	(6.296)
Firm FE	Yes		
Annuity Type × Year × Quarter FE	Yes		
Contract FE		Yes	Yes
Guarantee Type × Year × Quarter FE		Yes	Yes
Observations	167,568	934,224	176,152
Firms	100	86	75

Notes: *** p<0.01, ** p<0.05, * p<0.1. Standard errors clustered at the lead-state level. Contract hazard specifications report coefficients for all annuity types (Traditional VA, Fixed, Index-Linked). Benefits and fees columns report effects for both life and non-life guarantees in the corresponding subsamples. Fixed effects are the same as in the combined hazard table and are not repeated here.

Table 5 examines heterogeneity by shadow insurance usage, measured as an indicator for firms that employed shadow insurance transactions in 2010. We use shadow insurance as a proxy for the insurer's regulatory arbitrage sophistication and experience managing complex capital structures through affiliated reinsurance arrangements.

The baseline treatment effects for insurers with no shadow insurance (a) reduce TVA contract offerings by 27.7 percentage points while raising the RILA offerings by 18.2 percentage points (column (1)), (b) reduce life guarantee benefit offerings within the TVA product line by 17.5 percentage points (column (2)), and (c) raise fees on remaining life guarantees by 15.62 basis points.

The coefficients of the interaction variables reveal very limited differential responses for shadow insurance users across product offering decisions. From the contract hazard specification (column (1)), the coefficient of the TVA interaction term is 0.046 and statistically insignificant. The RILA interaction is similarly small (0.018) and insignificant. These results show that the impact of ORSA adoption on insurers with high sophistication was very similar to that on less sophisticated firms, indicating similar contract-level adjustments across shadow insurance status.

For benefit hazard (results shown in column (2)), the interaction is -0.032 and statistically insignificant, suggesting that shadow insurance users do not exhibit significantly different propensities to remove life guarantee benefits. However, the results in column (3) for fees show that the coefficient estimate for the interaction variable is -9.910 basis points and statistically significant at the 5% level, implying that shadow insurance users raise fees by 9.91 basis points less than non-shadow firms. This yields a total fee effect of $15.624 - 9.910 = 5.714$ basis points for shadow insurance users.

The pronounced fee heterogeneity indicates that shadow-insurance users experience smaller fee increases following ORSA adoption. This pattern is consistent with the role of shadow insurance as a form of regulatory arbitrage that provides capital relief through affiliated reinsurance arrangements. Because these firms had already shifted portions of their liabilities off balance sheet before ORSA, their effective capital costs were lower, and the incremental tightening imposed by ORSA was less binding. As a result, shadow-insurance users required smaller price adjustments to meet solvency requirements. In other words, firms that previously engaged in regulatory arbitrage were partially insulated from ORSA's impact on pricing.

These heterogeneity patterns provide insight into how firms' prior risk management practices and regulatory arbitrage strategies shape their adaptation to integrated risk management mandates. Contract-level TVA reductions are similar across ERM ratings and shadow insurance status, while shadow-insurance users exhibit more muted fee responses, reflecting differing mechanisms for meeting ORSA's enterprise risk standards.

5 Structural Model Set-Up

We analyze consumer welfare around ORSA adoption with a structural model that links firms' pricing, costs, and demand in the variable annuity market. We begin with the supply side, then describe demand and estimation.

5.1 Supply and Cost Structure

Firms choose prices to maximize profits given demand and perceived marginal costs. True marginal costs combine guarantee and rollup components with a natural-hedge term that depends on non-TVA volumes:

$$c_{jt}^{\text{true}} = \beta_{\text{guarantee}} \cdot \text{Guarantee}_{jt} + \beta_{\text{rollup}} \cdot \text{Rollup}_{jt} + \beta_{\text{outside}} \cdot \text{LogOutsideSales}_{ft}^{\text{std}} + X_{jt}\gamma_X + \theta_f + \tau_t + \xi_{jt} \tag{9}$$

where Guarantee_{jt} indicates lifetime withdrawal guarantees, Rollup_{jt} is the benefit base rollup rate, $\text{LogOutsideSales}_{ft}^{\text{std}}$ is the standardized log of firm f's outside (non-TVA) sales, X_{jt} contains controls, and θ_f, τ_t are firm and year fixed effects. The parameter β_{outside} captures portfolio diversification: $\beta_{\text{outside}} < 0$ implies cost-reducing natural hedges.

Perceived costs blend true components with possible misperceptions via two weights:

$$\begin{aligned} c_{jt}^{\text{perceived}} = {} & \alpha_i^{(1)} \cdot (\beta_{\text{guarantee}} \cdot \text{Guarantee}_{jt}) \\ & + \alpha_i^{(2)} \cdot \beta_{\text{outside}} \cdot \text{LogOutsideSales}_{ft}^{\text{std}} \\ & + \beta_{\text{rollup}} \cdot \text{Rollup}_{jt} \\ & + X_{jt}\gamma_X + \theta_f + \tau_t + \varepsilon_{jt} \end{aligned} \tag{10}$$

with $\alpha_i^{(1)}, \alpha_i^{(2)} \in [0, 1]$ mapping pre-ORSA awareness of guarantee costs and natural-hedge benefits to perceived costs. We parameterize these perceptions as

$$\alpha_i^{(k)} = \begin{cases} \frac{1}{1+\exp(-(\alpha_{\text{base}}^{(k)} + \alpha_{\text{ERM}}^{(k)} \cdot \text{ERM}_i + \alpha_{\text{Shadow}}^{(k)} \cdot \text{Shadow}_i))} & \text{if } t < \text{ORSA}_{ft} \\ 1 & \text{if } t \geq \text{ORSA}_{ft} \end{cases} \quad \text{for } k \in \{1, 2\} \tag{11}$$

so ORSA forces integrated cost recognition with $\alpha_i^{(1)} = \alpha_i^{(2)} = 1$.

Multiproduct firms f then set prices by Bertrand competition:

$$\max_{p_{jt} \in \mathcal{J}_{ft}} \sum_{j \in \mathcal{J}_{ft}} (p_{jt} - c_{jt}(p_{jt})) \cdot s_{jt}(p_t) \cdot M_t \tag{12}$$

with first-order conditions that incorporate hedge feedback:

$$s_{jt} + \sum_{k \in \mathcal{J}_{ft}} (p_{kt} - c_{kt}) \frac{\partial s_{kt}}{\partial p_{jt}} - \sum_{k \in \mathcal{J}_{ft}} \frac{\partial c_{kt}}{\partial p_{jt}} s_{kt} = 0 \tag{13}$$

where $\frac{\partial s_{kt}}{\partial p_{jt}}$ are elements of the market-by-market Jacobian matrices derived from the random coefficients demand model (PyBLP), ensuring that markup calculations fully account for consumer heterogeneity. and, for c_{jt} depending on $\log(1 + Q_{ft}^{\text{non-TVA}})$,

$$\frac{\partial c_{kt}}{\partial p_{jt}} = \alpha_i^{(2)} \cdot \frac{\beta_{\text{outside}}}{\sigma_{\text{out}}} \cdot \frac{1}{1 + Q_{ft}^{\text{non-TVA}}} \cdot \frac{\partial Q_{ft}^{\text{non-TVA}}}{\partial p_{jt}} \quad \text{with} \quad \frac{\partial Q_{ft}^{\text{non-TVA}}}{\partial p_{jt}} \approx M_t \cdot \frac{\partial s_{0t}}{\partial p_{jt}}. \tag{14}$$

Equilibrium prices satisfy the implicit system

$$p_t = c_t(\mathbf{p}_t) + \Omega_t^{-1}(s_t - \mathbf{b}_t), \tag{15}$$

where $\mathbf{b}_t$ has elements $b_{jt} = \sum_{k \in \mathcal{J}_{ft}} \frac{\partial c_{kt}}{\partial p_{jt}} s_{kt}$. In this specification, $\beta_{\text{guarantee}}$ and β_{rollup} capture the marginal costs of guarantees and rollups; β_{outside} quantifies how non-TVA portfolios reduce TVA costs through natural hedging; and the α parameters map ERM strength and Shadow Insurance usage into pre-ORSA perception, with ORSA forcing full integration.

Estimation proceeds by minimizing first-order-condition residuals implied by (13); full details for both demand (BLP) and supply estimation are provided in Appendix D.[23]

5.2 Demand Side

We model consumer demand for variable annuities using a discrete choice framework. Each consumer i in year-quarter t chooses among annuity products j to maximize contemporaneous utility. Throughout, t indexes year-quarter time periods.

$$u_{ijt} = \delta_{jt} + \mu_{ijt} + \varepsilon_{ijt} \tag{16}$$

where δ_{jt} is the mean utility from product j in year-quarter t, μ_{ijt} captures individual-specific deviations from mean utility, and ε_{ijt} is an i.i.d. Type I extreme value error term. The mean utility is specified as:

$$\delta_{jt} = \alpha^p p_{jt} + \lambda_{\text{guarantee}} \text{Guaranteed}_{jt} + \lambda_{\text{rollup}} \text{Rollup}_{jt} + \xi_{jt} \tag{17}$$

[23] Standard errors are computed with a state-level cluster bootstrap. For each of 100 iterations, we resample states (LeadState) with replacement, assign frequency weights to observations, re-estimate the objective, and report the standard deviation across the iterations.

where p_{jt} is the NPV of future expected fees, Guaranteed_{jt} is an indicator for products with Life Guarantees, Rollup_{jt} is the scaled rollup rate for benefit bases, and ξ_{jt} is an unobserved product characteristic not absorbed by year-quarter fixed effects. We convert fees from basis points to present value dollar amounts to properly account for the fact that VA fees are charged over the entire contract life on growing assets.[24] We allow for heterogeneity in both price and guarantee sensitivity by specifying:

$$\mu_{ijt} = \sigma_p \nu_i^p (-p_{jt}) + \sigma_g \nu_i^g \text{Guaranteed}_{jt} \tag{18}$$

where $\nu_i^p, \nu_i^g \sim N(0,1)$ are independent standard normal draws representing unobserved consumer heterogeneity in price and guarantee preferences. We parameterize the price random coefficient using $I(-p_{jt})$ in the formulation, so the individual-level price coefficient is $\alpha_i^p = \alpha^p - \sigma_p \nu_i^p$. Similarly, the guarantee coefficient is $\lambda_{g,i} = \lambda_{\text{guarantee}} + \sigma_g \nu_i^g$. To ensure economically sensible coefficient signs for nearly all consumers (negative price sensitivity, positive guarantee valuation), we impose an upper bound of 0.4 on both σ_p and σ_g during estimation.

Our model relies on three primary assumptions. First, price endogeneity is addressed with an instrument set that shifts costs but is excluded from utility. Following Koijen and Yogo (2022), we use the reserve valuation ratio and the reinsurance share as cost shifters. We also include the ORSA treatment indicator, a within-firm Hausman instrument (group life premium), and an ORSA×ERM interaction. We construct centered instruments via first-stage prediction differencing: for each endogenous variable, the instrument equals the difference between fitted values from a regression with and without the excluded cost shifters, controlling for market and firm fixed effects. For random coefficient identification, we additionally construct sigma instruments using the recentered shift-share approach, then residualize against controls and fixed effects before combining into the final instrument set.

Second, we assume ORSA roll-out is unrelated to consumer demand. This assumption has two parts: (1) we assume that consumers are either unaware of ORSA regulations or, if they are aware, indifferent towards them. This assumption is likely true, since consumers are largely protected from insurer-specific risk through state guarantee funds. We also assume (2) that there is no other event correlated with the ORSA roll-out that also impacted firm demand. This assumption is plausible because ORSA was adopted at the lead-state level and, once in force, applied to the insurer's entire enterprise. Because these firms are large and variable annuities are priced and marketed nationally with common menus and prices across states, the staggered lead-state roll-out is unlikely to be correlated with national demand shocks or with unobserved preference shifts that differentially load on Life or Non-Life Guarantees and rollups.

[24] We assume an expected contract duration of 12 years, apply a 3% annual discount rate, and use an expected asset return of 7% per year to reflect that fees are charged on growing assets under management. The calculations are based on a normalized initial investment of \$100,000. This yields: $p_{jt} = \text{Fee Rate}_{jt} \times \$100,000 \times \sum_{t=1}^{12} \frac{(1.07)^t}{(1.03)^t}$.

Finally, when examining welfare impacts, we adopt a rational-preferences benchmark: consumers have stable, well-behaved preferences over prices and attributes and choose to maximize the utility specified in equations (16)–(17). If, in practice, consumers make systematic mistakes or misperceive features, our estimated elasticities will still capture revealed substitution patterns and remain identified. However, welfare interpretations that treat the estimated utility as normative would then be incorrect. We estimate the demand parameters $\theta_D = \{\alpha^p, \lambda_{\text{guarantee}}, \lambda_{\text{rollup}}, \sigma_p, \sigma_g\}$ using the generalized method of moments (GMM) following Berry et al. (1995). The estimation proceeds through the Berry-Levinsohn-Pakes (BLP) algorithm, which inverts observed shares within each year-quarter to recover mean utilities δ_{jt}, then estimates the linear and nonlinear parameters using instrumental variables. The technical details of the BLP estimation algorithm are provided in Appendix D.

At the quarter level, expected consumer surplus is calculated using the simulation method of Small and Rosen (1981) to account for random coefficients:

$$CS_t = M_t \int \frac{1}{|\alpha_i^p|} \log\left(1 + \sum_{j \in \mathcal{J}_t} \exp(\delta_{jt}(\nu_i))\right) dP(\nu), \tag{19}$$

which we approximate using the same simulation draws employed in estimation. We aggregate these values to annual averages and, where noted, report separately by guarantee segment.

6 Structural Model Results

Table 6 presents the supply-side results from the natural hedge model estimated via FOC residual minimization. This model captures how firms' perceived marginal costs incorporate both guarantee costs and portfolio-level natural hedge benefits. The specification estimates seven key parameters: $\beta_{\text{guarantee}}$, β_{rollup}, and γ (cost parameters), plus $\alpha^{(1)}_{\text{base}}$, $\alpha^{(1)}_{\text{ERM}}$, $\alpha^{(2)}_{\text{base}}$, and $\alpha^{(2)}_{\text{ERM}}$ (dual perception weights). Standard errors, shown in parentheses, are computed via firm-level block bootstrap with 100 iterations.

We estimate the supply side using an innovative natural hedge model that allows firms to misperceive different cost components with different accuracy. Unlike traditional models with a single perception weight, our specification recognizes that firms may better understand some costs (like explicit guarantee features) while significantly underestimating others (like portfolio-level natural hedge benefits). Firms set prices based on perceived marginal costs that blend their understanding of guarantee costs and natural hedge effects:

$$c_{jt}^{\text{perceived}} = \alpha_{ft}^{(1)} \cdot \beta_{\text{guarantee}} \cdot \text{Guarantee}_{jt} + \alpha_{ft}^{(2)} \cdot \gamma \cdot \log(1 + Q_{ft}^{\text{non-TVA}}) + \beta_{\text{rollup}} \cdot \text{Rollup}_{jt} + \theta_f + \tau_t \tag{20}$$

Table 6: Natural Hedge Cost Model: Regression Results

	(1) True Cost Parameters		(2) Perception Parameters
log(1 + Non-TVA Sales)	-0.0010***	$\alpha^{(1)}_{base}$	1.9945***
	(0.0002)		(0.0855)
Life Guarantee	0.0055***	$\alpha^{(1)}_{ERM}$	0.6132***
	(0.0001)		(0.2199)
Rollup Rate	0.0005*	$\alpha^{(2)}_{base}$	0.1001***
	(0.0003)		(0.0154)
		$\alpha^{(2)}_{ERM}$	0.4934***
			(0.0179)
Observations	25,825		
R^2	0.4985		
Fixed Effects	Firm, Year-Quarter, Share Class		
Implied Perception Weights (Pre-ORSA):			
Guarantee: Weak ERM	88.0%		
Guarantee: Strong ERM	93.1%		
Natural Hedge: Weak ERM	52.5%		
Natural Hedge: Strong ERM	64.4%		

Notes: This table presents estimation results from the natural hedge cost model. The natural hedge effect (log(1 + Non-TVA Sales)) captures cost reductions from portfolio diversification. Life Guarantee captures the marginal cost of lifetime withdrawal benefits (GMWB/GLWB). Rollup Rate is the annual percentage increase in the benefit base. $\alpha^{(1)}$ parameters govern perception of guarantee costs, while $\alpha^{(2)}$ parameters govern perception of natural hedge benefits. Pre-ORSA perception weights vary by ERM strength (Weak = ERM rating ≤ 2, and Strong = ERM rating ≥ 3). Post-ORSA, both perception weights equal 1.0 (100%) for all firms. Standard errors in parentheses. $^{*}p < 0.1$, $^{**}p < 0.05$, and $^{***}p < 0.01$. Estimation uses the first-order condition residual minimization with firm and year-quarter fixed effects. Sample: 2008Q1–2016Q4.

The dual perception weights $\alpha^{(1)}_{ft}$ and $\alpha^{(2)}_{ft}$ each range from 0 to 1, measuring how accurately firms perceive guarantee costs and natural hedge benefits, respectively. Both map firm ERM capabilities through logistic links: $\alpha^{(k)}_{ft} = \text{expit}(\alpha^{(k)}_{\text{base}} + \alpha^{(k)}_{\text{ERM}} \cdot \text{ERM}_f)$ for $k \in \{1, 2\}$ pre-ORSA, with ORSA implementation forcing both $\alpha^{(1)} = \alpha^{(2)} = 1$.

Table 6 presents the supply-side parameter estimates. The guarantee cost coefficient $\beta_{\text{guarantee}} = 0.0055$ is positive and statistically significant, confirming that guarantee features increase marginal costs. The rollup rate coefficient $\beta_{\text{rollup}} = 0.0005$ is marginally significant, indicating modest additional costs from step-up provisions. Most notably, the natural hedge parameter $\gamma = -0.0010$ is negative and statistically significant, confirming that larger non-TVA portfolios reduce TVA

marginal costs through portfolio diversification. The economic magnitude indicates a 0.1% cost reduction per unit increase in $\log(1 + Q^{\text{non-TVA}})$, consistent with risk-offsetting benefits from product diversification.

The perception parameters reveal a striking asymmetry in how firms understand their cost structure pre-ORSA. For guarantee costs, firms with weak/adequate ERM ratings (2010) achieved $\alpha^{(1)} = 0.8802$ (88% accuracy), while firms with strong or better ERM achieved $\alpha^{(1)} = 0.9314$ (93% accuracy). In contrast, perception of natural hedge benefits was substantially worse: weak ERM firms had $\alpha^{(2)} = 0.5250$ (only 53% accuracy), while even strong ERM firms achieved just $\alpha^{(2)} = 0.6442$ (64% accuracy).

This differential misperception has important economic implications. Firms relatively accurately priced explicit guarantee costs—features directly visible in product design and actuarial models. However, they substantially underestimated natural hedge benefits arising from portfolio interactions between TVA and non-TVA products. This suggests that organizational silos prevented firms from recognizing cross-product cost interdependencies, even when their ERM frameworks adequately captured product-specific risks.

ORSA's requirement for integrated risk assessment forced both perception weights to unity ($\alpha^{(1)} = \alpha^{(2)} = 1$) for all firms, regardless of pre-ORSA ERM strength. This represents a discrete jump in cost perception, particularly dramatic for natural hedge awareness. Weak ERM firms' understanding of natural hedge benefits increased from 53% to 100%—a near-doubling of accuracy. Even strong ERM firms experienced a 36 percentage point increase (from 64% to 100%).

To assess the economic magnitude of the natural hedge effect of ORSA, we conduct a counterfactual analysis. Figure 7 plots the distribution of the implied marginal-cost increase that would occur if the hedge from non-TVA products were removed. The distribution indicates that the hedge is economically meaningful for a large share of products: the average cost increase is 1.58%, the median is 2.11%, and the 90th percentile is 2.39%, indicating that the impact of cross-product hedging on marginal costs for TVA offerings is economically significant.

The welfare implications of this correction are complex. By forcing firms to recognize natural hedge benefits they previously ignored, ORSA should theoretically reduce perceived costs and encourage TVA product offerings with guaranteed benefits. However, this perception correction occurred simultaneously with other ORSA effects documented in our reduced-form analyses: compliance costs, increased capital requirements, and heightened regulatory scrutiny. The net effect—combining improved perception of cost-reducing natural hedges with increased regulatory burdens—determines the overall market impact.

Figure 7: Marginal Cost Distribution Under Alternative Perception Scenarios

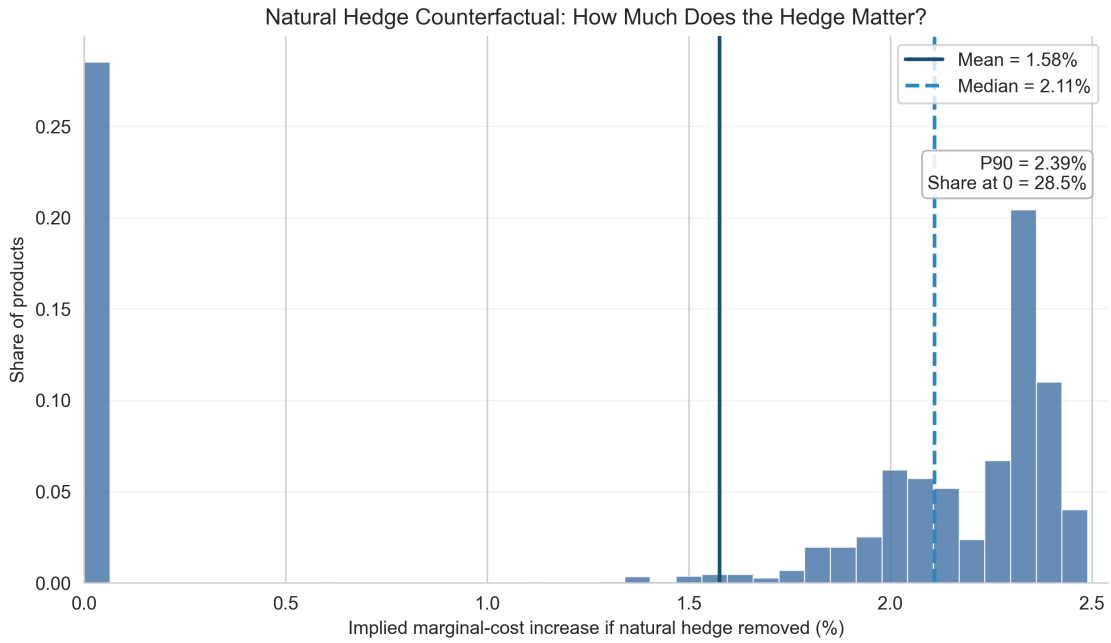


Table 7: BLP Demand Estimation Results

	Coefficient	Heterogeneity
Demand-Side Parameters		
Price	-16.709***	0.145
	(3.870)	(0.705)
Guaranteed	3.513***	0.192***
	(0.563)	(0.063)
Additional Product Chars	Yes	
Observations	24,831	
Firm FEs	Yes	
Year-Quarter FEs	Yes	

Notes: Table reports BLP demand estimation results with firm and market (year $\times$ quarter) fixed effects. Reported values are coefficients with standard errors in parentheses. Random coefficients are estimated for price and guarantee; their standard deviations are shown in the heterogeneity column. Additional product characteristics include A.M. Best rating, rollup rate, and share class indicators. Standard errors are computed using robust two-step GMM. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 7 reports the demand-side results. The price coefficient of -16.71 is negative and precisely estimated, indicating substantial fee sensitivity—a 1-cent increase in the NPV of fees per dollar invested reduces utility substantially. The guarantee coefficient (3.51) is positive and sta-

tistically significant, consistent with consumers valuing downside protection. Overall, consistent with our model, we find that demand responds strongly to prices, indicating consumers do not blindly purchase TVAs. The positive coefficient on guarantees suggests consumers value formal protection, which aligns with standard risk aversion.

Figure 8 summarizes how the structural model links ORSA adoption to changes in market structure and consumer welfare. As ORSA adoption rises in the model, TVA and non-TVA sales and the associated consumer surplus paths adjust to the cost and perception mechanisms documented in Tables 6 and 7.

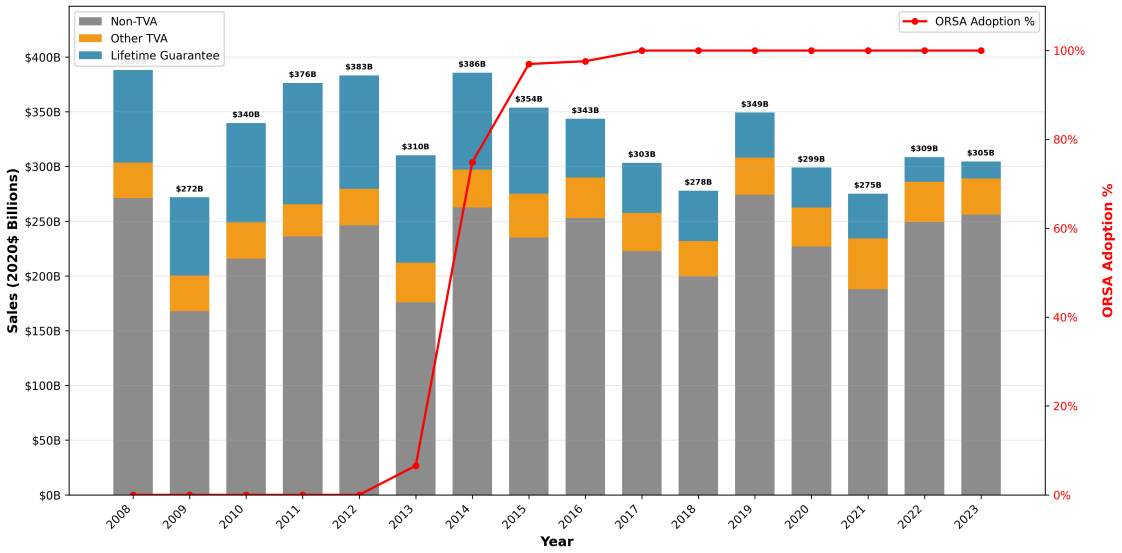
Figure 8 presents a comprehensive overview of how market structure and consumer welfare evolve as ORSA adoption rises. All values are inflation-adjusted to constant 2020 dollars. Panel (a) shows the composition of total annuity sales across product types. Non-TVA products (fixed and immediate annuities) dominate the market, remaining relatively stable at $200–270 billion annually throughout the period. TVA sales—comprising lifetime guarantee products (GMWB/GLWB) and other TVA products—contract substantially from approximately $120 billion in 2008 to around $50–70 billion by 2017–2023. Within the TVA segment, lifetime guarantee products experience the sharpest declines, falling from roughly $90 billion pre-ORSA to approximately $20–40 billion post-ORSA, while other TVA products remain relatively stable at $30–50 billion.

These quantity responses are consistent with the incentives implied by the natural hedge cost estimates. Guarantees are genuinely costly, so the direct guarantee channel pushes TVA prices up as firms more fully account for guarantee risk under ORSA. At the same time, the interaction between TVA guarantees and non-TVA portfolios gives firms an additional motive to tilt their product mix toward non-TVAs that hedge guarantee exposures. In equilibrium, keeping TVA prices high both covers guarantee costs and nudges marginal consumers toward non-TVA products, so the natural-hedge channel also works in the direction of higher TVA prices and a shift away from guaranteed TVAs.

Panel (b) shows how these market structure changes translate into consumer welfare from TVA products. Total consumer surplus peaks at $24.0 billion in 2011, then declines to $15.7 billion in 2014 and $15.2 billion in 2015 as ORSA adoption accelerates, with the largest losses concentrated in lifetime guarantee products. By 2016–2017, annual consumer surplus falls to $8.2–9.1 billion. After ORSA adoption reaches near-full coverage, consumer surplus partially rebounds to $10.8–11.8 billion in 2018–2019 and reaches $13.6 billion in 2021, before declining again to $7.5 billion by 2023. Overall, welfare remains well below the pre-ORSA peak, consistent with persistent contraction in lifetime guarantee products.

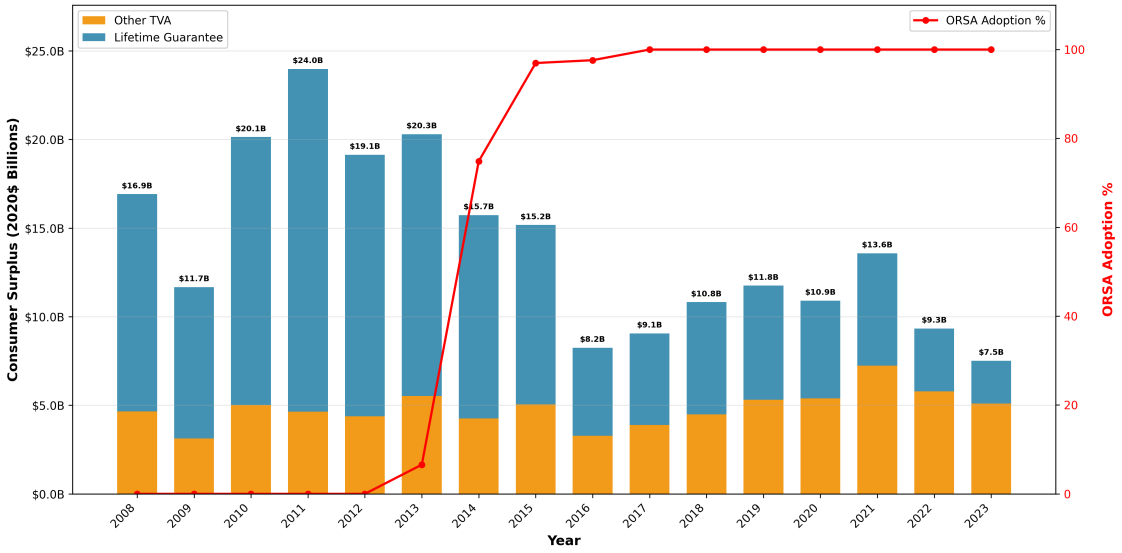
Figure 8: Market Trends and Consumer Welfare Effects During ORSA Adoption Period

(a) Annuity Sales by Type



Note: Values in constant 2020 dollars. TVA sales from Morningstar; Non-TVA from SNL (total annuity premiums minus TVA).

(b) Consumer Surplus by Benefit Type



Note: Values in constant 2020 dollars. Consumer surplus computed from BLP demand estimates for TVA products only.

7 Conclusions

This paper examines how ORSA requirements affect insurance companies' annuity product offerings and consumer welfare. Exploiting the staggered adoption of ORSA across U.S. states from 2013 to 2017, we find that risk management mandates reshape product markets through

the correction of firms' cost misperceptions, with substantial welfare consequences. Our analysis documents sharp changes in product offerings following ORSA adoption: the probability of offering traditional variable annuities declines by 24.9 percentage points, and, within those, lifetime guarantee riders fall by 20 percentage points, with the most complex guarantees experiencing the largest reductions, while index-linked annuities increase by 19.9 percentage points. These effects vary systematically with firms' pre-ORSA ERM capabilities, with weak ERM firms reducing TVA offerings more significantly than strong ERM firms. Welfare losses and substitution toward RILAs are consistent with our structural estimates indicating that non-TVA activity reduces TVA marginal costs through natural hedging.

Our structural estimates reveal that pre-ORSA, firms substantially under-recognized portfolio-level natural-hedge benefits (about 53% accuracy in weak-ERM firms versus about 64% in strong-ERM), while recognizing guarantee costs at roughly 88–93%. In our model, ORSA forces all firms to fully internalize the modeled cost components, eliminating these pricing distortions; our model does not separately estimate a post-ORSA compliance-cost shifter; price increases arise from corrected cost perception and natural-hedge interactions, absorbed alongside fixed effects.

The welfare implications are substantial and negative for consumers. Consumer surplus peaks at $24.0B in 2011 and falls to $8.2–9.1B by 2016–2017 as ORSA adoption reaches near-full coverage; it partially rebounds to $13.6B in 2021 but declines again to $7.5B by 2023. This implies annual welfare losses on the order of $6–17B under our logit welfare measure. This welfare loss reflects a fundamental tension: while the regulation successfully corrects firms' cost misperceptions and improves risk management, the resulting price increases more than offset any benefits consumers receive from better-priced products. The substantial reduction in guaranteed annuity offerings suggests meaningfully reduced systemic risk, achieving ORSA's macroprudential objectives. However, this stability comes at substantial cost to retirement savers who relied on guaranteed annuities for retirement security, transferring investment risk back to consumers at precisely the life-cycle stage when they are least equipped to bear it.

This paper contributes to the literature on financial regulation by providing the first comprehensive analysis of risk management mandates' effects on product markets. While prior work examines regulations targeting capital requirements or disclosure, we show how internal risk management mandates reshape firm behavior through the correction of cost misperceptions. Our structural framework provides a tractable method for quantifying these misperceptions and evaluating whether regulatory corrections improve outcomes. The heterogeneous effects by firm ERM strength suggest opportunities for more targeted regulatory design, imposing lighter compliance burdens on sophisticated firms while focusing intensive oversight on those with weaker risk management. As policymakers worldwide strengthen financial regulation in response to systemic risk concerns, the challenge lies in designing regulations that achieve an appropriate balance between

stability and consumer welfare. Future research should examine ORSA's long-run effects as the regulation matures and explore tiered regulatory structures that preserve stability benefits while minimizing compliance costs.

References

Abadie, A., S. Athey, G. W. Imbens, and J. M. Wooldridge (2023). When should you adjust standard errors for clustering? The Quarterly Journal of Economics 138(1), 1–35.

Acharya, V. V., A. N. Berger, and R. A. Roman (2018). Lending implications of us bank stress tests: Costs or benefits? Journal of Financial Intermediation 34, 58–90.

Aiyar, S., C. W. Calomiris, J. Hooley, Y. Korniyenko, and T. Wieladek (2014). The international transmission of bank capital requirements: Evidence from the uk. Journal of Financial Economics 113(3), 368–382.

Almazora, L. (2024, February). U.s. annuity sales hit a fresh record in 2023: Limra. InvestmentNews.

American Council of Life Insurers (2024). Life Insurers Fact Book 2024. Washington, DC: American Council of Life Insurers. Chapter 3 (Liabilities) and Chapter 8 (Annuities). In 2023, annuity and supplementary contract reserves equal 66% of total life insurer policy reserves (about two-thirds of liabilities).

Amiraslani, H., C. Deller, C. D. Ittner, and T. Keusch (2025). Board risk oversight and environmental and social performance. Journal of Accounting and Economics 79(2-3), 101754.

Ankirchner, S., J. C. Schneider, and N. Schweizer (2014). Cross-hedging minimum return guarantees: Basis and liquidity risks. Journal of Economic Dynamics and Control 41, 93–109.

Athey, S., M. Bayati, N. Doudchenko, G. Imbens, and K. Khosravi (2021). Matrix completion methods for causal panel data models. Journal of the American Statistical Association, 1–15.

Balasubramanyan, L., N. D. Daniel, J. G. Haubrich, and L. Naveen (2024). Impact of risk oversight functions on bank risk: Evidence from the dodd-frank act. Journal of Banking & Finance 158, 107049.

Barbu, A. (2023). Ex-post loss sharing in consumer financial markets. Available at SSRN 4079524.

Barbu, A., D. Humphry, and I. Sen (2025). The evolution of insurance markets, capital regulation and insurance provision. Working Paper.

Barbu, A. and I. Sen (2024, June). Hedging through product design. Working Paper.

Basel Committee on Banking Supervision (2010). Principles for enhancing corporate governance. Basel: Bank for International Settlements Communications. Accessed August 3rd, 2015.

Bauer, J. (2020). Hedging of variable annuities under basis risk. Asia-Pacific Journal of Risk and Insurance 14(2), 20190040.

Berger, A. N., C. H. Bouwman, T. Kick, and K. Schaeck (2016). Bank liquidity creation following regulatory interventions and capital support. Journal of Financial Intermediation 26, 115–141.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. Econometrica 63(4), 841–890.

Berry-Stölzle, T. R. and J. Xu (2018). Enterprise risk management and the cost of capital. Journal of Risk and Insurance 85(1), 159–201.

Bodnar, G. M., E. Giambona, J. R. Graham, and C. R. Harvey (2019). A view inside corporate risk management. Management Science 65(11), 5001–5026.

Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting event study designs: Robust and efficient estimation. Review of Economic Studies, rdae007.

Braumann, E. C., I. Grabner, and A. Posch (2020). Tone from the top in risk management: A complementarity perspective on how control systems influence risk awareness. Accounting, Organizations and Society 84, 101128.

Callaway, B. and P. H. Sant'Anna (2020). Difference-in-differences with multiple time periods. Journal of Econometrics.

Cengiz, D., A. Dube, A. Lindner, and B. Zipperer (2019). The effect of minimum wages on low-wage jobs. The Quarterly Journal of Economics 134(3), 1405–1454.

Chen, X., E. Higgins, H. Xia, and H. Zou (2020). Do financial regulations shape the functioning of financial institutions' risk management in asset-backed securities investment? The Review of Financial Studies 33(6), 2506–2553.

Cortés, K. R., Y. Demyanyk, L. Li, E. Loutskina, and P. E. Strahan (2020). Stress tests and small business lending. Journal of Financial Economics 136(1), 260–279.

COSO (2004). Executive summary: Enterprise Risk Management - Integrated Framework. New York: AICPA. Committee of Sponsoring Organizations of the Treadway Commission (COSO).

COSO (2017). Enterprise risk management – integrating with strategy and performance. New York: American Institute of Certified Public Accountants. Committee of Sponsoring Organizations of the Treadway Commission (COSO).

Eastman, E. M., A. C. Ehinger, and J. Xu (2024). Enterprise risk management and corporate tax planning. Journal of Risk and Insurance 91(3), 529–566.

Eastman, E. M. and J. Xu (2021). Market reactions to enterprise risk management adoption, incorporation by rating agencies, and orsa act passage. Risk Management and Insurance Review 24(2), 151–180.

Eber, M. and C. Minoiu (2016). How do banks adjust to stricter supervision? Available at SSRN 2662502.

Eckles, D. L., R. E. Hoyt, and S. M. Miller (2014). The impact of enterprise risk management on the marginal cost of reducing risk: Evidence from the insurance industry. Journal of Banking & Finance 49, 409–423.

Ellis, C., T. Moenig, and J. Volkman-Wise (2025). Registered index-linked annuities in qualified retirement plans. Journal of Risk and Insurance.

Ellul, A., C. Jotikasthira, A. Kartasheva, C. T. Lundblad, and W. Wagner (2022). Insurers as asset managers and systemic risk. The Review of Financial Studies 35(12), 5483–5534.

Ellul, A. and V. Yerramilli (2013). Stronger risk controls, lower risk: Evidence from us bank holding companies. The journal of finance 68(5), 1757–1803.

Frame, W. S., A. Mihov, and L. Sanz (2020). Foreign investment, regulatory arbitrage, and the risk of us banking organizations. Journal of Financial and Quantitative Analysis 55(3), 955–988.

Friedberg, L. and A. Webb (2021). A primer on variable, fixed indexed, and registered index-linked annuities: What economists and financial professionals need to know. Research Paper for the Retirement Income Institute at the Alliance for Lifetime Income.

Gardner, J. (2021). Two-stage differences in differences. Working Paper.

Giambona, E., A. Kumar, and G. M. Phillips (2025). Hedging, contract enforceability, and competition. The Review of Financial Studies 38(7), 2034–2087.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. Journal of Econometrics 225(2), 254–277.

Houston, J. F., C. Lin, and Y. Ma (2012). Regulatory arbitrage and international bank flows. The Journal of Finance 67(5), 1845–1895.

Hoyt, R. E. and A. P. Liebenberg (2011). The value of enterprise risk management. Journal of risk and insurance 78(4), 795–822.

Hufeld, F., R. S. Koijen, and C. Thimann (2017). The economics, regulation, and systemic risk of insurance markets. Oxford University Press.

Jakiela, P. (2021). Simple diagnostics for two-way fixed effects. arXiv preprint arXiv:2103.13229.

Jiang, L. and M. Ji (2024). Risk management committee and bank performance: Evidence from the adoption of the dodd–frank act. Journal of Business Finance & Accounting 51(7-8), 1762–1788.

Kashyap, A., R. Rajan, J. Stein, et al. (2008). Rethinking capital regulation. Maintaining stability in a changing financial system 43171.

Kim, D. and A. M. Santomero (1988). Risk in banking and capital regulation. The journal of finance 43(5), 1219–1233.

Koijen, R. S. and M. Yogo (2015). The cost of financial frictions for life insurers. American Economic Review 105(1), 445–475.

Koijen, R. S. and M. Yogo (2016). Shadow insurance. Econometrica 84(3), 1265–1287.

Koijen, R. S. and M. Yogo (2022). The fragility of market risk insurance. The Journal of Finance 77(2), 815–862.

Krupa, T. J. (2024). Do enterprise risk assessments affect tax aggressiveness? evidence from us orsa regulation. The Accounting Review 99(2), 307–339.

Laeven, L. and R. Levine (2009). Bank governance, regulation and risk taking. Journal of financial economics 93(2), 259–275.

Landier, A., D. Sraer, and D. Thesmar (2009). Financial risk management: When does independence fail? American Economic Review 99(2), 454–458.

Li, C., K. Stack, L. Sun, and J. Xu (2025). Enterprise risk management and management earnings forecasts. Management Science. Forthcoming.

Li, Z. (2024). Long rates, life insurers, and credit spreads. Life Insurers, and Credit Spreads (May 25, 2024).

Liu, S. and J. Xu (2024). Enterprise risk management, risk-taking, and macroeconomic implications: Evidence from bank mortgage loan management. Journal of Financial Services Research, 1–30.

Mezger, M. P. (2024). The appeal of registered index-linked annuities for myopic prospect theory investors. Available at SSRN 4193584.

Moenig, T. (2022). It's rila time: An introduction to registered index-linked annuities. Journal of Risk and Insurance 89(2), 339–369.

Moenig, T. and B. Samuelson (2024). A comparison of index-linked annuities. North American Actuarial Journal 28(1), 104–125.

Moenig, T. and C. Xu (2023). Valuing lifetime withdrawal guarantees in rilas. North American Actuarial Journal 27(4), 771–786.

NAIC (2012). Risk management and own risk and solvency assessment model act. https://content.naic.org/sites/default/files/model-law-505.pdf. NAIC, Accessed: 2024-11-05.

National Association of Insurance Commissioners (2025). Direct premiums written by line, life/annuity industry, 2022–2024. NAIC data as reported in Insurance Information Institute, "Facts + Statistics: Life insurance". Table "Direct Premiums Written by Line, Life/Annuity Industry, 2022–2024"; annuities account for 55.8% of life/annuity direct premiums written in 2024. Accessed 14 November 2025.

Neel, M. and J. Xu (2025). Does enterprise risk management bolster investor confidence? evidence from options-based restatement contagion, investment, and misstatements. Contemporary Accounting Research. Forthcoming.

Nocco, B. W. and R. M. Stulz (2022). Enterprise risk management: Theory and practice. Journal of applied corporate finance 34(1), 81–94.

Ongena, S., A. Popov, and G. F. Udell (2013). "when the cat's away the mice will play": Does regulation at home affect bank risk-taking abroad? Journal of Financial Economics 108(3), 727–750.

Rogers, J. B. C. and C. Tonetti (2023). Risky insurance: Life-cycle insurance portfolio choice with incomplete markets.

SEC (2007). Risk management for broker-dealers. https://www.sec.gov/news/speech/2007/spch112807mag.htm. Securities and Exchange Commission (SEC), Accessed: 2025-02-04.

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

Simonsohn, U., J. P. Simmons, and L. D. Nelson (2020). Specification curve analysis. Nature Human Behaviour 4(11), 1208–1214.

Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. Econometrica: Journal of the Econometric Society, 105–130.

Stulz, R. M. (2008). Risk management failures: What are they and when do they happen? Journal of Applied Corporate Finance 20(4), 39–48.

Wenning, D. (2024). Equal prices, unequal access: The effects of national pricing in the us life insurance industry. Working Paper.

Appendix A Detailed Annuity Benefit Descriptions

This appendix provides comprehensive descriptions of the various guaranteed benefit types commonly offered within variable annuity contracts. These benefits are central to understanding how insurers manage risk across different product lines and how regulatory changes affect product offerings.

Appendix A.1 Contract Phases and Risk Protection

To understand the purpose of these benefits, it is helpful to consider the three phases of a variable annuity contract:

1. **Accumulation phase** (pre-retirement): The period when the policyholder contributes funds and the contract value grows based on underlying investment performance.
2. **Annuitization phase**: The point at which the policyholder converts the accumulated value into a stream of income payments.
3. **Post-retirement phase** (retirement/withdrawal period): The period when the policyholder receives income from the contract, either through systematic withdrawals or annuitized payments.

Different guaranteed benefits protect against different risks across these phases, creating varying degrees of risk exposure for insurers.

Appendix A.2 Guaranteed Benefit Types

The following guaranteed benefits are commonly offered within traditional variable annuity (TVA) contracts:

Appendix A.2.1 Guaranteed Minimum Death Benefit (GMDB)

The GMDB provides a minimum benefit to a beneficiary upon the policyholder's death before retirement or annuitization. While it introduces some market risk during the accumulation phase, it is typically triggered only at death, so its risk exposure can be lower compared to lifetime guarantees like GLWB or GMIB. This benefit essentially ensures that beneficiaries receive at least a specified minimum amount regardless of market performance, protecting against the scenario where account values have declined due to poor investment returns.

Appendix A.2.2 Guaranteed Minimum Accumulation Benefit (GMAB)

The GMAB protects the policyholder during the accumulation phase by ensuring that the contract's account value will not fall below a specified amount at the end of a set holding period. Since it provides a floor on the contract value, it exposes the insurer to market risk if the underlying investments perform poorly. This benefit is particularly valuable during extended market downturns, as it guarantees that the policyholder will have a minimum accumulation value regardless of interim market volatility. The insurer bears the risk that market performance may require substantial payments to honor the guarantee at maturity.

Appendix A.2.3 Guaranteed Minimum Income Benefit (GMIB)

The GMIB protects the policyholder at the annuitization phase by allowing them to annuitize at a guaranteed rate, effectively locking in a minimum lifetime income that extends into the post-retirement phase. Because it combines longevity risk with market risk, GMIB can be more complex and costly for insurers to hedge. The benefit ensures that even if the account value has decreased substantially, the policyholder can still convert their contract into a lifetime income stream based on a predetermined formula, often referencing the highest historical account value or a guaranteed growth rate.

Appendix A.2.4 Guaranteed Minimum Withdrawal Benefit (GMWB)

The GMWB protects the policyholder during the post-retirement phase by guaranteeing a minimum level of withdrawals over a defined period, regardless of market fluctuations. The insurer bears the risk that the underlying account value may deplete faster than expected, particularly during market downturns. This benefit allows policyholders to withdraw a specified percentage of their benefit base annually without penalty, even if the account value falls to zero. The guarantee typically continues until the total guaranteed withdrawals have been made, providing income certainty over a fixed period.

Appendix A.2.5 Guaranteed Lifetime Withdrawal Benefit (GLWB)

The GLWB provides lifetime protection during the post-retirement phase by extending GMWB to a lifetime horizon, thereby adding significant longevity risk on top of equity market risk. As a result, GLWB is often considered one of the most challenging guarantee types to hedge cost-effectively. This benefit combines the withdrawal flexibility of GMWB with lifetime income pro-

tection, ensuring that policyholders can continue to receive guaranteed withdrawals for as long as they live, regardless of account performance or longevity. The open-ended nature of this commitment makes it particularly expensive and risky for insurers.

Appendix A.3 Risk Implications for Insurers

Among these benefits, the living benefits that continue throughout the policyholder's lifetime (notably GMIB and GLWB) tend to be the most complex and risky for insurers, since they must account for both uncertain market returns and uncertain lifespans. By contrast, GMDBs, although still risky, generally have more predictable triggers and shorter time horizons.

Although, in theory, insurers could hedge this risk through derivatives or reinsurance, practical and regulatory barriers often render such hedging strategies costly or incomplete. In the wake of the Great Financial Crisis and the COVID-19 Pandemic, many insurers faced severe financial distress on their TVA exposures, leading to higher prices, renegotiation of contract terms, and even exits from the market (Koijen and Yogo, 2022; Barbu, 2023).

Appendix A.4 Classification for Analysis

In our analysis, we focus on the "lifetime" benefits (as they are colloquially called in the industry) of GMWB and GLWB and treat the remaining benefits as "non-lifetime" benefits. This classification reflects the fundamental difference in risk exposure between benefits that provide protection for a fixed period versus those that continue for the policyholder's entire lifetime. Our results are robust to also treating GMAB and GMIB as "lifetime" benefits, though the magnitudes weaken somewhat, reflecting the intermediate risk profile of these benefits that combine elements of both accumulation protection and lifetime income features.

Appendix B Definition of Variables

This appendix lists the core variables used across our analyses.

Table B1: Variables used in the analyses

Variable Name	Definition	Source
Total assets	Firm total assets.	NAIC
Policyholder surplus	Policyholder surplus (capital and surplus).	NAIC
RBC ratio	Risk-based capital ratio, calculated as the (adjusted) statutory capital divided by the required risk-based capital.	NAIC
Capital and surplus to assets ratio	Capital plus surplus divided by assets.	NAIC
ERM rating	S&P ERM rating based on five categories—risk management culture, risk controls, emerging risk management, risk models, and strategic risk management—and mapped onto a five-point scale from Weak (1) to Very Strong/Excellent (5). A rating of 3–5 is classified as strong, and otherwise as weak.	S&P Global
Shadow insurance (indicator)	= 1 if the insurer uses affiliated and unauthorized reinsurers without an A.M. Best rating, and 0 otherwise, following Koijen and Yogo (2016).	NAIC
TVA sales	Quarterly sales of traditional variable annuity (TVA) contracts.	Morningstar
Non-TVA sales	Quarterly total sales of non-TVA products, including index-linked annuities and fixed annuities.	NAIC
Annuity product type	Product category for the contract, including TVAs, index-linked annuities, and fixed annuities.	Morningstar
Contract offered (indicator)	= 1 if the annuity contract (TVAs, index-linked annuities, or fixed annuities) is offered in a given year-quarter, and 0 otherwise, based on contract open and close dates.	Morningstar

Variable Name	Definition	Source
Benefit active (indicator)	= 1 if the benefit rider (lifetime vs. non-lifetime guarantees) of a TVA contract is offered in a given year-quarter, and 0 otherwise, based on benefit availability dates.	Morningstar
Benefit fee (current rate)	Current fee rate of a given benefit rider (lifetime vs. non-lifetime guarantees) on a TVA contract, derived from historical Morningstar prospectuses.	Morningstar
Guaranteed (indicator)	= 1 if the annuity product has lifetime guaranteed benefit riders, and 0 otherwise.	Morningstar
Roll-up rate	Annual rate at which an annuity's benefit base increases over time, regardless of investment performance.	Morningstar

Appendix C Additional Specification Curve Analyses

This appendix presents specification curve analysis following Simonsohn et al. (2020) for our main reduced-form results. Specification curves visualize how estimated treatment effects vary across reasonable analytical choices, providing a comprehensive robustness assessment beyond standard sensitivity tables.

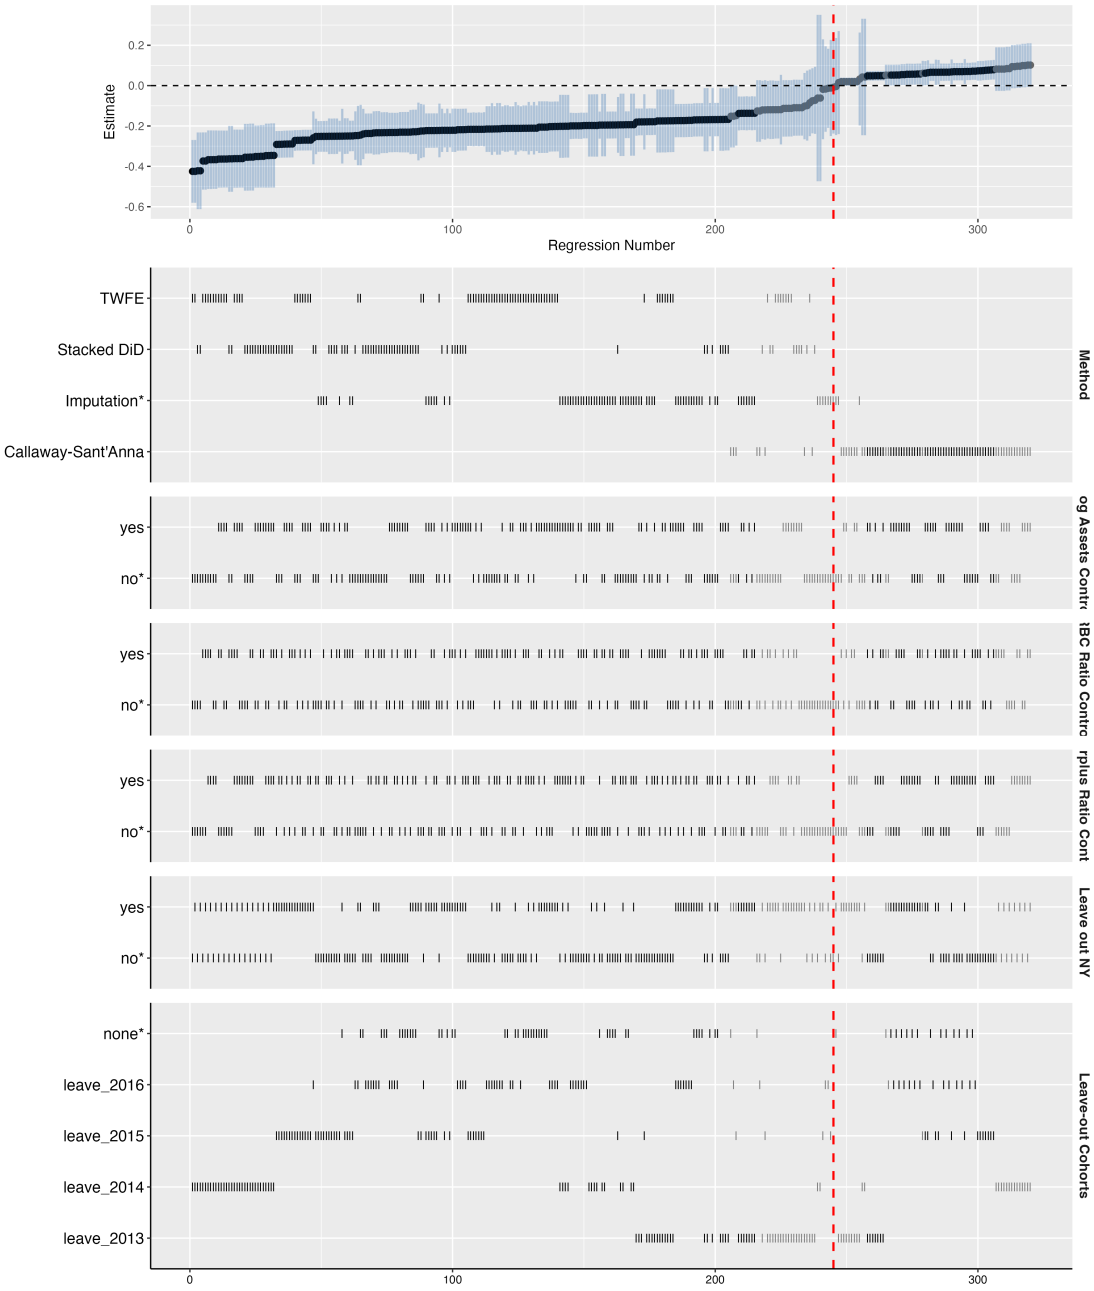
For each analysis, we systematically vary:

- **Control variables**: Each control (log total assets, risk-based capital ratio, capital surplus to assets ratio) is independently toggled on or off, creating $2^3 = 8$ control combinations
- **Estimation method**: Imputed difference-in-differences (Borusyak et al., 2024), standard two-way fixed effects, Callaway and Sant'Anna (Callaway and Sant'Anna, 2020), or stacked DiD
- **Lead state exclusion**: Including or excluding firms with New York as lead state
- **Leave-out cohorts**: All adoption cohorts included, or leave-one-out by adoption year (2013, 2014, 2015, or 2016) to assess sensitivity to specific cohorts

All specifications use unit and product type $\times$ time fixed effects, matching the main paper specification. This creates 320 specifications per coefficient ($8 \times 4 \times 2 \times 5$).

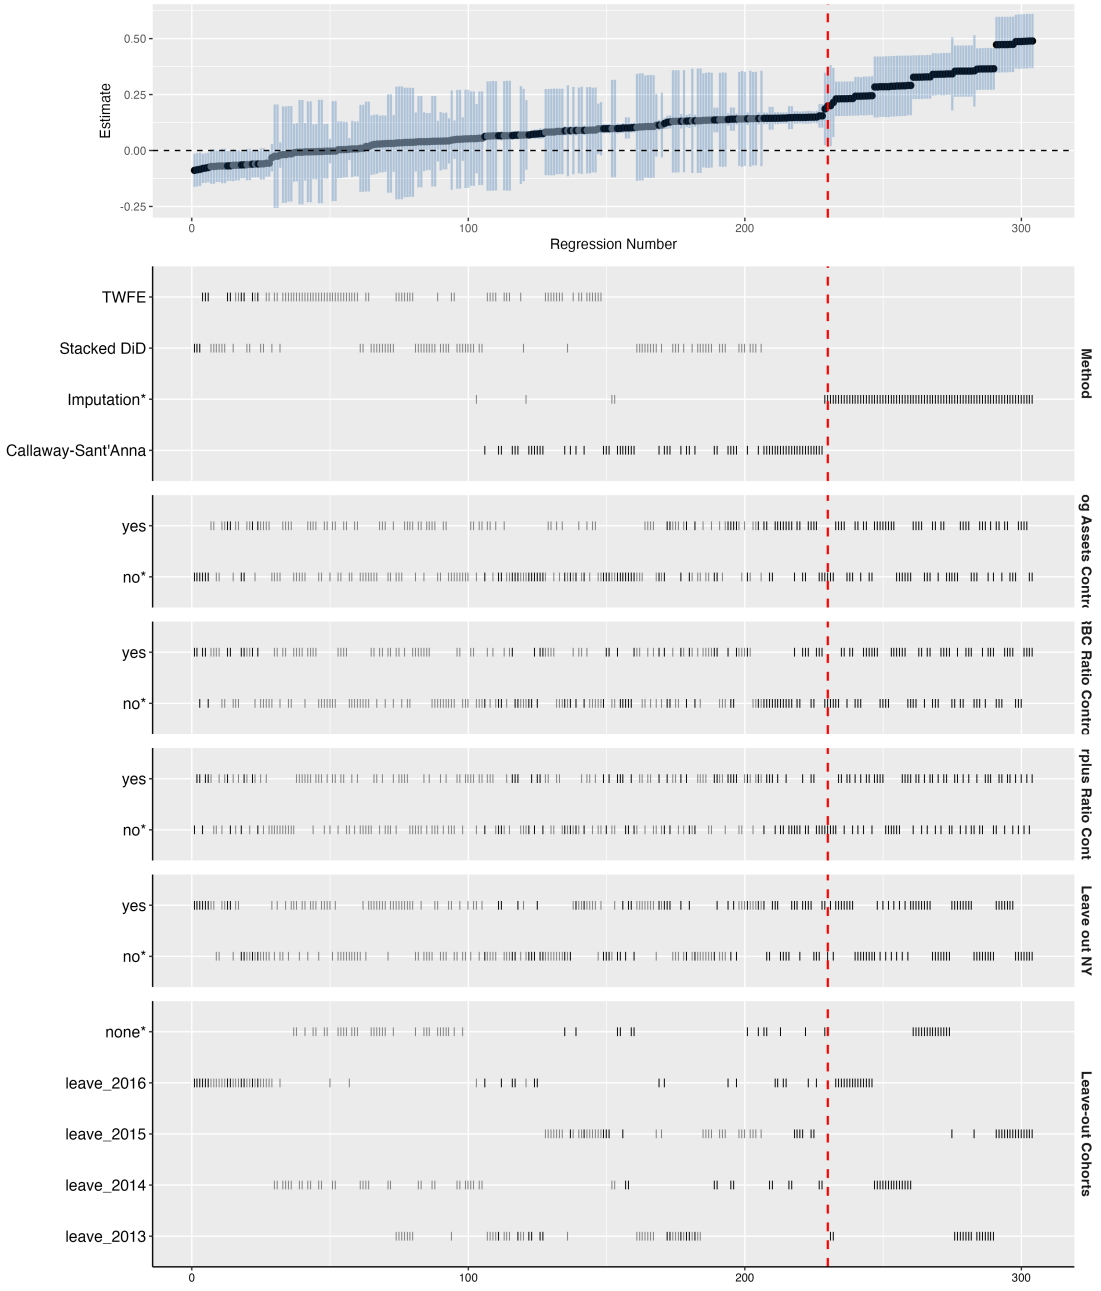
The figures display coefficient estimates sorted from smallest to largest, with 95% confidence intervals. The bottom panel indicates which specification choices correspond to each estimate. The preferred specification from the main text is highlighted.

Figure C1: Specification Curve: Fixed Annuity Contract Hazard



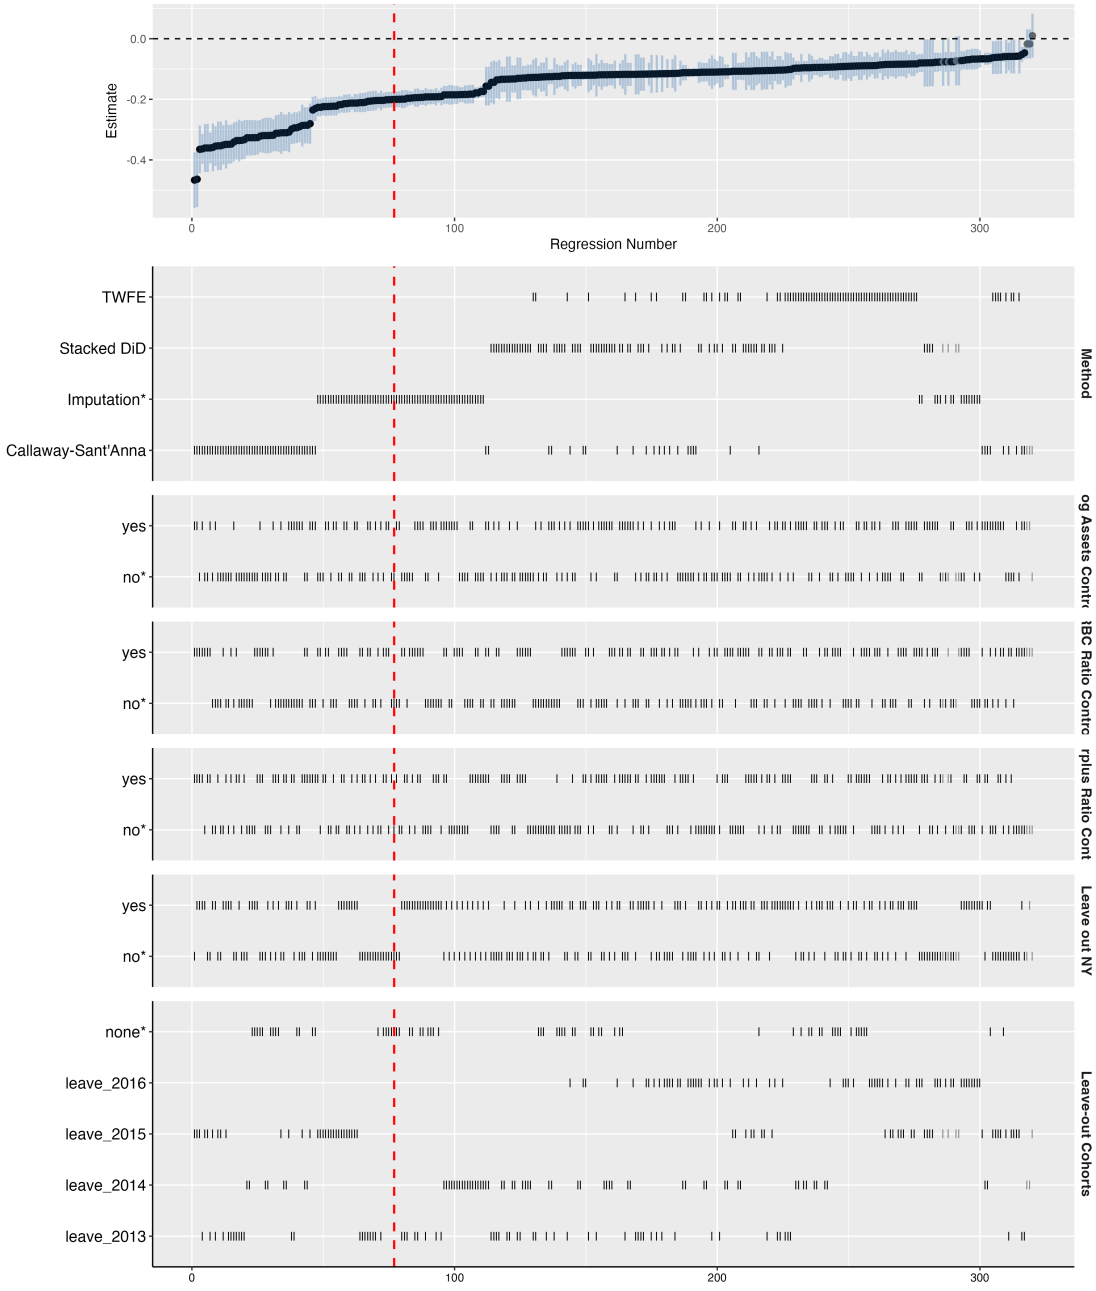
Note: Specification curve for the effect of ORSA adoption on fixed annuity contract exit hazard rates.

Figure C2: Specification Curve: Index-Linked Annuity Contract Hazard



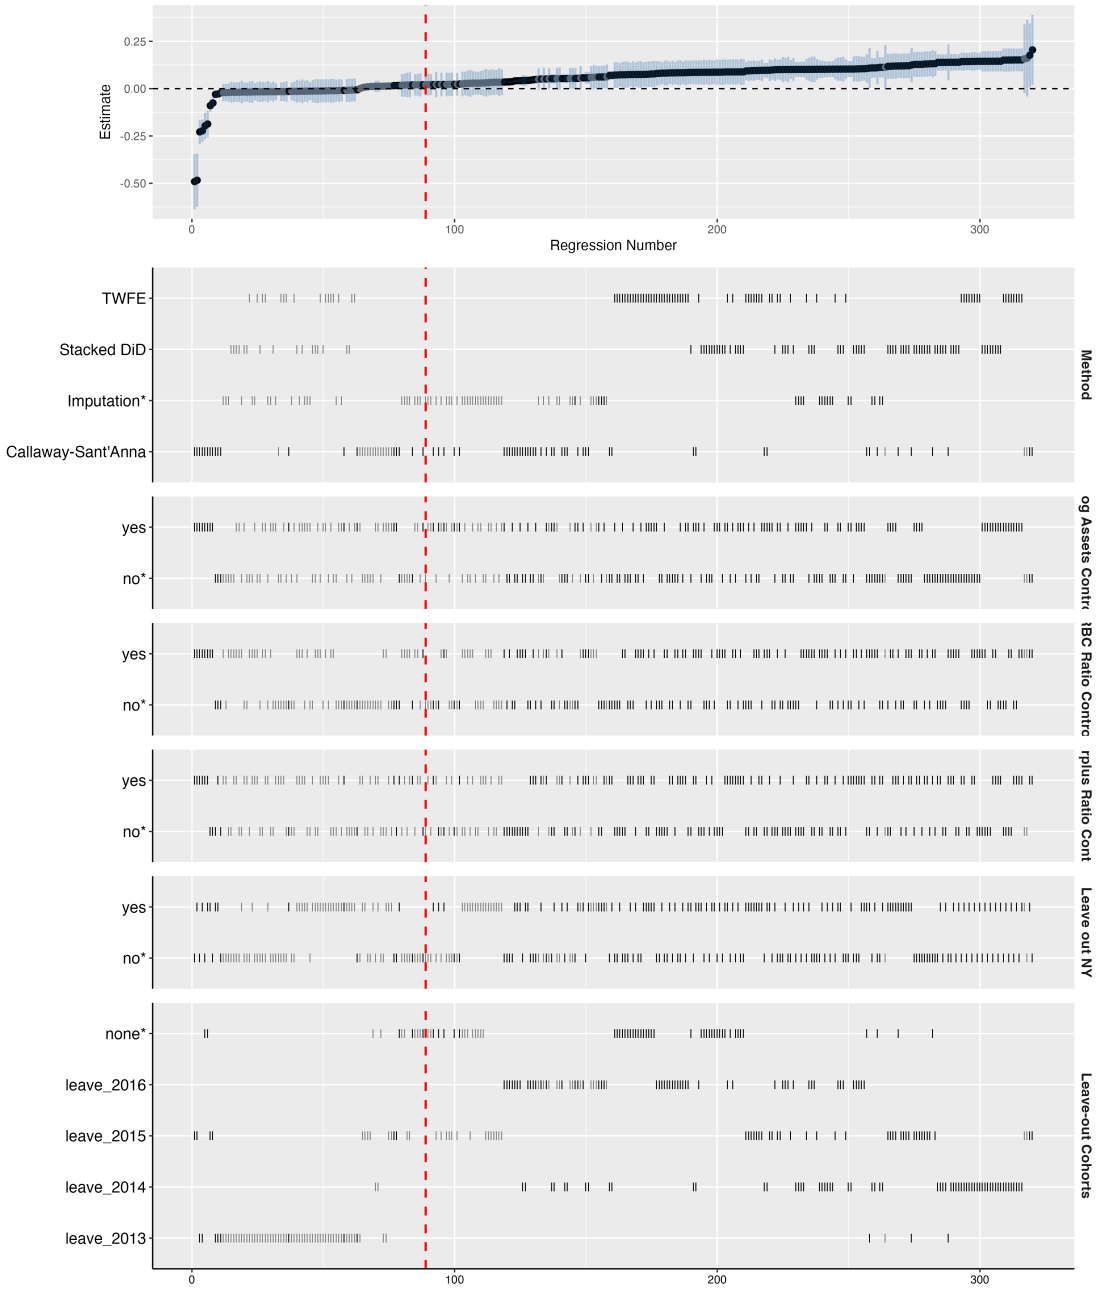
Note: Specification curve for the effect of ORSA adoption on index-linked annuity (ILA) contract exit hazard rates. ILA products show minimal effects consistent with their role as natural hedges.

Figure C3: Specification Curve: Lifetime Guarantee Benefit Hazard



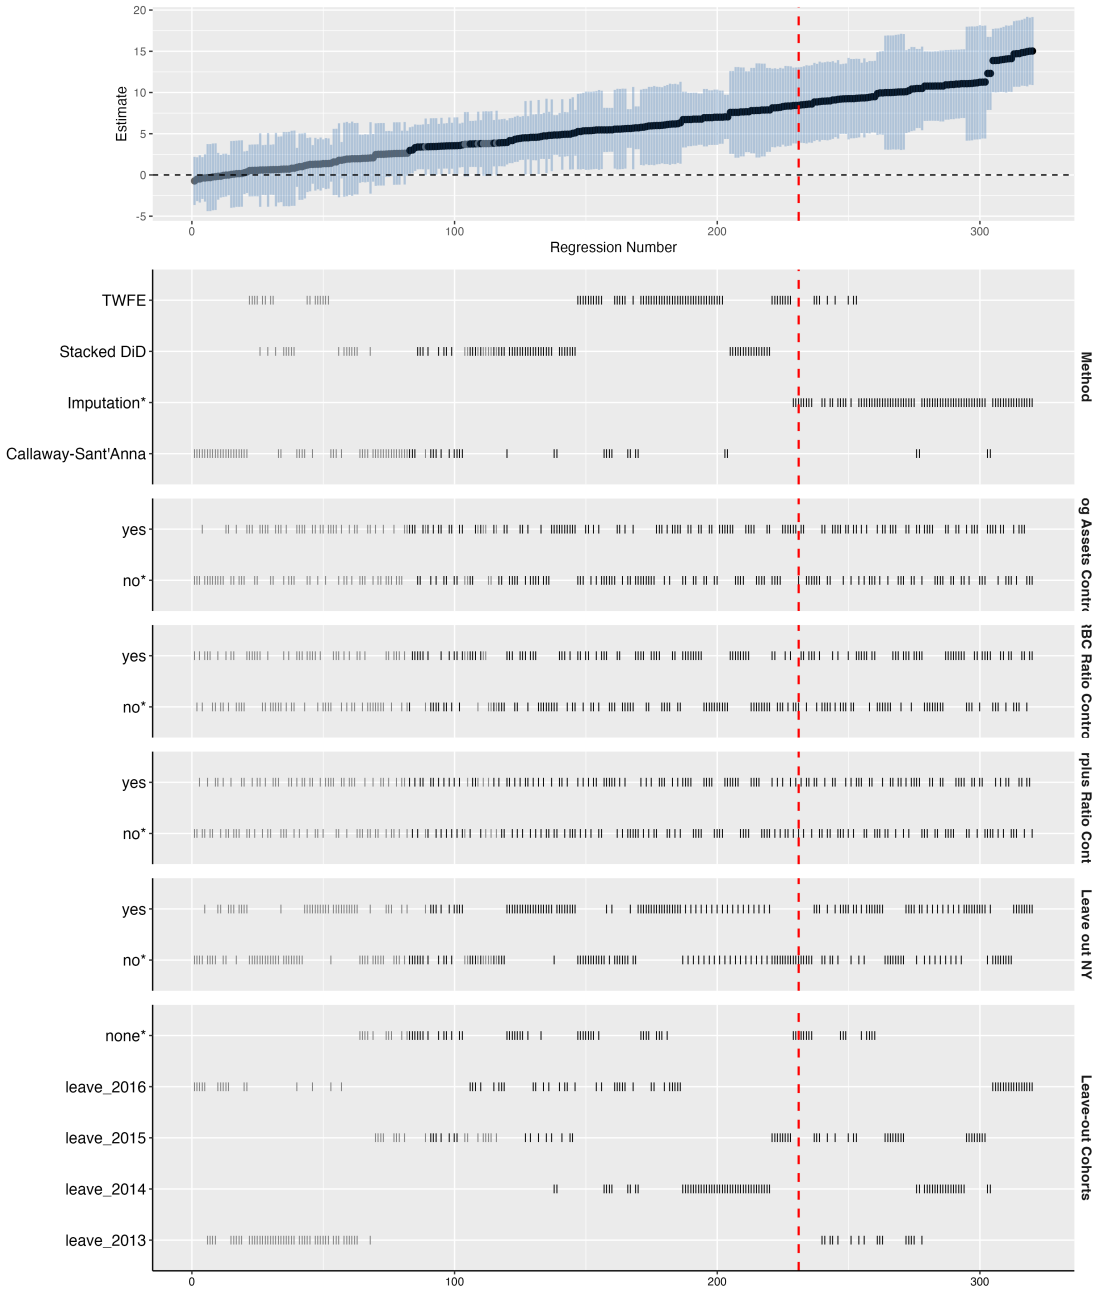
Note: Specification curve for the effect of ORSA adoption on lifetime guarantee (GMWB/GLWB) exit hazard rates. Lifetime guarantees exhibit the strongest response to ORSA.

Figure C4: Specification Curve: Non-Lifetime Guarantee Benefit Hazard



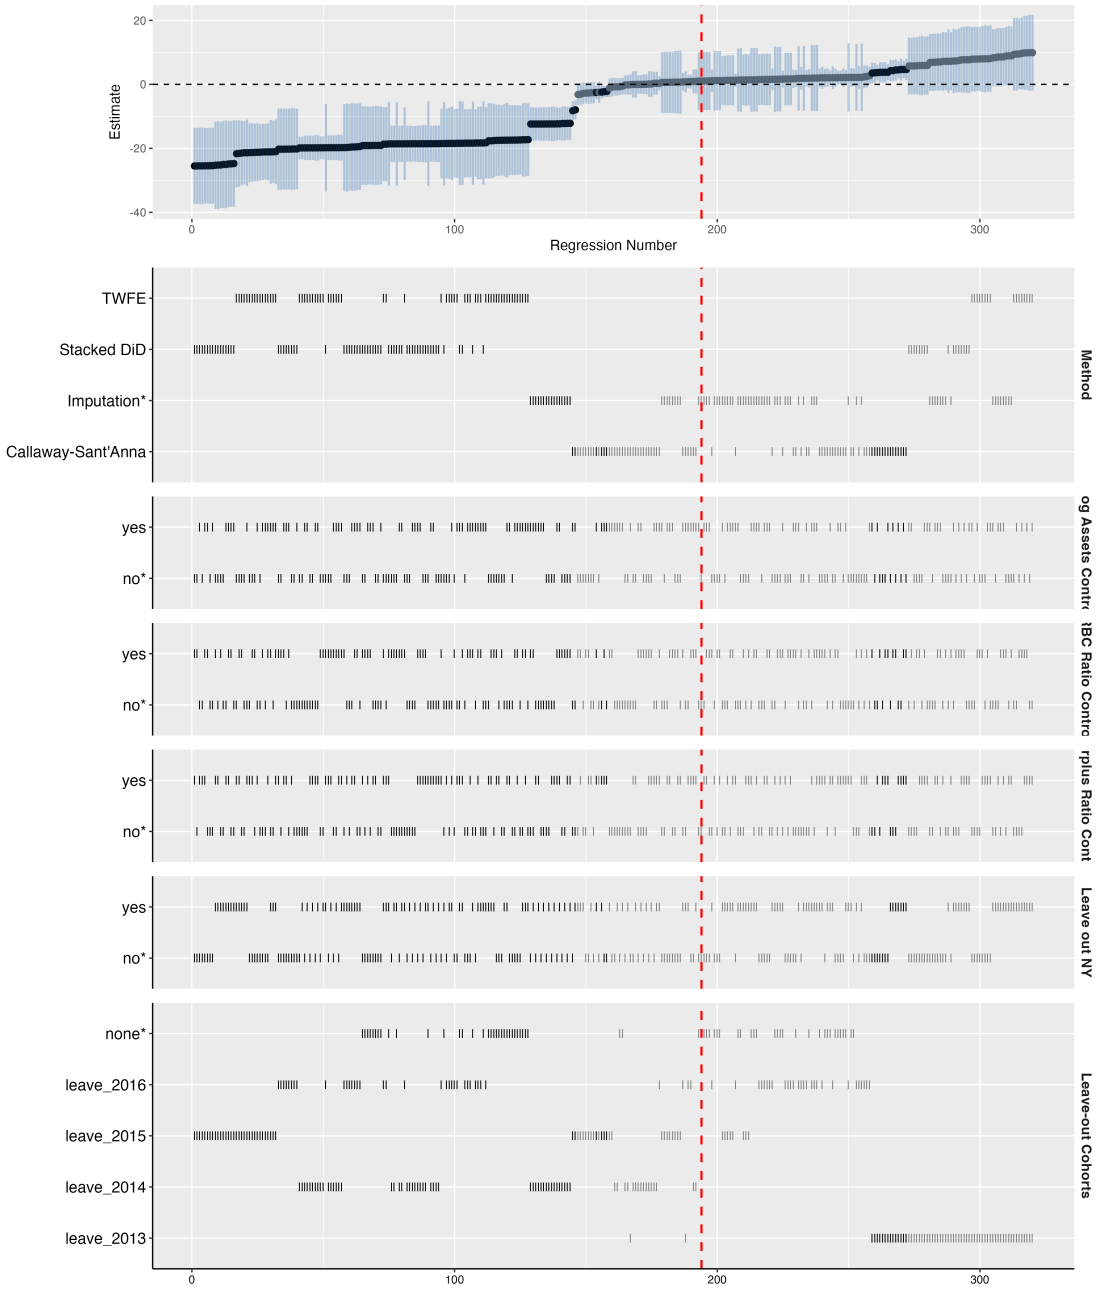
Note: Specification curve for the effect of ORSA adoption on non-lifetime guarantee (GMAB/GMIB) exit hazard rates.

Figure C5: Specification Curve: Lifetime Guarantee Fee Changes



Note: Specification curve for the effect of ORSA adoption on lifetime guarantee (GMWB/GLWB) fees.

Figure C6: Specification Curve: Non-Lifetime Guarantee Fee Changes



Note: Specification curve for the effect of ORSA adoption on non-lifetime guarantee (GMAB/GMIB) fees.

Appendix D BLP Estimation Details

This appendix documents the structural demand estimation procedure as implemented using PyBLP.

Appendix D.1 Random Coefficients (Two Dimensions)

We estimate a random coefficients logit model with exactly two random coefficients (normal heterogeneity):

Price heterogeneity is implemented by specifying $I(-\text{prices})$ in the random-coefficients formulation. With mean price coefficient β_{price} and standard deviation σ_{price}, the individual-specific price sensitivity is

$$\alpha_i = \beta_{\text{price}} - \sigma_{\text{price}}\nu_i, \quad \nu_i \sim \mathcal{N}(0,1). \tag{D1}$$

Using $I(-\text{prices})$ makes the sign convention explicit: economically, we expect $\beta_{\text{price}} < 0$ (utility decreases in price).

Guarantee heterogeneity is implemented by specifying the guarantee indicator in the random-coefficients formulation. With mean guarantee coefficient β_{guar} and standard deviation σ_{guar},

$$\gamma_i = \beta_{\text{guar}} + \sigma_{\text{guar}}\nu_i, \quad \nu_i \sim \mathcal{N}(0,1), \tag{D2}$$

and we expect $\beta_{\text{guar}} > 0$.

No additional random coefficients are estimated.

Appendix D.2 Sign-Preserving Bounds and Sigma Cap

To keep the implied distribution of individual-specific coefficients economically well-behaved, we bound the random-coefficient standard deviations (diagonal elements only). In the implementation, we impose:

$$\sigma_{\text{price}} \in [0, 0.4], \qquad \sigma_{\text{guar}} \in [0, 0.4], \tag{D3}$$

with off-diagonal elements restricted to zero.

This σ cap is motivated by a sign-preserving heuristic: keeping σ small relative to $|\beta|$ (e.g., $\sigma \lesssim |\beta|/3$) makes the probability of "wrong-signed" individual coefficients under normal heterogeneity negligible. With $|\beta| \geq 1.25$ and $\sigma \leq 0.4$, the probability of a wrong-signed coefficient is less than 0.1%.

Appendix D.3 Integration (Simulation)

We approximate the integral over heterogeneous tastes using **Modified Latin Hypercube Sampling (MLHS)** with **500 draws** per market (antithetic draws enabled in the PyBLP integration configuration).

Appendix D.4 Estimation, GMM Objective, and Optimal Instruments

Estimation is conducted in PyBLP via **two-step GMM** and uses an **L-BFGS-B** optimizer.

We follow an explicit **initial** $\rightarrow$ **optimal instruments** $\rightarrow$ **final** workflow:

1. **Initial Estimation (Current Instruments):** Solve the model using the baseline instrument set (centered instruments, plus sigma instruments when available), using a looser tolerance for stabilization.
2. **Optimal Instruments:** Compute PyBLP optimal demand instruments from the initial results using the approximate method and replace the baseline instrument matrix with these optimal instruments.
3. **Final Estimation (Optimal Instruments):** Re-solve the model with the optimal instruments using production-grade tolerances (1×10^{-10}).

Appendix E Risk-Based Capital Analysis

A natural concern is that ORSA's effects on product design operate through regulatory capital constraints rather than improved risk perception. If ORSA tightened effective capital requirements, firms might restructure product lines simply to reduce capital charges rather than because they better understand cross-product risk dependencies. We investigate this channel by examining the impact of ORSA on firms' risk-based capital (RBC) ratios and their components.

Table E1 reports imputed DiD estimates for three log-transformed outcomes: the RBC ratio itself, adjusted capital (the numerator), and RBC required capital (the denominator). The first row reports the overall ORSA effect; subsequent rows interact ORSA with indicators for RBC constraint status (bottom quartile of 2010 log RBC ratio versus top three quartiles) and firm size (above versus below median 2010 log total assets).

Table E1: Impact of ORSA on Risk-Based Capital Components

	log(RBC Ratio)	log(Adj. Capital)	log(Req. Capital)
ORSA	-0.130	0.010	0.137^{**}
	(0.091)	(0.078)	(0.067)
ORSA $\times$ Not Binding	-0.181^{**}	0.003	0.181^{**}
	(0.092)	(0.083)	(0.071)
ORSA $\times$ Binding (Q1 RBC)	0.049	-0.015	-0.065
	(0.105)	(0.077)	(0.088)
ORSA $\times$ Large	-0.132	-0.076	0.054
	(0.092)	(0.073)	(0.060)
ORSA $\times$ Small	-0.127	0.075	0.198^{**}
	(0.098)	(0.091)	(0.086)
Observations	918	918	918
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Clustering	Lead State	Lead State	Lead State

Notes: This table reports imputed DiD estimates of the effect of ORSA adoption on log risk-based capital components. Binding = bottom quartile of 2010 log(RBC ratio) (26 firms); Not Binding = top three quartiles (77 firms). Large/Small = above/below median 2010 log(total assets) (52/51 firms). Standard errors from Dirichlet bootstrap (100 replications) clustered at the lead-state level are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The overall results show that ORSA does not increase firms' RBC ratios. The point estimate on log(RBC ratio) is negative (-0.130) though not statistically significant, and log adjusted capital is essentially unchanged (0.010). Log RBC required capital increases by 13.7% (significant at the 5% level), reflecting higher capital charges driven by the shift toward RILAs documented in Table 2.

Splitting by RBC constraint status sharpens the picture. Firms in the top three quartiles of baseline RBC—those with ample capital headroom—experience a statistically significant 18.1% decline in their RBC ratio, driven entirely by higher required capital (also up 18.1%). Firms in the bottom quartile (closest to regulatory minimums) show no significant change in any component. This pattern is consistent with unconstrained firms pursuing the product reallocation strategy that raises regulatory capital charges, while constrained firms lack the headroom to do so. The size split tells a similar story: both large and small firms show negative but insignificant point estimates on the RBC ratio, with small firms showing a significant increase in required capital.

These results rule out the hypothesis that ORSA tightens binding capital constraints and thereby forces product restructuring. If anything, the RBC ratio declines for unconstrained firms, precisely because they pursue a product strategy—shifting toward RILAs—that the regulatory capital formula penalizes. The RBC framework does not recognize the natural hedging benefit that RILAs provide to TVA guarantees, so firms that act on improved risk perception are penalized in capital terms. This supports our interpretation that ORSA changes firms' economic assessment of risk rather than mechanically imposing capital pressure.