Accounting Under Pressure: Recognition Rules, Insurer Bond Sales, and Real Investment [*]

John M. Barrios[†] Andreas Neuhierl[‡] Linda Schilling[§]

August 13, 2026

Abstract

We study whether corporate bonds held by insurers suffer smaller price drops in crises because of the accounting rules applied to these bonds or because of the intrinsic characteristics of the holding institutions. In a three-period model, historical-cost (HCA) insurers can continue to report an impaired bond at amortized cost. In contrast, mark-to-market (MTM) insurers must immediately book the loss. Since only realized losses reduce regulatory capital and can trigger asset sales, insurers' selling decisions, and the prices they help sustain, differ precisely where their statutory accounting treatments diverge. Using bond-level holdings data, we find that a one-standard-deviation increase in insurer ownership reduces crisis-period drawdowns by 0.3 percentage points, about 3 percent of the average crisis decline. Even within the same insurance group, a P&C subsidiary is roughly 7 percent more likely than its Life affiliate to sell a speculative-grade bond they both hold. This gap vanishes when their accounting rules are aligned. Moreover, for the average firm in our sample, a 10-percentage point increase in insurer ownership corresponds to roughly $165 million more in annual post-crisis investment. These results indicate that during a crisis, the recognition rules applied to bondholders determine which investors are forced to sell into declining markets and which are able to hold, with significant implications for the firms whose bonds they hold.

Keywords: Fair-Value Accounting, Historical Cost Accounting, Recognition Rules, Accounting Standards, Bond Fire Sales, Credit Market Liquidity, Insurers, Financial Stability
JEL codes: M41, M48, M4, G22, G32, G12, G01, E44

[*]We thank Jeremy Bertomeu, Matthias Breuer, Emmanuel De George, Jannis Bischof, John Gallemore, Thomas Bourveau, Sehwa Kim, Christian Laux, Vincent Lin, and Zeqiong Huang for helpful comments, suggestions, and discussions. Any errors remain our own. This paper was previously circulated under the title "Accounting Under Pressure: Recognition Rules, Bond Prices, and Real Investment."

[†]School of Management, Yale University & NBER, john.m.barrios@yale.edu

[‡]Mitch Daniels School of Business, Purdue University, aneuhier@purdue.edu

[§]Olin Business School, Washington University in St. Louis, CEPR, FTG

1 Introduction

U.S. insurers hold roughly a third of the U.S. corporate bond market, more than any other investor class, and in each of the crises we study, bonds concentrated in insurer hands declined by less than otherwise-similar bonds held by mark-to-market investors (Coppola, 2025; Ellul, Jotikasthira, Lundblad, and Wang, 2015; Becker, Opp, and Saidi, 2022). There are two potential explanations for this pattern. Insurers might be inherently stable holders because of their structural characteristics such as long-term liabilities, low redemption pressure, and patient capital, or because of how they account for these assets. Under statutory accounting, most insurers carry corporate bonds at amortized cost, meaning that short-run price drops do not reduce reported capital and therefore do not force sales. This difference matters because the recognition rule, unlike an insurer's liability profile, is a policy choice. We examine whether it is the accounting regime applied to the bondholder, rather than the bondholder's identity, that drives how much a bond's price falls during a crisis and, in turn, how strongly that price change supports the issuing firm's investment activity after the crisis.

Accounting not only records value, it also shapes the choices made by those who report within its framework (Barrios, Fujiy, Lisowsky, and Minnis, 2026). The mechanism we study is a feedback loop that operates through reported capital: a shock reduces the value of a bond, the investor must recognize the loss, the recognized loss reduces reported capital relative to its regulatory requirement, the tighter constraint pushes the investor to sell, and those sales, in turn, drive down the market price faced by all remaining holders. An entity that carries the bond at amortized cost, historical-cost accounting (HCA), does not face this pressure: its reported capital does not move with the temporary price drop, allowing it to maintain the position until prices recover.

In contrast, an entity that marks the bond to market (MTM) must record the decline immediately and, if this recognized loss makes its capital constraint binding, is compelled to sell. Hence, which investors turn into sellers depends on the accounting recognition rule they follow, not only on the size of the shock. As a result, whether a shock dissipates benignly or escalates into a fire sale depends, in part, on the rule used by the marginal holder.

The insurance sector provides an ideal setting to examine this question. Corporate bonds have long durations and are highly sensitive to interest rates, so a given shock leads to sizable unrealized losses and makes the choice of whether to recognize those losses important. Other large holders, mutual funds and ETFs, mark every position to market under the Investment Company Act, providing a cross-investor comparison. Crucially for our identification, accounting treatment differs within the insurance sector itself: life insurers, which maintain an Asset Valuation Reserve (AVR), continue to carry the bond at amortized cost after it is downgraded into the speculative-grade (NAIC 3–5) tier, while property-and-casualty (P&C) insurers, which lack an AVR, must write the same bond down to the lower of amortized cost or fair value. The two accounting standards coincide for investment-grade bonds, but differ mechanically when the bond crosses the NAIC 2/3 threshold.

This within-sector variation allows us to disentangle the role of accounting rules from that of institutional characteristics. Relative to other bond investors, insurers differ in liability duration, vulnerability to investor redemptions, investment horizon, and regulatory constraints; any one of these factors alone could explain why they could be more stable holders, which prevents cross-institution comparisons from isolating the impact of accounting. Our identification strategy instead utilizes the variation in accounting treatment inside the insurance sector, focusing on differences between a Life and a P&C subsidiary belonging to the same holding company, holding the same bond in the same quarter, at the moment when the accounting rule changes but the underlying institution does not.

Theoretical studies have recognized the importance of recognition rules, yet their practical significance remains unclear. Allen and Carletti (2008) and Plantin, Sapra, and Shin (2008) show that MTM accounting can propagate shocks through commonly held assets while HCA can dampen the transmission, and Sapra (2008) formalizes the procyclicality this implies, in which write-downs and forced sales reinforce one another. Yet Laux and Leuz (2009) and Laux and Leuz (2010) are skeptical that fair value was a major cause of the 2008 crisis as it applied to only a subset of bank assets, and the reporting framework allowed ample discretion to avert write-downs. HCA, moreover, carries its own costs in tranquil periods: selective realization of gains, capital arbitrage, and postponed recognition of lasting impairments (Ellul, Jotikasthira, Lundblad, and Wang, 2015; Badertscher, Burks, and Easton, 2012; Becker, Opp, and Saidi, 2022). Whether recognition rules stabilize or destabilize during a crisis is therefore an empirical question, and the answer depends on how the rules interact with the regulatory constraints of the institutions that actually hold the assets.

Our model's contribution is to make the *composition of a bond's holders* the operative variable. Where Allen and Carletti (2008) and Plantin, Sapra, and Shin (2008) establish that mark-to-market accounting *can* propagate shocks and historical cost *can* dampen them, we characterize *when and by how much*: the price a bond sustains in a crisis is strictly monotone in two sufficient statistics—the insurer share and the historical-cost fraction within it—so an identical shock dissipates or cascades according to the recognition rule of the marginal holder.

To motivate our analysis, we build a three-period model in which HCA and MTM insurers hold a bond alongside diffuse noise traders, and an interim interest-rate shock forces each insurer to choose whether to sell. The distinctive object of the model is the *composition* of the bond's holders: η, the total insurer share, and μ, the fraction of that share that reports under HCA. MTM insurers are compelled to liquidate once the price falls to or below a precautionary buffer set just high enough to guarantee post-liquidation solvency; because this buffer grows with the MTM share, replacing one MTM insurer with one HCA insurer, holding total insurer mass fixed, makes forced liquidation both less likely and less severe, and in turn strictly raises the expected bond price after an identical interest rate shock. HCA stabilization is thus strictly monotone in the ownership composition.

The model generates a sharp set of cross-sectional predictions. Two arise directly from the ownership structure. First, bonds with a larger overall insurer share η should experience smaller price declines in crises (a level effect). Second, for a given η, bonds predominantly held by HCA insurers, Life subsidiaries that use an Asset Valuation Reserve should drop less in value than otherwise similar bonds held by MTM insurers (P&C subsidiaries in the same NAIC tier). This difference is entirely driven by forced sales and is most pronounced where such sales are most likely, among weakly capitalized insurer groups and in bonds with concentrated insurer ownership, where a dominant holder internalizes its own price impact.

The model also predicts where this stabilizing effect should be strongest. First, the stabilizing effect of HCA should be greater for bonds with longer maturities. Shorter-maturity bonds generate cash flows sooner, which naturally limits the need for fire sales. Second, the stabilization gap should expand as crises become more severe. When noise-trader sales intensify in deeper downturns, adverse shocks are more likely to hit regulatory sales thresholds under MTM, increasing the likelihood of forced liquidation. In contrast, the model predicts no such effect outside crisis periods and after positive shocks: in calm markets, insurer ownership should not help predict returns, and HCA should cushion downside risks without impeding subsequent recoveries.

We test these predictions with detailed data on bond holdings, transactions, and issuer characteristics: quarterly Schedule D insurer holdings from CapitalIQ, mutual-fund holdings from Morningstar, bond transactions from Enhanced TRACE, bond characteristics from Mergent FISD, and firm financials from Compustat. For the core within-holding-company test we build a new sample, matching NAIC Group Codes from S&P Global to Schedule D Securities Valuation Office (SVO) designations extracted from 193 annual regulatory filings. This yields 2.45 million subsidiary-bond-quarter observations, of which roughly 92,000 fall in the NAIC 3–5 cells where the statutory rules diverge, spanning 60 holding companies. We add a risk-based capital (RBC) panel that covers 97.3 percent of insurer-bond-quarter holdings for the mechanism tests. Identification proceeds in three layers: fully saturated bond-characteristic fixed effects interacted with crisis indicators, a shift-share instrument built from primary-market allocations at issuance, and group $\times$ bond $\times$ quarter fixed effects.

We find that bonds with more insurer ownership fall substantially less in crises. In a fully saturated specification that interacts crisis events with bond size deciles, duration categories, credit rating, issuer fixed effects, and coupon and seniority indicators, so the comparison is between otherwise-similar bonds in the same crisis, a 10 percentage-point increase in insurer share predicts a smaller drawdown of about 1.2 percent of the mean crisis drawdown. A shift-share IV, built from quasi-random primary-market allocations at issuance, yields a larger coefficient, roughly 10–11 percent of the mean drawdown per 10 percentage-point increase in instrumented insurer ownership. The IV exceeds the OLS estimate, a gap consistent with measurement error in ownership shares attenuating the cross-sectional association and with the instrument identifying the effect among the recently issued bonds from which it is built. Entering mutual-fund ownership alongside insurer share leaves the insurer coefficient essentially unchanged while the mutual-fund coefficient is indistinguishable from zero, so the stabilization is specific to insurer ownership rather than to institutional ownership in general.

Although cross-investor tests confirm the stabilization effect, they cannot fully separate accounting from institutional differences. As a result, we next look inside the insurance sector for identification. At the NAIC 2/3 downgrade threshold, the P&C accounting mechanically shifts from amortized cost to fair value, while the Life accounting treatment remains the same. At this stage, the behavior of the two sectors diverges according to the rule: P&C insurers dispose of the downgraded bond, while Life insurers keep it on their books. Bonds held by P&C insurers at the time of downgrade subsequently earn higher returns in the following quarters, consistent with the reversal of a fire-sale discount arising from the forced selling pressure on the price of the bond (Ellul et al., 2011). At the NAIC 1/2 threshold, where the accounting treatment is identical for both subsidiary types, both differentials are zero.

To further refine our identification strategy, we employ a fixed-effects design that compares entities with the same security within the same holding company.[1] Parent groups that own both a Life and a P&C subsidiary share investment philosophy, risk appetite, capital position, and liability composition across their subsidiaries. What differs for NAIC 3–5 bonds is the statutory accounting recognition rule. Comparing how the two subsidiaries manage the same bond in the same quarter isolates the rule from everything common to the parent. Additionally, because the rules agree for the NAIC 1–2 bonds and diverge for the NAIC 3–5 bonds, the design embeds its own placebo, since any channel other than accounting should produce the same gap in both rating bands. We find results consistent with the accounting interpretation: for speculative-grade bonds, a P&C subsidiary sells the bond about 1.4 percentage points (roughly 7 percent) more frequently than its Life sibling, while for investment-grade bonds, where the rules are identical, the differential is essentially zero. What identifies the impact of the accounting rule is the *difference* between the NAIC 3–5 versus 1–2 comparison, which is statistically significant. This gap arises only where the accounting rule diverges, remains when any single holding company is excluded, and is robust to using different sources for the NAIC classification.

This within-group comparison most clearly separates the accounting channel from the main alternative explanation. Coppola (2025) documents a similar stabilizing effect for insurers, but ascribes it to liability structure and redemption risk, focusing exclusively on investment-grade bonds, where statutory treatment is uniform across all holders. In contrast, our identification relies precisely on the margin that study leaves out: the speculative-grade (NAIC 3–5) segment, where statutory treatment differs across holders. The liability-structure explanation predicts no discontinuity at the NAIC 2/3 threshold, since neither liability duration nor redemption exposure changes when a bond moves across a rating boundary, whereas the accounting rule does. Life and P&C affiliates have distinct liability structures at all rating levels, but the selling gap appears only in the segment where accounting standards diverge, precisely where the liability-structure explanation makes no prediction.

Our cross-sectional tests also align with our proposed accounting mechanism. The stabilization

[1] Ellul, Jotikasthira, Lundblad, and Wang (2015), in one robustness subsample, restrict to Life and P&C insurers belonging to groups that contain both types (their Internet Appendix Table IAXIII) but do not implement a same-security, within-holding-company fixed-effects design, our central identification contribution.

effect is roughly four times stronger for long-duration bonds (remaining maturity > 5 years) than for short-duration bonds, consistent with the HCA shield tightening more when unrealized losses are larger. It is concentrated in the Great Recession and COVID, the two most severe crises, and is weak or absent in milder episodes. The effect operates through insurer regulatory capital along two dimensions. Across bonds, stabilization is strongest precisely where forced sales are a realistic risk: the stabilization slope is about four times steeper when a bond's insurer owners are within 300–500 percent of the regulatory action threshold than when they have ample slack, and the sparse below-300 bin points in the same direction. Within holding companies, the accounting gap is itself a weak-capital phenomenon: the extra selling of speculative-grade bonds is driven by weakly capitalized P&C subsidiaries; their weakly capitalized Life affiliates hold the same bonds, actively using the recognition shield, while well-capitalized affiliates of either type exhibit no gap at all, because neither is close enough to the constraint for forced sales to matter. Bond-level ownership concentration amplifies the effect, consistent with large holders internalizing their own price impact. A triple interaction of insurer share, long duration, and severe-crisis periods is substantial. Taken together, none of the leading alternatives matches all of these features at once: a patient-capital explanation does not imply the duration loading; a liability-driven-investment story does not generate the crisis-severity pattern; and a low-volatility-selection hypothesis predicts the opposite duration interaction and cannot account for the capital-ratio channel.

Three null results support the accounting-based mechanism: one functions as a scope check, and two operate as genuine falsification tests. First, for the scope condition, health insurers, who, like other insurers, use amortized-cost accounting for the investment-grade bonds that dominate insurer portfolios but themselves hold only negligible bond positions do not exert any stabilizing influence through their ownership stakes. This is finding is consistent with a channel that requires materially large amortized-cost holdings rather than the accounting convention in isolation. Second, outside of crisis periods, once the baseline fixed-effects specification is used, insurer ownership ceases to forecast bond returns, the apparent low-volatility tilt in simpler specifications is entirely absorbed by rating and issuer controls. Third, within a given holding company, Life and P&C subsidiaries exhibit essentially identical selling behavior for NAIC 1-2, where the statutory treatment is the same. The accounting mechanism fits all three null findings, whereas each alternative mechanism we consider is rejected by at least one of them.

Finally, we ask whether the accounting channel operates through the credit market to affect firms' investment behavior. Firms whose bonds are more heavily held by HCA-reporting insurers experience smaller crisis-period price declines on their debt, implying more limited increases in effective borrowing costs and, in turn, stronger post-crisis investment. At the firm level, a 10-percentage-point rise in insurer ownership is associated with roughly a 0.14-percentage-point increase in post-crisis capital expenditures and about a 0.63-percentage-point increase in total investment (capital expenditures plus acquisitions). Instrumental-variables estimates of firm-level insurer ownership continue to indicate positive investment effects under weak-instrument-robust inference. Overall, the same portfolio-composition logic that identifies the price effect also identifies the investment effect. When we decompose a firm's insurer ownership into the portions held by HCA-reporting (Life) and MTM-reporting (P&C) subsidiaries and interact each with the NAIC tier of the firm's bonds, the post-crisis investment response is concentrated in the HCA share within the NAIC 3–5 range, where accounting treatments diverge, and there is no difference across holder types in the NAIC 1–2 range, where the rules coincide. Thus, the real effect, like the price effect, arises exactly at the accounting boundary, where total ownership and liability structure remain smooth. Consequently, the impact of recognition rules extends beyond insurers' balance sheets and propagates through the credit market to shape issuing firms' post-crisis investment and acquisition activity.

Our paper makes three contributions to the accounting and finance literature. First, we provide what is, to our knowledge, the first direct identification of the accounting recognition rule, separately from insurers' contractual characteristics, as the mechanism behind crisis-period bond-price stability, the question at the center of the fair-value debate (Laux and Leuz, 2009; Bischof, Laux, and Leuz, 2021). Prior work documents the insurer–mutual-fund stabilization gap without isolating the accounting channel from the bundle of institutional differences (Coppola, 2025; Becker, Opp, and Saidi, 2022), or compares insurer types without holding group-level attributes fixed (Ellul, Jotikasthira, Lundblad, and Wang, 2015; Khan, Ryan, and Varma, 2019). Our within-holding-company design isolates the accounting rule.

Second, our theoretical contribution is to make the composition of the bondholders under HCA versus MTM the central determinant of the severity of the fire-sale, while Allen and Carletti (2008) fixes the relative size of the two sectors and Plantin, Sapra, and Shin (2008) studies the pure accounting regimes. This yields the comparative static we take to the data that stabilization is strictly increasing in the HCA share, with each marginal substitution of HCA for MTM reducing fire-sale severity, a prediction that prior models do not deliver. Two further additions, an exogenous fundamental-value shock and insurers that internalize their price impact, produce the capital-channel, duration, severity, and concentration predictions that our empirical findings support. Relative to the broader fire-sale literature (Ellul, Jotikasthira, and Lundblad, 2011; Becker and Ivashina, 2015; Diamond, Falasconi, and Xu, 2024), the framework supplies an endogenous, composition-dependent buffer that other mechanisms generate through collateral constraints, limited arbitrage, or liability maturity structure.

Third, recognition rules transmit through the credit market into firm investment: HCA-shielded price stability lowers the issuer's effective borrowing cost and supports post-crisis capital expenditure and acquisitions, extending the investor-composition literature's real-effects link (Coppola, 2025; Bischof, Brüggemann, and Daske, 2023) by showing that the composition that matters is the accounting treatment of the holder rather than the institution type. This clarifies a welfare tension in the earlier literature. In calm periods, HCA generates the gains-trading incentives that Ellul, Jotikasthira, Lundblad, and Wang (2015) and Becker, Opp, and Saidi (2022) document; in crises, the same discretion stabilizes prices and supports investment. The welfare calculus is state-contingent, and the case for HCA in corporate bond markets is strongest exactly when the risk of a fire sale is most acute.

The paper proceeds as follows. Section 2 describes the institutional setting, with particular attention to the SSAP rules that facilitate our identification strategy. Section 3 develops our theoretical framework and presents the empirical predictions that stem from the model. Section 4 describes our data and sample. Section 5 presents our empirical analysis, which progressively narrows the set of alternative explanations. We conclude in Section 6.

2 Institutional Context

2.1 Held-to-Maturity vs. Mark-to-Market in Financial Institutions

Two recognition rules govern how financial institutions value fixed-income assets, and the difference between them is the subject of this paper. Under historical-cost accounting (HCA), an institution records a bond at amortized cost: it books the bond at purchase price, amortizes any premium or discount over the bond's life, and does not mark the position to market. The rule fits institutions with long-horizon liabilities, such as insurers, that hold bonds to maturity to meet future claims. Its appeal is stability: because unrealized gains and losses never hit reported capital, the balance sheet is insulated from short-run price swings. The same insulation is its weakness: HCA can hide unrealized losses from investors and regulators (Granja et al., 2024) and delay the recognition of distress, since an institution can keep an impaired bond on the books at cost (Acharya et al., 2014). In calm periods it also invites gains trading: an insurer sells bonds with unrealized gains and keeps those with unrealized losses, managing reported earnings around regulatory thresholds and rating changes (Ellul et al., 2015). Left unchecked, this delays both regulatory and managerial responses in a downturn (Chircop and Novotny-Farkas, 2016).

Under mark-to-market accounting (MTM), or fair value, an institution instead reports each bond at its current market price at the end of every reporting period. Its appeal is transparency: investors, regulators, and counterparties see the current value of the portfolio, which supports more informed decisions (Laux and Leuz, 2009; Barth, 1994), and the rule limits managerial discretion over valuations, narrowing the room for earnings manipulation (Barth, 1994). The cost is volatility: because valuations move with the market, reported earnings and capital swing with them, and the swings are largest precisely when markets are under stress. Under that stress, MTM can feed a fire sale: a price decline cuts reported capital, a binding capital constraint forces sales, and the sales push prices down further (Allen and Carletti, 2008; Plantin et al., 2008; Sapra, 2008; Ellul et al., 2015; Merrill et al., 2021). Which rule serves markets better is an empirical question.

2.2 Insurers' Accounting Treatment of Fixed-Income Securities

U.S. insurers report under Statutory Accounting Principles (SAP), developed by the National Association of Insurance Commissioners (NAIC), which emphasize solvency and conservative valuation over income measurement.[2] The accounting treatment of a bond is not chosen by the insurer; it is determined mechanically by two observables: the bond's NAIC designation (1 through 6) and whether the holder maintains an Asset Valuation Reserve (AVR).[3] For the corporate bonds in our sample, both Life and property-and-casualty (P&C) insurers report under SSAP No. 26R; loan-backed and structured securities, which fall under SSAP No. 43R, lie outside our universe. The life-versus-P&C divergence we exploit arises within SSAP No. 26R's measurement provision, which keys carrying value to AVR status.[4] Life and fraternal insurers maintain an AVR; P&C and Health insurers do not. The two regimes align at the extremes of the rating distribution and diverge in the middle. For NAIC 1–2 bonds (investment grade), both AVR and non-AVR insurers carry the bond at amortized cost, so a Life and a P&C subsidiary treat the same bond identically. For NAIC 3–5 bonds (speculative grade, non-default), the AVR insurer (Life) keeps the bond at amortized cost while the non-AVR insurer (P&C) must report it at the lower of amortized cost or fair value (LOCOM); the same bond is carried differently by Life and P&C subsidiaries inside the same parent group. For NAIC 6 bonds (in or near default), both carry the bond at LOCOM again. A bond's path through the NAIC tiers therefore produces two empirically distinct shocks. A downgrade across the NAIC 2/3 boundary changes non-AVR (P&C) accounting from amortized cost to LOCOM while leaving AVR (life) accounting unchanged. A downgrade within the NAIC 1–2 band, or within the NAIC 3–5 band, changes nothing in either treatment. The variation is rule-determined rather than discretionary, and it occurs at observable, regulator-defined thresholds.

This rule structure is the foundation of the within-insurer tests in Section 5.4. At the bond level, NAIC 2/3 downgrades discretely shift the HCA fraction μ of the bond's insurer holders downward when those holders are concentrated in P&C. At the holding-company level, a Life and a P&C subsidiary of the same parent group can hold the same CUSIP in the same quarter and apply different accounting rules to it, variation that holds parent strategy, internal capital markets, and macro conditions fixed and isolates the rule itself. The NAIC 1–2 versus NAIC 3–5 contrast supplies a built-in placebo: where the rule does not diverge across AVR status, no within-group differential should appear.

[2] Under the McCarran-Ferguson Act of 1945, insurance regulation is delegated to the states. Insurers must file statutory financial reports with both state regulators and the NAIC using SAP. Although SAP has been codified by the NAIC since 2001 in the Accounting Practices and Procedures Manual, individual states retain limited discretion, resulting in modest jurisdictional variation.

[3] For the filing-exempt rated corporates that make up our sample, the NAIC designation is obtained by translating the issuer's NRSRO ("Credit Rating Provider") rating into its NAIC-designation equivalent using the NAIC's published crosswalk, rather than through a hands-on credit assessment by the NAIC Securities Valuation Office (SVO); this Filing-Exempt rule was adopted in 2004. The six designations were subdivided into twenty more granular NAIC Designation Categories effective year-end 2019, after most of our sample; we use the 1–6 grades throughout.

[4] SSAP No. 26R, "Balance Sheet Amount," ¶ 20.a in the editions covering our 2007–2020 sample (SSAP No. 26 before its 2017 revision as 26R, with the measurement provision we exploit unchanged; the section was renumbered again in the principles-based rewrite effective January 1, 2025). The same AVR-based split also governs structured securities under SSAP No. 43R, which confirms that the divergence is driven by AVR status (entity type) rather than by which standard applies. The principles-based redefinition of SSAP No. 26R and 43R effective January 1, 2025 post-dates our sample and does not affect the rule we exploit. SSAP 43R securities are empirically negligible in our data: none of the 13,465 bonds in the analysis sample is flagged asset-backed in Mergent FISD (0.0 percent by count, by par value, and by insurer-held market value) against 0.2 percent asset-backed in the full FISD corporate master.

Realization-timing discretion under amortized cost is a separate matter from the rule classification itself. Once a bond is carried at amortized cost, insurers retain some flexibility over when to recognize realized gains or losses, which generates the gains-trading patterns documented by Ellul et al. (2015) in calm periods. Fair-value reporting in turn imposes its own constraints: Laux and Leuz (2010) note that mark-to-market accounting becomes difficult to apply when market prices are thin or unobservable, forcing reliance on mark-to-model valuations that import additional measurement error during stress. Khan et al. (2019) document that insurers subject to fair-value recognition show more timely impairment recognition and more frequent portfolio rebalancing under market stress than insurers carrying the same exposures at amortized cost.

The first-order point for the analysis that follows is that the recognition rule is a property of the bond–holder pair rather than of the insurer's strategic decisions. Mutual funds and ETFs, by contrast, mark all fixed-income holdings to market under the Investment Company Act of 1940; insurers and mutual funds therefore differ in recognition treatment for every bond, while Life and P&C subsidiaries inside the same insurance group differ in recognition treatment only for NAIC 3–5 bonds. These two sources of variation (between institutions and within an institution) drive the empirical strategy of Section 5.1 onward.

2.3 Crises

We study how bond drawdowns vary with insurer ownership during the four largest bond-market downturns in our sample. Following Gilchrist and Zakrajšek (2012) (GZ), we define crises as the periods from local troughs to subsequent peaks in the GZ index, a measure of credit spreads on U.S. corporate bonds; the shaded areas in Figure 1 mark them.

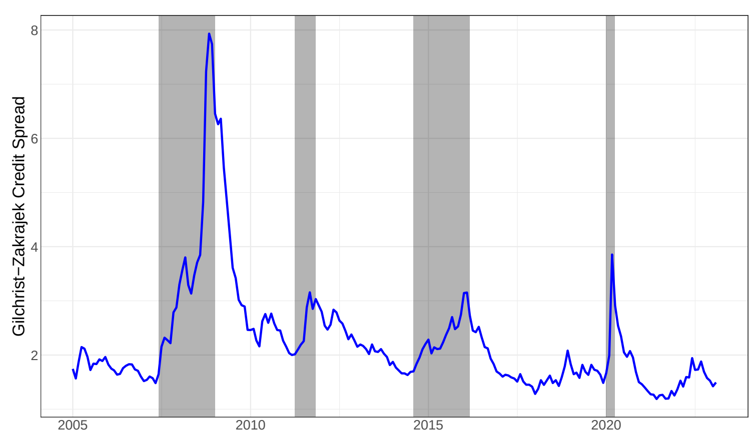

Figure 1: Crisis Events

Notes: This figure plots the Gilchrist and Zakrajšek (2012) index (the solid blue line), an indicator of credit spreads across U.S. corporate bonds. The shaded gray areas delineate the event windows used in our analysis for estimating bond drawdowns. The start and end dates of these windows correspond to the preceding trough and the subsequent peak in the Gilchrist and Zakrajšek index.

Four downturns fall within our sample. The Great Recession of 2008–09 is the largest: the GZ index rose more than fourfold, from a trough of 1.7 percent in June 2007 to a peak of 7.9 percent in December 2008. The second window, April to October 2011, coincides with Standard & Poor's downgrade of U.S. sovereign debt. The third, August 2014 to February 2016, captures the collapse in commodity prices, especially oil. The fourth, January to March 2020, covers the onset of the COVID-19 pandemic, which produced large price dislocations in corporate bonds (Haddad et al., 2021; Ma et al., 2022). The four episodes span distinct shocks—financial, sovereign, commodity, and pandemic—letting us test whether the accounting channel operates across crisis types rather than in a single episode.

3 Theoretical Framework

How does the accounting rule applied to a bond shape its behavior in a crisis? We model insurers that hold corporate bonds under either historical cost (HCA) or mark-to-market (MTM) accounting, take an interest-rate shock that lowers fundamental value, and ask how much the bond's price falls. The answer turns on *who owns the bond*: the share held by HCA insurers, relative to MTM holders, governs both the probability and the severity of a forced-sale event.

The framework builds on Allen and Carletti (2008) and Plantin, Sapra, and Shin (2008). Allen-Carletti's contagion channel and Plantin-Sapra-Shin's coordination channel both operate here: MTM can compel sales that spill into prices, and strategic complementarities among discretionary sellers amplify the spillover. We add two features absent from those settings. First, the shock directly alters fundamental value, so sales can be efficient, profit-maximizing in light of the new continuation value, rather than merely forced. Predictions therefore depend on shock magnitude, bond duration, and capital buffers, not just on the presence or absence of a solvency event. Second, insurers are strategic: each internalizes its price impact, producing an endogenous no-sale region under HCA that competitive or atomistic models cannot generate. The model delivers a Nash characterization of sales and prices as a function of ownership composition (Proposition 3.1) and a comparative static on expected prices (Proposition 3.2). Proofs of the main results are collected in Appendix OA-8.

3.1 Model Setup

Bond Market and Investor Groups. Time is discrete with three dates, $t = 0, 1, 2$. The economy contains N large risk-neutral investors of equal size, split into N_I insurers and N_o diffuse investors (noise traders), $N_I + N_o = N$. All discount at rate $\beta_0 = 1/(1 + r_0)$, $r_0 \in (0, 1)$. At $t = 0$ the issuer floats zero-coupon bonds of total size S_0 at initial price b_0, promising principal P at $t = 2$ with publicly observed repayment probability $p \in (0, 1)$.

Insurer and Insurance Contracts. Each insurer issues long-term insurance contracts to households of measure $H_I \geq 0$ at premium d_I and invests the proceeds, together with retained earnings

E_I, in bonds. All investors hold an equal balance sheet,

$$A_I := \frac{S_0}{N} = \frac{d_I H_I + E_I}{b_0}. \tag{1}$$

where A_I denotes each investor's quantity of bonds. Contracts span $t = 0$ to $t = 2$, matching bond maturity. Under U.S. statutory accounting (SAP), fixed-income securities are typically carried at amortized cost unless impairment is other-than-temporary.[5] Let D_I denote the $t = 1$ value of each insurer's liabilities.

Capital Regulation. Each insurer is held to a maximum debt-to-asset ratio $D_I/(b_0 A_I) \leq \alpha_E$, $\alpha_E \in (0,1)$, and operates exactly at the limit:

$$\frac{D_I}{b_0 A_I} = \alpha_E. \tag{2}$$

Under HCA, interim fair-value changes leave the reported asset base unchanged, so the capital ratio survives a price decline. Under MTM, price declines reduce measured assets and can force deleveraging. The recognition rule is what translates a valuation shock into a binding capital constraint.

Diffuse Investors. Diffuse investors collectively hold $A_o = S_0 - N_I A_I = N_o A_I$ and are nonstrategic. A fraction $\lambda \in [0,1]$ are impatient and liquidate at $t = 1$ regardless of price; the rest hold to maturity. Their exogenous sales set the floor of secondary-market supply,

$$\underline{S}_1 := \lambda N_o A_I. \tag{3}$$

Insurer Composition and Ownership Shares. Among the N_I insurers, n_H are subject to HCA and $n_M = N_I - n_H$ to MTM, holding $S_{\text{HCA}} := n_H A_I$ and $S_{\text{MTM}} := n_M A_I$ respectively. The composition (n_H, n_M) is taken as given. Let

$$\eta := N_I/N \in (0,1), \qquad \mu := n_H/N_I, \tag{4}$$

so that η is the total insurer ownership share of the bond and μ is the HCA fraction within the insurer sector. Because U.S. insurers are generally subject to HCA, the observed insurance share of a bond proxies empirically for $\mu\eta$, while within-insurer variation in μ (e.g., P&C vs. Life subsidiaries on the same bond) identifies the accounting channel separately from the institutional channel. These two margins, level and composition, are the bond-level variation our empirical work exploits.

The dependence of equilibrium on ownership composition (η, μ) is the model's main point of departure from Allen and Carletti (2008), where the relative size of the HCA and MTM sectors is held fixed.

[5] See Section 2.2 for SAP treatment and recognition asymmetries across insurer types.

3.2 Interest Rate Shocks, Accounting Standards, and Sales Decisions

Exogenous Systematic Shock. At $t = 1$, an exogenous shock $\varepsilon \sim U[0,1]$ raises the discount rate from r_0 to $r_1 = r_0 + \varepsilon$, depressing the bond's fair value to

$$b_2(r_1) := \frac{1}{1+r_1} p P. \tag{5}$$

A key departure from Allen and Carletti (2008), where sector-specific solvency shocks leave the common asset's value unchanged, is that our shock directly alters the bond's fundamental value. Bond sales can therefore be *efficient*: profit-maximizing when the interim price exceeds the updated continuation value, not merely forced by regulation. The issuer's repayment probability p remains unchanged, so the shock operates entirely through discounting. We consider only positive interest rate shocks here, and discuss negative shocks further below.

Although bond maturity is fixed in the model, the mechanism applies with greater force to longer-duration bonds, whose fair values decline more sharply for a given rate shock. The gap between book value b_0 and fair value b_2 widens with duration, and the accounting channel bites harder.

All agents agree on the fair value b_2. Fair value is independent of the quantity traded at $t = 1$, decoupling it from fire-sale dynamics. We model an updated demand curve $D_1(b_1 \mid r_1, b_2)$ that is continuous and strictly decreasing in b_1. The interest rate increase depresses interim demand,

$$D_1(b_1 \mid r_1, b_2) < D_0(b_0, r_0), \tag{6}$$

reflecting that tighter monetary conditions raise the opportunity cost of holding existing bonds.

Accounting Standards and Investors' Actions. Diffuse investors' sales are exogenous: impatient types liquidate at $t = 1$ regardless of price. Insurers are strategic. Each insurer $i = 1, \ldots, N_I$ chooses between full liquidation and full retention, $a_i \in \{0,1\}$, with $a_i = 1$ indicating sale; the choice is governed by regulation (for MTM insurers when the buffer is breached) or by a profit-maximizing comparison between selling at $t = 1$ and holding to maturity.

Historical Cost Accounting. HCA lets each insurer carry bonds at book value b_0 even when the interim price falls below liabilities. No HCA insurer is *required* to liquidate. Profit at $t = 1$ is

$$\Pi_{\text{HCA}}(a_i) = A_I(a_i b_1(S_1) + (1 - a_i) b_2(r_1)) - D_I, \tag{7}$$

where $S_1 = \underline{S}_1 + kA_I$ is total secondary supply when $k \le n_H + n_M$ insurers sell and $b_1(S_1)$ is the clearing price. Holding pays $A_I b_2(r_1) - D_I$; selling pays $A_I b_1(S_1) - D_I$. The recognition discretion mirrors the cross-insurance stabilizer in Allen and Carletti (2008): long-horizon investors smooth transitory shocks rather than realize losses at depressed interim prices. Let $k_H \le n_H$ denote the number of HCA insurers that sell.

Mark-to-Market Accounting. MTM insurers report at market prices. A drop in b_1 reduces measured assets directly and can breach solvency. We summarize the breach trigger by a precautionary buffer $\bar{c} > 0$: if the pre-sale clearing price $b_1(S_1)$ lies in $[0, \bar{c}]$, all MTM insurers are compelled to liquidate; if it lies above, they retain discretion.

Assumption 3.1 (MTM: Determining Precautionary Selling Buffer). *For given D_I/A_I and n_M, and every pre-MTM-sale supply $S_1 \in [\underline{S}_1, \underline{S}_1 + S_{\text{HCA}}]$, the buffer $\bar{c}$ satisfies*

$$\Big\{b_1(S_1) > \bar{c}\Big\} \iff \Big\{b_1(S_1 + S_{\text{MTM}}) > D_I/A_I\Big\}.$$

The assumption pins $\bar{c}$ at the price level above which the *post*-MTM-liquidation clearing price still meets the common solvency threshold D_I/A_I. Equivalently, solving $b_1(S_1^* + S_{\text{MTM}}) = D_I/A_I$ for the pre-sale supply level S_1^* (Figure 2) and reading back up the inverse demand curve gives $\bar{c} := b_1(S_1^*) > D_I/A_I$.

Lemma 3.1 (Behavior of Precautionary Selling Buffer). *The buffer $\bar{c}$ increases in the MTM supply S_{MTM} and increases in the debt-to-asset ratio D_I/A_I.*

Both comparative statics hold because the bond price b_1 is downward sloping in total supply S_1: a larger S_{MTM} shifts total interim supply outward, so the solvency threshold is reached at a smaller discretionary supply S_1^* and the read-back price $\bar{c} = b_1(S_1^*)$ is higher; a higher D_I/A_I raises the solvency threshold itself, which likewise shrinks S_1^* and raises $\bar{c}$. Symmetric capital structures imply that the buffer is either breached for all MTM insurers or for none.

The expected profit of an individual MTM insurer is:

$$\Pi_{\text{MTM}}(a_i) = [A_I(a_i\, b_1(S_1) + (1 - a_i)\, b_2) - D_I]\, \mathbf{1}_{\{b_1(S_1) > \bar{c}\}}, \tag{8}$$

where the indicator equals one only when the pre-sale price exceeds the buffer, giving the insurer discretion over its sales decision. When the indicator is zero, the insurer is compelled to liquidate with zero profits. Let $k_M \le n_M$ denote the number of MTM insurers that sell.

Clearing of the Bond Market. With $k_H \le n_H$ HCA insurers and $k_M \le n_M$ MTM insurers selling, total secondary supply is $S_1 = \underline{S}_1 + (k_H + k_M)A_I$, and the market clears at

$$b_1 = D_1^{-1}(S_1), \qquad D_1(\cdot \mid r_1, b_2) \text{ strictly decreasing.} \tag{9}$$

Supply and price are mutually endogenous: each insurer chooses its sale decision while internalizing its own price impact, a strategic feature absent from the competitive setting of Allen and Carletti (2008) and the atomistic setting of Plantin, Sapra, and Shin (2008). We assume the image of D_1 contains $(0, S_0]$, and do not impose clearing at $t = 2$.

Definition 3.1 (Equilibrium). *Given $(S_0, P, n_H, n_M, p, r_1, D_1)$, an equilibrium at $t = 1$ is a tuple $(b_1^*, k_H^*, k_M^*, S_1^*)$ such that each selling insurer prefers to sell and each holding insurer prefers to hold*

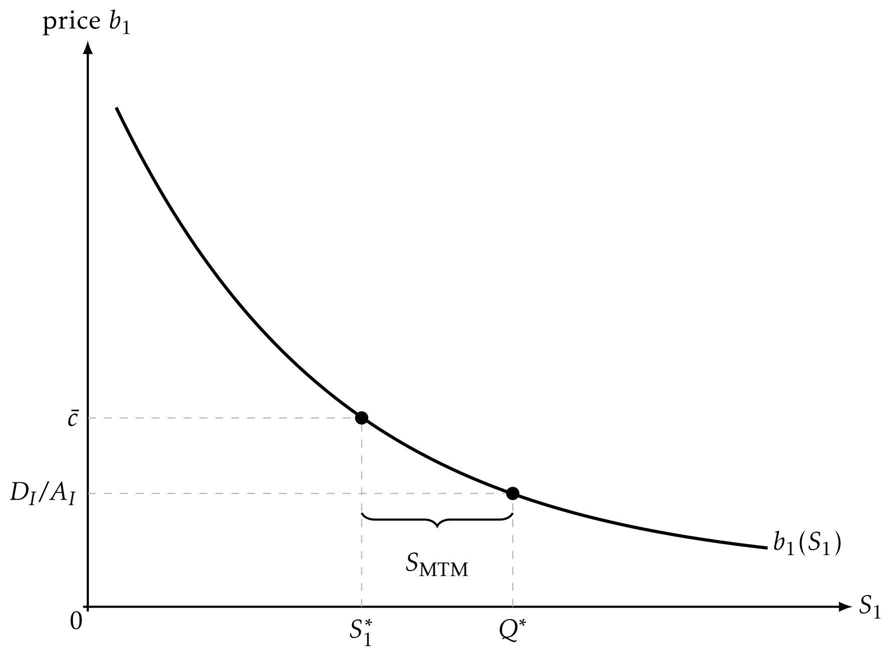

Figure 2: Precautionary Selling Buffer

Note: Construction of the precautionary selling buffer $\bar{c} := b_1\left(b_1^{-1}(D_I/A_I) - S_{\text{MTM}}\right)$. Reading down from the solvency threshold D_I/A_I locates the supply level $Q^* = S_1^* + S_{MTM}$ at which the post-MTM-sale price equals D_I/A_I; shifting left by the MTM sector's holdings S_{MTM} yields the pre-sale supply S_1^*; reading back up the inverse demand curve gives the buffer $\bar{c} = b_1(S_1^*) > D_I/A_I$.

(given the prevailing price, with MTM insurers compelled when their buffer is breached), and the market clears at $S_1^* = \underline{S}_1 + (k_H^* + k_M^*)A_I$, $b_1^* = D_1^{-1}(S_1^*)$.

A voluntary sale requires both profitability, $b_1(S_1) > b_2$, and solvency, $b_1(S_1) > D_I/A_I$. The solvency condition applies identically to HCA and MTM insurers since they share the same capital structure.

Remark 3.1 (Buffer Breach Precludes HCA Sales)**.** *Whenever the MTM buffer is breached and all MTM insurers liquidate, the post-sale clearing price necessarily satisfies* $b_1(\underline{S}_1 + S_{\text{MTM}}) \le D_I/A_I$*. The common solvency threshold then rules out any voluntary HCA sale: a buffer breach halts sales by both insurer types, even though only MTM insurers face a regulatory obligation to liquidate.*

3.3 Analysis

Given the demand curve $D_1(\cdot)$, the shock-induced continuation value b_2, and the composition (n_H, n_M), three qualitatively distinct equilibria can arise at $t = 1$, depending on whether the precautionary buffer is breached and on the size of the shock relative to b_2.

3.3.1 Optimal Coordinated Sales Decisions

Proposition 3.1 (Equilibrium Characterization)**.** *For given demand curve $D_1(\cdot)$, continuation value b_2, precautionary buffer $\bar{c}(\cdot)$, and insurer composition (n_H, n_M), the equilibrium at $t = 1$ is characterized as follows, where k_H^* and k_M^* denote the number of selling HCA and MTM insurers, respectively, S_1^* the equilibrium secondary market supply, and b_1^* the equilibrium bond price.*

(i) Case 1: All MTM insurers sell, all HCA insurers hold. *If $b_1(\underline{S}_1) \le \bar{c}(\underline{S}_1)$, then by Assumption 3.1, $b_1(\underline{S}_1 + S_{\mathrm{MTM}}) \le D_I/A_I$, and*

$$k_M^* = n_M, \quad k_H^* = 0, \quad S_1^* = \underline{S}_1 + S_{\mathrm{MTM}}, \quad b_1^* = b_1\Big(\underline{S}_1 + S_{\mathrm{MTM}}\Big) \le \tfrac{D_I}{A_I}.$$

(ii) Case 2: All insurers hold. *If $b_1(\underline{S}_1) > \bar{c}(\underline{S}_1)$ and $b_1(\underline{S}_1) \le b_2$, then*

$$k_M^* = 0, \quad k_H^* = 0, \quad S_1^* = \underline{S}_1, \quad b_1^* = b_1\Big(\underline{S}_1\Big).$$

(iii) Case 3: Profitable sales with MTM discretion. *If $b_1(\underline{S}_1) > \bar{c}(\underline{S}_1)$ and $b_1(\underline{S}_1) > b_2$:*

Case 3a: All insurers sell. If additionally $b_1(\underline{S}_1 + N_I A_I) > \max(b_2, D_I/A_I)$, then

$$k_M^* = n_M, \quad k_H^* = n_H, \quad S_1^* = \underline{S}_1 + N_I A_I, \quad b_1^* = b_1(S_1^*) > \max\Big(b_2, \tfrac{D_I}{A_I}\Big).$$

Case 3b: Partial selling, decomposition indeterminate. If additionally $b_1(\underline{S}_1 + N_I A_I) < \max(b_2, D_I/A_I)$, let

$$k^* := \sup\Big\{1 \le k \le N_I : b_1\Big(\underline{S}_1 + k \cdot A_I\Big) > \max\Big(b_2, \tfrac{D_I}{A_I}\Big)\Big\}. \tag{10}$$

The equilibrium supply and price are uniquely determined:

$$S_1^* = \underline{S}_1 + k^* A_I, \qquad b_1^* = b_1(S_1^*).$$

The decomposition of k^* into HCA and MTM sellers is indeterminate: since the MTM buffer is not breached at S_1^*, MTM insurers face no regulatory compulsion and behave identically to HCA insurers. Any pair (k_H^*, k_M^*) satisfying $k_H^* + k_M^* = k^*$, $k_H^* \le n_H$, and $k_M^* \le n_M$ constitutes an equilibrium.

Figure 5(a) depicts the case regions. Two economic forces drive the partition. Whether the equilibrium falls in Case 1 (forced MTM liquidation) or in the discretionary region (Cases 2 and 3) is governed by *ownership composition*: by Assumption 3.1, the buffer $\bar{c}$ rises with the MTM share n_M, so bonds with a high MTM share face an elevated buffer that exogenous diffuse-investor sales alone may suffice to breach. Whether the discretionary region itself falls in Case 2 or Case 3 is governed by *shock severity*: a small shock leaves b_2 high enough that no insurer finds selling profitable, while a large shock makes voluntary sales attractive for both types.

The first-order implication is that accounting matters only in Case 1. There, mandatory MTM liquidation drives the clearing price below the common solvency threshold D_I/A_I, which mechanically precludes any voluntary HCA sale. This is the contagion channel of Allen and Carletti (2008), but with one difference: its *probability* depends on ownership composition. In Cases 2 and 3, by contrast, the buffer is not breached, HCA and MTM insurers face identical incentives, and the accounting regime is immaterial. The decomposition of total sales across insurer types in Case 3 is therefore indeterminate.

Remark 3.2 (Asymmetry)**.** *For negative shocks ($\varepsilon < 0$), fair values rise, no capital pressure arises, and HCA and MTM outcomes coincide. The stabilization channel is one-sided: HCA dampens price declines but does not impede recoveries.*

Across all cases, HCA insurers sell weakly less often than MTM insurers. The corollary below makes this precise.

Corollary 3.1 (HCA versus MTM bond sales)**.** *For any fixed insurer composition (n_H, n_M) and in every case of Proposition 3.1:*

i If there is an equilibrium in which all HCA insurers sell, then either all MTM insurers also sell, or, if $n_H \leq n_M$, there is an equilibrium in which the same quantity of MTM insurers sell instead.

ii The converse does not apply: in Case 1, all MTM insurers are compelled to sell while all HCA insurers hold.

MTM's stricter constraints make HCA sales the weaker-or-equal event in every case. Because each sale endogenously depresses the clearing price, the ex-ante likelihood of a large drawdown is therefore strictly higher under MTM. By Remark 3.2, the HCA-MTM gap operates only on the downside.[6]

3.4 HCA versus MTM Ownership and Expected Prices

We now ask how the HCA/MTM composition shifts expected bond prices. To isolate accounting composition from aggregate insurer size, we hold N_I and N_o (and hence $\underline{S}_1$) fixed and compare two economies that differ by exactly one insurer's accounting standard.

Definition 3.2 (Economy 1 versus Economy 2)**.** *Economy 1 has MTM insurer count $n_M^{(1)}$ and Economy 2 has $n_M^{(2)} = n_M^{(1)} - 1$, with $n_H^{(2)} = n_H^{(1)} + 1$. Write $S_{\text{MTM}}^{(i)} := n_M^{(i)} A_I$ and $\bar{c}^{(i)}$ for the precautionary buffer in economy i, $i = 1, 2$.*

[6] On part (i): it is not generically true that whenever all HCA insurers sell, all MTM insurers also sell. This follows from Cases 2 and 3 of Proposition 3.1, where both types face identical incentives and, whenever the buffer is not breached, the decomposition of k^* is indeterminate. Consider the case $n_H < n_M$, and $k^* = 1$ with asymmetric ownership $n_H = 1$, $n_M = 3$. There is an equilibrium in which all HCA insurers sell, $k_H = n_H = 1$, but none in which all MTM insurers sell: at $k = n_M$ the clearing price would fall below the continuation value, making such sales suboptimal. There is, however, an equilibrium in which the same quantity of MTM insurers sell, $k_M = 1 \leq n_M$, confirming the disjunction in part (i). If instead $n_H > n_M$, the first disjunct always applies: whenever there is an equilibrium in which all HCA insurers sell, there exists a second equilibrium in which all MTM insurers sell. This is because with $k^* \geq n_H > n_M$, the split $(k_H, k_M) = (k^* - n_M, n_M)$ is feasible.

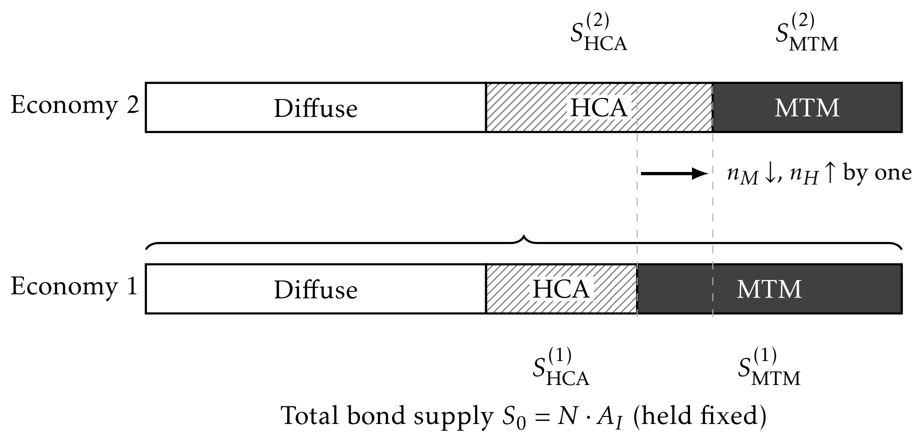

Figure 3: Economies with Different HCA vs MTM Composition

Note: Two economies differing only in HCA/MTM composition. Total insurer mass N_I, diffuse investor mass N_o, and aggregate supply S_0 are held fixed; Economy 2 differs from Economy 1 by reclassifying one MTM insurer as HCA, so $n_H^{(2)} = n_H^{(1)} + 1$ and $n_M^{(2)} = n_M^{(1)} - 1$. Each segment's width represents holdings $n_j^{(i)} A_I$ for $j \in \{H, M\}$ in economy i.

Lemma 3.2. *If $n_M^{(1)} > n_M^{(2)}$, then $\bar{c}^{(1)}(\underline{S}_1) > \bar{c}^{(2)}(\underline{S}_1)$.*

The buffer rises with both the MTM share n_M and the debt-to-asset ratio D_I/A_I (Figure 4). Liquidation by a larger MTM sector has a heavier price impact, so the buffer must be set higher to guarantee post-liquidation solvency; a more fragile capital structure raises the solvency threshold for the same reason. A higher buffer is also easier to breach through exogenous diffuse-investor sales alone, so the probability of forced MTM liquidation rises with n_M even before any insurer chooses to act. This composition dependence is the foundation for the main result below.

Proposition 3.2 (HCA Ownership and Bond Price Stability)**.** *Let Economies 1 and 2 be as in Definition 3.2. Then*

$$\mathbb{E}\big[b_1(S_1^*)\,|\,\textit{economy 1 (low HCA)}\big] < \mathbb{E}\big[b_1(S_1^*)\,|\,\textit{economy 2 (high HCA)}\big].$$

A marginal increase in the HCA ownership share, holding total insurer size and the number of insurers fixed, strictly raises the expected secondary market bond price.

Figure 5(b) illustrates: a marginal HCA-for-MTM swap shifts the buffer left from $\bar{c}^{(1)}$ to $\bar{c}^{(2)}$, shrinking Case 1 of forced MTM sales, and yielding a strictly higher expected price. The proof decomposes the expected price difference into two nonnegative pieces: in Case 1 (which occurs in both economies), Economy 2's smaller MTM sector generates strictly less forced supply; and on the intermediate strip $\bar{c}^{(2)} < b_1(\underline{S}_1) \le \bar{c}^{(1)}$, Economy 1 is in Case 1 while Economy 2 is in the discretionary region with a strictly higher clearing price.

The stabilization effect of HCA ownership runs entirely through Case 1 and combines two reinforcing channels: a smaller MTM sector lowers the buffer (so Case 1 occurs less often) and, conditional on Case 1 occurring, generates less forced supply (so the price impact is smaller). The

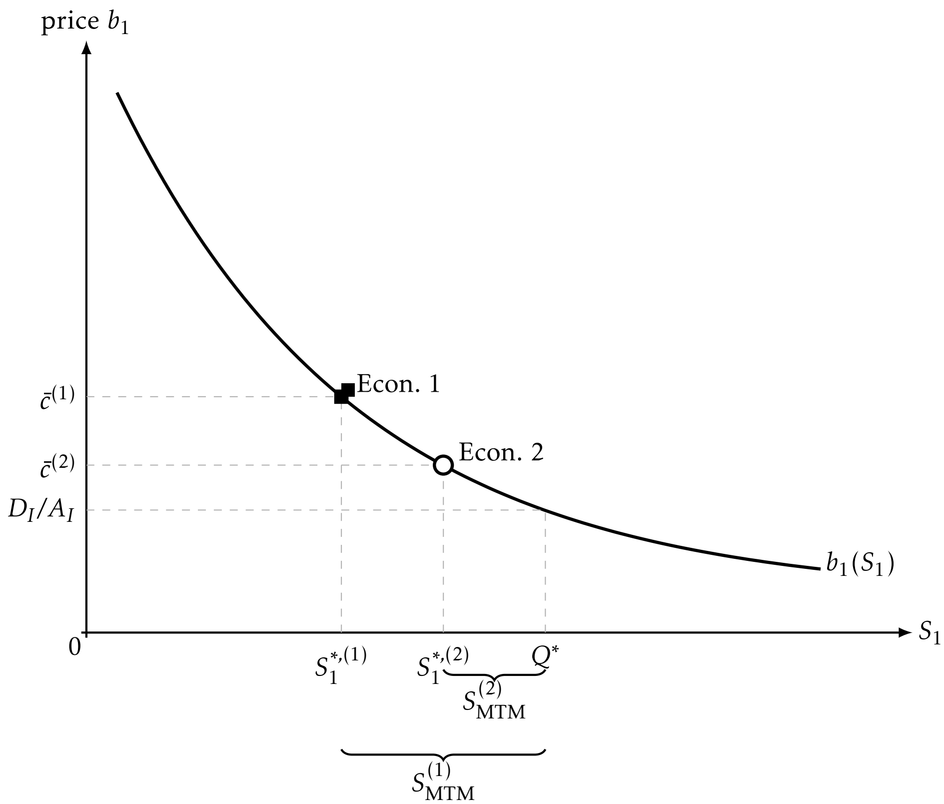

Figure 4: Comparative Static for the Precautionary Buffer

Note: Both economies share the same demand curve and solvency threshold D_I/A_I, hence the same anchor $Q^* = b_1^{-1}(D_I/A_I)$. Economy 1 has more MTM insurers ($S_{\mathrm{MTM}}^{(1)} > S_{\mathrm{MTM}}^{(2)}$), so its buffer construction reads back further from Q^*, yielding a smaller pre-sale supply $S_1^{*,(1)} < S_1^{*,(2)}$ and a higher buffer $\bar{c}^{(1)} > \bar{c}^{(2)}$. A higher buffer is easier to breach through exogenous diffuse-investor sales alone.

polar cases sharpen the intuition. In an all-MTM economy, the buffer is set high enough that diffuse sales alone routinely breach it, and the resulting forced liquidation is maximal. In an all-HCA economy, no insurer can be compelled to sell, the buffer is irrelevant, and Case 1 never arises. Proposition 3.2 says this is monotonic in between: every marginal HCA-for-MTM swap strictly reduces both the frequency and severity of forced liquidation. The stabilization properties of HCA are therefore continuous in the ownership share rather than binary. This refines the contagion channel of Allen and Carletti (2008) (the *probability* of accounting-induced contagion is itself a function of ownership composition) and extends Plantin, Sapra, and Shin (2008) beyond the polar cases of a purely MTM or purely HCA sector, generating a cross-sectional prediction in η and μ that maps directly onto bond-level data.

3.4.1 Shock Severity

Lemma 3.3. *As the shock grows and b_2 falls, equilibrium bond sales increase, driven both by regulatory compulsion of MTM insurers and by profit-maximizing sales of HCA insurers.*

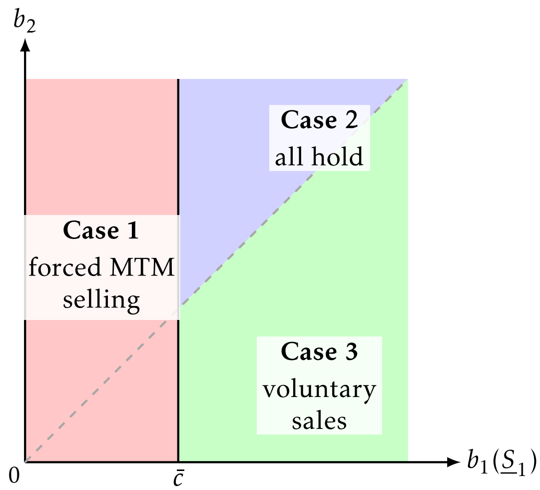

(a) Single economy: case regions.

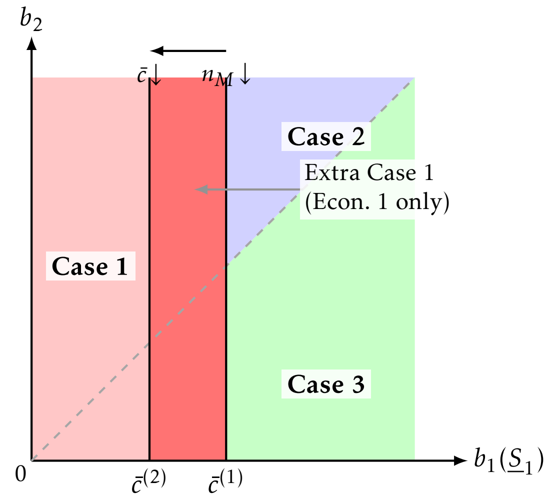

(b) Two economies: lower MTM share shifts $\bar{c}$ leftward.

Figure 5: Equilibrium Regions

Note: Equilibrium case structure as a function of the pre-sale price $b_1(\underline{S}_1)$ and the continuation value b_2. *Panel (a)*: For a single economy, the vertical line at $\bar{c}$ separates Case 1 (MTM buffer breached: all MTM insurers forced to sell, HCA insurers precluded by the common solvency constraint) from the discretionary region. The 45-degree line separates the quiescent Case 2 (selling unprofitable) from the voluntary-sales Case 3. Accounting standards matter only in Case 1. *Panel (b)*: A lower MTM ownership share shifts the buffer from $\bar{c}^{(1)}$ to $\bar{c}^{(2)}$, shrinking the Case 1 region. The darker red strip shows states that trigger forced MTM liquidation in Economy 1 but fall in Cases 2 or 3 in Economy 2, generating the expected price gap in Proposition 3.2.

A larger shock pushes b_2 further down, so the discretionary region tilts from Case 2 (selling unprofitable) toward Case 3 (voluntary sales attractive for both types), and within Case 3 from partial toward full participation. The probability of Case 1, however, is independent of shock magnitude: the buffer $\bar{c}$ depends only on $(n_M, D_I/A_I)$, neither of which moves with the shock. Severity therefore amplifies aggregate sales in the discretionary region by drawing in voluntary HCA sellers, but does not raise the probability of forced MTM liquidation.

3.4.2 Short versus Long Duration Bonds

The baseline model features a zero-coupon bond, which has duration of the time to maturity. We now vary duration by introducing a coupon bond and trace the effect on the equilibrium case structure. Throughout, both bonds face the same interim demand curve $b_1(S_1)$.

Let P^H denote the principal of the zero-coupon (high-duration) bond. The low-duration bond pays a safe coupon C at $t = 1$ and $C + P^L$ at $t = 2$ with the same repayment probability p, where P^L is principal of the low duration bond. Both bonds are normalized to a price of one at $t = 0$, which requires $C + P^L < P^H$.[7]

[7] With $\tilde{\beta} := \mathbb{E}[1/(1+r_1)]$, present-value pricing gives $1 = \tilde{\beta}\frac{pP^H}{(1+r_0)}$ and $1 = \frac{C}{(1+r_0)} + \tilde{\beta}\frac{p(C+P^L)}{(1+r_0)}$. Substituting the first equation into the second yields $(C + P^L)/P^H = 1 - C/(1 + r_0) \in (0, 1)$.

After the interest-rate shock at $t = 1$, the continuation values of the remaining claims are

$$b_2^H = \frac{pP^H}{1+r_1}, \qquad b_2^L = \frac{p(C+P^L)}{1+r_1} < b_2^H.$$

Duration operates through two separate margins.

Profitability margin (accounting-blind). Because $b_2^L < b_2^H$, the interim price b_1 exceeds the continuation value more often for the low-duration bond. Within the discretionary region, Case 3 (voluntary sales) therefore arises more often for low-duration bonds and Case 2 (all hold) more often for high-duration bonds. Since HCA and MTM insurers face identical incentives throughout Cases 2 and 3, this margin does not affect the accounting channel.

Solvency margin (accounting channel). The $t = 1$ coupon is cash in hand when the shock arrives. It offsets the insurer's liabilities dollar for dollar, independently of the bond's market price, so the solvency threshold relevant for the buffer construction falls from D_I/A_I to the effective threshold

$$\frac{D_I - CA_I}{A_I} < \frac{D_I}{A_I}.$$

By Lemma 3.1, the buffer is increasing in the debt-to-asset ratio, so $\bar{c}^L < \bar{c}^H$: mandatory MTM liquidation (Case 1) is triggered less often for low-duration bonds. Since the HCA-MTM divergence operates only through Case 1, this is the margin on which duration affects the accounting channel.

Corollary 3.2 (Duration and the Accounting Channel)**.** *Relative to the high-duration bond, the low-duration bond (i) enters Case 1 less often, since $\bar{c}^L < \bar{c}^H$, and (ii) within the discretionary region, enters Case 3 rather than Case 2 more often, since $b_2^L < b_2^H$.*

Part (i) carries the accounting content: forced MTM liquidation, and with it the HCA-MTM price gap, concentrates in high-duration bonds. The HCA stabilization effect is therefore more pronounced for longer-duration bonds. Part (ii) is accounting-blind: it raises voluntary sales of low-duration bonds identically across regimes.

3.4.3 Crisis Severity

The baseline model assumes that sales by diffuse investors are constant in the size of the shock: $\underline{S}_1(\varepsilon)$ is independent of ε. In severe crises, however, it is realistic that noise traders too respond to the size of the shock and sell more. In this extension, let $\underline{S}_1(\varepsilon)$ be increasing in ε, so the floor of secondary-market supply rises with the crisis.

As a consequence, MTM insurers breach the buffer more often. Consider a mild versus a severe shock, $\varepsilon_1 < \varepsilon_2$. Under the severe shock, noise-trader sales are larger, $\underline{S}_1(\varepsilon_2) > \underline{S}_1(\varepsilon_1)$, so for a fixed MTM sector the price that would prevail if only noise traders and the MTM sector sold is lower:

$$b_1\Big(\underline{S}_1(\varepsilon_2) + S_{\text{MTM}}\Big) < b_1\Big(\underline{S}_1(\varepsilon_1) + S_{\text{MTM}}\Big).$$

The solvency condition $b_1(S_1 + S_{\text{MTM}}) > D_I/A_I$ therefore holds less often the larger the shock, and by Assumption 3.1 the precautionary buffer is breached more often under severe shocks.

Corollary 3.3 (Crisis Severity and the Accounting Channel)**.** *If sales by noise traders rise with shock severity, then severe interest rate shocks trigger Case 1 more often than mild shocks.*

Under severe shocks the equilibrium thus falls more often into Case 1, the only region in which the accounting regime matters. HCA stabilization therefore bites more often in severe crises.

3.5 Empirical Predictions

The model has two ownership margins: the total insurer share η and, within the insurer sector, the HCA fraction μ, both of which are observable. Total insurer ownership of a bond is measured directly. Within-insurer variation in μ arises because U.S. statutory rules treat the same NAIC-3-5 bond as amortized cost in life subsidiaries (which maintain an Asset Valuation Reserve) but at the lower of amortized cost or fair value in P&C subsidiaries (which do not). From these two margins and the comparative statics derived above, the model delivers seven sharp predictions, which we state in turn and take to the data in Section 5.

Prediction 1: Drawdown–ownership relationship. Proposition 3.2 implies that, holding total insurer share fixed, a higher HCA share strictly raises expected bond prices through Case 1: a smaller MTM sector lowers the buffer $\bar{c}$ and reduces both the frequency and the severity of forced liquidation. The prediction has two empirical readings. *Level*: bonds with higher total insurer ownership η should fall less in crises, because greater insurer ownership proxies for a larger HCA investor base relative to MTM holders such as mutual funds and ETFs. *Composition*: holding η fixed, bonds with a higher within-insurer HCA fraction μ should fall less, which we test using P&C versus Life subsidiaries holding the same bond.

Prediction 2: Duration scaling. In the baseline model, duration has no bite on the accounting channel: the HCA stabilization effect of Proposition 3.2 runs entirely through Case 1, whose probability is governed by the buffer $\bar{c}$ and hence by ownership composition and capital alone. The coupon-bond extension of Section 3.4.2 changes this. A shorter-duration bond pays an interim coupon that offsets liabilities dollar for dollar, lowering the effective solvency threshold to $(D_I - CA_I)/A_I$ and with it the buffer, $\bar{c}^L < \bar{c}^H$ (Corollary 3.2). Mandatory MTM liquidation is therefore triggered less often for short-duration bonds and concentrates in long-duration bonds. The HCA stabilization effect, and thus the HCA-MTM price gap, should be more pronounced for longer-duration bonds. The extension's second margin, the lower continuation value of short-duration bonds tilting the discretionary region toward voluntary sales, is accounting-blind and generates no differential prediction across regimes.

Prediction 3: Capital proximity. The threshold $M(r_1) = \max\{b_2, D_I/A_I\}$ ties the HCA-MTM divergence to proximity to regulatory capital. A better capitalized insurer base (lower D_I/A_I) has two effects on price drawdowns, which point in different directions but are not equally decisive. First, and primarily, a lower D_I/A_I lowers the precautionary buffer $\bar{c}$ (Lemma 3.1), so the buffer is breached less often and Case 1 forced liquidation occurs less often, while the accounting-independent Cases 2 and 3 occur more often. Second, within the discretionary region and when capital is the binding constraint ($D_I/A_I > b_2$), a lower D_I/A_I lowers $M(r_1)$, making full voluntary liquidation (Case 3a) more likely, since insurers are less capital-constrained when sales are profitable. This second effect is immaterial to the accounting channel, since HCA and MTM insurers face identical incentives in Case 3a. The first effect is therefore decisive: because the HCA-MTM divergence operates only through Case 1, well-capitalized insurers are more resistant to forced liquidation regardless of accounting regime. The incremental stabilization attributable to HCA ownership should be *stronger* among weakly capitalized insurer bases.

Two remarks sharpen the empirical content of this prediction. First, it concerns the *increment*, the HCA–MTM gap, not the level of insurer stabilization. Because Case 1 grows rarer as capital improves under either regime, well-capitalized insurers of both types hold through the shock and the gap closes from above; an interaction of *total* insurer share with holder capital therefore mixes the two comparative statics and is not a test of Prediction 3. The accounting channel should instead be strongest where the probability that the buffer is breached is material: capital weak enough that forced liquidation is a live possibility. Second, the prediction is sharpest within holding companies, where the subsidiaries' own capital positions identify who is near the constraint: the P&C-minus-Life selling gap on divergently treated bonds should be concentrated among weakly capitalized subsidiaries and absent among well-capitalized ones. We test the level margin in Section 5.5 and the increment margin in Section 5.4.

Prediction 4: Crisis severity. The baseline model separates the role of crisis severity from the role of ownership structure. Whether Case 1 (forced MTM sales) or Cases 2 and 3 (discretionary sales) arises depends on the MTM ownership share and insurer capital, through the buffer $\bar{c}$, and is independent of the interest-rate shock magnitude, which affects only b_2; this independence reflects the baseline convention that the interim demand schedule, and with it the buffer $\bar{c}$ and the noise-trader supply $\underline{S}_1$, is held fixed across shock sizes. Crisis severity instead governs the split between Cases 2 and 3 and, within Case 3, the extent of voluntary participation (Lemma 3.3): fire-sale activity rises with the shock, but this amplification is accounting-blind, since HCA and MTM insurers face identical incentives throughout Cases 2 and 3. In the baseline model, larger shocks raise aggregate sales and depress prices by more, but do not widen the gap between HCA and MTM outcomes.

The extension of Section 3.4.3 overturns this. If sales by noise traders rise with shock severity, $\underline{S}_1(\varepsilon)$ increasing in ε, a more severe shock breaches the MTM selling buffer more easily and triggers Case 1 more often (Corollary 3.3). Since Case 1 is the only region in which the accounting regime matters, the HCA-MTM stabilization gap should be increasing in crisis severity and concentrated in severe episodes.

Prediction 5: Non-crisis null. For $\varepsilon \approx 0$, b_2 stays close to b_0 and no capital pressure arises under either regime. Insurer ownership should have no predictive power for returns outside crisis periods.

Prediction 6: Asymmetry. By Remark 3.2, negative shocks raise fair values and never trigger regulatory intervention. The stabilization channel is one-sided: HCA dampens price declines but does not impede recoveries. The insurer ownership effect should appear only in drawdowns.

Prediction 7: Ownership concentration. Each strategic insurer internalizes its own price impact, generating an endogenous no-sale region under HCA in Cases 2 and 3. This force is most relevant when ownership is concentrated: a larger individual holding sharpens each insurer's awareness of the price consequences of selling. Concentration among HCA insurers specifically reinforces the buffer channel as well, since a higher HCA share lowers $\bar{c}$ and reduces Case 1 probability. Because the model's strategic insurers are symmetric, we read this comparative static informally, through the price-impact force the model builds in, rather than as a derived proposition. Stabilization should therefore be stronger for bonds with more concentrated insurer ownership.

4 Data

4.1 Insurance and Mutual Fund Holdings Data

Institutional investors, including mutual funds, ETFs, and insurance companies, hold substantial shares of the corporate bond market. To analyze their bond holdings, we draw on data from multiple sources. For insurance companies, we use the quarterly "Schedule D" filings, submitted to the National Association of Insurance Commissioners (NAIC), and obtained via CapitalIQ. These filings provide detailed and standardized information on the bond holdings of U.S. insurers, allowing us to track their exposure and investment behavior over time.

Our dataset includes 5,621 distinct insurance companies, each with varying levels of bond holdings. The median dollar amount held by each insurer is approximately \$970 million, reflecting the sector's large presence in corporate bond markets. The empirical insurer ownership share of a bond, computed from Schedule D holdings, corresponds to the total insurer share η in the model of Section 3; bonds with high η are predominantly held by institutions subject to SSAP, while bonds with low η are predominantly held by institutions subject to fair-value reporting.

Bond holdings data for mutual funds and ETFs come from Morningstar. Our sample comprises 16,474 mutual funds and ETFs with an average fund size of \$231 million.[8] Mutual funds and

[8]We exclude funds designated as "Equity", "Equity - REIT", "Preferred Stock", "Undefined - Warrants/Rights", or "Cash - Currency" to retain funds with material fixed-income exposure.

ETFs report under U.S. GAAP and the Investment Company Act of 1940, both of which require fair-value reporting of fixed-income holdings, providing a natural counterfactual to the SSAP insurer treatment described in Section 2.2.

Within-Holding-Company Panel. To support the within-holding-company test of Section 5.4, we build a subsidiary-level panel that allows us to compare Life and P&C subsidiaries of the same parent insurance group holding the same bond in the same quarter. The panel is constructed from 193 annual Schedule D Excel statements released by the NAIC over our sample period, comprising approximately 14 million subsidiary-bond-year records. Because these statutory statements are filed annually whereas our holdings panel is quarterly, we assign each subsidiary-bond pair its modal SVO designation across the quarters in which the bond appears (an assignment that borrows information from later filings to maximize coverage) and carry it across the remaining quarters in which the subsidiary holds the position. This modal assignment covers 58 percent of overlapping bond-quarters; a strict contemporaneous (entity × CUSIP × quarter) merge recovers only about 23 percent, because a bond retained in the portfolio need not reappear in every annual statement.[9] Each NAIC subsidiary code is matched to its parent group via NAIC Group Codes obtained from S&P Global, yielding 4,349 matched entities. We retain the 107 holding companies that operate at least one Life and one P&C subsidiary over the sample. For each group-quarter, we flag bonds held simultaneously by at least one Life and at least one P&C subsidiary of the group. Bond NAIC designations come from two sources merged together: Mergent FISD credit ratings translated to NAIC designations using the NAIC's published rating crosswalk (the filing-exempt translation), and the NAIC designations reported directly in the Schedule D annual statements themselves. The redundancy lets us recover designations for bonds whose Mergent FISD ratings are missing or stale. The final panel contains roughly 2.45 million subsidiary-bond-quarter observations in which a Life and a P&C subsidiary of the same parent both hold the same CUSIP. The central test identifies off the NAIC 3–5 cells where the statutory rules diverge (about 92,000 subsidiary-bond-quarter observations across 60 holding companies), the within-group variation the design is built to isolate. This combination of S&P Global NAIC Group Codes with Schedule D SVO designations is, to our knowledge, new to the empirical accounting literature.

Risk-Based Capital Panel. The mechanism tests of Section 5.5 use insurer RBC ratios as a proxy for proximity to the regulatory capital constraint. We define the RBC ratio as total adjusted capital (TAC) divided by the Authorized Control Level (ACL) requirement, so that higher values indicate greater capital slack above the threshold at which regulators may seize control. We obtain RBC ratios from S&P Global at the legal-entity level and merge them to the Schedule D holdings panel using NAIC company codes. The merged panel covers 97.3 percent of insurer-bond-quarter

[9] The within-group result is robust to the designation source rather than to this assignment rule: an independent designation taken from the bond's contemporaneous Mergent FISD credit rating reproduces both the NAIC 3–5 selling gap and the NAIC 1–2 null (Online Appendix OA-2.6). Because the FISD rating is measured in the same quarter, this check embeds none of the look-ahead in the modal assignment.

holdings in our sample. RBC ratios are reported at fiscal year-end; throughout, an observation in year t carries its holder's ratio as of the end of year $t-1$, so capital is always measured before the outcomes it conditions. For the bond-level tests, we compute the par-weighted average RBC ratio across a bond's insurer holders, winsorized at the 1st and 99th percentiles, and sort holder bases into regulatory-proximity zones: below 300 percent of the ACL (encompassing the Company Action Level and the trend-test zones; the Life trend test reaches 250 percent of the ACL and the P&C test 300 percent), 300 to 500 percent, and above 500 percent.[10] For the within-group tests, each subsidiary carries its own prior-year-end RBC ratio; because the Life and P&C RBC formulas differ, we classify a subsidiary as weakly capitalized relative to subsidiaries of its own industry in the same year. Summary statistics for the insurance, mutual fund, and within-holding-company samples appear in Table 1.

4.2 Bond Data

We obtain transaction-level data for U.S. corporate bonds from the Enhanced Trade Reporting and Compliance Engine (Enhanced TRACE) database provided by FINRA. Following Dick-Nielsen (2014), we remove erroneous transactions. Weekly bond returns are computed as

$$R_{t+1} = \frac{P_{t+1}}{P_t} - 1,$$

where P_t and P_{t+1} denote the "clean" bond prices. We obtain bond characteristics (issuer cusip, offering date, offering amount, Moody's rating, S&P rating) from the Mergent Fixed Income Securities Database (FISD). To compute excess returns, we subtract the return on a one-month Treasury bill from Kenneth French's data library. Our sample period is from July 2002 to March 2020. However, since most of our analysis is confined to the crisis periods, we do not use the full time series. We keep bonds with an offering amount of more than \$25 million and require an uninterrupted trading record over the event window: weekly returns are computed only between adjacent weeks with observed transactions, and a bond enters the drawdown sample for a given crisis only if it trades in every week of that crisis window. Drawdowns are therefore measured from transaction prices throughout the window, so a bond that simply stops trading near the market trough cannot mechanically register a smaller drawdown. We then compute the drawdowns over the crisis periods outlined in Section 2.3. This leaves us with bond returns for 4,874 firms and 74,313 firm-bond pairs.

[10] We do not use the minimum RBC ratio across a bond's holders. The minimum is an order statistic that declines mechanically with the number of holders (the correlation between the minimum and the holder count is −0.29 in our sample), so it conflates holder capital with ownership breadth.

Table 1: Descriptive Statistics

Panel A: Bond Level Analysis

	N	Mean	SD	p25	Median	p75
Log drawdown	31,793	0.107	0.101	0.034	0.075	0.149
Ins. share	31,840	0.372	0.241	0.169	0.354	0.549
MF/ETF share	31,840	0.080	0.091	0.011	0.055	0.117
Bond size decile	31,840	5.827	2.816	4.000	6.000	8.000
Duration category	31,672	3.719	1.256	3.000	4.000	5.000
Credit rating (numeric)	31,840	8.412	6.335	2.000	6.000	15.000
Market value (millions)	31,840	58.504	54.600	26.153	41.102	70.569

Panel B: Firm Level Analysis

	N	Mean	SD	p25	Median	p75
CAPEX	5,908	0.054	0.062	0.017	0.039	0.069
CAPEX + Acquisitions	5,566	0.081	0.117	0.022	0.051	0.091
Firm size (total assets, millions)	5,908	26112.791	84086.717	1963.850	4785.400	15364.750
Total liabilities (millions)	5,908	20898.829	74955.607	1226.151	2997.900	10227.175
Average bond duration	5,908	5.298	2.622	3.693	4.999	6.543

Notes: This table reports descriptive statistics for the bond-level and firm-level variables used in the analysis. Panel A reports statistics for bond-level characteristics, including log bond drawdown, insurer and mutual-fund (MF/ETF) ownership shares, bond size decile, duration category, investment-grade rating, and bond market value. The bond drawdown reflects the percentage price decline of a bond during a financial crisis. Insurer and fund shares represent the percentage ownership of the bond by insurance companies and mutual funds, respectively. The decile bond size, duration category, and bond rating category reflect the relative size, duration, and creditworthiness of the bonds. The bond market value is the market price of the bond at the time of observation reported in millions of dollars. Panel A is based on the crisis-period bond-quarter sample (up to 31,840 observations). Panel B provides statistics for firm-level variables, including capital expenditures (CAPEX), the sum of capital expenditures and acquisitions (CAPEX+ACQ), firm size measured as total assets (TA), total liabilities, and the average bond duration for the firm. CAPEX reflects firms' investment in physical assets, while CAPEX+ACQ measures the total investment in both assets and acquisitions. Firm size (total assets) and total liabilities are reported in millions of dollars. Average duration refers to the weighted average duration of bonds held by the firm. Panel B is based on up to 5,908 firm-year observations.

4.3 Firm Outcomes

To analyze the real effects of insurer bond holdings on firm behavior, we examine firm-level outcomes using data from the Compustat database. Our sample includes firm-year observations from 2000 to 2020, focusing on firms with outstanding bonds as of the first quarter of 2007. The primary variables of interest are capital expenditures (CAPEX), acquisitions (AQC), total assets (AT), and total liabilities (LT). We use these to test whether insurer ownership of a firm's bonds shapes its investment and acquisition behavior around financial crises.

5 Empirical Analysis

Our empirical strategy follows the economics of the problem. The question—whether accounting recognition rules, rather than the institutions that report under them, drive crisis-period price stability—requires us to separate the rule from the investor type, and then from the firm behind the investor type. We first document that bonds with greater insurer ownership experience smaller crisis-period drawdowns (§5.1) and establish that the pattern is not driven by insurers sorting into safer bonds, using a shift-share instrument based on primary-market allocations (§5.2). We then compare insurer and mutual-fund ownership directly, since these are the two institutional holders that differ most sharply in accounting treatment, to show that the stabilization effect attaches to the HCA investor rather than to institutional ownership broadly (§5.3). Because insurers and mutual funds also differ along many non-accounting dimensions, we move inside the insurance sector, where accounting treatment varies across insurer types at the NAIC designation boundary (§5.4, Panel A) and, more sharply, between Life and P&C subsidiaries of the same holding company holding the same bond (§5.4, Panel B). Section 5.5 examines four mechanism dimensions (duration, regulatory capital, crisis severity, and ownership concentration) and §5.6 reports a scope condition (health insurers) and a further falsification exercise (non-crisis windows); a third null, the NAIC 1–2 within-group placebo, is built into the design of §5.4. Section 5.7 closes with the real consequences for firm capital investment.

5.1 Insurer Ownership and Crisis-Period Drawdowns

We first test the model's central comparative static (Proposition 3.2): replacing a marked-to-market holder with a historical-cost holder raises a bond's expected price in a crisis. Bonds with greater insurer ownership, and hence a larger historical-cost base, should therefore fall less. We measure each bond's drawdown over the four crisis windows defined in Section 2.3.

The baseline specification regresses the log drawdown ($\xi_{i,e}$) of bond i during event e on the ex ante insurer share $\phi_{i,t_{Se-1}}$ and interacted fixed effects:

$$\log(1+\xi_{i,e}) = \alpha + \beta\phi_{i,t_e^S-1} + \text{Interacted FE} + \epsilon_{i,e}, \tag{11}$$

The coefficient of interest, β, captures the semi-elasticity of drawdowns with respect to ex ante insurer shares. We measure insurer shares in the quarter before each event window (time $t_e^S - 1$) to rule out reverse causation. In the fully saturated specification, the interacted fixed effects absorb each crisis episode interacted with the bond's size decile of outstanding value, duration category, credit rating (S&P and Moody's), and coupon and seniority type, together with issuer fixed effects, so that identification comes from comparing otherwise-similar bonds of the same issuer within the same crisis.

Table 2 reports the estimates of equation (11). The unconditional relationship in column 1 is negative: a one-standard-deviation increase in insurer holdings (24 percentage points) is associated with a drawdown that is 0.56 percentage points smaller, roughly 5 percent of the mean crisis drawdown of 10.7 percent.

Table 2: Insurer Shares and Bond Drawdowns

	Dependent variable: Log Drawdown					
	(1)	(2)	(3)	(4)	(5)	(6)
Ins. share	-0.0234***	-0.0227***	-0.0648***	-0.0467***	-0.0156***	-0.0124***
	(0.0068)	(0.0067)	(0.0071)	(0.0066)	(0.0038)	(0.0038)
Crisis episode FE	No	Yes	Yes	Yes	Yes	Yes
Crisis × Size decile FE	No	No	Yes	Yes	Yes	Yes
Crisis × Duration FE	No	No	Yes	Yes	Yes	Yes
Crisis × Rating FE	No	No	No	Yes	Yes	Yes
Crisis × Issuer FE	No	No	No	No	Yes	Yes
Crisis × Coupon/Senior FE	No	No	No	No	No	Yes
Observations	31,793	31,793	31,672	31,671	31,246	31,245

Notes: This table reports estimates of equation (11) on the sample of actively traded U.S. corporate bonds. The dependent variable is the log bond drawdown over each of the four crisis events. Column (1) reports the unconditional relation; columns (2) through (6) progressively saturate the fixed-effect structure, interacting crisis events with the bond's size decile, duration category, credit rating, issuer, and coupon and seniority type. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Columns 2 through 6 progressively saturate the fixed-effect structure with crisis-event dummies. Column 2 adds crisis-event fixed effects to absorb variation in severity and market conditions across episodes; columns 3 and 4 add crisis-specific controls for bond size, duration, and rating, and the coefficient barely moves. Column 5 adds crisis × issuer fixed effects, absorbing all within-firm, within-crisis variation, and column 6 layers on crisis × coupon/senior interactions for a fully saturated specification. The coefficient attenuates but remains economically meaningful: a one-standard-deviation increase in insurer share is associated with a drawdown that is 0.30 percentage points smaller, about 3 percent of the mean crisis drawdown.

The relation looks modest in proportional terms, but it is economically large once scaled to the dispersion in ownership. Insurer holdings vary widely across bonds: moving across the interquartile range of insurer ownership, from 17 to 55 percent, compresses crisis drawdowns by 0.5 to 0.9 percentage points, or 4 to 8 percent of the mean crisis drawdown. Because insurers hold 37 percent of the bonds in our sample on average, a per-bond effect of this size dampens fire-sale dynamics across a large share of outstanding credit precisely when liquidity is scarcest.

That our estimate survives these increasingly demanding fixed effects rules out explanations tied to observable bond or issuer heterogeneity. Bonds held by insurers, who can defer loss recognition under HCA, decline less during crises than otherwise similar bonds held by MTM investors, the pattern the model predicts when HCA dampens insurers' incentive to sell into falling markets. While the OLS coefficient is consistent with the model, it cannot fully identify the accounting channel.

5.2 A Primary-Market Shift-Share Instrument

Even with saturated fixed effects, insurer ownership could be endogenous if insurers systematically sort into safer or less volatile bonds. To break this selection channel, we use a shift-share IV design (Bartik, 1991; Goldsmith-Pinkham et al., 2020; Breuer, 2022) that exploits quasi-random variation in how new bond issues are allocated to insurers in the primary market. The institutional details motivating the design come from Coppola (2025).

Large insurers buy bonds at issuance and hold them for long periods. Figure 6 confirms this: the share of a bond held by insurers at issuance predicts insurer ownership many quarters later, with pass-through coefficients close to one over horizons of up to seven years. A bond's initial investor composition is effectively sticky, so the allocation determined at issuance shapes insurer exposure for years.

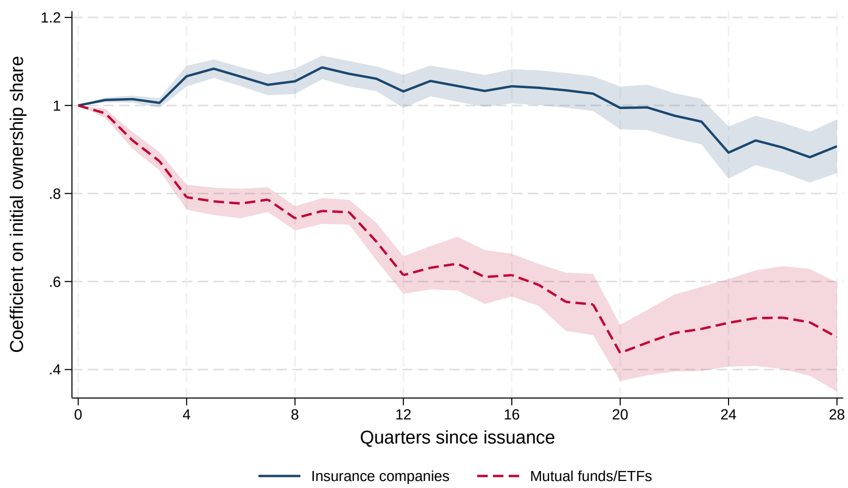

Figure 6: Insurer vs Mutual Fund Persistence of holding decisions for bonds acquired at issuance

Notes: This figure displays the persistence of bond ownership by investor type over time. For each bond, we measure the share held by insurers at issuance and track how that share evolves at quarterly intervals for up to 28 quarters. The solid blue line displays the estimated relationship between insurer holdings at issuance and insurer holdings at each subsequent quarter, based on the following regression:

$$\text{Share}_{i,\tau(i)+Qs} = \alpha_s + \beta_s \, \text{Share}_{i,\tau(i)} + \varepsilon_{i,s},$$

where $\text{Share}_{i,\tau(i)}$ denotes the insurer ownership share of bond i at issuance date $\tau(i)$, and $\text{Share}_{i,\tau(i)+Qs}$ is the ownership share s quarters later. A coefficient β_s close to 1 indicates high persistence. For comparison, the red line plots the same relationship using ownership by mutual funds and ETFs instead of insurers.

Two institutional features drive the instrument's quasi-random variation. First, insurers buy only a subset of new issues, and their participation is based on idiosyncratic liquidity needs and underwriter relationships, rather than on bond fundamentals. Second, primary-market allocations are highly concentrated within an offering: the median new issue is taken up by four insurer entities, and the largest buyer absorbs roughly half of the total insurer allocation, while smaller entities participate sporadically. Thus, issuance timing interacted with these participation patterns generates exogenous variation in which bonds end up in insurer portfolios.

We construct the instrument by comparing insurer ownership across bonds issued in the same quarter and maturity category but differing in credit rating, isolating deviations from expected ownership that stem from allocation rather than fundamentals. Formally, the shift-share instrument Z_i is defined as follows:

$$Z_i = \sum_{j \in J} a_{i,j} \bar{\phi}_j, \tag{12}$$

where $a_{i,j}$ equals one if insurer j purchased bond i at issuance and $\bar{\phi}_j$ is the average size of insurer j's primary-market bond purchases relative to the aggregate insurance sector, normalized so that $\sum_{j \in J} \bar{\phi}_j = 1$. The purchase indicators $a_{i,j}$ play the role of shares, and the size coefficients $\bar{\phi}_j$ play the role of shifters; Z_i therefore captures the collective size of the insurers acquiring bond i at issuance.

A decomposition of insurer ownership clarifies the sources of identifying variation. The insurer share $\phi_{i,t}$ decomposes into its value at issuance $\phi_{i,\tau(i)}$ and deviations from secondary-market trading, $\phi_{i,t}^{\text{sec}} = \phi_{i,t} - \phi_{i,\tau(i)}$. Summing over potential outcomes for individual insurers decomposes $\phi_{i,\tau(i)}$ further:

$$\phi_{i,t} = \sum_{j \in J} a_{i,j} \cdot \tilde{\phi}_{i,j,\tau(i)} + \phi_{i,t}^{\text{sec}} \tag{13}$$

where J is the set of insurers in the sample, $a_{i,j}$ indicates whether insurer j acquired bond i at issuance, and $\tilde{\phi}_{i,j,\tau(i)}$ is the share of bond i that insurer j would acquire on the primary market conditional on participating, a potential outcome shaped by the insurer's capital position and other factors. The three terms map to three distinct sources of variation: secondary-market trading ($\phi_{i,t}^{\text{sec}}$), the intensive margin of primary-market trading ($\tilde{\phi}_{i,j,\tau(i)}$), and the extensive margin of primary-market trading ($a_{i,j}$). The shift-share instrument isolates the extensive margin.

To build Z_i, we strip out the secondary-market component $\phi_{i,t}^{\text{sec}}$ and replace the intensive-margin terms $\tilde{\phi}_{i,j,\tau(i)}$ with time- and bond-invariant measures of insurer size. The resulting instrument takes higher values when the insurers participating at issuance are particularly large.

The exclusion restriction requires $E[Z_i \epsilon_{i,e} | F_{i,e}] = 0$, where $F_{i,e}$ contains all observables in the regression and $\epsilon_{i,e}$ is the residual. Conditional on these observables, the matching between insurers and bonds within a quarter must be orthogonal to unobserved drivers of future drawdowns. Because the purchase indicators $a_{i,j}$ are themselves insurers' equilibrium participation choices rather than predetermined exposure shares, the identifying assumption is this conditional quasi-randomness of the issuance-quarter matching, not exogeneity of the shares themselves. Evidence that allocations track insurer liquidity and underwriting relationships rather than ex ante bond risk supports this assumption (Coppola, 2025), and Goldsmith-Pinkham et al. (2020) and Breuer (2022) show that shift-share designs deliver plausibly exogenous exposure shocks when sectoral participation is concentrated and timing is quasi-random.[11]

5.2.1 Two-Stage Least Squares (2SLS) Estimation

We estimate the causal effect of insurer ownership on crisis-period drawdowns by two-stage least squares (2SLS). The first stage regresses insurer ownership on the shift-share instrument Z_i:

$$\phi_{i,t_{Se-1}} = \gamma Z_i + \text{Interacted FE} + \nu_{i,e}, \tag{14}$$

with first-stage residual $\nu_{i,e}$. The second stage then substitutes the fitted values $\widehat{\phi}_{i,t_e^S-1}$ for insurer shares in the baseline regression:

$$\log(1+\xi_{i,e}) = \alpha + \beta\widehat{\phi}_{i,t_{Se-1}} + \text{Interacted FE} + \epsilon_{i,e}, \tag{15}$$

where $\widehat{\phi}_{i,t_{Se-1}}$ denotes the first-stage fitted values.

Table 3 reports the 2SLS results. The first-stage coefficient on Z_i is positive and significant.

The IV sample is smaller than the OLS sample (4,468 versus 31,793 bond-crisis observations) because the instrument is defined only for bonds whose issuance-quarter insurer allocation we observe, a restriction that is not selection on outcomes.[12]

[11]Online Appendix OA-7 reports two diagnostics for the instrument. Rotemberg-style weights in the spirit of Goldsmith-Pinkham et al. (2020) show identification is diffuse across insurers (the largest entity carries 5.3 percent of the total weight and the top ten carry 27.5 percent), and balance tests show that conditional on the regression's observables, residual associations between the instrument and pre-determined bond characteristics are at most 0.06 standard deviations.

[12]The instrument is built from insurers' primary-market purchases in a bond's offering quarter, so roughly one bond in seven in the drawdown sample qualifies; seasoned issues predating the statutory holdings panel, private placements, and issues insurers did not buy at origination have no observed initial allocation. The subsample is not a selected slice of the market: the OLS coefficient on the restricted IV sample (−0.022, Table 3, column 2) is essentially identical to the full-sample estimate (−0.023, Table 2).

Table 3: Insurer Shift-Share IV Regression

	OLS		2SLS	
	(1)	(2)	(3)	(4)
Ins. share	-0.0264***	-0.0225**	-0.1122***	-0.1164***
	(0.0092)	(0.0094)	(0.0285)	(0.0361)
Crisis episode FE	Yes	Yes	Yes	Yes
Crisis × Size decile FE	Yes	Yes	Yes	Yes
Crisis × Duration FE	Yes	Yes	Yes	Yes
Crisis × Rating FE	Yes	Yes	Yes	Yes
Crisis × Issuer FE	Yes	Yes	Yes	Yes
Crisis × Coupon/Senior FE	No	Yes	No	Yes
Estimator	OLS	OLS	IV	IV
First-stage F			108.7	77.8
Observations	4,468	4,468	4,468	4,468

Notes: This table reports shift-share instrumental-variable estimates of the effect of insurer ownership on crisis-period bond drawdowns, alongside the corresponding OLS estimates, on the sample of actively traded U.S. corporate bonds. The dependent variable is the log bond drawdown. Columns (1) and (2) report OLS; columns (3) and (4) report 2SLS estimates that instrument insurer share with the primary-market shift-share instrument Z_i (first stage in equation (14)). The fixed-effect structure interacts crisis events with bond size, duration, credit rating, and issuer (columns (1) and (3)), and additionally with coupon and seniority type (columns (2) and (4)). First-stage F-statistics are reported for the instrumented columns. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

The second-stage coefficient on instrumented insurer share ranges from −0.112 (column 3) to −0.116 (column 4), substantially larger than OLS and significant at the 1 percent level. First-stage F-statistics of 108.7 and 77.8 clear the weak-instrument cutoffs. The 2SLS estimates thus assign a causal role to insurer ownership in dampening bond-price declines during downturns. Whether that role belongs to the accounting rule the insurer reports under or to some other feature of insurers is the question we take up in the remainder of this section.

5.3 The Rule versus the Institution

Mutual funds, unlike insurers, mark portfolios to market and report daily net asset values, a reporting regime that can force sales when prices fall.[13] If mark-to-market recognition amplifies fire-sale dynamics, bonds tilted toward mutual funds should drop further during crises than bonds tilted toward insurers. Entering insurer share and mutual-fund/ETF share jointly separates the HCA investor from institutional ownership in general. The two shares move together because both measure the concentration of institutional holdings, but they differ sharply in accounting treatment, so their relative coefficients indicate which channel drives the stabilization. Figure 7

[13] See Appendix OA-1 for details on mutual fund accounting.

graphically previews the behavior during the Great Recession, with insurers holding their positions while mutual funds sell into the decline.

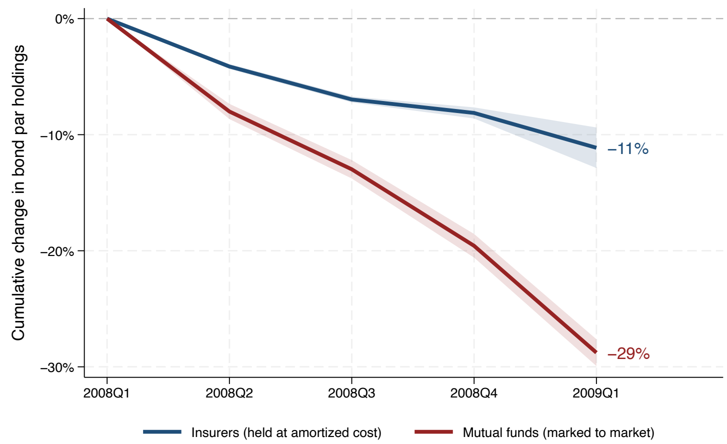

Figure 7: Insurer vs Mutual Fund Bond Sales During Great Recession

Notes: This figure displays how the average insurer's (blue) and average mutual fund or ETF's (red) corporate bond holdings changed during the Great Recession, starting from the first quarter of 2008 (2008Q1). Changes are measured in par values, so they show actual bond amounts held and are unaffected by market price changes. The shaded areas denote 95 percent confidence intervals.

Table 4 reports the joint estimates under the fully saturated fixed-effect structure (crisis × size decile, duration, rating, issuer, and coupon/seniority). Column (1) enters only standardized insurer share: the coefficient is −0.0030, significant at the 1 percent level. Column (2) enters only standardized MF/ETF share: +0.0008, indistinguishable from zero. Column (3) runs both simultaneously; the insurer coefficient is essentially unchanged at -0.0029^{***} while the MF/ETF coefficient remains indistinguishable from zero at +0.0004. Column (4) reports the same joint specification in levels, delivering insurer share = -0.0121^{***} and MF/ETF share = +0.0047.

The interpretation is sharp: the insurer coefficient barely moves when mutual-fund share enters as a competing regressor, while the MF/ETF coefficient remains near zero whether entered alone or jointly with insurer share. Once the fixed effects absorb rating and issuer heterogeneity, mutual-fund ownership has no explanatory power for crisis drawdowns, and stabilization is specific to insurer holdings.

Figure 8 visualizes the contrast by replacing the continuous shares with decile indicators. Higher insurer-share deciles experience smaller drawdowns while higher MF/ETF-share deciles show no such pattern.

Table 4: Insurer and Mutual Fund Shares: Joint Estimates

	Standardized			*Levels*
	(1)	(2)	(3)	(4)
Ins. share (std.)	-0.0030***		-0.0029***	
	(0.0009)		(0.0010)	
MF/ETF share (std.)		0.0008	0.0004	
		(0.0006)	(0.0006)	
Ins. share				-0.0121***
				(0.0039)
MF/ETF share				0.0047
				(0.0070)
Crisis episode FE	Yes	Yes	Yes	Yes
Crisis × Size decile FE	Yes	Yes	Yes	Yes
Crisis × Duration FE	Yes	Yes	Yes	Yes
Crisis × Rating FE	Yes	Yes	Yes	Yes
Crisis × Issuer FE	Yes	Yes	Yes	Yes
Crisis × Coupon/Senior FE	Yes	Yes	Yes	Yes
Observations	31,245	31,245	31,245	31,245

Notes: This table reports joint estimates of equation (11) comparing the effects of insurer share and mutual-fund/ETF share on crisis-period bond drawdowns, on the sample of actively traded U.S. corporate bonds. The dependent variable is the log bond drawdown. Columns (1) and (2) enter each ownership share individually in standardized form; column (3) enters both simultaneously in standardized form; and column (4) enters both in levels rather than standardized units. All specifications include the fully saturated fixed effects, interacting crisis events with the bond's size decile, duration category, credit rating, issuer, and coupon and seniority type. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

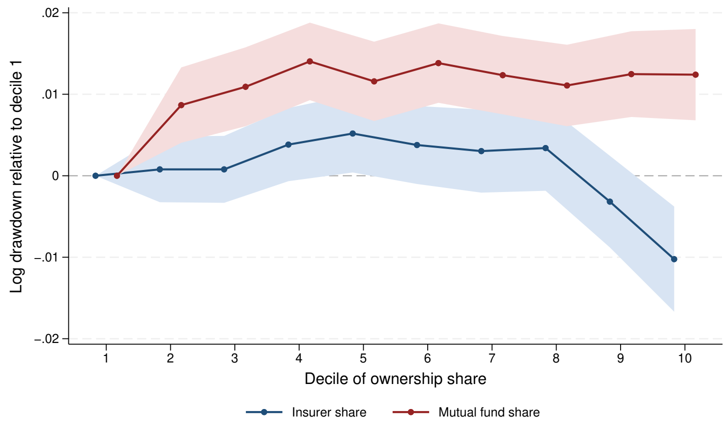

Figure 8: Drawdowns by Insurer and Mutual Fund Ownership Decile

Notes: This figure plots decile-indicator estimates from a modified version of equation (11). Decile indicators for insurer shares are plotted in blue; decile indicators for mutual fund and ETF shares are plotted in red. Higher insurer-share deciles are associated with smaller drawdowns; the mutual-fund relationship is flat after controls.

The joint specification pins the stabilization effect on the HCA investor rather than institutional ownership at large. The comparison holds across NAIC designations 1–5 because insurers carry all such bonds at amortized cost under statutory accounting, while mutual funds mark all holdings to market under the Investment Company Act of 1940. Insurers and mutual funds also differ in the liability structure, investment horizon, redemption risk, and regulatory environment, all of which could in principle drive the result. These remaining alternatives motivate the next step of the analysis, which moves *inside* the insurance sector.

5.4 Evidence from Within the Insurance Sector

The prior horse race pinned stabilization on insurers, but it cannot separate the recognition rule from institutional form: moving from a mutual fund to an insurer changes the total insurer share η, the within-sector HCA fraction μ, and the entire institutional bundle at the same time. The theory's Prediction 1 (§3.5) splits these objects cleanly. The level of stabilization scales with η, the share of the bond held under amortized cost; its composition, holding η fixed, scales with μ, the share of insurer holdings that fall under the lenient SSAP regime. The previous sections varied η across the institutional boundary. To identify the accounting channel itself, we now look inside the insurance sector, where η is effectively held fixed and μ varies across rating buckets, across subsidiary types, and across siblings of the same parent.

Two complementary tests follow, the first exploiting a discrete μ shock at the NAIC 2→3 downgrade boundary: when a bond crosses into NAIC 3, P&C subsidiaries must flip from amortized cost to the lower of amortized cost or fair value (LOCOM) because they do not maintain an Asset Valuation Reserve, while Life subsidiaries, which do, continue carrying it at amortized cost. The same rating action that changes μ for P&C-held bonds leaves μ unchanged for Life-held bonds. The second test exploits μ-variation within a single bond: a Life and a P&C subsidiary of the same parent holding the identical bond in the identical quarter differ only in their accounting treatment of that bond.

Accounting Divergence at the NAIC 2→3 Boundary

The NAIC 2→3 downgrade is the cleanest discrete μ shock in the data. The divergence is mechanical, triggered by the rating action rather than chosen by the insurer, and it does not apply at the NAIC 1→2 boundary, which therefore serves as a placebo.

Figure 9 plots event-study dynamics around the downgrade, separating log insurer holdings for bonds concentrated among Life (blue) versus P&C (red) insurers. Holdings track closely before the downgrade; afterwards both decline, but P&C holdings fall markedly further than Life holdings, consistent with the accounting mechanism.

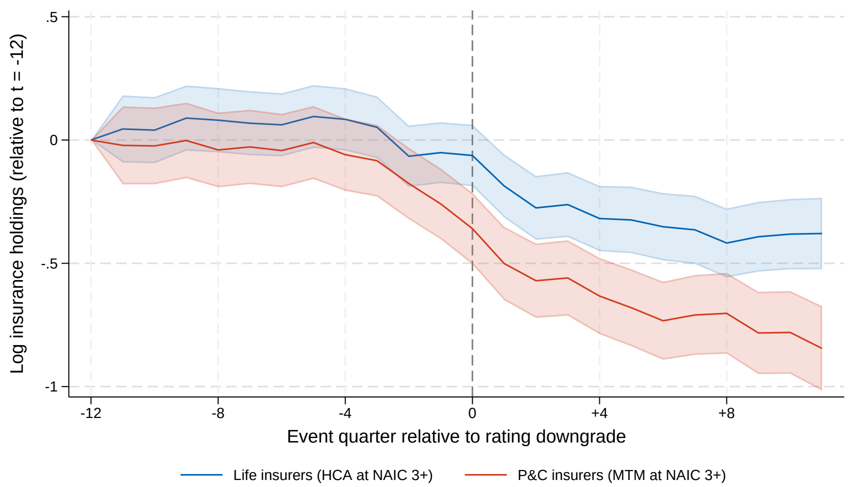

Figure 9: Insurer Holdings around NAIC 2→3 Downgrades: Life vs. P&C

Notes: This figure plots event-study coefficients for log insurer holdings around NAIC 2→3 rating downgrades, estimated separately for Life insurers (blue) and P&C insurers (red) and normalized to event quarter −12. The event window spans $[-12,+8]$ quarters around the downgrade. Specifications control for the bond price and absorb duration-category and calendar-quarter fixed effects. Pre-downgrade, the two holding paths track closely. Post-downgrade, P&C holdings fall sharply, consistent with the lower-of-amortized-cost-or-fair-value requirement that applies to non-AVR insurers; Life holdings decline far less, consistent with the amortized-cost treatment that AVR insurers retain. Shaded areas are 95 percent confidence intervals.

Panel A of Table 5 reports the price counterpart to these holdings paths. Let $\Delta\phi_i = \phi_i^{Life} - \phi_i^{P\&C}$ denote the pre-downgrade Life–P&C ownership gap for bond i. Column (1) regresses the bond's quarter-end return on $\Delta\phi$ interacted with a post-downgrade indicator. The coefficient is -0.0144^{***}: bonds that enter the downgrade in P&C-concentrated hands earn *higher* returns over the post-downgrade window than bonds in Life-concentrated hands. This is the signature of forced-sale price pressure and its reversal (Ellul et al., 2011): the pattern is consistent with the MTM holders' selling at the boundary pushing the price of P&C-concentrated bonds below fundamental value, with subsequent returns recovering the discount; Life-concentrated bonds, whose holders are not forced to sell (Figure 9), never overshoot and so have nothing to recover. Column (2) decomposes the gap into its two pieces: Life share × Post enters at −0.0146*** (Life-concentrated bonds do not participate in the rebound that the average downgraded bond experiences), while P&C share × Post is positive though imprecise (+0.0229). Column (3) is the placebo on NAIC 1→2 downgrades, where no accounting divergence occurs; here the coefficient is −0.0011, tightly estimated around zero: where no holder is forced to sell, there is no discount to reverse. The estimate remains significant at the 10 percent level under two-way clustering on bond and quarter and is robust to a wider event window. The cross-bond pre-trend it exhibits reflects issuer selection; it is absorbed by issuer × quarter fixed effects, under which the downgrade effect is preserved, and it is removed by construction in the within-group test below (Online Appendix OA-5.3). The boundary test identifies the μ-variation between bonds; our next test identifies it within a single bond, keeping the parent group constant.

Table 5: Between and Within Insurer Types

Panel A: NAIC 2→3 Boundary — Bond Returns around Downgrades

	NAIC 2→3 downgrades		NAIC 1→2 placebo
	(1)	(2)	(3)
(Life − P&C) share × Post	-0.0144***		
	(0.0042)		
Life share × Post		-0.0146***	
		(0.0043)	
P&C share × Post		0.0229	
		(0.0241)	
(Life − P&C) share × Post			-0.0011
			(0.0019)
Bond FE	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes
Clustering	Bond	Bond	Bond
Observations	8,824	8,824	14,302

Panel B: Within-Insurance-Group Selling Behavior

	Dependent variable: Sold		
	(1) NAIC 3–5	(2) NAIC 1–2 placebo	(3) Pooled
P&C subsidiary	0.0141*	-0.0012	-0.0012
	(0.0085)	(0.0012)	(0.0012)
NAIC 3–5			-0.0017
			(0.0034)
P&C × NAIC 3–5			0.0149*
			(0.0079)
Group × Bond × Qtr FE	Yes	Yes	Yes
Clustering	Group	Group	Group
Observations	92,159	1,318,120	1,426,103

Panel C: Within-Group Selling Gap by Subsidiary Capitalization

	Dependent variable: Sold			
	(1)	(2)	(3)	(4)
P&C subsidiary	-0.0015	-0.0019*	-0.0033*	-0.0006
	(0.0013)	(0.0011)	(0.0019)	(0.0011)
P&C × NAIC 3–5	-0.0036	0.0181	-0.0047	0.0040*
	(0.0085)	(0.0119)	(0.0096)	(0.0024)
NAIC 3–5 × Weak	-0.0199**	-0.0203	-0.0166	0.0046
	(0.0094)	(0.0227)	(0.0127)	(0.0033)
P&C × NAIC 3–5 × Weak	0.0274*	0.0095	0.0274	-0.0041
	(0.0145)	(0.0196)	(0.0274)	(0.0046)
Wild-cluster bootstrap p-value (triple)	0.045	0.681	0.342	0.437
Weak definition	Below median	RBC < 300%	Below median	Below median
NAIC designation	Modal	Modal	Modal	Contemp.
Sample	All SVO-rated	All SVO-rated	Crisis years	All SVO-rated
Group × Bond × Qtr FE	Yes	Yes	Yes	Yes
Observations	1,425,547	1,425,547	249,926	415,895

Notes: This table examines whether bond-price stability and insurer selling at the NAIC 2/3 rating boundary reflect the accounting rule applied to a bond's holder rather than the holder's institutional type. *Panel A* presents estimates of bond returns around NAIC 2→3 rating downgrades. At the NAIC 2/3 boundary, accounting treatment switches from amortized cost to the lower of amortized cost or fair value (LOCOM) for property-and-casualty (P&C) insurers, which do not maintain an Asset Valuation Reserve (AVR), whereas life insurers, which do, continue to carry the bond at amortized cost. The dependent variable is the bond's quarter-end return over the $[-4,+4]$ quarter event window around the downgrade, and pre-downgrade life and P&C ownership shares are measured at the quarter preceding the downgrade. A negative coefficient on (Life − P&C) share × Post indicates that bonds entering the downgrade in P&C-concentrated hands earn higher post-downgrade returns, the reversal of the fire-sale discount created by forced P&C selling at the boundary. Column (3) presents a placebo using NAIC 1→2 downgrades, at which no accounting change occurs for either insurer type. *Panel B* presents estimates of within-insurance-group selling behavior, comparing P&C and life subsidiaries of the same holding company that hold the same bond in the same quarter. The dependent variable is an indicator for a quarter-over-quarter decline in a subsidiary's holdings of the bond, and NAIC designations are obtained from Schedule D SVO filings. Column (1) restricts the sample to NAIC 3–5 bonds, for which the recognition rule diverges across holder types, and column (2) presents a placebo on NAIC 1–2 bonds, for which it does not. *Panel C* augments the Panel B specification with an indicator for a weakly capitalized subsidiary and its interactions. Weak equals one if the subsidiary's risk-based capital (RBC) ratio at the prior year-end is below the median among subsidiaries of its own industry (life or P&C) in that year (columns (1), (3)–(4)) or below 300 percent of the Authorized Control Level (column (2)); capital is classified within industry because the life and P&C RBC formulas are not level-comparable. Column (3) restricts the sample to crisis years (2008, 2009, and 2020), and column (4) assigns NAIC designations from contemporaneous-quarter SVO filings only rather than the modal carried-forward assignment. The Weak main effect and the P&C × Weak interaction are included in every specification but not reported (both are small and statistically indistinguishable from zero); main effects absorbed by the fixed effects are omitted. Panel C samples exclude subsidiary-quarters without a matched prior-year-end RBC ratio, which accounts for the small difference in observation counts relative to Panel B. Standard errors, clustered by bond in Panel A and by insurance holding company in Panels B and C, are reported in parentheses. Because Panels B and C have few clusters (60 holding companies in Panel B column (1) and 94 in the pooled specifications), we also report wild-cluster bootstrap p-values (Rademacher weights, 9,999 replications, clustered by holding company): $p = 0.10$ for the NAIC 3–5 selling gap in Panel B column (1), $p = 0.04$ for the P&C × NAIC 3–5 interaction in Panel B column (3), and the values reported in the Panel C table rows for the triple interaction. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively, based on analytic standard errors.

The Within-Holding-Company Test: Same Bond, Same Quarter, Different Accounting

Ellul et al. (2015) compare Life and P&C insurers using rating-group and asset-pool fixed effects and, in one robustness check, limit the sample to Life and P&C insurers that are part of groups containing both types, thereby holding asset-management expertise constant (their Internet Appendix Table IAXIII). However, they do not employ a *same-security*, within-holding-company fixed-effects framework, namely, a design that, by construction, keeps investment philosophy, capital position, liability structure, and parent-level risk culture constant while varying only the statutory accounting treatment of a given CUSIP. We build exactly this subsidiary-level sample and use it to clearly isolate the accounting channel.

Our sample begins with the 193 annual Schedule D Excel filings released by the NAIC, parsed into approximately 14 million subsidiary-bond-year holdings. Each subsidiary is matched to its parent group using NAIC Group Codes obtained from S&P Global, and we retain the 107 holding companies that operate at least one Life and one P&C subsidiary. For every group-quarter, we identify bonds that are concurrently held by both types of subsidiaries. Bond ratings are linked to NAIC designations using Mergent FISD, combined with the SVO-reported designations listed in the annual Schedule D filings. For each bond, we take the most frequently observed designation across all quarters in which it appears and then apply that designation to all quarters during which the subsidiary holds the bond. This redundancy allows us to infer designations for bonds that lack Mergent FISD ratings. The final panel contains roughly 2.45 million subsidiary-bond-quarter observations in which a Life and a P&C subsidiary of the same parent both hold the same CUSIP. Our primary test obtains its identifying variation from NAIC 3–5 cells in which statutory treatment differs, roughly 92,000 subsidiary-bond-quarter observations spanning 60 holding companies, providing us with the within-group comparison that the design is intended to isolate.

The identification rests on what the group × bond × quarter fixed effect $\alpha_{g\times i\times t}$ absorbs. Every characteristic of the bond (i.e., issuer, sector, duration, rating, coupon, seniority, callability) is held fixed at the bond fixed effect. Every attribute of the parent group, such as investment philosophy, internal capital market, liquidity policy, risk culture, capital position, is also held fixed by the group fixed effect. Every macro shock common to the quarter is held fixed by the time fixed effect. That leaves the identifying variation between two subsidiaries of the same parent, holding the same CUSIP in the same quarter. The siblings still differ along type-specific dimensions (i.e., liability duration, product-line cash flows, tax positions, and the RBC formula each faces) but those differences are present at every rating; what changes discretely at the NAIC 2/3 boundary, and only there, is the statutory accounting rule each subsidiary applies to the bond. It is the NAIC 1–2 placebo below, rather than the fixed effect alone, that disciplines these residual type-level differences. This fixed-effects structure is comparable to the within-bond-quarter saturation in Becker and Ivashina (2015) and the within-issue-event designs of Ellul et al. (2011). The estimating equation is

$$\text{Sold}_{i,g,s,t} = \beta_1 \cdot \mathbf{1}[\text{P\&C}]_s + \beta_2 \cdot \mathbf{1}[\text{P\&C}]_s \times \mathbf{1}[\text{NAIC 3–5}]_{i,t} + \alpha_{g\times i\times t} + \varepsilon_{i,g,s,t}, \tag{16}$$

where i indexes the bond, g the holding company group, $s \in \{\text{Life}, \text{P\&C}\}$ the subsidiary type, and t the quarter. The dependent variable equals one if subsidiary s reduces its position in bond i during quarter t. Standard errors are clustered at the holding-company-group level.

Before reading the numbers, note that the design includes a placebo of its own. For NAIC 1–2 bonds, both subsidiary types are required to carry the bond at amortized cost under their respective SSAP, so μ does not vary across siblings and the within-group selling differential should be zero. For NAIC 3–5 bonds, μ varies sharply across the two subsidiaries, and the model predicts a positive P&C-vs-Life selling differential. The split absorbs an alternative explanation that no fixed effect can address: any residual difference between Life and P&C operating cultures that survives the group fixed effect would have to manifest uniformly across rating buckets rather than appear discretely at the NAIC 2/3 boundary.

Panel B of Table 5 reports the estimates. Column (1), restricted to NAIC 3–5 bonds where the accounting rules diverge, yields a P&C coefficient of +0.0141, significant at the 10 percent level. Against a baseline Life-subsidiary selling rate of 20.7 percent, this is a 1.4-percentage-point increase, about a 7 percent higher propensity for a P&C subsidiary to sell the same bond in the same quarter than its Life sibling.[14] Column (2) runs the same specification on NAIC 1–2 bonds, where both subsidiaries use amortized cost: the coefficient is a precise zero at −0.0012. P&C and Life subsidiaries behave identically when the accounting rules treat them identically. Column (3) pools both rating buckets and adds the P&C × NAIC 3–5 interaction, which enters at +0.0149, significant at the 10 percent level; the P&C main effect is essentially zero. The differential selling behavior arises only where the accounting rules diverge.

[14]The dependent variable is constructed from quarter-over-quarter declines in reported holdings, which can reflect maturities, calls, and intra-affiliate transfers as well as market sales. Schedule D Part 4 disposal records, which name the counterparty of each disposal, let us strip these out: classifying as market sales only disposals to a named external purchaser—excluding maturities, calls, redemptions, tenders, exchanges, and transfers—over 2005–2021, the P&C–Life differential in the quarterly market-sale rate is 0.83 percentage points larger for NAIC 3–5 bonds than for NAIC 1–2 bonds (SE = 0.06 percentage points), against a baseline market-sale rate of 0.8 percent. The band contrast that identifies Panel B thus survives on transactions that verifiably reached the market. Because the Part 4 extract identifies insurer type but not the individual subsidiary, this validation operates at the type level rather than within holding company.

Identification rests on a small number of clusters (60 holding companies in the NAIC 3–5 subsample of column (1) and 94 in the pooled specification of column (3)), so we assess inference with a wild-cluster bootstrap (Rademacher weights, 9,999 replications, clustered by holding company), the appropriate procedure when clusters are few. Under the bootstrap, the interaction is significant at the 5 percent level, with a 95 percent confidence interval of [+0.0004, +0.0361] that excludes zero; the standalone NAIC 3–5 gap in column (1) stays marginal, with a bootstrap interval that just includes zero. The quantity our placebo logic targets, the NAIC 3–5-versus-1–2 *difference*, is therefore significant at the 5 percent level under inference robust to the modest cluster count, while we read the standalone magnitudes more cautiously and rely on the sharp on/off pattern across the NAIC boundary, a clean effect at NAIC 3–5 and a null at NAIC 1–2.

The pattern is not an artifact of how we assign NAIC designations: replacing the Schedule D SVO designation with an independent one (the contemporaneous Mergent FISD credit rating of the bond, which does not share input with the SVO extraction and, being measured in the same quarter, embeds no look-ahead) leaves the NAIC 3–5 point estimate essentially unchanged at +0.0152, though less precisely estimated, and the NAIC 1–2 placebo remains zero (Online Appendix OA-2.6).

A natural concern is that the NAIC 2/3 boundary triggers not only the recognition divergence but also an increase in the regulatory capital charge: a downgrade into speculative grade raises an insurer's required risk-based capital against the bond, whereas a downgrade within investment grade does not (Ellul et al., 2011). The differential selling might then reflect capital relief rather than the accounting rule, and the group × bond × quarter fixed effect does not absorb this channel: the fixed effect holds the bond's NAIC designation constant across the siblings, but the Life and P&C RBC formulas apply different bond factors to the same designation, so the charge itself varies by subsidiary type. The factor schedules, however, run against the confound. Over our sample period the Life bond factor jumps from roughly 1.3 to 4.6 percent of carrying value at the NAIC 2/3 boundary, while the P&C factor rises from roughly 1.0 to 2.0 percent, so the capital-relief motive for shedding a newly speculative-grade bond is about three times stronger for the Life subsidiary, yet it is the P&C subsidiary that sells.[15] A capital-charge story therefore predicts a within-group gap of the opposite sign. Panel C reinforces the point: weakly capitalized Life subsidiaries (the holders with both the strongest capital-relief motive and the larger factor jump) sell speculative-grade bonds *less* than their well-capitalized counterparts, the signature of the recognition shield rather than of charge minimization. What moves in the direction of the estimated gap at the boundary is the recognition treatment: amortized cost for the AVR-maintaining Life subsidiary versus the lower of amortized cost or fair value for the P&C subsidiary. The pattern accords with the recognition channel documented across insurers more broadly, where fair-value-constrained holders shed downgraded bonds while historical-cost holders retain them (Ellul et al., 2015).

[15] Pre-2021 NAIC risk-based capital bond factors for NAIC designations 1 through 5 are 0.4, 1.3, 4.6, 10.0, and 23.0 percent under the Life formula, against 0.3, 1.0, 2.0, 4.5, and 10.0 percent under the P&C formula; the 2021 revision of the Life factors post-dates our sample. Fixed-income charges are also a smaller share of total required capital for P&C insurers, whose RBC is dominated by underwriting risk, further muting the capital-relief motive on the P&C side.

The model's capital-proximity prediction refines the design one step further. The HCA–MTM divergence operates only through forced liquidation, so the within-group selling gap should be concentrated among the subsidiaries for which the capital constraint is close, and vanish where slack is abundant (Prediction 3). We test this by augmenting equation (16) with an indicator for a weakly capitalized subsidiary (a prior-year-end RBC ratio below the within-industry median for that year) and its interactions; because the Life and P&C RBC formulas are not level-comparable, capital is always classified within industry. Panel C of Table 5 reports the estimates.

In column (1), the triple interaction P&C × NAIC 3–5 × Weak enters at +0.0274, significant at the 5 percent level, while the P&C × NAIC 3–5 term, the selling gap among well-capitalized siblings, is flat (−0.0036). The entire selling gap of Panel B thus comes from the weakly capitalized siblings: a weak P&C subsidiary sells the same NAIC 3–5 bond about 2.3 percentage points more often than its Life sibling, while well-capitalized siblings of either type behave identically, as the model implies for holders far from the constraint.

The Life side moves the way the shield implies: weakly capitalized Life subsidiaries sell speculative-grade bonds *less* than their well-capitalized counterparts (−0.0199, significant at the 5 percent level), consistent with holders that can least afford to realize losses using the amortized-cost treatment to avoid doing so. Column (2) shows a directionally similar but imprecise pattern when Weak is instead defined as an RBC ratio below 300 percent of the Authorized Control Level, a zone containing few subsidiary-quarters; column (3) shows the crisis-year point estimate is identical to column (1) though the smaller sample cuts precision. Column (4) repeats the test under the contemporaneous-quarter-only SVO assignment, which retains 29 percent of the estimation sample: the base selling gap survives (+0.0040, significant at the 10 percent level) but the capital split is not identified at that coverage. We therefore read the capital heterogeneity as established under the modal assignment and suggestive elsewhere.

At the NAIC 2→3 boundary, P&C-concentrated bonds carry a fire-sale discount that subsequent returns reverse, consistent with P&C insurers recognizing mark-to-market losses and selling while Life insurers hold. Inside the same holding company, P&C subsidiaries shed NAIC 3–5 bonds faster than their Life siblings, the gap vanishes for NAIC 1–2 bonds where the rules match, and the gap is concentrated among the weakly capitalized siblings for whom the constraint the rule protects is close. No liability-duration, redemption-pressure, or institutional-form story predicts a differential that appears and disappears at the accounting-rule boundary; each would predict a uniform P&C/Life gap across rating categories. The design identifies off the holding companies that operate both a Life and a P&C subsidiary, and these groups are larger than single-type insurers. On the margin the mechanism turns on, however, they are unexceptional: their median risk-based capital ratio is indistinguishable from that of single-type Life and P&C insurers, and they hold speculative-grade bonds in the proportions a Life–P&C blend implies (Online Appendix OA-2.7). Because the recognition rule bites only through the capital constraint, and that constraint is no tighter inside these groups, the channel we identify plausibly extends to the stand-alone insurers that hold most insurer-owned debt. Having isolated μ as the operative margin, we now ask through what economic channels μ delivers stabilization.

5.5 Mechanism: Duration, Capital Slack, and Ownership Concentration

We have established that insurer-held bonds decline less in crises, that the effect is absent for mutual funds, and that within the insurance sector the differential tracks the accounting-rule boundary; what remains is to show *how* the shield operates. Our framework highlights four moderating factors. Two of these arise directly from the model's core building blocks: (1) the regulatory capital constraint that converts an unrealized loss into a compelled asset sale, and (2) the investor's recognition of its own impact on prices. These two features also set our framework apart from that of Allen and Carletti (2008) and Plantin et al. (2008). The remaining two determine how important the shield is: the magnitude of the unrealized loss, which increases with duration, and the intensity of the crisis, which determines whether the capital buffer is breached in the first place.

We take each to the data, augmenting the crisis-drawdown regression in equation (11) with an interaction between a bond's insurer ownership share and one moderator at a time, while keeping the full set of crisis × bond-characteristic fixed effects. We estimate but do not report the main effects, so each interaction is identified from bonds that are otherwise comparable. A negative coefficient means insurer ownership cushions the drawdown by more where the moderator is larger. Table 6 reports the results of the four tests.

Table 6: Mechanism Tests

	Dependent variable: Log Drawdown			
	(1) Duration	(2) Capital	(3) Severity	(4) Concentration
Ins. share × Long duration	-0.0405***			
	(0.0060)			
Ins. share × RBC < 300%		-0.0392		
		(0.0393)		
Ins. share × RBC 300–500%		-0.0128***		
		(0.0041)		
Ins. share × Severe crisis			-0.0221*	
			(0.0117)	
Ins. share × HHI				-0.0833***
				(0.0127)
Crisis episode FE	Yes	Yes	Yes	Yes
Crisis × Size decile FE	Yes	Yes	Yes	Yes
Crisis × Duration FE	Yes	Yes	Yes	Yes
Crisis × Rating FE	Yes	Yes	Yes	Yes
Crisis × Issuer FE	Yes	Yes	Yes	Yes
Crisis × Coupon/Senior FE	Yes	Yes	Yes	Yes
Observations	31,245	31,098	31,245	31,235

Notes: This table reports estimates of equation (11) augmented with an interaction between insurer share and a moderator X, on the crisis-period sample. The dependent variable is the log bond drawdown. The moderators are (1) a long-duration indicator for bonds with remaining maturity greater than five years; (2) regulatory-proximity zones of the par-weighted risk-based capital (RBC) ratio among the bond's insurer holders, measured at the pre-crisis year-end (below 300 percent of the Authorized Control Level and 300–500 percent, with above 500 percent omitted); (3) an indicator for the severe crisis episodes (the Great Recession and COVID-19); and (4) the bond-level Herfindahl index of insurer holdings. All specifications include the fully saturated crisis × bond-characteristic fixed effects; main effects are estimated but not reported. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Duration: HCA's stabilization benefit should rise with duration as a reduced-form matter: longer-maturity bonds lose more to a given shock under mark-to-market accounting, widening the gap between amortized cost and fair value. Column (1) interacts insurer share with a long-duration indicator and yields -0.0405^{***}: insurer stabilization is nearly four times stronger for long-duration bonds. A patient-capital story would predict uniform stabilization across maturities, and a low-volatility-selection story would predict the opposite (stronger stabilization at short durations). Of these candidates, only the accounting mechanism fits this pattern.

Capital proximity: The model ties the shield to how close insurers are to their regulatory capital constraint. When they have substantial slack, forced sales never bind, so the accounting treatment of unrealized losses is irrelevant; when insurers are near the constraint, that treatment determines whether the loss feeds through to reported capital. Column (2) evaluates this by sorting bonds according to the par-weighted RBC ratio of their insurer owners at the pre-crisis year-end and interacting insurer share with indicators for different regulatory-proximity bins (with ratios above 500 percent of the Authorized Control Level omitted as the baseline group). For bonds whose ownership is concentrated among insurers with RBC ratios in the 300–500 percent range, the interaction coefficient is -0.0128^{***}, indicating a stabilization slope roughly four times steeper than for bonds held primarily by insurers with ample slack. The sample below the 300 percent threshold is too small (570 bond-quarters) to estimate precisely with a coefficient of -0.0392. The point estimate in this thin below-300 zone is even larger in absolute value, consistent with the model's monotonic prediction, but is too imprecise to be statistically distinguished from the 300–500 zone. This is the key restriction implied by the model and corresponds to the same capital-proximity margin along which the within-group selling gap is concentrated (Table 5, Panel C).[16]

Severity: Column (3) interacts the insurer share with an indicator for crises severity. The coefficient is -0.0221^{*}, implying that the shield binds more tightly when unrealized losses are largest, exactly when the model says it should. Only a shock large enough to push prices toward the regulatory threshold breaches the buffer and forces the marked-to-market holders to sell.

Concentration: Column (4) interacts the insurer share with the bond-level Herfindahl index (HHI) of insurer holdings. A concentrated holder takes into account the effect its own trades have on prices, exactly as the model assumes because selling into an illiquid market pushes down the value of the position it still holds. The interaction term Ins. share $\times$ HHI is negative and statistically significant at -0.0833^{***}, indicating that more concentrated investors act as stronger stabilizers. This aligns with the RBC result that stabilization requires both the ability to hold capital and the incentive to refrain from selling, which arises from concentrated ownership.

5.6 Falsification and Scope Tests

The accounting channel operates only when two conditions hold: insurer holders with economically material bond positions under the amortized-cost rule, and the crisis capital pressure that makes the rule bind. The next two tests remove each condition in turn. In both, the accounting channel predicts no effect, while insurer-patience and low-volatility-selection stories still predict one.

Health insurers: Under statutory accounting, health insurers, like all other insurers, report investment-grade bonds at amortized cost, as non-AVR filers, they apply a lower-of-cost-or-fair-value rule to speculative-grade holdings, as do P&C insurers.[17] For the investment-grade bonds that make up the bulk of insurer portfolios, a generic insurance-based argument would thus suggest that Health share should stabilize bond prices in the same way that Life share does. However, Health

[16] A linear interaction of insurer share with the par-weighted RBC ratio is zero (-0.0004), consistent with the zone pattern. The relevant variation is proximity to the constraint, not capital differences among comfortably capitalized holders.

[17] The health test therefore serves as a scope check for the investment-grade segment, where all insurer types use amortized-cost accounting and where their bond portfolios are concentrated; it does not address the NAIC 3–5 divergence region.

insurers' bond portfolios are tiny relative to their overall balance sheets: the accounting treatment is in place, but the exposures are not. We re-estimate the baseline regression, decomposing the pooled insurer share into separate Life, P&C, and Health shares (Table 7). When included alone, Health share has a coefficient of +0.0101 a noisy null whose standard error is nearly seven times the Life estimate (-0.0139^{***}): given such small holdings, the test has little power to detect any stabilizing effect; the P&C share coefficient is +0.0138 and also insignificant. These null results persist when all three shares enter the regression together (column 4). The accounting rule by itself is insufficient; the mechanism requires bond positions that are large enough to matter economically.

Table 7: Type Decomposition: Life vs. P&C vs. Health

	Dependent variable: Log Drawdown				
	(1)	(2)	(3)	(4)	(5)
Life ins. share	-0.0139***			-0.0137***	-0.0136***
	(0.0038)			(0.0039)	(0.0038)
P&C ins. share		0.0138		0.0077	0.0070
		(0.0117)		(0.0112)	(0.0118)
Health ins. share			0.0101	-0.0293	
			(0.0924)	(0.0917)	
Crisis episode FE	Yes	Yes	Yes	Yes	Yes
Crisis × Size decile FE	Yes	Yes	Yes	Yes	Yes
Crisis × Duration FE	Yes	Yes	Yes	Yes	Yes
Crisis × Rating FE	Yes	Yes	Yes	Yes	Yes
Crisis × Issuer FE	Yes	Yes	Yes	Yes	Yes
Crisis × Coupon/Senior FE	Yes	Yes	Yes	Yes	Yes
Observations	31,245	31,245	31,245	31,245	31,245

Notes: This table decomposes the pooled insurer-share effect across Life, P&C, and Health insurers, on the crisis-period sample. The dependent variable is the log bond drawdown. Columns (1)–(3) enter the Life, P&C, and Health ownership shares separately; column (4) enters all three jointly; and column (5) enters the Life and P&C shares jointly. All columns include the fully saturated crisis × bond-characteristic fixed effects. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Non-crisis quarters: Outside crises there are no material unrealized losses and no binding capital pressure: book values and market prices track each other and no insurer faces forced selling to shore up regulatory surplus, so HCA has nothing to shield against. The model is explicit on this point: outside the region where the capital constraint binds, the accounting regime is immaterial and insurer ownership should carry no predictive power. If insurer share predicted returns in calm quarters, the main result would reflect permanent low-volatility selection rather than the accounting channel. In progressive-FE specifications on non-crisis quarters (349,000–405,000 bond-quarter observations), insurer share enters negatively and significantly under coarse controls, but the coefficient collapses to +0.0001, indistinguishable from zero, once rating and issuer fixed effects enter. Selection on observables exists, but the baseline fixed-effect structure absorbs it.[18]

No single alternative explanation proposed, whether patient capital, liability-driven investing, a preference for low volatility, or institutional inertia, can simultaneously explain both null findings and remains consistent with the placebo built into the within-group design of §5.4. The accounting mechanism appears in every context where the model predicts an effect and is absent from all contexts where the model predicts no effect. The remaining issue is whether the stability of the bond-price is transmitted to the real economy.

5.7 Real Effects on Firm Investment

The stabilization channel operates through lower borrowing costs in the crisis-period. If HCA dampens bond-price declines, firms whose debt is held on insurer balance sheets should be better positioned to invest during and after crises. Proposition 3.2 provides the first link, accounting discretion under HCA prevents forced liquidations and stabilizes crisis prices and the remaining links follow standard arguments: stable prices imply lower and more predictable borrowing costs, and borrowing costs shape investment (Campello et al., 2010; Chava and Roberts, 2008). Firms with higher insurer ownership should therefore sustain or expand capital expenditures and acquisitions after downturns (Ivashina and Scharfstein, 2010; Gilchrist et al., 2013). Coppola (2025) also trace the composition of the investors through to real results at firm-level during the credit cycle; our contribution here is to link the investment response specifically to the *accounting* treatment of the holder rather than to the type of investor, the same recognition-rule channel identified in the bond-price tests above.

To test this, we track firm-level CAPEX and acquisition activity in the years after the Great Recession, the severe crisis for which a long post-period is observable. Table 8 reports a difference-in-differences specification comparing firms with higher insurer bond ownership to those with lower ownership, before and after the crisis, where the coefficient of interest interacts with the insurer share with a post-crisis indicator. We add firm-level controls for size, total liabilities, and bond characteristics, each interacted with the post indicator, to absorb the confounding variation that might otherwise masquerade as an ownership effect.

Column 1 shows that a larger share of insurers forecasts higher post-crisis capital expenditure: a 10%-point increase in insurer ownership is linked to a 0.14-percentage-point statistically significant rise in the CAPEX-to-assets ratio. Controlling for pre-period firm size, total liabilities, and bond duration, each interacted with the post dummy, barely changes the estimate, which remains 0.0131** (column 2). To gauge the magnitude, scale this by the cross-sectional dispersion in ownership:

[18] Prediction 6 that HCA cushions declines without impeding recoveries we do not test directly: our outcome is defined over crisis drawdowns, so recovery-period behavior lies outside the estimation windows. The non-crisis null here is its closest empirical counterpart.

a one-standard-deviation increase in insurer share, about 26 percentage points, corresponds to roughly a 0.34-percentage-point increase in post-crisis CAPEX-to-assets, or around 6 percent of the sample's mean investment rate of 5.5 percent. This pattern aligns with the model's mechanism that firms whose bonds are held on insurer balance sheets experience more stable debt prices because HCA shields their creditors from market-price swings, consistent with price stability translating into lower effective borrowing costs and higher investment. The associated financing margin appears in the issuance data, during the 2008–09 freeze, the only period in our sample when primary markets shut down, a 10-percentage-point higher pre-crisis insurer share is linked to a 2.4-percentage-point greater likelihood of issuing a bond, about one-tenth of the freeze-year issuance rate, with no comparable differential once markets reopened in 2010–11 (Online Appendix Table OA-11).

Table 8: Real Effects: Firm Investment and Total Insurer Ownership

	CAPEX			*CAPEX + Acquisitions*		
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Ins. share (OLS)	0.0144***	0.0131**		0.0631***	0.0645***	
	(0.0054)	(0.0054)		(0.0127)	(0.0140)	
Post × Ins. share (IV)			0.5049**			1.1529
			(0.2538)			(0.7209)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Post × Size/Lev/Dur FE	No	Yes	Yes	No	Yes	Yes
Estimator	OLS	OLS	IV	OLS	OLS	IV
First-stage F			3.4			2.7
Anderson–Rubin p			0.001			<0.001
Observations	5,876	5,876	5,876	5,527	5,527	5,527

Notes: This table reports difference-in-differences estimates of firm investment around the Great Recession. The dependent variable is annual CAPEX (columns (1)–(3)) and CAPEX plus acquisitions (columns (4)–(6)). Each column enters total insurer share interacted with the post-crisis indicator; columns (3) and (6) instrument it with the firm-level analogue of the shift-share instrument Z_i (equation (14)) and, because that first stage is weak, report weak-instrument-robust Anderson–Rubin p-values, which reject the null of no effect at the 1% level. All columns include firm and year fixed effects and pre-period size, total liabilities, and duration interacted with the post indicator; columns (1) and (4) omit the latter. Standard errors, clustered at the firm level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

In Columns 4 and 5 we extend the outcome to total investment (CAPEX plus acquisitions). The effect is larger and more precisely estimated: a 10 percentage-point increase in insurer share is associated with total investment-to-assets higher by 0.63–0.65 percentage points, significant at the 1 percent level, about $165 million of additional investment per year at the mean sample firm, whose total assets average $26.1 billion (the size distribution is right-skewed, so the corresponding dollar figure at the median firm is smaller). Scaled to a one-standard-deviation move in ownership, that is about 1.6 percentage points of additional investment, roughly 20 percent of the 8.1 percent sample-mean total-investment rate, and a markedly larger response than for capital expenditures alone. The firms whose creditors were shielded from fire-sale pricing expanded, through both organic growth and acquisitions.

In Columns 3 and 6, we instrument firm-level insurer ownership using the firm-level counterpart of the shift-share instrument. The resulting IV coefficients are larger than the OLS benchmarks (0.505 for capital expenditures and 1.153 for total investment), but the firm-level first stage is weak (first-stage F-statistics of 3.4 and 2.7). This weakness arises because the quasi-random variation at the primary-market level is partially averaged out when bond-level allocations are aggregated to the issuer level. Since weak instruments bias both the 2SLS point estimates and their conventional standard errors, we conduct inference using the weak-instrument-robust Anderson–Rubin test (Anderson and Rubin, 1949; Andrews et al., 2019). The Anderson–Rubin test rejects the null hypothesis of no effect at the 1 percent level for both investment outcomes. However, because the first stage is weak, the corresponding confidence set is wide and does not place an upper bound on the size of the effect. Consequently, we rely on the OLS estimates to gauge the magnitude of the response and interpret the IV results as supporting evidence that the investment effect is statistically different from zero under inference procedures that do not require a strong first stage.

Figure 10 traces the dynamics underlying the average effects. The event-study graphs show, year by year, the insurer-share coefficients on CAPEX and on CAPEX plus acquisitions, normalized to 2007, for the sample of firms with bonds outstanding in 2007Q1. The pre-crisis pattern supports the identification strategy given that from 2004 to 2006, both coefficients hover around zero and their confidence intervals include zero, indicating that firms with greater insurer ownership were not already following a different investment path before the crisis hit. The divergence appears with the onset of the crisis: the total-investment coefficient climbs steeply in 2008, reaches a peak of about 0.08 in 2009, and remains elevated through 2011; the CAPEX coefficient follows a comparable pattern, though with smaller values. The response therefore persists rather than simply appearing in a single period and it is consistently stronger for overall investment than for CAPEX alone. This pattern aligns with insurer-protected firms accessing cheaper and more dependable credit to fund both organic growth and acquisitions throughout the recovery period.

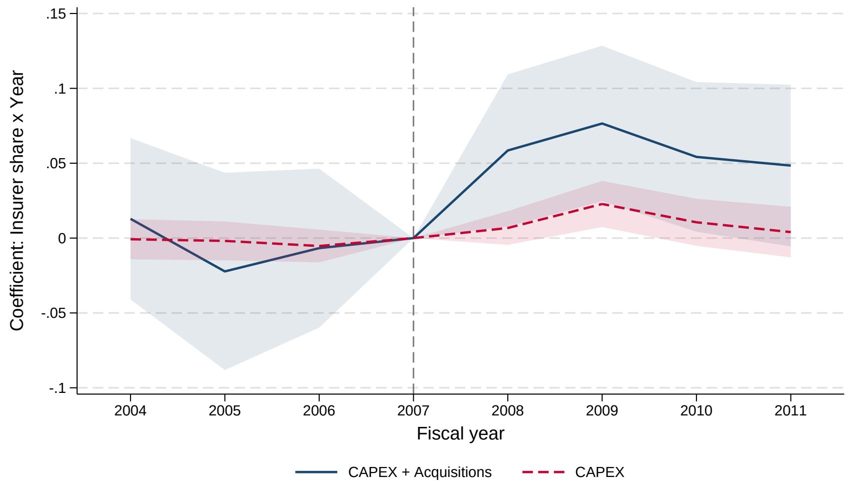

Figure 10: Investment Event Study

Notes: This figure displays the estimated coefficients derived from the difference-in-differences analysis reported in Table 8. The analysis is based on a dataset comprising U.S. firms that had outstanding bonds as of the first quarter of 2007 (2007Q1). These coefficients trace the dynamic relationship between the investment rates of firms and the proportion of their bonds held by domestic insurers in 2007Q1. The model includes firm-level controls, measured in the pre-period and interacted with time dummies: deciles of average bond duration, firm size, and total liabilities. The red line represents the estimates for capital expenditures (CAPX) relative to lagged total assets, whereas the blue line illustrates the estimates for the sum of CAPX and acquisitions relative to lagged total assets. Standard errors are clustered at the firm level, and the bars indicate 95 percent confidence intervals.

Table 9: Real Effects: Decomposition by Holder Accounting Type

	CAPEX	*CAPEX + Acquisitions*
	(1)	(2)
Post × HCA (Life) share	0.0163***	0.0849***
	(0.0061)	(0.0146)
Post × MTM (P&C) share	-0.0349	-0.0674
	(0.0311)	(0.0996)
p(HCA=MTM), wild cluster	0.147	0.182
Firm FE	Yes	Yes
Year FE	Yes	Yes
Post × Size/Lev/Dur FE	Yes	Yes
Observations	5,876	5,527

Notes: This table decomposes the total insurer-share effect on firm investment by holder accounting type. The dependent variable is annual CAPEX (column (1)) and CAPEX plus acquisitions (column (2)). Total insurer share is split into the share held by HCA-reporting (Life, AVR) subsidiaries and the share held by MTM-reporting (P&C, non-AVR) subsidiaries, aggregating across NAIC tiers, and each is entered interacted with the post-crisis indicator. The reported p-value tests equality of the HCA and MTM coefficients by the wild-cluster bootstrap (Rademacher weights, 9,999 replications, clustered on the firm). Table OA-10 in the Online Appendix sharpens this decomposition by further interacting each holder share with the bonds' NAIC tier: NAIC 3–5, where statutory treatment diverges across holder types, versus NAIC 1–2, where it agrees. All columns include firm and year fixed effects and pre-period size, total liabilities, and duration interacted with the post indicator. Standard errors, clustered at the firm level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively, based on analytic standard errors.

Table 9 examines which types of insurers drive the Table 8 findings by decomposing total insurer ownership into the fraction held by HCA-reporting (Life) subsidiaries and the fraction held by MTM-reporting (P&C) subsidiaries, each interacted with the post indicator. The results align with our accounting mechanism: the investment responds to the HCA share, not to the MTM share. A 10 percentage-point rise in HCA (Life) ownership is linked to a post-crisis CAPEX-to-assets ratio that is about 0.16 percentage points higher and a total investment-to-assets ratio that is about 0.85 percentage points higher, both significant at the 1 percent level. The latter effect is on the order of one-tenth of the sample-average investment rate. In contrast, the MTM (P&C) coefficients are negative and statistically indistinguishable from zero for both CAPEX (−0.0349) and total investment (−0.0674). Thus, the entire real-effects response is driven by the holders subject to more lenient recognition rules.[19] In line with this interpretation, the HCA coefficients exceed the total-ownership estimates in Table 8 (0.0131 and 0.0645). This is what we would expect if economically active ownership is concentrated in the HCA share, so that pooling it with the

[19] Although the point estimates differ economically, the pooled-band equality test is underpowered: the wild-cluster bootstrap does not reject HCA = MTM at conventional levels because P&C mostly holds investment-grade bonds, where the recognition rule does not bind, so pooling NAIC tiers combines the divergence band that identifies the channel with a placebo band that does not. The Online Appendix (Table OA-10) restores power by interacting each holder share with the bond's NAIC tier: in the speculative-grade NAIC 3–5 band, where statutory treatment diverges, the HCA and MTM coefficients differ at the 5 percent level for CAPEX and the 10 percent level for total investment, but in the investment-grade NAIC 1–2 placebo band, where the rule treats Life and P&C identically, they do not.

relatively passive MTM share dampens the effect measured on total shares.

The investment response thus matches the particular holder category that the rule distinguishes, and within that category it concentrates exactly at the accounting cutoff. Ownership-composition channels based on either total insurer positions or on overall liability structure cannot generate this pattern, because neither of those margins changes at the NAIC 2/3 breakpoint. Since these decompositions rely on OLS and firms' holder mixes are not randomly assigned, the credible identification of the mechanism comes from the within-group design in Table 5: group × bond × quarter fixed effects that contrast Life and P&C subsidiaries of the same parent holding the same bond. It does *not* rely on a shift-share instrument, because the firm-level instrument is informative about aggregate insurer ownership but has no explanatory power for the Life-versus-P&C split required for this decomposition. Figure 11 illustrates the on/off pattern by plotting the HCA − MTM investment differential within each rating band. In the speculative-grade NAIC 3–5 range, where the recognition rule differs across holder types, this gap is sizable and its analytic 90 percent confidence interval excludes zero for both CAPEX and total investment; in contrast, in the investment-grade NAIC 1–2 placebo range, where Life and P&C face identical treatment, the gap contracts toward zero and its confidence interval easily covers it. The figure reports only the difference and its analytic interval; the underlying four band-by-holder coefficients and the associated wild-cluster tests of equality are presented in Table OA-10 of the Online Appendix. Taken together, the two bands provide the sharpest illustration of the channel: the investment gap emerges precisely where the accounting rule diverges across holder types and disappears where the rule treats them the same.

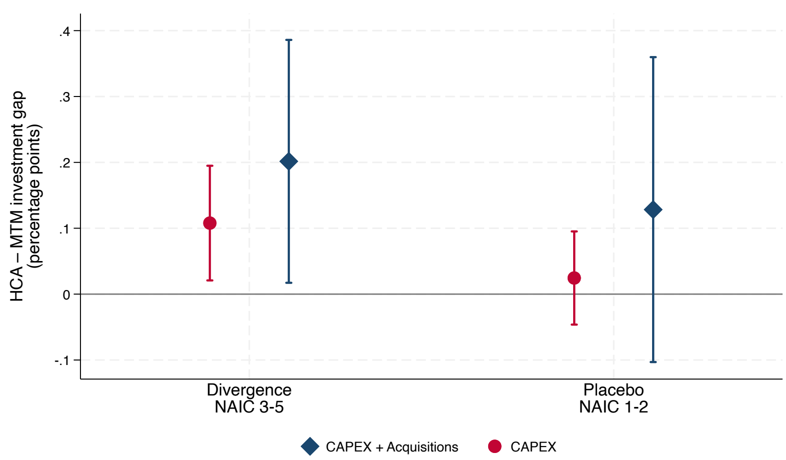

Figure 11: Real Effects by Holder Type and Rating Band

Notes: This figure plots, within each NAIC rating band, the difference between the HCA-reporting (Life) and MTM-reporting (P&C) Post × insurer-share coefficients, estimated from the band-level decomposition of Table OA-10 (firm and year fixed effects; pre-period size, total liabilities, and duration × post). The divergence band (NAIC 3–5) is where statutory treatment differs across holder types; the placebo band (NAIC 1–2) is where it agrees. Cranberry circles are CAPEX and navy diamonds are CAPEX + acquisitions; capped spikes are analytic 90 percent confidence intervals from firm-clustered standard errors. The HCA − MTM gap is large in the divergence band and collapses toward zero in the placebo band.

Accounting discretion is thus tied to real consequences: where HCA holders dampen crisis-period bond-price declines and limit the accompanying rise in borrowing costs, the issuing firms invest more heavily in capital expenditure and acquisitions during and after downturns. Firms whose debt is held on insurer balance sheets exit crises in a stronger financial position, and the reach of amortized-cost accounting extends beyond the bond market into firm-level growth and strategic decisions.

6 Conclusion

Accounting rules determine what an institution's own balance sheet reports, and in crises that reported position shapes the institution's own selling decisions as much as it informs outside observers. A Life insurer holding a downgraded bond at amortized cost and a P&C insurer holding the same bond under a lower-of-cost-or-market rule face the same market price, the same issuer, and the same regulator, but not the same recognition rule, and that difference alone is enough to generate different selling behavior in the quarters that matter most for price stability. We take that observation to the corporate bond market and ask whether the stabilization of insurer-held bonds in crises reflects the accounting rule applied to the bond or the type of institution that happens to hold it. Our results indicate that the rule, rather than the institution, generates the stability.

Our evidence comes from tests that progressively rule out alternatives explanations. Bonds more heavily held by insurers fall less in value during crises than bonds more heavily held by mutual funds, and a shift-share instrument shows this is caused by insurer ownership itself, not by insurers picking safer bonds. Within-sector tests then ask whether the stabilizing effect comes from the accounting rule or other insurer traits. A horse-race specification ties stabilization specifically to insurer ownership, not institutional ownership in general. Within insurance, bonds downgraded across the NAIC 2/3 cutoff, where P&C insurers switch to fair-value recognition but Life insurers do not show a sharp split in post-downgrade returns, while bonds downgraded across the NAIC 1/2 cutoff, where the rule is unchanged, show no difference. In addition, a Life and a P&C subsidiary in the same holding company, holding the same bond in the same quarter, behave differently for the NAIC 3–5 bonds and identically for the NAIC 1–2. The fixed-effect design absorbs the parent's investment philosophy, liability structure, and risk tolerance, but not the statutory rule.

Our mechanism tests support this channel: stabilization rises with bond duration, is concentrated in the two major crises, operates through insurer regulatory capital, and strengthens with more concentrated ownership. Three predicted null results also hold: Health insurers that apply the same accounting standard but do not in fact maintain large bond portfolios, and thus show no impact; the fraction of bonds owned by insurers does not forecast outcomes in calm times, when book and market values coincide and capital is not under pressure; and differences in sales behavior within the same holding company disappear exactly at the NAIC 1/2 threshold, where the regulatory rule does not differ.

Our contribution is to demonstrate that the accounting recognition rule is a key determinant of a pattern that prior work has attributed either to liability structure (Coppola, 2025) or to institutional fragility (Chen, Goldstein, and Jiang, 2010; Goldstein, Jiang, and Ng, 2017). We do not dispute that those mechanisms exist; rather, our findings identify an additional channel that works in parallel with them and that, in contrast to liability structure, can be directly modified via the design of recognition rules. The within-holding-company evidence implements a comparison that the insurance-markets literature has so far run only across companies, contrasting Life and P&C insurers without holding group-level attributes fixed (Ellul, Jotikasthira, Lundblad, and Wang, 2015; Khan, Ryan, and Varma, 2019), and combines it with the NAIC 1–2 falsification built into the same design to rule out residual differences in subsidiary-level risk preferences.

Our theoretical framework extends the accounting-based fire-sale models of Allen and Carletti (2008) and Plantin et al. (2008) along three dimensions. An exogenous fundamental-value shock lets sales be privately efficient rather than only regulator-forced. Insurers that internalize their price impact generate a no-sale region under HCA that atomistic settings do not. And the HCA and MTM sectors coexist with a tunable composition, so the precautionary buffer and the probability of forced liquidation become functions of the bond's ownership mix, making stabilization continuous in the HCA share rather than binary, and yielding a comparative static that maps directly onto observable bond-level ownership data. The cross-sectional patterns we document (namely the level and composition of insurer ownership, along with duration, capital, severity, and concentration) are difficult for the alternative theories of insurer behavior we consider to simultaneously match. And the real-effects evidence extends the accounting channel through the credit market to firm-level capital expenditures and acquisitions, completing the link from recognition rule to investment.

The bigger implication of our results is that accounting design is not a neutral exercise in information disclosure. In normal periods, historical cost accounting (HCA) generates the selective-realization incentives documented in prior research (Ellul, Jotikasthira, Lundblad, and Wang, 2015; Badertscher, Burks, and Easton, 2012). However, in crises, the same discretion over recognition helps stabilize prices and propagates through credit markets into investment decisions of firms. The welfare effects of HCA are therefore contingent on the state of the world, and its advantages in corporate bond markets are greatest when fire-sale risk is most severe. This state dependence also highlights what our paper does not resolve.

Our empirical strategy exploits a cutoff created by insurance regulation, but the core concept it captures, the fraction of an asset's investors that report it at amortized cost, arises in any setting where multiple accounting regimes coexist. This is especially relevant in banking, where securities are classified into held-to-maturity or available-for-sale portfolios (Granja et al., 2024). The model's key comparative static, that stabilization increases with the fraction of HCA owners, is testable in all such environments, and extending the analysis to them is a natural direction for future work. A more difficult issue is rule design, specifically whether recognition standards themselves should vary with market conditions (i.e., becoming tighter in normal periods, when HCA primarily induces selective realization and postpones loss recognition, and more permissive in crises, when its main benefit is the buffering mechanism we document) and whether such a regime could be implemented without introducing new types of discretion. Our estimates indicate that in the corporate bond market we examine, the recognition rule functions as a policy instrument that influences bond prices and investment. They do not, however, pin down how this instrument should be calibrated over the cycle, so additional research on this dimension of design is needed.

Generative AI Assistance Disclosure

The authors used Grammarly, ChatGPT (GPT-5), and Cursor AI to assist with copy editing, cleaning, and organizing coding scripts. The authors reviewed, verified, and revised all generated material and take full responsibility for the content of this manuscript.

References

Acharya, V., I. Drechsler, and P. Schnabl (2014). A pyrrhic victory? bank bailouts and sovereign credit risk. *The Journal of Finance 69*(6), 2689–2739.

Allen, F. and E. Carletti (2008). Mark-to-market accounting and liquidity pricing. *Journal of accounting and economics 45*(2-3), 358–378.

Anderson, T. W. and H. Rubin (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics 20*(1), 46–63.

Andrews, I., J. H. Stock, and L. Sun (2019). Weak instruments in instrumental variables regression: Theory and practice. *Annual Review of Economics 11*, 727–753.

Badertscher, B. A., J. J. Burks, and P. D. Easton (2012). A convenient scapegoat: Fair value accounting by commercial banks during the financial crisis. *The Accounting Review 87*(1), 59–90.

Barrios, J. M., B. C. Fujiy, P. Lisowsky, and M. Minnis (2026). Measurement matters: Financial reporting and productivity. *Working paper*.

Barth, M. E. (1994). Fair value accounting: Evidence from investment securities and the market valuation of banks. *The Accounting Review 69*(1), 1–25.

Bartik, T. J. (1991). Who benefits from state and local economic development policies?

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

Becker, B., M. M. Opp, and F. Saidi (2022). Regulatory forbearance in the us insurance industry: The effects of removing capital requirements for an asset class. *The Review of Financial Studies 35*(12), 5438–5482.

Bischof, J., U. Brüggemann, and H. Daske (2023). Asset reclassifications and bank recapitalization during the financial crisis. *Management Science 69*(1), 75–100.

Bischof, J., C. Laux, and C. Leuz (2021). Accounting for financial stability: Bank disclosure and loss recognition in the financial crisis. *Journal of Financial Economics 141*(3), 1188–1217.

Breuer, M. (2022). Bartik instruments: An applied introduction. *Journal of Financial Reporting 7*(1), 49–67.

Campello, M., J. R. Graham, and C. R. Harvey (2010). The real effects of financial constraints: Evidence from a financial crisis. *Journal of Financial Economics 97*(3), 470–487.

Chava, S. and M. R. Roberts (2008). How does financing impact investment? the role of debt covenants. *Journal of Finance 63*(5), 2085–2121.

Chen, Q., I. Goldstein, and W. Jiang (2010). Payoff complementarities and financial fragility: Evidence from mutual fund outflows. *Journal of Financial Economics 97*(2), 239–262.

Chircop, J. and Z. Novotny-Farkas (2016). The economic consequences of extending the use of fair value accounting in regulatory capital calculations. *Journal of Accounting and Economics 62*(2-3), 183–203.

Coppola, A. (2025). In safe hands: The financial and real impact of investor composition over the credit cycle. *The Review of Financial Studies 38*(8), 2275–2325.

Diamond, W., L. Falasconi, and C. C. Xu (2024). Collateralized loan obligations as fire-sale insulation. *SSRN*.

Dick-Nielsen, J. (2014). How to clean enhanced trace data. *Available at SSRN 2337908*.

Ellul, A., C. Jotikasthira, C. Lundblad, and Y. Wang (2015). Is historical cost accounting a panacea? market stress, incentive distortions, and gains trading. *Journal of Finance 70*(6), 2489–2538.

Ellul, A., C. Jotikasthira, and C. T. Lundblad (2011). Regulatory pressure and fire sales in the corporate bond market. *Journal of Financial Economics 101*(3), 596–620.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2013). Misallocation and financial market frictions: Some direct evidence from the dispersion in borrowing costs. *Review of Economic Dynamics 16*(1), 159–176.

Gilchrist, S. and E. Zakrajšek (2012). Credit spreads and business cycle fluctuations. *American economic review 102*(4), 1692–1720.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

Goldstein, I., H. Jiang, and D. T. Ng (2017). Investor flows and fragility in corporate bond funds. *Journal of Financial Economics 126*(3), 592–613.

Granja, J., E. X. Jiang, G. Matvos, T. Piskorski, and A. Seru (2024). Book value risk management of banks: Limited hedging, htm accounting, and rising interest rates. *Working paper, National Bureau of Economic Research*.

Haddad, V., A. Moreira, and T. Muir (2021). When selling becomes viral: Disruptions in debt markets in the covid-19 crisis and the fed's response. *The Review of Financial Studies 34*(11), 5309–5351.

Ivashina, V. and D. Scharfstein (2010). Bank lending during the financial crisis of 2008. *Journal of Financial Economics 97*(3), 319–338.

Khan, U., S. G. Ryan, and A. Varma (2019). Fair value versus amortized cost measurement and the timeliness of other-than-temporary impairments: Evidence from the insurance industry. *The Accounting Review 94*(6), 285–307.

Kim, S., S. Kim, and S. G. Ryan (2019). Economic consequences of the aoci filter removal for advanced approaches banks. *The Accounting Review 94*(6), 309–335.

Laux, C. and C. Leuz (2009). The crisis of fair-value accounting: Making sense of the recent debate. *Accounting, organizations and society 34*(6-7), 826–834.

Laux, C. and C. Leuz (2010). Did fair-value accounting contribute to the financial crisis? *Journal of economic perspectives 24*(1), 93–118.

Ma, Y., K. Xiao, and Y. Zeng (2022). Mutual fund liquidity transformation and reverse flight to liquidity. *The Review of Financial Studies 35*(10), 4674–4711.

Merrill, C. B., T. D. Nadauld, R. M. Stulz, and S. M. Sherlund (2021). Were there fire sales in the rmbs market? *Journal of Monetary Economics 122*, 17–37.

Plantin, G., H. Sapra, and H. S. Shin (2008). Marking-to-market: panacea or pandora's box? *Journal of accounting research 46*(2), 435–460.

Sapra, H. (2008). Do accounting measurement regimes matter? a discussion of mark-to-market accounting and liquidity pricing. *Journal of Accounting and Economics 45*(2-3), 379–387.

Online Appendix

Accounting Under Pressure: Recognition Rules, Insurer Bond Sales, and Real Investment

OA-1 Accounting Treatment for Mutual Funds

Unlike insurers, mutual funds report under U.S. GAAP and are required to mark all portfolio holdings to market. This means that all unrealized gains and losses on bond investments are reflected directly in the fund's net asset value (NAV), which updates continuously with market prices. As a result, mutual funds cannot use HCA to insulate their financial reports from transitory market fluctuations.

This MTM requirement makes mutual funds particularly sensitive to changes in bond valuations during crises. Any deterioration in asset prices immediately translates into a lower NAV, potentially triggering investor redemptions. These redemptions, in turn, force funds to liquidate assets, often in declining markets, amplifying price declines and creating a fire-sale feedback loop. This behavior parallels the GAAP treatment of available-for-sale (AFS) securities, where valuation volatility flows through financial statements. Kim et al. (2019) show that removing the AOCI filter, which buffered capital from AFS volatility, led banks to shift away from AFS classifications and into HCA. This shift allowed them to suppress valuation swings and reduce capital volatility. Unlike banks, mutual funds do not have access to this accounting discretion. Their inability to reclassify assets into amortized-cost treatment limits their flexibility and increases procyclicality, especially during downturns when redemption risk is high.

Under the fully fair value accounting framework mandated by U.S. GAAP, mutual funds must mark all bond holdings to market, leaving little room for managerial discretion in reporting asset values. This strict mark-to-market regime improves transparency and comparability between funds, but it also amplifies earnings and NAV volatility, especially during periods of market stress (Kim et al., 2019; Laux and Leuz, 2009). These structural differences in accounting rules offer a natural source of cross-institutional variation in crisis-period trading behavior, which motivates our empirical strategy.

OA-2 Within-Group Sample Construction

This appendix documents the construction of the within-insurance-group sample used in Panel B of Table 5 of the main text.

OA-2.1 Sample Universe

We identify mixed insurance holding companies (groups that own at least one Life subsidiary *and* at least one P&C subsidiary) using NAIC Group Codes from S&P Global's statutory-filings database. The initial universe contains 107 mixed holding companies spanning the 2007–2020 sample period.

OA-2.2 Bond Overlap and NAIC Designation Mapping

For each group-quarter, we identify bonds (CUSIPs) held simultaneously by both a Life and a P&C subsidiary within the same parent group. The overlap sample contains approximately 2.45 million subsidiary-bond-quarter observations. We map bond CUSIPs to NAIC designations using two complementary sources.

Source 1: Mergent FISD ratings. We use bond-level S&P and Moody's ratings from Mergent FISD, converted to NAIC designations following the official SVO crosswalk (NAIC 1 = AAA to A−, NAIC 2 = BBB+ to BBB−, NAIC 3 = BB+ to BB−, NAIC 4 = B+ to B−, NAIC 5 = CCC+ to CCC−, NAIC 6 = CC and below or in or near default). The NAIC 2/3 boundary on which our design rests is therefore the BBB−/BB+ (investment-grade/speculative-grade) line, and the NAIC 1/2 boundary is the A−/BBB+ line. This coverage averages 37 percent of overlapping bond-quarter observations.

Source 2: Schedule D SVO designations. We extract NAIC designations directly from Schedule D annual statements as reported by the Securities Valuation Office. This source covers 193 annual Excel files containing approximately 14 million holding records. Matching Schedule D CUSIPs to the overlap sample via a three-layer merge—(i) entity × CUSIP modal designation, (ii) CUSIP × insurer-type modal designation, (iii) CUSIP across any reporting entity—raises coverage to 58 percent of bond-quarter observations.

OA-2.3 Regression Sample

Of the approximately 2.45 million subsidiary-bond-quarter observations in the overlap sample, 58 percent carry a valid Schedule D SVO designation (Source 2 above), or roughly 1.42 million observations. After excluding five holding-company groups flagged during manual audit for data-quality issues and applying the group × bond × quarter fixed-effect restrictions (which drop singleton cells), the final regression samples are:

- 92,159 subsidiary-bond-quarter observations in the NAIC 3–5 sample used in column (1) of Panel B.
- 1.32 million observations in the NAIC 1–2 placebo sample used in column (2).
- 1.43 million observations in the pooled sample (NAIC 1–5) used in column (3).

The two single-tier samples sum to 1.41 million, consistent with the pooled sample of 1.43 million and with the 58 percent SVO-designation coverage of the 2.45-million overlap.

OA-2.4 Stickiness of Bond Holdings

Bond portfolios are sticky: 77 percent of subsidiary-bond-quarter observations show zero change in holdings between $t-1$ and t. Most of the variation in the dependent variable is concentrated on the extensive margin of whether any selling occurs at all. This stickiness is a feature of the institutional setting (insurers are regulatorily encouraged to hold bonds to maturity) rather than a deficiency of the research design. Although the stickiness limits statistical power, it does not weaken identification: the group × bond × quarter fixed effect is among the most demanding specifications in empirical asset pricing.

OA-2.5 Robustness to Group Exclusions

We assess sensitivity to the exclusion of individual holding companies by sequentially dropping each of the 107 groups and re-estimating Panel B specifications. The NAIC 3–5 coefficient is stable across all leave-one-out iterations, ranging from +0.0121 to +0.0168. The NAIC 1–2 placebo coefficient remains statistically indistinguishable from zero in every iteration. No single holding company drives the result.

OA-2.6 Robustness to the Designation Source

Because the Schedule D SVO designations are extracted from annual statements and carried across quarters by modal assignment (Online Appendix OA-2), classification error in the extraction or the modal merge could in principle generate the NAIC 3–5 selling gap. We address this directly by re-estimating Panel B with an *independent* source for each bond's NAIC tier: the bond's contemporaneous S&P/Moody's credit rating from Mergent FISD, mapped to NAIC designations through the SVO crosswalk. This source shares no inputs with the Schedule D SVO extraction (it is the market rating of the security, observed every quarter), so agreement between the two cannot reflect a common extraction or merge artifact.

Under the independent FISD-based designation (Table OA-1), the NAIC 3–5 P&C-versus-Life selling gap is +0.0152 (SE = 0.0104), essentially identical to the +0.0141 obtained under the Schedule D SVO designation, and the NAIC 1–2 placebo is unchanged (−0.0011 versus −0.0012). This pattern, a selling gap in the NAIC 3–5 band and a null in the NAIC 1–2 band, is thus reproduced by a designation constructed entirely independently of the SVO filings, and is not an artifact of the SVO extraction or the modal-merge procedure.

Table OA-1: Within-Group Selling Gap by NAIC Designation Source

	Schedule D SVO (baseline)	Mergent FISD rating (independent source)
P&C subsidiary, NAIC 3–5	+0.0141	+0.0152
	(0.0085)	(0.0104)
P&C subsidiary, NAIC 1–2 (placebo)	−0.0012	−0.0011
	(0.0012)	(0.0019)
Observations, NAIC 3–5	92,159	155,697
Observations, NAIC 1–2	1,318,120	911,263
Holding-company clusters (NAIC 3–5)	60	74

Notes: This table reports within-holding-group regressions of the subsidiary selling gap by the source of each bond's NAIC designation. The dependent variable is an indicator for a quarter-over-quarter decline in a subsidiary's holdings, regressed on a P&C-subsidiary indicator, with group × bond × quarter fixed effects. The two columns differ only in how each bond's NAIC designation is assigned: from the Schedule D SVO filings (carried across quarters by modal designation) versus from the bond's contemporaneous Mergent FISD credit rating mapped through the SVO crosswalk. Standard errors, clustered by holding company, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Robustness to the holdings measure. Replacing the amortized-cost-based selling indicator with one based on the bond's *par* value held—which moves only on sale, maturity, or paydown and is therefore immune to premium and discount amortization—leaves the within-group result essentially unchanged: the NAIC 3–5 P&C-versus-Life selling gap is +0.0129 (wild-cluster bootstrap $p = 0.12$), the pooled interaction is +0.0139 ($p = 0.05$), and the NAIC 1–2 placebo remains null (−0.0016).

OA-2.7 Representativeness of the Dual-Type Groups

The within-group test identifies off the holding companies that operate both a Life and a P&C subsidiary. These are a selected set, larger and more diversified than single-type insurers, which raises a question of external validity: does an estimand recovered from dual-type conglomerates speak to the broader insurer population that produces the market-wide stabilization? Table OA-2 compares the three group types at a pre-crisis (2007Q1) snapshot.

The dual-type groups are indeed larger, carrying a median of five subsidiaries, 841 distinct bonds, and $1.9 billion in bond holdings, against two subsidiaries and $88–356 million for single-type groups. On the margins that drive the accounting mechanism, however, they are ordinary. The mechanism runs through the regulatory capital constraint (Section 5.5), and median risk-based capital ratios are economically indistinguishable across the three types (544, 519, and 508), so the capital pressure that makes historical-cost accounting bind operates no differently inside a dual-type group than in a stand-alone Life or P&C insurer. Their portfolios are representative in composition as well: the speculative-grade share of dual groups (1.2 percent) lies between that of Life-only (2.6 percent) and P&C-only (0.2 percent) groups, the natural blend of the two business lines, so the NAIC 3–5 cells that identify the accounting channel are held in normal proportions rather than disproportionately concentrated in speculative grade. The within-group estimate is thus a local effect for larger groups, but for groups whose capital position and speculative-bond exposure, the two features the mechanism depends on, match those of the single-type insurers that account for the bulk of insurer ownership.

Table OA-2: Holding Groups by Type: Dual vs. Single-Type Insurers

	Dual (Life + P&C)	Life-only	P&C-only
Number of groups	103	152	301
Subsidiaries	5	2	2
Bonds held	841	295	155
Bond holdings ($M)	1,896	356	88
Speculative-grade share	1.2%	2.6%	0.2%
Risk-based capital ratio	544	519	508

Notes: This table reports median values across NAIC holding groups active in 2007Q1, by group type. A group is classified Dual if it operates at least one Life and one P&C subsidiary over the sample, and Life-only or P&C-only otherwise; the same group-code filter as the within-group sample applies. Bond holdings are the sum of statutory amortized cost. Speculative-grade share is the cost-weighted fraction of bond holdings designated NAIC 3 or lower, using the same modal SVO designation as the main test. The risk-based capital ratio is the company-level RBC ratio from S&P Global; higher values indicate better capitalization.

OA-3 Residual MTM and Rating Tiers

OA-3.1 Composition of the Non-Insurer Holdings

The baseline specification in Table 4 enters insurer share and MF/ETF share side-by-side. Because insurer and MF/ETF ownership do not sum to one, a cleaner decomposition enters the "residual MTM" share, the share of non-insurer holders operating under mark-to-market rules, as an explicit regressor.

Table OA-3: Residual MTM Composition and the Drawdown Gap

	(1) Baseline	(2) Both	(3) Crisis Int.	(4) MF/ETF Only	(5) Horse Race
Insurer Share	-0.0124***	-0.0129***	-0.0134***		-0.0199***
	(0.0035)	(0.0035)	(0.0035)		(0.0042)
Fund Share				0.0089	-0.0698***
				(0.0068)	(0.0184)
Residual MTM (MF Share / (1 - Ins Share))		0.0140***			0.0580***
		(0.0048)			(0.0126)
Resid MTM x Crisis 1			0.0327**		
			(0.0161)		
Resid MTM x Crisis 2			-0.0140***		
			(0.0051)		
Resid MTM x Crisis 3			-0.0177**		
			(0.0074)		
Resid MTM x Crisis 4			0.0460***		
			(0.0090)		
Observations	31245	31245	31245	31245	31245
Adjusted R^2	0.817	0.817	0.818	0.817	0.817
Specification	Baseline	Both	Crisis Interact	MF/ETF Only	Horse Race
Fixed Effects	Full	Full	Full	Full	Full
Clusters	4,120	4,120	4,120	4,120	4,120

Notes: This table reports how the composition of mark-to-market ownership relates to the crisis-period drawdown gap. The dependent variable is the log bond drawdown. Residual MTM = MF/ETF share / (1 − insurer share), i.e. the share of non-insurer holders subject to MTM. Column (1) is the insurer-only baseline. Column (2) adds Residual MTM. Column (3) interacts Residual MTM with crisis-event dummies. Column (4) reports the MF/ETF-only baseline. Column (5) enters insurer share, Residual MTM, and MF/ETF share jointly. All columns include the fully-saturated fixed-effect structure. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

The Residual MTM coefficient (Table OA-3) is positive and significant at +0.0140*** in column (2) and rises to +0.0580*** in the fully joint specification. Once insurer share and Residual MTM are controlled for, the Fund Share coefficient turns from +0.0089 (insignificant, column 4) to −0.0698*** (column 5): the apparent destabilizing role attributed to mutual funds is driven by the broader class of MTM-constrained holders, not mutual funds specifically.

OA-3.2 Investment Grade vs. Speculative Grade

Table OA-4: Investment-Grade vs. Speculative-Grade Subsamples

	(1)	(2)	(3)	(4)	(5)
Insurer Share	-0.012***	-0.011		-0.011***	-0.068*
	(0.004)	(0.037)		(0.004)	(0.041)
Investment Grade		0.000			
		(.)			
Ins Share x Inv Grade		-0.001	-0.012***		
		(0.037)	(0.004)		
Ins Share x Spec Grade			-0.011		
			(0.037)		
Adjusted R^2	0.82	0.82	0.82	0.81	0.91
Sample	Full	Full	Full	IG Only	HY Only
Event Dummy	Yes	Yes	Yes	Yes	Yes
Bond Size	Yes	Yes	Yes	Yes	Yes
Bond Duration	Yes	Yes	Yes	Yes	Yes
Bond Rating	Yes	Yes	Yes	Yes	Yes
Issuer FE	Yes	Yes	Yes	Yes	Yes
Coupon Indicator	Yes	Yes	Yes	Yes	Yes
Senior Indicator	Yes	Yes	Yes	Yes	Yes
p(IG=Spec)			1		
Identifying Issuers	4,120	4,120	4,120	3,822	334
Obs	31,245	31,245	31,245	30,077	1,151

Notes: This table reports estimates of the baseline drawdown regression split by rating class. The dependent variable is the log bond drawdown. Column (1) is the pooled baseline. Column (2) adds the Ins × Inv Grade interaction. Column (3) reports interaction terms with main effects suppressed. Column (4) restricts to investment-grade bonds; column (5) to high-yield. All columns include the fully-saturated fixed-effect structure with issuer fixed effects. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

Table OA-4 splits the sample by rating class. The IG subsample yields -0.011^{***}; the HY subsample yields -0.068^{*}, larger in magnitude but less precisely estimated given the smaller sample (1,151 observations). A formal test does not reject equality of the two point estimates.

OA-4 Crisis Heterogeneity and Severity

The baseline specification pools across the four crisis events. The model predicts the stabilization effect should be stronger when the unrealized-loss gap is larger, that is, during severe crises with large aggregate price dislocations.

OA-4.1 Crisis-by-Crisis Subsamples

Table OA-5: Crisis Heterogeneity: Insurer Stabilization across Individual Crisis Events

	(1) Pooled	(2) Crisis 1	(3) Crisis 2	(4) Crisis 3	(5) Crisis 4	(6) Interaction
Insurer Share	-0.012***	-0.030**	0.004	-0.003	-0.023***	
	(0.004)	(0.012)	(0.003)	(0.004)	(0.008)	
Ins Share x Crisis 1						-0.030***
						(0.012)
Ins Share x Crisis 2						0.004
						(0.003)
Ins Share x Crisis 3						-0.003
						(0.004)
Ins Share x Crisis 4						-0.023***
						(0.008)
Observations	31245	6087	7085	8212	9861	31245
Adjusted R^2	0.82	0.78	0.57	0.86	0.67	0.82
Sample	Pooled	Crisis 1	Crisis 2	Crisis 3	Crisis 4	Pooled (Interacted)
Clusters	4,120	959	1,009	1,046	1,106	4,120

Notes: This table reports insurer stabilization estimated separately across individual crisis events. The dependent variable is the log bond drawdown. Column (1) reports the pooled specification. Columns (2)–(5) re-estimate separately in each of the four crisis windows: GFC (2008–2009), 2011 Euro/US downgrade, 2014–2016 commodity collapse ("Taper"), and the 2020 COVID-19 crisis. Column (6) runs the pooled specification with insurer share interacted with crisis-event dummies. All specifications include the fully-saturated fixed-effect structure. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

The pattern across columns of Table OA-5 tracks crisis severity: the GFC yields -0.030^{**} and COVID -0.023^{***}, the two severe crises, while the milder episodes are insignificant, with the 2014–2016 commodity collapse at -0.003 and the 2011 downgrade at $+0.004$; the stabilization effect is thus concentrated in the two large dislocations.

OA-4.2 Triple Interaction: Insurer Type $\times$ Duration $\times$ Severity

The model's sharpest prediction is that stabilization should be strongest along three dimensions simultaneously: (a) long-duration bonds, (b) during severe crises when unrealized losses are largest, and (c) among insurer types subject to HCA.

Table OA-6: Triple Interaction: Insurer Type × Bond Duration × Crisis Severity

	(1)	(2)	(3)	(4)	(5)
Insurer Share	0.021***	0.018***			
	(0.005)	(0.004)			
Ins Share x Long Dur	-0.048***	-0.026***			
	(0.006)	(0.005)			
Ins Share x Severe Crisis		0.006			
		(0.009)			
Ins Share x Long Dur x Severe		-0.043***			
		(0.011)			
Life Ins Share			0.021***		0.020***
			(0.004)		(0.004)
Life Share x Long Dur			-0.025***		-0.025***
			(0.005)		(0.005)
Life Share x Severe Crisis			-0.005		-0.007
			(0.009)		(0.009)
Life Share x Long Dur x Severe			-0.037***		-0.032***
			(0.012)		(0.012)
P&C Ins Share				0.017*	0.009
				(0.010)	(0.009)
P&C Share x Long Dur				-0.091***	-0.085***
				(0.016)	(0.016)
P&C Share x Severe Crisis				0.103***	0.076***
				(0.028)	(0.028)
P&C Share x Long Dur x Severe				-0.062*	-0.055
				(0.037)	(0.037)
Clusters	4,120	4,120	4,120	4,120	4,120
Obs	31,245	31,245	31,245	31,245	31,245

Notes: This table reports the triple interaction of insurer type, bond duration, and crisis severity in the drawdown regression. The dependent variable is the log bond drawdown. Columns (1)–(2) use pooled insurer share. Column (3) isolates Life-insurer share. Column (4) isolates P&C share. Column (5) enters Life and P&C shares jointly with their full interaction sets. Long Dur is an indicator for remaining maturity > 5 years; Severe Crisis is an indicator for the Great Recession and COVID-19. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

In Table OA-6, the triple interaction Ins. share × Long Dur × Severe Crisis enters at -0.043^{***} in column (2), and the Life-insurer analogue is -0.037^{***} in column (3). The P&C triple interaction in column (4) is -0.062^{*} (SE = 0.037), estimated far less precisely and off a much smaller P&C ownership base. The HCA-consistent dimensions—long duration, severe crisis, and Life ownership—load together as a reduced-form pattern; we read the marginally significant P&C coefficient cautiously given its imprecision.

OA-5 Supplementary Tests

OA-5.1 Asymmetry: Deep vs. Mild Drawdowns

Table OA-7: Asymmetry: Insurer Effect in Deep vs. Mild Drawdown Bonds

	(1) Baseline	(2) Deep DD	(3) Mild DD	(4) Ins x Deep	(5) MF x Deep	(6) Joint
Insurer Share	-0.012***	-0.011**	0.000	-0.010***		-0.006*
	(0.004)	(0.005)	(0.003)	(0.003)		(0.003)
Deep Drawdown (above crisis median)				0.073***	0.072***	0.079***
				(0.002)	(0.002)	(0.003)
Ins Share x Deep Drawdown				-0.008**		-0.014***
				(0.004)		(0.005)
Fund Share					0.036***	0.037***
					(0.007)	(0.007)
MF/ETF Share x Deep Drawdown					-0.037***	-0.048***
					(0.010)	(0.011)
Observations	31245	15463	15367	31245	31245	31245
Adjusted R^2	0.82	0.88	0.73	0.87	0.87	0.87
Sample	Full	Deep	Mild	Full	Full	Full
Fixed Effects	Full	Full	Full	Full	Full	Full
p(Ins base=0)						0.095
Clusters	4,120	2,853	2,873	4,120	4,120	4,120

Notes: This table reports the asymmetry in the insurer effect between deep- and mild-drawdown bonds. The dependent variable is the log bond drawdown. Deep Drawdown is an indicator for bonds with above-median drawdown within the crisis event. Column (1) is the baseline. Columns (2)–(3) restrict to deep- and mild-drawdown bonds. Column (4) interacts insurer share with the deep-drawdown indicator; column (5) runs the analogous test for MF/ETF share; column (6) is the full joint specification. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

The Ins × Deep Drawdown interaction (Table OA-7) is -0.008^{**} in column (4), rising to -0.014^{***} in column (6). In mild-drawdown bonds, the insurer coefficient is a near-zero +0.000, consistent with an accounting mechanism that binds only when unrealized losses are material.

OA-5.2 Life vs. P&C Shares by Rating Tier

Table OA-8 splits insurer ownership into separate Life and P&C shares and re-estimates the baseline across rating tiers.

Table OA-8: Life vs. P&C Insurer Shares by Rating Tier

	(1)	(2)	(3)	(4)	(5)
Life Ins Share	-0.014***	-0.007	-0.010**	-0.076**	
	(0.004)	(0.005)	(0.005)	(0.038)	
P&C Ins Share	0.007	-0.004	0.023	-0.016	
	(0.012)	(0.014)	(0.018)	(0.111)	
Life Share x High IG					-0.017***
					(0.004)
Life Share x BBB					-0.008
					(0.005)
Life Share x Sub-IG					-0.013
					(0.035)
P&C Share x High IG					0.017
					(0.014)
P&C Share x BBB					-0.023
					(0.017)
P&C Share x Sub-IG					-0.006
					(0.121)
Sample	All	High IG	BBB	Sub-IG	All
Clusters	4,120	1,785	2,127	334	4,120
Obs	31,245	16,371	13,676	1,151	31,245

Notes: This table reports the baseline drawdown regression with insurer ownership split into separate Life and P&C shares across rating tiers. The dependent variable is the log bond drawdown. Column (1) pools Life and P&C shares. Columns (2)–(4) restrict to High IG (AA or higher), BBB, and Sub-IG subsamples. Column (5) interacts Life and P&C shares with rating-tier indicators in the pooled sample. All columns include the fully-saturated fixed-effect structure. Standard errors, clustered at the issuer-event level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

In column (5), Life $\times$ High IG is -0.017^{***}, while Life $\times$ BBB (-0.008) and Life $\times$ Sub-IG (-0.013) are not individually significant, and all P&C interactions are statistically indistinguishable from zero. The Life stabilization effect is concentrated in the highest-quality investment-grade tier; the P&C null holds in every tier.

OA-5.3 Robustness of the NAIC 2→3 Boundary Test

Table OA-9 subjects the boundary-test estimate of §5.4 (Table 5, Panel A) to a battery of robustness checks. The coefficient on $\Delta\phi \times \text{Post}$ is -0.0144 across specifications. Two-way clustering on bond and quarter (column 2) roughly doubles the standard error, leaving the estimate significant at the ten-percent level. Widening the event window to $[-8, +8]$ quarters (column 4) leaves it essentially unchanged (-0.0120, $p < 0.05$). Adding issuer × quarter fixed effects (column 3), comparing bonds of the same issuer in the same calendar quarter, leaves the point estimate intact (-0.0148).

Table OA-9: Robustness of the NAIC 2→3 Boundary Test

	(1)	(2)	(3)	(4)
$\Delta\phi$ (Life − P&C) × Post	-0.0144***	-0.0144*	-0.0148*	-0.0120**
	(0.0042)	(0.0082)	(0.0079)	(0.0055)
Issuer-Quarter FE	No	No	Yes	No
Clustering	Bond	Bond & Qtr	Bond & Qtr	Bond & Qtr
Window	[-4,4]	[-4,4]	[-4,4]	[-8,8]
Obs.	8,824	8,824	7,889	14,629

Notes: This table reports robustness checks for the NAIC 2→3 boundary test. The dependent variable is the monthly bond return (RET_{EOM}) around NAIC 2→3 downgrades; $\Delta\phi$ is the pre-downgrade Life − P&C ownership gap and Post indicates event quarters ≥ 0. All columns include bond and quarter fixed effects. Column (1) reproduces the baseline of Table 5, Panel A, with standard errors clustered at the bond level; columns (2)–(4) cluster two-way on bond and quarter. Column (3) adds issuer × quarter fixed effects; column (4) widens the event window to $[-8, +8]$. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

A natural concern is whether bonds that end up more heavily Life-held were already on a different return trajectory before the downgrade. Figure OA-1 addresses this with an event study that interacts $\Delta\phi$ with leads and lags of event time (reference period $t = -1$). Under the baseline fixed-effect structure (red), the pre-downgrade interactions drift upward and jointly reject the no-pre-trend null ($F = 3.96$, $p = 0.008$). This pre-trend reflects issuer selection (Life insurers cluster in particular issuers) rather than a violation of the accounting mechanism: adding issuer × quarter fixed effects (green) flattens the pre-period coefficients ($F = 0.85$, $p = 0.47$) while the downgrade-quarter return differential is preserved and, if anything, sharpens (from -0.027 to -0.041; negative values indicate higher downgrade-window returns for P&C-concentrated bonds, the fire-sale-reversal pattern discussed in the main text). The within-holding-company test of §5.4 removes this confound by construction, holding the identical bond fixed across Life and P&C siblings of the same parent.

Figure OA-1: Returns Event Study around NAIC 2→3 Downgrades: Pre-Trend and Its Resolution

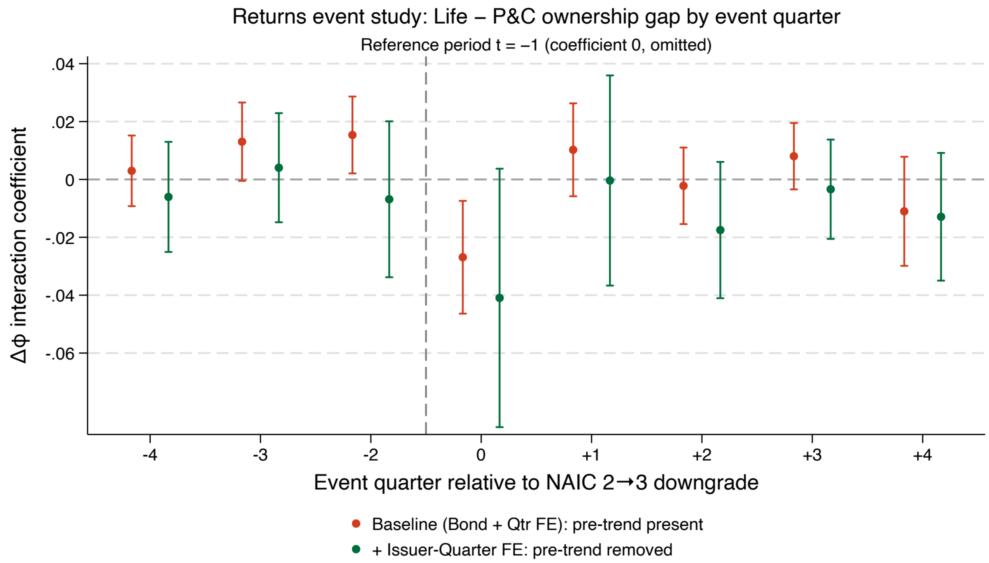

Notes: This figure plots event-study coefficients interacting the pre-downgrade Life − P&C ownership gap $\Delta\phi$ with event-quarter indicators (reference $t = -1$), for monthly bond returns around NAIC 2→3 downgrades. Red: bond and quarter fixed effects (baseline). Green: adds issuer × quarter fixed effects. Capped spikes are 95 percent confidence intervals from bond-clustered standard errors. The baseline pre-trend (joint $p = 0.008$) is absorbed by issuer × quarter fixed effects (joint $p = 0.47$); the downgrade-quarter return differential is preserved.

The decomposition in Table 5, Panel A column (2) (Life share × Post at −0.0146 and P&C share × Post at an insignificant +0.0229) is robust to two-way clustering ($p = 0.32$ for the P&C term), confirming that the differential loads on the Life share (Life-concentrated bonds do not participate in the post-downgrade rebound because their prices never carried a fire-sale discount) rather than on a statistically distinguishable P&C-side effect. The NAIC 1→2 placebo stays at zero under the same clustering (−0.0011, $p = 0.74$).

OA-5.4 Real Effects by Holder Type and Rating Band

Table 9 establishes that post-crisis firm investment loads on the HCA-reporting (Life) insurer share and not on the MTM-reporting (P&C) share, but it pools holdings across rating tiers and so cannot reject equality of the two holder coefficients at conventional levels. Table OA-10 sharpens the contrast by interacting each holder share with the NAIC tier of the underlying bonds: the speculative-grade NAIC 3–5 band, where statutory recognition diverges across the two holder types, and the investment-grade NAIC 1–2 band, where the two types are treated identically and the rule cannot bind. The accounting channel predicts that the HCA–MTM gap should appear in the divergence band and vanish in the placebo band, which is the necessary-condition logic that runs through the bond-level tests of §5.4.

Table OA-10: Real Effects by Holder Accounting Type and NAIC Rating Band

	CAPEX	*CAPEX + Acquisitions*
	(1)	(2)
Post × HCA (Life) share × NAIC 3–5	0.0235***	0.0911***
	(0.0090)	(0.0189)
Post × MTM (P&C) share × NAIC 3–5	-0.0843*	-0.1104
	(0.0468)	(0.1034)
Post × HCA (Life) share × NAIC 1–2	0.0129**	0.0820***
	(0.0066)	(0.0160)
Post × MTM (P&C) share × NAIC 1–2	-0.0116	-0.0463
	(0.0402)	(0.1338)
p(HCA=MTM), NAIC 3–5 (divergence)	0.044	0.080
p(HCA=MTM), NAIC 1–2 (placebo)	0.608	0.581
Firm FE	Yes	Yes
Year FE	Yes	Yes
Post × Size/Lev/Dur FE	Yes	Yes
Observations	5,876	5,527

Notes: This table reports difference-in-differences estimates of firm investment around the Great Recession by holder accounting type and NAIC rating band. The dependent variable is annual CAPEX (column 1) and CAPEX + acquisitions (column 2). Each firm's insurer ownership is split into the share held by HCA-reporting (Life, AVR) and MTM-reporting (P&C, non-AVR) subsidiaries, and each holder share is further interacted with the post indicator and the bonds' NAIC tier: NAIC 3–5, where statutory treatment diverges across holder types, versus NAIC 1–2, where it agrees. The reported p-values test equality of the HCA and MTM coefficients within each band by the wild-cluster bootstrap (Rademacher, 9,999 reps, clustered on the firm). All columns include firm and year fixed effects and pre-period size, leverage, and duration interacted with post. Standard errors, clustered at the firm level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively, based on analytic standard errors.

The estimates follow the predicted pattern: in the divergence band, a 10 percentage-point increase in HCA (Life) ownership raises post-crisis investment-to-assets by about 0.24 percentage points for CAPEX (0.0235***) and 0.91 percentage points for total investment (0.0911***), while the MTM (P&C) coefficient is negative and imprecise; the two coefficients are statistically distinguishable under wild-cluster-robust inference ($p = 0.044$ for CAPEX, $p = 0.080$ for total investment). In the investment-grade placebo band, where the recognition rule treats Life and P&C identically, the HCA and MTM coefficients are statistically indistinguishable ($p = 0.61$ and 0.58). The investment response thus appears precisely at the accounting boundary and only for the holder type the rule treats differently: the within-band p-values reconcile the cleanly separated but formally insignificant pooled test of main-text Table 9, since the P&C share concentrates in investment-grade bonds and pooling across tiers dilutes the divergence-band contrast that identifies the channel.

OA-5.5 Primary-Market Access During the Issuance Freeze

The investment responses of Table 8 presuppose a financing margin: if statutory accounting stabilizes the prices of insurer-held bonds in a crisis, issuers with a larger insurer base should have found it easier to return to the primary market when new issuance froze in 2008–09. Table OA-11 takes this prediction to the issuance data directly. We match every corporate, non-convertible, U.S.-dollar, fixed-coupon issue in Mergent FISD offered between 2004 and 2011 to the Table 8 firm panel through the six-digit issuer CUSIP (3,813 issues across the 876 sample firms) and regress an indicator for issuing any bond in a given year (column 1), and the inverse hyperbolic sine of the amount raised (column 2), on the firm's pre-crisis (2004–07 average) insurer share interacted with separate indicators for the freeze (2008–09) and the recovery (2010–11), retaining the firm and year fixed effects and the size, leverage, and duration controls of the main specification, interacted here with the freeze and recovery indicators.

Table OA-11: Primary-Market Access During the 2008–09 Issuance Freeze

	Issued any bond (1)	asinh(Amount issued) (2)
Freeze (2008–09) × Pre-crisis ins. share	0.2421***	3.5186***
	(0.0570)	(0.7719)
Recovery (2010–11) × Pre-crisis ins. share	0.0163	0.2880
	(0.0616)	(0.8438)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	5,867	5,867

Notes: This table reports difference-in-differences estimates of bond-market access on the Table 8 firm-year panel (2004–2011). The dependent variable is an indicator for the firm issuing any bond in the fiscal year (column 1) and the inverse hyperbolic sine of the total amount issued (column 2), constructed from all Mergent FISD corporate, non-convertible, U.S.-dollar, fixed-coupon issues matched to sample firms by six-digit issuer CUSIP. Pre-crisis ins. share is the firm's average insurer ownership share over fiscal years 2004–2007, interacted with indicators for the issuance freeze (2008–09) and the recovery (2010–11). Both columns include firm and year fixed effects and pre-period size, leverage, and duration deciles interacted with the freeze and recovery indicators. Standard errors, clustered at the firm level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

The access advantage appears exactly when the mechanism predicts: during the freeze, a 10 percentage-point higher pre-crisis insurer share is associated with a 2.4 percentage-point higher probability of issuing (column 1, 0.2421***)—roughly a tenth of the freeze-year issuance rate, which fell to 24 percent of sample firms in 2008 from 32 percent pre-crisis—while the recovery interaction is a near-zero 0.016. The amount margin moves the same way: the freeze interaction is 3.519*** against a near-zero recovery interaction of 0.288. Untabulated year-by-year estimates concentrate the effect in 2008 itself (+0.32, $t = 3.8$), with largely flat pre-period interactions. We read the table as suggestive rather than conclusive: the freeze-specific timing and the null recovery interaction are consistent with insurer ownership supporting market access when issuance was hardest to obtain, but because insurer holdings concentrate in investment-grade issuers, this margin cannot, on its own, separate the accounting channel from rating-level flight to quality.

OA-6 Placebo Tests

This section reports the full tables underlying the three placebo tests referenced in §5.6 of the main text. Each placebo targets a distinct scope condition of the accounting mechanism and produces a clean null exactly where the model predicts zero effect.

OA-6.1 Health Insurer Placebo

Health insurers, like Life insurers, carry bond holdings at amortized cost under statutory accounting. If the stabilization result reflected a generic feature of insurance-company behavior—patient capital, liability-driven investment horizons, or institutional inertia—then health-insurer ownership share should stabilize bond prices in the same way Life-insurer share does. But Health insurers hold negligible bond portfolios relative to their balance sheets, so the accounting rule is present while the material holdings are not.

The full type decomposition is reported in Table 7 of the main text, which splits the pooled insurer share into separate Life, P&C, and Health shares and enters them both separately and jointly. Entered on its own, Life share is −0.0139*** (SE = 0.0038), P&C share is +0.0138 (insignificant), and Health share is +0.0101 with a standard error of 0.0924, nearly seven times the magnitude of the Life point estimate. Entering all three jointly leaves Life negative and significant while P&C and Health remain indistinguishable from zero. The health-insurer coefficient is an uninformative null; with holdings this small the test has no power to detect stabilization, which is the scope condition at work: the accounting rule alone is insufficient; the mechanism requires both the amortized-cost classification and economically meaningful bond holdings.

OA-6.2 Non-Crisis Placebo

Outside of crisis periods, there are no material unrealized losses for historical-cost accounting to insulate and no binding capital pressure to make that insulation valuable. If insurer share predicted returns in calm quarters, the result would reflect permanent selection into low-volatility bonds rather than the accounting mechanism.

Table OA-12: Non-Crisis Placebo: Insurer Share and Bond Returns in Calm Periods

	(1)	(2)	(3)	(4)	(5)
	Dependent variable: Quarterly excess return (non-crisis sample)				
Insurer share	−0.0012	0.0011	−0.0066***	−0.0013	0.0001
	(0.0009)	(0.0012)	(0.0010)	(0.0010)	(0.0009)
Year-Quarter FE	Yes	Yes	Yes	Yes	Yes
Bond Size FE	No	Yes	Yes	Yes	Yes
Duration FE	No	No	Yes	Yes	Yes
Rating FE	No	No	No	Yes	Yes
Issuer FE	No	No	No	No	Yes
Observations	404,989	404,989	399,808	349,157	349,026

Notes: This table reports a non-crisis placebo relating insurer share to bond returns in calm periods. The dependent variable is the quarterly excess return for bond-quarter observations outside the four crisis windows defined in the main text. All specifications include year-quarter fixed effects and progressively add bond size, duration, rating, and issuer fixed effects. Standard errors, clustered at the bond level, are reported in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

In column (1) with only year-quarter fixed effects, insurer share enters at −0.0012 (n.s.). Adding size flips the sign (+0.0011, n.s.), adding duration produces −0.0066***, adding rating attenuates it to −0.0013 (n.s.), and adding issuer fixed effects collapses the coefficient to +0.0001 ($p = 0.91$), a near-exact zero. Selection on observables is thus present in the raw comparison, but the baseline fixed-effect structure absorbs it.

OA-6.3 Within-Group NAIC 1–2 Placebo

The third falsification is embedded in Table 5, Panel B of the main text: for NAIC 1–2 bonds, where both Life and P&C subsidiaries use amortized cost, the within-group selling differential is −0.0012 ($p = 0.32$), statistically indistinguishable from zero. The accounting divergence between insurer types—present for NAIC 3–5 bonds, absent for NAIC 1–2 bonds—is the *necessary condition* for the differential selling documented in Panel B.

No single alternative explanation—patient capital, liability-driven investment, low-volatility selection, or institutional inertia—predicts all three nulls simultaneously.

OA-7 Diagnostics for the Shift-Share Instrument

Two diagnostics probe the instrument behind Table 3. The first asks which insurers supply the identifying variation. Because the instrument sums participating insurers' historical allocation intensities, $Z_i = \sum_j \bar{\phi}_j D_{ij}$, it admits a decomposition in the spirit of the Rotemberg weights of Goldsmith-Pinkham et al. (2020): each insurer's weight is proportional to its allocation intensity times the covariance of its participation indicator with the residualized treatment. Panel B of Table OA-13 reports the concentration of these weights. Identification is diffuse: the largest single entity carries 5.3 percent of the total weight, the top five carry 17.5 percent, and the top ten carry 27.5 percent, spread across 3,298 weighted entities: no single insurer's allocation behavior drives the estimate. As is standard in shift-share settings, a subset of entities (1,404) carries negative weights summing to −0.398, so under heterogeneous effects the two-stage least squares estimand is not a strict convex combination of insurer-level effects; the diffuseness of the positive mass limits the scope for any single entity's heterogeneity to dominate.

The second diagnostic asks whether the instrument is balanced on pre-determined bond characteristics. Unconditionally it is not, and should not be: allocations scale with deal size and tilt toward safer issues, so a one-standard-deviation increase in Z is associated with lower offering yields (−0.165 standard deviations), longer-dated issues sorting negatively (−0.183), and lower gross underwriting spreads (−0.181), all significant at the 1 percent level (Panel A, first column). This is precisely why the exclusion restriction is stated conditionally and why the Table 3 specifications absorb size, rating, duration, and issuer. Conditional on that observable structure (second column), the residual associations shrink by roughly an order of magnitude: offering yield at −0.033 standard deviations ($p = 0.050$), coupon at +0.058 ($p = 0.030$), and maturity and underwriting spread indistinguishable from zero. Two of the four retain marginal significance, but at magnitudes of at most 0.06 standard deviations per standard deviation of the instrument, the residual imbalance is economically negligible. Together with the non-crisis placebo above (the instrumented ownership measure predicts nothing when no insurer faces capital pressure), these diagnostics support reading the instrument as quasi-random allocation exposure conditional on the regression's observables.

Table OA-13: Shift-Share Instrument Diagnostics: Balance and Weight Concentration

Panel A: Balance of pre-determined bond characteristics				
	Unconditional		Conditional	
	β	N	β	N
Offering yield (%)	-0.165***	1,908	-0.033*	1,622
	(0.018)		(0.017)	
Coupon (%)	-0.067***	2,341	0.058**	2,019
	(0.023)		(0.027)	
Maturity (years)	-0.183***	2,351	-0.012	2,019
	(0.020)		(0.010)	
Gross underwriting spread	-0.181***	1,863	0.014	1,598
	(0.022)		(0.021)	
Panel B: Rotemberg-style insurer weights on the instrument				
Largest single entity's share of total weight		0.053		
Top 5 entities' share of total weight		0.175		
Top 10 entities' share of total weight		0.275		
Insurer entities receiving weight		3,298		
Entities with negative weight		1,404		
Aggregate negative-weight mass		-0.398		

Notes: Panel A reports balance tests of pre-determined bond characteristics against the shift-share instrument on the bond-level cross-section of the Table 3 estimation sample. Each cell is the coefficient from a regression of the standardized characteristic on the standardized instrument. Unconditional specifications absorb offering-year fixed effects; conditional specifications additionally absorb the size-decile, duration-category, rating-group, and issuer fixed effects that enter the instrumental-variables specifications. Standard errors, clustered by issuer, are reported in parentheses. Panel B summarizes the distribution of Rotemberg-style insurer weights on the instrument, computed as each insurer's allocation intensity times the sum of the residualized treatment over the bonds in which it participated at issuance, normalized to sum to one. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels (two-tailed), respectively.

OA-8 Proofs of Theoretical Results

Proof of Proposition 3.1.

Case 1: All MTM Insurers Sell, All HCA Insurers Hold. Suppose the interim price at the minimum supply level satisfies $b_1(\underline{S}_1) \le \bar{c}(\underline{S}_1)$, so that the precautionary buffer is breached before any insurer sells. By Assumption 3.1, all MTM insurers are compelled to liquidate, $k_M^* = n_M$, raising secondary market supply to at least $\underline{S}_1 + S_{\text{MTM}}$. Assumption 3.1 further implies that the post-MTM clearing price falls at or below the solvency threshold common to all insurers: $b_1(\underline{S}_1 + S_{\text{MTM}}) \le D_I/A_I$. Because solvency is a necessary condition for any voluntary sale, HCA insurers are precluded from selling regardless of the comparison between b_1 and b_2. They therefore hold, $k_H^* = 0$, and the equilibrium supply and price are

$$S_1^* = \underline{S}_1 + S_{\text{MTM}}, \qquad b_1^* = b_1\left(\underline{S}_1 + S_{\text{MTM}}\right) \le \frac{D_I}{A_I}. \tag{17}$$

Case 2: All Insurers Hold. Suppose the initial price exceeds the buffer, $b_1(\underline{S}_1) > \bar{c}(\underline{S}_1)$, so that MTM insurers retain discretion over their sales decision. If additionally $b_1(\underline{S}_1) \le b_2$, selling at $t = 1$

is unprofitable for all insurers: even absent any insurer sales, the interim price does not exceed the continuation value, and additional supply would only depress b_1 further. Both HCA and MTM insurers hold, $k_H^* = k_M^* = 0$, and

$$S_1^* = \underline{S}_1, \qquad b_1^* = b_1\big(\underline{S}_1\big). \tag{18}$$

Case 3: Profitable Sales with MTM Discretion. Suppose $b_1(\underline{S}_1) > \bar{c}(\underline{S}_1)$ and $b_1(\underline{S}_1) > b_2$, so that the MTM buffer is not breached and bond sales are profitable for both MTM and HCA insurers. Let $k_H \le n_H$ and $k_M \le n_M$ denote the number of selling HCA and MTM insurers, respectively, and define $k^* := k_H + k_M$ as the total equilibrium number of selling insurers. In any equilibrium, k^* must satisfy $b_1(\underline{S}_1 + k^* \cdot A_I) > \max(b_2, D_I/A_I)$; if this condition fails, a sale is either unprofitable or renders selling insurers insolvent, so at least one seller would prefer to hold ex post.

Case 3a: All Insurers Sell. If $b_1(\underline{S}_1 + (n_M + n_H) \cdot A_I) > \max(b_2, D_I/A_I)$, the interim price remains above both the continuation value and the solvency threshold even when every insurer liquidates. The unique equilibrium features full participation, $k^* = n_M + n_H$, with

$$S_1^* = \underline{S}_1 + (n_M + n_H) \cdot A_I, \qquad b_1^* = b_1(S_1^*) > \max\Big(b_2, \tfrac{D_I}{A_I}\Big). \tag{19}$$

Case 3b: Partial Selling. If $b_1(\underline{S}_1 + (n_M + n_H) \cdot A_I) < \max(b_2, D_I/A_I)$, full participation is infeasible: it would either be unprofitable or render selling insurers insolvent. The equilibrium total number of selling insurers is therefore $k^* < n_M + n_H$, defined as

$$k^* := \sup\Big\{1 \le k \le n_H + n_M : b_1\big(\underline{S}_1 + k \cdot A_I\big) > \max\Big(b_2, \tfrac{D_I}{A_I}\Big)\Big\}, \tag{20}$$

with equilibrium supply $S_1^* = \underline{S}_1 + k^* \cdot A_I$ and price $b_1^* = b_1(S_1^*)$. Which constraint stops further sales depends on the relative magnitude of b_2 and D_I/A_I.

Case 3bi: Solvency binds ($D_I/A_I \ge b_2$). At $k^* + 1$ sellers, the clearing price falls below the solvency threshold, $b_1(\underline{S}_1 + (k^* + 1) \cdot A_I) < D_I/A_I$, rendering every selling insurer insolvent. Both HCA and MTM insurers therefore abstain from the marginal sale. In particular, no cascade arises in which HCA sales depress b_1 below the MTM buffer and thereby compel additional MTM sales: because HCA and MTM insurers are subject to the same solvency condition, any price realization that forces MTM insurer insolvency simultaneously forces insolvency on selling HCA insurers. Both types of insurer refrain from the solvency-violating sale for the same reason.

Case 3bii: Profitability binds ($D_I/A_I < b_2$). At $k^* + 1$ sellers, the clearing price falls below the continuation value, $b_1(\underline{S}_1 + (k^* + 1) \cdot A_I) < b_2$, making the marginal sale unprofitable. The solvency threshold remains slack. Both HCA and MTM insurers prefer to hold and collect the higher continuation value b_2.

In both Cases 3bi and 3bii, the aggregate quantity k^* is uniquely determined by (20), but its decomposition into HCA and MTM sellers is indeterminate: since the MTM buffer is not breached at S_1^*, MTM insurers face no regulatory compulsion and behave identically to HCA insurers. Any pair (k_H^*, k_M^*) satisfying $k_H^* + k_M^* = k^*$, $k_H^* \le n_H$, and $k_M^* \le n_M$ constitutes an equilibrium, reflecting coordination indeterminacy among discretionary sellers. □

Proof of Lemma 3.2. Define $Q^* := D_1(D_I/A_I)$ as the total supply level at which the clearing price equals the solvency threshold, so that $b_1(Q^*) = D_I/A_I$. By Assumption 3.1, the buffer satisfies $\bar{c}^{(i)} = b_1(S_1^{*,(i)})$, where $S_1^{*,(i)}$ solves $b_1(S_1 + S_{\mathrm{MTM}}^{(i)}) = D_I/A_I$, which requires $S_1^{*,(i)} + S_{\mathrm{MTM}}^{(i)} = Q^*$, and hence $S_1^{*,(i)} = Q^* - S_{\text{MTM}}^{(i)}$. If $n_M^{(1)} > n_M^{(2)}$, then $S_{\text{MTM}}^{(1)} > S_{\text{MTM}}^{(2)}$, and therefore $S_1^{*,(1)} < S_1^{*,(2)}$. Since $b_1(\cdot)$ is strictly decreasing in supply, it follows that $\bar{c}^{(1)} = b_1(S_1^{*,(1)}) > b_1(S_1^{*,(2)}) = \bar{c}^{(2)}$. $\square$

Proof of Proposition 3.2. We define $M := \max(b_2, D_I/A_I)$.

We write the expected price difference directly from the case structure of Proposition 3.1. In economy i, the equilibrium price equals $b_1(\underline{S}_1 + S_{\text{MTM}}^{(i)})$ in Case 1 and the discretionary price in Cases 2 and 3. We therefore have

$$\begin{aligned}
&\mathbb{E}[b_1(S_1^*) \mid \text{econ } 1] - \mathbb{E}[b_1(S_1^*) \mid \text{econ } 2] \\
&= \mathbb{E}\Big[b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le \bar{c}^{(1)}\}} - b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(2)}\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le \bar{c}^{(2)}\}}\Big] \qquad (21)\\
&\quad + \mathbb{E}\Big[\mathbf{1}_{\{\bar{c}^{(2)} < b_1(\underline{S}_1)\le \bar{c}^{(1)}\}}\Big(-b_1\Big(\underline{S}_1\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le b_2\}} - b_1\Big(\underline{S}_1 + N_I A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)>M\}} \\
&\qquad\qquad - b_1\Big(\underline{S}_1 + k^* A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)\le M\}}\Big)\Big].
\end{aligned}$$

The second block captures the price that economy 2 receives on the intermediate region $\{\bar{c}^{(2)} < b_1(\underline{S}_1) \le \bar{c}^{(1)}\}$, where economy 1 is in Case 1 but economy 2 is in Cases 2 or 3. We now expand the first block using $\{b_1(\underline{S}_1) \le \bar{c}^{(1)}\} = \{b_1(\underline{S}_1) \le \bar{c}^{(2)}\} \cup \{\bar{c}^{(2)} < b_1(\underline{S}_1) \le \bar{c}^{(1)}\}$ and collect terms:

$$\begin{aligned}
&\mathbb{E}[b_1(S_1^*) \mid \text{econ } 1] - \mathbb{E}[b_1(S_1^*) \mid \text{econ } 2] \\
&= \underbrace{\mathbb{E}\Big[\Big(b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big) - b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(2)}\Big)\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le \bar{c}^{(2)}\}}\Big]}_{\text{Term 1}} \qquad (22)\\
&\quad + \mathbb{E}\Big[\mathbf{1}_{\{\bar{c}^{(2)} < b_1(\underline{S}_1)\le \bar{c}^{(1)}\}}\Big(b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big) - b_1\Big(\underline{S}_1\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le b_2\}} \\
&\qquad\qquad - b_1\Big(\underline{S}_1 + N_I A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)>M\}} \\
&\qquad\qquad - b_1\Big(\underline{S}_1 + k^* A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)\le M\}}\Big)\Big]. \qquad =: \text{Term 2}
\end{aligned}$$

Term 1 is strictly negative. Since $S_{\text{MTM}}^{(1)} > S_{\text{MTM}}^{(2)}$ and b_1 is strictly decreasing in supply, $b_1(\underline{S}_1 + S_{\text{MTM}}^{(1)}) - b_1(\underline{S}_1 + S_{\text{MTM}}^{(2)}) < 0$ pointwise. Term 1 is therefore strictly negative.

Term 2 is strictly negative. We show that on $\{\bar{c}^{(2)} < b_1(\underline{S}_1) \le \bar{c}^{(1)}\}$, the expression inside the expectation is strictly negative. We proceed in three steps.

Step 1: On $\{\bar{c}^{(2)} < b_1(\underline{S}_1) \le \bar{c}^{(1)}\}$, Case 3a is ruled out. Since $b_1(\underline{S}_1) \le \bar{c}^{(1)}$, Assumption 3.1 implies $b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big) \le \frac{D_I}{A_I} \le M$. Since $N_I \ge n_M^{(1)}$ and b_1 is decreasing in supply, $b_1\Big(\underline{S}_1 + N_I A_I\Big) \le b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big) \le M$, so $\mathbf{1}_{\{b_1(\underline{S}_1+N_I A_I)>M\}} = 0$ on this region. Therefore $\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)>M\}} = 0$, and $\mathbf{1}_{\{b_1(\underline{S}_1)>b_2,\, b_1(\underline{S}_1+N_I A_I)\le M\}} = \mathbf{1}_{\{b_1(\underline{S}_1)>b_2\}}$, since $\mathbf{1}_{\{b_1(\underline{S}_1+N_I A_I)\le M\}} = 1$ on this region. Term 2 therefore simplifies to

$$\text{Term 2} = \mathbb{E}\Big[\mathbf{1}_{\{\bar{c}^{(2)} < b_1(\underline{S}_1)\le \bar{c}^{(1)}\}}\Big(b_1\Big(\underline{S}_1 + S_{\text{MTM}}^{(1)}\Big) - b_1\Big(\underline{S}_1\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le b_2\}} - b_1\Big(\underline{S}_1 + k^* A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2\}}\Big)\Big]. \qquad (23)$$

Step 2: The economy 2 discretionary price exceeds D_I/A_I. By definition of k^* in (20), the equilibrium price in Case 3b satisfies $b_1(\underline{S}_1 + k^* A_I) > D_I/A_I$. Since b_1 is decreasing and $k^* \ge 0$, $b_1(\underline{S}_1) \ge b_1(\underline{S}_1 + k^* A_I)$. Using $\mathbf{1}_{\{b_1 \le b_2\}} + \mathbf{1}_{\{b_1 > b_2\}} = 1$, we obtain

$$\begin{aligned}
&b_1\Big(\underline{S}_1\Big)\mathbf{1}_{\{b_1(\underline{S}_1)\le b_2\}} + b_1\Big(\underline{S}_1 + k^* A_I\Big)\mathbf{1}_{\{b_1(\underline{S}_1)>b_2\}} \\
&\ge b_1\Big(\underline{S}_1 + k^* A_I\Big)\Big(\mathbf{1}_{\{b_1(\underline{S}_1)\le b_2\}} + \mathbf{1}_{\{b_1(\underline{S}_1)>b_2\}}\Big) = b_1\Big(\underline{S}_1 + k^* A_I\Big) > \frac{D_I}{A_I} \ge b_1(\underline{S}_1 + S_{\text{MTM}}^{(1)}).
\end{aligned}$$

where the last inequality follows from Step 1, because on $\{\bar{c}^{(2)} < b_1(\underline{S}_1) \leq \bar{c}^{(1)}\}$ the precautionary selling buffer in economy 1 is breached. The expression inside Term 2 is therefore strictly negative on $\{\bar{c}^{(2)} < b_1(\underline{S}_1) \leq \bar{c}^{(1)}\}$, making Term 2 strictly negative. Since both Term 1 and Term 2 are strictly negative, $\mathbb{E}[b_1(S_1^*) \mid \text{econ } 1] < \mathbb{E}[b_1(S_1^*) \mid \text{econ } 2]$, completing the proof. $\square$