The Cost of Maturity: Do Firms Overpay for Long-Term Debt?*

Henri Servaes† Dian Xu‡

August 27, 2026

Abstract

We estimate the realized cost of corporate bond maturity by comparing long-term bond costs with rating-updated short-rolling counterfactuals. Across 20,501 U.S. fixed-rate corporate bonds issued over January 1997–July 2025, long-term debt cost 1.85 percentage points more per year than a strategy of rolling over debt every 3 years, with a larger wedge among investment-grade issuers. A persistent maturity term premium, not falling short rates, explains the wedge. At the firm level, these costs translate into meaningful additional interest expense, while the occasional years in which rolling over short-term debt is more expensive can generally be funded from cash savings or by cutting share repurchases. The evidence supports debt-maturity conservatism: managers (over)pay for stable interest expense rather than retaining cheaper financing options.

Keywords: cost of debt, debt maturity, corporate bonds, quiet life, debt conservatism
JEL Classification: G12, G30, G32

*We would like to thank Anna Pavlova and Paolo Surico for helpful comments, and Peter Feldhütter and Stephen Schaefer for their help with the historical bond data.

†London Business School, CEPR, and ECGI. Email: hservaes@london.edu.

‡London Business School. Email: dxu@london.edu.

1 Introduction

> *"Alphabet, which is rated double A plus by S&P, this week seized on the massive demand by raising a rare 100-year bond."*
>
> — Financial Times, *February 12, 2026*

In February 2026, Alphabet—rated AA+ by S&P—issued a 100-year bond at a coupon of 6.125%, materially above the 4.125% coupon on the 3-year note that the firm priced on the same day.[1] To lock in a fixed coupon for a century, the company accepted a borrowing rate substantially above what it could obtain at shorter maturities. For a firm with near-sovereign credit quality and uninterrupted market access, this decision raises questions that go beyond any single issuance: *Do firms actually overpay, ex-post, for long-term fixed-rate debt, after accounting for fees and prepayments? If so, how much do they overpay, and why do they accept this cost?*

This paper speaks to a central question in corporate finance: how do firms choose the maturity of their debt? Seminal theoretical work casts maturity as a balance among taxes, debt-overhang and agency costs, adverse selection, endogenous default, and liquidity or rollover risk: short debt can mitigate some incentive problems or signal borrower quality, but it exposes firms to refinancing conditions and can amplify default incentives when market liquidity deteriorates (Myers, 1977; Brick and Ravid, 1985; Flannery, 1986; Diamond, 1991; Leland and Toft, 1996; He and Xiong, 2012; Diamond and He, 2014; He and Milbradt, 2016). Empirical issue- and firm-level evidence links maturity to growth opportunities, asset maturity, firm size, credit quality, managerial incentives, debt-structure heterogeneity, and broader bond-market supply conditions (Barclay and Smith, 1995; Guedes and Opler, 1996; Stohs and Mauer, 1996; Johnson, 2003; Datta, Iskandar-Datta, and Raman, 2005; Rauh and Sufi, 2010; Badoer and James, 2016). Survey, issuance, and hedging evidence indicate that managers explicitly look at interest rates, yield-curve conditions, fixed-versus-floating exposure, and risk-management practices when choosing maturity and rate exposure (Bodnar, Hayt, and Marston, 1998; Graham and Harvey, 2001; Faulkender, 2005; Servaes and Tufano, 2006; Banerjee, Caballero, Kharroubi, Spigt, and Zakrajšek, 2023), while the market-timing and gap-filling literatures show that firms tilt issuance maturity as expected bond returns, the term structure, and Treasury supply change (Baker, Greenwood, and Wurgler, 2003; Greenwood, Hanson, and Stein,

[1] Alphabet notes due 2126, prospectus supplement filed with the SEC on 12 February 2026 (424B2), `https://www.sec.gov/Archives/edgar/data/1652044/000119312526048730/d108423d424b2.htm`. The accompanying shorter tranche was priced on the same day, making the long-short coupon gap immediately apparent. Both tranches are senior unsecured debt.

2010). More recent empirical work links maturity choices to refinancing-risk management, cash buffers, credit-cycle liquidity risk, and the dispersion of maturity dates within firms (Custódio, Ferreira, and Laureano, 2013; Harford, Klasa, and Maxwell, 2014; Mian and Santos, 2018; Choi, Hackbarth, and Zechner, 2018). This literature explains how firms choose debt maturities—why they shorten, lengthen, or stagger them. Crisis evidence shows why this protection can matter: firms with large amounts of long-term debt due in 2008 cut investment more after the 2007 credit shock (Almeida, Campello, Laranjeira, and Weisbenner, 2011), while directly constrained firms reported giving up attractive projects and using cash, credit lines, and asset sales to cope (Campello, Graham, and Harvey, 2010). This literature does not measure the *realized* cost firms pay for longer maturities relative to a counterfactual rolling strategy that refinances at the issuer's prevailing credit quality. That is what we study in this paper: bond by bond, we compare the all-in realized cost of long-term debt with a rolling shorter-maturity strategy the issuer could have executed, with each rollover priced at the firm's prevailing credit rating. This yields a direct measure of the realized *cost of bond maturity*.

For each of the 20,501 unique U.S. corporate bonds in our headline sample, issued over January 1997–July 2025, we compute the **cost wedge**: the difference between the bond's realized annual borrowing cost (the all-in yield adjusted for fees and prepayment) and the cost it would have incurred under a short-term rolling strategy that refinances at the issuer's *prevailing* credit rating at every hypothetical rollover date. The counterfactual uses monthly Merrill/TRACE corporate yield curves by rating, so it reflects the borrowing cost the firm would have faced at each rollover date. Online Appendix A.1 describes the full yield curve construction methodology.

Main findings. We report five main findings.

First, the cost wedge between firms' actual issuance decisions and the counterfactual is large for both investment-grade (IG) and high-yield bonds (HY). The mean annual cost wedge is 1.85 pp/yr (median 1.70 pp/yr) under the 3-year counterfactual. The effect for IG bonds is somewhat larger than for HY bonds: the IG bond-year mean is 1.96 pp/yr (median 1.75 pp/yr), against a high-yield mean of 1.64 pp/yr (median 1.61 pp/yr). This is an IG puzzle: the issuers with the least rollover-risk exposure pay the larger long-minus-short wedge. The wedge rises with contractual tenor and translates into more than $1.5 trillion in aggregate additional interest payments over the sample period in 2025 dollars, with IG issuers bearing the bulk of the cost because they issue larger and longer-maturity bonds and have more public debt outstanding. The HY figure is best read as an *upper bound* on the cost wedge, because our counterfactual assumes the firm

can roll over its debt at the prevailing yield for bonds with the same rating in every period. That assumption is realistic for IG issuers, who retained bond-market access even during the 2008 Global Financial Crisis (GFC) and 2020 COVID crisis periods, but it omits the distress and market-access costs HY firms would face if a rollover failed during market disruptions. Adding those costs would raise the cost of the short-rolling benchmark and therefore reduce the measured HY long-minus-short wedge. Avoiding these costs, which are difficult to quantify, is one of the reasons why HY firms may issue debt with extended maturities. For IG firms, on the other hand, these costs are virtually nil.

Second, much of the long-short cost wedge is explained by a persistent maturity term premium, not by falling short rates over many years during our sample period. That is, our findings are not due to the fact that short rates happened to be lower than expected when firms chose debt maturity ex-ante. Our cost-wedge decomposition shows that the IG wedge is priced mainly at issuance in the long-over-short yield spread. The largest component is the risk-free term premium followed by the credit term premium. The realized base rate change explains some of the year-to-year movement: it raises the wedge after Treasury rates fall and lowers it after rate hikes, but its full-sample contribution is only 0.21 pp/yr for IG firms. Systematic spread changes and firm-specific rating migration also contribute little to the overall cost wedge. This result complements Baker, Greenwood, and Wurgler (2003), who find that firms adjust aggregate issuance maturity when bond returns are predictable. Our results indicate that this response is not large enough to eliminate the positive realized wedge for bonds with longer maturities relative to shorter-term counterfactuals.

Third, we document that preventing rollover risk is not what firms pay for when they issue long-term fixed-rate debt by comparing *IG fixed-rate* and *IG floating-rate bonds*.[2] As in a strategy of rolling over short-term fixed-rate debt, floating-rate bonds reset to a short benchmark and therefore have interest-rate exposure similar to that of a strategy of rolling over short-dated fixed-rate debt. However, unlike that strategy, they do not require repeated refinancing at prevailing credit spreads or after rating migration, thereby removing rollover risk. For IG floating-rate bonds, we actually find a small negative wedge when comparing the interest cost of the actual long-term floating-rate issues to the counterfactual short-term fixed-rate rollovers. This implies that, in the realized IG data, the large fixed-rate wedge is not primarily the price of removing repeated market-access, credit-spread, or rating-migration risk; it is mainly the price of fixing interest rate exposure.

Fourth, the increased volatility of interest payments in the counterfactual scenarios

[2]There are not enough HY floating-rate bonds to make a meaningful comparison.

does not cause IG firms any difficulties. While the short-rolling counterfactual strategy sometimes requires extra cash in years when short rates are high, these shortfalls can easily be covered. For example, for the 3-year counterfactual benchmark for IG firms, only 795 of 11,843 IG firm-years face an additional cash requirement, which averages $40.4 million. 94.8% of these firm-year shortfalls can be covered in any given year by using existing cash reserves and cutting share repurchases; employing positive-wedge savings from prior years and cutting dividends raise coverage to 99.1%, leaving only 7 IG firm-years uncovered. These few remaining cases would likely not mechanically lead to distress because the calculation excludes ordinary margins that IG firms can reasonably use to deal with cash shortfalls, including cutting costs, drawing on revolving credit lines, releasing working capital, timing capex, and short-term bank or commercial paper borrowing. The practical implication is that firms should reserve part of the prior positive-wedge savings as a buffer for future short-rate volatility, consistent with the short-maturity exposure documented by Custódio, Ferreira, and Laureano (2013) and the precautionary-cash mechanism in Harford, Klasa, and Maxwell (2014). The same calculation for HY firms leaves 246 of 2,404 positive-need firm-years uncovered after all four modeled sources of liquidity are exhausted, with final coverage of 89.8%. However, this assumes that these firms can always refinance maturing debt, which may not be the case.

Fifth, the wedge is economically large in aggregate dollars and relative to cash flows. The aggregate excess interest cost over the sample period exceeds $1.5 trillion in 2025 dollars when measured across the January 1997–July 2025 universe of bonds issued by nonfinancial corporations, and $1.2 trillion in 2025 dollars for the subset of those bond issuers that can be matched to firms listed on Compustat. In operating flow terms, the mean (median) 3-year dollar wedge is 5.1% (2.4%) of EBIT, 2.9% (1.7%) of EBITDA, and 23.5% (15.8%) of interest expenses among IG firm-years with positive values for each of these measures.

Explaining the Cost of Maturity: Debt-maturity conservatism. These findings are jointly consistent with a behavioral-agency interpretation. Our back-of-the-envelope calculations suggest that the price paid to avoid both rollover risk and the inability to service debt when interest rates spike by locking in fixed-rate long-term debt is many times higher than the expected losses for IG firms. And even for HY firms, the borrowing cost savings from shortening maturities outweigh expected losses, but the gap is much smaller than for IG firms.

We call the resulting pattern *corporate debt-maturity conservatism*: a repeated issuer tendency to pay for stable long-term fixed-rate financing when measured rollover-insurance

benefits alone do not explain the full price. This interpretation is related to two distinct elements: an agency element of management enjoying a quiet life, and a behavioral element of over-insurance against rollover risk.

(A) Management enjoying a quiet life. This is related to but distinct from the *quiet life* literature pioneered by Bertrand and Mullainathan (2003), who document that weakly governed managers shy away from difficult investment and labor decisions. We identify an analogous *financial-side* pattern in which managers shy away from the rate-uncertainty exposure that rolling over short-term debt would impose. We replicate and extend the agency result of Datta, Iskandar-Datta, and Raman (2005) and find that CFO ownership is negatively related to contractual bond tenor. The behavioral-agency interpretation is consistent with career-concern models (Holmstrom, 1999) and with the broader behavioral-corporate-finance evidence that managerial traits transfer across decision domains (Cronqvist, Makhija, and Yonker, 2012). A CFO who locks in a high rate and watches rates fall likely faces no career penalty, while one who rolls over short-term debt and faces a transient spike risks a board inquiry.

(B) Over-insurance of rollover risk. The second element concerns how firms justify such conservatism internally. The downside of a rollover failure is easy for boards and investors to see, while the term-premium saving from rolling over short-term debt is spread out over time and can be dismissed as the result of a risky strategy. The conclusion is that IG firms pay mainly for base-rate lock-in, while the policy is often justified as protection against bad rollover states.

For IG firms, the ex-post estimates imply that issuing shorter-term debt rather than long-term fixed-rate debt would have enhanced value over our sample period. For HY firms, on the other hand, the counterfactual short-rolling strategy omits potentially large rollover-failure costs in stressed primary markets. Shortening debt maturities is not a suggested strategy for these firms.

Contribution. This paper contributes to three strands of research. First, it extends the literature that quantifies the economic magnitude of corporate financing choices. Graham (2000) estimates the tax benefits of debt, and van Binsbergen, Graham, and Yang (2010) and Korteweg (2010) estimate the costs and net benefits associated with firms' leverage choices. These studies focus primarily on how much debt firms use. We study a distinct policy margin—how long firms borrow for and how long they fix their interest rate, conditional on issuing public debt. Our bond-level rolling counterfactual measures the realized financing cost of that policy after accounting for issuance costs, actual repayment timing, prevailing refinancing conditions, and changes in issuer credit quality. The resulting mea-

sure is not a complete NPV calculation; rather, it identifies the cost side of the maturity choice against which its benefits can be assessed.

Second, we contribute to the debt-maturity literature by supplying a direct empirical measure of the price of maturity. Theory shows how maturity affects adverse selection, debt overhang, liquidity and rollover risk, and endogenous default (Flannery, 1986; Diamond, 1991; He and Xiong, 2012; Diamond and He, 2014; He and Milbradt, 2016). Empirical work documents how firms choose, time, and stagger maturities and how refinancing exposure affects liquidity management and real decisions (Baker, Greenwood, and Wurgler, 2003; Almeida, Campello, Laranjeira, and Weisbenner, 2011; Custódio, Ferreira, and Laureano, 2013; Harford, Klasa, and Maxwell, 2014; Mian and Santos, 2018; Choi, Hackbarth, and Zechner, 2018). We complement this literature by measuring, bond by bond, the recurring realized premium associated with long maturity and relating that premium to the incidence and severity of the states against which long maturity provides protection. This supplies a quantitative benchmark for evaluating both rational and behavioral explanations of maturity choice.

Third, the paper separates funding maturity from interest-rate exposure. Survey, issuance, and hedging evidence shows that firms choose maturity, fixed-versus-floating exposure, and derivatives jointly (Graham and Harvey, 2001; Faulkender, 2005; Servaes and Tufano, 2006; Chava and Purnanandam, 2007). Long-term floating-rate debt preserves funding maturity while retaining short-rate exposure, whereas interest-rate swaps alter rate exposure without changing legal maturity. Combining these instruments with fixed-rate bonds and a yield-curve decomposition allows us to distinguish the price of avoiding repeated refinancing from the price of locking in base rates. This distinction adds a debt-structure dimension to the literature on financial conservatism, which has focused primarily on firms' leverage choices (Graham, 2000; Strebulaev and Yang, 2013).

2 Data and Sample Construction

2.1 Data Sources

Our analyses combine various datasets.

FISD. The LSEG Mergent Fixed Income Securities Database (FISD) provides bond-level characteristics: maturity date, coupon rate, coupon type (fixed or floating), offering yield, offering amount, gross spread, and callability indicators. We use the FISD call schedule file and the amount-outstanding file to identify prepayments and to compute effective tenors net of calls, partial calls, sinking fund redemptions, and maturity. For our

main bond-year analyses, we employ data from January 1997 through July 2025.

S&P Capital IQ. We use date-stamped entity-level S&P credit ratings from Capital IQ to identify the issuer's prevailing rating at each hypothetical rollover date in the counterfactual exercise.[3]

Compustat and ExecuComp. Firm-level financial data (total assets, total debt, interest expense, EBITDA, and market capitalization) are obtained from Compustat, and CEO/CFO equity holdings from ExecuComp.

Merrill Lynch/TRACE yield curves. We construct monthly corporate yield curves at five headline rating buckets (AAA/AA, A, BBB, BB, B). The pre-TRACE segment uses Merrill Lynch bond-day yields from January 1997 to January 2003; from February 2003 through July 2025, the benchmark uses FINRA TRACE Enhanced transaction data. Bond-days enter rating cells using the issuer's prevailing CIQ/S&P long-term rating at the trade date, following Dick-Nielsen, Feldhütter, and Lando (2012) and Bao, Pan, and Wang (2011). The risk-free benchmark is the Gürkaynak–Sack–Wright (GSW) Treasury zero-coupon curve (Gürkaynak, Sack, and Wright, 2007). The Merrill segment is level-spliced to TRACE using the adjacent 5.5-year overlap window, February 2003–July 2008, matching the approximate length of the Merrill-only extension. Online Appendix A.1 gives the full procedure.

Swap disclosures. We extract interest-rate swap exposures from 10-K footnotes discussing derivatives for all firms in our sample over January 1997–July 2025. We scrape the notional amounts, explicit pay-fixed/pay-floating direction, and the textual rationale. Online Appendix A.3 provides details.

2.2 Bond Issuance Sample

The rated fixed-rate bond sample is used for the main bond-level descriptive statistics and counterfactual wedge estimates. It is restricted to USD-denominated, non-convertible, non-putable, fixed-rate U.S. corporate bonds issued outside the financial sector. The sample also requires an issue rating from AAA through B and a contractual tenor longer than one year. We exclude floating-rate bonds (6.6% of the issues) from this sample so that the headline results measure the cost of long-term fixed-rate borrowing. Floating-rate bonds are studied in Section 4.3, where we compare the cost of floating-rate debt with

[3]After an acquisition, the target's old bonds can retain their issuer CUSIP even though the acquirer becomes the relevant source of ratings. Of the 3,390 bonds outstanding when their issuer was acquired, Capital IQ links 1,945 to a rated acquirer; the other 1,445 enter the first exclusion check. Removing those bonds changes the full-sample, IG, and HY means by at most 2.6 basis points; removing all acquisition-affected bonds changes them by at most 4.9 basis points. Online Appendix A.2 documents the links, rating changes, and both exclusions.

short-term fixed-rate counterfactuals. These screens share the usual U.S.-issuer, USD, nonconvertible, and fixed-coupon restrictions used in TRACE–FISD studies such as Dickerson, Mueller, and Robotti (2023). Unlike their return sample, we retain Rule 144A bonds because they are a material part of corporate issuance. We additionally exclude financial issuers.

Table 1 reports the sequential filter cascade from the raw FISD universe, which covers issues offered in January 1997–July 2025, to the fixed-rate issuance sample. The broad prepayment-adjusted counterfactual panel contains 21,110 unique bonds and 161,388 bond-year observations spanning January 1997–July 2025. 66.2% of these broad-panel bond-year observations correspond to bonds with an investment-grade rating at issuance. Requiring a valid 3-year comparison leaves the headline sample of 20,501 unique bonds and 152,531 bond-year observations.

[Table 1 about here.]

Table 2 reports data on the composition of the rated fixed-rate bond sample (21,110 U.S. nonfinancial bonds issued in January 1997–July 2025). The sample is entirely fixed-rate by construction: 61.3% of bonds are IG at issue, 91.8% are callable, 79.8% are senior unsecured, and 29.4% are issued in the Rule 144A market.[4]

[Table 2 about here.]

Table 3 reports summary statistics for various characteristics of the bonds in our sample: *Offering Amount*, *Tenor*, *Coupon Rate*, *Yield-to-Maturity at Issue*, and *Gross Spread*. *Offering Amount* is the par face value in millions of U.S. dollars. The median bond issue is $400 million. *Tenor* is the calendar-year difference between contractual maturity and offering date, averaging 11.4 years in our sample. *Coupon Rate* is the stated annual coupon. *Yield-to-Maturity at Issue* is the annualized yield that solves the present-value relation $P_0 = \sum_t c\,F/(1+y)^t + F/(1+y)^T$ at the offering date, evaluated against the offering price. *Gross Spread* is the underwriter's compensation as a percentage of par. Gross spread is available for 48.4% of fixed-rate issuance-sample bonds, and offering yield is available for 78.1%. Coverage is incomplete because Rule 144A placements, medium-term-note (MTN) issues, and other source-reporting gaps omit this information. Firms pay underwriters 70 basis points, on average, for each bond issue with a reported gross spread.

[Table 3 about here.]

[4]Rule 144A bonds are private debt securities that can be traded among large institutional investors without full SEC registration.

2.3 Issuance-Cost Model

Firms pay underwriting spreads to the investment banks that distribute their bonds to investors. To obtain projected costs in our counterfactual analysis, we estimate the following issuance cost regression model using the bond sample with available spread data (10,222 bond-level observations):

$$\log(g_i) = \alpha + \text{FE}_{r(i)} + \text{FE}_{t(i)} + \beta \log(T_i) + \eta_1 \text{Size}_i + \eta_2 \text{Size}_i^2 + \lambda_{r(i)} \log(T_i) + \delta' X_i + \varepsilon_i, \quad (1)$$

where g_i is gross spread as a percentage of par, winsorized at the 1st and 99th percentiles. $r(i)$ is the rating bucket[5], T_i is contractual tenor, $Size_i$ is the offering amount in units of \$100 million, X_i contains callable, senior-secured, and subordinated indicators, and the two FE terms are rating-bucket and issue-year fixed effects. For prediction, we refit this specification to log gross spread and multiply the exponentiated fitted value by the sample mean of $\exp(\widehat{\varepsilon}_i)$, following Duan (1983). This is one common log-to-level factor (1.049 in the current fit), not a residual draw for each bond. The log specification keeps predicted spreads positive without a floor. Online Appendix A.4 reports the same right-hand-side specifications in levels for coefficient interpretation and gives full prediction details.

2.4 Effective-Tenor Adjustment

91.8% of the bonds in our sample are callable, while 28.0% are actually prepaid, at least partially.[6] FISD records contractual call schedules and realized amount-outstanding histories. Online Appendix A.5 reports coverage of this information.

We use all principal repayments reported by FISD for each bond. Let t_i^0 be the offering date, $t_{i,k}$ a principal-repayment date, $\Delta P_{i,k}$ the principal repaid on that date, and $P_{i,0}$ the original offering amount. Effective tenor is the principal-weighted average time to repayment:

$$\widetilde{T}_i = \sum_k \frac{\Delta P_{i,k}}{P_{i,0}} (t_{i,k} - t_i^0), \qquad \sum_k \Delta P_{i,k} = P_{i,0}. \quad (2)$$

Elapsed time is measured in years using exact calendar dates. We assign each principal reduction to its event date and any remaining principal to the contractual-maturity date. A reported residual determines the reduction when valid; an action amount is used only as a flagged fallback, never as a second subtraction.

[5] We include six buckets: AAA, AA, A, BBB, BB, and B.

[6] We classify a bond as prepaid when FISD reports a fall in principal outstanding between issuance and contractual maturity.

Let $p^R_{i,k}$ be the price paid per dollar of principal retired on date $t_{i,k}$. The redemption premium and its annualized cost are

$$\rho_i \equiv \sum_k (p^R_{i,k} - 1)\frac{\Delta P_{i,k}}{P_{i,0}}, \qquad \xi_i \equiv \frac{\rho_i}{\widetilde{T}_i}.$$

For bonds held to maturity or repaid at par, $\rho_i = \xi_i = 0$. If a repayment has no usable price, the baseline ignores it but still uses any priced repayments for the same bond. The appendix reports a check that excludes bonds with an unpriced repayment.

Online Appendix A.5 contains additional details. The effective-tenor adjustment is what allows the bond-year wedge construction in Section 3 to use *realized* repayment timing rather than contractual maturity.

Table 4 reports the contractual tenure and the effective tenor. The mean principal-weighted effective tenor is 10.22 years, compared with 11.43 contractual years. For prepaid bonds, Panel B additionally reports contractual tenor, effective tenor, and observed action price. Tenor is reduced by 4.32 years on average, or 39.6% of contractual tenor.[7] Firms pay a mean redemption premium of 4.46% of par to redeem their debt.

[Table 4 about here.]

Table 5 reports early-retirement events in the FISD data. It shows which events have the balance and price needed for the adjustment and how many distinct bonds they cover.

[Table 5 about here.]

3 Annual Cost of Debt Methodology

3.1 Definition

The **annual cost of debt** (ACD) is defined as the annualized all-in borrowing rate that equates the issuer's net proceeds at issuance to the stream of coupon and principal payments over the bond's life, accounting for gross underwriting fees, call timing, and early-retirement events. Expressed as a fraction of par, let $p_{i,0}$ denote the offer price, g_i the gross underwriting spread, $\mathcal{C}_i$ the set of coupon-payment dates, $C_{i,m}$ the coupon cash flow at elapsed time $\tau^C_{i,m}$, and $\mathcal{R}_i$ the set of principal-retirement events. For each retirement event k, let $a_{i,k}$ be the share of original par retired, $\tau^R_{i,k}$ the elapsed time, and $p^R_{i,k}$ the redemption

[7] Online Appendix A.5 shows how each reported repayment ends exposure for the principal repaid on that date and compares the full effective- and contractual-tenor distributions.

price per dollar of retired par, including returned principal. The exact ACD is the internal rate r_i solving

$$p_{i,0} - g_i = \sum_{m \in \mathcal{C}_i} \frac{C_{i,m}}{(1+r_i)^{\tau^C_{i,m}}} + \sum_{k \in \mathcal{R}_i} \frac{p^R_{i,k} a_{i,k}}{(1+r_i)^{\tau^R_{i,k}}}, \qquad \sum_{k \in \mathcal{R}_i} a_{i,k} = 1. \tag{3}$$

The combined FISD history and current tables populate this equation with separate partial calls, sinking-fund payments, tenders, repurchases, and other early retirements. The remaining share is repaid at contractual maturity. A reported residual and an action amount are never counted twice.

3.2 Annual Cost Decomposition

For measurement and decomposition, we use a transparent undiscounted annualization of (3). Let $\rho_i = \sum_{k \in \mathcal{R}_i} (p^R_{i,k} - 1) a_{i,k}$ be the redemption premium defined in Section 2.4, where $p^R_{i,k}$ is the redemption price and $a_{i,k}$ is the share of principal retired at each retirement date. The remaining maturity payment is at par. Annualizing the underwriting spread and the realized redemption premium over the principal-weighted effective tenor gives

$$\mathrm{ACD}_i \approx \underbrace{y_i}_{\text{borrowing cost}} + \underbrace{\frac{g_i}{\widetilde{T}_i}}_{\text{ann. issuance cost}} + \underbrace{\xi_i}_{\text{ann. redemption premium}}, \tag{4}$$

where y_i is the observed borrowing rate and $\xi_i = \rho_i / \widetilde{T}_i$. This is an annual cash-cost decomposition, not the exact IRR or its Taylor expansion: at positive discount rates the exact yield effect also depends on cash-flow timing. Non-par redemptions are retained in ρ_i and enter (4) through ξ_i. If FISD records a principal reduction but no usable action price, that event is not applied and adds no redemption premium. Online Appendix A.6 gives the compact derivation. Equation (3) defines the exact IRR; the implemented annual measure and cost-wedge decomposition use Equation (4).

3.3 Counterfactual: k-Year Rolling

For each bond i with effective tenor $\widetilde{T}_i$, we construct a k-year *rolling* counterfactual. Starting on the offering date, the firm refinances every k years until the original bond is fully repaid. If less than k years remain, the final bond has a shorter maturity. Each new bond is priced on its issue date using the issuer's *prevailing* credit rating. The counterfactual ACD is the principal-exposure-weighted average borrowing cost, including the imputed

issuance cost at each refinancing. The k-year comparison is defined only for bonds whose contractual tenor T_i is strictly longer than k. Bonds with $T_i \leq k$ are already at or below the proposed counterfactual maturity; we exclude them from the k-specific wedge calculation rather than forcing a counterfactual rollover at a longer tenor. We report the 3-year counterfactual as the primary benchmark; it avoids putting all weight on the very short end of the curve, and corresponds to a realistic medium-term-note treasury strategy. But we also conduct analyses for counterfactual maturities of 1, 2, 4, and 5 years.[8]

Formal definitions. Index the counterfactual bonds by j. Bond j is issued on date $d_{i,j}$ and normally matures k calendar years later; the final bond matures when the original bond ends and may have a shorter tenor. Its annual all-in rate, $r^{\text{CF}}_{i,j}(k)$, is the yield for the issuer's rating and the new bond's tenor at $d_{i,j}$, plus the predicted issuance fee annualized over its maturity. Let $\omega_{i,j,t}$ be bond j's share of bond i's principal exposure in year t. The annual counterfactual ACD is

$$\text{ACD}^{\text{CF}}_{i,t}(k) \;=\; \sum_j \omega_{i,j,t}\, r^{\text{CF}}_{i,j}(k), \qquad \sum_j \omega_{i,j,t} = 1. \tag{5}$$

The weights use exact dates and balances. A year containing a rollover therefore averages the costs of the old and new bonds. Let $\omega_{i,j}$ be the same principal-exposure share over the bond's full realized life. The bond-level counterfactual ACD is

$$\text{ACD}^{\text{CF}}_{i}(k) \;=\; \sum_j \omega_{i,j}\, r^{\text{CF}}_{i,j}(k), \qquad \sum_j \omega_{i,j} = 1. \tag{6}$$

The counterfactual contains *no* redemption-premium term: each new bond is assumed held to maturity at par.

The cost wedge. The **cost wedge** measures the ex post cost of maturity choice. At the bond-year level,

$$\mathit{wedge}_{i,t}(k) \;=\; \text{ACD}^{\text{obs}}_{i,t} \;-\; \text{ACD}^{\text{CF}}_{i,t}(k), \tag{7}$$

where $\text{ACD}^{\text{obs}}_{i,t}$ is the observed ACD in year t from (4). At the bond level, the lifetime wedge is the exact principal-exposure-weighted average of (7), equivalently

$$\mathit{wedge}_{i}(k) \;=\; \overline{\text{ACD}^{\text{obs}}_{i}} \;-\; \text{ACD}^{\text{CF}}_{i}(k). \tag{8}$$

[8]The final counterfactual bond is often shorter than k so that it ends on the observed repayment date. As a check, Online Appendix A.6 prices it as a full k-year bond while leaving its repayment date and principal exposure unchanged; the average wedge changes little.

A formal lifetime value in Equations (6) and (8) is reported only after the bond is fully repaid by the sample end and every counterfactual bond with positive principal exposure has a valid yield and issuance-cost quote. Annual bond-year values remain available through the last observed year for bonds still outstanding at the sample end. These partial histories are labeled "to date" and are not included in the lifetime tables or figures.[9] A positive wedge indicates the issuer paid more for the long-term bond than rolling k-year bonds would have cost with perfect hindsight; the wedge is a realized, ex-post measure, not a forecast.

Counterfactual yield curves (Merrill/TRACE). The yield $y^{\mathrm{CF}}(k, b, t)$ at counterfactual tenor k, rating bucket $b \in \{\mathrm{AAA/AA, A, BBB, BB, B}\}$ in the headline sample, and calendar month t is read off a Merrill/TRACE corporate curve at monthly frequency. The curve combines the Gürkaynak–Sack–Wright calendar-month mean of daily Treasury zero rates with a duration-matched corporate spread, and is interpolated across its 2-, 3-, 5-, 7-, 10-, 15-, 20-, and 30-year nodes by PCHIP (Piecewise Cubic Hermite Interpolating Polynomial) within each rating-month cell.[10] The rating-month cells use the prevailing CIQ/S&P long-term rating at the Merrill or TRACE trade date. The Merrill segment before February 2003 is level-spliced with a rating-by-tenor median adjustment estimated over the adjacent 5.5-year overlap with TRACE over the period February 2003–July 2008. The full Merrill extension, TRACE cleaning, VWAP pricing, FISD merge, GSW benchmarking, cell-aggregation, PCHIP smoothing, counterfactual rebuild, validation, and splice-window sensitivity are in Online Appendix A.1. The monthly frequency caps the mismatch between the scheduled rollover date and the yield curve used at 15 days. We also estimate a CCC-and-below curve for diagnostics and for post-issuance downgrades of otherwise eligible bonds, but bonds rated CCC or below at issuance do not enter the headline wedge sample because their sparse transaction cells make the fitted curve noisy.[11]

Time-varying credit ratings (CIQ). The credit rating employed at each counterfactual rollover date is the issuer's *prevailing* S&P Local-Currency Long-Term rating at that time, rather than the rating at original issuance, so a firm that is downgraded between rollovers correctly faces a higher counterfactual borrowing rate at the next rollover date. We match

[9] After a recorded bankruptcy or default, we stop annual observations but do not treat the event as the bond's maturity.

[10] If the required tenor is shorter than the first curve point, which happens occasionally in the earlier years in our sample period, we use that first curve point.

[11] The baseline uses broad rating buckets, without the + or – ratings, so adjacent S&P notches can share the same yield curve. Online Appendix A.7 interpolates between neighboring bucket curves and the results remain virtually unchanged.

each bond to its issuer via the S&P Capital IQ entity-rating mapping. For each EntityID we build a chronologically sorted rating-history array; the rating at $d_{i,j}$ is then the most recent observation prior to $d_{i,j}$. Rating-revision frequencies during each bond's contractual life are reported in Online Appendix A.8, which shows that notch changes are rare for the AAA bonds and much more common below investment grade. A related concern that a counterfactual strategy of repeated short-term issuance might itself trigger downgrades, and thus undo part of the wedge through a *higher* borrowing cost rather than a lower one, is addressed in Online Appendix A.8.

3.4 Illustration

Table 6 illustrates our procedure using a 30-year fixed-rate bond rated BBB+ issued by Oklahoma Gas & Electric Co. in August 2004 and later called above par. Its realized components are the fixed coupon, amortized issuance cost, and amortized redemption premium. The 3-year rolling counterfactual replaces the realized life with three 3-year bonds and a shorter final bond, priced off the TRACE segment of the Merrill/TRACE yield curve at the issuer's contemporaneous rating (BBB+ at issuance, A- after upgrade).

Table 6 reports one row for each counterfactual bond. The realized bond ACD is constant because the bond is fixed-rate, while the counterfactual ACD resets at each three-year refinancing date. The last row reports the principal-exposure-weighted bond-level mean.

[Table 6 about here.]

4 Bond-Level Results

4.1 Headline Results

Figure 1 reports the bond-year-weighted mean cost wedge separately for IG and HY bonds against the n-year rolling counterfactual, where n ranges from 1 to 5. The wedge is positive in most calendar years, but turns slightly negative at the mid-2025 endpoint as the IG fixed-rate mean falls below zero. Across the full sample, the average is 1.85 pp/yr (1.96 pp/yr for IG; 1.64 pp/yr for HY), and is visibly countercyclical: it widens around 2001–2003, 2009–2012, and 2020, when short-term yields decline substantially relative to long-term yields.

[Figure 1 about here.]

Table 7 shows the aggregate average wedge for counterfactual rollover horizons of 1 to 5 years. The wedge generally declines with the rollover horizon: shorter rollover horizons widen the wedge because the counterfactual short rate is further below the locked-in long-term coupon, but the annual means are not uniformly positive: short-tenor counterfactuals can turn negative. Each k-specific column excludes bonds whose contractual tenor is less than or equal to k. Choosing $k = 3$ as a rollover window for the headline specification is therefore conservative; at $k = 1$ or $k = 2$ the implied cost savings relative to long maturities are even larger.

[Table 7 about here.]

Table 8 splits the bond-level wedge by credit rating, weighting each bond by its effective tenor after averaging it until the bond is fully repaid. A bond enters a tenor panel only if it is fully repaid by the sample end and every counterfactual bond with positive principal exposure is valid. The 3-year wedge rises from AAA and AA through BB and is somewhat lower for B-rated bonds. Relative to observed borrowing cost, however, the wedge is larger for IG than for HY bonds and generally falls as ratings worsen. This is because higher-rated bonds have lower borrowing costs.

[Table 8 about here.]

Table 9 reports formal bond-level lifetime averages by callability under the same full-repayment and complete-pricing rule.[12] The count column is the number of unique securities. In this bond-level split, callable bonds have higher observed borrowing costs and a larger average wedge for two reasons: firms pay a higher interest rate for the right to call, and a called bond may be redeemed above par. For the 3-year benchmark, the mean wedge is roughly 19% of observed cost for non-callable bonds and 34% for callable bonds.

[Table 9 about here.]

Table 10 reports the full distribution of the bond-year wedges for various subperiods. The mean and median wedges are both reliably positive in all subperiods and the share of negative-wedge bond-years is below half in every subsample, confirming that firms issuing long-term debt systematically pay above the short-rolling counterfactual. The GFC and rate-hike rows have higher negative-wedge shares than the adjacent low-rate expansion rows, but negative-wedge bond-years remain a minority in each era.

[12] When FISD records an early principal reduction, the effective-tenor adjustment stops counting the repaid principal on that date. We do not treat a missing announced-call price as evidence that the bond lacked a call option: FISD's announced-call file records call notices, not the full contract. Online Appendix A.5 explains the repayment rule and the available announced and realized prices.

[Table 10 about here.]

The split across time periods in Table 10 indicates that the wedge is positive during periods with very different interest rate environments. Figure 2 shows the distribution of bond-level lifetime wedges by issuance year separately for IG and HY bonds and for each counterfactual rollover tenor.

[Figure 2 about here.]

Table 11 provides the split between interest cost and issuance cost for the 3-year counterfactual, after first averaging each bond over its life. Because firms come to market much more frequently in the counterfactual scenario, the counterfactual issuance costs are generally higher than the actual issuance costs, but the difference is not substantial enough to outweigh the interest cost differential.

[Table 11 about here.]

The bond-level wedge is not driven by a few outlier issues. For each bond, Figure 3 plots the lifetime-average observed ACD against the corresponding k-year-rolling counterfactual ACD, separately at $k \in \{1, 2, 3, 5\}$. The mass of bonds sits above the 45° line at every counterfactual tenor, and the red star, which represents the equally weighted average, is uniformly above the line, confirming that the positive aggregate wedge is a pervasive cross-sectional pattern rather than the result of a few outliers.

[Figure 3 about here.]

We next examine annual observed and counterfactual ACD for the *outstanding stock* of bonds in each calendar year. Figure 4 shows the results for both investment-grade and high-yield issuers. The solid line is the observed ACD, while the dotted line represents the counterfactual. The IG wedge is positive in essentially every year of the sample, while the HY wedge is also positive, on average, but smaller and more cyclical.

[Figure 4 about here.]

Figure 5 displays the annual stock-based wedge by S&P rating bucket. For issuers with low credit risk, the wedge rises as ratings worsen from AAA to A, followed by a small decrease at the BBB level. For non-investment-grade issuers, the wedge is smaller. Thus, the wedge is generally higher for firms that face lower rollover risk. The post-2022 rate-hiking cycle compresses the wedge across all six ratings buckets as counterfactual short rates close in on locked-in long-term coupons.

[Figure 5 about here.]

4.2 A Cost-Wedge Decomposition

The baseline results show that long fixed-rate debt costs more than the short-rolling counterfactual, especially for IG bonds. The decomposition asks whether that gap was already built into the bond's issue yield or emerged only because financing conditions changed later. At issuance, a long bond can carry both a higher Treasury rate and a higher credit spread than a short bond. After issuance, the counterfactual refinancing cost changes as short Treasury rates, market credit spreads, and the issuer's rating change. Separating these two sources tells us whether the wedge is mainly the ex-ante price of long maturity or the ex-post result of a favorable path for short-term borrowing.

For compact notation, set the issuance date to 0 and let t denote a counterfactual rollover date. Write $r_t^{\mathrm{f}}(m)$ for the Treasury rate at tenor m, $s_t(R, m)$ for the credit spread at rating R and tenor m, and R_{it} for issuer i's rating at t. The observed issue yield is $y_i = r_0^{\mathrm{f}}(T_i) + s_0(R_i, T_i) + \nu_i$, while the counterfactual k-year rollover yield is $\widetilde{y}_{it}(k) = r_t^{\mathrm{f}}(k) + s_t(R_{it}, k)$. The yield difference telescopes exactly into five channels. Adding the observed-minus-counterfactual issuance-cost differential gives the all-in decomposition used in the paper:

$$
\begin{aligned}
w_{it}(k) \;=\; & \underbrace{\left[r_0^{\mathrm{f}}(T_i) - r_0^{\mathrm{f}}(k)\right]}_{\text{A: RF term premium}} + \underbrace{\left[s_0(R_i, T_i) - s_0(R_i, k)\right]}_{\text{B: Credit term premium}} \\
& + \underbrace{\left[r_0^{\mathrm{f}}(k) - r_t^{\mathrm{f}}(k)\right]}_{\text{C: Base-rate change}} + \underbrace{\left[s_0(R_i, k) - s_t(R_i, k)\right]}_{\text{D: Systematic spread change}} \\
& + \underbrace{\left[s_t(R_i, k) - s_t(R_{it}, k)\right]}_{\text{E: Rating migration}} + \underbrace{I_{it}(k)}_{\text{I: Issuance cost}} + \varepsilon_i. \qquad (9)
\end{aligned}
$$

The five yield-curve channel labels in (9) are:

- **A: Risk-free term premium.** The extra Treasury yield for fixing the base rate from tenor k through the bond's actual tenor T_i.
- **B: Credit term premium.** The extra credit spread for borrowing at T_i rather than k, holding the issuance rating fixed.
- **C: Base-rate change.** The change in the k-year Treasury rate between issuance and rollover. This channel is positive when the short Treasury rate is lower at rollover.
- **D: Systematic spread change.** The change in the k-year spread for the original rating. This channel is positive when market-wide spreads for that rating tighten before rollover.

- **E: Rating migration.** The additional spread change caused by the issuer's own rating movement. An upgrade contributes positively; a downgrade contributes negatively.[13]

The separate issuance-cost component is

$$I_{it}(k) \equiv \frac{g_i}{\widetilde{T}_i} - \iota_{it}^{\mathrm{CF}}(k),$$

where $\iota_{it}^{\mathrm{CF}}(k)$ is the annualized issuance cost of the counterfactual bonds active in that bond-year, weighted by principal exposure. It is negative when repeated short-term issuance costs more per year than the observed long-term issue. Channels A and B are known when the bond is sold: they are the market price of extending the Treasury rate and the credit spread from k to T_i. Channels C–E are realized later. They measure whether short Treasury rates fell, market spreads tightened, or the issuer's rating improved before a rollover. If A and B dominate, the firm paid most of the wedge up front for long maturity. If C–E dominate, the wedge is mainly an ex-post consequence of how rates, spreads, or ratings happened to evolve.

The five channels use the GSW Treasury and Merrill/TRACE corporate zero-coupon curves. The bond-specific residual is the actual issue yield minus the fitted yield for the same rating and maturity at issuance, a difference often called the new-issue premium, plus any annualized call-premium term. Issuance costs no longer enter that residual: they are reported separately as I. Figure 6 and Table 12 summarize the decomposition. For both IG and HY bonds, the risk-free term premium is the largest component: 1.03 pp/yr for IG and 0.99 pp/yr for HY. The credit term premium adds 0.32 and 0.30 pp/yr, respectively. Because both are fixed at issuance, the result is not mainly a by-product of unexpectedly low short rates later in the sample. For IG bonds, the three realized channels C, D, and E contribute only 0.21, 0.08, and -0.14 pp/yr, respectively. The issuance-cost differential contributes -0.07 pp/yr for IG and -0.09 pp/yr for HY; both are negative because the short-rolling strategy pays underwriting costs more often.

[Figure 6 about here.]

[Table 12 about here.]

Figure 7 shows the same decomposition by calendar year. The issuance-date term premiums remain the largest components, while the base-rate channel explains most of

[13]The decomposition restricts the sample to fixed-rate bonds with contractual tenor between 2 and 30 years, matching the constructed counterfactual yield curve range. Bonds with longer maturities are dropped because sparse coverage at those tenors prevents reliable yield-curve construction.

the year-to-year movement. It is positive in the post-GFC low-rate period, when rollover Treasury rates were below their issuance-date levels, and negative in 2022–2025, when short rates rose. The systematic-spread and rating-migration channels remain economically small.

[Figure 7 about here.]

4.3 Fixed versus Floating-Rate Bonds

Long fixed-rate debt bundles two choices: the firm secures funding for a long period and fixes its benchmark interest rate for that period. A long floating-rate bond separates those choices. It preserves long-term funding—so the firm need not refinance whenever the coupon resets—but its coupon moves with a short-term benchmark. Comparing fixed-rate and floating-rate bonds with the same short fixed-rate rollover benchmark therefore helps distinguish the price of long funding from the price of rate lock-in. A large positive wedge for both types would be consistent with long funding and rollover protection carrying the central premium. A large wedge only for fixed-rate bonds would instead point to fixing the benchmark rate. This comparison does not imply that rollover protection has no value; it asks whether that protection can explain the large realized IG fixed-rate wedge. Because firms choose which instrument to issue, this is mechanism evidence rather than a causal estimate of the effect of choosing a floater.

Floating rate bond sample. The fixed-versus-floating comparison uses a separate floating-rate bond sample. The floating-rate coupon-history sample starts from the cleaned FISD issuance universe and applies the same issuer, bond-type, rating, and outlier filters as the fixed-rate issuance sample, but keeps floating-rate coupon bonds instead of fixed-rate coupon bonds. This produces 2,214 floating-rate bonds.[14] For each floating-rate bond, we parse the FISD coupon-formula text to extract the reference index, the add-on spread, and any floor or cap. The realized annual coupon is then computed at reset dates from historical benchmark-rate series. Table 13 summarizes the floating-rate coupon-history sample. We have sufficient information to compute the annual coupons for 1,508 of the 2,214 floating-rate bonds (68.1% coverage), yielding 6,031 bond-year observations over January 1997–July 2025. The dominant benchmark is 3-month LIBOR, followed by 1-month LIBOR, 6-month LIBOR, and SOFR; few bond coupons are linked

[14] For a floating-rate bond, observed borrowing cost is its realized coupon in each year. We do not estimate a separate underwriting-fee model from the much smaller floating-rate sample. When the observed fee is missing, and for every hypothetical rollover fee, we apply the fixed-rate issuance-cost model; Online Appendix A.4 reports its sample, specification, and validation.

to the Prime rate, the Federal funds rate, Constant Maturity Treasury (CMT) rates, and commercial paper.

[Table 13 about here.]

Table 14 makes this comparison for IG bonds. Panel A reports the fixed–floating wedge difference across rollover tenors, and Panel B shows the 3-year observed and counterfactual cost components. IG fixed-rate bonds have a large positive wedge, whereas IG floaters have a slightly negative wedge against the same fixed-rate corporate rollover benchmark.

[Table 14 about here.]

The economic contrast is direct. A floater protects the issuer from returning to the bond market at every reset date and keeps its contractual credit spread in place, but it does not protect the issuer from movements in the short benchmark rate. If repeated market access, wider future credit spreads, or rating deterioration were the main source of the IG fixed-rate wedge, IG floaters should also have a sizable positive wedge. They do not.

Table 15 reports the corresponding decomposition. For this fixed-versus-floating comparison, we use the same five (A through E) yield-curve channels as Equation (9) and add one floating-rate component:

$$w_{it}(k) = A_i + B_i + C_{it} + D_{it} + E_{it} + F_{it} + I_{it}(k) + \varepsilon_i, \tag{10}$$

$$F_{it} \equiv \begin{cases} 0, & \text{fixed-rate bonds,} \\ r_{it}^{\text{float}} - r_{i0}^{\text{float}}, & \text{floating-rate bonds.} \end{cases}$$

For fixed-rate bonds, component F is zero. For floaters, it is the change from the first observed coupon, r_{i0}^{float}, to the coupon at the counterfactual rollover date. It therefore captures the part of the floater's cost that moves when its benchmark resets. I retains the issuance-cost definition above. Observed gross spread is reported in 100.0% of the 1,107 bond-years, covering 268 floaters; their hypothetical rollover fees use the fixed-rate model because the floater sample is too small to estimate a separate one. For both rate types, the residual contains primary–secondary pricing and call-premium differences left after the named components. For floaters, it also contains the first coupon's discount relative to the fitted fixed-rate yield for the same maturity and credit risk. The later coupon path enters F, not the residual.

The results reinforce the simple comparison. For IG fixed-rate bonds, the 1.72 pp/yr wedge is dominated by the risk-free term premium (1.03 pp) and credit term premium (0.32 pp). For IG floaters, the wedge is -0.51 pp/yr, largely because the residual is -0.68 pp. Their issuance-cost differential is only -0.05 pp/yr. Even excluding the residual, channels A–F and I sum to only 0.16 pp/yr. Table 15 also reports HY bonds, but the HY floating-rate sample is too small for meaningful inference.

Overall, long funding without a fixed coupon does not carry a comparable IG wedge. The fixed-versus-floating evidence therefore points to benchmark-rate lock-in, rather than market-access insurance, as the main priced difference in long-term IG debt.

[Table 15 about here.]

5 Firm-Level Results

Up to this point, the analysis has been conducted at the individual bond level. We now aggregate bonds to the firm-year level for three reasons. First, the aggregation measures total dollar savings across all of a firm's outstanding bonds. Second, when the wedge is negative, it lets us investigate whether the firm has the funds available to cover the extra interest cost. This test is essential because firms may issue long-term fixed-rate debt to avoid interest-rate volatility. If firms in the counterfactual cannot pay the higher interest expense when rates rise, that would justify locking in long-term fixed rates. Third, the firm-level panel lets us adjust for interest-rate swaps that change firms' effective rate exposure.

5.1 Methodology

The firm-level analysis maps each FISD-listed bond to its Compustat issuer. Each cost-of-debt measure is weighted by within-year average principal across all of a firm's active bonds, and dollar wedges are summed across bonds to produce the firm-year dollar cost.

For firm f in year t, the firm-level wedge is

$$\text{Wedge}_{f,t}^{(k)} = \frac{\sum_{i \in \mathcal{B}_{f,t}} N_{i,t}\, w_{i,t}^{(k)}}{\sum_{i \in \mathcal{B}_{f,t}} N_{i,t}}, \tag{11}$$

where $\mathcal{B}_{f,t}$ is the firm's active bond set and $N_{i,t}$ is bond i's average outstanding notional during that year, using exact active days and the principal outstanding during each part of the year. The term $w_{i,t}^{(k)}$ is the bond-year wedge against the k-year rolling counterfactual.

In dollar calculations, wedge rates are written directly as percent rates; for example, $w = 1.5\%$ means 1.5 percent of notional. To convert the rate into dollars, we account for both the principal outstanding and how long it remains outstanding. For bond i in year t, $E_{i,t}$ is the sum of each reported principal balance multiplied by its days outstanding and divided by 365.25. The firm-year dollar wedge is

$$\text{DollarWedge}_{f,t}^{(k)} = \sum_{i \in \mathcal{B}_{f,t}} w_{i,t}^{(k)} \, E_{i,t}. \tag{12}$$

Aggregate dollar magnitudes sum this annual object across firm-years and convert nominal dollars to 2025 dollars using annual-average headline CPI from FRED (CPIAUCSL). The affordability tests use the same firm-year dollar wedge. On a flow basis, the negative-wedge need is

$$\text{Need}_{f,t,k} = \max\{0, -\text{DollarWedge}_{f,t}^{(k)}\}. \tag{13}$$

Coverage is assessed sequentially. In each year, the firm first eliminates share repurchases (Compustat PRSTKC), then uses cash and short-term investments (CHE), then cuts dividends (DVT), and finally it uses any unused savings accumulated from prior positive-wedge years. The prior-savings stock uses only positive wedge years that occur before the shortfall year and the stock is depleted when prior savings are exhausted.

5.2 Headline Results

Table 16 shows summary statistics for the firm-year panel after aggregating the bond-year data to the firm level. This panel also includes the Compustat financial measures used in the paper's firm-level analyses, two CRSP equity-risk measures, and the year-end S&P rating from the S&P Capital IQ credit rating file.

[Table 16 about here.]

The average firm-year has 5.1 bonds outstanding, totaling $2.3 billion. Its mean 3-year wedge is 1.72 percentage points per year, equivalent to $41 million annually. Additional summary statistics are reported in Table 16.

Panel A of Table 17 shows the firm-level wedge across counterfactual tenors $k \in \{1, 2, 3, 4, 5\}$ years. For each counterfactual horizon we report the firm-year sample size, mean observed cost of debt, mean and median wedge, the wedge as a share of the observed cost of debt, the cross-firm-year standard deviation of the wedge, and the annual wedge in U.S. dollars (based on each reported principal balance and its exact days outstanding, using days/365.25; in millions per firm-year and summed across firm-years

to billions). The 2-year counterfactual delivers the largest mean wedge, but the wedge remains substantial for 3- to 5-year counterfactuals. We continue employing the 3-year counterfactual as the base case. Panel B reports nominal and inflation-adjusted totals for all bonds, IG bonds, Compustat-matched firms, and IG-rated firms.

[Table 17 about here.]

The dollar magnitudes are large. In 2025 dollars, the annual 3-year firm-year wedges sum to $1.2 trillion ($943.3 billion without inflation adjustment) across Compustat-matched firm-years; the corresponding number across all bonds (including those not matched with Compustat) is $1.5 trillion. Table 17 reconciles these amounts with their nominal values and IG subsets. On an annual flow basis, the same 3-year benchmark equals 2.4% of EBIT for the median positive-EBIT IG firm-year, 1.7% of EBITDA for the median positive-EBITDA IG firm-year, and 15.8% of annual interest expense for the median positive-interest-expense IG firm-year.

Table 18 reports the annual 3-year firm-year wedge, splitting firms by the weighted average of the contractual tenor of their outstanding bond portfolios. The wedge is positive in all three tenor bins, and it is largest for firm-years whose bonds have original maturities beyond 10 years.

[Table 18 about here.]

5.3 Interest Rate Swap Adjustment

A number of firms alter their interest rate exposure by entering into interest rate swaps. Pay-floating, receive-fixed swaps convert fixed-rate debt into floating-rate exposure; thus, the bond-level fixed-rate wedge reported above overstates the actual wedge for these firms. Pay-fixed, receive-floating swaps do the opposite: they convert floating-rate exposure into fixed exposure, such that the bond-only wedge understates the issuer's true fixed-rate lock-in.

We gather data on interest rate swaps outstanding at the firm level and adjust the wedge computation accordingly. The adjusted firm-year wedge is

$$\text{Wedge}_{f,t}^{\text{swap}-\text{adj}} = \text{Wedge}_{f,t} - \phi_{f,t}\left(y_{f,t}^{\text{fix}} - r_t^{\text{float}}\right), \tag{14}$$

where $\phi_{f,t}$ is the fraction of debt swapped to floating and r_t^{float} is the matched-tenor floating-leg reference rate.

We determine swap direction from 10-K footnote disclosures; Online Appendix A.3 describes the full construction.[15] Across the full sample, 1,017 of 11,153 firm-years with observed long-bond balances disclose fixed-to-floating interest-rate swaps. The aggregate fixed-to-floating notional is $1,186 billion, or 5.14% of aggregate FISD public-bond balances with more than five years remaining. Among the 2,339 firm-years that report swap direction, the net fixed-to-floating notional is $-201.5 billion, or -0.87% of that same aggregate long-bond denominator. Thus, fixed-to-floating and floating-to-fixed positions largely net out in aggregate.

Pay-floating stress test. As a conservative stress test, we set ϕ equal to the disclosed gross fixed-to-floating swap notional divided by debt with more than five years remaining.[16] Table 19 shows how much the headline fixed-rate bond-year mean and median wedge would fall if a share ϕ of fixed-rate debt were synthetically converted to floating through pay-floating swaps. At the disclosed gross fixed-to-floating share from the 10-K disclosures, $\phi = 5.14\%$, the table's interpolation implies an adjusted mean wedge of about 1.76 percentage points per year and an adjusted median of about 1.60, down from the unadjusted headline values of 1.85 and 1.70. This evidence indicates that much of the wedge remains after adjusting for interest-rate swaps.

[Table 19 about here.]

5.4 Affordability of the Counterfactual Strategy

A natural concern with the rolling counterfactual is that short-rate spikes occasionally make a given year's counterfactual payment *higher* than the locked-in fixed rate, and that firms may lack the liquidity or other resources to absorb the shortfall. This may be one of the reasons why firms issue longer-dated debt in the first place.

As discussed in Section 5.1, we employ various sources of liquidity to address cash shortfalls. If need be, the firm first eliminates share repurchases (Compustat PRSTKC), then uses cash and short-term investments (CHE), then cuts dividends (DVT), and finally it uses any unused savings accumulated from prior positive-wedge years.

[Table 20 about here.]

[15] Disclosure of material derivative positions is mandated under ASC 815 (*Derivatives and Hedging*), codified at `https://asc.fasb.org/815`, which requires recognition at fair value on the balance sheet and tabular footnote disclosure of notional amounts, fair values, and hedge designations. Therefore, firms that do not disclose interest-rate-swap positions in a given year likely held no material swap exposure.

[16] This makes the conservative assumption that all disclosed fixed-to-floating swaps apply to this long-term debt and cover all of its remaining interest payments.

Table 20 analyzes the shortfalls and the affected firms' ability to cover them. For each counterfactual tenor and separately for IG and HY firms, we report the number of positive-need firm-years, mean and median need, and the share of those firm-years fully covered after each cumulative liquidity step. For example, for the 3-year counterfactual, 795 of 11,843 firm-years have a positive cash need, with a mean shortfall of $40.4 million and a median shortfall of $5.4 million.

The coverage columns report the share of positive-need firm-years that each cumulative set of liquidity sources fully covers. For the 3-year counterfactual, cutting share repurchases alone fully covers 61.5% of positive-need IG firm-years; adding cash raises coverage to 94.8%. Adding dividend cuts and savings from prior positive-wedge years raises coverage to 99.1%.[17] Overall, IG firms are better positioned than HY firms to address the shortfall at every counterfactual tenor. Moreover, the fraction of IG firm-years having any shortfall is also much smaller than the fraction of HY firm-years. The residual uncovered set is small for IG firms: only 7 IG firm-years (0.9%) remain uncovered after all four liquidity buffers are used up. For HY firms, on the other hand, 246 firm-years (10.2%) remain uncovered.

[Table 21 about here.]

Table 21 uses the above analysis to compute a direct interest-coverage distress indicator. A firm-year is classified as distressed if EBIT is below annual interest expense. Under the counterfactual, interest expense is adjusted by subtracting the firm-year dollar wedge, so positive-wedge years lower the counterfactual interest burden and negative-wedge years raise it. The good/bad split in the table is therefore the counterfactual cash-flow state: good years are counterfactual-savings years, while bad years are counterfactual-cost-spike years. For IG firms, the counterfactual does not increase distress frequency in the full sample or in either state. Across all firms, the higher bad-state distress rate is more than offset by the lower good-state distress rate; the bad-state increase is entirely concentrated among HY firms. The few uncovered observations do not necessarily imply that these firms are severely distressed. The four-buffer exercise deliberately omits current operating cash flow, undrawn revolvers and committed bank lines, working-capital release, capex deferral, and short-term bank or commercial-paper borrowing. Adding those margins would make the test less conservative. The right interpretation, therefore, is that the residual cases identify years requiring outside or operational liquidity, not that

[17] This computation does not take into account fixed-to-floating swaps. Under the gross fixed-to-floating swap adjustment, mean need falls to $38.3 million but the final progressive coverage remains the same at 99.1%.

these cases lead with certainty to distress. Affordability does not appear to be the binding margin, which raises the question of why firms pay the wedge at all. This question is addressed in Section 6.

5.5 Robustness: Ratings under the Short-Rolling Counterfactual

Another concern with the bond-level and firm-level wedges computed previously is that adopting the short-rolling counterfactual strategy raises the firm's level of short-term debt, which could worsen the borrower's credit profile even when keeping the level of debt unchanged. That is, in addition to taking into account the level of debt, ratings agencies may also pay attention to its maturity profile. Higher levels of short-term debt could trigger a rating downgrade, which, in turn, would raise the counterfactual borrowing cost and shrink the wedge, potentially overturning our results. We address this concern in Online Appendix A.8 by estimating the ratings-determinants specification of Baghai, Servaes, and Tamayo (2014), augmented with the Short-Debt Share to quantify how ratings respond to a higher reliance on short-term debt; the implied counterfactual rating shift is then translated into a per-notch borrowing-cost increment that is netted out of the headline wedge. The preferred firm-fixed-effect specification does not predict a downgrade. Mean and median rating changes are slightly negative, indicating a small predicted improvement, and no firm-year worsens by more than half a notch. The signed borrowing- and issuance-cost adjustments are therefore slightly negative, so the adjusted wedge is marginally larger. This sensitivity does not support a rating-deterioration explanation for the cost-of-maturity result, although it should not be read as a causal benefit of short-term debt.

6 Explaining Debt Maturity Conservatism

The results of the previous sections suggest that firms leave a lot of money on the table by not using short-term debt in their capital structure more diligently. Long fixed-rate debt costs firms more than a short-rolling counterfactual, and the gap is largest for IG firms where standard rollover-risk logic is the weakest. The mean annual wedge is 1.85 percentage points per year for all firms, increasing to 1.96 percentage points per year for IG firms. Thus, while IG issuers have the lowest rollover risk, they pay the highest cost of maturity. Moreover, when we compare floating rate debt to the counterfactual of rolling over fixed rate debt with shorter maturities, the price wedge disappears or becomes negative. This suggests that by issuing long-dated *fixed-rate* debt, which is more expensive,

firms are not paying primarily to avoid rollover risk; rather, they pay for interest rate stability. For HY firms, the rollover risk argument is more compelling because they did lose access to capital markets during the GFC.

The goal of this section is to understand why firms, and IG firms in particular, purchase rollover risk insurance, while they never lost access to capital markets in the first place, and why they may value interest rate stability. For HY firms, the narrative is more nuanced, but it is still useful to compute how much they benefit from this insurance.

We argue that this debt conservatism could be driven by both shareholder preferences and managerial preferences.

6.1 Shareholder preferences

One possibility is that rollover risk is difficult to estimate ex-ante. If managers overestimate the possibility that they will not be able to roll over their debt in any given year and/or overestimate the potential loss associated with a failure to roll over their debt, they would be willing to pay a very high price to avoid such a state, and they would do so in the best interest of the shareholders. In reality, IG firms never lost access to the primary market over our sample period. During the nine-month GFC stress window (July 2008–March 2009), average monthly IG issuance rose from $11.7 billion in the 42-month pre-GFC period (January 2005–June 2008) to $20.2 billion, while average monthly HY issuance fell from $8.3 billion to $4.0 billion. During March 2020, average monthly IG issuance rose from $36.6 billion in the 26-month pre-COVID period (January 2018–February 2020) to $155.7 billion, while HY issuance fell from $11.1 billion to $3.6 billion.

[Figure 8 about here.]

Moreover, even though IG firms never lost access, some IG firms may worry about getting downgraded to HY–thereby becoming so called fallen angels–exactly around the time when HY markets contract severely or shut down, which would still prevent them from rolling over their debt. Indeed, during our sample period, 10.6% of bonds that were rated IG at issuance were downgraded to high yield. In addition, such downgrades are more likely during periods of market stress.

A second possibility is that managers are worried about having to roll over debt when interest rates spike to very high levels. Subsequently, they may not be able to service this debt, leading to financial distress. The computed wedge already takes into account the direct cost of those spikes. In addition, Table 20 shows that only 7 of 11,843 eligible IG firm-years remain uncovered after using up cash holdings, cutting share repurchases and

dividends, and exhausting prior cash savings. Thus the relevant unconditional annual probability is 0.059%, not the conditional percentage among firm-years that happen to have a positive need. For HY firms, 246 of 11,204 eligible firm-years remain uncovered, an unconditional annual probability of 2.196%.

The above figures allow for a back-of-the-envelope calculation of what distress probability would justify the wedge observed in the data. Let W denote the annual cost of choosing long fixed-rate public debt instead of the short-rolling counterfactual, expressed as a fraction of debt. Let p be the objective probability, over a three-year counterfactual cycle, of either a rollover-failure state or an inability to service debt after an interest-rate spike; let L be the state-contingent loss as a fraction of firm value; and let l be long-term debt as a fraction of firm value. Assuming a three-year counterfactual cycle, total interest savings are $3W$. The risk-neutral shareholder rule is

$$3W \leq p \cdot L \cdot \frac{1}{l}. \tag{15}$$

We use long-term-debt shares of 26.3% for IG firms and 46.1% for HY firms, together with a distress cost of 20% of firm value. This deliberately severe benchmark lies between the 10–20% range for distressed highly leveraged transactions in Andrade and Kaplan (1998) and the 21.7% mean market-implied default cost in Davydenko, Strebulaev, and Zhao (2012). An annual wedge of 1.96 percentage points for IG firms then requires a three-year distress probability of 7.72%. The 1.64 percentage-point HY wedge requires 11.35%.

Based on the discussion above, the historical distress proxy combines two probabilities: (1) inability to roll over counterfactual debt in a crisis and (2) inability to service the counterfactual debt after an interest-rate spike. To compute the first probability for IG firms, we collapse 25,531 issuer-mapped bond rollovers to 9,904 rollover issuer-years. Of these, 41 occurred after the issuer had been downgraded to HY and while HY markets were severely stressed (July 2008–March 2009 and March 2020), a probability of 0.414%. The issuer-year denominator is essential because one issuer-year can contain several bonds or rollover dates. As discussed above, the second probability is 0.059% annually for IG firms, which becomes 0.177% over three independent years. Adding up the two probabilities yields 0.591%.

For HY firms, 224 of 7,391 rollover issuer-years occur during HY market stress, an access probability of 3.031%. Combining this with the three-year liquidity probability of 6.443% gives 9.279%. Thus, the historical probability needed to justify the IG wedge is 13.1 times what materialized under these screens, which appears implausibly high. For

HY firms, the required probability is only 1.2 times the historical proxy, so the rollover-insurance explanation is much more reasonable.

One additional factor at play could be risk aversion due the ownership structure of the firm. The above computations assume risk neutrality. However, if decision makers are risk averse, perhaps because shareholders are poorly diversified, a gap could open between the price they are willing to pay to avoid distress and the expected loss. External benchmarks provide useful scales. Catastrophe-bond evidence puts market prices near 2.7 times expected loss over the long run (Lane and Mahul, 2008); homeowner deductible choices imply that households often pay roughly 5.0 times expected value for marginal insurance (Sydnor, 2010).

A more direct behavioral scale comes from probability weighting. The loss-domain weighting function estimated by Tversky and Kahneman (1992),

$$w^{-}(p) = \frac{p^{0.69}}{[p^{0.69} + (1-p)^{0.69}]^{1/0.69}},$$

overweights small objective probabilities. This function maps the 0.591% probability into a 2.8% decision weight, still far below the required 7.72% distress probability to justify the observed wedge for IG firms. Thus, both the probability-weighting benchmark and the observed insurance-price benchmarks remain well below the 13.1-times IG cost-to-loss ratio.

In sum, for IG firms, the price paid to lock in long-term fixed-rate debt appears extremely high if default and rollover risk are the only concerns. For HY firms, on the other hand, the price paid appears more reasonable.

6.2 Managerial preferences

It is also possible that, in addition to avoiding distress costs, managers obtain private benefits from issuing longer-term fixed-rate debt rather than rolling over short-term debt. This is what we examine in this subsection. If we write this management-side private payoff, expressed as a fraction of debt value, as Q, the decision becomes

$$3W \leq p \cdot L \cdot \frac{1}{l} + Q. \tag{16}$$

Here Q captures the private payoff managers receive from avoiding both funding stress and volatile interest expenses.

We start by documenting that, not surprisingly, the time-series standard deviation of borrowing costs is much higher in the counterfactual scenario than in what actually

materialized. Table 22 reports the results. The average within-firm standard deviation of IG borrowing costs is 0.71 percentage points under observed debt choices and 1.64 percentage points under the short-rolling counterfactual, a volatility ratio of 2.31. For HY firms, the same comparison is 0.65 versus 1.94 percentage points, a ratio of 2.99. The fixed-versus-floating evidence in Section 4.3 shows that issuers pay for this rate stability, rather than merely for long legal maturity.

[Table 22 about here.]

There are three reasons why managers may derive private benefits from the stability provided by long-term fixed-rate debt. First, long-dated debt reduces the time and effort required to access the bond market regularly, as well as the monitoring associated with short-term rollovers (Datta, Iskandar-Datta, and Raman, 2005). Second, superiors may view the additional interest-expense volatility from frequent rollovers as poor performance, particularly when rollovers occur during market stress and at high interest rates. Third, executives may have much of their own wealth tied to the firm and therefore value the lower total volatility obtained by locking in long-term rates.

The first and third explanations have opposite predictions regarding the relation between ownership and bond maturity. If monitoring and effort are central, greater managerial ownership should be associated with shorter debt. If undiversified wealth is central, managers with more skin in the game should issue longer-term debt to reduce interest-rate volatility. Prior evidence from Datta, Iskandar-Datta, and Raman (2005) supports the first interpretation. Table 23 repeats their exercise for CEO and CFO ownership. In model (5), which includes CFO ownership, CFO tenure, and CEO ownership jointly, the CFO-ownership coefficient is negative and significant. A one-standard-deviation increase in CFO ownership is associated with a 2.2% decline in the maturity of newly issued debt. This reduced-form association is not a causal estimate of CFO preferences, but it suggests that CFOs with more skin in the game are more willing to accept interest-rate volatility when shorter debt lowers the cost of maturity. If firms issued longer debt because stable interest payments protected shareholder value, this negative relation would not be expected. Instead, the relation is more consistent with managers attaching private value to smooth financing conditions and shareholders not fully internalizing how much of the wedge buys that stability rather than rollover-loss protection. CFO tenure is positively associated with maturity, consistent with longer-tenured CFOs placing more value on smooth financing conditions.

[Table 23 about here.]

6.3 Relation to Maturity Timing

Our findings complement the maturity-timing literature. Baker, Greenwood, and Wurgler (2003) show that firms adjust maturity with debt-market conditions, and Greenwood, Hanson, and Stein (2010) formalize a gap-filling theory in which the corporate sector changes maturity in response to Treasury supply. Our evidence does not claim that firms ignore the yield curve. It shows that the adjustment is incomplete on the cost-relevant margin: among firms that issue long fixed-rate bonds, especially IG firms, the realized cost of maturity remains large and persistent. This incomplete adjustment is consistent with the cost-wedge decomposition above: the issuer pays for rate certainty while the long-duration investor earns an ex-post term premium.

Debt-maturity conservatism is therefore a residual interpretation, but a disciplined one. It respects the rational insurance benchmark, allows catastrophic rollover failure to be costly, and recognizes that the term premium is a real price paid to long-duration investors. The evidence nevertheless points to a costly issuer-side demand for stability: low-rollover-risk firms repeatedly pay for smooth interest expense and financing certainty, while the realized tail events that would justify the purchase are rare and concentrated. Separating the exact shares due to shareholder probability weighting, managerial probability weighting, and managerial stability value requires additional quasi-experimental variation. The contribution here is to measure the price of that conservatism and show why a pure rollover-insurance model is too weak to explain it.

7 Conclusion

We measure the realized cost of choosing long-term corporate bonds over a counterfactual short-rolling strategy, bond by bond, for a comprehensive sample of U.S. corporate bonds issued over January 1997–July 2025. The mean annual cost wedge is 1.85 percentage points per year across all bonds. It is large for both IG and HY firms: the IG bond-year wedge is 1.96 pp/yr, versus 1.64 pp/yr for high-yield issuers. We interpret the HY comparison cautiously because the counterfactual assumes uninterrupted market access at the prevailing rated yield; that assumption is reasonable for IG issuers, who retained primary-market access through both the 2008–2009 and 2020 disruptions, but not for HY issuers. The interest-cost savings associated with shortening debt maturities are therefore more evident for IG firms.

The cost-wedge decomposition attributes most of the persistent IG wedge to the market price of maturity at issuance, i.e., the risk-free term-premium channel, with a smaller

credit term-premium component. The base-rate-change channel explains part of the annual movement: it is positive after rate declines and negative after rate increases, but its full-sample contribution for IG firms is only 0.21 pp/yr; issuer-specific rating migration is small. IG floating-rate bonds—which share short-benchmark exposure with rolling debt but eliminate repeated refinancing at uncertain credit spreads or after rating changes—have a wedge that is zero or even negative. Fallen-angel evidence points in the same direction: only 10.6% of IG-at-issuance bonds ever cross to HY during their life, and only 0.414% of IG rollover issuer-years coincide with both an HY rating and a stressed HY market.

Adjusting the wedge for interest-rate swaps, or for the possibility that issuing shorter-maturity debt would itself lower credit ratings and increase credit spreads, has only a small effect on our results.

We explore various mechanisms for the observed debt-issuance strategy. It is unlikely that managers are simply overestimating rollover risk or the inability to service debt when interest rates spike. Instead, our results are better explained by private benefits that managers derive from having longer-dated debt. Consistent with this explanation, CFOs with greater ownership issue shorter-dated debt. Exploring these private benefits in more detail is an interesting direction for future work.

References

Almeida, Heitor, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner, 2011, Corporate debt maturity and the real effects of the 2007 credit crisis, *Critical Finance Review* 1, 3–58.

Andrade, Gregor, and Steven N. Kaplan, 1998, How costly is financial (not economic) distress? evidence from highly leveraged transactions that became distressed, *Journal of Finance* 53, 1443–1493.

Badoer, Dominique C., and Christopher M. James, 2016, The determinants of long-term corporate debt issuances, *Journal of Finance* 71, 457–492.

Baghai, Ramin P., Henri Servaes, and Ane Tamayo, 2014, Have rating agencies become more conservative? Implications for capital structure and debt pricing, *Journal of Finance* 69, 1961–2005.

Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler, 2003, The maturity of debt issues and predictable variation in bond returns, *Journal of Financial Economics* 70, 261–291.

Banerjee, Ryan, Julián Caballero, Enisse Kharroubi, Renée Spigt, and Egon Zakrajšek, December 2023, Interest rate risk of non-financial firms: Who hedges and does it help? BIS Bulletin 81, Bank for International Settlements. URL https://www.bis.org/publ/bisbull81.htm.

Bao, Jack, Jun Pan, and Jiang Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance* 66, 911–946.

Barclay, Michael J., and Clifford W. Smith, Jr., 1995, The maturity structure of corporate debt, *Journal of Finance* 50, 609–631.

Bertrand, Marianne, and Sendhil Mullainathan, 2003, Enjoying the quiet life? corporate governance and managerial preferences, *Journal of Political Economy* 111, 1043–1075.

Bodnar, Gordon M., Gregory S. Hayt, and Richard C. Marston, 1998, 1998 wharton survey of financial risk management by US non-financial firms, *Financial Management* 27, 70–91.

Brick, Ivan E., and S. Abraham Ravid, 1985, On the relevance of debt maturity structure, *Journal of Finance* 40, 1423–1437.

Campello, Murillo, John R. Graham, and Campbell R. Harvey, 2010, The real effects of financial constraints: Evidence from a financial crisis, *Journal of Financial Economics* 97, 470–487.

Chava, Sudheer, and Amiyatosh Purnanandam, 2007, Determinants of the floating-to-fixed rate debt structure of firms, *Journal of Financial Economics* 85, 755–786.

Choi, Jaewon, Dirk Hackbarth, and Josef Zechner, 2018, Corporate debt maturity profiles, *Journal of Financial Economics* 130, 484–502.

Cronqvist, Henrik, Anil K. Makhija, and Scott E. Yonker, 2012, Behavioral consistency in corporate finance: Ceo personal and corporate leverage, *Journal of Financial Economics* 103, 20–40.

Custódio, Cláudia, Miguel A. Ferreira, and Luís Laureano, 2013, Why are US firms using more short-term debt?, *Journal of Financial Economics* 108, 182–212.

Datta, Sudip, Mai Iskandar-Datta, and Kartik Raman, 2005, Managerial stock ownership and the maturity structure of corporate debt, *Journal of Finance* 60, 2333–2350.

Davydenko, Sergei A., Ilya A. Strebulaev, and Xiaofei Zhao, 2012, A market-based study of the cost of default, *Review of Financial Studies* 25, 2959–2999.

Diamond, Douglas W., 1991, Debt maturity structure and liquidity risk, *Quarterly Journal of Economics* 106, 709–737.

Diamond, Douglas W., and Zhiguo He, 2014, A theory of debt maturity: The long and short of debt overhang, *Journal of Finance* 69, 719–762.

Dick-Nielsen, Jens, Peter Feldhütter, and David Lando, 2012, Corporate bond liquidity before and after the onset of the subprime crisis, *Journal of Financial Economics* 103, 471–492.

Dickerson, Alexander, Philippe Mueller, and Cesare Robotti, 2023, Priced risk in corporate bonds, *Journal of Financial Economics* 150, 103707.

Duan, Naihua, 1983, Smearing estimate: A nonparametric retransformation method, *Journal of the American Statistical Association* 78, 605–610.

Faulkender, Michael, 2005, Hedging or market timing? selecting the interest rate exposure of corporate debt, *Journal of Finance* 60, 931–962.

Flannery, Mark J., 1986, Asymmetric information and risky debt maturity choice, *Journal of Finance* 41, 19–37.

Graham, John R., 2000, How big are the tax benefits of debt?, *Journal of Finance* 55, 1901–1941.

Graham, John R., and Campbell R. Harvey, 2001, The theory and practice of corporate finance: Evidence from the field, *Journal of Financial Economics* 60, 187–243.

Greenwood, Robin, Samuel G. Hanson, and Jeremy C. Stein, 2010, A gap-filling theory of corporate debt maturity choice, *Journal of Finance* 65, 993–1028.

Guedes, Jose, and Tim Opler, 1996, The determinants of the maturity of corporate debt issues, *Journal of Finance* 51, 1809–1833.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Harford, Jarrad, Sandy Klasa, and William F. Maxwell, 2014, Refinancing risk and cash holdings, *Journal of Finance* 69, 975–1012.

He, Zhiguo, and Konstantin Milbradt, 2016, Dynamic debt maturity, *Review of Financial Studies* 29, 2677–2736.

He, Zhiguo, and Wei Xiong, 2012, Rollover risk and credit risk, *Journal of Finance* 67, 391–430.

Holmstrom, Bengt, 1999, Managerial incentive problems: A dynamic perspective, *Review of Economic Studies* 66, 169–182.

Johnson, Shane A., 2003, Debt maturity and the effects of growth opportunities and liquidity risk on leverage, *Review of Financial Studies* 16, 209–236.

Korteweg, Arthur, 2010, The net benefits to leverage, *Journal of Finance* 65, 2137–2170.

Lane, Morton, and Olivier Mahul, 2008, Catastrophe risk pricing: An empirical analysis. Policy Research Working Paper 4765, World Bank.

Leland, Hayne E., and Klaus Bjerre Toft, 1996, Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads, *Journal of Finance* 51, 987–1019.

Mian, Atif, and João A. C. Santos, 2018, Liquidity risk and maturity management over the credit cycle, *Journal of Financial Economics* 127, 264–284.

Myers, Stewart C., 1977, Determinants of corporate borrowing, *Journal of Financial Economics* 5, 147–175.

Rauh, Joshua D., and Amir Sufi, 2010, Capital structure and debt structure, *Review of Financial Studies* 23, 4242–4280.

Servaes, Henri, and Peter Tufano, February 2006, The theory and practice of corporate debt structure. Technical report, Deutsche Bank, Liability Strategies Group.

Stohs, Mark Hoven, and David C. Mauer, 1996, The determinants of corporate debt maturity structure, *Journal of Business* 69, 279–312.

Strebulaev, Ilya A., and Baozhong Yang, 2013, The mystery of zero-leverage firms, *Journal of Financial Economics* 109, 1–23.

Sydnor, Justin, 2010, (over)insuring modest risks, *American Economic Journal: Applied Economics* 2, 177–199.

Tversky, Amos, and Daniel Kahneman, 1992, Advances in prospect theory: Cumulative representation of uncertainty, *Journal of Risk and Uncertainty* 5, 297–323.

van Binsbergen, Jules H., John R. Graham, and Jie Yang, 2010, The cost of debt, *Journal of Finance* 65, 2089–2136.

Table 1: Fixed-rate bond-sample filter cascade

Filter Step	Bonds Remaining	Removed	% Removed
Raw FISD issues (January 1997–July 2025)	652395		
Conventional U.S. nonfinancial corporate filters	35313	617082	94.6%
Offering amount above $10 million	34197	1116	3.2%
Exclude convertible and putable bonds	30336	3861	11.3%
Fixed-rate coupon bonds	28086	2250	7.4%
Rated AAA–B; maturity longer than 1 year	21110	6976	24.8%
Valid annual 3-year wedge (all inputs)	20501	609	2.9%

Note: Sequential application of filters applied to the FISD source file. "Raw FISD issues (January 1997–July 2025)" is the dated source universe used for subsequent filters. Each later row reports the bonds remaining after the corresponding filter, the number removed, and the share removed relative to the previous step. The conventional-filter row keeps U.S. nonfinancial corporate issuers, standard FISD corporate bond types, and valid issue and maturity dates. The next rows apply the size, convertible, putable, and fixed-rate screens. The issue-rating row keeps AAA through B and maturity longer than one year. A rollover comparison requires a contractual maturity longer than its counterfactual tenor. The final row groups all data needed for the annual 3-year wedge. It uses issue yield when reported and the stated coupon otherwise. If a reported principal reduction has no usable price, the baseline ignores that reduction but still uses any priced reductions for the same bond. The sensitivity check removes every bond with an unpriced reduction. The final row contains 20,501 unique bonds and 152,531 bond-years.

Table 2: Rated fixed-rate FISD U.S. nonfinancial bond sample composition (21,110 bonds, January 1997–July 2025)

Metric	Count	%
Sample Composition		
Total bonds	21,110	
Unique issuers	2,967	
Issue date range	January 1997–July 2025	
Investment grade	12,932	61.3%
High yield	8,178	38.7%
Bond Features		
Redeemable (callable)	19,372	91.8%
Rule 144A	6,197	29.4%
Medium-term note	1,004	4.8%
Seniority		
Senior unsecured	16,843	79.8%
Senior secured	1,760	8.3%
Subordinated	2,470	11.7%
Subordinated (SUB)	9	0.0%
Junior subordinated	19	0.1%
Other / unclassified	5	0.0%
Missing	4	0.0%

Note: Fixed-rate FISD U.S. nonfinancial bonds rated AAA through B at issuance, with contractual tenor longer than one year, issued over the period January 1997–July 2025, after the filter cascade in Table 1. Composition shares are bond-level (unweighted by face value). Rating-at-issuance is constructed from S&P Capital IQ entity-level long-term ratings using the latest available Local Currency LT rating observed on or before the offering date. Bonds without any prior valid CIQ rating and bonds rated below B are excluded.

Table 3: Rated fixed-rate bond-sample summary statistics

Variable	N	Mean	Std	Min	P25	Median	Max
Offering Amount ($M)	21,110	554.34	565.23	10.50	250.00	400.00	15,000
Tenor (Years)	21,110	11.43	9.17	1.00	6.62	9.67	100.05
Coupon Rate (%)	21,110	6.04	2.54	1.12	4.20	5.80	13.00
YTM at Issue (%)	16,490	5.71	2.44	1.02	3.99	5.49	13.00
Gross Spread (%)	10,222	0.70	0.41	0.08	0.45	0.65	3.00

Note: *N* is the number of unique bonds with a nonmissing value in that row. Summary statistics on the fixed-rate FISD U.S. nonfinancial bonds rated AAA through B at issuance, with contractual tenor longer than one year, issued in January 1997–July 2025, after applying the filter cascade displayed in Table 1. Sample sizes vary across rows because of variable-specific FISD coverage gaps. Offering amount and tenor are universally available; YTM at issue requires a nonmissing offering price and is available for 78.1% of the sample; the gross spread is available for 48.4% of the sample. Gross spread is a deal-specific underwriter fee, not a required field for every bond record. Rule 144A offerings are privately marketed, so their offering materials need not be registered; MTN issues are drawdowns from a standing program, where compensation can be recorded at the program, agent, or pricing-supplement level. Consequently, the source record often has no issue-level gross-spread figure. The sample requires offering amounts strictly above $10 million.

Table 4: Contractual versus balance-adjusted effective tenor

Variable	*N*	Mean	Std Dev	Median	Min	Max
Panel A: Rated AAA–B Sample (21,110 bonds)						
Contractual Tenor (years)	21,110	11.43	9.17	9.67	1.00	100.05
Effective Tenor (years)	21,110	10.22	8.96	8.03	0.08	100.03
Bonds with Principal Reductions (%)	5,921 (28.0)					
Panel B: Prepaid Bonds (5,921 bonds)						
Contractual Tenor (years)	5,921	10.36	7.81	9.32	1.01	100.05
Effective Tenor (years)	5,921	6.05	4.98	4.92	0.08	69.84
Tenor Reduction (years)	5,921	4.32	5.57	3.02	0.00	98.61
Tenor Reduction (%)	5,921	39.57	26.26	37.96	0.00	99.24
Action Price (% par)	5,921	104.47	9.06	102.50	50.00	190.35
Redemption Premium (%)	5,921	4.46	7.79	2.50	−15.00	37.72

Note: *N* is the number of unique bonds in each row. The sample contains fixed-rate FISD U.S. nonfinancial bonds rated AAA through B at issuance, with contractual tenor longer than one year. Effective tenor is the principal-weighted average time to repayment defined in Equation (2), computed from FISD repayment records. Panel B reports prepaid bonds only. Action price is the realized FISD price on retired principal when reported. Redemption premium in the table is the action-price premium; the ACD contribution also weights it by the share of principal retired.

Table 5: Coverage of FISD early-retirement events

Code	Action	Reported events	Usable events	Reported bonds	Usable bonds	Usable (%)
T	Tender	3553	3365	2822	2677	94.9
E	Entire Call	2332	2226	2332	2226	95.5
B	Balance Call	1901	1858	1901	1858	97.7
IRP	Repurchase	1052	93	634	79	12.5
P	Partial Call	761	715	553	521	94.2
S	Sinking Fund	264	255	20	20	100.0
F	Refunding	4	3	4	3	75.0
	Total	9867	8515	6063	5673	93.6

Note: Sample: fixed-rate FISD U.S. nonfinancial bonds rated AAA through B at issuance, with contractual tenor longer than one year, FISD issued in January 1997–July 2025. Event columns count events; bond columns count distinct bonds and therefore do not sum across rows. Usable (%) is Usable Bonds divided by Reported Bonds. A usable event is strictly inside the bond's life, reduces principal, and has a usable action price. Short labels follow the official meanings; for example, P is a partial call. Other FISD codes may appear in the sample but are not shown because they do not describe prepayment; for example, X is an exchange.

Table 6: Worked-example bond: counterfactual refinancing calculation (Oklahoma Gas & Electric Co., CUSIP 678858BE0, BBB+, 30Y issued 2004, called 2014; 3-year rolling counterfactual)

Bond	Start	End	Notional years	Prevailing rating	CF yield / curve tenor	CF iss. cost (%)	CF ACD (%)	Obs. bond ACD (%)	Bond wedge (%)
1	2004-08	2007-08	3	BBB+	3.82 / 3Y	0.133	3.95	6.92	2.97
2	2007-08	2010-08	3	BBB+	5.63 / 3Y	0.146	5.78	6.92	1.14
3	2010-08	2013-08	3	BBB+	2.26 / 3Y	0.161	2.42	6.92	4.50
4	2013-08	2014-08	1	A-	0.56 / 1Y	0.203	0.76	6.92	6.16
Principal-weighted lifetime wedge			10		3.57	0.152	3.72	6.92	3.20

Note: Each row is one counterfactual bond. Start and end dates are shown by month, and notional years and curve tenor are rounded to whole years for display. Calculations use exact dates and principal exposure. The final bond has a shorter maturity. *Obs. bond ACD* aggregates the observed coupon (y), amortized gross spread ($g/\widetilde{T}$), and amortized call premium (ξ) per Equation (4). *CF Yield / curve tenor* uses the Merrill/TRACE yield in the refinancing month for the issuer's prevailing rating. *CF Iss Cost* comes from the issuance-cost model in Online Appendix A.4. *Bond Wedge* is observed ACD minus counterfactual ACD. The final row uses exact principal-exposure weights; Online Appendix A.6 reports the final-bond pricing sensitivity.

Table 7: Annual cost wedge by counterfactual tenor

	N	Mean 1Y	Med 1Y	Mean 2Y	Med 2Y	Mean 3Y	Med 3Y	Mean 4Y	Med 4Y	Mean 5Y	Med 5Y
1997	936	1.16	0.88	1.14	0.81	1.00	0.64	0.93	0.57	0.88	0.52
1998	2,020	0.81	0.63	0.90	0.65	0.82	0.55	0.76	0.50	0.73	0.47
1999	2,683	0.76	0.60	0.93	0.63	0.89	0.56	0.85	0.50	0.81	0.47
2000	2,875	-0.36	-0.41	0.28	0.20	0.57	0.38	0.75	0.48	0.73	0.43
2001	3,237	-0.53	0.09	-0.17	0.16	-0.05	0.25	0.51	0.51	0.56	0.43
2002	3,538	1.58	2.23	0.61	1.12	0.13	0.55	0.40	0.68	0.56	0.53
2003	4,158	2.67	3.05	1.97	2.36	1.11	1.25	0.88	1.14	0.84	0.88
2004	4,531	3.64	3.73	2.93	2.98	2.09	2.02	1.27	1.33	1.05	1.16
2005	4,663	2.68	2.63	2.81	2.72	2.45	2.26	1.55	1.46	1.20	1.18
2006	4,708	1.12	0.98	1.83	1.67	2.09	1.92	1.68	1.48	1.28	1.20
2007	4,797	0.64	0.52	1.14	0.95	1.39	1.25	1.50	1.24	1.27	1.18
2008	4,666	0.71	0.81	0.92	0.86	0.99	0.94	1.10	0.96	1.22	1.07
2009	4,817	0.52	1.36	0.84	1.27	0.98	1.13	0.85	0.96	0.87	0.93
2010	5,122	3.32	3.51	2.05	2.52	1.82	1.98	1.47	1.52	1.12	1.12
2011	5,380	4.13	4.10	3.49	3.51	2.57	2.72	2.06	2.14	1.41	1.42
2012	5,655	3.89	3.95	3.87	3.87	3.27	3.25	2.48	2.66	1.69	1.70
2013	5,959	4.11	4.15	3.95	3.99	3.60	3.59	2.97	2.92	2.11	2.03
2014	6,116	4.36	4.34	4.09	4.10	3.68	3.65	3.19	3.04	2.49	2.27
2015	6,316	3.72	3.64	3.76	3.66	3.44	3.33	3.08	2.87	2.59	2.32
2016	6,351	2.76	2.73	3.07	2.94	2.95	2.82	2.80	2.53	2.42	2.15
2017	6,503	2.82	2.59	2.75	2.55	2.61	2.36	2.54	2.21	2.29	1.95
2018	6,443	2.22	1.90	2.44	2.14	2.25	1.94	2.19	1.85	2.00	1.66
2019	6,642	1.53	1.20	1.88	1.52	1.93	1.61	1.86	1.49	1.75	1.37
2020	7,277	2.00	1.81	1.88	1.57	1.81	1.48	1.79	1.42	1.67	1.30
2021	7,398	3.05	2.86	2.62	2.41	2.16	1.93	1.94	1.66	1.75	1.34
2022	7,212	2.33	2.08	2.59	2.46	2.29	2.11	1.95	1.68	1.65	1.22
2023	7,241	-0.76	-0.88	0.87	0.69	1.38	1.10	1.53	1.27	1.32	1.02
2024	7,426	-1.22	-1.32	-0.53	-0.49	0.24	0.23	0.82	0.61	1.11	0.79
2025	7,348	-0.49	-0.48	-0.45	-0.39	-0.26	-0.14	0.32	0.32	0.84	0.55

Note: *N* is the valid 1Y bond-year count; each longer-tenor column applies its own tenor-eligibility and data-availability restrictions. **Yearly wedge**: for each bond-year, the wedge is the difference between the observed cost of debt and the counterfactual cost of debt *in that specific calendar year*. Mean and median columns are in percentage points per year. The table covers the full paper sample, January 1997–July 2025, including the pre-TRACE Merrill segment before February 2003. The sample is restricted to bond-years with a valid annual wedge.

Table 8: **Effective-tenor-weighted lifetime average**: wedge by S&P rating at issue (pp/yr)

	N	Obs	Mean	Median	Std	Rel. (%)
Panel: 1Y CF Tenor						
AAA	99	3.71	1.65	1.63	1.43	44.54
AA	522	4.50	1.88	1.74	1.46	41.73
A	2,315	5.19	2.14	1.99	1.61	41.25
BBB	3,412	5.72	2.29	2.13	1.98	39.98
BB	2,455	7.71	2.12	1.87	2.16	27.47
B	2,222	9.59	2.21	1.93	2.79	23.01
Panel: 2Y CF Tenor						
AAA	74	3.75	1.59	1.59	1.31	42.27
AA	512	4.50	1.72	1.53	1.34	38.21
A	2,284	5.20	2.00	1.88	1.50	38.56
BBB	3,378	5.73	2.21	2.06	1.85	38.47
BB	2,476	7.70	2.20	2.02	2.11	28.62
B	2,290	9.60	2.07	1.87	2.74	21.56
Panel: 3Y CF Tenor						
AAA	68	3.77	1.32	1.17	1.23	34.92
AA	487	4.55	1.48	1.32	1.23	32.39
A	2,144	5.23	1.76	1.58	1.38	33.65
BBB	3,228	5.76	1.94	1.74	1.73	33.69
BB	2,512	7.72	2.00	1.84	2.02	25.89
B	2,381	9.62	1.73	1.55	2.80	18.04
Panel: 4Y CF Tenor						
AAA	57	3.90	1.15	1.00	1.13	29.38
AA	406	4.73	1.36	1.19	1.13	28.83
A	1,919	5.33	1.54	1.36	1.29	28.90
BBB	2,989	5.84	1.70	1.53	1.67	29.15
BB	2,477	7.73	1.73	1.64	2.03	22.32
B	2,389	9.62	1.33	1.27	2.87	13.84
Panel: 5Y CF Tenor						
AAA	55	3.92	0.87	0.92	1.06	22.26
AA	381	4.67	1.08	0.96	1.10	23.09
A	1,777	5.33	1.30	1.12	1.18	24.32
BBB	2,766	5.84	1.36	1.20	1.62	23.21
BB	2,434	7.75	1.51	1.37	1.94	19.48
B	2,263	9.58	1.14	1.04	2.78	11.88

Note: N is the number of unique bonds in each rating row. The 3Y rows sum to 10,820 bonds (5,927 IG and 4,893 HY), versus 20,501 bonds in the annual headline sample. The table is smaller because each formal lifetime observation requires the bond to be fully repaid by the sample end and a valid price for every counterfactual bond with positive principal exposure. For each counterfactual tenor, the table uses one observation per bond. Exact principal exposure weights costs within a bond; effective tenor weights bonds within each rating group. Obs = weighted mean lifetime observed cost (%/yr), reconstructed as counterfactual cost plus the wedge. Rel. = weighted mean wedge as % of weighted mean observed cost. Median and Std are also effective-tenor weighted. Std is the cross-sectional standard deviation across bonds in each rating group. Only bonds with valid observed and counterfactual costs and a valid S&P rating at issue are included. Ratings: AAA through B, with finer notch ratings grouped to the letter level. The baseline counterfactual uses a 3Y rolling yield curve strategy; all five tenors shown for robustness.

Table 9: **Bond-level lifetime average**: mean wedge by callable status (pp/yr)

	Bonds	Obs	1Y	2Y	3Y	4Y	5Y	Rel. 3Y (%)
Non-Callable	1,043	6.67	1.41	1.37	1.24	1.16	1.00	18.59
Callable	9,773	6.58	2.51	2.47	2.26	2.06	1.81	34.38

Note: Callable = FISD redeemable flag = Y (including make-whole and fixed-price call features). Bonds is the number of unique securities with a valid 3Y lifetime-average wedge, not bond-year observations. Table 2 describes all 21,110 bonds in the rated issue sample. This table starts from the 21,110 bonds that also enter the security-year sample, then keeps bonds fully repaid by the sample cutoff with every counterfactual refinancing priced. Of those bonds, 8,562 are still outstanding and 1,728 are fully repaid but lack complete counterfactual pricing. This leaves 10,820 retired bonds with valid 3Y lifetime wedges; 10,816 have a reported Y/N callability flag and enter this table. Entries average bond-level lifetime wedges (Obs Cost of Debt − CF Cost of Debt) by CF tenor, in percentage points per year; each security contributes once to each tenor for which it is eligible. Obs = mean observed cost of debt (%/yr). Rel. 3Y = mean 3Y wedge as % of mean observed cost.

Table 10: Distributional properties of the yearly wedge by economic era

	N	Obs	Mean	Rel. (%)	Median	Std	% Neg	% Neg IG	% Neg HY
1997–02	15,221	8.75	0.40	4.62	0.47	2.24	29.60	19.33	41.37
2003–07	23,164	7.92	1.81	22.86	1.73	2.71	14.11	9.30	20.62
2008–09	9,651	7.52	0.98	13.02	1.03	2.33	23.95	18.55	32.40
2010–12	16,291	7.34	2.57	35.07	2.66	2.92	10.83	6.66	17.41
2013–19	44,400	5.81	2.90	49.88	2.70	2.33	5.15	4.01	7.56
2020–21	14,605	4.88	1.97	40.42	1.69	2.13	9.49	6.92	17.54
2022–25	29,199	4.81	0.88	18.34	0.67	2.37	33.28	31.99	38.50

Note: N is the number of bond-years in each era. Yearly wedge = Obs ACD − CF ACD (3Y rolling), pp/yr. Obs = mean observed cost of debt (%/yr). Rel. = mean wedge as % of mean observed cost. % Neg = fraction of bond-year observations with negative wedge. IG = investment grade ($\geq$ BBB−); HY = high yield. Eras: Pre-TRACE extension (1997–02), Pre-crisis (2003–07), GFC (2008–09), Post-GFC (2010–12), ZLB era (2013–19), COVID-QE (2020–21), Rate hike (2022–25).

Table 11: ACD components by issuance cohort (retired fixed-rate bonds)

Group/ cohort	*N*	Obs Borr	Obs Iss	Obs ACD	CF Borr	CF Iss	CF ACD	Wedge	Med Tenor
1997	509	8.19	0.28	8.79	7.57	0.27	7.84	0.95	9.9
1998	591	7.54	0.30	8.04	6.95	0.26	7.21	0.82	9.8
1999	513	8.30	0.87	9.38	7.21	0.38	7.59	1.79	9.6
2000	283	8.94	0.98	10.19	7.51	0.43	7.94	2.25	8.1
2001	541	7.93	0.41	8.82	6.30	0.32	6.62	2.20	8.4
2002	497	7.48	0.21	8.17	6.14	0.25	6.38	1.78	9.5
2003	676	6.69	0.30	7.62	5.37	0.29	5.66	1.96	9.5
2004	501	6.67	0.22	7.22	5.37	0.28	5.65	1.56	9.5
2005	407	6.82	0.24	7.51	5.89	0.28	6.17	1.34	9.6
2006	349	7.31	0.22	8.19	6.03	0.27	6.30	1.90	9.6
2007	358	7.23	0.21	7.95	4.92	0.26	5.18	2.77	9.5
2008	336	6.91	0.15	7.47	3.99	0.19	4.17	3.30	9.6
2009	553	7.53	0.20	8.70	4.51	0.24	4.76	3.94	7.7
2010	643	6.90	0.33	8.42	3.89	0.31	4.20	4.22	8.0
2011	528	5.87	0.22	6.72	3.20	0.24	3.45	3.27	8.0
2012	584	4.96	0.20	5.56	3.26	0.25	3.51	2.05	9.0
2013	625	4.79	0.24	5.49	3.12	0.27	3.39	2.10	8.0
2014	523	4.54	0.18	5.05	2.72	0.22	2.95	2.11	8.0
2015	502	4.43	0.19	5.04	3.09	0.22	3.30	1.74	7.3
2016	320	4.37	0.21	5.06	3.16	0.23	3.39	1.67	6.1
2017	333	4.43	0.23	5.25	3.29	0.23	3.52	1.73	6.1
2018	183	4.73	0.21	5.44	3.42	0.22	3.65	1.79	5.0
2019	178	4.17	0.18	5.18	3.12	0.18	3.30	1.89	5.0
2020	238	4.51	0.28	6.03	3.79	0.23	4.02	2.01	5.0
2021	36	2.38	0.19	2.62	1.52	0.19	1.72	0.90	3.0
2022	12	3.52	0.14	3.50	3.19	0.17	3.37	0.13	3.0

Note: *N* is the number of unique retired bonds in each issuance-year row, not bond-years. Sample: fixed-rate bonds fully repaid by the sample end, with valid observed ACD and counterfactual pricing for every 3-year bond with positive principal exposure. The table first aggregates each bond's counterfactual refinancing costs using principal exposure over all years in which it is outstanding, then assigns the bond once to its issuance-year cohort. *Obs Borr* is observed borrowing cost; *Obs Iss* is annualized observed issuance cost; *Obs ACD* is total observed ACD; the next three columns report the corresponding counterfactual quantities; *Wedge* is observed minus counterfactual ACD. *CF Iss* is predicted by the issuance-cost model in Equation (1), reported in Online Appendix A.4.

Table 12: Cost-wedge decomposition, IG vs. HY

	Bond-years	Total	A	B	C	D	E	I	ε
Investment Grade	75,667	1.72	1.03	0.32	0.21	0.08	-0.14	-0.07	0.28
High Yield	50,053	1.64	0.99	0.30	0.11	-0.10	-0.11	-0.09	0.53

Note: Mean bond-year wedge (pp/yr, 3Y rolling strategy) decomposed into five additive yield-curve channels, issuance-cost differential I, and residual ε. A = RF term premium $[r_f(0,T) - r_f(0,s)]$; B = Credit term premium $[\hat{s}(b_0,0,T) - \hat{s}(b_0,0,s)]$; C = Base-rate change $[r_f(0,s) - r_f(t_k,s)]$; D = Systematic spread change $[\hat{s}(b_0,0,s) - \hat{s}(b_0,t_k,s)]$; E = Rating migration $[\hat{s}(b_0,t_k,s) - \hat{s}(b_{t_k},t_k,s)]$. I is observed minus counterfactual annualized issuance cost. Fixed-rate bonds only. All yield-curve components from Merrill/TRACE corporate and GSW Treasury zero-coupon curves. The total wedge includes observed and counterfactual issuance-cost components; counterfactual issuance cost is predicted by the issuance-cost model in Equation (1); estimates are reported in Online Appendix A.4. Call-premium terms and primary–secondary market residuals load into ε. The decomposition requires contractual tenor in $[2, 30]$ years and complete observations for all five yield-curve channels. The headline sample has 152,531 bond-years. The 30-year limit drops 26,313 IG and 498 HY bond-years; no other headline rows are lost. The dropped IG rows average 2.64 pp/yr, which lowers the IG mean from 1.96 in the headline sample to 1.72 here.

Table 13: Floating rate bond sample and benchmark composition

Panel A: Coupon-history panel size

Metric	Value
Total bond-year observations	6,031
Unique bonds	1,508
Year range	January 1997–July 2025

Panel B: Benchmark reference-rate composition

Benchmark	Bonds	Bond-Years	Avg Coupon	Std	Avg Spread
USD_LIBOR_3M	1,194	4,903	4.04%	2.77%	+169 bp
USD_LIBOR_1M	146	340	3.51%	1.79%	+22 bp
USD_LIBOR_6M	79	558	13.33%	51.16%	+1058 bp
SOFR	67	174	3.95%	2.27%	+66 bp
PRIME_RATE	10	23	3.76%	2.04%	−283 bp
FED_FUNDS	6	17	3.75%	2.40%	+22 bp
CMT_2Y	2	6	5.99%	0.46%	+42 bp
CMT_3M	2	7	4.69%	1.56%	+168 bp
CMT_1Y	1	2	4.22%	0.62%	−7 bp
CP_30D	1	1	5.28%	–	+0 bp

Note: Sample: floating-rate FISD bonds, January 1997–July 2025. The coupon-history block reports the annual floating-coupon panel after parsing the FISD coupon-formula text and mapping each index name to a benchmark-rate series. Each bond-year observation aggregates all reset-level coupons within a calendar year, so bond-years exceed unique bonds. The benchmark block groups the same coupon histories by reference rate and sorts them by the number of unique bonds. Two bonds switch benchmarks, so benchmark-row bond counts are not mutually exclusive. Average coupon and standard deviation are annual coupon-rate moments; average spread is the contractual add-on converted to basis points. LIBOR entries denote USD LIBOR at the listed tenor; SOFR is the Secured Overnight Financing Rate; Fed Funds is the Federal Funds Effective Rate; CMT is Constant Maturity Treasury; CP is 30-day AA financial commercial paper. Prime-linked bonds can have negative quoted spreads because Prime is typically above LIBOR/SOFR-style reference rates.

Table 14: Investment-grade fixed-rate and floating-rate ACD evidence

Panel A: Mean wedge by counterfactual tenor

Tenor	IG Fixed *N*	IG Fixed Obs. ACD	IG Fixed Wedge	IG Floating *N*	IG Floating Obs. ACD	IG Floating Wedge	Fixed–floating gap
1Y	102,953	5.36	2.03	2,143	3.02	−0.15	2.18
2Y	102,985	5.36	2.07	1,831	2.97	−0.12	2.19
3Y	101,980	5.38	1.96	1,322	2.98	−0.04	2.00
4Y	99,525	5.44	1.84	637	3.48	0.46	1.38
5Y	97,460	5.44	1.64	575	3.42	0.25	1.39

Panel B: Observed and counterfactual ACD components, 3Y counterfactual

Group	*N*	Obs Borr	Obs Iss	Obs Redem	Obs ACD	CF Borr	CF Iss	CF ACD	Wedge	Med Tenor
IG Fixed	101,980	5.17	0.06	0.15	5.38	3.29	0.14	3.43	1.96	10.0
IG Floating	1,322	2.90	0.09	−0.01	2.98	2.88	0.14	3.02	−0.04	3.0

Note: *N* is bond-years. Sample: investment-grade annual bond-year observations, January 1997–July 2025. Panel A reports mean observed ACD and observed-minus-counterfactual ACD wedges by counterfactual tenor. Panel B reports the 3-year counterfactual component view. *Obs Borr* is observed borrowing cost, *Obs Iss* is annualized observed issuance cost, and *Obs Redem* is the annualized redemption premium. *Obs ACD* is their observed all-in cost. The CF columns report the corresponding counterfactual borrowing, issuance-cost, and ACD quantities. *Wedge* is observed minus counterfactual ACD. *Med Tenor* is the bond-level median contractual tenor in years, with each bond counted once.

Table 15: Fixed-rate versus floating-rate wedge decomposition, IG and HY

	Bond-years	Total	A	B	C	D	E	F	I	ε
Investment Grade Fixed	75,667	1.72	1.03	0.32	0.21	0.08	−0.14	0.00	−0.07	0.28
Investment Grade Float	1,075	−0.51	0.23	0.07	0.02	−0.02	−0.01	−0.08	−0.05	−0.68
High Yield Fixed	50,053	1.64	0.99	0.30	0.11	−0.10	−0.11	0.00	−0.09	0.53
High Yield Float	32	−0.48	0.53	0.30	0.00	−0.11	−0.11	−0.54	0.06	−0.61

Note: Mean bond-year wedge components (pp/yr, 3Y rolling strategy). The table implements Equation (10). Columns A–E are the original five yield-curve channels in order: risk-free term premium, credit term premium, base-rate change, systematic spread change, and rating migration. F is zero for fixed-rate bonds; for floating-rate bonds, it is the subsequent floating coupon minus the first observed floating coupon. I is observed minus counterfactual annualized issuance cost. The residual ε is the part left after A–F and I. For fixed-rate bonds, it captures primary–secondary pricing differences and call-premium terms. For floating-rate bonds, it is mainly the first observed floating coupon minus the matching fixed-rate yield-curve rate at issuance, plus small remaining pricing and call-premium differences. The observed gross spread is reported in 100.0% of floating-rate bond-years in this decomposition; counterfactual fees are predicted by the fixed-rate issuance-cost model because the floater sample is too small to estimate a separate model. The fixed rows use the same 30-year limit as Table 12: it drops 26,313 IG and 498 HY bond-years from the headline sample. The dropped IG rows average 2.64 pp/yr, so the IG fixed mean is 1.72 here rather than the headline 1.96. The HY floating row is small and should be read as descriptive.

Table 16: Firm-year panel summary statistics

	N	Mean	Std Dev	P25	Median	P75
Bonds per Firm-Year	23,575	5.05	6.38	2.00	3.00	6.00
Avg Bonds Outstanding ($M)	23,575	2338.00	6272.00	285.00	700.00	2077.00
Original Par of Bonds ($M)	23,575	2759.00	7413.00	350.00	825.00	2408.00
Wtd Obs Cost of Debt (%/yr)	23,575	7.20	2.54	5.37	6.90	8.78
Wtd 3Y CF Cost of Debt (%/yr)	19,034	5.25	2.65	3.16	4.83	6.73
Annual 3Y Wedge (pp/yr)	23,047	1.72	2.14	0.61	1.71	2.92
Annual 3Y Dollar Wedge ($M/yr)	23,047	41.00	131.00	2.00	10.00	37.00
Annual 3Y Wedge / Interest Exp. (%)	22,674	21.35	192.14	2.31	13.61	34.09
Market Leverage (Debt / EV)	18,810	0.39	0.42	0.19	0.34	0.52
Book Leverage	23,334	0.41	0.22	0.27	0.37	0.51
Convertible Debt / Assets	23,334	0.01	0.03	0.00	0.00	0.00
Rent / Assets	23,334	0.02	0.03	0.00	0.01	0.02
Cash / Assets	23,334	0.07	0.09	0.01	0.04	0.10
Short Debt Share	23,292	0.10	0.15	0.01	0.05	0.13
Debt / EBITDA	23,575	3.85	3.32	1.88	3.18	4.83
Negative EBITDA Indicator	23,575	0.03	0.16	0.00	0.00	0.00
Interest Coverage	23,575	14.94	46.28	3.04	5.41	9.97
Profitability	23,252	0.21	0.16	0.11	0.18	0.29
Profitability Volatility	23,117	0.04	0.07	0.01	0.02	0.04
Size (log AT)	23,334	8.45	1.55	7.39	8.45	9.57
PPE / Assets	23,334	0.40	0.27	0.15	0.36	0.64
CAPEX / Assets	23,334	0.06	0.06	0.02	0.04	0.07
Cash and Short-Term Inv. ($M)	23,575	1269.00	5692.00	26.00	140.00	667.00
Share Repurchases ($M)	23,575	414.00	2418.00	0.00	0.00	120.00
Dividends ($M)	23,575	363.00	1180.00	0.00	32.00	206.00
Interest Expense (XINT, $M)	23,190	270.00	705.00	38.00	90.00	243.00
Dimson Beta (std)	13,574	0.99	0.50	0.63	0.94	1.27
Idiosyncratic Risk (std)	13,574	0.94	0.53	0.59	0.80	1.12
S&P Rating (annual, numeric)	23,575	10.68	3.45	8.00	10.00	14.00

Note: N is the number of nonmissing issuer-years in each row. The full Compustat-matched firm panel has 23,575 issuer-years and 2,161 issuers. N varies because financial and rating variables have different data availability. Bond portfolios contain fixed-rate bonds from the security-year sample that are active in the issuer-year and are weighted by their within-year average face amount. Before a reported repayment, the balance uses the principal then outstanding; after it, the reported remaining principal. For a repayment year, the average reflects the days at each balance. Bond and Compustat dollar values are in millions of nominal U.S. dollars. The dollar wedge is an annual flow: each bond-year wedge is multiplied by each reported principal balance and by the exact fraction of the year for which that balance is outstanding (days/365.25). Thus a partial principal reduction enters the dollar denominator exactly once. The annual 3Y wedge uses only bond-years with valid observed and 3-year counterfactual costs. Debt/EV equals (DLTT + DLC) divided by market equity plus debt less cash, and is defined only for positive EV. Ratios are reported as fractions unless a percent sign is shown.

Table 17: Firm-level wedge by counterfactual tenor and aggregate dollar totals

Panel A: Firm-year wedge by counterfactual tenor

	N firm-yrs	Obs (pp)	Mean (pp)	Rel. (%)	Median (pp)	Std (pp)	$ Mean ($M)	$ Total ($B)
1Y	22,847	7.15	1.86	26.01	1.96	2.53	42.12	962.37
2Y	22,974	7.16	1.90	26.48	1.95	2.32	44.09	1012.99
3Y	23,047	7.17	1.72	23.92	1.71	2.14	40.93	943.30
4Y	23,039	7.18	1.53	21.28	1.50	2.02	37.21	857.38
5Y	22,864	7.17	1.34	18.71	1.27	1.81	32.23	736.79

Panel B: Aggregate 3-year dollar wedge ($B)

Sample	Bonds	Bond-yrs	Firms	Firm-yrs	Nominal	Real (2025)
All eligible bonds	20,501	152,531	2,963	34,186	1172.5	1574.1
IG-at-issue bonds	12,389	101,980	1,072	17,271	891.2	1191.3
Compustat-matched firms	16,830	116,277	2,152	23,047	943.3	1264.2
IG-rated firms	11,364	79,767	925	11,843	761.0	1012.3

Note: In Panel A, N firm-yrs is the number of issuer-years with a valid wedge for that counterfactual tenor. Panel B reports unique bonds, bond-years, unique firms, and firm-years in separate columns. In its first two rows, firm counts are the issuers in the stated bond sample. In its last two rows, bond counts are the eligible bonds held by the stated Compustat-matched firm-years. The wedge is the amount-outstanding-weighted average across fixed-rate bonds with valid observed and counterfactual annual costs. The 3-year N is exactly the nonmissing Annual 3Y Wedge N in Table 16; the larger full-panel N there also includes issuer-years without a valid annual wedge. Both tables use the same sample years. Obs is mean observed cost of debt; Rel. is mean wedge divided by mean observed cost. Dollar wedges multiply each reported principal balance by the exact fraction of the year for which it is outstanding (days/365.25) and are nominal in Panel A. Panel B uses the 3-year counterfactual. IG bonds are classified at issuance; IG firms use the year-end S&P issuer rating. Nominal totals sum annual dollar wedges. Real totals first convert each annual flow to 2025 dollars using headline CPI.

Table 18: Annual 3-year firm-level bond wedge by portfolio tenor

	N firm-years	Mean (pp/yr)	Median (pp/yr)	P25	P75
Short-term ($\leq$5 Y)	557	1.42	0.70	0.02	2.89
Medium-term (5–10 Y)	12,475	1.48	1.54	0.39	2.75
Long-term (>10 Y)	10,015	2.02	1.96	0.91	3.11

Note: N is the number of issuer-years in each portfolio-tenor bin in the Compustat-matched firm panel. Portfolio tenor is the within-year-average-balance-weighted contractual tenor of all fixed-rate bonds in the security-year sample with positive amount outstanding in that issuer-year, including bonds too short to support a 3-year counterfactual. The annual 3-year wedge is the within-year-average-balance-weighted observed-minus-counterfactual cost across only the bonds eligible for the annual 3-year counterfactual.

Table 19: Swap-adjustment sensitivity: mean wedge under alternative pay-floating shares

Swap Share (%)	Adj. Mean Wedge (%)	Adj. Median Wedge (%)	Reduction vs. Baseline
0%	1.85	1.70	0.0%
5%	1.76	1.61	5.1%
10%	1.66	1.51	10.2%
15%	1.57	1.42	15.2%
20%	1.48	1.32	20.3%
25%	1.38	1.23	25.4%
30%	1.29	1.13	30.5%
40%	1.10	0.94	40.7%
50%	0.91	0.75	50.8%
5.14%	**1.76**	**1.60**	**5.2%**

Note: Illustrative mean and median fixed-rate-bond wedge (pp/yr) under the assumption that a uniform share ϕ of the headline fixed-rate sample is synthetically converted to floating via interest-rate swaps. This sensitivity is motivated by Equation (14); it does not implement a direct firm-year adjustment. "Reduction vs. Baseline" is the percent decline relative to the unadjusted ($\phi = 0$) mean reported in the first row. The **bold** row at the bottom corresponds to $\phi = 5.14\%$, the gross fixed-to-floating share detected in 10-K filings (Section 5.3). The full detection methodology is reported in Online Appendix A.3.

Table 20: Affordability of negative yearly wedge: flow basis

	Eligible N	Positive Need N (%)	Mean	Median	Buyback	+CHE	+DVT	+Prior W^+	Uncovered N (% eligible)
Panel: 1Y CF Tenor									
All	22,847	4,440 (19.4%)	39.36	7.99	36.08	85.50	90.38	93.63	283 (1.239%)
IG	11,829	1,502 (12.7%)	77.05	18.03	58.66	90.08	97.87	99.93	1 (0.008%)
HY	11,018	2,938 (26.7%)	20.10	6.63	24.54	83.15	86.56	90.40	282 (2.559%)
Panel: 2Y CF Tenor									
All	22,974	3,512 (15.3%)	29.08	6.98	34.05	85.96	89.72	93.28	236 (1.027%)
IG	11,842	1,033 (8.7%)	55.38	11.65	60.70	92.64	98.06	99.42	6 (0.051%)
HY	11,132	2,479 (22.3%)	18.11	6.31	22.95	83.18	86.24	90.72	230 (2.066%)
Panel: 3Y CF Tenor									
All	23,047	3,199 (13.9%)	21.92	5.62	31.92	85.68	88.47	92.09	253 (1.098%)
IG	11,843	795 (6.7%)	40.41	5.37	61.51	94.84	97.99	99.12	7 (0.059%)
HY	11,204	2,404 (21.5%)	15.80	5.65	22.13	82.65	85.32	89.77	246 (2.196%)
Panel: 4Y CF Tenor									
All	23,039	2,868 (12.4%)	19.08	4.77	29.60	85.77	88.63	91.63	240 (1.042%)
IG	11,808	511 (4.3%)	37.10	2.74	63.80	94.13	97.26	98.04	10 (0.085%)
HY	11,231	2,357 (21.0%)	15.18	5.14	22.19	83.96	86.76	90.24	230 (2.048%)
Panel: 5Y CF Tenor									
All	22,864	2,941 (12.9%)	17.17	3.03	31.15	89.22	91.57	93.17	201 (0.879%)
IG	11,715	536 (4.6%)	31.89	1.23	59.33	95.52	97.76	98.13	10 (0.085%)
HY	11,149	2,405 (21.6%)	13.89	3.91	24.86	87.82	90.19	92.06	191 (1.713%)

Note: Eligible N is the total number of firm-years with a valid annual wedge for that tenor and rating group. Positive Need N is the number with a negative wedge; its percentage is relative to Eligible N. Uncovered is the number still not covered after all four buffers; its percentage is also relative to Eligible N and is therefore the unconditional historical frequency. Need is the dollar magnitude of negative yearly wedge: $\text{Need}_{it} \equiv \max(0, -W_{it})$, where W_{it} is the yearly dollar wedge for firm i in year t. It does not count consecutive blocks. Mean and Median are in \$M across positive-need firm-years. Coverage columns report the cumulative share (%) of positive-need firm-years covered after sequentially cutting share repurchases (Compustat PRSTKC), then using cash and short-term investments (CHE), then dividends (DVT), and finally the needed portion of prior-year positive wedge savings W^+. For the 3Y IG benchmark, 7 of 11,843 eligible firm-years (0.059%) remain uncovered; the conditional share among 795 positive-need firm-years is 0.9%.

Table 21: Firm-years with EBIT below interest expense: actual versus counterfactual strategy

	Firm-Years N	Actual Distress (%)	Counterfactual Distress (%)	Change (pp)	Distress Added N	Distress Avoided N
All firm-years						
All valid firm-years	22,649	15.1	14.1	-1.0	122	342
Counterfactual saves interest	19,532	12.7	11.0	-1.8	0	342
Counterfactual raises interest	3,117	30.2	34.1	3.9	122	0
Investment-grade firm-years						
All valid firm-years	11,733	2.7	2.5	-0.2	0	23
Counterfactual saves interest	10,958	2.7	2.5	-0.2	0	23
Counterfactual raises interest	775	3.4	3.4	0.0	0	0
High-yield firm-years						
All valid firm-years	10,916	28.4	26.6	-1.8	122	319
Counterfactual saves interest	8,574	25.5	21.8	-3.7	0	319
Counterfactual raises interest	2,342	39.1	44.3	5.2	122	0

Note: Firm-Years N is the number of Compustat-matched firm-years in each row; Distress Added N and Distress Avoided N are also firm-year counts. Each observation has valid EBIT, nonnegative interest expense (XINT), and a valid 3Y firm-year dollar wedge. Distress means EBIT is below annual interest expense. Counterfactual interest expense is actual interest expense minus the 3Y dollar wedge, so counterfactual distress means EBIT is below $\max\{0, \text{XINT} - W_{it}^{(3)}\}$, where $W_{it}^{(3)}$ is the 3Y firm-year dollar wedge. A positive wedge means the counterfactual saves interest; a negative wedge means it raises interest. Change is counterfactual minus actual distress in percentage points. Distress Added counts firm-years distressed only under the counterfactual; Distress Avoided counts firm-years distressed only under the actual strategy.

Table 22: **Firm-level volatility**: observed vs. yearly counterfactual cost of debt (pp/yr)

	N firms	Obs Mean	CF Mean	Obs Std	CF Std	Vol Ratio	Obs IQR	CF IQR
IG	817	5.82	3.68	0.71	1.64	2.31	1.02	2.43
HY	1,268	8.87	7.59	0.65	1.94	2.99	0.82	2.29

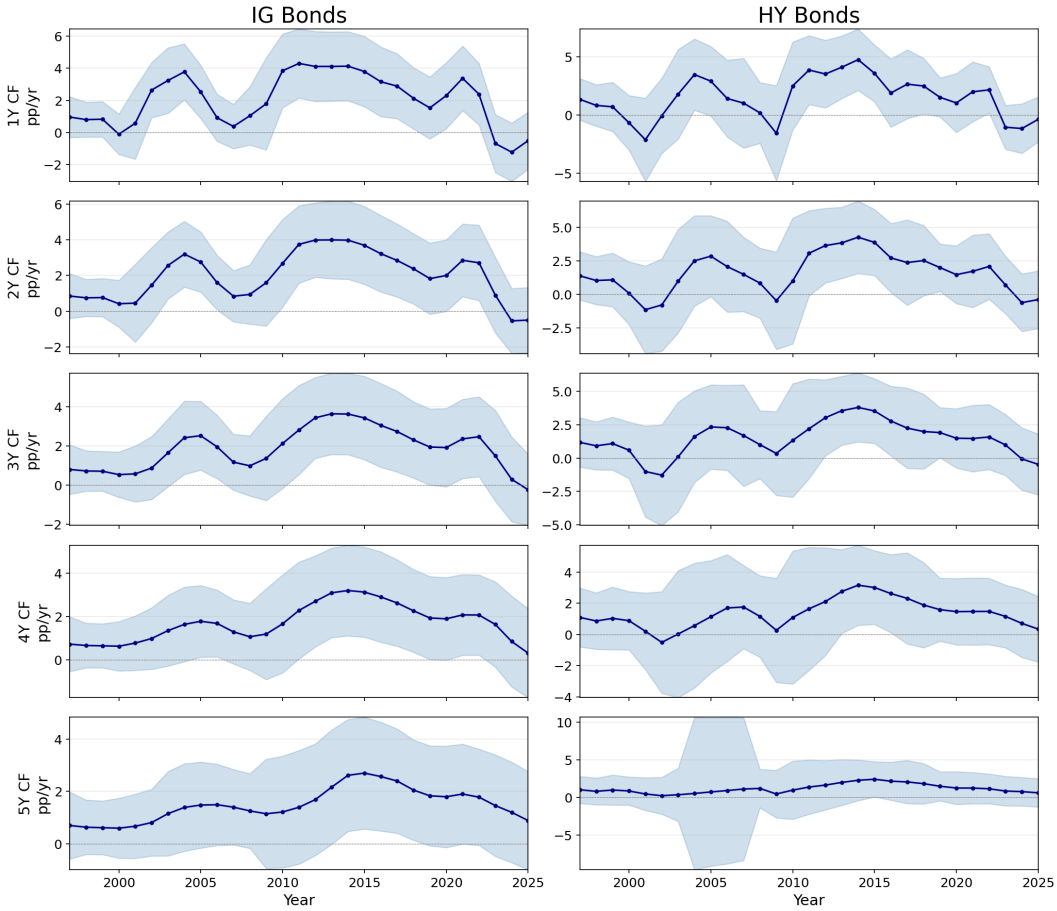
Note: N firms is the number of firms used in each row. Within-firm time-series volatility of the weighted-average bond cost of debt, computed over the full Compustat-matched firm-year sample (January 1997–July 2025). **Minimum-year criterion**: a firm is included only if it has at least 3 firm-year observations with non-missing observed cost and counterfactual cost. This ensures each within-firm standard deviation is computed over a meaningful number of data points. Obs Mean and CF Mean are cross-firm averages of within-firm means. Obs Std (CF Std) is the cross-firm average of within-firm standard deviations. Vol Ratio = CF Std / Obs Std. IQR is the cross-firm average of within-firm interquartile ranges. IG/HY is assigned separately in each firm-year using that year's S&P rating (BBB− or better is IG), so a firm can appear in both rows if it has at least three valid years in each group.

Table 23: Managerial ownership and contractual bond tenor

	(1)	(2)	(3)	(4)	(5)
Dep. var: Log Contractual Tenor					
Governance					
CEO Ownership (%)	-0.0023			0.0067	0.0085
	(0.0068)			(0.0070)	(0.0072)
CFO Ownership (%)		-0.0524		-0.0858	-0.1379**
		(0.0477)		(0.0591)	(0.0653)
CFO Tenure (yrs)			0.0066		0.0080*
			(0.0044)		(0.0045)
Datta, Iskandar-Datta, and Raman (2005) Controls					
IG (d)	0.1746***	0.1643***	0.1625***	0.1652***	0.1622***
	(0.0492)	(0.0524)	(0.0521)	(0.0525)	(0.0522)
Log Offering Amt	0.1461***	0.1455***	0.1458***	0.1460***	0.1468***
	(0.0253)	(0.0261)	(0.0256)	(0.0265)	(0.0259)
Book Leverage	-0.2800***	-0.3036***	-0.3215***	-0.3137***	-0.3293***
	(0.0934)	(0.1016)	(0.0989)	(0.1029)	(0.1008)
Log Total Assets	-0.1025***	-0.1139***	-0.1233***	-0.1147***	-0.1295***
	(0.0226)	(0.0249)	(0.0249)	(0.0249)	(0.0254)
Profitability	0.1721**	0.1827**	0.1776**	0.1778**	0.1706**
	(0.0758)	(0.0823)	(0.0825)	(0.0839)	(0.0838)
Profit Volatility	-0.1504	-0.0383	-0.0581	-0.0660	-0.0875
	(0.3137)	(0.3338)	(0.3408)	(0.3594)	(0.3671)
PPE / Assets	0.0301	0.0976	0.1143	0.1108	0.1220
	(0.1879)	(0.2016)	(0.1970)	(0.2067)	(0.2013)
CAPEX / Assets	0.2801	0.2437	0.2405	0.2512	0.2589
	(0.3988)	(0.4576)	(0.4568)	(0.4698)	(0.4678)
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	No	No	No	No	No
R^2	0.016	0.016	0.017	0.017	0.017
N	7,987	7,384	7,384	7,283	7,283
Mean CFO Ownership (%)	–	0.071	–	0.071	0.071
SD CFO Ownership (pp)	–	0.161	–	0.160	0.160
Effect of +1 SD CFO Ownership (%)	–	-0.8	–	-1.4	-2.2

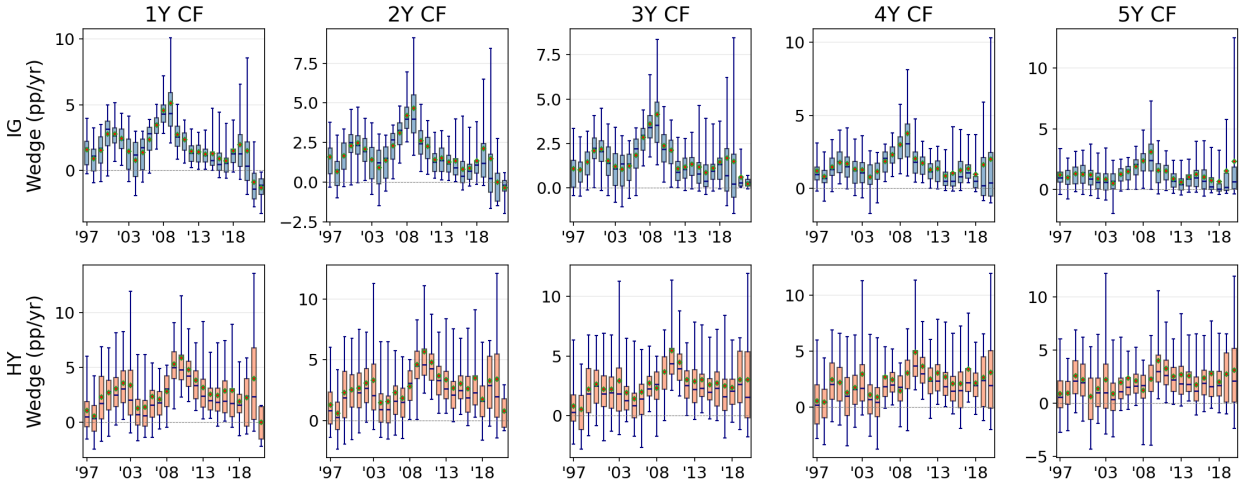
Note: N is the number of bond issues in each regression. The OLS regressions use the governance, bond, and firm controls shown, with firm fixed effects and no macro or year fixed effects. A bond's full life need not be observed, and no maturity restriction is applied. Ownership is the percentage of shares held. The last row converts the CFO coefficient into the percentage change in tenor for a one-standard-deviation increase in CFO ownership. Standard errors are clustered by GVKEY. $^{***}\ p < 0.01$, $^{**}\ p < 0.05$, $^{*}\ p < 0.10$.

Figure 1: Aggregate cost wedge over time (annual, bond-year weighted).



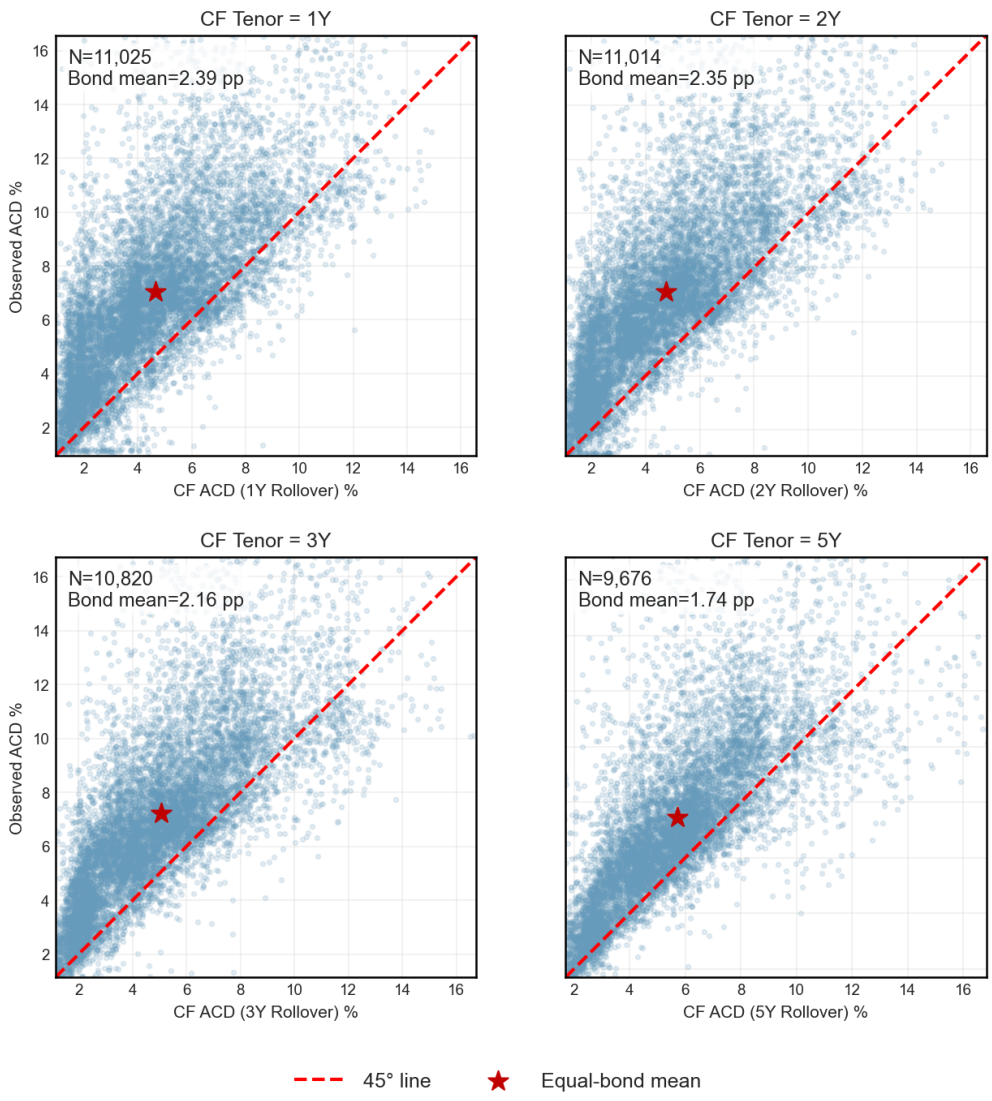
Note: **Sample.** Fixed-rate USD non-convertible, non-putable corporate bonds observed at the bond-year level, January 1997–July 2025. Bonds enter the annual sample in each year they are observed, including bonds still outstanding at the sample end. **Construction.** For bond i outstanding in calendar year t, the annual *cost wedge* is $w_{i,t} = \text{ACD}^{\text{obs}}_{i,t} - \text{ACD}^{\text{CF}}_{i,t}$, using the costs defined in Section 3. The counterfactual is evaluated at each exact rollover date using the issuer's prevailing rating and the monthly Merrill/TRACE rating-by-tenor curve. If two counterfactual bonds are outstanding during the same calendar year, the annual counterfactual ACD is their exact-day-weighted blend as in Equation (5). The displayed line is the equal-weighted mean across valid bond-year observations, so each outstanding bond contributes once to its calendar year. The shaded band is the annual mean plus or minus one cross-sectional standard deviation within the same rating group and counterfactual-tenor panel.

Figure 2: Lifetime wedge by issuance year and counterfactual tenor (IG vs. HY).



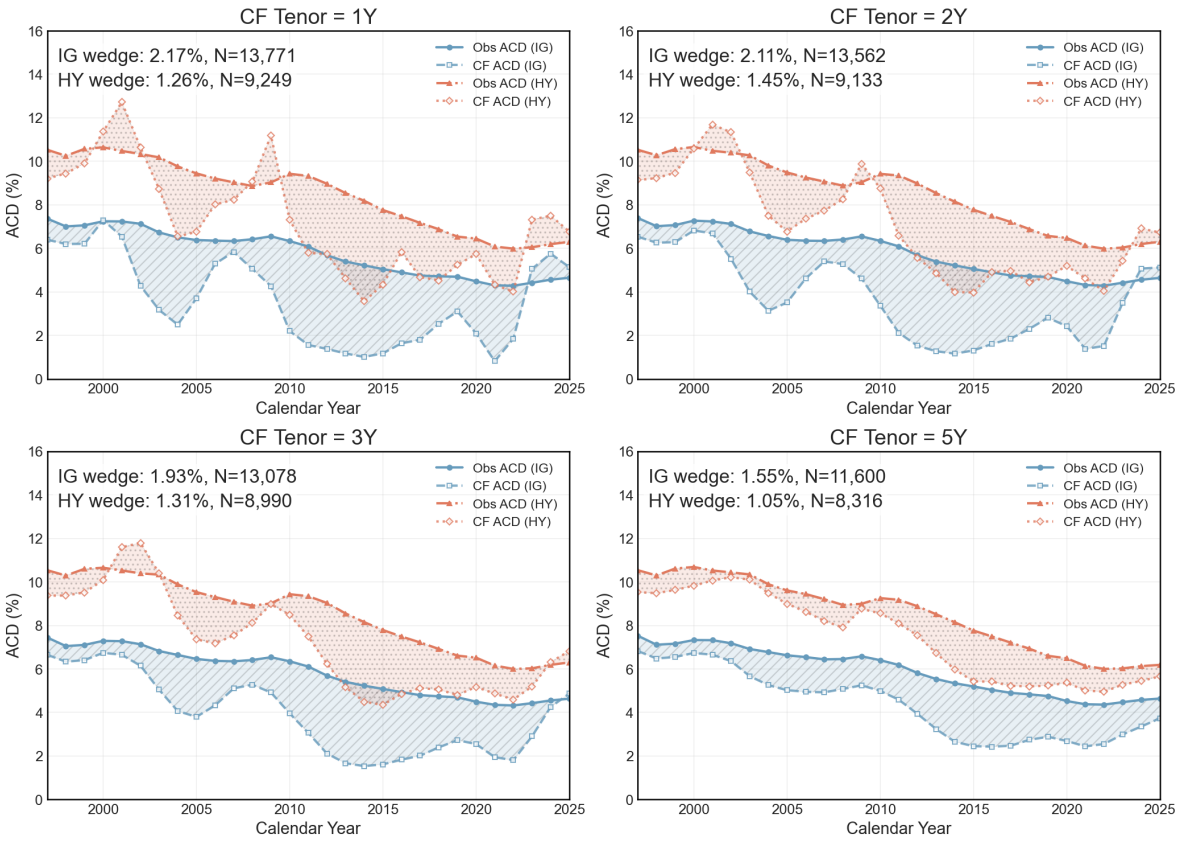
Note: **Sample.** Fixed-rate USD non-convertible, non-putable bonds fully repaid by the sample end, with valid observed cost and counterfactual pricing for every counterfactual bond with positive principal exposure, sorted by S&P long-term rating at issue into investment grade (top row, blue) and high yield (bottom row, coral). **Construction.** A bond's lifetime wedge is its observed lifetime ACD minus its counterfactual lifetime ACD, with both weighted by exact principal exposure as in Equations (6)–(8). Columns are counterfactual rollover tenors $k \in \{1, 2, 3, 4, 5\}$; within each panel, boxes group bonds by issuance year and require at least 10 bonds in the year cell. Boxes show the interquartile range, whiskers the 5th–95th percentile, the central line the median, and the diamond the cohort mean.

Figure 3: Observed vs. counterfactual ACD by counterfactual tenor.



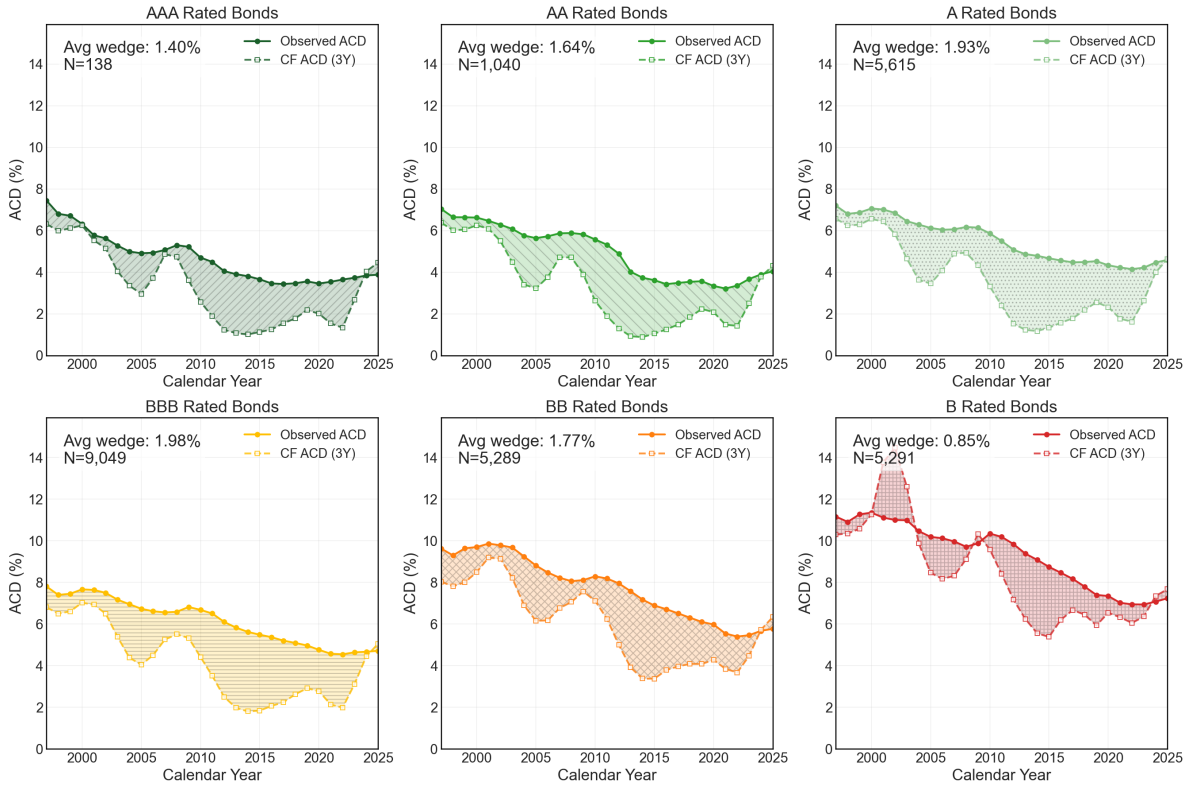
Note: **Sample.** Fixed-rate bonds fully repaid by the sample end, with valid observed ACD and counterfactual pricing for every counterfactual bond with positive principal exposure; one panel per counterfactual rollover tenor $k \in \{1, 2, 3, 5\}$. **Construction.** The *Annual Cost of Debt* (ACD) is the all-in annualized cost defined in Equation (4). For fixed-rate bond i with effective life $\widetilde{T}_i$, the observed ACD on the vertical axis is $\overline{\text{ACD}_i^{\text{obs}}} = y_i + g_i/\widetilde{T}_i + \xi_i$, where y_i is the issue yield when reported and the stated coupon otherwise, g_i is the gross spread, and ξ_i is the annualized redemption premium. The counterfactual ACD is $\text{ACD}_i^{\text{CF}}(k)$ from Equation (6): the exact principal-exposure-weighted mean over the fully observed life of the k-year rolling Merrill/TRACE borrowing cost plus predicted rollover issuance costs. Each dot is one bond plotted as $\left(\text{ACD}_i^{\text{CF}}(k), \overline{\text{ACD}_i^{\text{obs}}}\right)$. The dashed 45° line marks break-even; points above the line indicate the long-term bond was *ex post* more expensive than the hypothetical short-rolling strategy. Red stars mark equally weighted bond lifetime means, such that each bond receives equal weight.

Figure 4: Stock-based ACD of outstanding bonds: IG vs. HY (annual, prepayment-adjusted).



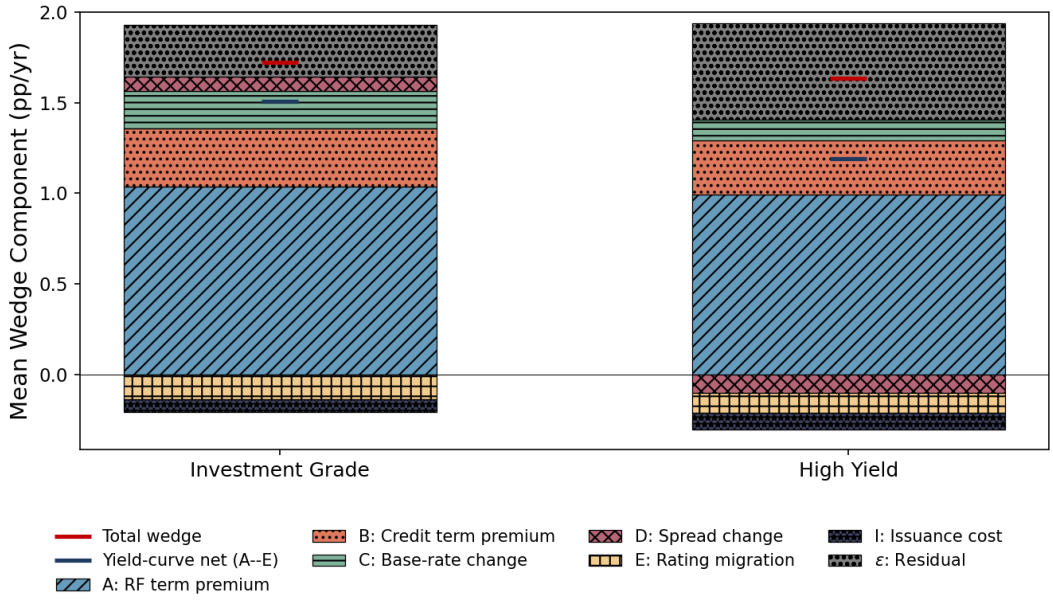
Note: **Sample.** Fixed-rate bonds outstanding in each calendar year of the counterfactual wedge sample, January 1997–July 2025, including bonds still outstanding at the sample end. One panel is shown for each counterfactual rollover tenor $k \in \{1, 2, 3, 5\}$. **Construction.** For year t and rating group $g \in \{\text{IG}, \text{HY}\}$, the stock-based observed and counterfactual ACDs are equal-weighted means across valid bond-years in g. Within a bond-year, counterfactual ACD is principal-exposure-weighted across all counterfactual bonds active during the year. The bond-year is assigned to the IG/HY status used for the largest share of its within-year principal exposure; the rating at issue is used only when no linked CIQ history is available. Thus, rating assignment does not alter the underlying blended annual ACD. Counts printed in different rating panels are nonexclusive because a bond can migrate across groups over time. Solid lines show observed ACD; dashed lines show counterfactual ACD; the shaded area is the annual stock-based wedge. A bond leaves the annual sample when all principal is repaid; after a partial repayment, only the remaining principal is counted.

Figure 5: Stock-based ACD of outstanding bonds by rating bucket (annual, 3Y rollover, prepayment-adjusted).



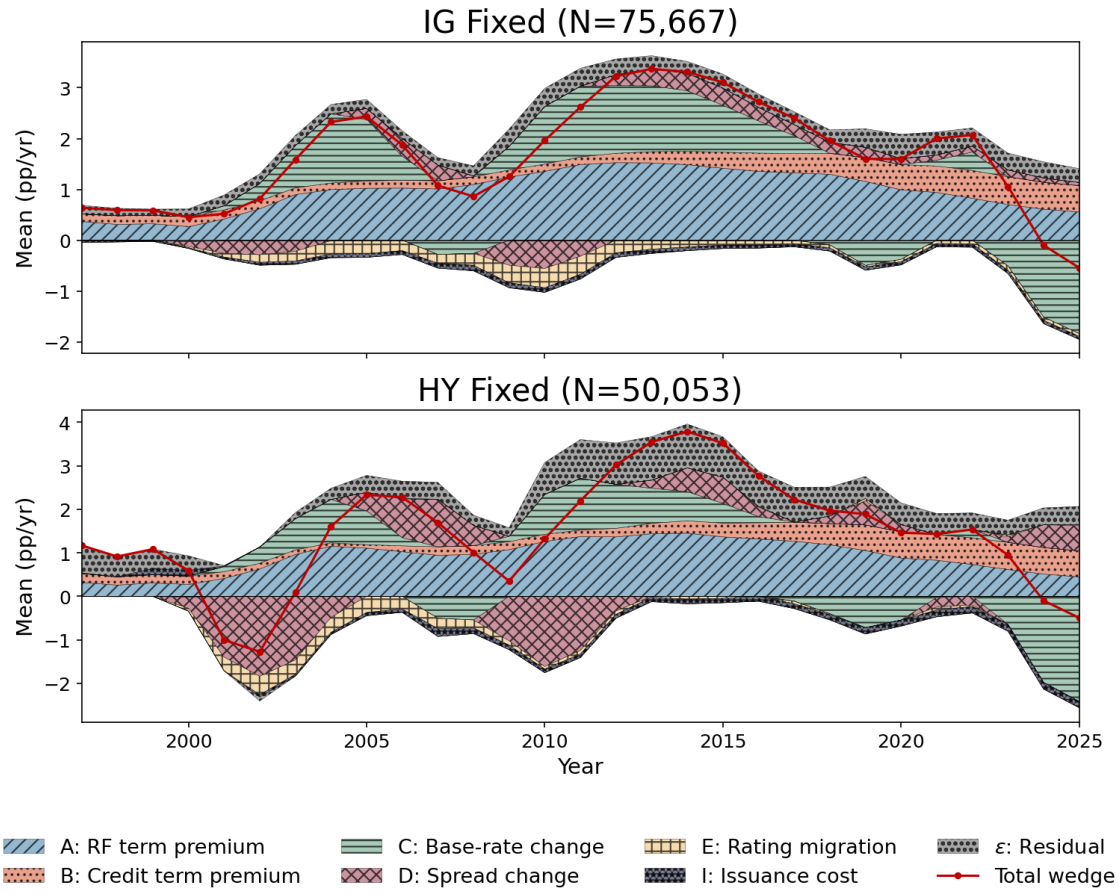
Note: **Sample.** Fixed-rate bonds outstanding in each calendar year of the counterfactual wedge sample, January 1997–July 2025, including bonds still outstanding at the sample end. One panel is shown for each S&P long-term rating bucket (AAA, AA, A, BBB, BB, B). **Construction.** The annual observed and counterfactual ACDs mirror Figure 4. Counterfactual ACD is principal-exposure-weighted across every counterfactual bond active during the year. The bond-year is assigned for display to the prevailing rating used for the largest share of its within-year exposure. The rating at issue is used only when no linked CIQ history is available. Counts across rating panels are nonexclusive because a migrating bond can appear in more than one panel over time. Solid lines show observed ACD; dashed lines show counterfactual ACD; the shaded area is the annual wedge. After a partial repayment, only the remaining principal is counted.

Figure 6: Cost-wedge decomposition, IG vs. HY.



Note: **Sample and unit.** Fixed-rate bond-years in the 3-year rollover sample with contractual tenor between 2 and 30 years and complete observations for all five curve channels; bars report equal-weighted mean percentage points per year for IG and HY bonds at issuance. **Construction.** Equation (9) splits the wedge into the risk-free term premium (A), credit term premium (B), base-rate change (C), systematic spread change (D), rating migration (E), and the residual ε_i. I separately reports observed minus counterfactual annualized issuance cost. Channels A and B are set at issuance; C–E depend on the subsequent paths of rates, spreads, and ratings. The GSW Treasury and Merrill/TRACE credit curves supply the fitted yields.

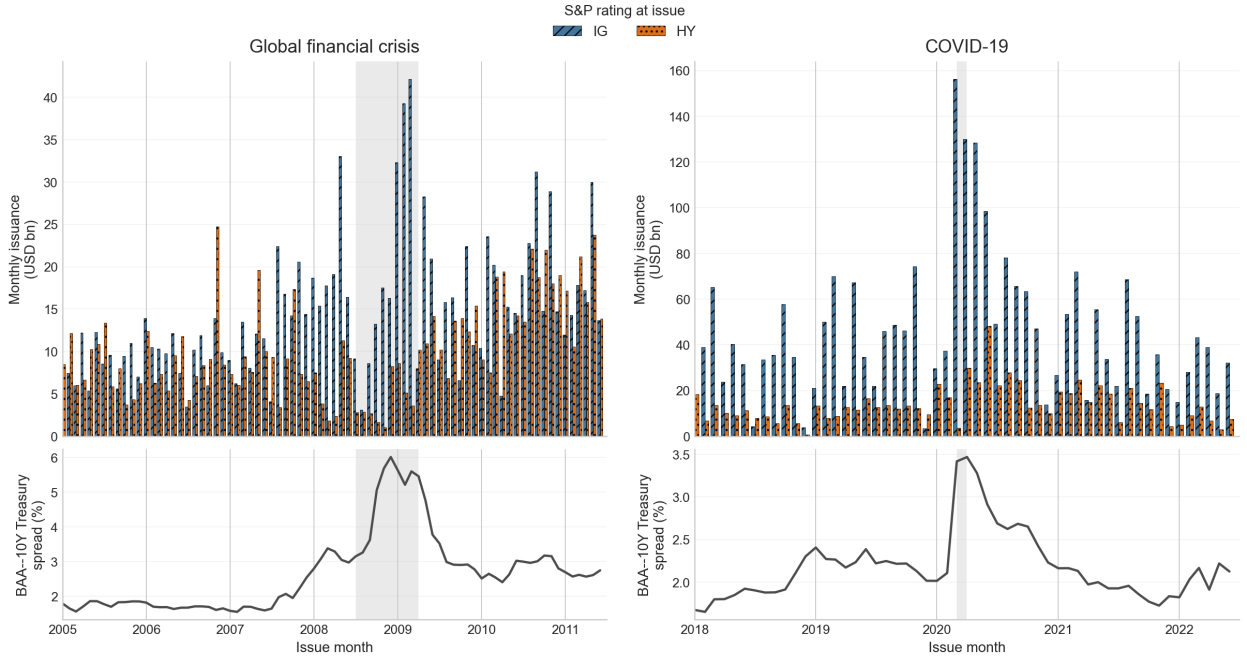
Figure 7: Cost-wedge decomposition over time: IG vs. HY fixed-rate bonds (annual).



Note: **Sample and unit.** Fixed-rate bond-years in the decomposition sample, shown separately for IG (top) and HY (bottom). Each bar is the equal-weighted mean in a calendar year, in percentage points per year; the red line is the realized wedge, and the number of bond-years varies by year. **Construction.** Equation (9) defines channels A–E, issuance-cost differential I, and the residual. Credit-spread curves are fitted under the isotonic non-decreasing prior described in Online Appendix A.1. Bonds with contractual tenor outside the supported 2–30-year curve grid are excluded because the benchmark would otherwise require extrapolation.

the

Figure 8: Monthly IG and HY primary issuances and credit spreads



Note: **Sample.** FISD nonfinancial U.S. dollar corporate bond issues, one observation per issue; unrated bonds are dropped. Fine S&P ratings at issue, including plus/minus notches, are collapsed to IG (BBB− or above) and HY (BB+ or below). The columns cover January 2005–June 2011 around the GFC and January 2018–June 2022 around COVID. The shaded stress windows are July 2008–March 2009 and March 2020. Average monthly IG issuance rose from $11.7 billion in the unshaded January 2005–June 2008 period to $20.2 billion in the GFC window, while HY issuance fell from $8.3 billion to $4.0 billion. Average monthly IG issuance rose from $36.6 billion in the unshaded January 2018–February 2020 period to $155.7 billion in March 2020, while HY issuance fell from $11.1 billion to $3.6 billion. **Construction.** Upper panels report monthly issuance volume in USD billions, with IG and HY bars shown side by side. Lower panels report the monthly mean Moody's seasoned-BAA yield minus the 10-year Treasury yield (FRED `BAA10Y`, percent). This pinned series is a credit-spread proxy, not an option-adjusted spread.

Online Appendix

The Cost of Maturity: Do Firms Overpay for Long-Term Debt?

Henri Servaes* Dian Xu†

August 27, 2026

A Supplementary Material

A.1 Yield Curve Construction

The bond-year counterfactual yields are constructed from monthly Merrill Lynch and Trade Reporting and Compliance Engine (TRACE) rating-specific corporate yield curves from January 1997 to July 2025. The main paper defines the counterfactual. This appendix explains how we construct the TRACE-based yield curves, extend them back to 1997 with Merrill Lynch data, and validate them against an independent source.

Methodology overview. For the baseline corporate yield curves using TRACE data, the construction proceeds as follows: (1) clean TRACE transactions; (2) build bond-day volume-weighted prices and yields-to-maturity; (3) merge bond characteristics from FISD; (4) benchmark to a fitted Treasury zero-coupon curve; (5) compute duration-matched corporate spreads at the bond-day level; (6) aggregate to rating $\times$ maturity cells using cell-month medians; (7) smooth the spread term structure with shape-preserving interpolation; (8) rebuild counterfactual yield curves at standard tenors for use in the bond-year panel.

TRACE cleaning (Step 1). We follow the Dick-Nielsen, Feldhütter, and Lando (2012) cleaning procedure: keep newly reported trades only; drop corrections, cancels, and reversals; remove inter-dealer buy-side duplicates; drop prices outside $[10, 300]$; drop zero

*London Business School, CEPR, and ECGI. Email: hservaes@london.edu.
†London Business School. Email: dxu@london.edu.

or missing quantities; collapse same-timestamp duplicates. The remaining transactions enter the subsequent steps.

Bond-day prices and yields (Step 2). For each CUSIP-day we compute the volume-weighted average price (VWAP) across all valid intra-day trades, following the market-microstructure convention of Bao, Pan, and Wang (2011). The median price is retained as a robustness check. Yield-to-maturity is then computed from the clean price, accrued interest, coupon, and maturity date.

Bond characteristics and sample (Step 3). TRACE is merged by CUSIP to the Mergent Fixed Income Securities Database (FISD). The baseline sample retains USD-denominated, fixed-rate, non-convertible, non-putable bonds. Callable bonds are kept.[1] Ratings used for yield-curve cells are assigned from S&P Capital IQ credit rating (Local Currency LT) at the TRACE trade date.

Treasury benchmark (Step 4). The risk-free reference is the Gürkaynak–Sack–Wright (GSW) fitted zero-coupon Treasury curve (Gürkaynak, Sack, and Wright, 2007), updated daily by the Federal Reserve Board. It is the conventional risk-free benchmark choice as it gives smooth zeros at any maturity and avoids the kinks and forward-rate discontinuities of linear Constant Maturity Treasury (CMT) interpolation.

Duration-matched spreads (Step 5). For each bond-day, the corporate spread is the bond yield minus the GSW zero yield at the bond's modified duration:

$$s_{i,t} \;=\; y_{i,t} - r_t^{\mathrm{GSW}}\big(\mathrm{Dur}_{i,t}\big).$$

Duration-matching (rather than maturity-matching) controls for coupon and embedded-option timing of the bond's cash flows.

Cell aggregation (Step 6). Bond-day spreads are aggregated within each rating $\times$ maturity $\times$ month cell using the *median* (volume-weighted average is the robustness alternative). The median is preferred at the cell level because it is robust to outlier trades,

[1] Callable bonds account for 70–80% of FISD issuances, so excluding them would substantially reduce coverage, especially for investment-grade (IG) bonds. Pooling different call structures within each rating-by-maturity cell reduces the influence of any one call provision, but the fitted spread may still contain an average call premium.

particularly in thin high-yield (HY) buckets. Rating buckets follow the standard mapping (AAA, AA, A, BBB, BB, B, CCC and below); maturity buckets cover 1, 2, 3, 5, 7, 10, 15, 20, and 30 years, as in Friewald, Jankowitsch, and Subrahmanyam (2012).

Smooth interpolation (Step 7). Within each rating-month, an isotonic (non-decreasing) constraint is first applied to the node spreads, then cell medians are interpolated across maturity using shape-preserving *Piecewise Cubic Hermite Interpolating Polynomial* (PCHIP; Fritsch and Carlson, 1980).[2] PCHIP is monotonic at observed nodes, prevents the spurious oscillation typical of cubic splines on sparse data, and does not impose a parametric Nelson–Siegel functional form on the credit-spread component. Due to sparse data beyond 30 years, the bond-year panel restricts the contractual-tenor sample to $T \in [2, 30]$ years.[3]

Counterfactual yields (Step 8). For each issuance date t, rating b, and target tenor τ, the counterfactual issuance yield is rebuilt as $y^{\mathrm{CF}}_{b,t}(\tau) = r^{\mathrm{GSW}}_t(\tau) + \hat{s}(b, t, \tau)$, where $\hat{s}$ is the PCHIP-fitted rating-month spread curve evaluated at τ. We use these yields to calculate the bond-year wedge. The main bond-year panel sample uses bonds rated AAA through B at issuance; bonds rated CCC-and-below are excluded due to sparse data and distressed transactions that make their yield curves too noisy for stable headline pricing.[4]

Pre-TRACE extension using Merrill Lynch data. The pre-TRACE extension uses the Merrill Lynch bond-day price/yield panel, which contains 26,729,731 bond-day observations before screening. After the bond filters and trade-date rating assignment, 17,041,306 observations on 15,290 bonds remain, spanning January 1997 to October 2012. Because the Merrill Lynch data do not contain TRACE transaction volume, dealer-side, correction, cancel, or timestamp fields, the transaction-level filters are not applicable, and the construction begins at the bond-day stage of the procedure. We match each Merrill Lynch bond to its FISD terms using the FISD issue identifier or, when needed, the full CUSIP. We

[2]The isotonic (non-decreasing) step limits maturity-composition noise in sparse HY cells. Helwege and Turner (1999) show that pooled speculative-grade bond yield curves can appear downward sloping because issuer composition changes with maturity, even though matched bonds from the same issuer typically have upward sloping yield curves. In practice, applying an isotonic prior affects the estimated yield curve almost exclusively on the HY short end and leaves IG curves essentially unchanged—the headline IG results are therefore robust to this choice.

[3]Extrapolation beyond the longest observed node is guarded: the curve is held flat at the endpoint rather than extrapolated.

[4]cells for ratings of CCC and below have few observations and many distressed-price prints, making the fitted yield curve unreliable for the main counterfactual. We nevertheless use the CCC-and-below curve at rollover dates when an issuer that entered the sample above CCC is later downgraded.

then apply the same fixed-rate, non-convertible, non-putable, U.S.-dollar filters, assign Capital IQ Local Currency LT ratings at each Merrill observation date, compute GSW-matched spreads, form monthly rating $\times$ maturity cells, and fit the same isotonic PCHIP term structure.

The Merrill Lynch–TRACE overlap sample contains 6,309 matched rating-tenor-month observations from February 2003 to October 2012. Figure OA1 plots the overlap for both IG (BBB) and HY (B) bonds at the 5- and 10-year tenors. The series generally move together, but the gaps differ by rating and period: the BBB gap is most visible around the 2008–2009 financial crisis, while the B gap is largest in the early 2000s.

[Figure OA1 about here.]

The extended yield curves use adjusted Merrill values before February 2003 and TRACE values from February 2003 onward. For each rating bucket and tenor, the baseline adjustment is the median TRACE-minus-Merrill spread difference in the adjacent 5.5-year overlap window, February 2003–July 2008. This window matches the approximate length of the Merrill-only extension and ends before the global financial crisis. The resulting median difference aligns each pre-TRACE Merrill rating–tenor curve with the TRACE scale.

Coverage and validation. The combined yield-curve file spans January 1997 to July 2025: 3,600 adjusted Merrill Lynch rating-tenor-month rows covering January 1997 to January 2003 and 13,914 TRACE-based rows from February 2003 onward, yielding 17,514 rating-tenor-month observations across the six broad rating groups. Table OA1 summarizes the baseline specification; Figures OA2 and OA3 below report the external validation of the TRACE-based segment against LSEG (formerly Refinitiv) corporate yield curves. The two sources produce similar IG term structures for most ratings buckets.

[Table OA1 about here.]

The specification in Table OA1 records the construction choices behind every counterfactual yield used in the paper. The two figures that follow benchmark this construction against an external commercial pricing source. Figure OA2 reports the headline IG validation at the 10-year tenor: across the 2010–2025 overlap window, the average Pearson correlation between the TRACE-fitted and LSEG evaluated dealer-quoted IG curves is 0.984, and the average mean absolute error is 23 bp. Speculative buckets show wider gaps that we trace to thinner speculative-grade TRACE coverage and LSEG's reliance on dealer evaluations rather than transactions.

[Figure OA2 about here.]

Figure OA2 shows that, at a single tenor, the two pricing sources have similar numbers. The cross section in Figure OA3 extends the comparison to the entire term structure at four representative year-ends. For both A and BBB ratings, which are the IG buckets that drive most of the headline wedge results, the curves have similar shapes with small level differences consistent with the MAE bands of Figure OA2. The mean-reversion shape and the post-2022 inversion of the BBB curve appear in both sources.

[Figure OA3 about here.]

A.2 Mergers and Counterfactual Ratings

This appendix discusses how we deal with bond rollovers after the issuer has been acquired. A CUSIP is a bond's serial number, and its first six characters are an issuer label. An old bond can keep its issuer label after the company is sold, while later acquirer debt can use a different label. FISD preserves the bond record but does not identify the buyer.

We match WRDS SDC control acquisitions to the 4,506 FISD issuer groups represented in the issuance sample. The target CUSIP6 must match a sample issuer, the transaction must be a completed merger or control acquisition ending with at least 50% ownership, and a bond is affected only if it was outstanding on the completion date. SDC contains 1,844 such transactions involving a sample target; 1,090 occurred while at least one sample bond was outstanding. Those transactions affect 3,390 bonds and 1,018 distinct issuer groups (22.6% of the sample issuer groups). The deal, bond, and issuer-group counts therefore measure different units. At the bond level, a rated-acquirer link is available for 1,945 of the 3,390 affected bonds; no such link is available for the other 1,445 bonds.

The key question is whether we can identify the surviving rated firm. For 324 acquired issuer groups, Capital IQ identifies a rated acquirer that later issues another bond passing our fixed-rate sample screens. Those later bonds are priced with the acquirer's rating. In 319 of these groups, the new bond uses a CUSIP6 different from the target's. Thus a CUSIP change creates no ambiguity once the acquirer is observed: the new bond follows the acquirer's identifier and rating. The cases that remain uncertain are the 1,445 outstanding target bonds for which no rated-acquirer link can be recovered.

We next ask whether using the surviving firm's identity can change the rating. We observe both the target's and the acquirer's rating on the completion date for 420 deals. The two ratings differ in 290 deals, including 141 in which one firm is IG and the other is HY. This is why the counterfactual rating should follow the surviving firm rather than the target's old CUSIP. At the bond level, both completion-date ratings are observable for 1,476 linked bonds: 571 have the same notch and 905 have different notches.

We address this remaining uncertainty with two sensitivity checks. First, we discard the entire history of those 1,445 unresolved bonds. This removes 7,024 3-year bond-years; the full-sample mean moves from 1.85 to 1.88 percentage points per year, and the largest change among the full, IG, and HY samples is 2.6 basis points. An even stronger check avoids assuming that a corporate-control link identifies the bond's legal obligor or guarantor: it discards all 3,390 SDC-identified acquired-company bonds. This removes 18,562 bond-years and changes the group means by at most 4.9 basis points. The headline result is therefore robust to acquisition-related rating assignment.

A.3 10-K Interest-Rate Swap Evidence

Some issuers can use pay-floating interest-rate swaps to convert fixed-coupon debt into floating-rate exposure. If this practice were widespread relative to the long-bond portfolio, a bond-level fixed-rate wedge would overstate the issuer's effective fixed-rate exposure. This appendix describes the information we collect to address this concern.

Sample and measurement. We search the 10-K filings of the nonfinancial issuers in the FISD bond sample for disclosed interest-rate-swap notional and explicit pay-fixed or pay-floating language. Direction is recorded only when the filing states it; it is not inferred from hedge labels. The panel contains 21,802 company filing-years over January 1997–July 2025. We identify positive interest rate swap (IRS) notional in 7,993 company-years (36.7%) and explicit direction in 3,857 of those observations (48.3%). We exclude 58 company filing-years whose gross notional exceeds 200% of interest-paying debt, a screen

against obvious unit or extraction errors. This approach extends the binary swap-use evidence in Banerjee, Caballero, Kharroubi, Spigt, and Zakrajšek (2023) by recovering the direction needed to quantify how swaps change fixed-rate exposure.

Swap-adjusted result. The relevant denominator is aggregate FISD public-bond principal with more than five years remaining in matched nonfinancial company-years. Disclosed pay-floating notional equals 5.14% of that denominator. This gross share gives swaps the largest possible reduction in measured fixed-rate exposure: it treats every disclosed pay-floating dollar as offsetting long fixed-rate bonds and does not subtract a firm's simultaneous pay-fixed positions. Once the two directions are netted, the disclosed position is only -0.87% of the same denominator; the negative sign means that disclosed pay-fixed notional is slightly larger than pay-floating notional. We therefore use 5.14% as a conservative stress-test adjustment. Even under that convention, swaps remove only a small part of the measured fixed-rate wedge.

Figure OA4 shows that this conclusion does not come from one unusual year. Swap disclosure becomes more common over time, but the net position remains small relative to long public bonds.

[Figure OA4 about here.]

Figure OA5 explains why gross swap use does not imply a large reduction in fixed-rate exposure. Firms report both directions, and floating-to-fixed observations are more common over much of the sample. The lower panel reports gross fixed-to-floating notional and the net swap position, each as a percentage of long public bonds.

[Figure OA5 about here.]

A.4 Issuance Cost Model

This appendix documents the issuance-cost model used when the underwriting gross spread g_i is missing and for each hypothetical counterfactual bond in the annual cost of debt (ACD) calculation. Gross spread is observed for 48.4% of bonds in the filtered FISD issue sample. The reported regressions describe gross spreads in levels; predictions use the same preferred covariate specification with a log outcome, as described below.

Why is the gross spread missing for some bond issues? The gross spread is the underwriter's deal-specific fee. Rule 144A offerings are privately marketed, so their offering

materials need not be filed as publicly registered offering documents. Medium-term note (MTN) issues are drawdowns from a standing program, where compensation can instead be recorded at the program, agent, or pricing-supplement level. Thus the source record often has no issue-level gross-spread figure; a missing value is not treated as a zero fee.

Sample and outcome. The model is estimated only for fixed-rate U.S.-listed non-financial corporate bonds in FISD issued over January 1997–July 2025 with a valid (non-missing, positive) gross spread ($N = 10,222$ after requiring the controls used in the regressions). This fixed-rate restriction matches the bonds for which the model supplies issuance costs in the headline ACD calculation. Under the same issue filters, the sample construction retains 21,110 fixed-rate issues and separately identifies 1,489 floating-rate issues. Thus floating-rate issues are only 6.6% of the combined 22,599 fixed/floating universe and are analyzed separately rather than pooled into this pricing model. The reported dependent variable is gross spread as a percentage of face value. The prediction fit instead uses its natural log so that retransformed issuance-cost predictions are strictly positive.

Preferred covariate specification (Model (4)). We use the same covariate specification as the main paper's issuance-cost equation and estimate the following model:

$$h(g_i) = \alpha + \mathrm{FE}_{r(i)} + \mathrm{FE}_{t(i)} + \beta \log(T_i) + \eta_1 \mathrm{Size}_i + \eta_2 \mathrm{Size}_i^2 + \lambda_{r(i)} \log(T_i) + \delta' X_i + \varepsilon_i.$$

Here $h(g_i) = g_i$ in the level regressions reported in Table OA2, whereas $h(g_i) = \log(g_i)$ in the prediction model. The explanatory variables are as follows: $r(i)$ indexes rating-bucket fixed effects, and $t(i)$ indexes issue-year fixed effects. T_i is contractual tenor, and Size_i is the offering amount in units of \$100M. The vector X_i contains callable, senior-secured, and subordinated indicators. Rating and year fixed effects capture level differences, while $\lambda_{r(i)} \log(T_i)$ allows the tenor slope to vary by rating. The regressions use heteroskedasticity-robust standard errors.

Table OA2 reports five specifications of this model. Model (4) replaces log issue size with linear and quadratic terms in the offering amount. Model (5) adds the VIX market-volatility index only as a check for the years for which VIX is available; VIX is not part of the preferred model and does not define the regression sample.

[Table OA2 about here.]

Interpretation and prediction. Because the model allows the tenor slope to vary by rating, each rating coefficient must be read together with its rating–tenor interaction. For counterfactual predictions, we use the controls in Model (4) but refit the model with log gross spread as the outcome; this produces positive predicted spreads without imposing a floor or cap. Exponentiating a log prediction returns a median on the original scale. To obtain the mean gross spread, we apply the standard Duan adjustment (Duan, 1983): every prediction is multiplied by the sample mean of the exponentiated residuals, $\exp(\widehat{\varepsilon}_i)$. The estimated factor is 1.049, so the adjustment raises each prediction by 4.9%. We use the same rule to fill missing gross spreads and to price counterfactual issues; no residual is drawn for an individual bond.

Robustness: cell-median lookup grid. A natural concern is that gross spreads are negotiated fees anchored to round-number conventions, not market-clearing prices, and that a parametric model may overfit. As a nonparametric robustness check we construct a rating $\times$ tenor cell-median lookup grid: within each cell the predicted gross spread is the within-cell median of winsorized (1%/99%) actual spreads, with tenor- or rating-median fallbacks when fewer than 20 observations are available. This check is motivated by the pronounced round-number clustering of underwriting spreads, the corporate analog of the IPO underwriting-spread clustering in Chen and Ritter (2000).

Cell-median comparison. We compare the positive parametric predictions with the cell-median lookup on their common sample of 10,222 fixed-rate issues. The two predictions have a correlation of 0.897 and differ by only 0.093 percentage points (9.3 basis points) on average. Their mean absolute errors against observed gross spreads are 0.134 percentage points for the parametric model and 0.112 percentage points for the cell-median lookup. The lookup is transparent and robust to discrete fee quoting; the parametric model is continuous and accommodates characteristics outside densely populated cells. We use the log-outcome Model (4) prediction in the baseline and retain the cell-median method as a nonparametric robustness check.

A.5 Prepayment and Effective Tenor

This appendix documents the FISD prepayment and price records used in the prepayment adjustment. Each balance-change record reports the date of the change, the price paid, and the principal that remains outstanding. We use these realized balance changes, rather than a callable-bond flag, to measure how long each dollar of principal remains

outstanding.

Announced and realized call prices. FISD separates future *announced* calls from realized balance changes. The announced-call file reports the call date and price for announced partial or full calls; it is not the bond's complete contractual call-option schedule. In that file, 2,924,874 of 2,944,293 rows carry valid call prices. Within the filtered FISD issue sample, schedule coverage is less than universal: 11,812 of 25,496 callable bonds have at least one valid announced call price. Coverage rises to 65.3% among callable prepaid bonds and 72.9% among callable bonds that experience an entire-call action. The balance-change file contains the realized price used for the redemption-premium component. Table OA3 reports coverage separately for each FISD action type.

[Table OA3 about here.]

Effective tenor. The contractual tenor at issuance can overstate principal exposure when FISD reports a principal reduction before maturity. For original principal F_i, let $B_{i,j}$ be the balance outstanding from date $t_{i,j-1}$ to the next reported reduction date $t_{i,j}$, with the last interval ending at contractual maturity t_i^M. We define

$$\widetilde{T}_i = \frac{1}{F_i} \sum_j B_{i,j}(t_{i,j} - t_{i,j-1}).$$

Thus each repaid share stops on its reported reduction date and any remaining principal stays exposed through contractual maturity. A zero balance ends the bond; in the occasional cases where the prepayment event doesn't specify a price and amount, effective tenor is assumed to equal contractual tenor. Within FISD action codes that indicate prepayment types, a lower reported balance identifies a reduction. We use the reported remaining balance when available; otherwise, we infer the balance change from the reported transaction amount. Mean effective tenor is 10.22 years against 11.43 contractual years. The bond-year panel and ACD wedge use this adjustment, so repaid principal is not treated as outstanding through contractual maturity.

[Figure OA6 about here.]

In Figure OA6, the left histogram shows that most mass in $\widetilde{T}_i$ lies between three and ten years, with a long right tail from long-dated bonds held to maturity. The right-panel scatter shows that observations either sit on the 45° line (no usable principal reduction) or below it (a principal reduction shortens principal-time exposure), and the gap between T_i

and $\widetilde{T}_i$ widens with contractual tenor. The figure therefore shows how reported principal reductions shorten the time for which the issuer pays borrowing costs on the original bond.

A.6 Exact ACD and Annual Cost Decomposition

Expressed as cash flows per dollar of original par, the issuer receives net proceeds $p_{0i} - g_i$, pays coupons $c_{i,m}$ at times $\tau^C_{i,m}$, and retires share $a_{i,k}$ at price $p^R_{i,k}$ and time $\tau^R_{i,k}$. The exact ACD is the internal rate of return r_i solving

$$p_{0i} - g_i = \sum_{m \in \mathcal{C}_i} \frac{c_{i,m}}{(1+r_i)^{\tau^C_{i,m}}} + \sum_{k \in \mathcal{R}_i} \frac{p^R_{i,k} a_{i,k}}{(1+r_i)^{\tau^R_{i,k}}},$$
$$\sum_{k \in \mathcal{R}_i} a_{i,k} = 1.$$

This equation allows maturity payments, calls, partial calls, sinking-fund payments, tenders, and buybacks.

Additive annual measure. Define principal-weighted effective tenor, total redemption premium, and its annualized value as

$$\widetilde{T}_i = \sum_{k \in \mathcal{R}_i} a_{i,k} \tau^R_{i,k}, \qquad \rho_i = \sum_{k \in \mathcal{R}_i} (p^R_{i,k} - 1) a_{i,k}, \qquad \xi_i = \rho_i / \widetilde{T}_i.$$

Let y_i be the issue yield when reported (otherwise, we use the stated coupon for fixed-rate bonds, and the realized annual coupon for floating-rate bonds). The additive measure treats $\widetilde{T}_i y_i$ as the regular borrowing-cost allocation and allocates the two one-time dollar-of-par costs over the same principal-time exposure:

$$\mathrm{ACD}_i^{\mathrm{approx}} = \frac{\widetilde{T}_i y_i + g_i + \rho_i}{\widetilde{T}_i} = y_i + \frac{g_i}{\widetilde{T}_i} + \xi_i.$$

The equality follows because regular borrowing cost, issuance cost, and the redemption premium are separate dollar amounts per dollar of original par and are all divided by the same exposure measure. It requires that (i) g_i and ρ_i are measured per dollar of original par, (ii) $\widetilde{T}_i$ is the common principal-weighted exposure denominator, and (iii) the reported y_i is used as the rate for regular borrowing cost over that exposure. It does *not* assume that fees or redemption premiums occur uniformly through time.

Accuracy relative to the exact IRR. The same additive expression is also the first-order approximation to the exact IRR under standard small-cost conditions. Define the signed premium on retirement event k as $d_{i,k} = (p^R_{i,k} - 1)a_{i,k}$, so $\rho_i = \sum_k d_{i,k}$. Hold retirement shares and dates fixed, assume $\widetilde{T}_i > 0$, and consider a local expansion around a par, zero-cost bond. Specifically, take the regular yield and coupon rate, the underwriting fee g_i, and $\sum_k |d_{i,k}|$ to be small quantities of order η_i, and assume the gross issue price satisfies $p_{0i} = 1 + O(\eta_i^2)$. The last condition says that the offering-price discount or premium is negligible at first order, as is typical for new bonds issued near par. With no underwriting fee and par repayment at every retirement date, the regular cash-flow yield is y_i up to second-order terms. At the zero-rate, par-repayment benchmark, the change in present value from a one-unit change in the annual rate is the principal-time exposure $\sum_k a_{i,k}\tau^R_{i,k} = \widetilde{T}_i$.

A first-order expansion of the exact present-value equation around that benchmark therefore gives

$$\widetilde{T}_i(r_i - y_i) = g_i + \sum_k d_{i,k} + O(\eta_i^2), \qquad \text{or} \qquad r_i = y_i + \frac{g_i}{\widetilde{T}_i} + \frac{\rho_i}{\widetilde{T}_i} + O(\eta_i^2).$$

The omitted terms are second order under these assumptions: they consist of products between rates and one-time costs, discounting adjustments that depend on payment dates, squares of the fee and redemption-premium perturbations, and the near-par issue-price deviation. If that price deviation were instead a first-order quantity, the expansion would also contain a timing term because the issue discount or premium would be spread over realized rather than contractual life. The additive measure does not estimate that term separately. The worked example below therefore checks the total approximation error using the actual issue price and dated cash flows rather than relying only on the local assumptions.

The additive measure is useful even apart from this local approximation. It keeps regular borrowing cost, issuance cost, and redemption premium separately visible; uses the same principal-exposure denominator for bonds and counterfactual refinancing periods; aggregates linearly across observations; and avoids re-solving an IRR for every component. At-par bonds have $\rho_i = 0$; if $g_i = 0$ as well, the measure reduces to y_i.

Worked-example check. Table OA4 uses the same Oklahoma Gas & Electric bond as the main paper's worked example. It applies the exact cash-flow equation to the issue price net of the 87.5-bp gross spread, the semiannual 6.5% coupons, and the 103.25 call price on August 1, 2014. The annual allocation is 7.55 bp below the exact IRR (1.08% of the

total annual cost and 2.4% of that bond's 3Y wedge). This small difference is the expected consequence of discounting and payment timing, while the additive measure remains appropriate for decomposition and aggregation.

[Table OA4 about here.]

Final-rollover tenor sensitivity. The baseline counterfactual ends on each bond's maturity or prepayment date. When the remaining time is shorter than the benchmark tenor, the final counterfactual bond has a shorter maturity. This occurs for bonds ending at contractual maturity and for prepaid bonds. As a sensitivity, we keep the comparison period and principal exposure unchanged but price the final counterfactual bond at the full benchmark tenor. Table OA5 compares this convention with the baseline.

[Table OA5 about here.]

Table OA6 applies the same sensitivity to the Oklahoma Gas & Electric worked example used in the main paper. The table isolates the final counterfactual bond, so the comparison shows the mechanical effect of pricing it at the full benchmark tenor while leaving the realized call date fixed.

[Table OA6 about here.]

A.7 Rating-Notch Interpolation Sensitivity

This appendix examines whether broad rating buckets in the counterfactual yield curve drive the cost-of-maturity wedge. The baseline maps each fine S&P rating to one of the broad Merrill Lynch/TRACE yield-curve buckets (AAA/AA, A, BBB, BB, B, CCC). We build a notch-level sensitivity that linearly interpolates between adjacent anchor curves using the fine S&P rating from Capital IQ at each rollover date. We map ratings to a numeric notch scale (AAA$=1$, AA$^{+}=2$, …, C$=21$, D$=22$). A rating between two anchors receives a linear combination of their curves for the same month and tenor. Ratings above AA retain the pooled AAA/AA curve, while ratings below CCC retain the CCC-and-below curve. The sensitivity changes only the counterfactual yield; observed ACD, issuance costs, rollover timing, and sample membership are held fixed.

Table OA7 shows that this fine-notch interpolation has little effect on the aggregate wedge. The table reports the baseline and notch-interpolated counterfactual yields and wedges for the full, fixed-rate, IG, and HY samples. The subgroup adjustments have offsetting signs, and the full-sample change is small relative to the level of the headline

maturity-cost wedge. Using broad rating buckets that combine plus and minus ratings therefore does not drive the headline result.

[Table OA7 about here.]

Table OA8 explains where the small aggregate change comes from. Anchor ratings are unchanged by construction because they already sit on an estimated benchmark curve. The adjustments are concentrated in plus/minus notches between anchors, with positive and negative changes offsetting in the aggregate. The table reports the sample size and wedge change for every notch, making the offsetting changes transparent.

[Table OA8 about here.]

A.8 Short-Term Debt, Ratings, and the Adjusted Wedge

A possible cost of the short-rolling counterfactual is a weaker credit rating. If rating agencies pay attention to upcoming maturities in addition to the overall level of debt, the may punish firms if too much of their debt is short-dated. We assess this channel by estimating how much debt ratings are affected by the Short-Debt Share ($\text{SDS} = \text{DLC}/(\text{DLC} + \text{DLTT})$), where DLC is debt due within one year and DLTT is long-term debt. We then predict the rating change under each rollover tenor and reprice borrowing and issuance costs. To translate the bond strategy into a firm-year measure, we hold total debt fixed and assume that the firm's k-year counterfactual bonds are spread evenly across issue years. A fraction $1/k$ of long-term debt then comes due within the next year, so the counterfactual share is

$$\text{SDS}_{it}^{\text{CF}}(k) = \frac{\text{DLC}_{it} + \text{DLTT}_{it}/k}{\text{DLC}_{it} + \text{DLTT}_{it}} = \text{SDS}_{it} + \frac{1 - \text{SDS}_{it}}{k}, \qquad k \in \{1, 2, 3, 4, 5\}. \tag{1}$$

Following Baghai, Servaes, and Tamayo (2014), we augment a standard ratings-determinants model with SDS and its interaction with book leverage. Numeric ratings run from 1 (AAA) to 21 (C):

$$\text{Rating}_{it} = \alpha + \beta_1 \,\text{Lev}_{it} + \beta_2 \,\text{SDS}_{it} + \beta_3 \,(\text{Lev} \times \text{SDS})_{it} + \gamma' \mathbf{X}_{it} + \delta' \mathbf{D} + \varepsilon_{it}.$$

The control vector contains leverage, coverage, profitability, volatility, size, and asset composition. Columns (3) and (5) add Dimson beta and idiosyncratic risk, whose CRSP inputs now cover the full 1997–2025 window. Specifications vary the estimator and fixed effects.

[Table OA9 about here.]

Book leverage is associated with worse ratings. By contrast, neither the Short-Debt-Share level effect nor its interaction with leverage is stable across estimators and control sets: signs change, and the interaction is not consistently significant in the models that include firm fixed effects. Nevertheless, we employ the augmented ratings model to determine the effect of counterfactual maturities on the firm's rating. The underlying annual series covers 23,575 rated firm-years across 1,892 firms. Ratings are year-end Capital IQ Local Currency LT issuer ratings, matched to firms through issuer identifiers and supplemented, when needed, with ratings linked through the Compustat firm identifier or the bond issue. The preferred firm-fixed-effect model in column (5) uses 18,466 firm-years, has $R^2 = 0.397$, and an in-sample mean absolute error of 0.72 notches. 75.8% of predictions are within one notch, 94.7% are within two, and in-sample IG/HY classification accuracy is 100.0%.

For each firm-year and rollover tenor k, we use equation (1), apply the preferred model's SDS coefficients to obtain a continuous predicted rating shift.

[Table OA10 about here.]

Table OA10 shows that mechanically shortening maturity raises SDS and actually predicts a small rating improvement (a decline in the numeric predicted rating); however, not a single firm-year worsens by more than half a notch. The rating-level impact is therefore economically small across all counterfactual tenors.

References

Baghai, Ramin P., Henri Servaes, and Ane Tamayo, 2014, Have rating agencies become more conservative? Implications for capital structure and debt pricing, *Journal of Finance* 69, 1961–2005.

Banerjee, Ryan, Julián Caballero, Enisse Kharroubi, Renée Spigt, and Egon Zakrajšek, December 2023, Interest rate risk of non-financial firms: Who hedges and does it help? BIS Bulletin 81, Bank for International Settlements. URL `https://www.bis.org/publ/bisbull81.htm`.

Bao, Jack, Jun Pan, and Jiang Wang, 2011, The illiquidity of corporate bonds, *Journal of Finance* 66, 911–946.

Chen, Hsuan-Chi, and Jay R. Ritter, 2000, The seven percent solution, *Journal of Finance* 55, 1105–1131.

Dick-Nielsen, Jens, Peter Feldhütter, and David Lando, 2012, Corporate bond liquidity before and after the onset of the subprime crisis, *Journal of Financial Economics* 103, 471–492.

Duan, Naihua, 1983, Smearing estimate: A nonparametric retransformation method, *Journal of the American Statistical Association* 78, 605–610.

Friewald, Nils, Rainer Jankowitsch, and Marti G. Subrahmanyam, 2012, Illiquidity or credit deterioration: A study of liquidity in the US corporate bond market during financial crises, *Journal of Financial Economics* 105, 18–36.

Fritsch, F. N., and R. E. Carlson, 1980, Monotone piecewise cubic interpolation, *SIAM Journal on Numerical Analysis* 17, 238–246.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, The U.S. Treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Helwege, Jean, and Christopher M. Turner, 1999, The slope of the credit yield curve for speculative-grade issuers, *Journal of Finance* 54, 1869–1884.

Table OA1: Baseline yield-curve specification.

Component	**Specification**
Transaction data	Merrill Lynch pre-TRACE bond-days; TRACE Enhanced from February 2003 onward
Bond-day price	Volume-weighted average price (VWAP)
Bond filters	Fixed-rate, non-convertible, non-putable; callable included
Risk-free curve	GSW fitted Treasury zero-coupon curve
Spread definition	Duration-matched: $y_{i,t} - y_t^{\mathrm{RF}}(\tau_{i,t})$
Cell aggregation	Median spread within rating $\times$ maturity bucket
Curve interpolation	PCHIP under isotonic non-decreasing prior
Rating assignment	S&P Capital IQ Local Currency LT rating at trade date
Frequency	Monthly
Period	January 1997 – July 2025

Note: **Sample.** Merrill Lynch bond-days before TRACE and TRACE Enhanced bond-days thereafter, subject to the bond filters and period shown in the table. **Unit.** This is a specification table rather than an observation-level estimate, so it does not report an N. All choices—VWAP pricing, median cell aggregation, PCHIP interpolation under the isotonic prior—are paired with the validation checks reported in this appendix.

Table OA2: Issuance-cost panel: gross-spread regressions

	(1) Baseline	(2) + Bond Features	(3) + Rating × Tenor	(4) Quadratic Size	(5) + VIX
Credit Rating (ref: A-rated)					
AAA	−0.034	−0.013	−0.023	−0.009	−0.009
	(-1.61)	(-0.64)	(-0.33)	(-0.13)	(-0.13)
AA	−0.098***	−0.095***	−0.076***	−0.069***	−0.069***
	(-11.18)	(-10.97)	(-3.19)	(-2.91)	(-2.91)
BBB	0.077***	0.069***	0.142***	0.137***	0.137***
	(20.45)	(18.25)	(12.39)	(12.14)	(12.14)
BB	0.728***	0.653***	1.050***	1.041***	1.041***
	(34.32)	(31.85)	(10.55)	(10.55)	(10.55)
B	1.333***	1.182***	0.560	0.539	0.539
	(28.26)	(25.88)	(1.09)	(1.06)	(1.06)
Bond Characteristics					
log(Tenor)	0.217***	0.208***	0.227***	0.228***	0.228***
	(71.75)	(74.58)	(55.86)	(56.43)	(56.38)
log(Size)	−0.024***	−0.031***	−0.030***		
	(-5.62)	(-6.56)	(-6.50)		
Size ($100M)				−0.0094***	−0.0094***
				(-11.91)	(-11.89)
Size^2 ($100M)				0.000077***	0.000077***
				(4.76)	(4.75)
Rating × log(Tenor)					
AAA × log(Tenor)			0.004	0.004	0.004
			(0.15)	(0.17)	(0.17)
AA × log(Tenor)			−0.007	−0.004	−0.004
			(-0.57)	(-0.34)	(-0.34)
BBB × log(Tenor)			−0.031***	−0.031***	−0.031***
			(-6.15)	(-6.17)	(-6.18)
BB × log(Tenor)			−0.179***	−0.177***	−0.177***
			(-4.11)	(-4.10)	(-4.10)
B × log(Tenor)			0.294	0.302	0.302
			(1.23)	(1.27)	(1.27)
Bond Feature Controls	No	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
N	10,222	10,222	10,222	10,222	10,222
R^2	0.6071	0.6380	0.6416	0.6460	0.6460
Adj. R^2	0.6057	0.6366	0.6401	0.6445	0.6444

Notes: **Sample and unit.** Fixed-rate U.S.-listed nonfinancial corporate bond issues with a valid positive gross spread and the required issue controls. Under the same issue filters, 1,489 floating-rate issues are excluded, equal to 6.6% of the combined 22,599 fixed/floating issues. Each observation is one bond issue; N is the number of issues used in each regression column. The dependent variable is gross spread as a percentage of par. t-statistics (HC1 robust) in parentheses. ***, **, * denote significance at the 1%, 5%, 10% levels. Reference rating category: A. Gross spread is winsorized at the 1st/99th percentiles only to limit extreme tails and likely data errors. D-rated bonds are excluded (< 10 observations); CCC and below are excluded because of small sample size. VIX availability is not a sample restriction; Model (5) uses the VIX-available subsample only as a robustness check. Bond feature controls are callable, senior secured, and subordinated indicators. Model (4) defines the preferred covariate specification. For predictions, the same specification is refit with log gross spread. The residual-smearing correction multiplies its exponentiated fitted value by the single sample mean of exponentiated residuals, with no residual draw for individual bonds; the coefficients shown here are level-OLS estimates for interpretation. Model (5) adds VIX only as a non-preferred robustness check.

Table OA3: FISD announced and realized call-price availability

Record type	Unit	Denominator	Valid price	Share (%)
Panel A: Announced-call file				
All announced-call records	Rows	2,944,293	2,924,874	99.3
Callable sample bonds	Bonds	25,496	11,812	46.3
Callable prepaid bonds	Bonds	8,097	5,285	65.3
Callable entire-call bonds	Bonds	3,241	2,363	72.9
Panel B: Realized balance-change file				
Balance of Issue Called (B)	Bonds	2,685	2,621	97.6
Entire Issue Called (E)	Bonds	3,260	3,118	95.6
Issue Refunded (F)	Bonds	1	0	0.0
Issue Repurchased (IRP)	Bonds	78	13	16.7
Part of Issue Called (P)	Bonds	43	41	95.3
Sinking Fund Payment (S)	Bonds	22	22	100.0
Issue Tendered (T)	Bonds	1,046	992	94.8

Note: **Sample and unit.** Panel A uses FISD's announced-call file and reports either announced-call rows or unique bonds, as shown in the Unit column. Panel B uses unique bonds in the filtered FISD issue sample with the stated action type. There is no separate *N* column: "Denominator" is the relevant row or bond count, and "Valid price" is the subset with an admissible price. A valid announced-call price is between 50 and 200 percent of par on a row dated between issuance and contractual maturity. Panel B uses the FISD balance-change file for bonds in the filtered FISD issue sample with action dates strictly between issuance and maturity. A valid realized price is between 50 and 200 percent of par. Panel B includes every entire-call record, whereas the corresponding Panel A row begins with bonds that FISD labels callable. The two denominators therefore differ slightly.

Table OA4: Worked-example check of the annual-cost allocation

Measure	Value
Reported offering yield	6.510
Gross-spread allocation (0.875%/9.990 years)	0.088
Redemption-premium allocation (3.250%/9.990 years)	0.325
Additive annual cost	6.923
Exact cash-flow IRR	6.998
Difference: additive minus exact	-0.075
Absolute difference (bp)	7.55

Note: **Sample and unit.** One Oklahoma Gas & Electric bond and its dated cash flows; this worked example has one bond observation and therefore does not report a separate N. All values are annual percentage points except the final row. The exact IRR uses dated semiannual coupon payments from issuance on August 4, 2004 through the August 1, 2014 call, and the stated call price. The allocation uses the reported offering yield, gross spread, redemption premium, and the 9.990-year effective tenor.

Table OA5: Same-tenor final-rollover sensitivity

Sample	CF bonds	Bonds	Baseline wedge	Same-tenor wedge	Δ
All eligible 3-year bonds	30,447	11,024	2.12	1.97	-0.142
Prepaid bonds	11,319	4,839	3.13	2.91	-0.217
Short final counterfactual maturity	30,396	10,996	2.12	1.98	-0.143

Note: **Sample and unit.** Eligible observed bonds under the 3-year rollover counterfactual. "Bonds" counts unique observed bonds; "CF bonds" counts their hypothetical issues, where CF means counterfactual. The baseline prices the final counterfactual bond at its remaining life. The sensitivity instead uses the full 3-year curve and issuance-model tenor but leaves the repayment date, principal exposure, fee annualization, and every other counterfactual bond unchanged. Bond results use principal-exposure weights. Wedges are annual percentage points; Δ is same-tenor minus baseline.

Table OA6: Worked-example final-rollover sensitivity

CF bond	Start	End	Notional years	Base tenor	Base CF ACD	Same-3Y CF ACD	Base wedge	Same-3Y wedge	Δ
0	2004-08-04	2007-08-04	2.998	3.000	3.95	3.95	2.97	2.97	0.000
1	2007-08-04	2010-08-04	3.001	3.000	5.78	5.78	1.14	1.14	0.000
2	2010-08-04	2013-08-04	3.001	3.000	2.42	2.42	4.50	4.50	0.000
3	2013-08-04	2014-08-01	0.991	0.991	0.76	1.61	6.16	5.31	-0.848
Principal-weighted			9.990		3.72	3.81	3.20	3.12	-0.084

Note: **Sample and unit.** One Oklahoma Gas & Electric bond, with one row per hypothetical issuance under its 3-year rollover path; the table therefore has one observed bond and four counterfactual bonds rather than a separate N. CF means counterfactual. This table repeats the Oklahoma Gas & Electric worked example from the main paper and changes only the final 2014 counterfactual bond. That bond has a remaining tenor of 0.99 years in the baseline. The same-tenor sensitivity prices it as a 3-year counterfactual issuance at the issuer's contemporaneous A− rating, while keeping the comparison window and exact principal exposure through the observed call date. "Notional years" is the principal share multiplied by the years for which that counterfactual bond remains outstanding; it is the weight used in the final row.

Table OA7: Rating-notch interpolation sensitivity for counterfactual yields

Sample	Bonds	Rollovers	Base CF yield	Notch CF yield	Δ yield	Base wedge	Notch wedge	Δ wedge
All eligible 3-year bonds	10,877	30,021	4.75	4.80	0.045	2.12	2.07	-0.045
Fixed-rate bonds	10,673	29,628	4.79	4.83	0.046	2.16	2.12	-0.046
Floating-rate bonds	204	393	3.08	3.11	0.029	-0.41	-0.43	-0.029
IG at issuance	6,103	18,091	3.42	3.49	0.063	1.65	1.59	-0.063
HY at issuance	4,774	11,930	6.45	6.47	0.022	2.71	2.68	-0.022

Note: **Sample and unit.** Eligible observed bonds under the 3-year rollover counterfactual, partitioned by rate type or issuance rating. "Bonds" is the number of unique observed bonds and "Rollovers" is the number of hypothetical bond issues generated across them; these two columns replace a generic N. The baseline counterfactual yield uses broad rating buckets. The notch sensitivity linearly interpolates between adjacent benchmark bucket curves in S&P notch space, with the pooled AAA/AA curve anchored at AA. Wedges are annual percentage points. Δ is notch interpolation minus baseline.

Table OA8: Rating-notch interpolation sensitivity by rollover rating

Prevailing rating	Bonds	Rollovers	Base CF yield	Notch CF yield	Δ yield	Base wedge	Notch wedge	Δ wedge
AAA	70	166	2.60	2.60	0.000	0.88	0.88	0.000
AA+	56	118	2.16	2.16	0.000	0.74	0.74	0.000
AA	232	418	3.01	3.01	0.000	1.07	1.07	0.000
AA-	439	731	2.89	2.98	0.092	1.08	0.99	-0.092
A+	719	1,216	3.47	3.38	-0.085	1.25	1.33	0.085
A	1,318	2,402	3.11	3.11	0.000	1.55	1.55	0.000
A-	1,496	2,408	2.90	3.08	0.178	1.87	1.69	-0.178
BBB+	1,982	3,276	3.42	3.24	-0.184	1.62	1.80	0.184
BBB	2,438	4,347	3.39	3.39	0.000	1.96	1.96	0.000
BBB-	1,816	2,850	3.51	4.08	0.571	2.34	1.77	-0.571
BB+	1,209	1,714	4.96	4.44	-0.519	1.05	1.57	0.519
BB	1,645	2,346	5.07	5.07	0.000	1.85	1.85	0.000
BB-	1,719	2,406	5.26	6.13	0.866	2.64	1.77	-0.866
B+	1,935	2,733	8.04	7.17	-0.874	1.38	2.26	0.874
B	1,414	1,926	7.30	7.30	0.000	2.37	2.37	0.000
B-	579	692	7.54	9.71	2.166	2.09	-0.08	-2.166
CCC+	143	157	13.93	11.63	-2.299	-5.01	-2.72	2.299
CCC	65	72	12.82	12.82	0.000	-3.35	-3.35	0.000
CCC-	15	15	15.88	15.88	0.000	-5.79	-5.79	0.000
CC	18	19	15.83	15.83	0.000	-7.32	-7.32	0.000
D	9	9	12.84	12.84	0.000	-3.11	-3.11	0.000

Note: **Sample and unit.** Eligible observed bonds under the 3-year rollover counterfactual, grouped by the fine S&P rating prevailing at each hypothetical issuance. "Bonds" counts unique observed bonds represented in a row and "Rollovers" counts the hypothetical issues priced at that rating; these columns replace a generic N. Anchor ratings (AA, A, BBB, BB, B, and CCC) reproduce the broad bucket curve; plus/minus notches are interpolated toward the adjacent anchor. AAA and AA+ also use the pooled AAA/AA curve because separate AAA and AA benchmark curves are not estimated in the Merrill/TRACE extended baseline.

Table OA9: Credit rating determinants with short-term debt interaction

	(1) *OLS*	(2) *OLogit*	(3) *OLS*	(4) *FE*	(5) *FE*
Book Leverage	2.487***	2.198***	1.572***	1.865***	1.542***
	(0.228)	(0.088)	(0.230)	(0.219)	(0.224)
Short Debt Share	-0.958**	-1.587***	-0.913***	0.335	-0.154
	(0.372)	(0.165)	(0.318)	(0.235)	(0.239)
Leverage × Short Debt	0.357	1.423***	-1.196**	0.290	-0.340
	(0.524)	(0.294)	(0.564)	(0.365)	(0.439)
Debt / EBITDA	0.251***	0.274***	0.169***	0.117***	0.081***
	(0.013)	(0.006)	(0.012)	(0.009)	(0.009)
Interest Coverage	-0.005***	-0.004***	-0.005***	-0.003***	-0.003***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
Negative EBITDA (d)	2.781***	2.841***	1.499***	1.264***	0.619***
	(0.271)	(0.120)	(0.249)	(0.221)	(0.218)
Profitability	0.115	-0.285**	0.587*	-0.375	-0.096
	(0.395)	(0.129)	(0.327)	(0.335)	(0.306)
Profit Volatility	5.595***	5.192***	2.627***	1.622***	0.511
	(0.652)	(0.277)	(0.544)	(0.515)	(0.453)
Log Total Assets	-1.107***	-1.107***	-1.048***	-0.864***	-0.759***
	(0.047)	(0.013)	(0.043)	(0.071)	(0.069)
Cash / Assets	-2.248***	-2.721***	-2.682***	-0.388	-0.338
	(0.427)	(0.172)	(0.395)	(0.339)	(0.327)
Conv. Debt / Assets	4.597***	2.238***	4.442***	0.971	1.321
	(0.973)	(0.401)	(0.893)	(0.913)	(0.827)
Rent / Assets	11.324***	3.711***	5.187***	7.004***	5.086**
	(1.921)	(0.760)	(1.760)	(2.224)	(2.160)
PPE / Assets	-1.459***	-1.561***	-0.999***	-1.389***	-1.618***
	(0.319)	(0.103)	(0.302)	(0.366)	(0.361)
CAPEX / Assets	-1.593**	-0.739**	-1.008	-3.600***	-2.298***
	(0.774)	(0.320)	(0.649)	(0.809)	(0.656)
Beta (Dimson)			0.449***		0.140***
			(0.051)		(0.038)
Idiosyncratic Risk			1.766***		1.225***
			(0.088)		(0.074)
Year FE	Yes	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	No	No
Firm FE	No	No	No	Yes	Yes
Pooled R^2	0.698		0.759		
Pseudo R^2		0.230			
Within R^2				0.299	0.397
N	22,983	22,983	18,466	22,983	18,466

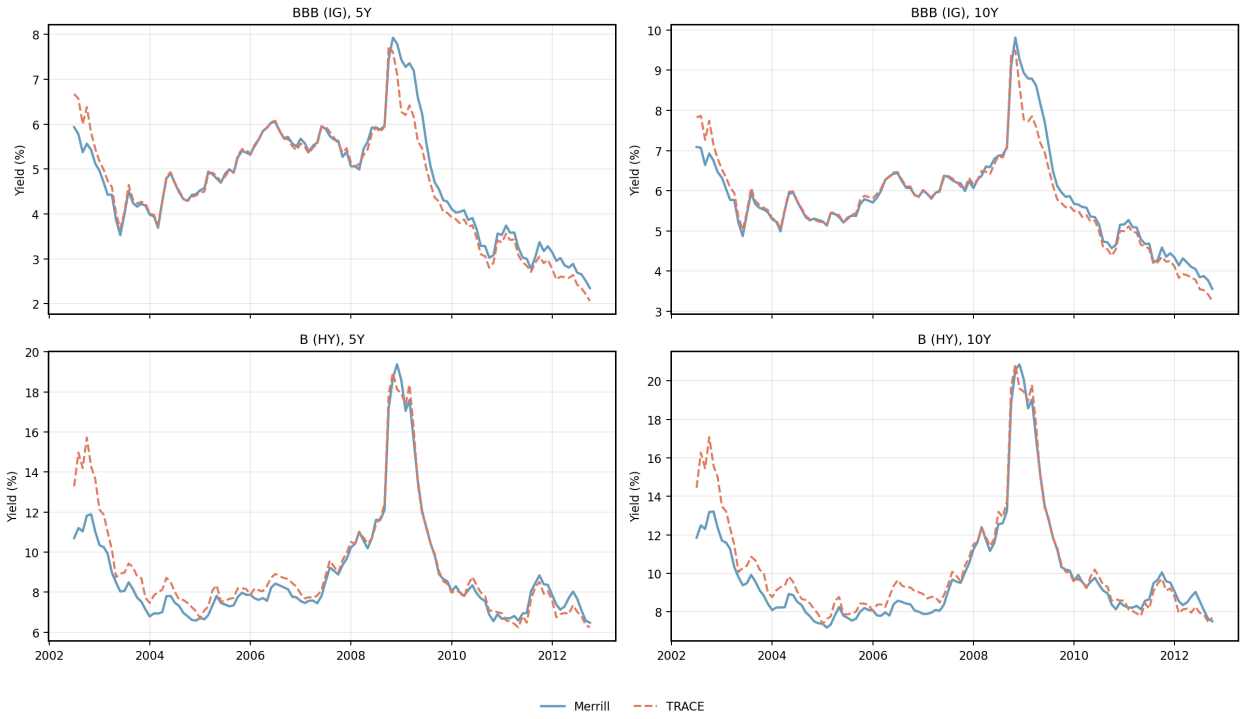
Note: **Sample and unit.** Compustat-matched rated firm-years over 1997–2025; each observation is one firm-year, and N is the number of firm-years used after each specification's control requirements. Dependent variable: numeric issuer credit rating (1 = AAA, ..., 21 = C). The estimating equation is $\text{Rating}_{i,t} = \alpha + \beta_1 \text{Lev}_{i,t} + \beta_2 \text{SDS}_{i,t} + \beta_3 (\text{Lev} \times \text{SDS})_{i,t} + \Gamma' \mathbf{X}_{i,t} + \delta' \mathbf{D}_{i,t} + \varepsilon_{i,t}$, where SDS ≡ DLC/(DLC+DLTT), **X** contains the controls from Baghai, Servaes, and Tamayo (2014), and **D** denotes the fixed effects or trend reported at the bottom of the table. OLS means ordinary least squares, OLogit means ordered logit, and FE means firm fixed effects. Columns (3) and (5) add Dimson beta and idiosyncratic risk; their CRSP inputs cover 1997–2025. Pooled, ordered-logit pseudo-, and within-firm R^2 are reported separately because they are not directly comparable. Standard errors are clustered by firm in parentheses. $^{***}\, p < 0.01$, $^{**}\, p < 0.05$, $^{*}\, p < 0.10$.

Table OA10: Predicted rating change at counterfactual rollover

		Δ Short Debt Share		Δ Rating (notches)		% Worse By	
CF Tenor	N	Mean	Med	Mean	Med	>0.5	>1.0
1Y	23,292	0.898	0.953	-0.266	-0.253	0.0	0.0
2Y	23,292	0.449	0.476	-0.133	-0.127	0.0	0.0
3Y	23,292	0.299	0.318	-0.089	-0.084	0.0	0.0
4Y	23,292	0.225	0.238	-0.067	-0.063	0.0	0.0
5Y	23,292	0.180	0.191	-0.053	-0.051	0.0	0.0

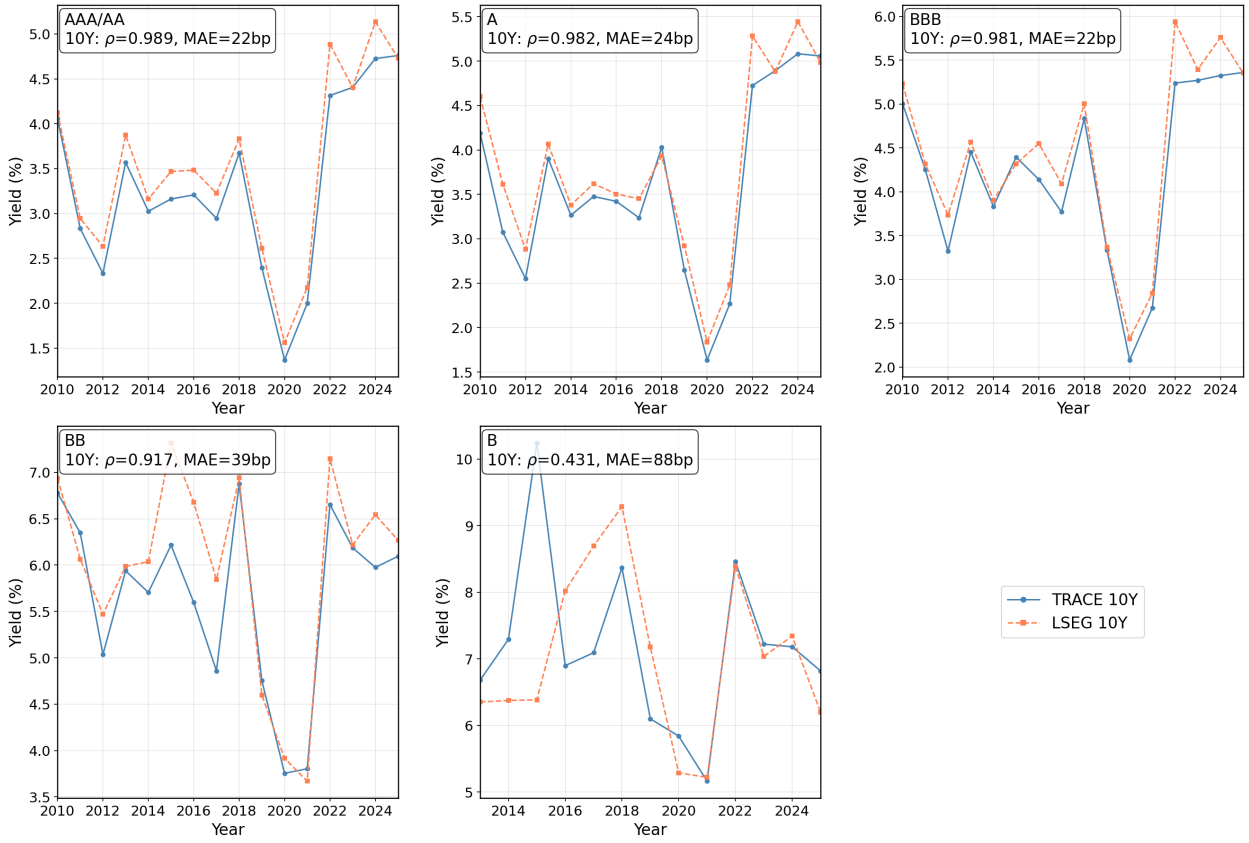
Note: **Sample and unit.** Compustat-matched rated firm-years; each observation is one firm-year and N is the number of firm-years simulated at each counterfactual tenor. For each firm-year the counterfactual short-term debt share (SDS) assumes that $1/k$ of long-term debt (DLTT) becomes due within one year under the k-year rollover strategy: $\text{SDS}^{CF} = \text{SDS} + (1 - \text{SDS})/k$. Rating changes predicted using coefficients on *Short Debt Share* and *Leverage* $\times$ *Short Debt* from the preferred firm-FE model (column 5 of Table OA9). Positive values mean worse ratings; negative values mean better ratings. Sample: 23,292 Compustat-matched firm-years.

Figure OA1: Merrill and TRACE yields in the overlap window: IG (BBB) and HY (B).



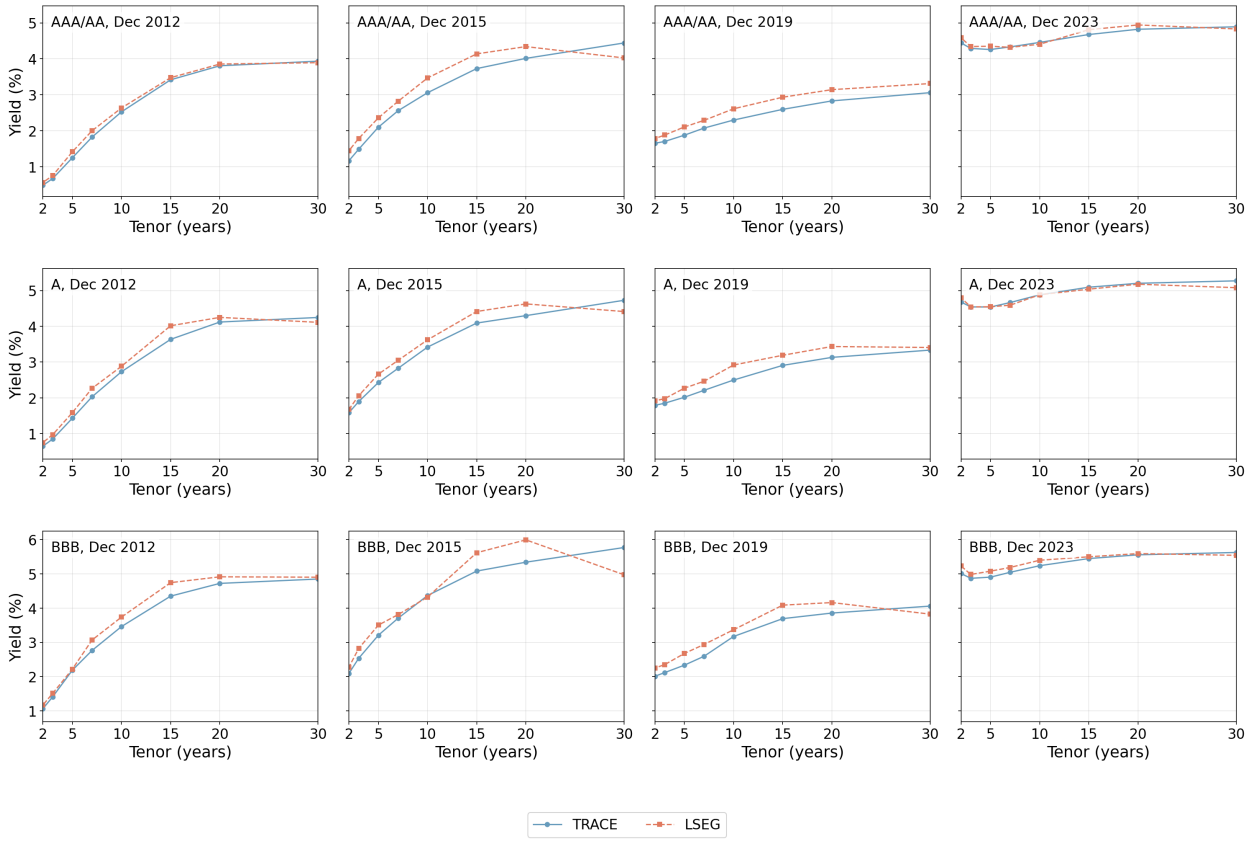
Note: **Sample.** February 2003–October 2012 overlap between trade-date Merrill and TRACE counterfactual yield curves at the 5- and 10-year tenors, for BBB-rated (IG, top row) and B-rated (HY, bottom row) bonds. **Unit.** Each plotted observation is a monthly rating–tenor fitted yield in percent; the figure does not report a separate N. **Construction.** Both series use the same broad rating bucket, GSW risk-free curve, isotonic PCHIP spread interpolation, and standard-tenor counterfactual-yield reconstruction. The series generally move together. The BBB gap is most visible around the 2008–2009 financial crisis, while the B gap is largest in the early 2000s.

Figure OA2: TRACE vs. LSEG: quantitative validation at the 10-year tenor.



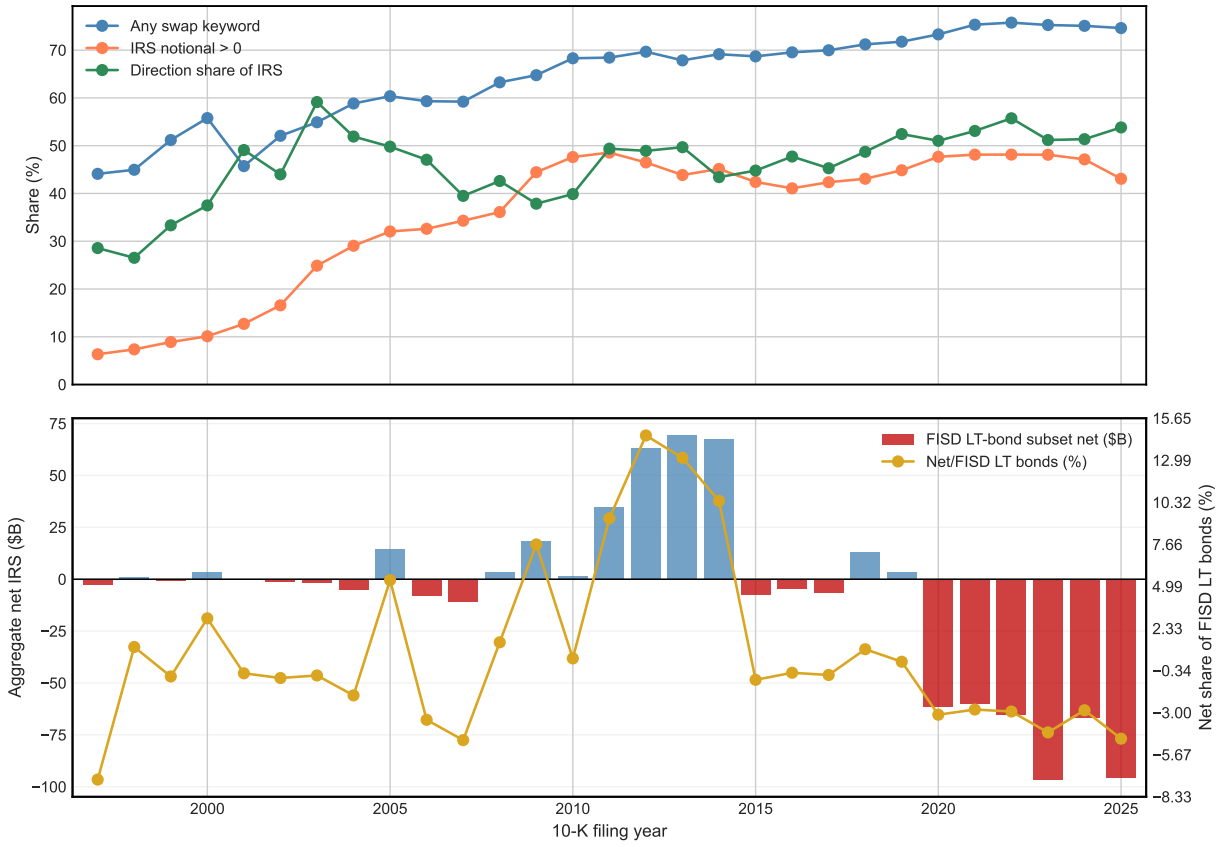
Note: **Sample.** 2010–2025 overlap window between TRACE and LSEG corporate-bond pricing at the 10-year tenor, by S&P long-term rating bucket (AAA, AA, A, BBB, BB, B). **Unit and count.** Each observation is one rating–year-end pair. For each rating, T is the number of year-ends with both fitted yields; the figure reports T rather than a separate N. **Construction.** For each rating g and year-end t, let $y_{g,t}^{\text{TRACE}}$ be the PCHIP-fitted 10-year corporate yield from TRACE VWAP transaction prices and $y_{g,t}^{\text{LSEG}}$ the corresponding LSEG evaluated dealer-quoted yield. The reported statistics are the Pearson correlation $\rho_g = \text{corr}\left(y_{g,t}^{\text{TRACE}}, y_{g,t}^{\text{LSEG}}\right)$ and the mean absolute error (MAE), $\text{MAE}_g = (1/T)\sum_t \left|y_{g,t}^{\text{TRACE}} - y_{g,t}^{\text{LSEG}}\right|$ in basis points. IG ratings show high correlation and low MAE between the two data sources, supporting the TRACE-fitted yield curve at the IG tenors used in the wedge; speculative-grade buckets differ more because TRACE coverage is thinner and LSEG uses dealer evaluations rather than transactions.

Figure OA3: TRACE vs. LSEG: term-structure snapshots, AAA/AA, A, and BBB.



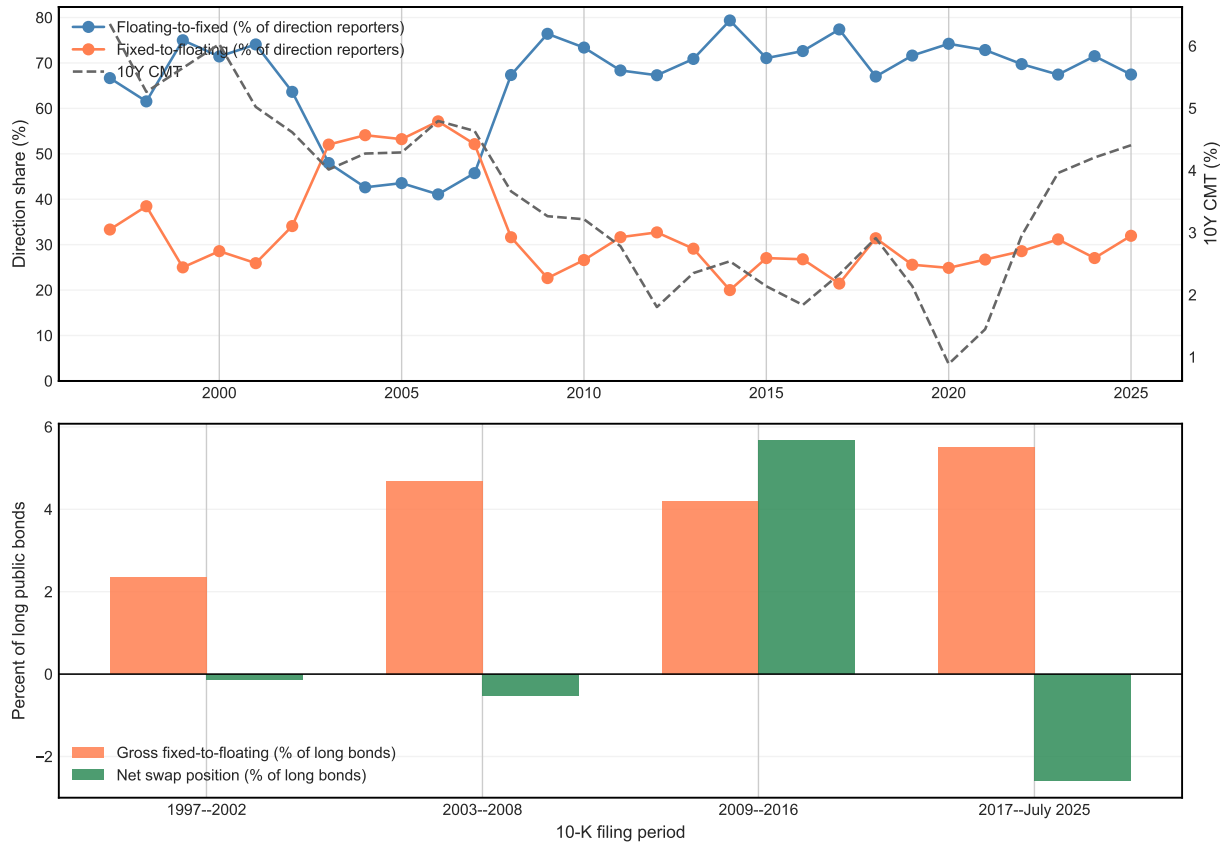
Note: **Sample.** Bond-year panel universe restricted to AAA/AA-rated, A-rated, and BBB-rated bonds at year-end 2012, 2015, 2019, and 2023. **Unit.** Each plotted observation is one rating–year–tenor fitted yield in percent; the figure does not report a separate N. **Construction.** For each rating $g \in \{\text{AAA/AA}, \text{A}, \text{BBB}\}$ and year-end t, the panels plot the cross-sectional term structure $\{(\tau, y_{g,t,\tau})\}_{\tau=1}^{30}$, where $y_{g,t,\tau}$ is the rating-by-tenor corporate yield at maturity τ years. TRACE uses PCHIP-interpolated VWAP transaction prices fitted under an isotonic prior on tenor; LSEG uses evaluated dealer-quoted yields. Both data sources produce similar IG term-structure shapes up to a small parallel level shift (within the MAE bands of Figure OA2); the mean-reversion shape and the post-2022 inversion of the BBB curve appear in both data sources.

Figure OA4: Nonfinancial swap disclosure and net IRS position over time



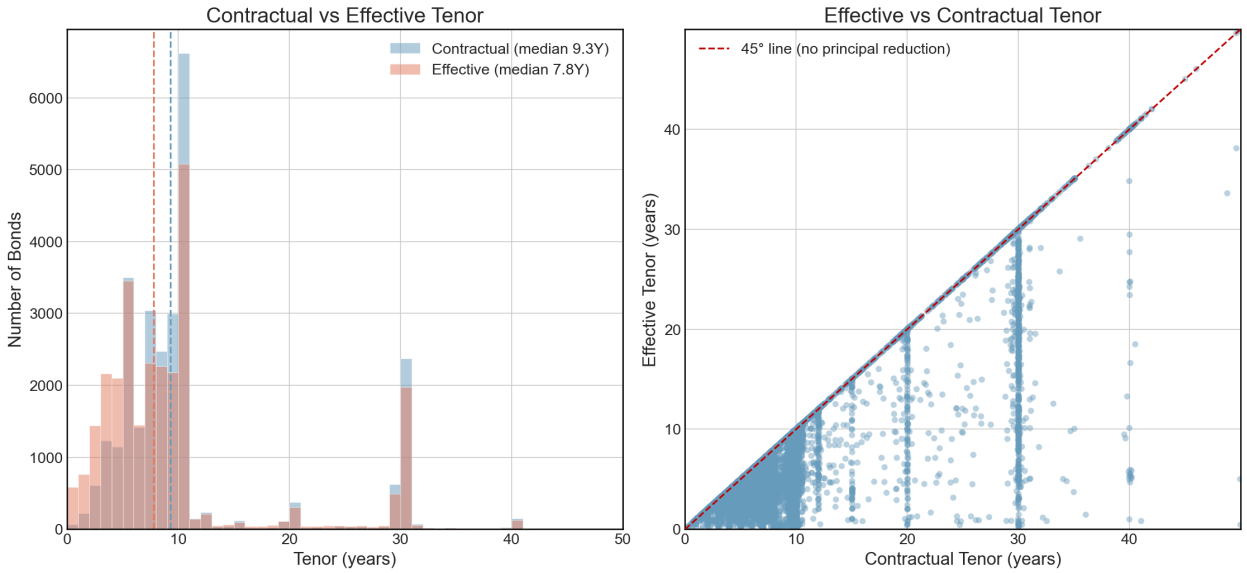
Note: **Sample and unit.** Matched nonfinancial-company filing-years in the EDGAR swap-disclosure sample; the top panel's observation unit is a company filing-year and the bottom panel aggregates those observations to calendar years. The figure reports shares and aggregate notionals, not a separate N. The top panel reports annual shares of company filing-years with any swap keyword, positive IRS notional, and explicit direction among IRS users. The bottom panel reports aggregate net IRS notional and the same net divided by aggregate FISD public bonds with more than five years remaining. Positive net means floating-to-fixed conversion (pay fixed − pay floating > 0). Direction is never imputed.

Figure OA5: Swap direction and the long-maturity debt benchmark



Note: **Sample and unit.** Matched nonfinancial-company filing-years with swap disclosures; the top panel's observation unit is an explicit-direction company filing-year and the bottom panel aggregates to the reported periods. The figure reports shares and aggregate notionals, not a separate N. The top panel shows the annual shares of explicit direction reporters that are net floating-to-fixed or net fixed-to-floating; the 10-year Constant Maturity Treasury (CMT) yield is a descriptive interest-rate benchmark. The bottom panel reports gross fixed-to-floating notional and net swap notional as percentages of aggregate FISD public-bond principal with more than five years remaining. Company-years without an explicit pay-floating disclosure contribute zero to the disclosed gross measure.

Figure OA6: Contractual vs. effective tenor.



Note: **Sample.** 21,110 fixed-rate U.S.-dollar, non-convertible, non-putable corporate bonds in the effective-tenor sample (90 bonds with $T_i > 30$Y excluded for visualization only). **Unit and count.** Each plotted observation is one bond, so the stated bond count is the figure's N. **Construction.** Effective tenor is the average time to repayment, weighted by the amount of principal repaid in each FISD record. **Left panel.** Overlaid histograms of contractual tenor T_i (steelblue) and effective tenor $\widetilde{T}_i$ (coral). **Right panel.** Scatter of $\widetilde{T}_i$ against T_i; the dashed 45° line marks no principal reduction. Reported principal reductions move observations below the diagonal; any remaining principal stays exposed through contractual maturity.