Research Department

Technical Appendix: Why Mortgage Rates Exceed Treasury Yields

Paul S. Willen

This appendix provides additional detail on the methodology used in the accompanying Current Policy Perspectives brief. It is organized as a series of frequently asked questions.

1. What is the coupon spread, and how does it differ from the primary–secondary spread?

The coupon spread measures how much more borrowers pay relative to what they would pay if they could borrow like the government—specifically, if they could issue a government bond with the same predictable (non-refinancing) prepayment pattern as a typical mortgage.

Computing the coupon spread

Step 1: Define the primary mortgage coupon. The primary coupon is the interest rate borrowers pay, minus the guarantee fee (g-fee) charged by Fannie Mae or Freddie Mac:

$$c = r_{\text{mortgage}} - \text{gfee}$$

The g-fee compensates the GSEs for credit risk. We subtract it because we want to isolate the spread attributable to prepayment risk, not credit risk.

Step 2: Construct the benchmark coupon. We ask: What coupon would make an investor indifferent between holding a government bond and holding a mortgage that prepays only for "life events" (death, divorce, job relocation)—not for refinancing? This prepayment pattern follows the Public Securities Association (PSA) standard, which assumes prepayment rates rise gradually to 6 percent CPR over 30 months and remain constant thereafter.

Using the Treasury yield curve, we calculate the fair coupon c^* that prices such a bond at the current mortgage-backed securities (MBS) market price:

$$c^* = \text{coupon such that } P_{\text{PSA}}(c^*, \text{Treasury curve}) = P_{\text{MBS}}$$

Step 3: Compute the spread. The coupon spread is:

$$\text{Coupon Spread} = c - c^*$$

This spread represents the additional coupon borrowers pay to compensate investors for bearing refinancing risk.

How does this differ from the primary–secondary spread?

The standard primary–secondary spread compares the primary mortgage rate to a secondary market yield. This approach faces a fundamental problem: Computing the yield on an MBS requires a prepayment model. Different prepayment assumptions produce different yields, and these models are proprietary, complex, and often opaque.

One common shortcut assumes that mortgages sell in the secondary market at par (face value). Under this assumption, the MBS yield equals the MBS coupon rate, and the spread is simply the difference between the primary rate and the coupon. But MBS rarely trade at par—they frequently

trade at premiums or discounts—so this assumption introduces systematic measurement error that varies with market conditions.

My coupon spread avoids both problems. Rather than converting prices to yields (which requires a prepayment model), I work directly with coupons. And rather than assuming par pricing, I use actual market prices. The benchmark c^* answers a transparent question: What coupon would a government-like bond need to offer to match the current MBS price, given only predictable (non-refinancing) prepayment? The calculation uses the publicly available Treasury yield curve and a standard prepayment assumption, making it fully replicable.

Figure 1 compares the coupon spread to two conventional measures of the MBS–Treasury spread. The first is the current coupon minus the 10-year Treasury yield, where the current coupon is the MBS coupon rate that would price at par—this measure assumes par pricing. The second is the interpolated MBS yield (at the production coupon) minus the 10-year Treasury yield—this avoids the par assumption but still requires computing a yield, which depends on a prepayment model. The coupon spread is consistently larger than either alternative, because it captures the full cost of the prepayment option rather than just the gap between MBS and Treasury rates.

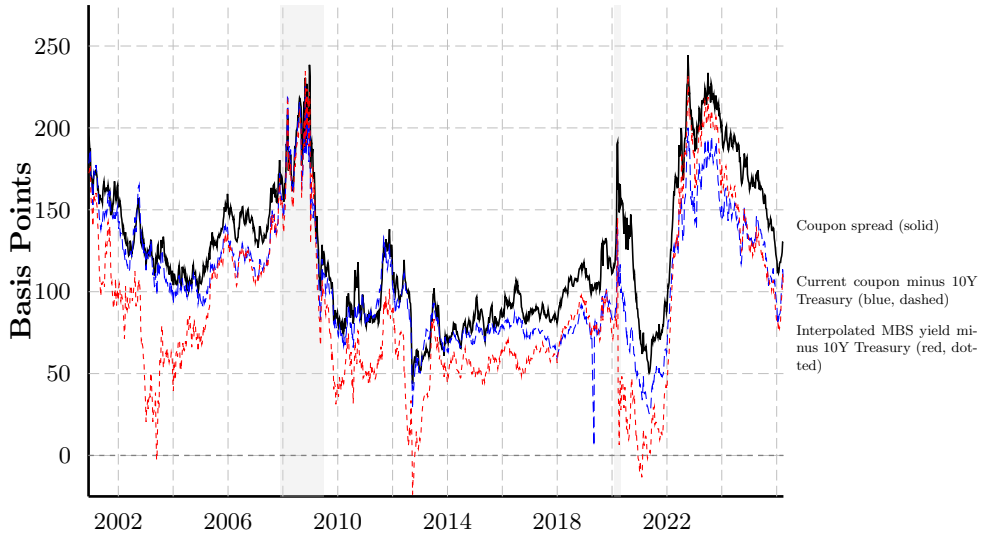

Figure 1: Coupon spread (black) vs. current coupon minus 10-year Treasury (blue, dashed) and interpolated MBS yield minus 10-year Treasury (red, dotted), 2000–present.

2. How do we control for the prepayment option?

The regression framework

We decompose the coupon spread using a regression with three explanatory variables that capture the value of the prepayment option:

$$\text{Coupon Spread}_t = \alpha + \beta_1 \cdot \text{Slope}_t + \beta_2 \cdot \text{Vol}_t + \beta_3 \cdot \text{OPUC}_t + \varepsilon_t$$

- **Slope**: The difference between 10-year and two-year Treasury rates (percentage points). A steeper yield curve signals expected rate increases, reducing refinancing option value.

- **Vol**: The seven-year into three-year ATM payer swaption implied volatility (basis points). Higher volatility increases the probability that rates could fall enough to make refinancing attractive.

- **OPUC**: Originator profits and unmeasured costs, measured as deviation from a three-year (756-day) moving average.[1] When the fees and expenses of getting a mortgage are unusually high, markets anticipate they will normalize, increasing future refinancing probability.

Why a rolling regression?

The relationship between these factors and the coupon spread is not necessarily constant over time. Market structure evolves, the Fed's policy framework changes, and the composition of the mortgage market shifts. A fixed-coefficient model estimated over the full sample would impose the assumption that, say, a 1 percentage point change in the yield curve slope had the same effect on spreads in 2006 as it did in 2025.

We instead use a 260-week (five-year) rolling window, re-estimating coefficients each week using only data from the preceding five years. This approach has two advantages. First, it allows the relationships to evolve over time—if the sensitivity of spreads to volatility has changed, the rolling regression captures this. Second, it ensures that predictions are genuinely out-of-sample: The predicted spread for any given week is based entirely on coefficients estimated from earlier data.

The model explains approximately 80 to 85 percent of the variation in coupon spreads. Residuals typically fall within ±20 basis points during normal periods, with larger deviations during crises.

Regression results

Table 1 reports regression coefficients from a fixed-window specification estimated over the 2018–2022 period, which is then used for out-of-sample prediction. Panel A shows the estimated coefficients and their interpretation. Panel B compares model performance in-sample versus out-of-sample.

Table 1: Regression coefficients and model performance.

Panel A: Regression Coefficients

Variable	Coefficient	Interpretation
Intercept	64.9	Base spread level (bps)
YC Slope (10Y–2Y)	−70.4	Steeper curve → lower spread
Swaption Vol (7Y×3Y)	+1.03	Higher vol → higher spread
Moneyness	+18.4	More in-the-money → higher spread
R^2	0.891	
Residual SD	13.5 bps	

Panel B: Model Performance

Period	Mean Residual	SD Residual
In-sample (2018–2022)	0.0 bps	13.5 bps
Out-of-sample (2023–2026)	+0.8 bps	9.7 bps

The coefficients have intuitive signs: A steeper yield curve reduces the spread (borrowers are

less likely to refinance when rates are expected to rise), higher volatility increases the spread (more uncertainty means refinancing is more likely), and higher intermediation costs above the historical average increase the spread. The model has an R^2 of 0.89 in-sample and performs even better out-of-sample (residual standard deviation of 9.7 basis points versus 13.5 basis points in-sample), suggesting the relationships are stable and not overfit.

How much does each factor contribute?

Because the three explanatory variables are correlated, a simple comparison of partial R^2 values can be misleading—the order in which variables enter the model affects how much explanatory power is attributed to each. Shapley values provide a principled solution: They average each variable's marginal contribution across all possible orderings, ensuring that the three contributions sum to the total R^2.

Figure 2 shows how these contributions have evolved over time. The yield curve slope (blue) is consistently the largest contributor, typically accounting for 30 to 50 percent of the model's explanatory power. Swaption volatility (red) becomes more important during periods of monetary policy uncertainty, such as the 2008–2010 and 2022–2023 periods. OPUC deviation (green) contributes modestly in most periods but becomes more important when intermediation costs are elevated, as during the COVID-19–era refinancing boom.

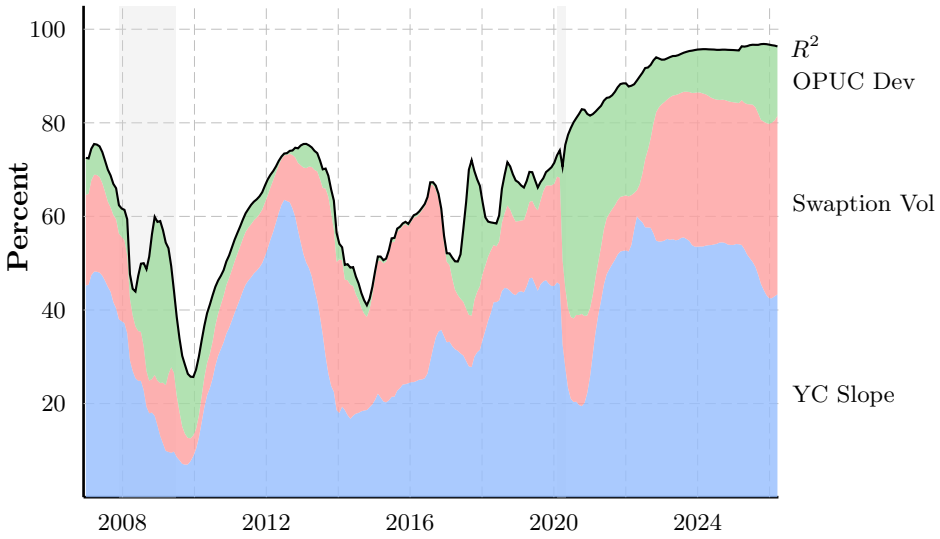

Figure 2: Shapley value decomposition of model R^2 over time. Each colored area shows one variable's contribution; the black line shows total R^2. Computed using 260-week rolling windows.

3. How does our residual differ from option-adjusted spread (OAS)?

Option-adjusted spread (OAS) is the standard measure of the mortgage risk premium in fixed-income markets.[2] OAS starts with the market price of a mortgage-backed security and asks: After accounting for the value of the prepayment option, what residual spread are investors earning? If OAS is high, investors are being compensated for something beyond prepayment risk—credit concerns, liquidity, or market dislocations. If OAS is low, the prepayment option accounts for most of the spread.

Our regression residual answers a similar question. After the yield curve slope, interest rate volatility, and intermediation costs explain as much of the coupon spread as they can, what is left over? Both measures are trying to isolate the same thing: the portion of the spread that is not attributable to the prepayment option. They differ in how they model the option.

Two approaches to the same question

OAS uses a *structural model*. A stochastic interest rate model generates thousands of possible rate paths. At each path, a prepayment model predicts how borrowers will respond—who refinances, who moves, who stays. The model prices the mortgage's cash flows along each path, and OAS is the constant spread added to the discount rates that makes the model price equal the observed market price. Whatever the structural model cannot explain shows up in OAS.

Our approach uses a *statistical model*. We regress the coupon spread on three observable factors that drive the value of the prepayment option, and whatever the regression cannot explain shows up in the residual.

The logic is the same: Model the prepayment option and define the residual as the gap between model and data. The structural approach models prepayment behavior directly; the statistical approach captures the same forces indirectly through the macroeconomic variables that drive them.

Why do results differ across sources?

The structural approach requires choosing an interest rate model and a prepayment model, and different dealers make different choices. Figure 3 compares the regression residual to OAS estimates from two major dealers: J.P. Morgan and Citi. All three series are demeaned to facilitate comparison of dynamics rather than levels.

The three series share broad contours—all spike during the financial crisis and COVID-19 pandemic—but diverge meaningfully in timing and magnitude. The two dealer OAS estimates correlate with each other at only about 0.4, and each correlates with the residual at a similar level. The disagreement reflects different modeling assumptions: different prepayment functions, different interest rate processes, different treatments of borrower heterogeneity. These are not minor calibration differences—they produce meaningfully different answers about what is left over after accounting for the option.

Tradeoffs

The structural approach is more detailed—it models borrower behavior explicitly and can price individual securities. But that detail comes at a cost: The models are proprietary, opaque, and can break down during crises, precisely when accurate measurement matters most.

My statistical approach is coarser—three variables in a linear regression cannot capture every nuance of prepayment behavior. But it has the advantage of transparency. The inputs are publicly available, the specification is simple, and the results are fully replicable. And because the regression does not require assumptions about how borrowers respond to rate changes, it is less vulnerable to model misspecification during unusual episodes.

Explanatory variables

Figure 4 shows the three explanatory variables used in my regression: the yield curve slope (10-year minus 2-year Treasury rates), swaption implied volatility, and OPUC deviation from its three-year moving average.

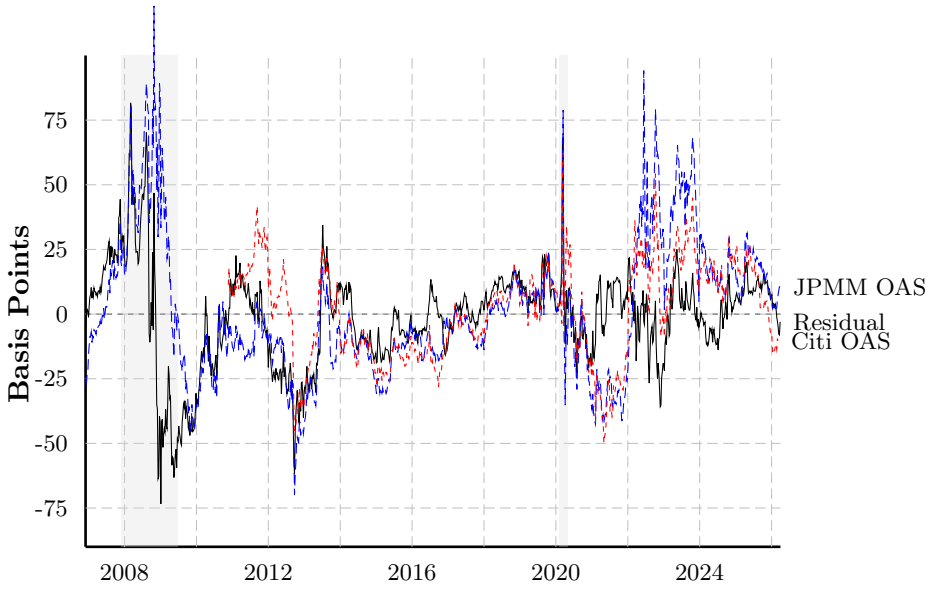

Figure 3: Regression residual (black) vs. demeaned OAS from J.P. Morgan (blue, dashed) and Citi (red, dotted), 2004–present. All series shown as deviations from their respective means.

Data Sources

- **Primary mortgage rate**: Optimal Blue Mortgage Market Indices (OBMMI), a daily index based on actual rate locks
- **MBS prices and yields**: J.P. Morgan Markets API (TBA prices for 30-year FNMA securities)
- **Treasury yields**: Federal Reserve H.15 release and Gürkaynak-Sack-Wright zero-coupon yield curve
- **Swaption volatility**: Bloomberg (7Y×3Y ATM payer swaption implied volatility)
- **Guarantee fees**: Fannie Mae and Freddie Mac published g-fee schedules
- **OPUC**: Author's calculations following Fuster et al. (2013)

4. What about policy announcements?

Policy announcements provide a different test of the model. Unlike normal periods when fundamentals change gradually, announcements are discrete events. And unlike ordinary macroeconomic episodes, major policy announcements typically do not coincide with large changes in the yield curve or volatility—in some cases, those fundamentals actually move in the opposite direction.

Table 2 presents event study evidence around major Fed announcements. For each event, I report both one-day and one-month changes in the coupon spread and in the regression residual.

QE1 and the pandemic had large negative residuals on the announcement day (-49 and -33 basis points), indicating spread compression beyond what the yield curve and volatility would predict. These were crisis interventions when market dislocations were severe. The pandemic residual had reverted entirely within one month, while the QE1 residual expanded to -78 basis points, suggesting the dislocation took longer to resolve.

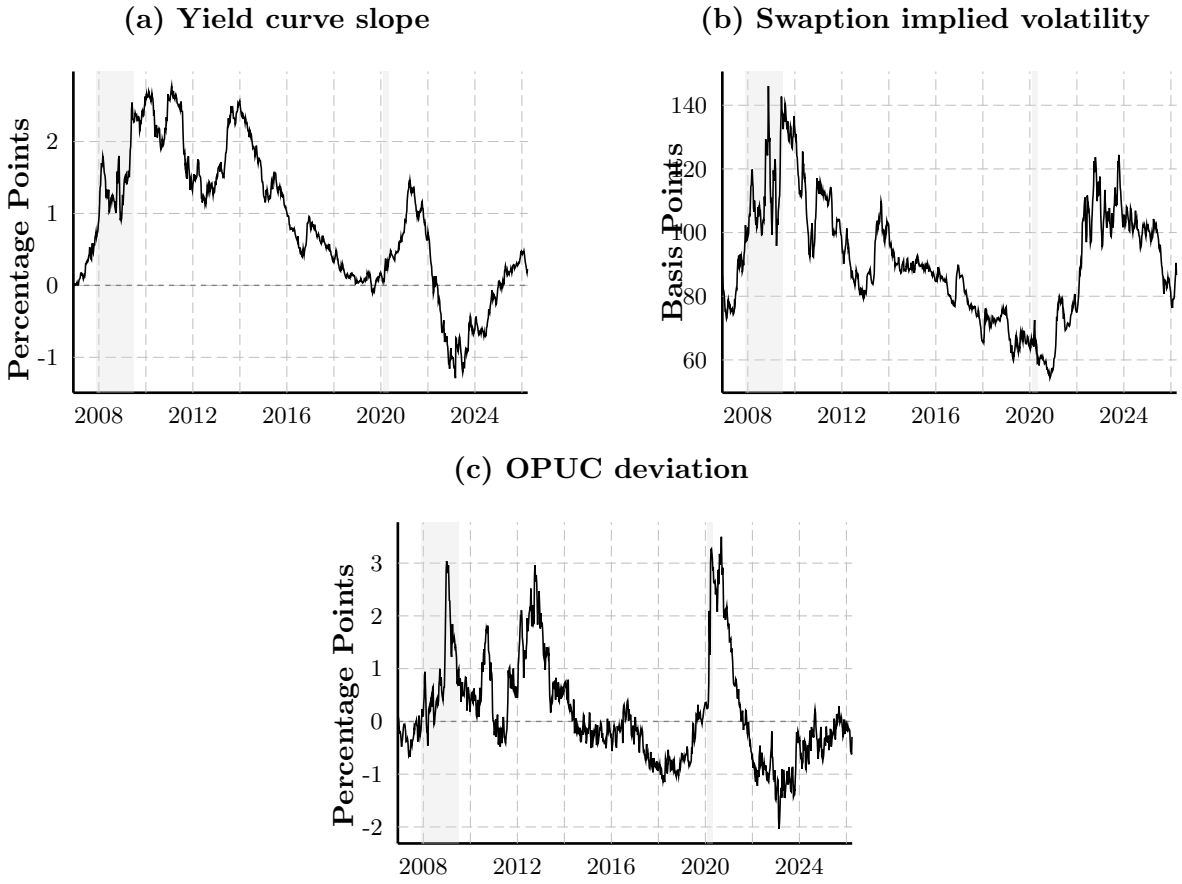

Figure 4: Explanatory variables in the coupon spread regression, 2004–present. (a) Yield curve slope: 10-year minus 2-year Treasury rates (percentage points). (b) Swaption implied volatility: seven-year into three-year ATM payer (basis points). (c) OPUC deviation from three-year moving average (percentage points). Shaded areas indicate the financial crisis and COVID-19 pandemic.

Table 2: Changes in spread and residual around major Fed announcements.

		1-Day		1-Month	
Event	**Date**	Spread	Residual	Spread	Residual
QE1 (Fed MBS purchases)	Nov 2008	−40	−49	−16	−78
QE2 (Treasury purchases)	Nov 2010	−11	−9	−5	+10
Op. Twist (maturity extension)	Sep 2011	−10	−17	+2	−4
QE3 (open-ended purchases)	Sep 2012	−19	−20	−40	−15
Taper (slowing purchases)	Dec 2013	+1	+3	+8	+6
COVID (emergency purchases)	Mar 2020	−18	−33	+28	+0
FHFA (GSE purchases)	Jan 2026	−10	−12	−16	−5

Non-crisis episodes (QE2, Twist, QE3, Taper) show smaller residuals. The Federal Housing Finance Agency (FHFA) announcement in January 2026 had a modest residual (−12 basis points), consistent with incremental effects from GSE purchases in normal market conditions.

5. GSE Purchases in 2025

In late 2025, the FHFA announced that Fannie Mae and Freddie Mac would expand their retained portfolios by as much as $200 billion. From September 30, 2025, to January 8, 2026, coupon spreads fell 21 basis points. Was this the effect of GSE purchases? Table 3 decomposes the change.

Table 3: Decomposition of spread change, September 30, 2025, to January 8, 2026.

	Sept 30	**Jan 8**	**Change**	**Contribution**
Yield curve slope	0.27	0.48	+0.21	−8 bps
Volatility	85	78	−6	−9 bps
Intermediation	3.4	3.2	−0.2	−1 bps
Spread	104	83		−21 bps
Predicted	97	79		−18 bps
Residual	7	4		−4 bps

The yield curve steepened by 21 basis points (contributing −8 basis points to spreads), volatility fell 7 basis points (contributing −9 basis points), and intermediation costs declined modestly (contributing −1 basis point). Together, these factors explain 18 of the 21 basis point decline. The residual fell only 4 basis points. The last row of Table 2 shows that on the day of the FHFA announcement itself, the one-day residual was −12 basis points, a modest effect consistent with GSE purchases in normal market conditions. Most of the spread compression reflected improving macro conditions, not the direct effect of GSE purchases.

6. Implications

The gap between mortgage rates and Treasury yields is large, but it is not primarily a reflection of market inefficiency or lender profits. Most of the gap's variation is explained by the same factors that determine the value of the prepayment option. For borrowers, a wide spread is not necessarily a bad time to buy—it often signals conditions favorable to refinancing. For policymakers, monetary policy affects spreads mainly through fundamentals: Clear forward guidance reduces volatility, and a steeper yield curve reduces refinancing incentives. Direct interventions such as MBS purchases can compress spreads during crises but have diminishing effects in normal conditions.

References

[1]See Fuster, Goodman, Lucca, Madar, Molloy, and Willen, "The Rising Gap between Primary and Secondary Mortgage Rates," *Economic Policy Review*, Federal Reserve Bank of New York, December 2013.

[2]For foundational treatments of OAS and prepayment modeling, see Richard and Roll, "Prepayments on Fixed-Rate Mortgage-Backed Securities," *Journal of Portfolio Management*, Spring 1989. For recent analysis of MBS risk premia and OAS, see Boyarchenko, Fuster, and Lucca, "Understanding Mortgage Spreads," *Review of Financial Studies*, 2019; and Diep, Eisfeldt, and Richardson, "The Cross Section of MBS Returns," *Journal of Finance*, 2021.