Module 3 — Slide Notes

Decomposing Sovereign Bond Yields

Prepared for M2 Finance Technology Data (M2 FTD), Sorbonne — "Financial Crises and Emerging Risks", April 14, 2026.

Antoine Camous — acamous.github.io — Banque de France

The big question

When we observe a sovereign bond yield, what is driving that number?

The lecture builds the answer in three movements: the theory of why yields must be decomposed (Part A) → a step-by-step construction of a tractable model that actually does it (Part B) → what the decomposition reveals in the data, for EMEs and the euro area (Part C). The red thread is the equation $s_t(\tau) = EC_t(\tau) + RP_t(\tau)$: introduced conceptually in Part A, derived from first principles in Part B, and read empirically in Part C.

The opening move — showing the CS/ECB figure before any theory — is deliberate. The answer appears before the question, which creates the pull that sustains the subsequent derivation.

Opening: Motivation

Slide 2 (Motivation — CS/ECB figure)

The ECB figure — showing sovereign spreads in the left panel and their decomposition into components in the right — opens the lecture before any theory. The right panel poses the question directly: how do we know which part of the Italian spread is default risk, which is redenomination risk, and which is expected policy divergence? The answer requires a model, and the rest of the lecture builds it.

Transition: the roadmap for getting there.

Slide 3 (Roadmap)

Three parts. Part B is the intellectual core; Part A provides the vocabulary that makes it legible. The derivation in Part B is not a detour — it is the lecture.

Part A — Bond Pricing Basics & the Yield Decomposition Problem

Part A builds the working vocabulary — zero-coupon bonds, yield curve, short rate, expectations hypothesis, term premium, risk premium in spreads — and establishes why a model is needed to decompose yields. The CS (2023) result provides the destination; Part B provides the route.

Slide 4 (Part A divider)

Slide 5 (Zero-coupon bonds and yields)

The continuously-compounded convention $P_t(\tau) = e^{-y_t(\tau)\cdot\tau}$ is not universal, and the equivalence $y_t(\tau) = -\frac{1}{\tau}\ln P_t(\tau)$ is the one that makes all subsequent formulas clean. A zero-coupon bond is the simplest possible debt instrument — one promise, one payment, no coupons — and all coupon bonds can be decomposed into zero-coupon bonds (strip pricing). The yield curve $\tau \mapsto y_t(\tau)$ is the full term structure: a snapshot at date t of what the market charges to lend for each horizon. The short rate r_t is the limit as $\tau \to 0$, in practice proxied by the overnight rate (Fed Funds, €STR). *Transition*: what theory governs the relationship between short rates and long yields?

Slide 6 (The Expectations Hypothesis)

The EH is the simplest possible theory: long yields equal the average of expected future short rates. The no-arbitrage argument rests on two strategies for investing over τ periods, equalized by risk-neutral investors. The numerical example grounds the formula: $y_t(3) = (3\%+3.5\%+4\%)/3 = 3.5\%$. The upward slope reflects expected rate hikes, not a premium. Under EH, the yield curve is a pure expectations machine — it contains no information about risk, only about future monetary policy. *Transition*: does this hold in the data?

Slide 7 (Why the Expectations Hypothesis fails)

Three empirical facts, escalating in force. (1) Long bonds earn excess returns on average — if EH held, all maturities would earn the same return. (2) The slope of the yield curve predicts excess returns — the slope should be uninformative under EH but empirically it is the single best predictor of bond returns at horizons of 1–5 years (Fama-Bliss, Cochrane-Piazzesi). (3) The term premium is time-varying and co-moves with macro conditions — it rises in recessions (flight to quality) and falls when bond supply shrinks (QE). The yield curve therefore contains two distinct signals — expectations about monetary policy, and the price of bearing duration risk. Any serious analysis must separate them.

Technical note: the term premium $TP_t(\tau)$ can be negative — during QE episodes, central bank purchases compressed it well below zero for German Bunds and US Treasuries. This can seem paradoxical: investors accept a yield below expected future short rates, exchanging a certain loss for safety. That is precisely what happened 2012–2021.

Transition: for sovereign spreads, the decomposition gets even richer.

Slide 8 (For sovereign spreads, the decomposition gets richer)

Moving from yield levels to spreads over a benchmark, the "risk premium" component fractures into distinct pieces: credit/default risk, FX risk (for local-currency bonds held by foreign investors), redenomination risk (specific to the eurozone — the possibility of exit and forced currency conversion), liquidity risk, and segmentation/convenience effects. Each requires separate identification. The same spread — say 200 bps on a 10-year Brazilian bond — can mean very different things depending on which component dominates. A spread driven by FX volatility calls for a different policy response than one driven by default risk. *Transition*: the most complete decomposition available for the euro area is Corradin & Schwaab (2023).

Slide 9 (Corradin & Schwaab (2023): the frontier for the euro area)

CS decompose EZ sovereign yields into five country-specific components plus a common (EZ-wide) component, using a clever identification strategy based on the difference between CDS protocols. The identification logic deserves careful attention.

Background: what is a CDS? A credit default swap is an insurance contract on sovereign debt. The protection buyer pays a periodic premium (the CDS spread, in basis points per year) to the protection seller. In return, if a credit event occurs, the seller compensates the buyer for the loss. The CDS spread is therefore a market price for sovereign default risk — it moves with perceived creditworthiness in real time, unlike bond yields which reflect both credit risk and the term structure of interest rates.

The ISDA 2003 protocol is the legal standard that governed most sovereign CDS contracts until 2014. It defines which events trigger a payout. Under ISDA 2003, the covered credit events are: *failure to pay* (the government misses a coupon or principal payment) and *restructuring* (debt is rescheduled, principal written down, or maturity extended under duress). Crucially, redenomination — a government converting euro-denominated bonds into a new national currency at a depreciated exchange rate — was **not** a covered credit event under ISDA 2003. The drafters in 2003 did not anticipate eurozone break-up as a realistic scenario.

Why does ISDA 2003 not cover redenomination? When Greece was contemplating exit in 2012, a legal question arose: if Greece reintroduced the drachma and redenominated all its bonds from euros into drachmas overnight, would that trigger CDS? Under ISDA 2003, the answer was arguably no — because redenomination into a new national currency is not the same as a failure to pay or a restructuring in the traditional sense. This created a dangerous gap: investors holding Greek bonds could be severely harmed by redenomination without their CDS insurance paying out.

The ISDA 2014 protocol was introduced precisely to close this gap. It adds *governmental intervention* and, for European reference entities, *redenomination* as explicit credit events. A CDS written under ISDA 2014 pays out if the sovereign redenominates its debt into a currency other than the euro. This made ISDA 2014 CDS a more comprehensive — but also more expensive — form of insurance for eurozone sovereigns.

The ISDA basis is defined as:

$$\text{ISDA basis}_t = \text{CDS spread (2014 protocol)}_t - \text{CDS spread (2003 protocol)}_t \geq 0$$

Both contracts reference the same sovereign, same maturity. The only difference is what triggers a payout. Since the 2014 contract covers more events (in particular, euro exit), it commands a higher spread. The difference — the basis — is the market's price for the *additional* risk that the 2003 contract does not cover: the redenomination risk premium. This is a clean, market-based identification: same issuer, same maturity, same counterparty risk, different legal definition of credit event. The basis isolates one specific risk in an almost pure way.

With this in mind, the CS identification strategy becomes clear: - **Default RP**: ISDA 2003 CDS spread. Covers only outright default (failure to pay, restructuring). Since redenomination is excluded, this prices pure credit/default risk — the probability that the government cannot or will not repay in full. - **Redenomination RP**: The ISDA basis (2014 minus 2003 spread). Isolates the probability of euro exit. This widened dramatically for Italy in 2018 (Lega/M5S coalition forming a eurosceptic government), in 2011–12 (peak of the EZ crisis), and around the French 2017 election. - **Liquidity RP**: KfW bonds are German agency bonds — explicitly guaranteed by the Federal Republic of Germany, carrying the same credit risk as Bunds. But KfW bonds are less liquid: smaller issuance volumes, fewer market makers, higher bid-ask spreads. The KfW-Bund spread is therefore a pure liquidity premium, with credit risk netted out by construction. - **Segmentation**: the residual after accounting for the above. It captures regulatory frictions (banks must hold certain sovereign bonds for liquidity requirements), home bias (domestic investors over-weight their own sovereign), and the "convenience yield" — the extra demand for the safest asset in the system. For German Bunds this is negative (investors accept yields below what default and liquidity would justify), reflecting Bunds' status as the eurozone's ultimate safe haven.

CS is too complex to derive fully (state-space model with 6 latent states, Kalman filter, 7 observables per country). It serves as the destination; Part B provides the tools to understand how such a decomposition is constructed. *Transition*: CS in action.

Slide 10 (CS decomposition in action — ECB, 2026)

The same figure as the opening slide, now readable with the vocabulary of slide 9. The colour coding in the right panel maps to the five components, making it possible to identify which force is driving current dynamics. *Transition*: the most famous single application of this methodology.

Slide 11 (CS headline result: PEPP announcement, March 18, 2020)

Italian 5-year yields had spiked to 196 bps as COVID fears triggered a flight from peripheral debt. The ECB's PEPP announcement (€750bn, no issuer limits) triggered a 78 bp fall in two days. The CS decomposition attributes this almost entirely to the credit and redenomination channels (−35 and −14 bps respectively) rather than to the expected-rates channel (−8 bps). PEPP worked not by signalling future rate cuts, but by directly eliminating the bad equilibrium of Module 2 — the self-fulfilling panic where Italian spreads rose because investors feared Italy couldn't borrow cheaply enough to stay solvent. The decomposition makes "whatever it takes" legible as a quantitative fact. *Transition*: building the tools.

Part B — A Teachable Model: Step-by-Step Derivation

Part B constructs the full model from scratch — from the spread definition to the closed-form decomposition — with every step explicit and every assumption stated.

Slide 12 (Part B divider)

Part B is the engine room. The structure is sequential — each step builds on the previous — and Step 4 is the payoff. The steps: define the object (Step 0) → reduce its dimensionality (Step 1) → add external information (Step 2) → impose no-arbitrage pricing (Step 3) → decompose (Step 4) → estimate (Step 5).

Slide 13 (Step 0: Working with spreads)

The object is $s_t(\tau) \equiv y_t^{LC}(\tau) - y_t^{US}(\tau)$, the spread of the local-currency sovereign zero-coupon yield over US Treasuries. Three reasons to work with spreads rather than yield levels: (1) it avoids estimating two full yield curve models jointly — one for the local economy and one for the US — which would roughly double the parameter space; (2) spreads are not subject to the zero lower bound that constrained US yields 2009–2015 and makes standard affine models misbehave; (3) the decomposition $s_t = EC_t + RP_t$ directly answers the policy question — "how much of this spread reflects expected monetary divergence versus risk compensation?" — without having to net out the benchmark contribution separately.

Technical note: the spread short rate $r_t^s = r_t^{LC} - r_t^{US}$ is the object that anchors the model. In practice it is proxied by the shortest available spread maturity (typically 3 months), since true overnight spread data can be noisy for EMEs.

Slide 14 (The data — Burban & Golinski figure)

The figure shows 3-month and 10-year spreads for Brazil, Colombia, Mexico from 2009 to 2025. Three observations stand out: (1) spreads widened sharply during COVID (March 2020), the Fed taper tantrum (2013), and the global financial crisis (2008–09); (2) the term structure of spreads — whether the 10-year spread is above or below the 3-month — varies over time, which is why the full spread curve matters, not just one maturity; (3) there are idiosyncratic country episodes (Brazilian political crisis 2015–16, Mexican energy reform reversals, Colombian fiscal debates). This cross-country and cross-maturity variation is what the model is designed to explain.

Slide 15 (Step 1: PCA on the spread curve)

With $N = 7$ maturities at each date, PCA is the natural dimensionality reduction: collect the spread curve into a vector $S_t \in \mathbb{R}^7$ and extract the leading principal components. Retaining $K = 3$ PCs captures more than 99% of spread variation across all countries — a well-known stylized fact of fixed income that holds for yield curves, credit curves, and spread curves alike. The three factors have a canonical interpretation: level (parallel shifts — the overall spread goes up or down), slope (the short end moves more than the long end, or vice versa), and curvature (the middle of the curve bows relative to the ends). Knowing the three PCs at any date is essentially equivalent to knowing the full spread curve.

Technical note: PCA is run on the covariance matrix of S_t, demeaned. The loadings W (the eigenvectors) are country-specific — the shape of the spread curve differs between Brazil and Mexico.

Slide 16 (Step 2: Adding risk factors beyond the spread curve)

The three PCs describe today's spread curve almost perfectly. But the goal is not just to fit today's curve — it is to decompose today's spread into "what the market expects future short rates to be" versus "what compensation the market requires for risk." To compute the expectations component, future short-rate differentials must be forecast. The question is whether today's spread curve already contains all the information needed, or whether other observables carry additional predictive power.

This is the **spanning hypothesis**: the yield (or spread) curve spans all relevant information for forecasting if no external variable has incremental predictive power for future rates. Joslin, Priebsch

& Singleton (2014) show that spanning is often rejected — macro and credit variables help predict yield changes above and beyond what the current curve reveals. Ignoring these unspanned variables would attribute to the expectations component variation that actually belongs to the risk premium, biasing the decomposition.

The two additional factors are intuitive: **credit risk** (the first PC of the CDS term structure — a summary of how much the market charges to insure against sovereign default at each horizon) and **FX volatility** (6-month implied vol — capturing currency uncertainty as a wedge between local and foreign investors' risk perceptions). They are orthogonalized to avoid double-counting: CR_t is the residual of a regression of raw credit risk on FX volatility, so CR_t captures "pure" credit risk independent of FX conditions. *Transition*: do these variables actually help forecast?

Slide 17 (Step 2 cont.: spanning test via VAR)

A VAR(1) on the joint state (X_t, U_t) is estimated. The key block is A_{XU}: the coefficient on U_t (credit and FX) in the equation for X_{t+1} (the spread PCs). If $A_{XU} \neq 0$, then today's credit risk and FX volatility improve the forecast of tomorrow's spread curve beyond what today's curve already predicts — the spanning hypothesis is rejected, and these variables belong in the model. Burban & Golinski find this for most countries at the 5% level. The test disciplines the choice of risk factors — CDS and FX vol are included because the data demands it.

Slide 18 (Step 3: The affine no-arbitrage term structure model)

The full state vector is $Z_t = (X_t', U_t')' \in \mathbb{R}^5$: three PCs plus two risk factors. Three assumptions make the model tractable:

(i) **Affine short rate**: $r_t^s = \delta_0 + \delta_1' Z_t$. The spread short rate is a linear function of today's state, estimated directly by OLS. (ii) **VAR(1) dynamics under** $\mathbb{P}$: the state evolves as a first-order vector autoregression under the physical (real-world) probability measure, estimated by OLS. (iii) **Affine pricing**: $s_t(\tau) = A(\tau) + B(\tau)' Z_t$. Under no-arbitrage, all maturity spreads must be affine in the state. This is a consequence, not an assumption — it follows from (i) and (ii) plus the no-arbitrage condition.

The term "affine" refers to the fact that yields are linear in the state variables — as opposed to quadratic or nonlinear models. The affine class is the workhorse of modern term structure modelling because bond prices and yields have closed-form expressions.

Slide 19 (Step 3 cont.: No-arbitrage recursions)

No-arbitrage requires a probability measure $\mathbb{Q}$ (the "risk-neutral" or "pricing" measure) under which all assets earn the risk-free rate. Under $\mathbb{Q}$, the state follows a different VAR with parameters $(\mu^{\mathbb{Q},\Phi^{\mathbb{Q}})}$.

The difference between $\mathbb{P}$ and $\mathbb{Q}$ dynamics is captured by the **market price of risk** — it quantifies how much extra return the market demands for exposure to each factor.

The Riccati recursions for $B(\tau)$ and $A(\tau)$ are standard and computable in milliseconds. The starting conditions $B(0) = 0$ and $A(0) = 0$ reflect that a bond maturing today is worth exactly 1 (zero spread). At each maturity τ, the recursion adds one more period of factor exposure. The **Joslin-Singleton rotation** is the insight that one can always choose the state vector so that its first K components coincide exactly with the observed PCs — eliminating the need to estimate latent states and making the model directly estimable by OLS. *Transition*: all the ingredients for the decomposition are in place.

Slide 20 (Step 4: The decomposition — the payoff)

The decomposition $s_t(\tau) = EC_t(\tau) + RP_t(\tau)$ is now a theorem, not a definition. The expectations component $EC_t(\tau)$ is the spread that would prevail if investors were risk-neutral — the average expected future short-rate differential over horizon τ. The risk premium $RP_t(\tau)$ is the residual: compensation for bearing uncertainty about how those future rates will evolve, plus credit, FX, and other risks embedded in the pricing measure $\mathbb{Q}$.

Transition: computing $EC_t(\tau)$ explicitly.

Slides 21–22 (Computing $EC_t(\tau)$ in closed form)

$E^{\mathbb{P}_t[r^s_{t+j}]}$ is needed for each horizon j. Since the state follows a VAR(1) under $\mathbb{P}$, iterated expectations give $E^{\mathbb{P}_t[Z_{t+j}]=\Phi^j Z_t+\sum_{k=0}^{j-1}\Phi^k\mu}$. The proof by induction is one line, illustrating the mechanics of linear forecasting. The key intuition: Φ^j governs how fast the state mean-reverts. For large j, if Φ is stable (all eigenvalues inside the unit circle), $\Phi^j \to 0$ and $E^{\mathbb{P}_t[Z_{t+j}]\to(I-\Phi)^{-1}\mu}$ — the unconditional mean of the VAR.

Substituting into $r^s_{t+j} = \delta_0 + \delta_1' Z_{t+j}$ and averaging over $j = 0, \dots, \tau - 1$ gives the closed-form expression for $EC_t(\tau)$. It is affine in Z_t — directly observable once the model is estimated. Since $\hat{s}_t(\tau) = A(\tau) + B(\tau)' Z_t$ is also affine, $RP_t(\tau) = \hat{s}_t(\tau) - EC_t(\tau)$ is affine in Z_t as well. The entire decomposition is linear in the state.

Slide 23 (Interpretation)

The table and the two scenarios ($EC = 180, RP = 20$ vs. $EC = 40, RP = 160$) make the decomposition's policy stakes concrete. The same 200 bps spread has completely different implications depending on which component dominates. In Scenario A, the spread reflects fundamentals — expected monetary policy divergence. It will narrow as monetary cycles converge, without any specific sovereign risk event. In Scenario B, the spread is driven by risk aversion, credit fear, or currency uncertainty. Policy action targeting the spread (asset purchases, IMF programme, currency intervention) is more likely to be warranted and effective. The decomposition is the analytical backbone of central bank communication about sovereign risk.

Slide 24 (Step 5: Estimation)

The estimation follows Golinski & Spencer (2024)'s factor extraction method, which avoids the computational complexity of maximum likelihood estimation for the full system. OLS for the physical dynamics, cross-sectional fitting for the risk-neutral parameters. No Kalman filter, no numerical optimisation over a high-dimensional parameter space. Every step is analytically explicit.

Part C — Empirical Results & Policy Implications

Part C shows what the decomposition reveals in practice for EMEs (Burban & Golinski) and for the euro area (CS results), connecting back to Modules 1 and 2.

Slide 25 (Part C divider)

Slide 26 (EME application: the decomposition in action)

The Burban & Golinski figure shows the decomposition of 10-year spreads for Brazil, Colombia, and Mexico from 2009 to 2025. Three results stand out: (1) the risk premium component is consistently larger and more volatile than the expectations component at long maturities; (2) the expectations component tracks the monetary policy cycle — widening when the Fed hiked in 2022–23, narrowing since the Fed pivot in 2024; (3) country-specific shocks appear in the risk premium but not in the expectations component: Brazilian political crises, Colombian fiscal strains, and Mexican energy policy changes leave clean signatures in the RP. The decomposition separates the global (Fed cycle) from the local (country risk).

Slide 27 (What drives spread variance? Credit risk vs. FX volatility)

The variance decomposition table quantifies the relative importance of credit and FX factors within the risk premium. Three results stand out: (1) **risk premium dominates long-term spread variance everywhere** — even in the most "fundamental" country, EC explains less than 50% of 10-year spread variation; (2) the composition of risk premia differs by country — Brazil is balanced between credit and FX, while Mexico is almost entirely FX-driven (credit risk co-movement with spreads near zero); (3) the expectations component is persistent, slow-moving, and correlated across countries — driven by the Fed cycle, it behaves as the common factor of EME spreads.

Technical note: the variance share is computed from the model-implied variance decomposition, using the VAR structure to attribute forecast error variance to each factor at each horizon. At short horizons (1–3 years), EC matters more; at long horizons (10 years), RP dominates almost everywhere.

Slide 28 (Back to the euro area: CS (2023) results)

The CS results read through the lens of the model just derived. For Italy and Spain, the dominant components are default and redenomination risk premia — the features specific to an incomplete monetary union (no lender of last resort, redenomination fear). For France and Germany, expected rates (driven by ECB policy) and the segmentation/convenience premium dominate. The negative German segmentation premium (-36 bps) reflects the "flight to safety" premium on Bunds — investors accept below-market yields to hold the eurozone's safest asset. The French redenomination spike before May 2017 (30 bps) is a clean natural experiment: a purely political shock visible in the redenomination premium but not in the default premium.

The PEPP example ties Modules 2 and 3 together: the "whatever it takes" logic of Module 2 (eliminating the bad equilibrium) is quantified in Module 3 as a -49 bps reduction in default and redenomination premia in two days.

Slide 29 (Summary)

Four takeaways mapping back through the lecture:

1. Sovereign yields/spreads embed multiple unobservable components — decomposition requires a model. [Part A]
2. The affine DTSM with unspanned risks provides a tractable, closed-form decomposition with all steps transparent. [Part B]
3. Empirically, risk premia dominate long-term spread variance, but their composition varies across countries. [Part C]
4. This framework is the analytical tool behind central bank yield monitoring (TPI, PEPP, OMT). [Connecting thread]

Slide 30 (References)

Connections to Modules 1 and 2

Module 3 concept	Module 1 connection	Module 2 connection
$EC_t(\tau)$ — expected short-rate differential	The monetary policy channel from Sargent-Wallace: fiscal dominance forces the CB to accommodate, raising expected future short rates	ECB rate expectations during the EZ crisis: why spreads widened even for "sound" countries
$RP_t(\tau)$ — risk premium	Default risk in the historical record (R&R): the premium investors demand reflects 800 years of partial defaults	Self-fulfilling component in the EZ crisis: the "bad equilibrium" shows up as an elevated RP above what fundamentals justify
Redenomination RP	Not present in Module 1 (historical defaults = outright)	The distinctive feature of the EZ crisis — the risk of euro exit that drove the ISDA basis
OMT/PEPP mechanism	Lender of last resort eliminates fiscal dominance dynamics	The commitment alone (never implemented OMT) collapsed the bad equilibrium — quantified as −49 bps in default + redenomination RP

References

- Corradin, S. & Schwaab, B. (2023). Euro area sovereign bond risk premia before and during the Covid-19 pandemic. *European Economic Review*, 153.
- Joslin, S., Priebsch, M. & Singleton, K.J. (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *Journal of Finance*, 69(3), 1197–1233.
- Golinski, P. & Spencer, P. (2024). Unconventional monetary policies and the yield curve. *Review of Asset Pricing Studies*, 14(1), 119–152.
- Burban, B. & Golinski, P. (2025). Credit risk and bond yield spreads in EMEs. *Banque de France internal note.*
- Fama, E.F. & Bliss, R.R. (1987). The information in long-maturity forward rates. *American Economic Review*, 77(4), 680–692.
- Cochrane, J.H. & Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1), 138–160.