



UNIVERSITÀ
DI PAVIA

Department of Economics and Management

Master of Science in Finance

**The French Sovereign Yield Curve across
Regimes: Political Risk, Monetary
Tightening and Market Dislocation**

**La Curva dei Rendimenti Sovrani Francesi
attraverso i Regimi: Rischio Politico, Stretta
Monetaria e Dislocazione di Mercato**

Supervisor:

Prof. Patrizio Tirelli

Candidate:

Andrea Faraon

Student ID: 555184

Academic Year 2025–2026

ABSTRACT – ENGLISH VERSION

This thesis replicates and extends the parametric yield-curve methodology that Grishchenko, Moraux and Pakulyak (2020) developed for the French nominal sovereign market. The Svensson/GSW estimation is re-run day by day on an independently assembled Bloomberg cross-section of 155 fixed-coupon OAT and BTAN lines over 10,077 trading days from October 1987 to June 2026, of which 2,129, some 21% of the sample, lie beyond the original study's April 2018 horizon. The extension is validated against the published results on the overlapping window and against the Banque de France's official zero-coupon estimates and it is then used to examine three forces that the original sample could not observe. First, political risk: the OAT-Bund safety premium, measured against a German curve fitted with the identical engine, roughly doubled around the June 2024 dissolution of the National Assembly; a daily decomposition attributes at most a fifth of its variation to common euro-area factors, leaving an idiosyncratic French premium that rose from a two-decade norm of zero to approximately 40 basis points, priced at medium- and long-run fiscal horizons rather than as near-term rollover risk, externally consistent with the official INSEE measure and transmitted to the cost of new corporate loans at a pass-through of about two-thirds. Second, the monetary regime: the slope factor share of daily curve variation more than tripled to 23.6% during the 2022–2024 inversion, a result visible only in changes and certified regime-by-regime by the Lord and Pelsser sign-change criterion. Third, pricing efficiency: a cross-section-normalized comparison identifies March 2020 as the most severe dislocation on record in the 5–10-year belly alone, fully repaired by the PEPP within a year and distinguishes it from a structurally higher noise floor in 2023–2025. A Markov-switching, Bai-Perron and CUSUM layer dates each regime formally, placing the premium structural break in the week of the dissolution. A single thread runs through all three results: the Eurosystem balance sheet has become a first-order determinant of the French yield curve.

ABSTRACT – ITALIAN VERSION

Questa tesi replica ed estende la metodologia parametrica di stima della curva dei rendimenti sviluppata da Grishchenko, Moraux e Pakulyak (2020) per il mercato sovrano nominale francese. La stima di Svensson/GSW viene rieseguita su base giornaliera su una sezione trasversale di dati Bloomberg assemblata in modo indipendente, comprendente 155 linee di OAT e BTAN a tasso fisso per 10.077 giorni di negoziazione, dall'Ottobre 1987 al Giugno 2026. Di questi, 2.129 giorni, pari a circa il 21% del campione, si collocano oltre l'orizzonte temporale dello studio originale, conclusosi nell'Aprile 2018. L'estensione viene validata confrontandola con i risultati pubblicati sulla finestra temporale sovrapposta e con le stime ufficiali zero-coupon della Banque de France, per poi essere impiegata nell'analisi di tre dinamiche che il campione originale non poteva osservare. In primo luogo, il rischio politico: il premio al rischio OAT-Bund, misurato rispetto a una curva tedesca stimata con il medesimo modello econometrico, è approssimativamente raddoppiato in concomitanza dello scioglimento dell'Assemblea Nazionale del Giugno 2024. Una scomposizione giornaliera attribuisce al massimo un quinto della sua variazione a fattori comuni dell'area euro, assegnando il restante a un premio idiosincratico francese che è salito da una media ventennale pari a zero a circa 40 punti base. Tale premio, prezzato su orizzonti fiscali di medio e lungo periodo piuttosto che come rischio di rifinanziamento a breve termine, risulta esternamente coerente con le misurazioni ufficiali dell'INSEE e si trasmette al costo dei nuovi prestiti societari con un pass-through di circa due terzi. In secondo luogo, il regime monetario: la quota di variazione giornaliera della curva spiegata dal fattore pendenza è più che triplicata, raggiungendo il 23,6% durante l'inversione del 2022–2024, un risultato statisticamente visibile esclusivamente analizzando le variazioni e certificato, regime per regime, tramite il criterio di cambio di segno di Lord e Pelsser. In terzo luogo, l'efficienza di prezzo: un confronto normalizzato per la sezione trasversale identifica Marzo 2020 come la più grave dislocazione mai registrata, limitatamente al segmento intermedio a 5–10 anni, completamente risanata dal PEPP entro un

anno e la distingue da un livello strutturalmente più elevato di rumore del mercato nel triennio 2023–2025. Un ulteriore livello di analisi basato su modelli Markov-switching, test di Bai-Perron e CUSUM data formalmente ciascun regime, collocando la rottura strutturale del premio politico nell'esatta settimana dello scioglimento parlamentare. Un unico filo conduttore lega tutti e tre i risultati: il bilancio dell'Eurosistema è divenuto una determinante di primo ordine per la curva dei rendimenti francesi.

INDEX

EXECUTIVE SUMMARY	7
INTRODUCTION.....	13
1. FRENCH SOVEREIGN DEBT MARKET	17
1.1 HISTORICAL OVERVIEW OF THE MARKET	17
1.2 MICROSTRUCTURE AND ARCHITECTURE OF THE MARKET.....	19
1.3 SOUCHE SYSTEM	20
1.4 MARKET MAKERS, REPO MARKET AND INSTITUTIONS	21
1.5 INVESTOR BASE AND VULNERABILITY TO EXOGENOUS SHOCKS.....	23
1.6 LIQUIDITY DEBATE: CORROBORATING AND COMPETING EVIDENCE	26
1.7 EMPIRICAL PREDECESSORS ON FRENCH DATA	28
1.8 THE CASE AGAINST SVENSSON	29
1.9 DYNAMIC TERM STRUCTURE MODELS	31
1.10 MACHINE-LEARNING FRONTIERS	33
2. METHODOLOGY	35
2.1 FROM NELSON-SIEGEL TO SVENSSON	35
2.2 THE PARAMETRIC CURVE.....	35
2.3 DATASET: SOURCES AND CONSTRUCTION.....	37
2.4 FILTERS APPLICATION	38
2.5 EXTENDED ESTIMATION.....	42
2.6 GOODNESS OF FIT METRICS	46
2.6.1 MEAN ABSOLUTE ERROR	47
2.6.2 HPW NOISE MEASURE	51
2.6.3 DIRTY PRICE RMSE.....	54
2.7 FIT OF THE NEW PERIOD	55
2.8 VALIDATION AGAINST THE PAPER	56
2.8.1 THE FITTED CURVE ACROSS REGIMES	58
2.8.2 HISTORY OF FRENCH TERM STRUCTURE.....	62
2.9 RESULTS AND INTERPRETATION.....	63
3. CASE STUDY 1 – POLITICAL RISK FACTOR	68
3.1 FACTS AND LITERATURE REVIEW	68
3.2 TWO COMPLEMENTARY LENSES	69

3.3	OAT-BUND SAFETY PREMIUM AND BREXIT EVENT STUDY	70
3.3.1	THE SAFE BENCHMARK: FITTING THE BUND CURVE.....	70
3.3.2	THE SAFETY-PREMIUM SPREAD AND TERM STRUCTURE	71
3.3.3	DATA-DRIVEN SHOCK DETECTION	75
3.3.4	CUMULATIVE ABNORMAL SPREAD CHANGE (CASC)	77
3.3.5	ROLLING OAT-BUND BETA	80
3.3.6	BREXIT-STYLE REGRESSION	81
3.4	DECOMPOSITION AND INSEE CROSS-CHECK	82
3.4.1	RELATED LITERATURE AND THE OFFICIAL BENCHMARK	83
3.4.2	METHODOLOGY.....	85
3.4.3	THE DECOMPOSITION DESIGN.....	86
3.4.4	THE TRANSMISSION DESIGN.....	87
3.4.5	RESULTS: EXTERNAL VALIDATION	88
3.4.6	RESULTS: REDUCED STAGE-1	90
3.4.7	THE CORE DECOMPOSITION.....	92
3.4.8	THE TERM STRUCTURE OF THE POLITICAL PREMIUM.....	96
3.4.9	TRANSMISSION TO THE COST OF CORPORATE CREDIT	99
3.4.10	ROBUSTNESS CHECKS	103
3.5	CONCLUSIONS	103
4.	CASE STUDY 2 – PCA OF THE OAT ACROSS REGIMES	106
4.1	METHODOLOGY	106
4.2	LORD AND PELSSER SIGN-CHANGE CERTIFICATION	108
4.3	THE INVERSION.....	109
4.4	FACTOR SHAPES: LOADINGS AND SIGN-CHANGE RULE	112
4.5	ROLLING PC2 SHARE AND LORD TEST FAILURE	115
4.6	ROBUSTNESS CHECKS.....	118
4.7	CONCLUSION AND ECB INTERPRETATION	119
5.	CASE STUDY 3 – MARKET DISLOCATION	121
5.1	THE METHODOLOGICAL PREMISE: INVALIDITY OF COMPARISON....	122
5.2	RELATED LITERATURE	123
5.3	PER-BIN SEVERITY TABLES.....	124
5.4	OVERALL NOISE: PEAK AND RECOVERY (ONE-YEAR WINDOW)	128
5.4.1	DISLOCATION AND RECOVERY IN EVENT TIME	130

5.4.2	THE ANATOMY OF THE TWO RECENT DISLOCATIONS.....	132
5.5	THE COVID EPISODE AND THE PEPP	134
5.6	ROBUSTNESS: LEAVE-ONE-OUT	137
5.7	FORMAL STRUCTURAL-BREAK ANALYSIS	139
5.7.1	THE ECONOMETRIC METHODS	140
5.7.2	DATED RESULTS AND CORROBORATION	141
5.7.3	COMPARISON WITH THE OTHER CASE STUDIES	147
5.8	CONCLUSIONS	150
6.	CONCLUSIONS	153
	REFERENCES.....	157

EXECUTIVE SUMMARY

INTRODUCTION AND OBJECTIVE

The French nominal sovereign yield curve is one of the main pricing benchmarks in the euro area. The most complete academic study of the market, Grishchenko, Moraux and Pakulyak (2020), ends in April 2018, before a sequence of events that materially changed the environment in which French government bonds are priced: negative interest rates, the March 2020 pandemic dislocation and the Eurosystem's response, the 2022–2024 monetary tightening cycle and yield-curve inversion and the domestic political instability that followed the June 2024 dissolution of the National Assembly.

The thesis has two objectives. First, it replicates the Svensson/GSW yield-curve methodology of Grishchenko, Moraux and Pakulyak on an independently assembled Bloomberg dataset and validates the resulting estimates against both the published study and the Banque de France's official zero-coupon curve. Second, it extends the analysis through June 2026 and uses the resulting curve and diagnostics to address three questions that the original sample could not study: whether recent political instability has generated a persistent French sovereign-risk premium; whether the 2022–2024 inversion changed the factor structure of the curve; and how severe the March 2020 market dislocation was, relative to earlier crises, and how fully pricing efficiency recovered afterwards.

METHODOLOGY AND EMPIRICAL DESIGN

The central methodological issue is that the object being measured can easily be confounded with the instrument used to measure it. A very flexible curve may absorb cross-sectional dislocation into the fitted curve itself; a French-German spread constructed from curves estimated with different methods may mix economic differences with methodological ones; a principal-component analysis in yield levels can be dominated by persistence and common trends; and aggregate noise measures are difficult to compare across crises when the number and maturity composition of available bonds change over time. The thesis therefore

adopts measurement choices tailored to each of these problems rather than relying on a single undifferentiated diagnostic.

The core instrument is the parsimonious six-parameter Svensson/GSW specification, estimated separately for each trading day from October 1987 to June 2026. The final dataset contains 10,077 trading days and a filtered universe of 155 fixed-coupon OAT and BTAN lines; 2,129 trading days, about 21% of the sample, lie beyond the April 2018 endpoint of the original study. Four features are important. First, there is no intertemporal smoothing, so parameter stability is an empirical outcome rather than an imposed restriction. Second, the German benchmark curve used in the political-risk analysis is estimated with the same engine as the French curve, preserving methodological comparability. Third, robustness exercises are implemented as full-sample re-estimations. Fourth, the data filters and outlier rules are codified and reproducible, so the complete estimation pipeline is auditable.

The replication is successful on the dimensions that matter for the subsequent analysis. On the overlapping sample, the fitted curve closely reproduces the published yield levels, fitting-error statistics and principal-component structure. The estimates are also cross-checked against the Banque de France's official zero-coupon curve. This validation matters because the extension is used not simply as a forecasting exercise but as a measurement framework for the three case studies that form the empirical core of the thesis.

KEY ECONOMIC FINDINGS

The three case studies deliver a consistent set of results. The political-risk analysis identifies a persistent country-specific French premium of about 40 basis points by the end of the sample, with common euro-area factors explaining at most about one-fifth of daily spread variation. The factor analysis shows that the slope factor rose to 23.6% of daily yield-curve variation during the 2022–2024 inversion when PCA is performed on changes rather than levels. The market-functioning analysis finds that the March 2020 episode generated the largest relative pricing

dislocation in the 5–10-year segment, while overall noise fell halfway back towards its pre-crisis baseline within 17 trading days.

Case Study 1 examines the political repricing of the OAT-Bund safety premium. Around the June 2024 dissolution, the average 10-year spread more than doubled, from about 28.6 basis points in the pre-dissolution regime to about 69.4 basis points afterwards, while the OAT continued to co-move closely with the Bund. The change is therefore best described as a repricing in level rather than a decoupling from the euro-area core. A daily decomposition using common euro-area factors – including peripheral sovereign dispersion, systemic stress and the Bund level – explains at most roughly one-fifth of daily spread variation. The residual French component, close to zero over much of the earlier sample, rose to about 40 basis points by June 2026 and peaked at 47.7 basis points in October 2025. The repricing is concentrated at medium and long maturities: the 5–30-year spread moved roughly twice as much as the 2-year spread around the dissolution, consistent with a reassessment of medium-term fiscal and political risk rather than an immediate refinancing scare. The thesis also documents a sizeable association with corporate borrowing costs: the estimated pass-through to rates on new loans to French non-financial corporations is about two-thirds, implying an effect of roughly 0.2 percentage points on marginal borrowing costs. This interpretation is externally consistent with the INSEE evidence, while remaining explicitly reduced-form and associational.

Case Study 2 studies the 2022–2024 inversion through principal-component analysis across regimes. The main result appears only when the PCA is performed on daily yield changes. In normal regimes, the slope factor accounts for about 7–8% of daily curve variation; during the inversion its share rises to 23.6%, while the level factor becomes correspondingly less dominant. The Lord and Pelsser sign-change criterion validates the economic interpretation of the first three factors as level, slope and curvature in the relevant regime. The slope loading is concentrated at the short end – approximately -0.95 at one year against small positive loadings at longer maturities – which is consistent with a tightening-driven inversion in

which policy-sensitive short rates moved much more sharply than the long end. The result therefore supports the interpretation developed in the thesis and in the ECB-related literature: the euro-area inversion was not a clean recession signal, because it reflected both the rapid monetary tightening and a long end still affected by the inherited stock of asset purchases.

Case Study 3 analyses market dislocation and recovery using cross-section-normalized, maturity-bin measures of pricing noise. This qualification is essential: because the OAT cross-section is small and changes over time, raw aggregate fitting errors are not directly comparable across crises. Once maturities are held comparable, March 2020 emerges as the most severe dislocation on record only in the 5–10-year belly, at about 11 basis points, roughly twice the severity of earlier episodes in that segment; it was not the most severe event at every maturity. Overall noise reached 6.8 times its pre-crisis baseline and its recovery half-life was 17 trading days, as fast as the Lehman episode and much faster than the euro-area sovereign crisis. By late 2021 pricing noise had returned below its pre-pandemic level. The thesis therefore distinguishes the rapid repair of the pandemic dislocation, associated with the PEPP backstop, from the structurally higher noise floor that emerged after 2022. The latter is treated as a separate regime, consistent with quantitative tightening and tighter post-crisis constraints on dealer intermediation rather than as lingering COVID-related dysfunction.

A common structural-break framework – Markov switching, Bai-Perron tests and CUSUM diagnostics – is then applied to the main variables. It places the monetary-regime change in late 2022 and the high-premium political regime in June 2024, around the dissolution. This is important for the interpretation of the whole thesis: the monetary and political developments are sequential and independently timed rather than two manifestations of a single shock. Read together, the three case studies also point to a broader unifying theme: the Eurosystem balance sheet has become an important determinant of the French yield curve, affecting its slope, the restoration of market functioning during the

pandemic and the post-2022 environment in which cross-sectional pricing noise remains elevated.

IMPLICATIONS, LIMITATIONS AND EXTENSIONS

The evidence suggests that the recent French sovereign premium should be treated as persistent over the horizon observed in the thesis rather than assumed to revert immediately to its pre-2022 norm. Its concentration at medium and long maturities is more consistent with concerns about fiscal consolidation capacity and long-run debt sustainability than with near-term refinancing risk. At the same time, the results do not imply that France has ceased to be a core euro-area sovereign: the OAT continues to co-move closely with the Bund and the market remains deep and liquid. The more precise conclusion is that core status now coexists with a persistent country-specific premium.

The methodological implications are equally important. Single-tenor spreads can measure the size of a repricing but do not reveal its term-structure composition. For that purpose, tenor-resolved French and German curves estimated with an internally consistent methodology are preferable. Likewise, market-functioning comparisons should rely on maturity-specific or otherwise cross-section-normalized noise measures rather than raw aggregate fitting errors, which can confound changes in market stress with changes in the composition of the bond universe.

The interpretation of the results must remain within the limits of the empirical designs. The event studies, decompositions and reduced-form regressions are associational and do not identify structural causal parameters. The idiosyncratic French premium is a residual and therefore combines political risk with any remaining liquidity and measurement components. The corporate-loan transmission exercise omits a US long-yield driver whose relevance is visible in the diagnostics. In addition, the dataset is quote-based and the underlying cross-section is relatively small and heterogeneous, while the politically relevant 2025–

2026 subsample is still short. These limitations call for caution in interpreting magnitudes and for additional data as the post-2024 regime lengthens.

The thesis nevertheless provides a useful framework for continued monitoring. The fitted curves and regime diagnostics can be updated as new observations arrive to assess whether the high-premium state eventually mean-reverts, as earlier political episodes did, or becomes a more persistent feature of French sovereign pricing. The broader economic implication is therefore not one of an acute funding crisis, but of a market that has become structurally noisier, materially higher-yielding and more exposed to both domestic political risk and the configuration of the Eurosystem balance sheet.

Three natural extensions follow. First, a local-projection instrumental-variable design using political-event innovations could move from the associations documented here towards causal dynamic responses. Second, arbitrage-free Nelson-Siegel or shadow-rate term-structure models could be used to price and decompose the premium rather than only measure it. Third, the regime-classification layer can be re-estimated recursively as new data arrive, allowing the persistence of the post-2024 premium and the post-2022 noise regime to be assessed in real time.

Overall, the contribution of the thesis is to combine a validated long-run yield-curve reconstruction with three purpose-built empirical applications. The extension of the French curve through June 2026 is not merely an update of the original sample: it reveals a sovereign market in which political repricing, monetary-regime change and market-functioning stress are distinct phenomena that require different empirical lenses, but are linked by the increasingly important role of the Eurosystem in shaping French sovereign yields and market functioning.

INTRODUCTION

The term structure of interest rates, commonly referred to as the yield curve, represents a foundational construction in both macroeconomic policy and financial asset pricing. By mapping the relationship between the maturity of risk-free debt securities and their corresponding yields, the curve synthesizes market expectations regarding future monetary policy, inflation trajectories and sovereign credit risk. Within the European Economic and Monetary Union, the French sovereign public debt market, managed by the Agence France Trésor (AFT), stands as a critical pillar. Comprising primarily Obligations Assimilables du Trésor (OAT), the French market is one of the largest sovereign debt markets in Europe and serves as a vital benchmark for pricing fixed-income assets globally. Historically, the French yield curve has been characterized by a high degree of structural liquidity, a broad and diverse international investor base and a negligible on-the-run premium, which significantly distinguishes its microstructure from the US Treasury market.

These features are not incidental: they are the product of an issuance architecture, the souche system of fungible re-openings, a designated network of primary dealers and a deep repo market [5], [8], that concentrates liquidity in benchmark lines and makes the French curve estimable with unusual precision. It is precisely this combination of size, liquidity and institutional transparency that has made the OAT market a recurrent laboratory for term-structure estimation, from the official Banque de France curves [10], [24] to the comparative euro-area work of the ECB [27]. The most complete academic treatment remains Grishchenko, Moraux and Pakulyak [1], who fit a Svensson curve to 179 nominal French bonds from 1987 to April 2018 and deliver, alongside the estimates, a diagnostic apparatus, namely fitting-error metrics, a Hu-Pan-Wang-style noise measure, an on-the-run premium test and a principal-component decomposition, that turns the fitted curve into an instrument for reading the market rather than a mere interpolation device.

That study, however, closes at a singular moment: its sample ends in April 2018, before every structural episode that has since redefined the French market. In the eight years that follow, the OAT curve descended into negative territory across most of its maturities, absorbed the March 2020 pandemic dislocation and the Eurosystem's response to it, inverted under the fastest monetary tightening in the euro history and, uniquely among core euro sovereigns, was repriced by a domestic political crisis that produced five prime ministers in under two years, an Excessive Deficit Procedure and a sequence of rating downgrades. The main research motivation of this thesis follows directly: discover, analyse and understand what happened to the French sovereign yield curve, and to the diagnostics built on it, when the methodology of [1] is carried, on a like-for-like basis, through these last vibrant eight years.

In order to address this question, it was first necessary to re-establish the instrument. The Svensson/GSW pipeline is rebuilt from primary Bloomberg data and re-estimated daily over 10,077 trading days on a filtered universe of 155 fixed-coupon securities. Three commitments are fundamental. Each day is estimated in statistical isolation, with no intertemporal smoothing imposed, so that the stability of the parameter paths is an outcome to be verified rather than an assumption built in [16], [18]; every robustness check is a complete re-run of the forty-year sample under the alternative configuration, never a spot check, and the extension is validated externally too, against the published results of [1] on the overlapping window and against the official Banque de France zero-coupon curves [24], before any inference is drawn from it. The closing diagnostics of that replication – a level that completes a full monetary cycle to 3.74% by June 2026, a slope that inverts on 233 trading days between 2022 and 2024 against a single day in the entire original sample, and noise and fitting-error spikes clustered on identifiable crisis dates – are what motivate and delimit the three case studies that constitute the empirical core of the thesis.

Each case study isolates one force acting on the curve and answers a piece of the original motivation. Case Study 1 asks whether political instability demoted the OAT from a pure core asset: the safety premium is measured as the spread over a German curve fitted with the identical engine, the design choice that guarantees the spread is a credit and liquidity differential rather than a methodological artefact and is examined through two complementary lenses, an event study modelled on the Brexit design of [34] and a decomposition into common euro-area and idiosyncratic French components in the spirit of [35], externally validated against the official INSEE measure and extended, following the same source, to the transmission of the premium to the cost of new corporate loans [36], [32]. Case Study 2 asks whether the 2022–2024 inversion changed the factor structure of the curve, applying the Principal Component Analysis (PCA) across regimes under the sign-change certification of Lord and Pelsser [23] and reading the result against the ECB’s own interpretation of the euro-area inversion [41], [26], [43]. Finally, Case Study 3 asks how severe the March 2020 dislocation was and how completely the PEPP repaired it, insisting on a cross-section-normalized, per-maturity-bin comparison because the small and heterogeneous OAT cross-section makes a raw aggregate comparison invalid and closing with a formal regime-dating layer, with Markov-switching, Bai-Perron and CUSUM, applied identically to the noise, the slope and the political premium.

The contribution is therefore threefold. Empirically, the thesis documents a French sovereign market that is, on the evidence of its own curve, structurally noisier, materially higher-yielding, politically repriced and more dependent than ever on the central bank at its centre. Methodologically, it argues that each of these claims requires a purpose-built measurement, an internally consistent benchmark curve for the premium, a decomposition in changes rather than levels for the factor structure, a per-bin normalization for the dislocation and a dated regime layer to certify that what is being described is a regime at all and it delivers those measurements as a single auditable pipeline in which every modelling choice lives

in one configuration file. Analytically, reading the three studies together yields what none delivers alone: the monetary and political forces acting on the French curve are shown to be distinct, independently timed regimes rather than two faces of one shock, while the Eurosystem balance sheet emerges as the common determinant running through all three.

The thesis proceeds as follows. Chapter 1 describes the institutional architecture of the French sovereign debt market and reviews the literature on its liquidity, on the estimation methodologies available and on their known limitations [10], [11], [15], [17], [19] and [21]. Chapter 2 presents the methodology, the dataset, the filters, the goodness-of-fit metrics and the validation of the extended estimation against [1]. Chapters 3, 4 and 5 develop the three case studies. Chapter 6 delivers the general conclusions drawing the threads together and sets out the limitations and the directions this work leaves open.

1. FRENCH SOVEREIGN DEBT MARKET

Extensive econometric modelling of the French nominal yield curve has been conducted for the period spanning the late 1980s to early 2018. However, to accurately interpret the econometric movements in the curve, it is imperative to first understand the underlying microstructure of the French public debt market, as its specific architecture profoundly influences liquidity, pricing dynamics and susceptibility to macroeconomic shocks.

1.1 HISTORICAL OVERVIEW OF THE MARKET

The modern French government bond market is the result of a process of reform undertaken by the Trésor in the mid-1980s. Before that reform, French state debt was a heterogeneous collection of bond lines and Treasury bills with non-uniform conventions, placed substantially with a captive domestic clientele and, consequently, neither especially deep nor especially liquid in the secondary market. The decisive institutional innovation was the introduction of a small number of standardized, fungible instruments, the Obligations Assimilables du Trésor (OAT) for long-term financing, the Bons du Trésor à Taux Fixe et à Intérêt Annuel (BTAN) for the medium-term segment, first issued on February 11th, 1986, and the Bons du Trésor à Taux Fixe et à Intérêts Précomptés (BTF) for short-term cash management, together with a designated network of primary dealers and a futures market. The Marché à Terme International de France (MATIF), opened in 1986, introduced exchange-traded interest-rate derivatives, notably the notional bond future, and gave the cash market a hedging and price-discovery complement that materially deepened it.

The subsequent three decades can be read as a progressive consolidation of credibility and liquidity, punctuated by exogenous macro-financial shocks. The legal infrastructure for collateralized funding was modernized with the pension livrée, which is the delivery-versus-payment repo, framework in 1988; coupon stripping authorized in 1991, allowing arbitrage between coupon bonds and the

resulting zero-coupon certificates; the floating-rate OAT TEC 10 was launched in 1996 and the two segments of primary dealers, the *Spécialistes en Valeurs du Trésor* (SVT) and the *Spécialistes en Pensions sur Valeurs du Trésor* (SPVT), were merged in 1994 into a single panel of roughly twenty market makers, unifying cash and repo market-making obligations. Index-linked issuance followed with the first OATi, indexed to French CPI, in September 1998 and the first OAT€i, indexed to euro-area HICP, in October 2001.

Two structural breaks dominate the sample. The first is the January 1st, 1999, introduction of the euro, which redenominated the stock into a single currency, eliminated intra-EMU exchange-rate risk and opened the market to a large pool of non-resident investors. Both the empirical literature and the diagnostics reported further on in this thesis find that market functioning improved discontinuously thereafter to the point that the pre-euro and euro periods are best treated as distinct regimes. The institutional counterpart to this maturation was the creation of the Agence France Trésor (AFT) on February 8th, 2001, as a dedicated debt-management office. The second break is the sequence of crises that tested the resilience of the market: the 2008–2009 Global Financial Crisis, the 2011–2012 euro-area sovereign debt crisis, during which France lost its AAA rating from Standard & Poor's on January 9th, 2012, and at the long end of the policy spectrum, the prolonged regime of unconventional monetary policy that drove yields to and below the zero lower bound between roughly 2014 and 2022.

The simplification of the product range in 2013, when the AFT ceased issuing BTANs under their own name by folding the medium-term segment into the OAT line and, in compliance with the post-crisis euro-area framework, began embedding Collective Action Clauses (CACs) in new issues, closes the period studied. The last BTAN was redeemed on July 25th, 2017. The extension carried out in this thesis covers precisely the episodes not yet studied in the literature: the negative-rate environment, during which France issued 0% coupon OATs and, in January 2021, a fifty-year line maturing in 2072, the violent return of inflation in


2022–2024 and one of the most aggressive tightening cycles in European Central Bank (ECB) history, the associated inversion of the curve and the 2024–2025 episode of domestic political fragmentation characterized by successive minority governments, a public-debt ratio near 114% of GDP and a Fitch downgrade to A+ in September 2025, which reintroduced a visible political-risk component into OAT pricing. These episodes are the subject of three case studies contained in this thesis.

1.2 MICROSTRUCTURE AND ARCHITECTURE OF THE MARKET

The architecture of the market is organized around a single issuer-agent, the AFT, a state agency operating under the authority of the Directorate General of the Treasury with a mandate to raise necessary funds at the lowest possible cost while maintaining maximum market safety and transparency. This dual objective, minimizing the expected funding cost subject to a risk constraint, is operationalized through a predictable, transparent and rules-based issuance calendar of competitive auctions, namely the *adjudications*, intermediated by the primary dealers, complemented by occasional syndications for new long-dated or specialized lines and by an active buy-back and exchange programme used to smooth the redemption profile and manage the average maturity of the stock.

The instrument set is deliberately parsimonious and is composed of mainly three instruments spanning the maturity spectrum which are, by design, close analogues of the US Treasury Bill, Note and Bond and of the German Bubill, Schatz, Bobl and Bund triplets, namely: BTF, discount bills redeemed at par, with maturities expressed in weeks, typically 13, 26 and 52, and original tenor at or below one year, used for short-term cash management; BTAN, fixed-coupon medium-term notes with original maturities of two to five years, no longer issued since 2013 and fully extinguished by 2017, but present in the early part of the sample; OAT, fixed-coupon annual-paying long-term bonds, now the main actor of the entire medium to long segment. The maturity structure of the French market over the sample period is plotted in Figure 1.

Figure 1



Source: author's elaborations on Bloomberg data

Around this nominal core the AFT has developed a set of specialized lines that constitute the current product range: inflation-linked bonds indexed to French CPI and to euro-area HICP and, since 2017, also sovereign green OATs, effectively making France the first sovereign nation to issue a benchmark-size green bond whose proceeds are used for environmentally targeted expenditure; all of these notes are euro-denominated. The legacy of foreign-currency lines, mainly in US Dollars and British Pounds, that historically represented less than 3% of the stock, the inflation-linked segment and the green segment are excluded from the nominal curve estimation, consistent with the literature.

1.3 SOUCHE SYSTEM

A defining and unique characteristic of the French debt market is the empirical absence of a distinct on-the-run premium, a phenomenon that sharply contrasts with the largest and most important debt market in the world, the US Treasury

market. In the United States, newly issued bonds, designated as on-the-run, are highly sought after by market participants for their superior liquidity, causing them to trade at a premium and consequently a lower yield, relative to older, off-the-run bonds of similar maturities. This scarcity premium complicates the yield curve estimation, often requiring researchers to exclude on-the-run securities from their cross-sectional dataset to avoid pricing distortions.

In France, this liquidity premium is effectively neutralized by the implementation of the *souche* issuance system. When the French government requires additional financing, it frequently refrains from creating a completely new debt security. Instead, the AFT reopens an existing *souche*, or primary issue matrix. These newly auctioned OAT are perfectly fungible with the seasoned bonds already trading in the secondary market, sharing identical coupon rates, maturity dates, face value and International Securities Identification Numbers (ISIN). This continual reopening process ensures an abundant and continuous supply of specific maturities, thereby preventing the scarcity effects that typically generate liquidity premia and allowing for a highly efficient pricing mechanism across the entire term structure, without a proliferating set of small, heterogeneous securities.

1.4 MARKET MAKERS, REPO MARKET AND INSTITUTIONS

Market liquidity and continuous price discovery are further guaranteed by a highly regulated network of primary dealers, the SVT. Selected by the AFT under a published charter, the SVTs commit to take up a minimum share of each auction, to quote continuous two-way prices in the secondary market, to provide market intelligence and placement to the AFT and to support the liquidity of the repo market in individual lines. As the obligatory market makers on the inter-dealer electronic platforms on which the benchmark OATs trade, their quoting obligations are the proximate mechanism that translates the elastic supply of the *souche* system into observable, executable prices and it is precisely this uninterrupted stream of firm two-way quotes that later constitutes the cross-section of prices to which the yield curve is fitted. The cash market is complemented by a deep listed OAT

futures market, now cleared on Eurex, which provides hedging capacity, a second venue for price discovery and an arbitrage anchor for the cash curve.

For that cross-section to be internally consistent, however, quoting obligations alone are not sufficient. A dealer will keep a relatively expensive bond in line with the rest of the curve only if that bond can be sold short and a relatively cheap one bought and both legs of such relative-value trade must be funded. This is the economic role of the repo (repurchase agreement) market. By allowing dealers and investors to borrow specific securities against cash or other collateral, the repo market lets short positions be covered and settlement fails be avoided; it is therefore the channel through which arbitrage capital actually reaches the bond cross-section. Where repo funding is deep and cheap, relative mispricings between individual lines are competed away and every OAT trades close to a single, common discount function; where it is impaired, individual securities can drift rich or cheap on their own funding conditions, injecting idiosyncratic noise into the cross-section that no smooth curve can reproduce.

The price of this funding liquidity is directly observable in the repo rate itself. A line that becomes scarce trades special, commanding a repo rate below the general-collateral (GC) rate at which OATs are otherwise financed; the holder of a special bond can lend it and fund cheaply, earning that spread, which is capitalized into a higher cash price and hence a lower yield. Persistent specialness is therefore a genuine pricing distortion, the funding-market analogue of the on-the-run premium discussed in Section 1.3, since it moves a bond's yield for reasons that have nothing to do with the term structure the curve is meant to capture. In France this distortion is normally small and short-lived, because the continual reopening of *souches* replenishes the supply of any given line and the AFT itself stands ready, through its primary dealers, to lend scarce securities against other government collateral, curtailing specialness and settlement fails even under stress. Taken together, the *souche* system and a deep, well-supplied repo market keep every OAT priced off the same curve and this is precisely the condition under which the

Svensson estimation of Chapter 2 is well posed. The fit treats each observed price as the model discount function plus a small pricing error and reads the residual as noise rather than as structural liquidity premia, which is exactly what the HPW noise measure of Section 2.6.2 quantifies. It is because French repo specialness is contained that the entire eligible cross-section can be retained, without excluding the most recently issued lines as is customary on US Treasury data and that the resulting fitting errors can be interpreted as a gauge of market functioning rather than as an artefact of segmented liquidity.

This microstructure rests on a standard post-euro legal and operational infrastructure of master repo agreements, settlement through the central securities depository Euroclear France and central clearing at LCH. Two institutions sit above it and increasingly shape the very prices the curve is fitted to: the Banque de France, which acts as fiscal and settlement agent and, within the Eurosystem, conducts the monetary-policy operations for which OATs are eligible collateral and the ECB, whose asset-purchase programmes, the APP and the pandemic-era PEPP and more recent backstops such as the Transmission Protection Instrument, have become first-order determinants of OAT supply, liquidity and spreads. The interaction between ECB intervention, sovereign credit risk and market liquidity is taken up in the case studies of this thesis.

1.5 INVESTOR BASE AND VULNERABILITY TO EXOGENOUS SHOCKS

The composition of the investor base has been transformed over the market's lifetime and this transformation is essential to interpret the curve's sensitivity to shocks. In the 1980s the OAT clientele was overwhelmingly domestic and institutional, namely the *zinzins*, which were composed of French institutional investors such as insurers, pension and provident funds, mutual funds and large state-owned entities, many of them legally required to hold long-term assets and structurally inclined to buy-and-hold with very little incentive to trade actively or to lend in the repo market. The introduction of the euro changed this ecosystem decisively: by eliminating currency risk and integrating the EMU bond markets it

drew in a large pool of non-domestic investors; their number continued to grow relentlessly over the years peaking at over 70% in June 2010. They have receded since then, as Banque de France and AFT figures show, as the Eurosystem's own purchases and a partial repatriation to domestic insurers and banks reshaped the holder structure. The modern weight of this domestic institutional segment is documented by Ahado, Barnava, Chantrelle, Chevallier and Desseaux [2]: at end-2024 French insurers and pension funds held $\approx$ €2,950 billion of investment assets, of which some 57% in bonds, overwhelmingly euro-denominated, concentrated in highly rated issuers and with only a marginal share in low-rated assets. This composition deepens domestic demand for OATs, but it also embeds a conditional fragility that the raw holdings figures conceal: portfolios governed by prudential regimes and investment mandates that reference credit ratings are exposed to threshold effects, so that a deterioration of the sovereign rating, of the kind associated with the political-instability episode examined in Case Study 1, could convert today's stable buy-and-hold demand into rating-driven selling pressure. The forced-sale channel is an implication of the mandate structure instead of a documented episode, but it is precisely the kind of latent asymmetry that the internationalized holder base already displays on the foreign side. However, France remains, among large euro-area sovereigns, exceptionally reliant on foreign demand relative to more domestically funded markets such as Italy or Japan.

This internationalization is a double-edged property. In normal times it deepens demand, compresses yields and reinforces the OAT's status as a core euro safe asset, though it also makes the curve structurally exposed to portfolio decisions of investors who are not anchored by domestic regulatory frameworks or home-bias considerations and who can therefore reallocate quickly in response to credit-rating actions, fiscal news or political instability. Two main dynamics of the euro-area bond market can be distinguished: flight-to-quality and flight-to-liquidity. In particular, in situations of stress that afflict all European countries the OAT has at times benefited from safe-haven inflows out of the periphery, principally PIIGS

(Portugal, Ireland, Italy, Greece and Spain) countries, while in idiosyncratic French stress situations, as in 2011–2012 and again in 2024–2025, the same mobile foreign capital can withdraw and widen the OAT-Bund spread, temporarily repricing the OAT away from a pure core asset.

The concentration and foreign inclination of the holder base is shown in Table 1, which captures the twenty-five largest reported institutional holders of French government debt.¹ Two main features stand out. Firstly, ownership is highly concentrated: these twenty-five holders account for roughly 80% of the disclosed-holdings universe and furthermore the three largest alone, namely Japan Government Pension Investment Fund (~18%), Vanguard (~13%) and BlackRock (~7%), represent close to 40% of it. Secondly, the base is overwhelmingly non-domestic: of the portion that can be classified by nationality, foreign holders represent over 69% of the universe against only 11% for French holders, a ratio of almost six to one, or about 85% foreign versus 15% domestic once re-based on the total. The largest single position is itself a foreign public pension fund while US asset managers occupy the two very next ranks. The principal domestic names that do appear, Crédit Agricole, BNP Paribas, AXA and AG2R La Mondiale, are the large French banks and insurers, meaning the modern descendants of the already mentioned *zinzins*.

¹ These are Bloomberg disclosed institutional positions, not the full holder universe: they over-represent reporting asset managers and under-represent the Eurosystem, official reserve managers and untraced custody holdings, so the sample ~85% foreign share is an upper bound on the true non-resident fraction, which the aggregate Banque de France statistics put nearer ~50%. Even under the conservative aggregate figure, however, the conclusion is unchanged: a large, mobile and predominantly foreign investor base is the structural reason why the French curve, liquid and premium-free in normal times, can nonetheless reprice abruptly when an exogenous shock alters the perceived safety of the OAT.

Table 1

Bondholder Name	Share owned (K €)	% Tot
Govt Pension Inv Fund Japan	33.145.166,19 €	18,71%
Vanguard Group Inc	24.042.716,31 €	13,57%
Blackrock	13.350.893,50 €	7,54%
Credit Agricole Group	11.035.563,61 €	6,23%
Intesa Sanpaolo SPA	7.917.126,07 €	4,47%
BNP Paribas	6.001.594,50 €	3,39%
Allianz SE	5.309.508,47 €	3,00%
Deutsche Bank AG	4.801.516,51 €	2,71%
Caixbank Asset Management	3.972.963,74 €	2,24%
HSBC	3.212.955,79 €	1,81%
Nordea Bank ABP	3.078.315,19 €	1,74%
State Street Corp	2.480.852,61 €	1,40%
Royal Bank of Canada	2.469.677,95 €	1,39%

Bondholder Name	Share owned (K €)	% Tot
Banco Bilbao Vizcaya Argentaria	2.455.792,43 €	1,39%
JP Morgan Chase & CO	2.432.621,99 €	1,37%
KBC Group Inv	2.399.048,24 €	1,35%
UBS AG	2.198.450,99 €	1,24%
AG2R La Mondiale Gestion d'Actif	2.137.847,90 €	1,21%
Danske Bank A/S	2.037.513,82 €	1,15%
AXA	1.549.684,95 €	0,87%
Marsh & McLennan Cos Inc	1.543.286,19 €	0,87%
Banco Santander SA	1.534.466,79 €	0,87%
FMR LLC	1.443.288,56 €	0,81%
Banca Mediolanum SPA	1.400.266,41 €	0,79%
Aegon LTD	1.102.061,94 €	0,62%

TOT - sample	143.053.180,65 €	
TOT	177.152.144,25 €	80,75%

TOT - French holder - sample	20.724.690,96 €	11,70%
TOT - Foreigner holder - sample	122.328.489,69 €	69,05%

Source: author's elaborations on Bloomberg data

1.6 LIQUIDITY DEBATE: CORROBORATING AND COMPETING EVIDENCE

The institutional architecture described so far is not merely descriptive background: it embodies a precise and contested position in the literature on how liquidity enters sovereign bond prices. The premise inherited from [1] and maintained throughout this thesis is that the design of the French market compresses liquidity-driven price differentials across OATs to the point where a single frictionless discount function is an adequate pricing benchmark, so that deviations of individual bonds from the fitted curve can be read as noise rather than as structural premia. This premise enjoys independent empirical support. Ejsing and Sihvonen [3] examine the German market, whose architecture of fungible reopenings and dominant futures trading resembles the French one far more closely than the US Treasury market does and find that, in sharp contrast to the US evidence, on-the-run status has a negligible effect on both the liquidity and the pricing of German government bonds once other factors are controlled for, with liquidity spillovers accruing instead to the bonds deliverable into the futures

contracts; tellingly, the French 10-year OATs included in their traded-volume comparison display the same flat volume profile over the bond life cycle, corroborating on independent data the absence of the on-the-run phenomenon documented in Section 1.3. Christensen, Lopez and Shultz [4] complete the picture from the American side, showing that even within the US the premium is a creature of market design rather than a universal law: recently issued TIPS command a premium of merely one to four basis points, an order of magnitude below that of their nominal counterparts. The issuer's own doctrine, finally, is explicit: Copin and Dalbard [5] describe the liquidity of the debt as "the main endogenous factor in bond pricing" and present the entire AFT issuance strategy as designed around its preservation.

A competing strand of the literature argues, however, that liquidity in euro-area sovereign markets is neither negligible nor safely relegated to the error term. Monfort and Renne [6] model euro-area sovereign spreads with credit and liquidity as distinct, stochastic risk factors whose intensities follow regime-switching dynamics and show that liquidity premia are time-varying, priced and occasionally dominant, particularly in crisis regimes; in their reading, a static frictionless fit risks misattributing a systematic liquidity component to idiosyncratic noise. The empirical anatomy of this channel is provided by Beber, Brandt and Kavajecz [7], who show that, while credit quality explains the bulk of yield-spread levels, in episodes of market stress order flow chases liquidity almost exclusively and that liquidity matters most precisely for low-credit-risk issuers, which is the segment France occupies. In the same spirit, the post-2022 evidence assembled by the ICMA Bond Market Liquidity Taskforce [8] documents that liquidity in the core European sovereign markets, France included, has become markedly more sensitive to volatility than bursts of volatility alone can explain, the signature of intermediaries retreating from liquidity provision; Duffie, Fleming, Keane, Nelson, Shachar and Van Tassel [9] supply the micro foundation, showing on US

data that market functionality deteriorates sharply once dealer balance-sheet capacity becomes binding.

These two readings are less irreconcilable than they first appear and the diagnostics constructed in Chapter 2 are precisely the point where they meet. In tranquil regimes the corroborating evidence prevails: the fitting errors of the euro era are small, unsystematic and mean-reverting, exactly what a frictionless benchmark predicts. In stressed regimes the stochastic-liquidity view of Monfort and Renne [6] takes over: residuals widen, become systematically signed along the curve and cluster in the least liquid segments, as the 2008 snapshot, represented in Figure 5, and the post-2022 ultra-long end analysed in Chapter 2 will show. The methodological stance of this thesis is therefore not that liquidity premia do not exist, but that measuring their episodic emergence requires a rigid, frictionless measuring rod from which they can be seen to deviate: the MAE and noise diagnostics of Chapter 2 are exactly such an instrument and their spikes mark the moments in which the competing literature is right.

1.7 EMPIRICAL PREDECESSORS ON FRENCH DATA

The estimation of smooth parametric curves on French government bond data predates [1] by a quarter of a century and its early history already contains the methodological comparison at the heart of this thesis. Within the BIS survey of central bank estimation practices, Ricart, Sicsic and Jondeau [10] document the Banque de France procedure on franc-era data: both the original Nelson-Siegel function and the Svensson extension are estimated by minimizing a weighted sum of squared price errors, with weights equal to the interest-rate sensitivity of prices so that the price criterion approximates a criterion on yields to maturity, an inverse-duration logic identical in substance to the one later canonized by Gürkaynak, Sack and Wright (2007) and employed in this thesis. Their estimation strategy is deliberately hierarchical: the four-parameter Nelson-Siegel model is fitted first, its solution is used to initialize the augmented Svensson form; they further note that the two additional parameters are valuable whenever the curve exhibits complex

intermediate and long-maturity shapes and that on thin cross-sections they are fragile and must be disciplined by careful initialization.

This hierarchy is also the cleanest way to state why Svensson is preferred to Nelson-Siegel in this thesis. The two models are nested: the second curvature term costs two parameters and buys the ability to accommodate a second hump in the forward curve, a shape that the single-hump specification structurally cannot reproduce. On the French sample this flexibility is not ornamental: the flat, inverted and double-humped configurations traversed between 2022 and 2026, like the crisis shape of 2008, are exactly the environments in which the Nelson-Siegel restriction binds. The choice is corroborated institutionally, since the Banque de France retains the Svensson specification for its official curve after testing it against the simpler alternative and empirically by [1], whose diagnostics this thesis reproduces and extends. What the pre-euro fragility documented in Chapter 2, with its masked early short end, adds to the argument is a caveat rather than an objection: where the cross-section is too thin to identify six parameters, the remedy adopted here, warm-starting each day from the previous solution with a multi-start fallback, is the modern descendant of the two-stage initialization of Ricart, Sicsic and Jondeau [10].

1.8 THE CASE AGAINST SVENSSON

The Svensson model is not, however, the consensus it once was and an honest defence of its adoption must begin from the reason for which several central banks abandoned it. The critique rests on two structural defects. The first is multicollinearity: the two curvature terms share nearly identical functional forms and when the two decay parameters drift towards each other their loadings become nearly collinear, so that very different parameter vectors produce almost indistinguishable curves; the estimated parameters can then jump erratically from one day to the next even when the fitted yields barely move, which is fatal whenever the parameters themselves, rather than the curve, are the object of interest. The second is the twin risk of local rigidity and overfitting: because the

basis functions are global, every observation influences the whole curve, so an idiosyncratic quote at one maturity deforms the fit everywhere, while on days with sparse or unbalanced cross-sections the six parameters are free to chase noise. Anderson and Sleath [11] articulate both points in explaining why the Bank of England replaced Svensson, which it had used for five years, with the spline-based method of Waggoner (1997), whose variable roughness penalty (VRP) lets the curve be flexible where data are dense, at the short end and increasingly smooth at long maturities, while confining the influence of each price to its own segment of the curve. The Bank of Canada followed the same route with different tooling; Bolder and Gusba [12] compare exponential, polynomial and Fourier bases and conclude in favour of the Merrill Lynch Exponential Spline (MLES) model, whose production implementation, nine linear parameters and a single non-linear one estimated by maximum likelihood on a universe that includes Treasury bills from three months of residual maturity, is described in the technical note of Bolder, Gusba and Strélski [13].

It is essential to notice what these adopters have in common: a mandate that begins at the very short end of the curve. The Bank of England explicitly motivates its change with the behaviour of the curve at the short maturities, to the point of importing general-collateral repo rates to pin down that segment and the Canadian MLES universe starts at a quarter of a year. Spline methods earn their advantage where the cross-section is dense and the required flexibility is local, which is precisely the money-market segment. This thesis, however, replicating [1], deliberately excludes that segment: BTFs and all residual maturities below twelve months are filtered out exactly because institutionally driven bill yields carry idiosyncratic variation that is informative about neither the medium- nor the long-term discount function. Once the short end is out of scope, the comparative advantage of the spline shrinks materially, while its costs for the present purpose remain intact.

Those costs are decisive for the object of this thesis. The MAE and noise diagnostics of Chapter 2 presuppose a rigid, parsimonious benchmark: a heavily parameterized spline absorbs the local richness and cheapness of individual bonds into the curve itself, mechanically shrinking the very residuals that constitute the measurement, so that a more flexible model would not measure market functioning better, it would erase the signal. A six-parameter curve with economically interpretable parameters, an asymptotic long rate, a short rate and two humps, is also what makes the 40-year parameter paths readable as monetary history rather than as fitting artefacts and what guarantees comparability: like-for-like with [1] on the overlapping window, with the official Banque de France curve and with the Svensson-based production curves of the Federal Reserve and the ECB. The practical defects, finally, are managed rather than ignored: this thesis targets fitted yields, which are robust to the parameter trade-offs that multicollinearity induces and stabilizes the daily estimation through warm-started trust-region least squares with a multi-start fallback, obtaining smooth parameter paths over roughly ten thousand days. The claim defended here is thus deliberately modest: not that Svensson dominates the alternatives in general, but that for measuring the medium-to-long French nominal curve and its fitting errors over four decades, on a like-for-like basis with the study being extended, it remains the preferable instrument.

1.9 DYNAMIC TERM STRUCTURE MODELS

A further and more radical departure from the static framework questions not how each day should be fitted, but whether single days should be fitted in isolation at all. Diebold and Li [14] reinterpret the Nelson-Siegel exponential components as a dynamic factor model, letting level, slope and curvature evolve as autoregressive processes: the resulting dynamic Nelson-Siegel (DNS) model turns the curve into a three-dimensional time series and produces term-structure forecasts that outperform standard benchmarks, especially at long horizons. Christensen, Diebold and Rudebusch [15] close its main theoretical gap by deriving the affine arbitrage-free class of Nelson-Siegel models (AFNS): imposing

absence of arbitrage on the DNS factor structure yields a model that retains the tractability of DNS while satisfying the restrictions of the canonical affine framework, at the cost of a deterministic yield-adjustment term and the class has since become a workhorse of term-structure research.

The Zero Lower Bound (ZLB) decade that followed the Global Financial Crisis (GFC) in 2008 forced a further evolution of this dynamic family, motivated by the same macroeconomic environment that the French sample traverses after 2014. Affine Gaussian models assign positive probability to arbitrarily negative yields; the shadow-rate construction of Black (1995), in which the observed rate is the maximum between zero and an unobserved shadow rate, and quadratic policy-rate specifications, are the two principal remedies. Andreasen and Meldrum [16] compare systematically on US data and find that the shadow-rate model best captures the risk-neutral dynamics near the bound, while Opschoor and van der Wel [17] develop a smooth shadow-rate DNS specification that improves both fit and forecasting at the zero lower bound. A separate, purely econometric warning comes from Bauer, Rudebusch and Wu [18]: because interest rates are highly persistent, maximum-likelihood estimates of the factor dynamics in dynamic term structure models suffer from severe small-sample bias, spuriously understating persistence and thereby distorting the decomposition of long yields into expected short rates and term premia, a bias for which the authors provide corrections.

None of this apparatus is adopted here, for a reason that is methodological rather than dismissive. Dynamic models answer dynamic questions: forecasting, the decomposition of long rates into expectations and term premia, the persistence of shocks. The object of this thesis is different, namely the daily cross-sectional measurement of how tightly the market prices a single curve and for that object the statistical independence of each day's fit is a feature: imposing time-series structure on the parameters would smooth across precisely the daily dislocations, in March 2020 above all, that the diagnostics are designed to detect. Nor is a shadow-rate floor appropriate for a market whose yields genuinely and durably

crossed zero, as the French curve did for nearly eight years: the deliberately unconstrained short-end parametrization recalled in Section 2.2 handles the negative-rate regime without auxiliary machinery. The dynamic literature remains, however, the natural continuation of this work: estimating an AFNS or a bias-corrected dynamic model on the extended 1987–2026 French dataset constructed here and confronting its factor dynamics with the Principal Component Analysis (PCA) of Chapter 2, is an obvious avenue for future research.

1.10 MACHINE-LEARNING FRONTIERS

The most recent strand of the literature asks whether the functional form should be specified at all or rather learned from the data. Three contributions illustrate the range of possibilities. Gümbel and Schmidt [19] apply Gaussian process regression (GPR), a non-parametric Bayesian technique closely related to Kalman filtering, to the calibration of bond prices within a Gaussian affine framework: calibration proves fast and accurate in single-curve markets, though the authors document substantial difficulties in the multi-curve setting, where the small-data nature of term-structure problems prevents the method from exploiting its usual advantages. Akiyama and Matsuyama [20] address yield-curve extrapolation beyond the last observable tenor, a problem made pressing by the market-consistent valuation requirements of the major EU Solvency II reform, which principally overhauls capital requirements, reduces reporting burdens by up to 35% for smaller firms and introduces macroprudential tools, with a long short-term memory (LSTM) neural network, and show on US and Japanese data that it extrapolates with better generalization, stability and consistency with observed curve dynamics than the Nelson-Siegel and Svensson families; notably, they confront the black-box objection directly, using the backpropagation algorithm itself to recover the sensitivity of outputs to inputs and hence a measure of interpretability. Kauffmann, Takada, Terada and Stern [21] pursue a hybrid route, embedding a neural network that generates smooth factor-loading functions inside a Gaussian linear state-space model, so that an arbitrary number of factors and their loadings are learned jointly

by Kalman filtering and automatic differentiation, with shrinkage priors guarding against overfitting; on fourteen years of Brazilian data the learned decompositions forecast competitively against fixed-loading benchmarks.

The appeal of these methods is real: loadings learned rather than imposed, nonlinearities captured, extrapolation stabilized, calibration accelerated. Their limitations, however, at the current state of the literature, are equally concrete. They are data-hungry in a field whose daily cross-sections are small, as Gmbel and Schmidt [19] themselves find; their protection against overfitting is statistical regularization chosen by the researcher rather than an economically grounded functional form; their interpretability, even where actively engineered as in Akiyama and Matsuyama [20], remains partial compared with parameters that are a long rate, a short rate and two humps; and none of them has yet been adopted for production curve estimation by a central bank or debt-management office, so that their robustness across regimes, the property that the present 40-year sample tests most severely, is still undemonstrated. For a study whose contribution rests on strict comparability with a published benchmark across four decades, adopting an experimental estimator would confound methodological novelty with the economic signal of interest. These methods are best read, therefore, as a promising complement to, and a future robustness check on, the parametric apparatus this chapter has defended: an apparatus chosen not by default, but against explicit alternatives, on the grounds that for this market, this maturity range and this measurement purpose, the Svensson curve of [1] remains the right instrument.

2. METHODOLOGY

2.1 FROM NELSON-SIEGEL TO SVENSSON

The empirical term-structure literature estimates the yield curve not point by point but by fitting a small, smooth, parametric function to the cross-section of observed bond prices. The main framework is the Nelson and Siegel (1987) family, which models the instantaneous forward rate as the solution of a second-order differential equation yielding a curve flexible enough to be monotonic, humped or S-shaped with only four parameters. Svensson (1994) augmented this specification with a second curvature parameter and a second decay parameter, allowing the forward curve to accommodate two humps and thereby fit the more complex shapes observed at intermediate and long maturities. Gürkaynak, Sack and Wright (2007) (GSW) operationalized the Svensson form into the daily production methodology now standard at the main central banks, such as the Federal Reserve and the European Central Bank, fitting the six parameters each day to the cross-section of treasury prices, weighting by the inverse of duration and publishing the resulting zero-coupon curve. The [1] paper transplants this GSW production methodology onto French data; the choice is not arbitrary: the Banque de France itself estimates its official curve with the Svensson model, which is retained after testing against the Nelson-Siegel alternative. This thesis reproduces this exact methodology and then carries it forward to June 5th, 2026.

2.2 THE PARAMETRIC CURVE

With parameter vector $\theta = \beta_0, \beta_1, \beta_2, \beta_3, \tau_1, \tau_2$, the instantaneous forward rate at horizon m is: $f(m) = \beta_0 + \beta_1 e^{-\frac{m}{\tau_1}} + \beta_2 \frac{m}{\tau_1} e^{-\frac{m}{\tau_1}} + \beta_3 \frac{m}{\tau_2} e^{-\frac{m}{\tau_2}}$.

Integrating over $[0, m]$ we obtain the continuously-compounded zero-coupon yield:

$$f(m) = \beta_0 + \beta_1 G\left(\frac{m}{\tau_1}\right) + \beta_2 \left[G\left(\frac{m}{\tau_1}\right) - e^{-\frac{m}{\tau_1}}\right] + \beta_3 \left[G\left(\frac{m}{\tau_2}\right) - e^{-\frac{m}{\tau_2}}\right],$$

$$G(x) = \frac{1 - e^{-x}}{x}.$$

The parameters carry economic meaning: β_0 is the asymptotic long-horizon level, to which the forward curve converges as $m \rightarrow \infty$; $\beta_0 + \beta_1$ is the instantaneous short rate; β_2 and β_3 govern the magnitude of the first and second humps, located through the decay constants τ_1 and τ_2 . The zero curve is the discount-rate object used to price every bond, while the forward curve disentangles maturity-specific information, with short-horizon forwards reflecting cyclical and monetary policy factors, long-horizon forwards reflecting more fundamental risk attitudes.

For each trading day t the six parameters are estimated by repricing the whole eligible cross-section off the model curve. A coupon bond is priced by no-arbitrage as the sum of its cash flows discounted on the Svensson zero curve, the observed price being the dirty bid price. Following GSW the loss function is the duration-weighted sum of squared price errors: $\hat{\theta}_t = \arg \min_{\theta} \sum_{k=1}^{N_t} \left[\frac{p_k^{obs} - p_k^{model}(\theta)}{D_k} \right]^2$, where the weight is the inverse of each bond's modified duration D_k . Weighting by $1/D_k$ converts price errors into approximately yield errors, as a long-duration bond price moves more for a given yield error, so the fit is not dominated by the longest bonds; GSW adopt it precisely to avoid converting prices to yields at each iteration, which speeds the calibration.

Two estimation choices matter for the extended sample and are inherited from the literature. First, the constraints $\tau_1, \tau_2, \beta_0 > 0$ are imposed, but $\beta_0 + \beta_1$ is left unconstrained, so the fitted short rate is allowed to go negative, which is indispensable once French yields turn negative after 2014. Second, each day is solved by trust-region least squares warm-started from the previous day's solution, as in GSW, with a small multi-start grid invoked only when the warm-started fit is poor; this yields smooth parameter paths and keeps the full $\sim 10,000$ -day run to a few minutes.

2.3 DATASET: SOURCES AND CONSTRUCTION

The estimation rests on two Bloomberg exports for French government securities: a security master built up on one row per ISIN, carrying issuer name, coupon, issue and maturity dates, the maturity type field, series and currency of issuance; a wide bid-ask panel built up on two rows per bond (Ask, then Bid) against 10,077 daily date columns running from October 22nd, 1987, to June 5th, 2026. The two files are merged into a per-ISIN reference table and a date times ISIN matrix of clean bid prices, on which the static filters below operate.

The panel stores both sides of the market, using bid-ask spread as a liquidity descriptor. Computed across the whole panel, the cross-sectional daily mean spread averages ≈ 0.076 price points, meaning ≈ 7.6 euro-cents per 100€ of face, over the euro window, essentially identical to the value reported by the paper, “around 7 euro cents” with a comparable ≈ 8 cents in the pre-euro period. In the post-2018 extension the median spread widens only modestly to ≈ 11.5 cents, while the mean is inflated by a fat tail of wide quotes on illiquid and ultra-long lines; the gap between median and mean is itself informative about the heterogeneity of liquidity across the modern curve and motivates the robust screens of Section 2.4.

The parsed master contains 173 fixed-coupon OAT/BTAN lines; after the static filters and the requirements of usable price data, the estimation universe is made up of 155 securities, 152 of which are OATs and 3 legacy BTANs. Their characteristics span the entire monetary history of modern France: i. currency of issuance: of the 155 fitted bonds, 147 are euro-denominated (EUR), 7 are French-franc (FRF) and one is denominated in European Currency Units (XEU), the last two categories being pre-1999 vintages, while the euro lines account for everything issued since January 1st, 1999. This mirrors the paper’s statement that pre-euro fixed-coupon bonds were denominated in FRF and, in a few cases, in XEU. ii. Coupon and maturities: coupons range from 0% in lines under the negative-rate

regime to 11.60% in the early franc-era issues. Residual maturities extend to May 25th, 2072, and the original terms range from 2.6 to 51 years. The average life of the negotiable debt is 8 years and 200 days. iii. Tenor distribution and the cross-section per day: issuance is concentrated in the medium-to-long segment; the number of bonds quoted on a given day rises from a handful in the late 1980s to over a hundred late in the sample, between 6 and 155 per day, with a mean of ~33, though always above the six observations needed to identify the six Svensson parameters. The maturity-bin scheme used for diagnostics is 0–2, 2–5, 5–10, 10–20, 20–30 and 30–50 years.

The original paper fitted 179 nominal bonds over the 1987–2018 period; the present dataset holds 173 lines over 1987–2026. The count is not simply the paper’s count plus eight more years of issuance and the reason is a genuine data-availability fact that must be stated honestly. The current Bloomberg extract no longer retrieves some of the oldest franc era lines that had already matured decades ago and have dropped out of the accessible history; at the same time, it has been augmented with every nominal OAT issued through June 2026. The net effect is a slightly thinner early cross-section than the paper, whose daily minimum number of quotes is 7 against the 6 of this dataset, and a much richer recent one. There are two main visible examples of the extension: the first is the ultra-long segment, where the paper observed three 50-year bonds, issued in 2005, 2010 and 2016, the updated data contain a fourth, issued in January 2021 and maturing in 2072, as presented in Figure 1, and Figure 2 shows that the overall fitting errors for the initial segment are greater than the paper’s owing to the absence of some bonds now expired.

2.4 FILTERS APPLICATION

The five [1] filters are implemented as follows: i. regular fixed-coupon bullets only. Special feature securities are dropped; crucially inflation-linked bonds (OATi and OAT€i) and foreign-currency bonds (USD, GBP) are excluded, as are floating-

rate OATs and STRIPS, whose cash flows are not priced off the nominal euro discount function and the paper removes them for exactly this reason. The currency screen admits only EUR, FRF and XEU. ii. No BTFs: short-term discount bills are excluded, following the paper's caution about idiosyncratic bill-yield variation; they are in any case absent from this extract. iii. Always above 12 months to maturity, applied dynamically each day, to keep institutional short-end effects out of the fit. iv. On-the-run bonds are kept: unlike GSW for the US, the most recent and first-off-the-run bonds are not excluded, because the French on-the-run premium is negligible. v. Abnormal-quote screen: the two hand-checked ISINs flagged in the paper are blanked on the exact dates listed. Extending to 2026 required augmenting this: the new years contain gross Bloomberg quote errors the paper never dealt with, such as a frozen ~ 40 price on two 2021–2022 lines and $\sim 1/10$ -scaled 1988 franc quotes, implying yields of 24–31%. The extension to 2026, however, exposed the estimation to a class of data defects that the original sample never contained and dealing with them required two additional screens that constitute one of the genuine methodological contributions of this thesis: a static plausibility filter on the implied yield-to-maturity and a robust per-day outlier rejection. Since these screens are additions to and not replications of the protocol of [1], their design is justified in full.

The governing principle is that a data-quality screen must never be allowed to shade into an economic filter. The two kinds of rule have opposite loss functions: an economic filter is designed to bind on genuine observations, whereas a data screen must be constructed so that it cannot, under any regime the market might plausibly visit, touch a legitimate quote. The asymmetry follows from the object of this thesis: the MAE and the noise diagnostics of Section 2.6 treat the dispersion of observed yields around the fitted curve as the signal, so that deleting a genuine but extreme observation would not merely lose a data point, it would mechanically suppress the very stress episodes the diagnostics exist to measure. A screen that could conceivably bind on real data would therefore contaminate the measurement;

a screen that errs on the side of leniency costs nothing, because errors that survive the static band are precisely what the second, per-day robust screen is designed to catch. The two tiers are complementary rather than redundant: gross scale errors, if left in the sample, are large enough to poison the daily robust statistics themselves, so the static band exists first and foremost to protect the breakdown point of the adaptive screen that follows it. In adopting explicit, codified thresholds this thesis also departs deliberately from the practice, acknowledged for instance in the Bank of Canada documentation of [13], of handling problematic quotes by discretionary, judgement-based exclusion after case-by-case investigation: judgement does not scale to a ten-thousand-day panel and, more importantly, it is not reproducible; a fixed band is.

The specific values of the band, -5% and $+20\%$, are then anchored from both sides rather than chosen by taste. An audit of the full panel, in which the implied yield-to-maturity of all 338,080 eligible quotes was computed with the plausibility band switched off, shows that genuine French yields form a compact cluster reaching its maximum at $+11.53\%$, in the high-coupon franc era, while the erroneous quotes begin at $+21.33\%$ and drift upward to beyond $+36\%$. The drift is itself diagnostic of the error mechanism: a price frozen on a line whose residual maturity keeps shortening implies a mechanically rising yield, so each frozen-price episode enters the sample from below, at the bottom of its eventual range. The admissible interval for the upper threshold is therefore $[11.5\%, 21.3\%]$: any cut inside it removes exactly the erroneous quotes and nothing else and $+20\%$ is placed near its upper end to maximize the headroom over genuine data. Consider now the lower bound: the deepest yield ever fitted on French data is the $\approx -0.84\%$ touched by the 2-year rate in 2019, at the trough of the negative-rate regime and below the ECB deposit facility rate of -0.50% ; no major sovereign market has ever recorded yields materially below -1% and the arbitrage against holding physical currency, whose cost of carry is measured in fractions of a percent per annum, makes yields several points below zero economically unattainable rather than merely

unobserved. A threshold at -5% therefore sits more than five times deeper than the most negative genuine French yield in four decades, in territory that the cash-arbitrage bound renders unreachable, while the error modes it targets, decimal-shifts and scaling errors that place prices far above par, generate implied yields an order of magnitude beyond it.

Within the admissible region, the exact placement of the threshold is deliberately immaterial and this insensitivity is itself part of the design. The baseline band excludes 1,229 quotes, barely 0.36% of the panel, all of them above the upper cut; re-running the classification with alternative bands, namely $[-3\%, +15\%]$ and $[-8\%, +15\%]$, excludes the identical set, quote for quote. The round figures adopted here should therefore be read as a signal of robustness, not as the product of a tuning exercise: the results of the thesis are invariant to any band drawn through the empty region that separates the genuine cluster from the gross errors. A final consideration concerns forward validity: because the pipeline is designed to be re-run as new data accrue, the thresholds were also required not to bind in any regime the future might plausibly deliver and bounds that comfortably contain both the negative-rate trough and the franc-era peaks of the past forty years are the natural operationalization of that requirement.

The per-day robust screen that operates inside the band follows the same philosophy at finer resolution and its calibration is therefore stated exactly. After each daily fit, the absolute yield error of every bond, observed minus fitted yield-to-maturity, in basis points, is compared with the following: $\text{threshold} = \max(75 \text{ bp}, \text{med} + 6 \times 1.4826 \times \text{MAD})^2$, where the median and the median absolute deviation (MAD) are taken across that day's cross-section and the factor 1.4826 rescales the MAD into a consistent estimator of the standard

² For a Gaussian variable the median absolute deviation satisfies $\text{MAD} = \phi^{-1}(0.75) \cdot \sigma \approx 0.6745\sigma$, so that $1.4826 \times \text{MAD}$, with $1.4826 = 1/0.6745$ is a consistent, outlier-resistant estimator of the standard deviation. The rule $\text{median} + k \cdot 1.4826 \cdot \text{MAD}$ is the classical Hampel identifier for outlier detection; its conventional multiplier $k = 3$ is doubled here to $k = 6$ so that, under normality, a genuine observation breaches the threshold with probability of the order of 10^{-9} , ensuring the screen cannot bind on legitimate quotes.

deviation under normality; bonds beyond the threshold are treated as abnormal quotes, removed and the curve is re-estimated, with at most two such passes so that the pruning cannot cascade. Each component serves the same non-binding requirement stated above. The median and the MAD, unlike the mean and the standard deviation, are essentially unaffected by the very outliers being hunted, so the threshold cannot be dragged upwards by the contamination it is meant to detect. Six robust standard deviations is a deliberately extreme distance: across the entire sample of observations the expected number of legitimate quotes falsely rejected is far below one. The adaptive term makes the rule regime-aware: in the noisy pre-euro years, when the median fitting error alone reached $\approx 18\text{--}20$ bp, the MAD inflates the threshold well above any fixed level that would otherwise have amputated genuine dispersion. The 75 bp floor governs the opposite regime: in the ultra-calm euro area the daily MAD shrinks to a few basis points and a purely adaptive threshold would descend towards levels that genuine stress dislocations can reach; the floor guarantees that any deviation below 75 bp, however anomalous it may look on a quiet day, is always measured rather than deleted. The ≈ 60 bp cheapening of the ultra-long visible in the June 2026 snapshot of Figure 5 survives the screen by construction, which is exactly as intended.

2.5 EXTENDED ESTIMATION

The original sample stops on April 10th, 2018, while the present thesis re-runs the identical Svensson and GSW pipeline, day by day, all the way to June 5th, 2026, for a total of 10,077 trading days, of which 2,129 ($\sim 21.13\%$ of the total sample) lie beyond the paper's horizon. This is the methodological core of the thesis: documenting, on a like-for-like basis, how the fitted curve and its diagnostics behave over the eight years that the paper could not observe, filled with all the macroeconomic, political, monetary, pandemic and geopolitical shocks.

The extension is also computationally non-trivial and the architecture that makes it tractable is part of the contribution. Each trading day poses a six-

dimensional bound-constrained nonlinear weighted least-squares problem in which every evaluation of the objective reprices the entire eligible cross-section, discounting every future cash flow of every bond off the candidate curve. With an average of roughly ~ 33 bonds quoted per day and up to 600 objective evaluations allowed per solve, the full 10,077-day run entails well over 10,000 trust-region optimizations and, behind them, tens of millions of individual bond repricings. Two implementation choices keep this affordable. First, the cash-flow schedules of each day's cross-section are flattened into a single array, so that the model prices of all bonds are computed simultaneously in vectorized form rather than bond by bond. Second, the warm start inherited from the previous trading day, described in Section 2.2, places the solver so close to the solution on most days that convergence requires only a handful of iterations. The run is inherently sequential, since each day's starting point is the previous day's solution, yet it completes in the few minutes already noted in Section 2.2 on ordinary hardware. This economy has methodological, not merely practical, significance: because a full forty-year re-estimation is cheap, every robustness exercise reported in this thesis is a complete re-run of the entire sample under the alternative configuration, never a spot check on selected dates.

Convergence over $\sim 10,000$ heterogeneous days cannot be taken for granted, which is why three distinct failure modes are addressed. The first is the local-minimum risk inherent in the Svensson objective: as discussed in Section 1.8, when the two decay parameters drift towards each other the curvature loadings become nearly collinear and the objective develops long, flat valleys along which very different parameter vectors price the cross-section almost identically, so that a solver may stall far from the best solution. The second is specific to warm-starting: an undetected poor solution would not remain confined to its own day, but would seed the next day's search and could propagate along the chain. Both risks are policed by the convergence tripwire described in Section 2.6.3: whenever the warm-started fit is poor by the dirty-price RMSE criterion, the day is re-solved

from the multi-start grid and the lowest-cost solution retained, which breaks any incipient chain of contamination at the first link. The third mode is near-boundary identification: on the thinnest early days the cross-section exceeds the six-parameter minimum by only one or two securities and their maturities cluster around the 10-year point, so that portions of the curve are formally identified but practically unconstrained by the data. Deliberately, no attempt is made to repair this by tightening the measurement; instead, the affected tenors are flagged as extrapolations beyond the traded maturity range and masked in all downstream analysis, as described in Section 2.8.2.

Finally, parameter stability is in this design an outcome to be verified rather than an assumption to be imposed. No intertemporal smoothing links one day's parameters to the next. Beyond the purely numerical device of the warm start, each day is estimated in statistical isolation, the methodological choice defended in Section 1.9. Whatever smoothness the resulting parameter paths display is therefore evidence of a well-posed estimation, not an artefact of construction and the paths do behave exactly as this reading requires: they evolve slowly through calm regimes and move sharply precisely at the known breaks, the 2008 crisis, the descent below zero after 2014, the 2022 tightening. Stability is assessed in yield space instead of in parameter space: as noted in Section 1.8, the near-collinearity of the specification permits occasional offsetting movements among the curvature and decay parameters that leave the fitted curve essentially unchanged and such episodes are harmless here because every result of this thesis is a functional of fitted yields, never of an individual coefficient. The optimizer box constraints are set generously, wide enough to contain any economically conceivable configuration, so that the reported solutions are determined by the data and not by the box.

All these safeguards are complemented by a set of robustness checks described in various sections of this chapter, namely: the validation of the extended pipeline against the published results on the overlapping window with matching snapshots,

error statistics and factor structure in Section 2.8; the invariance audit of the additional data-quality screens and the verification that alternative plausibility bands classify the panel identically in Section 2.4; the daily re-estimation of the curve without the on-the-run and first-off-the-run securities, which underlies the on-the-run premium results in Section 2.9; the continuous convergence monitoring provided by the dirty-price RMSE in Section 2.6.3. Two further sensitivities were re-estimated in full over the entire 1987–2026 sample, at the trivial cost that the runtime economics described above make possible. Raising the minimum-maturity filter from twelve to eighteen months, the variation that [1] reports to be immaterial on the original sample, proves equally immaterial on the extended one: the euro-era mean MAE moves from 2.39 to 2.28 bp and the mean noise from 3.39 to 3.17, both marginally lower as the stricter screen removes the noisiest short bonds, while the post-2018 diagnostics are unchanged to within 0.05 bp. In yield space, the fitted 10-year zero rates of the two runs differ by a median of 0.3 bp across the common sample and the only appreciable divergences, reaching about 37 bp on a handful of days in March 1995, are confined to the thin franc-era cross-section, where removing the 12- to 18-month segment materially alters the information available to the short end: precisely the region already flagged as fragile and masked in Section 2.8.2. The second check, excluding the 17 retail France-Physiques tranches so as to re-estimate on the wholesale cross-section alone, resolves into an identity rather than a sensitivity: the retail lines, though present in the security master, carry no usable quotes in the Bloomberg extract and therefore never enter the estimation universe, so the wholesale and baseline runs coincide day for day. The estimated cross-section of this thesis is thus wholesale by construction, a fact worth recording because it turns a potential robustness concern into a property of the data. Every modelling choice of the estimation, from the day-count convention to the thresholds, lives in a single configuration file, so each of the checks above amounts to editing one line and re-running the forty-year sample: the pipeline is, by construction, an auditable object and this auditability is as much a part of the extension contribution as the eight new years of estimates themselves.

2.6 GOODNESS OF FIT METRICS

Three daily diagnostics, all comparing observed against fitted yields-to-maturity, with a bond's fitted YTM being the yield that reprices its model dirty price, are computed: i. the overall mean absolute error (MAE) and its within-bin counterpart, as shown in Figures 2 and 3; ii. the Hu-Pan-Wang (2013) (HPW) "noise" measure, the RMS of the daily yield errors, read here as an indicator related to the availability of arbitrage capital, an interpretation whose conditions are made explicit below; iii. the unweighted dirty price RMSE, an internal convergence check.

Before the three measures are presented individually, a caveat on their interpretation must be stated clearly, because it applies to all of them. A fitting error is the distance between one bond price and its counterpart produced by the parsimonious model; such distance is the joint product of several mechanisms, of which the withdrawal of arbitrage capital is only one. At least four competing sources deserve explicit acknowledgment. First, model approximation error: a six-parameter curve cannot reproduce every admissible configuration of several dozen yields and part of any day's error is the structural residual of the functional form itself, the very limitation on which the spline critique of [11] rests; this component is largest where curvature demands are extreme, as at the ultra-long end, where convexity makes even genuinely priced bonds deviate systematically from what a smooth parametric curve implies [22]. Second, data quality: the panel consists of indicative dealer quotes instead of transaction prices and although the screens of Section 2.4 remove gross errors, stale or carelessly refreshed quotes below those thresholds survive by design. Third, issuance concentration and liquidity segmentation: the cross-section is not homogeneous, since the souche system concentrates outstanding amounts in benchmark lines while leaving other segments thin and lumpy, futures deliverability privileges specific bonds [3], regulatory reporting dates distort quotes at predictable calendar points [8] and clientele effects of the buy-and-hold kind described in Section 1.5 can hold

individual bonds away from the curve persistently, the systematic liquidity premia modelled by Monfort and Renne [6] and documented empirically by Beber, Brandt and Kavajecz [7]. Only after these are accounted for does the fourth mechanism, the availability of arbitrage capital, carry the residual explanatory burden and the evidence of Duffie, Fleming, Keane, Nelson, Shachar and Van Tassel [9] shows how demanding this attribution is even in the deepest market in the world: most of the variation in the US Treasury market is explained by yield volatility alone and intermediation capacity becomes the binding explanation only in episodes of extreme dealer balance-sheet utilization. Throughout this chapter, therefore, the arbitrage-capital reading inherited from Hu, Pan and Wang (2013) [51] and [1] is treated as an interpretation to be earned case by case, through the coincidence of error spikes with independently documented stress and through their cross-sectional signature, not as a fact established by the diagnostics themselves.


2.6.1 MEAN ABSOLUTE ERROR

Once the six Svensson parameters have been estimated for a given day, a single scalar is needed to judge how faithfully that smooth, six-parameter curve reprices the entire cross-section of bonds observed on that day. Following [1], this thesis adopts the mean absolute error (MAE) of the fit, expressed in basis points of yield.

For each security the model dirty price implied by the estimated curve is inverted back into its yield-to-maturity and the fitting error is the difference between the observed and the fitted YTM. The overall daily MAE averages the absolute value of these errors across all N bonds quoted that day:

$$\text{MAE}_t = \frac{1}{N_{b,t}} \sum_{i=1}^{N_{b,t}} |y^{\text{obs}}(i) - y^{\text{fit}}(i)| . \text{ Results are reported in Figure 2.}$$


Figure 2



Source: author's elaboration in Python

The same statistic, restricted to the securities whose residual maturity falls in a given bin, defines the maturity-specific error, so that the overall figure is simply the bond-count-weighted average of the six bin errors, as shown in Figure 3. The errors are deliberately measured on yields, not prices, so that they are directly comparable across securities of different coupon and maturity.

Figure 3



Source: author's elaboration in Python

The MAE answers a precise question, namely, on average, by how many basis points does each traded bond deviate from the single arbitrage-free curve that best summarizes the market that day. Two properties make it the natural metric. First, being an absolute rather than a squared loss, it is robust: a handful of stale or erroneous quotes inflate it far less than they would a root-mean-square error, a material advantage over a 39-year sample that inevitably contains occasional data glitches. Second, the MAE lends itself to an economic reading as a diagnostic of market functioning, provided the reading is stated with the care that Section 2.6 requires. In a frictionless market with abundant arbitrage capital every bond must lie tightly on one smooth discount curve, so a small MAE is genuinely informative: it certifies that no mechanism, frictional or otherwise, is holding prices apart. A large MAE is more ambiguous, because dispersion can widen for any of the

reasons catalogued above, of which scarce arbitrage capital is only one candidate. What strengthens the economic reading is not the level of the MAE, which embeds the time-invariant approximation error of the functional form, but its variation: when the errors of the same bonds, priced off the same model and filtered by the same screens, widen abruptly and revert as documented stress arrives and recedes, the mechanisms that are constant by construction, the functional form, the quoting convention, the issuance structure, cancel out of the comparison and the time-varying candidates, volatility and arbitrage capacity foremost, remain. The per-bin decomposition adds spatial resolution, localizing misfit along the curve and revealing whether stress concentrates at the short end, suffering from institutional and monetary policy effects, in the belly or in the thin, illiquid long end, where model approximation and scarce issuance overlap and attribution is correspondingly harder.

Figure 2 confirms the picture of [1] on the overlapping window and extends it: the overall MAE settles at roughly 2–3 bp through the calm 2004–2007 and 2013–2019 phase and spikes at every episode conventionally associated with impaired market functioning, as shown by the early-euro learning period, up to ~17 bp, the 2008–2009 GFC and the 2011–2012 sovereign crisis and then, beyond the paper’s horizon, a sharp COVID spike in March 2020 and a sustained elevation to ~9 bp through the 2022–2024 inflation and tightening regime, before easing to ~3.5–4 bp by 2026.

Figure 3 shows that this recent deterioration is not uniform across the curve: while the short and intermediate bins remain well-behaved, the 30–50-year bin blows out from ~2–3 bp to ~15–23 bp after 2022. This is the segment where the competing explanations of Section 2.6 are hardest to separate: a handful of lines, lumpy issuance, extreme convexity demands on the functional form and a genuinely violent repricing all coincide, so the blow-out is best read as the joint failure of a thin market and a parsimonious model to agree on the ultra-long end,


rather than as evidence of arbitrage-capital scarcity specifically. In either reading, it is the honest caveat that qualifies the otherwise excellent post-2018 fit.

2.6.2 HPW NOISE MEASURE

The Hu-Pan-Wang “noise” measure is here adopted as the second diagnostic and read, in keeping with its origin and [1], as an indicator related to the quality of market functioning and, more specifically, to the availability of arbitrage capital, subject to the qualifications set out in Section 2.6. It is computed each day as the root-mean-square of the yield fitting errors across all quoted OAT/BTAN securities, in basis points: $\text{NOISE}_t = \sqrt{\frac{1}{N_{b,t}} \sum_{i=1}^{N_{b,t}} (y^{\text{obs}}(i) - y^{\text{fit}}(i))^2}$. It is thus the RMS twin of the MAE, built from the very same observed minus fitted yield errors, but squared rather than absolute, so it weights large deviations more heavily and is the most sensitive gauge of cross-sectional dispersion.

The economic logic runs cleanly in one direction: when arbitrage capital is abundant, relative-value traders keep every bond glued to a single smooth curve, so the errors, and hence the noise, are small. The converse inference, from high noise to scarce arbitrage capital, is weaker and must be drawn with care. Duffie, Fleming, Keane, Nelson, Shachar and Van Tassel [9] show on US data that most of the variation in market liquidity is explained by yield volatility alone, with balance-sheet capacity becoming the binding constraint only in exceptional episodes; a noise spike in a violently repricing week may therefore reflect the speed at which new information is being incorporated across a heterogeneous cross-section in place of an inability to arbitrage and the ICMA [8] evidence adds that part of the measured widening at predictable calendar dates is regulatory as opposed to frictional in origin. In this thesis, accordingly, high noise is read as consistent with mispricing and impaired market functioning, but is interpreted as such only where the timing, persistence and cross-sectional signature of the spike line up with independently documented stress, as in the case studies examined later in the following chapters.

Figure 4



Source: author's elaboration in Python

Figure 4 plots the measure over the full 1987–2026 sample. The single most striking feature is the regime break at the euro launch: the series is high and violently volatile before 1999 and low and stable afterwards. In the pre-euro window the noise averages ≈ 20 bp, with a median of ≈ 18 bp, and spikes as high as ≈ 84 bp in 1988; at the introduction of the euro it collapses almost instantaneously to the ≈ 3 bp average of the 1999–2018 benchmark, rising above that floor only during genuine stress, namely the early-euro adjustment of 2001–2002 at ~ 18 bp, the GFC and 2011–2012 sovereign crisis yet modest, ~ 5 – 6 bp and, in the extension, COVID-19 at ~ 11 bp in 2020 and the 2022–2023 inflation shock, ~ 16 bp in March 2023. It ends the sample at ~ 9 bp.

The economic heart of Figure 4 is the reason why the pre-euro period is so volatile and why it then stabilizes, which is fundamental to understanding the arbitrage-capital logic. The pre-euro French market was thin and immature: far

fewer bonds were quoted each day, with a mean of ~ 22 versus ~ 33 and more in the euro era, and in the late 1980s their maturities clustered around the 10-year point, so the cross-section barely identified a six-parameter curve and any idiosyncratic quote translated into large yield dispersion. This is the identification component of the caveat at work: a large part of the pre-euro noise level reflects what the cross-section could not pin down and would inflate the measure even in a market with no arbitrage friction at all. More fundamentally, the investor base was domestic and buy-and-hold, with little incentive to trade or to arbitrage relative-value discrepancies, so mispricings persisted; the paper [1] also notes that “arbitrage opportunities were not infrequent” in the first decade. Layered on top were franc-denominated currency and political risk, like the 1992–1993 ERM crises, EMU-convergence uncertainty, and a still-maturing market infrastructure. The euro removed the currency risk, integrated the EMU bond markets and drew in a large pool of outside investors and arbitrage capital; the paper [1] attributes the improvement in fit to the influx of European investors. With ample arbitrage capital, bonds are priced tightly to one curve and the noise collapses, spiking thereafter only when that capital is temporarily withdrawn in a crisis.

The qualitative result is reproduced exactly, with a high, volatile pre-euro noise collapsing to a low, stable euro-era level. The magnitudes also broadly agree in the euro era, ~ 3 bp mean of the model against the “around or below 5 bp” of the paper [1]. However, differences are worth signalling. First, the replication finds higher pre-euro peaks, up to ~ 84 bp in 1988, while the paper never exceeded 35 bp; this is attributable to the thinner earliest-sample cross-section in the current Bloomberg extract, due to fewer surviving 1987–1988 lines, which makes the curve harder to identify on those specific days, whereas the paper [1] figure effectively begins around 1990. Second, and this is the extension contribution, the post-2018 window reopens a modest gap: the noise mean roughly doubles to ~ 7 bp relative to the 2004–2019 lows, with a 2023 peak near 16 bp, still an order of magnitude below the pre-euro chaos, but a visible sign that the higher-rate, higher volatility and

ultra-long-supply regime of 2022–2026 has made the French market marginally less frictionless than at its 2004–2007 and 2013–2019 best. This partial re-elevation of the noise is precisely the market-functioning signal that Case Study 3 examines at daily resolution around the March 2020 dislocation and the ECB PEPP response.

2.6.3 DIRTY PRICE RMSE

The third quantity monitored on every day is the unweighted root-mean-squared of the dirty price errors, computed in price points as:

$$\text{RMSE}_t = \sqrt{\frac{1}{N_{b,t}} \sum_{i=1}^{N_{b,t}} (P_i^{\text{obs}} - P_i^{\text{fit}})^2}.$$

Unlike the MAE and the HPW noise, which live in yield space, in basis points and carry an economic interpretation, namely pricing accuracy and arbitrage-capital availability, the price RMSE lives in the model native price space and serves a purely numerical purpose: it is the raw unweighted counterpart of the very objective the optimizer minimizes, namely the duration-weighted sum of squared price errors. It therefore measures how tightly the Svensson curve reprices the bonds before the inverse-duration weights are applied.

Its role is that of an internal convergence check on the daily calibration. Because each day is warm-started from the previous day’s solution, as in GSW, there is a small risk that the trust-region solver settles into a poor local minimum. A well-converged day sits at only ~0.3–0.8 price points, which translates into a few basis points in yield; if the warm-started fit RMSE instead exceeds a threshold of 2 price points, the day is treated as a probable convergence failure and the optimizer is re-run from a small multi-start grid, keeping the lowest-cost solution. In this way the metric never enters the economic analysis, so it is not plotted or tabulated as a result, but it silently guarantees the smooth, well-identified parameter paths on which every figure and table of this chapter rests. The check also plays a quiet part in the interpretive discipline of Section 2.6: among the competing explanations for a spike in the yield-space diagnostics, the purely numerical one, a solver settling

on a poor local minimum, is the only candidate that can be eliminated outright rather than argued against and this is the check that eliminates it. Whatever survives in the MAE and noise series is therefore a property of the data and of the model, never of the optimizer.

2.7 FIT OF THE NEW PERIOD

The headline result is that the methodology continues to fit the French curve very well after 2018, with a modest, economically interpretable deterioration concentrated at the long end.

Table 2

Sample	MAE (bp)			Mean Noise (bp)	N days
	Mean	Std	Max		
Full (1987–2018, paper)	7,00	9,12	53,48	-	7.949
Euro (1999–2018, paper)	2,39	1,98	17,13	3,39	5.028
Pre-euro (1987–1998)	14,93	10,96	53,48	-	2.921
Post-paper (2018–today)	3,58	1,59	9,24	6,92	2.129
Euro extended (1999–today)	2,75	1,95	17,13	4,44	7.156
Full extended (1987–today)	6,28	8,25	53,48	-	10.077

Source: author’s elaboration in Python

As can be seen from Table 2, the post-2018 mean MAE of 3.58 bp is only ~1.2 bp above the paper-window figure and, significantly, its maximum of 9.24 bp is lower than the paper window’s, which settled at 17.13 bp. There is no single post-2018 day on which the fit breaks down as badly as the worst days of the 2008–2012 crises. The noise measure roughly doubles (from 3.39 to 6.92 bp), signalling a genuine, if moderate, rise in average rich or cheap dispersion in the higher-rate, higher-volatility regime. The deterioration is almost entirely a long-end phenomenon: the post-2018 bin MAEs are 2.9 (0–2y), 2.1 (2–5y), 3.0 (5–10y), 2.8 (10–20y), 3.6 (20–30y) and 9.9 bp (30–50y). The ultra-long bin is the only one materially worse than in the paper, consistent with a thin, lumpy supply of 30y–50y lines being repriced violently in 2022–2024 following the escalation of the

war in Ukraine and the resulting increase in prices of energy, which resulted in an inflation spike and the restrictive monetary policy carried out by the ECB.

2.8 VALIDATION AGAINST THE PAPER

Validation against a single study, however carefully replicated, would be a thin foundation for eight years of new estimates; this section therefore triangulates the extended pipeline against four independent anchors: the replicated paper itself, the official curve of the Banque de France, the estimation framework of the ECB and, adversarial, the spline-based curves of commercial providers. Against [1], three checks confirm that the pipeline reproduces the published results on the overlapping window: i. three snapshot dates; for example, on April 2nd, 2018, the fitted 10-year zero yield is 0.69%, matching the paper's $\approx 0.70\%$. ii. The euro-sample mean MAE of 2.39 bp and maximum of 17.13 bp bracket the paper's " ≤ 3 bp on average, 13 bp maximum". iii. The factor structure; the first principal component of the 1–10 year zero yields explains 96.5% of the variance during the paper window and 96.7% in the extended sample, against the paper's 97.55%. The robustness of this three-factor reading itself, against the classic critique of Lord and Pelsser [23] that level, slope and curvature may be artefacts of the covariance structure rather than facts of the market, is examined at length in Case Study 2.

The second anchor is institutional: the issuer's own central bank. The Banque de France has estimated zero-coupon curves on French data since the mid-1990s under the procedure of Ricart, Sicsic and Jondeau [10] discussed in Section 1.7, retains the Svensson specification for its official curve after testing it against the Nelson-Siegel alternative and publishes the resulting series on its Webstat statistical portal [24], so that a public, independently estimated Svensson curve exists for the very market studied here over most of the sample. The concordance is one of method as much as of output: this thesis's pipeline coincides, in model, weighting scheme and data type, with the estimator that the French authorities themselves regard as the appropriate description of the OAT term structure. Nor is

the official curve a statistical artefact: it feeds a research tradition that reads the French term structure for macro-financial signals, from the information-content comparison of the French and German curves in Jondeau and Ricart [25] to the recent reading of the euro-area inversion by Bussière and Lhuissier [26], who caution that the ECB's accumulated asset holdings depress long yields and make an inverted slope overstate recession risk, a caveat that carries directly to the French inversion of 2022–2024 documented in Section 2.9. It is precisely this use of the fitted curve, as a lens on expectations, term premia and policy, that the parsimonious specification defended in Sections 1.8 and 1.9 exists to serve.

The third anchor is the ECB. Nymand-Andersen [27] documents the conceptual framework with which the ECB selected the model behind its daily euro-area yield curves: candidate spline-based and parsimonious specifications are confronted over in-sample and out-of-sample fit, robustness, smoothness and the Svensson family emerges as the production choice. Two aspects of that exercise matter here. First, it reaches independently, on euro-area data and under central-bank production constraints, the same conclusion this thesis argued for its own purpose in Section 1.8. Second, its evaluation criteria are essentially the ones this chapter had already applied: goodness of fit in Section 2.6, full-sample robustness re-runs in Section 2.5, smoothness by construction of the functional form. The extended pipeline is thus validated not only against one paper's numbers, but against the test protocol of the institutions whose curves are the euro-area standard.

The fourth anchor is deliberately adversarial. Commercial providers, Bloomberg foremost, fit sovereign curves with spline-based methods whose objective is the closest possible mark-to-market of every traded security and in periods of acute stress such curves track dislocated prices better than any six-parameter form can, since their flexibility lets them follow the cross-section wherever it goes. That is exactly why they are unsuited to the object of this thesis, for the reasons established in Section 1.8: a curve that absorbs dislocations cannot measure them and spline coefficients carry no macro-financial reading of the kind

exercised by [25] and [26]. The comparison therefore sharpens the methodological claim of this work, which has never been that Svensson is the best curve in every sense, but that it is the right instrument for measuring the pricing coherence of the French market and for interpreting its factors.

Where the replication differs, it does so transparently: the short-end (0–2y) bin is noisier than the paper’s, with a mean of ≈ 3.6 versus 1.62 bp on the euro window, reflecting the thinner early short-end cross-section and the known fragility of the Svensson short end when few short bonds are quoted.


2.8.1 THE FITTED CURVE ACROSS REGIMES

Whereas the MAE compresses the quality of each day’s fit into a single number, it is instructive to open the results with the opposite view: a security-level snapshot of the curve on a few representative dates. Following the paper [1], Figure 5 plots the fitted par-yield curve on the three dates used there, meaning March 25th, 2003, June 10th, 2008 and April 2nd, 2018 and adds the last available day of the extended sample, namely June 5th, 2026, which is the specific contribution of this thesis.

The par yield at maturity m is the annual coupon rate at which a bond maturing in m years would trade exactly at par, meaning with a price equal to 100; computed from the fitted zero curve, it re-expresses the estimated term structure in the coupon-yield space that practitioners actually quote, and is the natural coupon-space complement to the zero curve. Each row of Figure 5 pairs two panels. The left panel superimposes three objects: the smooth fitted par curve, the black line, the observed yields to maturity of the individual OATs quoted that day, the blue circles, and the model-predicted YTM obtained by repricing each bond off the Svensson curve, the red crosses. The right panel reports, security by security, the fitting error in basis points, meaning the observed minus the predicted YTM, whereby a positive error means the bond observed yield lies above the fitted curve, meaning it trades cheap, at a lower price than the model implies, and a negative error means it is traded rich.

The fit is judged by the proximity of the crosses to the circles in the left panel and by the size and randomness of the fitting errors in the right panel, not by the vertical distance between the points and the black par line, which is expected: a seasoned high or low coupon bond trades at a premium or discount and so has a YTM that legitimately differs from the par yield of a hypothetical new par bond at the same maturity. Furthermore, the right-hand fitting errors are exactly the per-security quantities whose daily cross-sectional absolute average defines the MAE; each snapshot therefore decomposes one day's MAE across the maturity spectrum, showing not just how large the misfit is, but where along the curve it sits. These snapshots matter, then, because they simultaneously localize the fit at the individual bond level on a representative day and display the shape of the term structure, upward-sloping, humped, flat or inverted, which encodes the prevailing macroeconomic and monetary regime.

Figure 5



Source: author's elaboration in Python

As shown, the first date is a textbook upward sloping curve rising from $\approx 2.5\%$ at one year to $\approx 5\%$ at 30 years, signalling a “before the crisis” moment. Observed and predicted YTM are nearly coincident and the fitting errors are small and unsystematic, almost all within ± 2 bp, a few reaching ± 4 bp, with a single ~ 30 -year line near -6 bp. A clean, well-identified fit. June 10th, 2008, clearly signals a crisis, with a downward sloping, twice-humped term structure, where yields are

compressed in the narrow 4.55–4.9% band. This is the loosest of the four fits: the fitting errors are the widest, spanning roughly -6 to $+6$ bp and, tellingly, they are systematically signed along the curve, positive at the very short end, negative around 5–9 years, positive again around 13–18 years, rather than random. That wave pattern is the security-level fingerprint of a day on which the smooth six-parameter curve cannot fully absorb the cross-section, the counterpart of the crisis-time MAE spikes and of the reading in [1] of a shortage of arbitrage capital. April 2nd, 2018, represents the “after the crisis” period, characterized by the Zero Lower Bound (ZLB). The curve is smooth and upward sloping, running from $\approx -0.5\%$ at one year to $\approx 1.75\%$ at 48 years. The fit is again tight, with errors mostly within ± 2 bp and none beyond about ± 4 to 6 bp, and a 10-year of $\approx 0.7\%$, which is the lowest rate at which the original sample closed. Finally, June 5th, 2026, is the latest date, signalling an upward-sloping and slightly concave curve at a normalized level, from $\approx 2.8\%$ at the short end to $\approx 4.6\%$ at the long end.

In the core of the curve the 2026 fit is as good as, or better than, the calm 2003 and 2018 dates, and far cleaner than the 2008 snapshot: the fitting errors are essentially flat and clustered within ± 2 to 3 bp all the way out to about 25 years, where the overwhelming majority of liquidity and outstanding notional resides and the observed and predicted YTM_s are visually indistinguishable there. The misfit is instead entirely confined to the ultra-long tail: the errors widen to roughly -12 to -17 bp beyond 30–40 years and a single ≈ 46 -year line sits about 60 bp below the fitted curve, signalling a stale quote on an illiquid ultra-long OAT. This is precisely the 30–50-year segment flagged in Figure 3 as the locus of the post-2022 deterioration and isolating it to a handful of 40-year-plus points is itself reassuring. The economic message is therefore twofold and central to the thesis: eight years beyond the original sample, in a regime the paper never observed, the Svensson/GSW methodology still returns an economically sensible, well-identified curve, upward sloping, smoothly shaped and repricing the liquid 1- to 30-year cross-section to within a few basis points, so that the model remains an


effective and faithful description of the French government bond market all the way to today, with the sole, well-understood caveat of the illiquid ultra-long end.

2.8.2 HISTORY OF FRENCH TERM STRUCTURE

Figure 6 plots the fitted zero-coupon yields at the 2-, 5-, 10- and 30-year tenors across the whole sample, with the euro launch marked and the unreliable pre-1993 short end masked out. It is the time series view of the estimated term structure: the vertical position of the lines is the level of rates and the gap between the 30-year in red and the 2-year in blue lines is the slope. Because the four series move together but not identically, being highly correlated yet fanning in and out over time, the picture already previews the level and slope factor structure formalized by the Principal Component Analysis (PCA).

It traces, in a single frame, four decades of French monetary and macro-financial history and it is here that the contribution of the extension is most visible. Rates begin near 8–10% in the late 1980s franc era of high inflation and EMS credibility premia, then decline secularly through the EMU convergence of the 1990s and the launch of the euro, settling into the 2–5% band of the 2000s. The Global Financial Crisis and the euro-area sovereign crisis accelerate the descent, which the ECB’s negative-rate and asset-purchase policies then push below zero: the 2-year yield is negative for roughly eight years and the entire curve bottoms out at the COVID-19 trough of 2020. The paper [1] closed on this downward trend; the extension captures its violent reversal, with the return of inflation and the fastest tightening cycle in the ECB’s history driving a near-vertical repricing of the whole curve, during which the lines compress and briefly cross during the 2022–2023 period, before re-steepening into the normalized configuration of the last years. In short, what the paper [1] documented as a secular decline is revealed as the trough of a full disinflation and normalization cycle and the long end, robust throughout, provides the reliable backbone of that narrative while the masked short end honestly acknowledges where the early-sample data cannot support the curve.

Figure 6




Source: author's elaboration in Python

2.9 RESULTS AND INTERPRETATION

The extension does not merely add data; it overturns the paper's central stylized facts and sets up the three case studies. The paper documented a secular decline in French rates and a flattening slope, with level becoming "slightly negative at the end of the period". Running the same model to 2026 traces the reversal of both. i. The level reverses: the 10-year zero yield, $\approx 0.7\%$ at the paper's close, falls to a trough of -0.39% on August 15th, 2019, then rises monotonically through the 2022–2024 tightening to 3.74% on June 5th, 2026. The 30-year ends at 4.62% and the asymptotic level β_0 at $\approx 5.1\%$. The "downward trend" is thus revealed as the trough of a longer cycle, not a trend. ii. Negative rates, then their exit: in the euro area the fitted 2-year yield is negative on 1,947 trading days, from August 25th, 2014, to March 28th, 2022, with the deepest 2y $\approx -0.84\%$ in 2019 and the 10-year is negative on 424 days, a regime that the paper only glimpsed at its very end. The unconstrained $\beta_0 + \beta_1$ design is what allows the model to track it. iii. The slope

inverts: the paper's flattening continues past zero; the 10y – 2y slope turns negative in 2022–2023, reaching about –0.32 pp in mid-2023, meaning that the yield of 2y bonds was higher than that of 10y bonds, and oscillates around zero through the June 2024 snap election before re-steepening to 0.94 pp in June 2026. Correspondingly, while the level factor still dominates the PCA, as shown in Figure 7, explaining $\approx 96.7\%$ of variance, the slope factor share rises relative to the paper, the statistical fingerprint of the inversion. Level, slope and curvature follow the paper's definitions throughout: level = 10y, slope = 10y – 2y, curvature = $2 \cdot 5y - (2y + 10y)$. iv. The on-the-run premium remains negligible: re-estimating the curve each day without the on-the-run and first-off-the-run bonds of every maturity range and comparing the synthetic predicted yield to the observed one, the premium stays economically trivial through 2026: the 10-year premium averages ≈ 2 bp with 94% of days inside ± 5 bp, the 5-year ≈ 3.2 bp, the same order as the fitting error, exactly as the paper found and well within the bounds it set. The souche and repo architecture documented in Chapter 1 therefore continues to suppress the premium that pervades the US Treasury market.

Figure 7



Source: author's elaboration in Python

Figure 7 plots the three classic factors of the fitted curve over the euro sample. These three summarize almost all the information in the curve, so the panels are the economically labelled view of the factor structure formalized with the PCA. Level is the overall height of the curve, long-horizon nominal rate expectations, composed of expected inflation plus the real rate plus a term premium. It moves with the monetary and inflation regime. Slope is the steepness, the gap between long and short rates. It encodes the stance of monetary policy relative to its long-run anchor and the business cycle outlook: a steep curve signals accommodation or expansion ahead, a flat or inverted curve signals tight policy and elevated recession risk. The slope is a well-known leading indicator of GDP growth.

Curvature is the hump, the belly relative to the two wings. It reflects medium-term, policy-cycle expectations and is the smallest and least persistent factor, explaining under 1% of the variance.

On the overlapping window the reproduction matches the paper closely; the extension then reveals qualitatively new behaviour that the 1999–2018 sample could not contain, as reported in Table 3.

Table 3

Sample	PCA		
	PC1	PC2	PC3
Paper euro (1999–2018)	96,46%	2,95%	0,55%
Euro extended (1999–today)	96,72%	2,81%	0,44%
Full extended (1987–today)	89,16%	9,76%	1,03%

Source: author’s elaboration in Python

Table 4

Sample	PC2 levels	PC2 changes	Years
Convergence (1999–2007)	15,14%	7,32%	9
GFC & sovereign crisis (2008–2012)	4,49%	7,37%	5
Low or negative rates (2013–2021)	1,21%	7,95%	9
Inflation & ECB tightening (2022–2024)	1,70%	23,57%	3
Normalization (2025–today)	8,57%	10,44%	1,4

Source: author’s elaboration in Python

Three points stand out. First, the level confirms and then breaks the paper’s narrative: [1] documented a secular decline to a 10-year yield that stayed just positive, with their sample minimum of $\approx 0.11\%$; the extension shows it turning genuinely negative in 2019–2021 and then reversing to 3.74% by June 2026, representing a full monetary cycle rather than a one-directional trend. Second, and most importantly for what follows, the slope moves from a regime the paper described as merely flattening, being negative on a single day in 1999–2018, to one of sustained inversion: 233 trading days below zero in 2022–2024, bottoming at -0.33% in June 2023, when 2-year yields rose above 10-year ones under the ECB’s aggressive tightening. This inversion is a feature the original study never

observed. Third, the curvature behaves exactly as the paper found, small, highly time-varying and mostly slightly negative, deepening to -0.87% at end-2023, the belly signature of the inversion episode. Together these panels preview the PCA results, with the level still dominating the curve variance, $\sim 96.7\%$, but the slope factor share rises relative to the paper precisely because of the 2022–2024 inversion, and it does so not in levels but in changes, as shown in Table 4 and as it will be addressed later in the thesis.

These shifts are precisely the raw material of the three studies that follow. The level and slope repricing and the foreign-investor sensitivity motivate Case Study 1, which deals with the OAT-Bund spread and political risk. The inversion and the rising slope-factor share motivate Case Study 2, a PCA of level, slope and curvature across the 2022–2024 tightening. And the noise and MAE spikes as a gauge of arbitrage-capital scarcity motivate Case Study 3, which regards the March 2020 dislocation and the PEPP response, where the doubling of the noise measure and the long-end MAE blow-out documented above are examined at daily resolution.

3. CASE STUDY 1 – POLITICAL RISK FACTOR

The paper [1] measures France's safety premium as the spread of the OAT curve over the German Bund and interprets it as the compensation investors require for holding French rather than the euro area's ultimate safe asset. Over the 1987–2018 sample of [1] that premium was small and stable, consistent with the OAT's status as a core euro asset. The motivation for this case study is that, in the closing years of the extended sample, France experienced an episode of institutional instability without precedent in the euro era, which offers a natural experiment on how political risk is priced into a core sovereign yield curve.

3.1 FACTS AND LITERATURE REVIEW

The narrative facts are well documented. Following the June 2024 European elections, French President Emmanuel Macron dissolved the National Assembly on June 9th, 2024; the snap legislative elections of June 30th and July 7th, 2024, returned a hung parliament with no working majority. There followed an unusually rapid succession of governments, namely Gabriel Attal, then Michel Barnier, whose government collapsed on a no-confidence vote over the 2025 budget on December 4th, 2024, then François Bayrou and finally Sébastien Lecornu, such that France cycled through its fifth prime minister in under two years, with a resignation as late as October 2025 [44]. This paralysis was inseparable from a fiscal problem: French public debt was above €3.3 trillion, $\approx 114\%$ of GDP and the 2024 general-government deficit reached $\approx 5.8\%$ of GDP, far above the EU's 3% reference value, which placed France under the EU Excessive Deficit Procedure in 2024 [45]. The rating agencies responded in sequence: S&P cut France to AA– in 2024, Moody's moved to a negative outlook and Fitch downgraded France from AA– to A+ on September 12th, 2025, explicitly linking the action to political fragmentation and the weakened capacity to deliver fiscal consolidation [44].

Analytically, that such events should move a sovereign spread is exactly what the political-risk literature predicts and this case study is framed against it: i. elections and partisan turnover reprice sovereign bonds; bondholders demand higher credit spreads when incumbents are more likely to be replaced by challengers with less market-friendly platforms [28] and a nation's political leaning is itself a determinant of its spread [29]. ii. Policy and political uncertainty propagate into the whole term structure and into bond risk premia [30] and geopolitical and political risk shocks are known to move sovereign yields and to reshape the slope of the curve [31]. iii. Most directly, adverse political risk shocks, identified from the movement of the sovereign yield spread around political and policy dates, have significant, quantitatively large effects on domestic financial markets and generate cross-border spillovers with the euro area, an effect that is larger when the shock is "populism-related" [32]. iv. Political risk also feeds back into debt sustainability through ratings, willingness to pay and the cost of financing, an effect that is especially large for high-debt countries during periods of high interest rates [33], precisely France's 2024–2026 configuration.

The objective of Case Study 1 is therefore straightforward: to extract the 10-year OAT-Bund spread from the two internally consistent Svensson curves of this thesis and to quantify how the 2024–2026 political instability repriced France's safety premium, addressing, in particular, whether the OAT was temporarily downgraded from a pure core asset, co-moving one-for-one with the Bund, to a middle-ground asset carrying a persistent risk premium.

3.2 TWO COMPLEMENTARY LENSES

The theme is examined from two methodologically distinct angles, each anchored to a different strand of the literature. The first is the OAT-Bund safety premium and a Brexit-style event study; here, France's safety premium is measured as the OAT minus Bund zero-coupon spread and its reaction to the political calendar is quantified with an event study and a spread-change regression that carries a market-wide risk-aversion control. The event-study design and the

reduced-form spread equation follow [34], which studies the impact of the Brexit referendum on UK and European sovereign risk premia and documents a structural change in risk pricing after the vote. This provides a benchmark: a transparent, high-frequency mapping from political events to the premium. The second is the decomposition and the Institut national de la statistique et des études économiques (INSEE) cross-check. The spread is decomposed into a common-euro component and an idiosyncratic French residual, along the whole 2–30-year term structure, which follows a country-specific versus common decomposition, in the spirit of the macro-financial analysis of euro-area sovereign spreads of [35]. The result is validated externally against the French statistical office’s own measure, namely the 3-year OAT-Bund spread of [36], and is extended to the transmission of the premium to the real cost of new loans to French firms, replicating INSEE’s two-stage pass-through in reduced form.

These two are complementary: the first delivers a transparent event-to-premium mapping, while the second attributes and validates the premium, namely how much is genuinely French, whether it prices structural fiscal risk rather than near-term rollover stress and whether it reaches the real economy.

3.3 OAT-BUND SAFETY PREMIUM AND BREXIT EVENT STUDY

What follows is a description, section by section, in the same spirit as the walkthrough of the base methodology in Chapter 2.

3.3.1 THE SAFE BENCHMARK: FITTING THE BUND CURVE


The premium is a spread, so it needs a safe benchmark. The paper [34] used a German curve downloaded from Bloomberg; here the benchmark is constructed endogenously, by fitting the German Bund zero-coupon curve with the very same Svensson and GSW engine used for the OATs, which is the single most important design choice of the study, because it guarantees that the two curves are internally consistent and that any spread is a genuine credit or liquidity differential rather than an artefact of two different estimation methods.

Concretely, the Bund curve reuses the entire base pipeline unchanged: it reads two parallel Bloomberg exports assembled for Germany through the same data loading functions as the OATs, classifies the German nominal universe into its three segments, Bundesschatzanweisungen or Schätze, ~2y; Bundesobligationen or Bobls, ~5y; Bundesanleihen or Bunds, 10–30y and applies the same static filters, fixed-coupon EUR bullets with at least 12 months to maturity and the abnormal quote and plausibility screens and the same robust daily fit as the OAT curve. The German cross-section is entirely euro-era, starting from 1999, so no pre-euro caveats apply.

3.3.2 THE SAFETY-PREMIUM SPREAD AND TERM STRUCTURE

With both curves in hand, the premium is assembled as follows: loading the OAT zero curve and the Bund zero curve, aligning them on their common trading days and returning the OAT minus Bund spread in basis points at the 2-, 5-, 10- and 30-year tenors. On this thesis data the spread panel runs from 1999 to 2026, for a total of 7,156 common days, and its 10-year element, as represented in Figure 8, already tells the story: the premium averages ≈ 36 bp in 2019 and ≈ 33 bp in 2021, rises to ≈ 62 bp across 2024 and ≈ 70 bp across 2025 and stands at 67.8 bp on June 5th, 2026, with a full sample range from ≈ -11 bp in the early 1999 convergence to a 197 bp peak on November 15th, 2011, the height of the euro-area sovereign crisis. The plot of Figure 8 deliberately begins in 2002 for data-honesty reasons at the short end: in the early euro-convergence years there were no short-dated OATs, so the fitted 2-year OAT is a Svensson extrapolation and its spread is unreliable, whereas the German 2-year is always pinned by Schätze. Starting in 2002 ensures that all four plotted tenors rest on actual traded bonds rather than extrapolation, so every line in the figure is data-constrained.

Figure 8



Source: author's elaboration in Python

Figure 8 also shows us that through the pre-crisis great moderation, 2002–2007, the spread is essentially zero at every tenor, with a 10-year mean ≈ 2 bp: France and Germany were both AAA core and priced as near-perfect substitutes. The GFC opens a modest wedge bringing the 10-year mean to ≈ 18 bp, which then explodes during the euro-area sovereign crisis, where the highest spread is observed, with the 10-year premium peaking at ≈ 197 bp and the 30-year at ≈ 218 bp on November 15th, 2011, around France's loss of its S&P AAA, before Draghi's July 2012 “whatever it takes” compresses it again. From 2015 to 2021 the premium is low and stable, 10-year mean ≈ 34 bp, under ECB asset purchases, punctuated only by the 2017 election scare and the March 2020 COVID spike. It then rises through the 2022–2023 inflation and tightening phase and steps up to a persistently higher plateau in 2024–2026, with the tenors clearly fanned out: at the latest date the spread is ≈ 17 bp for the 2-year, ≈ 68 bp for the 10-year and ≈ 108 bp for the 30-year, signalling a steeply upward-sloping spread curve.


A notable feature of Figure 8 is that the 30-year premium turns negative for a short spell on 216 days, meaning $\approx 3.4\%$ of the window, concentrated in 2004 and, most deeply, in 2009–2010, reaching ≈ -20 bp on November 19th, 2009. This is

not a credit signal that France was safer than Germany; it is a supply-and-demand technicality specific to the ultra-long segment. France is the euro area's dominant issuer of 30- and 50-year nominal paper, while ultra-long Bunds are comparatively scarce and structural, price-inelastic demand from euro insurers and pension funds for long-dated assets, liability-driven and duration-matching investment, concentrates precisely there, especially in 2009, when such investors rushed to lock in long yields after the GFC. The net effect richened the ultra-long OAT relative to the Bund and pushed the spread below zero. The episode is a useful caution: at the very long end the safety premium is partly a supply-demand price, not a pure credit spread.

The premium is not a constant but a regime-dependent variable: near zero and even negative at the long end, mainly for technical reasons, when France trades as a pure core asset and explosive under systemic or idiosyncratic stress.

The section closes with a Bund against AAA robustness panel that justifies using the exact Bund rather than the AAA proxy, plotted in Figure 9. In calm periods and in the current episode the two benchmarks give a nearly identical 10y premium; however, during the 2008–2012 crisis the OAT-Bund spread is far wider than the OAT-AAA, peaking around 200 bp against roughly 75 bp, because flight to quality compressed Bund yields specifically, whereas the AAA basket, which blends the Netherlands, Austria and others with Germany, understates the premium to the ultimate safe asset. Using the exact Bund therefore matters precisely in stress episodes, specifically the ones this event study is about.

Figure 9



Source: author's elaboration in Python


The two lines measure the same French premium against two different notions of safeness. The Bund is a single sovereign, the euro area's deepest, most liquid instrument and the destination of flight-to-quality flows, the ultimate safe asset. The ECB AAA curve is instead a basket: a Svensson spot curve fitted on all euro-area central government bonds rated AAA, meaning Germany together with the Netherlands, Austria, Finland, Luxembourg and, crucially, France itself until it lost the AAA rating in 2012–2013. Because the non-German members trade a few basis points above the Bund, the AAA curve sits above the pure Bund curve, so the OAT-AAA premium is mechanically smaller than the true OAT-Bund premium. In calm regimes the two premia are nearly indistinguishable, which mutually cross-validates the two independent benchmarks. The divergence appears exactly in stress: during the sovereign crisis the true OAT-Bund premium averages ≈ 63 bp against the ≈ 26 bp for the OAT-AAA and, on the worst day, namely November 15th, 2011, the gap reaches ≈ 100 bp, with the Bund-based ≈ 197 bp versus the AAA-based ≈ 97 bp, almost double. Over the whole sample the two are correlated but the Bund spread is wider by ≈ 12 bp on average.

The crisis divergence has a clear cause and implication. Flight to quality compresses the Bund specifically, not the whole AAA basket, so the true premium to the ultimate safe asset widens far more than the OAT-AAA proxy. Moreover, pre-2012 the AAA basket contained France itself, biasing the proxy toward zero and in a few cases producing negative premia, visible in the figure as the orange line dipping below zero in 2009–2010. The AAA proxy therefore understates the genuine safety premium precisely when the premium is most informative. This is the main reason why the study adopts the exact German Bund and not the ECB AAA proxy.

3.3.3 DATA-DRIVEN SHOCK DETECTION

Because the 2025–2026 leg of the political calendar is still being finalized, the study does not rely solely on hand-coded event dates. Shock detection flags every day whose one-day spread change exceeds $k = 3$ rolling standard deviations, meaning a 60-day rolling window, so that the market’s own verdict on which days were shocks emerges independently of any calendar. Run on the 10-year premium, the largest data-driven shocks are dominated by the 2011 sovereign crisis, for example the +28 bp on November 15th, 2011, $z \approx 3.8$ and include the COVID “dash for cash” with a +19 bp on March 12th, 2020, $z \approx 6.0$ and the April-2017 election relief accounting for a –20 bp, $z \approx -4.9$. The study then zooms into 2024–2026, overlaying the hand-coded political events and the detection shocks on the 10-year premium, so that the calendar and the data corroborate each other. The output can be seen in Figure 10.

Figure 10



Source: author's elaboration in Python

Figure 10 is very useful as it zooms the safety premium onto the political episode and overlays two independent readings of when the market moved: the hand-coded calendar, event line and the data-driven shock detector, red dots. The premium sits at roughly 45–50 bp through early 2024, already elevated by 2022–23 repricing, then steps up sharply around the June 9th, 2024, dissolution to about 70 bp, climbs to its window peak of ≈ 85 bp on December 2nd, 2024, the eve of the fall of the Barnier government and thereafter oscillates on a persistently higher plateau of ≈ 65 –75 bp through 2025–2026, ending at 67.8 bp. The detector flags cluster exactly where the narrative says they should: four consecutive shock days in the dissolution week, a visual proof that the repricing was gradual rather than a single jump and again during the Barnier crisis. Tellingly, the largest 2025 cluster is in late August, before Bayrou's confidence vote, signalling that the market moved in anticipation, whereas the Fitch downgrade and Lecornu's 24-hour resignation produced no detected shock at all, sitting quietly on the plateau. For a better numerical understanding of the figure, see Table 5.

3.3.4 CUMULATIVE ABNORMAL SPREAD CHANGE (CASC)

The core statistical object is a classic event study. For each political event, the study computes the cumulative abnormal change in the spread over three post-event windows, $[0, +1]$, $[0, +5]$ and $[0, +10]$ trading days: normal daily drift μ and volatility σ are estimated on the 60 trading days ending 6 days before the event and the abnormal cumulative change over an n -day window is: $CASC_{[a,b]} = \sum_{\tau=a}^b \Delta Spread_{e+\tau} - n\mu$, with a t-statistic of $t = \frac{CASC_{[a,b]}}{\sigma\sqrt{n}}$.

$CASC_{[a,b]}$ is the cumulative abnormal spread change over the post event window $[a, b]$: the spread actual cumulative move over the window, net of the move it would normally have made. It is the object of interest, computed in basis points. $\Delta Spread_{e+\tau}$ is the one-day change in the 10-year OAT-Bund spread on a given trading day, where e is the event day, meaning the trading day on, or nearest to, which the political event occurs. Formally it is the position of that day in the ordered sequence of trading days, namely the origin of event time. τ corresponds to the summation index, measured in trading days relative to the event; it runs over the window, from a to b . $[a, b]$ represent the lower and upper bounds of the post-event window, in trading days relative to the event. n is the number of trading days in the window. μ is the normal expected daily drift of the spread: the mean of the daily spread changes over the estimation window, that are the 60 trading days ending 6 days before the event. The term $n\mu$ is therefore the drift the spread would normally accumulate over n days; subtracting it is what makes the change abnormal. Finally, σ is the normal daily volatility: the standard deviation of the daily spread changes over the same estimation window. Since $\sigma\sqrt{n}$ is the standard deviation of an n -day cumulative change under normal conditions, treating daily changes as roughly independent, the t-statistic expresses the abnormal change in units of its normal standard error, meaning how many normal standard deviations the post-event move represents.

This is the standard abnormal return design transported to a spread and it is the exact logic by which [34] isolates the Brexit-referendum effect on UK versus Europe risk premia.

On the full 2024–2025 calendar, the result is unambiguous and sharp. The June 9th, 2024, dissolution is, by far, the single clearest political shock in the entire sample, with a cumulative abnormal widening of +24.6 bp over the following week, $[0, +5]$ and $t \approx 12.0$; no other event, in 2024 or 2025, comes close. The remaining 2024 dates show small, mostly negative CASCs, retracements around already-repriced levels, with only the December 4th, 2024, fall of the Barnier government registering a secondary, weaker move with $t \approx -2.6$. The decisive new evidence is that the 2025 sequence produced no fresh widening: Bayrou's lost confidence vote on September 8th produced a CASC of -7.8 and a $t \approx -2.2$; Lecornu's appointment produced a CASC of -5.5 and his same-day resignation, on October 6th, resulted in a CASC of $+1.2$ with a $t \approx 0.3$, the Lecornu-II reappointment, again a CASC of -5.0 and even the Fitch downgrade to A+ on September 12th produced a CASC of $+0.7$ with a $t \approx 0.2$. All left the premium essentially unmoved; the only marginally significant reaction to Bayrou's fall is a narrowing, consistent with a fully anticipated event resolving into relief. The economic reading is that the safety premium was repriced once, at the June 2024 dissolution, onto a new ≈ 70 bp plateau, after which the market had so thoroughly internalized chronic French political instability that a sovereign downgrade and a 24-hour prime-ministerial resignation passed almost without a trace. However, these remain *ceteris paribus* associations, a limitation acknowledged and maintained throughout the case study.

Table 5

Event	Date	Spread-pre (bp)	CASC [0,1]	t [0,1]	CASC [0,5]	t [0,5]	CASC [0,10]	t [0,10]
EU elections; dissolution of the National Assembly	09/06/2024	48,9	11,5	9,74	24,6	12,01	21,2	7,63
Legislative elections, 1st round	30/06/2024	73,7	-8,3	-3,08	-15,6	-3,33	-15,8	-2,50
Legislative elections, 2nd round	07/07/2024	62,3	0,3	0,12	-3,3	-0,65	-3,9	-0,57
Barnier appointed Prime Minister	05/09/2024	66,9	-1,8	-0,49	-3,5	-0,57	-1,1	-0,13
Barnier government falls	04/12/2024	82,4	-7,5	-3,39	-9,9	-2,58	-7,2	-1,39
Bayrou appointed Prime Minister	13/12/2024	74,8	2,1	0,73	0,3	0,05	1,8	0,26
2025 budget adoption; no-confidence motion fails	05/02/2025	71,9	-0,7	-0,21	-0,5	-0,08	-3,0	-0,41
Bayrou loses confidence vote on debt plan	08/09/2025	80,7	-5,0	-2,47	-7,8	-2,19	-6,9	-1,44
Lecornu appointed Prime Minister	09/09/2025	79,0	-3,4	-1,68	-5,5	-1,56	-5,2	-1,10
Fitch downgrade France	12/09/2025	73,9	0,1	0,03	0,7	0,20	0,8	0,16
Lecornu resignation, <24h	06/10/2025	76,7	4,5	2,19	1,2	0,33	-4,9	-1,02
Lecornu reappointed	10/10/2025	77,8	1,0	0,51	-5,0	-1,41	-3,0	-0,61


Source: author's elaboration in Python

Table 5 is to be read alongside the time-series Figure 10 and converts the visual step and plateau into hard numbers. Exactly one event repriced the premium: the National Assembly dissolution. From a pre-event level of ≈ 49 bp it produced a positive, overwhelmingly significant abnormal widening in every window, namely +11.5 bp in the first day, +24.6 bp over the following week and +21.2 bp over two weeks, producing a visible jump in the figure, far beyond every fluctuation. Every subsequent event is either a retracement or a non-event. The first-round elections and, strikingly, the fall of the Barnier government show significant negative CASCs: the spread had run up into those episodes and narrowed on the outcome; anticipated events resolving into relief, the strategy of buying the rumour and selling the fact. The 2025 sequence barely registers: Bayrou's lost confidence vote is a widely anticipated relief, the Fitch downgrade is essentially zero and Lecornu's dramatic 24-hour resignation is only a one-day blip that fully reverses within the week. In one line: the premium was set once, at the dissolution, and every later shock was already priced.

3.3.5 ROLLING OAT-BUND BETA

To ask whether the OAT still behaves as a core asset, the study runs a 60-day rolling regression of daily 10-year OAT yield changes on daily 10-year Bund yield changes, reporting the rolling beta and correlation. The interpretation is explicit: a beta and correlation near one mean the OAT co-moves one-for-one with the safe benchmark, a pure core asset; a fall in beta and correlation would signal decoupling, a genuine risk repricing that breaks the co-movement. This distinguishes two very different ways a spread can widen: a level repricing with intact co-movement, $\beta \approx 1$, versus a structural decoupling, $\beta \downarrow$.

Figure 11



Source: author's elaboration in Python

A spread can widen in two very different ways and Figure 11 separates them. Beta answers how much the OAT yield moves when the Bund yield moves by one unit, while correlation tells how tightly they move together. When $\beta \approx 1$ and $\text{corr} \approx 1$, the OAT is a pure core asset moving one-for-one with the euro-area safe benchmark; a persistent fall in both, toward zero or below, signals decoupling. This chart is the direct test of the core versus middle-ground question that the safety-premium level alone cannot settle. Across 2015–2026 both lines hover near 1, with

the full sample mean being ≈ 0.92 for each and with two conspicuous decoupling dips: in spring 2017, around the Le Pen presidential-election scare, beta and correlation collapsed to ≈ 0.35 , signalling a genuine, if brief, severing of the OAT from the Bund on redenomination or “Frexit” fears and a milder dip in March 2020, when beta was around 0.6 during the COVID dash for cash. From 2021 onward the co-movement is tight and close to one. The decisive observation concerns the red event lines. Around every major 2024–2025 shock the co-movement does not break: beta and correlation stay ≈ 1 at the dissolution and the Barnier collapse and through the entire September-October 2025 sequence, with beta even drifting above one into 2026, having a mean ≈ 1.19 . This is the exact complement to the safety-premium figures: the level of France’s premium more than doubled, yet its co-movement with the core never fractured. France was therefore repriced in level, not in co-movement. It became a more expensive core asset, not a decoupled peripheral one. The contrast with 2017 is the whole point. The Le Pen scare was existential, euro-exit risk that briefly cut the OAT loose from the Bund; the 2024–2026 fiscal-political crisis, though it lifted the premium far higher and far more persistently, was priced as credit and fiscal deterioration within the core.

3.3.6 BREXIT-STYLE REGRESSION

The section that most directly reproduces [34] estimates a reduced-form model of the daily spread change: $\Delta \text{Spread}_t = \alpha + \gamma \text{Event}_t + \delta \Delta \text{CISS}_t + \phi \text{Spread}_{t-1} + \varepsilon_t$, with Newey-West (HAC) standard errors. Here Event_t is a dummy equal to one on each event day and the following two days; ΔCISS_t is the daily change in the ECB Composite Indicator of Systemic Stress, the euro-area systemic-stress and risk-aversion proxy, downloaded from the ECB Data Portal, whose role becomes more important after a political regime break; Spread_{t-1} captures mean reversion. On the thesis data, 7,156 days, the ΔCISS_t coefficient is strongly significant, $t \approx 4.5$, while the narrow event dummy is insignificant, $t \approx -0.08$ and mean reversion is weak, meaning that at the daily frequency the premium is driven by systematic risk aversion and the

political effect does not show up as a one-day jump. This is itself an informative contrast with the Brexit case: French repricing was more gradual and anticipated than the single-day Brexit surprise of [34], which is why the event-study cumulative window captures it, but the one-day dummy does not.

The decisive evidence is the structural pre- and post-comparison around the June 9th, 2024, dissolution, which the same function reports. Between the pre- and post-dissolution regimes the average 10-year OAT-Bund spread more than doubled, from ≈ 28.6 bp to ≈ 69.4 bp, while the OAT beta to the Bund stayed essentially at one, edging up from ≈ 0.91 to ≈ 1.02 . The combination is the study's central finding and the precise sense in which the OAT became a middle-ground asset: France continued to co-move with the core, so without decoupling, $\beta \approx 1$, but now carried a persistently higher premium: it was repriced in level, not disconnected. This is consistent with the political-risk shock evidence of [32], which finds that a durable spread widening is the response to political instability, and with the ratings and debt-sustainability channel of [33], which shows that sovereign repricing is largest under high rates; it sets up the next insight, which attributes this premium to genuinely French factors and traces its transmission to the real economy.

3.4 DECOMPOSITION AND INSEE CROSS-CHECK

This study established that the French safety premium was repriced: around the June 2024 dissolution the average 10-year OAT-Bund spread more than doubled, from ≈ 28.6 bp to ≈ 69.4 bp, while the OAT beta to the Bund stayed at one, a repricing in level without decoupling. What an event study cannot establish, by construction, is whose premium it is. A spread over the Bund is a joint statement about two curves and its widening can reflect three conceptually distinct forces: a genuinely French repricing of credit and political risk, a common repricing of euro-area sovereign risk in which France merely participates, as in the systemic episodes of 2011–2012 and, lastly, movements of the safe benchmark itself, whose scarcity value and monetary-policy exposure change the spread with no French news at all.

The Brexit-style regression already hinted that the daily premium loads heavily on the euro-area systemic-stress factor; however, a single control in a single-tenor regression is measurement, not attribution. This second and concluding analysis of Case Study 1 is built to deliver that attribution and it does so with three specific objectives.

First, attributing along the whole term structure: the daily OAT-Bund spread at the 2, 3, 5, 10 and 30-year tenors is decomposed into a common-euro component, the part explained by observable euro-area factors common to all member sovereigns and an idiosyncratic French residual, which is the political premium in the strict sense. Second, external validation: both curves underlying the spread are this thesis's own Svensson constructions, so before any interpretation the spread is anchored to the only official French measure available, namely the 3-year OAT-Bund spread published by the French national statistical office, the Institut National de la Statistique et des Études Économiques (INSEE), in its March 2025 *Note de conjoncture* [36], which reports a widening from ≈ 23 bp at end-May 2024 to ≈ 41 bp at end-January 2025. A replication of the level and, more importantly, of the widening of an externally produced statistic is the most direct check that the premium studied here is a market fact rather than an artefact of the estimation pipeline. Third, transmission: a sovereign premium that never left the bond market would remain a market curiosity. INSEE's own interest in the spread stems precisely from its pass-through to the borrowing costs of the real economy and this study closes the chapter by replicating the two-stage scheme through which [36] quantifies that transmission to the interest rate on new loans to French non-financial corporations (NFCs).

3.4.1 RELATED LITERATURE AND THE OFFICIAL BENCHMARK

The decomposition follows the country-specific-versus-common logic of the macro-financial literature on euro-area sovereign spreads, whose closest reference is [35]. Estimating a multi-market, no-arbitrage affine term-structure model with a single pricing kernel, the canonical representation of [46] augmented with un-

spanned macro factors as in [47], on Belgian, French, German, Italian and Spanish yield curves over 2005–2013, they split each sovereign spread into a fundamental component, justified by country-specific factors, euro-area fundamentals and international influences and a non-fundamental remainder capturing liquidity, political uncertainty and redenomination fears. Their headline result is double-edged: economic fundamentals are the dominant drivers of euro-area spreads, yet shocks unrelated to fundamentals played an important role from the intensification of the sovereign debt crisis in the summer of 2011 [35]. The present analysis transposes that decomposition to the 2024–2026 French episode in a deliberately lighter, daily-frequency reduced form: instead of a latent-factor pricing model, the common component is spanned by three fully observable euro-area factors and the residual inherits the interpretation of the non-fundamental component of [35], with the same honest caveat that a residual is a bundle, conflating idiosyncratic credit news, liquidity and political uncertainty rather than isolating any one of them. The daily-frequency reading of that residual as predominantly political over this sample rests on the identification logic of [32], which extracts political-risk shocks precisely from sovereign-spread movements around political and policy dates and on the tight alignment between residual jumps and the political calendar documented below, the same calendar already used in the event study above.

The external benchmark is the Focus “Since mid-2024, the transmission of base interest rate cuts to French corporate lending rates has been hampered by the reduction in the ECB’s balance sheet, political uncertainty and bond yield movements in the United States”, published by Adjera and Vermersch in INSEE Economic Outlook of March 18th, 2025 [36]. Several elements matter here. First, the measure: INSEE works at the 3-year tenor, chosen because three years is the average maturity of French corporate debt and documents the 23 bp to 41 bp widening of the France-Germany sovereign spread between May 2024 and January 2025, attributing it explicitly to domestic political uncertainty, “notably the dissolution of the National Assembly and the tensions surrounding the adoption of

the 2025 budget” [36]. Second, the model: a two-stage design in which Stage 1 explains the 3-year Bund, the euro-area risk-free anchor, with the 3-month Euribor, the same-maturity US Treasury yield and the size of the ECB balance sheet, an increase worth one point of euro-area GDP lowers the 3-year Bund by 0.72 bp at term in their estimate, while Stage 2 maps the French sovereign yield, decomposed by accounting identity into Bund plus spread, onto the rate on new NFC loans. Lastly, the results this study will be measured against: a 100 bp cut in base rates passes through to French NFC lending rates by about 50 bp over two years, 40 of which within the first year and the political widening of the spread alone raised NFC borrowing rates by about +0.2 percentage points over the year to end-2024 [36]. The INSEE Focus is macro-financial conjuncture analysis, not term-structure econometrics; the value added of the present study is to rebuild its objects from internally consistent estimated curves, to extend the decomposition from one tenor to the whole 2–30-year structure and to separate explicitly, as INSEE’s accounting identity does not, the common-euro from the idiosyncratic driver of the spread it feeds into Stage 2.

3.4.2 METHODOLOGY

The object of study is the zero-coupon OAT-minus-Bund spread, in basis points, computed on common trading days from the two Svensson curves of this thesis at the 2, 3, 5, 10 and 30-year tenors, with the 3-year tenor being carried along the whole study precisely so that every result can be read against INSEE’s official measure. Reusing the internally consistent curve pair is, as argued in the event study, the design choice that guarantees the spread is a credit-and-liquidity differential and not an artefact; the validation section makes that guarantee testable, comparing the end-of-month value and the monthly mean of the 3-year spread in May 2024 and January 2025 with the ≈ 23 bp and ≈ 41 bp of [36], the monthly mean being reported because a point-in-time institutional figure and an end-of-period estimate can legitimately differ by a few basis points on volatile weeks.

The common-euro component is spanned by three daily, fully ECB-sourced factors, deliberately chosen so that no country-level data for third sovereigns are needed. The first is a euro periphery dispersion factor: the spread, per tenor, between the ECB euro-area all-bonds government spot curve and its AAA-only counterpart, both estimated by the ECB with the same Svensson family used throughout this thesis [27] and both retrieved from the ECB Data Portal. Since the AAA bucket is dominated by Germany while the all-bonds curve embeds the whole ranking spectrum, their distance is a clean daily proxy for systemic peripheral risk, at exactly the tenors of the thesis spread. The second is the ECB Composite Indicator of Systemic Stress (CISS), the euro-area risk-aversion measure already used as the market-wide control in the Brexit-style regression above, following [34]. The third is the level of the study's own Bund curve at the matching tenor, which absorbs movements of the safe benchmark itself, the monetary and safe-asset-scarcity channel that INSEE's Stage 1 attributes to policy rates, the US cycle and the Eurosystem balance sheet [36]. One caveat is stated upfront: France itself is inside the euro-area all-bonds aggregate, so the dispersion factor partially contains the French repricing it is meant to control for. The bias runs against the study's thesis: it inflates the common component and shrinks the residual, so the idiosyncratic French premium estimated below is, if anything, a lower bound.

3.4.3 THE DECOMPOSITION DESIGN

The core attribution is a regression in daily changes, estimated tenor by tenor, by OLS with Newey-West (HAC) standard errors: $\Delta\text{Spread}_t = \alpha + \beta_1\Delta\text{dispersion}_t + \beta_2\Delta\text{CISS}_t + \beta_3\Delta\text{Bund}_t + \varepsilon_t$. Working in first differences keeps the analysis on near-stationary ground and makes the reading of the fit immediate: the regression R^2 is the share of daily spread variation explained by common-euro forces, its complement is the idiosyncratic share and the residual ε_t is the daily idiosyncratic French shock. Cumulated over time, the residuals trace the path of the idiosyncratic premium in basis points; the series answers, day by

day, how much of the widening was genuinely French. Two companions complete the block. A level projection of the spread on dispersion and CISS, namely an HAC with longer lags, given the persistence of levels, splits the spread level into a fitted common-euro component and a residual French premium: this is reported as a descriptive decomposition, so a projection and not a structural model, with the change regression remaining the rigorous attribution. The term-structure response module computes, around each date of the political calendar of the event study, the change in the average spread between the pre-event window $[-5, -1]$ and the post-event window $[0, +10]$ at every tenor, complemented by the 10y–2y spread-curve slope: a jump concentrated at the short end signals near-term rollover risk, one concentrated at the long end signals perceived structural fiscal deterioration, the term-structure reading of political risk anticipated by the literature on uncertainty and the yield curve slope [30], [31]. This is also where the case studies close their loop: the idiosyncratic premium series constructed here is the same series whose high-premium regimes are formally dated and classified by the Markov-switching analysis of Section 5.7.2 and Figure 23, so the regime evidence of Case Study 3 stands on the decomposition built in this section.

3.4.4 THE TRANSMISSION DESIGN

The final block replicates INSEE’s transmission scheme in reduced form. A reduced Stage 1 regresses the monthly 3-year Bund level on the 3-month Euribor and the Eurosystem balance sheet, composed of total assets, in trillions of euros, from the ECB weekly financial statements, plus a post-2010 dummy absorbing the unconventional-policy regime, with HAC standard errors; its purpose is not to rival INSEE’s error-correction specification but to verify that, on this thesis data, the monetary anchors of the risk-free leg carry the expected signs and magnitudes, the balance-sheet coefficient being read against INSEE’s -0.72 bp per point of euro-area GDP [36]. The US Treasury driver of INSEE Stage 1 is omitted for lack of data, a stated limitation of a reduced form. The reduced Stage 2 then regresses, at quarterly frequency, the rate on new loans to French NFCs, from the ECB

harmonized Monetary Financial Institutions Interest Rate (MIR) statistics, the dependent variable of INSEE's Stage 2, on the Euribor, the balance sheet and the OAT-Bund spread converted to percentage points. The design differs from [36] in one deliberate respect: where INSEE enters the French sovereign yield as a whole, here the common monetary drivers and the spread enter separately, so that the spread coefficient directly measures the marginal pass-through of the French premium to corporate borrowing costs at a given monetary stance, which is the reduced-form analogue of INSEE's +0.2 percentage-point attribution. The caveats are stated where they bind: the frequency is quarterly, the politically driven portion of the sample is short and the estimates are reported as conditional associations with fitted-versus-observed diagnostics, not as causal effects.

The presentation of the results mirrors the architecture of the methodology. The external anchor comes first: before any decomposition, the 3-year spread produced by the curve pair of this thesis is confronted with the official INSEE figures, since every subsequent claim about the premium is only as credible as the spread it is computed from. The attribution results follow, namely the variance decomposition of daily changes, the cumulated idiosyncratic French premium and its term structure around the political calendar and the reduced two-stage pass-through to the cost of corporate credit closes the study.

3.4.5 RESULTS: EXTERNAL VALIDATION

Table 6 performs the external cross-check. For the two reference dates of the INSEE Focus, end-May 2024, just before the dissolution of the National Assembly and end-January 2025, in the aftermath of the fall of the Barnier government and the budget standoff, it reports the 3-year OAT-Bund spread from the Svensson pair of this thesis in two variants, the end-of-period value and the monthly mean, against the published INSEE counterpart [36]; the third row is the widening between the two dates, the quantity of economic interest.

Table 6

Dates / Widening	End-of-period value (bp)	Monthly mean (bp)	INSEE (bp)
End-May 2024	17,8	19,7	23,0
End-Jan 2025	34,8	38,3	41,0
Widening	17,0	18,6	18,0

Source: author's elaboration in Python

The thesis end-of-month spread reads 17.8 bp in May 2024 and 34.8 bp in January 2025, with monthly means of 19.7 bp and 38.3 bp, against INSEE's ≈ 23 bp and ≈ 41 bp. Two features stand out. In levels, the spread of this thesis sits systematically 3–5 bp below the official measure at both dates. In changes, however, the gap disappears: the widening is 17.0 bp end-of-period and 18.6 bp in monthly means, against INSEE's 18 bp, a match within about one basis point, achieved with a fully independent data pipeline.

The meaning of this pattern is precisely a successful external validation. The small level wedge is the expected footprint of methodological differences: in this work the distance is between two fitted zero-coupon Svensson curves, estimated under continuous compounding on a filtered bond universe, whereas an institutional figure is typically read off quoted benchmark yields, which embed the on-the-run liquidity premia and idiosyncrasies of individual securities that the curve fitting deliberately smooths away. A wedge of this origin is approximately constant over an eight-month window and a constant cancels in differences, which is exactly what the third row shows.

The interpretation is therefore twofold. First, the premium studied in this Case Study is a market fact, not an artefact of the estimation choices of this thesis: an independent construction reproduces the dynamics of the official statistic to within a basis point. Second, the validation licenses the finer use to which the spread is put next, namely the daily decomposition along the whole term structure, since it is the variation of the spread, the very dimension certified here, that the change regressions exploit.

3.4.6 RESULTS: REDUCED STAGE-1

Table 7 reports the reduced Stage-1 regression: the monthly level of the 3-year Bund from this work's own curve on the 3-month Euribor, the Eurosystem balance sheet in trillions of euros and a post-2010 dummy absorbing the unconventional-policy regime, over January 1999 to May 2026 (329 months), with HAC standard errors. Its role, as stated in the methodology, is a sanity check of the risk-free leg, namely whether the monetary anchors of INSEE's Stage 1 do carry the expected signs and magnitudes on the data of this thesis, not an inferential exercise in its own right.

Table 7

	Coefficient	Std. Err.	t		
Constant	0,982	0,130	7,564	R2	0,935
Euribor	0,770	0,043	17,762	n	329
Balance sheet	-0,012	0,038	-0,330	Euribor pass-through	0,77
Post-2010	-0,978	0,203	-4,809	Balance sheet coef	-0,012

Source: author's elaboration in Python

Reading the table through the t-statistic first prevents any misleading equivalence between its rows. Three coefficients are estimated sharply and are significant at any conventional level: the constant, $t \approx 7.6$, the Euribor, $t \approx 17.8$ and the post-2010 dummy, $t \approx -4.8$. Only one is not: the balance-sheet coefficient, with a $t \approx -0.3$, is statistically indistinguishable from zero and nothing in what follows should be read as evidence of a balance-sheet effect on the Bund, only its sign is commented on, for comparison with [36].

Among the significant anchors, the base-rate pass-through is the one that works best: the Euribor coefficient is 0.77, with a standard error of 0.043, meaning that a 100 bp move in the 3-month rate is associated with a 77 bp move in the 3-year Bund. A coefficient below one is exactly what the expectations hypothesis prescribes for this maturity: a 3-year yield averages expected short rates over a horizon in which policy is expected to partially mean-revert, so it moves with the policy stance but less than one-for-one. The post-2010 dummy enters at -0.978

and with a Std. Err. of 0.203: conditional on the Euribor, the unconventional-policy era, composed of QE, forward guidance, negative deposit rates, is associated with a 3-year Bund about one percentage point lower, a compact and strongly significant summary of the lower-for-longer regime.

The insignificant coefficient deserves its own interpretation due to its very nature. The point estimate of -0.012 per trillion euros, converted to INSEE units, with one point of euro-area GDP roughly €150 billion, about -0.2 bp per GDP point, carries the same sign as, but sits well below the -0.72 bp that [36] estimates at term, itself below recent ECB estimates, as INSEE notes; yet with a standard error of 0.038 the data simply do not pin it down. The imprecision has a transparent mechanical reason: the dummy and the balance sheet are strongly collinear, since the regime the dummy marks is defined by the balance-sheet expansion, so the statistical specification cannot sharply apportion the lower-for-longer effect between the stock of asset purchases and everything else that changed after 2010, a separation that INSEE's error-correction dynamics and its US-yield regressor, omitted here, are designed to achieve.

A word on the fit is needed precisely because it looks strong. The R^2 of 0.935 should not be read as strong explanatory power at the monthly frequency: this is a regression in levels of highly persistent series and the coefficient of determination mostly records the fact that the Bund and the Euribor travelled the same two-decade journey, from above 4% at the euro birth, through the negative-rate decade, to the post-2022 tightening, together. Persistent regressors mechanically inflate the R^2 of level regressions, which is why the standard errors are HAC and why no claim in this study rests on this fit; the rigorous attribution work is done where the statistics are trustworthy, in daily changes. What Table 7 certifies is narrower and sufficient: on the data of this work the risk-free leg of the spread responds to the policy stance with the textbook sign and magnitude, so the movements of the OAT-Bund spread analysed next are measured against a benchmark whose own monetary behaviour is well understood.

3.4.7 THE CORE DECOMPOSITION

This is the section the whole study is built around. For each tenor, the daily change in the OAT-Bund spread is regressed on the daily changes of the three common-euro factors, namely periphery dispersion, CISS and Bund level, by OLS with Newey-West standard errors, on 5,557 trading days from September 2004, which marks the start of the ECB curve sample, to June 2026. The regression performs a clean division of labour: everything the three factors capture is, by construction, systemic, repricing that France shares with the euro area as a whole, while the residual is the daily idiosyncratic French shock, the raw material of the political premium. Table 8 collects the five tenor-by-tenor estimates; the R^2 of each regression is the share of daily spread variance attributable to common-euro forces and its complement the idiosyncratic share.

Table 8

	Coefficient	Std. Err.	t
Constant	0,009	0,025	0,336
Dispersion	0,281	0,024	11,892
CISS	6,650	2,964	2,244
Bund	-8,947	1,264	-7,098

Tenor	R2	Idiosyncratic	Dispersion t	Bund t
2-year	0,111	88,9%	6,3	-7,0
5-year	0,157	84,3%	12,2	-7,6
10-year	0,185	81,5%	11,9	-7,1
30-year	0,111	88,9%	6,8	-7,9

Source: author's elaboration in Python

The significance triage is unambiguous and, this time, uniform: all three factors are significant at the 1% level at every tenor, while the constant is never distinguishable from zero, the spread has no daily drift once the factors are accounted for. The dispersion loading rises from 0.135 at 2 years, $t \approx 6.3$, to a peak of 0.281 at 10 years, $t \approx 11.9$, before falling back to 0.120 at 30 years: for every 10 bp of widening in the euro periphery dispersion, the 10-year French spread widens by ≈ 2.8 bp on the same day. The CISS coefficient declines monotonically with maturity, from 9.4 at 2 years to 6.7 at 30 years, with t ranging from 3.8 to 2.1: a large systemic-stress day, a 0.1 jump in the indicator, adds roughly 0.7–0.9 bp to the spread, more at the short end where flight-to-safety is most violent. The Bund coefficient is negative throughout, from -7.1 at 2 years to


−9.8 at 30 years, with $t \approx -7$ to -8 : a 100 bp rise in the Bund compresses the spread by 7–10 bp. Read structurally, this last coefficient is a daily co-movement statistic in disguise: since $\Delta Spread = \Delta OAT - \Delta Bund$, a coefficient of −7 bp to −10 bp per percentage point says the OAT tracks the Bund with a daily beta of 0.9–0.93, fractionally below one-for-one, as in the event study, with the small shortfall reflecting the Bund’s superior safe-haven bid on the most stressed days.

The signs compose a coherent systemic profile. France loads positively on peripheral credit repricing, when Italy- and Spain-type risk is repriced, the OAT is repriced with it and most strongly at the 10-year benchmark tenor, where sovereign credit is priced most cleanly and the market is deepest; positively on systemic risk aversion, on risk-off days investors do not treat the OAT as the terminal safe asset; negatively on the safe benchmark itself, through the mechanical flight-to-quality channel. This is the factor signature of what can be called a middle-ground asset: systematically exposed to euro-area stress, unlike the Bund, yet with loadings an order of magnitude below what peripheral spreads themselves would display.

The R^2 column, however, inverts the emphasis and this is the central quantitative finding of the section. The common-euro component explains between 11.1% (2 and 30 years) and 18.5% (10 years) of daily spread variance, with the familiar hump at the benchmark tenor; equivalently, the idiosyncratic share is 81.5% to 88.9% at every maturity. At the daily frequency, the OAT-Bund spread is overwhelmingly not a systemic object: more than four-fifths of its day-to-day movement is orthogonal to euro-area credit dispersion, systemic stress and the safe benchmark and must therefore originate in France-specific news. This is the statistical licence for the study’s central object, namely cumulating those residuals into an idiosyncratic French premium, examined thereafter. Two qualifications keep the reading honest. First, idiosyncratic is a bundle, not a label for politics alone: at daily frequency the residual also collects microstructure noise, asynchronous quote updates between the two markets and estimation noise from both curve fits, so the 81%-89% is an upper bound on the French signal and it is

the cumulated, low-frequency drift of the residuals, not their daily variances, that carries the political interpretation. Second, the finding is complementary rather than contradictory to [35], who conclude that economic fundamentals are the dominant drivers of euro-area spreads: their decomposition operates on monthly levels, where slow-moving fundamentals accumulate, whereas daily differencing deliberately strips persistence out and leaves the news component, the frequency at which political shocks, in the spirit of [32], are identified. The two results together say that the spread trend is disciplined by fundamentals while its innovations are made in Paris; Figure 12 shows what those innovations sum to.

Figure 12



Source: author's elaboration in Python

Figure 12 plots the descriptive companion of the change regression announced in the methodology. The projection is estimated on the full 2004–2026 sample but the panel zooms on the period from 2015, where the story lives; because the projection includes a constant, the zero line is the average France-specific component over the whole sample, so the red series reads as deviations from France's two-decade norm.

Three phases are visible. Until 2021 the premium oscillates in a narrow band around zero, its 2015–2021 mean being -5.4 bp with a standard deviation of 8.6 and its only notable excursions are textbook events that promptly reverse: a spike to $+26$ bp in February 2017, the redenomination scare ahead of that spring presidential election, fully unwound after the run-off and a brief flare in March 2020 at the height of the dash for cash. In this phase the blue line, the CISS, essentially is the spread: whatever the OAT paid over the Bund was euro-systemic in nature and France, if anything, traded slightly tighter than its systemic exposure would imply, which is exactly the quantitative signature of the core status. From 2022 the premium turns durably positive and begins to drift, $+3.5$ bp on average in 2022, $+6.7$ bp in 2023 and $+9.7$ bp in the pre-dissolution months of 2024, a slow repricing that predates the political rupture and is consistent with the fiscal deterioration documented in the case study opening. The rupture itself then shows as a ratchet: the premium jumps at the June-July 2024 event dates and averages $+29$ bp over the second half of 2024, steps up again around the fall of the Barnier government and climbs through the autumn 2025 cluster, regrouping the Bayrou confidence vote, the sovereign downgrade and the Lecornu resignation, to a sample peak of $+47.7$ bp on October 7th, 2025, the first trading day after the prime minister's resignation. It has not come back: the premium averages $+36.4$ bp in 2025, $+36.6$ bp in 2026 and stands at $+39.4$ bp at the sample end in June 2026.

The meaning of Figure 12 is a statement about composition and it contains the study's sharpest result. Over 2024–2026 the two components of the spread cross: the common-euro component eases, with its annual mean falling from 40.6 bp in 2024 to 34.1 bp in 2025 and 28.2 bp in 2026, as euro-area systemic conditions normalized, while the French premium rises along the path just described. By the end of the sample the decomposition reads $67.3 = 27.9 + 36.6$: the majority of what France pays over Germany is no longer compensation for shared euro-area risk but a country-specific premium. The widening of 2024–2026, in other words,

is not a European risk-off episode that France happened to be caught in, but an idiosyncratic repricing of France, priced in on political dates.


Two features give the interpretation its weight. The first is persistence: the pre-2022 premium behaved like the non-fundamental spread component of [35] in calm times, episodic, mean-reverting and disciplined by fundamentals; the 2017 scare, twice the size of a typical fluctuation, was gone within months. The post-2024 premium does not mean-revert on any horizon visible in the sample: each political shock resets its level and two years after the dissolution it sits an order of magnitude above the old norm. This is what can be seen as the demotion from pure core to middle-ground asset, now measured net of everything systemic. The second is that its persistence is not asserted from visual inspection alone: the red series in this figure is precisely the input whose regimes are formally dated by the Markov-switching analysis of Section 5.7.2, where Figure 23 classifies the post-2024 window in the high-premium state. The usual caveats carry over unchanged from the methodology, the decomposition is a projection, not a structural model and the residuals remain a bundle in which liquidity and estimation noise coexist with politics, but the low-frequency drift documented here is exactly the component of the residual that the noise interpretation cannot generate and it is on that drift that the term-structure analysis of the next section and ultimately the pass-through of the premium to corporate borrowing costs are built.

3.4.8 THE TERM STRUCTURE OF THE POLITICAL PREMIUM

This section exploits the one dimension in which this study is constitutionally richer than its official benchmark. INSEE's measure lives at a single maturity, the 3-year spread [36]; a single tenor can register the size of a repricing but is silent on its shape and the shape is where the diagnosis lies. A sovereign shock concentrated at the short end signals near-term rollover risk, namely doubts about the state's ability to refinance imminent maturities, the profile of the peripheral crises of 2010–2012; a shock concentrated at the long end signals a repricing of the structural fiscal trajectory, with near-term repayment not in question. This is,

consequently, the horizon at which political risk feeds debt sustainability, especially for high-debt sovereigns in a high-rate environment [33], and at which uncertainty shocks are known to reshape the slope of the curve [30], [31]. The method reuses the event-window logic of the first study at every tenor simultaneously: for each tier-1 date, the average spread over the post-event window $[0, +10]$ is compared with the pre-event window $[-5, -1]$ at 2, 3, 5, 10 and 30 years; the companion diagnostic is the spread-curve slope, the 10-year minus 2-year spread, whose daily path classifies the risk type continuously rather than event by event.

Figure 13



Source: author's elaboration in Python

The left panel of Figure 13 plots the resulting term structure of the shock for the two 2024 watershed dates. The dissolution profile is the section headline: +9.4 bp at 2 years, +18.3 bp at 3 years, +20.6 bp at 5 years, +20.1 bp at 10 years and +18.6 bp at 30 years, signalling a hump peaking in the belly and staying essentially flat out to 30 years, with the short end moving less than half as much as the 5–10-year sector. The market's verdict on the dissolution was thus remarkably specific: no near-term funding scare, the 2-year, the tenor that would carry a rollover premium, was the least affected, but a durable deterioration of the medium- and long-horizon

fiscal outlook, priced almost uniformly from 5 years out. It is worth noting where INSEE's tenor sits in this profile: the 3-year, at +18.3 bp, correctly captures the magnitude of the repricing but, read alone, cannot reveal the tilt that identifies its nature. The Barnier no-confidence profile makes the complementary point: it is negative at every tenor, -6 bp to -8 bp at the short end and belly, -3 at 30 years. This is not evidence that the government's fall was good news; it is evidence of anticipation. By the pre-event window the 10-year spread already averaged 81.4 bp, its widening had occurred in the weeks of the budget standoff leading into the vote, so the event window records the partial relief that followed the resolution of uncertainty. The same signature extends to the autumn 2025 cluster, computed but not plotted: the Bayrou confidence vote shows small negative changes, -1 bp to -5 bp, and the Lecornu resignation essentially zero, ± 1 bp, the term-structure counterpart of the declining event-sensitivity already documented by the shock-detection analysis of the first study, where the later political shocks were found to be progressively pre-priced by a market that had internalized French political dysfunction as a state of the world rather than news.

The right panel turns the diagnostic into a time series. The 10y-2y spread slope, around 28 bp on average in 2019 and 33.5 bp in the pre-dissolution months of 2024, jumps visibly at the June-July 2024 event dates and averages 47.9 bp over the second half of 2024, peaking at 69.3 bp on December 12th, 2024, in the depth of the budget crisis; it then settles on a plateau, 53.3 bp on average in 2025, 50.6 bp at the sample end in June 2026, some 15-20 bp above its pre-dissolution norm, with no visible tendency to revert. The persistence matters as much as the level: a rollover scare produces a flattening or inversion of the spread curve as the short end blows out, which is the anatomy of the peripheral episodes and then unwinds; France's spread curve did the opposite, steepening on impact and staying steep. Combined with the previous section's finding that the underlying premium is idiosyncratic French, the two panels deliver the section's conclusion: what the market repriced in 2024-2026 is France's long-run fiscal trajectory as damaged by

politics, a structural premium charged at the horizon where consolidation would have to happen, not a doubt about the Treasury's next auction. This is precisely the diagnosis that the single 3-year spread of [36], sitting in the rising segment of the shock term structure, is structurally unable to deliver and it is the tenor-resolved premium that the internally consistent curve pair of this thesis was built to measure.

3.4.9 TRANSMISSION TO THE COST OF CORPORATE CREDIT

This closing block is where the case study apparently leaves the bond market. A sovereign premium, however well measured and attributed, matters for welfare only insofar as it changes the conditions at which the domestic economy borrows, which represents the very concern that motivates INSEE's interest in the spread [36] and the channel through which political risk becomes a real-economy variable rather than a market curiosity [32], [33]. Table 9 reports the reduced Stage-2 regression: the quarterly average rate on new loans to French non-financial corporations on the 3-month Euribor, the Eurosystem balance sheet and the 10-year OAT-Bund spread, over the first quarter of 2000 to the second quarter of 2026, meaning $n = 106$ quarters, with HAC standard errors.

Table 9

	Coefficient	Std. Err.	t		
Constant	1,656	0,052	31,946	R2	0,982
Euribor	0,723	0,019	38,219	n	106
Balance sheet	-0,027	0,010	-2,714		
Spread pp	0,667	0,082	8,139		

Source: author's elaboration in Python

The three regressors deserve a word each, because each carries a distinct economic channel. The dependent variable is the actual average rate signed on newly originated NFC loans, a transaction price, not a market quote and therefore the most direct available measure of what French firms pay at the margin. The Euribor captures the monetary stance: bank lending is predominantly priced off short-term money-market conditions, so this coefficient is the conventional pass-through of policy to loan rates. The balance sheet captures the quantitative arm of


policy, since asset purchases compress term and risk premia and ease bank funding beyond what the short rate records. The spread, converted to percentage points, is the study's own object: the French sovereign premium, entering separately from the monetary drivers so that, as argued in the methodology, its coefficient reads as the marginal effect of sovereign risk on corporate borrowing costs at a given monetary stance. The economic mechanism behind that coefficient is the sovereign-bank nexus: banks fund themselves and value their collateral off the sovereign curve, so a repricing of the state's credit migrates into the price of private credit even when policy rates are unchanged.

The significance triage is, for once, clean across the board: all four coefficients are significant at the 1% level. The Euribor enters at 0.723, $t \approx 38$, meaning that French firms receive roughly 72 bp of every 100 bp of money-market movement, the standard incomplete pass-through of the euro-area bank lending channel. The balance sheet enters at -0.027 , $t \approx -2.7$: unlike in Stage 1, where its effect on the Bund could not be separated from the post-2010 regime, here each trillion of Eurosystem assets is associated with NFC rates about 2.7 bp lower at given Euribor, small, yet correctly signed and now statistically identified, a trace of the credit-easing channel of QE. The headline coefficient is the spread: 0.667, with a standard error of 0.082 and a $t \approx 8.1$. Two-thirds of any movement in the French sovereign premium passes through to the rate on new corporate loans. The fit, with a $R^2 = 0.982$ is again a levels statistic and inherits the persistence caveat discussed for Stage 1, stated once and not re-litigated; more informative is the fitted-versus-observed diagnostic, which tracks the actual NFC rate within about 10 bp in recent quarters, with a systematic feature worth flagging: from 2025 the model underpredicts throughout, for example 3.52% observed against 3.40% fitted in the first quarter of 2026. A persistent positive residual is exactly what is expected given the regressor this reduced form omits, the US long yield, the third factor to which [36] attributes the sluggish transmission of ECB cuts, so the model's one visible failure independently corroborates INSEE's diagnosis.

The macro-financial interpretation closes the arithmetic of the whole study. The pass-through of 0.667, applied to the idiosyncratic French premium estimated in the decomposition section, some 30–40 bp above its pre-2022 norm since 2025, attributes on the order of +0.20 to +0.27 percentage points of French firms’ marginal borrowing costs to the political premium alone, squarely matching the $\approx$ +0.2 points that INSEE’s full Stage 2 attributes to the spread widening over 2024 [36]. What the preceding sections make it possible to say, and the official measure cannot, is the composition of that cost: it is charged on a premium that is idiosyncratically French, priced at long fiscal horizons, and persists, like the premium itself, years after the shock that created it, precisely in the configuration, high debt, high rates, weakened consolidation capacity, in which the political-risk-to-financing-cost loop binds hardest [33]. The caveats bind equally and are restated where they matter most: the frequency is quarterly and the politically driven subsample short; the attribution assumes the common and idiosyncratic components of the spread pass through alike; and the estimates are conditional associations in the sense maintained throughout the case study, so that the design measures how much firms pay when the premium moves, not a structural causal parameter.

Figure 14 closes the block by plotting, over the full 2000–2026 sample, the observed quarterly rate on new loans to French NFCs against the rate fitted by the Stage-2 regression. The observed series traverses every monetary regime of the euro history, the $\sim$ 5% world of the early 2000s, the 5.6% peak of 2008, the post-crisis collapse, the decade-long glide to the $\sim$ 1.3% floor of the QE-and-negative-rates era, the near-vertical ascent of 2022–2023 back to 4.9% and the descent from 2024 that stalls around 3.5% and the dashed fitted line tracks it closely through all of them. This breadth is the figure’s first message: the coefficients read in the table are not an artefact of the recent politically charged episode but a relationship disciplined by a quarter-century of regime variety, which is what licenses using them to price the effect of the premium at the sample end.

Figure 14



Source: author's elaboration in Python

The instructive part of the figure lies in its two visible misses. The largest is in the 2020 trough, where the observed rate plunges to 0.89% while the model floors near 1.4%: the pandemic emergency pushed actual lending below any market-consistent price through the state-guaranteed loan scheme, a fiscal intervention operating outside all three regressors and a useful reminder that the model prices market transmission, not administered credit. The second is the thin but systematic gap of 2025–2026, where the observed rate sits a few basis points above the fitted one, the omitted-US-driver signature already discussed with the table. Neither miss disturbs the reading of the final segment, which is where the figure earns its place in the case study: from mid-2024 the monetary drivers point firmly down, the Euribor falling with the ECB cuts, partially offset by the shrinking balance sheet, yet the descent of the French corporate lending rate flattens out near 3.5% and even ticks up in the last two quarters. The wedge holding it there is the spread term: with the sovereign premium parked at ~ 0.7 percentage points, the model attributes to it the difference between the floor French firms actually found and the lower one the monetary easing alone would have delivered. Seen from this figure, the political premium is not an abstraction of the bond market but a visible shelf in the borrowing costs of the corporate sector, the real-economy imprint, one loan

contract at a time, of the repricing documented across this case study and the reduced-form confirmation of the very transmission failure that titles the INSEE Focus [36].

3.4.10 ROBUSTNESS CHECKS

Three checks guard the study's conclusions and the first is the one with teeth. The entire construction rests on a benchmark choice, the thesis's own fitted Bund curve as the safe leg of the spread, so the decomposition is re-run replacing it wholesale with the ECB's official AAA euro-area curve: an externally estimated safe benchmark, produced by a different institution on a different bond universe, sharing only the Svensson functional family [27]. The result barely moves: at 10 years the idiosyncratic share of daily spread variance goes from 81.5% under the Bund benchmark to $\approx 85\%$ under the AAA curve and at 5 years from 84.3% to 85.6%, if anything marginally higher, since the AAA aggregate is a slightly less demanding benchmark than pure Germany. The premium is a property of the OAT, not of the ruler used to measure it. The two remaining checks are stated for completeness: the quarterly averaging that the transmission block requires to match the MIR frequency leaves the daily decomposition untouched by construction; the causal reticence maintained throughout, every estimate in this study is a *ceteris paribus* association, in the same spirit as [36], is not a robustness deficiency but a design choice, matched to what event windows and reduced-form regression can honestly deliver.

3.5 CONCLUSIONS

Assembled, the two studies of this chapter answer the question posed at its opening, whether the 2024–2026 political instability demoted the OAT from pure core asset to something in between and the answer arrives with an unusual degree of internal agreement. The first study established the fact of the repricing: a safety premium that roughly doubled around the dissolution, moving on the political calendar and surviving a battery of event-study designs, while the OAT beta to the

Bund never left one, a repricing in level, not a decoupling. This second study established its nature, in four steps that each sharpened the previous one. The spread measured by the curves of this thesis is externally certified: an independent pipeline reproduces the official INSEE widening to within a basis point. The premium is idiosyncratically French: common-euro factors, namely peripheral credit repricing, systemic stress and the safe benchmark itself, are significant with textbook signs at every tenor, yet they account for at most a fifth of daily spread variation and over 2024–2026 the common component eased while the French premium ratcheted from a two-decade norm of zero to roughly 40 bp, until the majority of what France pays over Germany was country-specific. The premium is structural, not a funding scare: the dissolution moved the 5–30-year spread twice as much as the 2-year and the spread-curve slope settled 15–20 bp steeper with no reversion, the market repriced the long-run fiscal trajectory, at exactly the horizons where the political-risk-and-debt-sustainability loop binds [33], not the Treasury’s next auction. And the premium reaches the real economy: two-thirds of it passes through to the rate on new corporate loans, putting on the order of +0.2 percentage points of French firms’ borrowing costs on account of political dysfunction, the reduced-form twin of INSEE’s own attribution [36].

The verdict on the opening question is therefore quantified rather than merely asserted: over the extended sample the OAT behaved as a middle-ground asset, still co-moving one-for-one with the Bund, still nowhere near peripheral pricing, but carrying a persistent, idiosyncratic, long-horizon premium that the small and stable spread of the 1987–2018 sample of [1] never showed and that two years of calm have not eroded. The limitations are those declared at each step – associational designs, a residual that bundles politics with liquidity and noise, an omitted US driver whose footprint the Stage-2 diagnostics made visible – and they mark the boundary of what this chapter claims. What the premium persistence means in the longer history of the French curve is taken up from here by the complementary angles of the thesis: the factor-structure perspective of the next

chapter and the formal regime-dating of Section 5.7.2, which classifies the post-2024 idiosyncratic premium constructed here as a distinct high-premium state, the statistical seal on the story this chapter has told event by event, tenor by tenor and loan by loan.

4. CASE STUDY 2 – PCA OF THE OAT ACROSS REGIMES

The [1] paper decomposes the OAT curve with a Principal-Component Analysis (PCA) of the levels of the 1–10-year yields and finds that the first component, namely the level, explains $\approx 97.55\%$ of the variance over 1999–2018, while the other two, slope and curvature, account for the small remainder. This is the canonical PCA reading of a yield curve [37]; however, the extension of the sample to 2026 adds an episode the original paper could not observe: the 2022–2024 return of inflation and the most aggressive tightening cycle in the ECB’s history, which drove the curve into its first sustained inversion of the euro era.

The objective of this case study is to test a single, sharply stated hypothesis, meaning that the slope factor (PC2) acquired materially greater importance during the 2022–2024 regime. The result is not a foregone conclusion, because, as will be shown below, the answer depends entirely on the statistical object decomposed and because the decomposition itself must be certified as economically meaningful before its variance shares can be trusted.

4.1 METHODOLOGY

The methodological core of the study is the distinction between running PCA on yield levels and on yield changes and it is the source of the study’s first and most important finding.

Yield levels at adjacent maturities share a strong common stochastic trend: over any sample the whole curve drifts up and down together, so the first principal component, a near-uniform loading across maturities, trivially absorbs almost all the variance and the role of the slope is mechanically hidden. Factor analysis of a term structure is therefore properly conducted on changes, where the common trend is differenced out and the genuine co-movement structure is exposed [37]. Concretely, the module builds the tenor block, optionally differences it, daily or weekly, demeanes it, forms the sample covariance matrix and takes its eigendecomposition; the ordered eigenvalues give the explained-variance shares

and the eigenvectors give the loadings. The same covariance-PCA-on-changes convention is used in applied euro-curve work [38]; it is a different statistical object from the PCA of standardized monetary policy surprises used elsewhere in the literature [39], a distinction the study is careful to preserve.


The empirical contrast is stark and justifies the methodology, as is shown in Table 4, reported here once more:

Table 4

Sample	PC2 levels	PC2 changes	Years
Convergence (1999–2007)	15,14%	7,32%	9
GFC & sovereign crisis (2008–2012)	4,49%	7,37%	5
Low or negative rates (2013–2021)	1,21%	7,95%	9
Inflation & ECB tightening (2022–2024)	1,70%	23,57%	3
Normalization (2025–today)	8,57%	10,44%	1,4

Source: author’s elaboration in Python

Figure 15



Source: author’s elaboration in Python

On levels the slope share in 2022–2024 is a negligible 1.7%, meaning that the hypothesis is invisible, because the massive common upward repricing of 2022–2024 makes the level factor dominate even more. On changes the slope share

triples to 23.6%, against a stable 7–8% in every other quiescent regime and the level factor correspondingly falls from $\approx 89\%$ to 73.2% of the variance. Reading the PC2 on changes confirms the hypothesis, while PC2 on levels misses it entirely.

In order to have a clear visual representation, Figure 15 pairs, regime by regime, the slope factor importance under two statistical objects: on changes, in red and on levels, in blue. The chart is the visual proof of both the study’s hypothesis and its methodological premise. Yield levels at neighboring maturities share a strong common trend, so a levels-PCA is dominated by the level factor and its PC2 mostly reflects persistent slope differences rather than slope dynamics, which is the main reason why the massive common upward repricing of 2022–2024 pushes the levels-slope share down to $\approx 2\%$, burying the very effect this study is testing and why the correct object is changes [37]. Read on changes, the economic interpretation is clear: in normal regimes the curve moves largely as a block, in which the daily variation is $\approx 90\%$ parallel level shifts and only $\approx 7\text{--}8\%$ slope, whereas in 2022–2024 the ECB’s tightening drove the policy-sensitive front end independently of the long end, so non-parallel steepening and flattening became a first-order driver, absorbing nearly a quarter of daily curve variation.

4.2 LORD AND PELSSER SIGN-CHANGE CERTIFICATION

A variance share is only meaningful if the component it belongs to is genuinely a level, slope or curvature factor rather than a statistical mixture. The study certifies this with the sign-change test of Lord and Pelsser (2007) [23]: the level-slope-curvature interpretation is valid if and only if the first three loading vectors have, respectively, 0, 1 and 2 sign changes across maturities; a flat level factor never changes sign, a monotone slope factor changes sign once, so the short end opposes the long end and a hump-shaped curvature factor changes sign twice [23]. This is the study’s quality measure: not a measure of fit, but a validity certificate for the economic labels. The results are shown in Table 10.

Table 10

Sample	PC1	PC2	PC3	Years	Lord signs	Lord validity
Convergence (1999–2007)	90,81%	7,32%	1,68%	9	[1, 0, 2]	False
GFC & sovereign crisis (2008–2012)	89,45%	7,37%	2,68%	5	[0, 1, 2]	True
Low or negative rates (2013–2021)	88,66%	7,95%	3,09%	9	[0, 1, 2]	True
Inflation & ECB tightening (2022–2024)	73,17%	23,57%	2,73%	3	[0, 1, 2]	True
Normalization (2025–today)	86,51%	10,44%	2,88%	1,4	[0, 1, 2]	True

Source: author's elaboration in Python

The test is therefore exact and transparent: it counts the sign flips along a loading vector and then returns the triple of counts together with the Boolean verdict. The certification results are decisive for reliability: the two most relevant comparison regimes pass; both low and negative rates, between 2013 and 2021 and, crucially, inflation and ECB tightening in 2022–2024 return the loading sign signature with, respectively, 0, 1 and 2 changes in PC1, PC2 and PC3. It stands out as a certified valid result, meaning that the 23.6% slope share is a genuine slope result, not an artefact. The full multi-regime change PCA, including the whole sample starting in 1999 until 2026, fails, with sign changes being 2, 1 and 2: the first component itself crosses zero twice, so on the pooled sample PC1 is no longer a clean level factor. This failure is not a nuisance but the study's justification. Because a single global PCA is not economically interpretable across such heterogeneous regimes, the regime-by-regime design is the statistically correct approach rather than an arbitrary choice [23].

4.3 THE INVERSION

The rise of the slope factor is the statistical shadow of an economic event: the inversion of the curve, defined here as the 10-year yield falling below the 2-year.

Table 11


Sample	Slope mean	Slope min	Days inverted	Slope volatility
Convergence (1999–2007)	101,2	0,1	-	7,29
GFC & sovereign crisis (2008–2012)	186,0	1,8	-	3,89
Low or negative rates (2013–2021)	110,9	30,2	-	2,78
Inflation & ECB tightening (2022–2024)	34,9	-33,3	29,80%	3,96
Normalization (2025–today)	119,5	85,9	-	2,57

Source: author's elaboration in Python

The per-regime diagnostics isolate it cleanly: 2022–2024 is the only regime with a sustained inversion in the entire euro-era OAT sample, inverted on $\approx 30\%$ of trading days, reaching -33 bp, whereas every other regime is strictly upward-sloping. By 2025 the curve had re-steepened to a $+120$ bp mean. An inverted yield curve has historically been one of the most reliable leading indicators of recession, since the slope forecasts future GDP growth [40]; both euro-area analyses in the case-study literature open on exactly this point [41], [26]. All the results are effectively shown in Table 11.

The 2022–2024 euro-area inversion was the deepest in decades, when the 10y minus the 2y Overnight Index Swap (OIS) spread reached ≈ -0.8 pp, its lowest since 1992 [41], and the German government-bond slope reached -37 bp in November 2023 [26], entirely consistent with Figure 16.

Figure 163



Source: author's elaboration in Python

Figure 16 shows that across the entire euro era the French slope is positive and follows a cyclical pattern. The slope diagnostics also let one attribute why the curve took each shape: between 1999 and 2007 it was a normal, positively sloped curve with a healthy term premium without inversion. The GFC, between 2008–2012, is the steepest regime, mainly because the ECB's emergency rate cuts collapsed the short end while long yields stayed high on crisis and sovereign-risk premia, signalling crisis-driven steepening, not an inversion. Between 2013 and 2021 the curve is persistently positive, characterized by very low volatility: Negative Interest Rate Policy (NIRP) and Quantitative Easing (QE), two unconventional monetary policies used by central banks when standard measures, like cutting interest rates, are exhausted, pinned the front end near or below zero while long rates remained positive. In the period on which the case study focuses, between 2022 and 2024, the inversion is clearly visible. This happens when the 2-year yield rises above the 10-year yield, which means that the market is pricing

short rates that are high now but expected to fall; this is a classic configuration that has historically preceded recessions, since the slope forecasts future GDP growth [40]. The ECB's fastest-ever hiking cycle drove the policy-sensitive front end up far faster than long rates, flattening and then inverting the curve. Critically, the two ECB and Banque de France studies stress that this negative slope was not a clean recession signal: it was driven mainly by the real curve rather than by inflation compensation and the ECB Asset Purchase Programme (APP) and Pandemic Emergency Purchase Programme (PEPP) asset stocks compressed long-term rates, so the risk-free slope "would have stayed positive longer absent QE" [41]; accordingly the inversion overstates recession risk, a statistical model conditioning on the November-2023 curve put the one-year-ahead recession probability at only 8%, against the 22% on the naive curve, once financial conditions and oil prices are accounted for [26]. Finally, between 2025 and today, it is possible to observe a re-steepening as the ECB moved past its peak and cut, the front end fell back and long yields stayed elevated on term and fiscal premia.

The 2022–2024 inversion identified here is formally dated and certified as a distinct regime in Section 5.7: a Markov-switching model assigns it a 0.83 probability of the flat and inverted state, against the 0.39 before 2022, a Bai-Perron break brackets it between November 2022 and the September-2024 re-steepening and a Ploberger-Krämer CUSUM test rejects structural stability at $p < 0.0001$ [48], [49], [50].


4.4 FACTOR SHAPES: LOADINGS AND SIGN-CHANGE RULE

The Lord and Pelsser test is made concrete by plotting the loadings themselves. For the 2022–2024 regime the change-PCA loadings are the following: PC1, level; all ten maturities load with the same sign, ≈ -0.26 to -0.34 , signalling a near-flat factor that shifts the whole curve in parallel, with zero sign changes. PC2, slope; the 1-year loads -0.95 against small positive loadings, between $+0.09$ and $+0.12$, at 2–10 years. The short end is opposite to the rest; there is exactly one sign change. Tellingly, PC2 is dominated by the front end, the econometric fingerprint of a

tightening-driven inversion in which the policy-sensitive short rate moves against a sluggish long end. PC3, the curvature; it loads positive, then negative and turns again positive across the curve. Two sign changes, the hump-shaped butterfly factor.

The signature is precisely $[0, 1, 2]$ so the loadings look like what the Lord test requires and the regime is certified valid [23]. This is the visual counterpart of the numerical certificate.

Figure 17



Source: author's elaboration in Python

Figure 17 shows the loadings; a loading is the shape of a factor: it is the amount by which each maturity yield change moves when that principal component moves by one unit. Reading a loading vector across the tenor axis therefore tells what kind of curve movement the factor represents, namely a parallel shift, a tilt or a bend. The chart plots the three loading vectors of the 2022–2024 change-PCA, so it is the visual representation of how the French curve actually moved during the inflation and tightening regime and it is the geometric complement of the sign change certificate. In particular it is possible to see that PC1, in blue, is flat; it loads

by almost the same amount on every maturity, ≈ -0.30 to ≈ -0.34 , nearly a horizontal line. This means that when PC1 moves, all yields move together by the same amount in the same direction, a parallel shift of the whole curve up or down. That is actually the definition of the level factor [37]: its flatness is its meaning. The loadings are all negative here, however the overall sign of a principal component is arbitrary, since an eigenvector is defined only up to ± 1 , so what matters is that they share one sign, giving therefore 0 sign changes. PC2, in red, is deeply negative at the start. It loads -0.95 on the 1-year and then jumps to a small, flat $+0.11$ across 2–10 years. Two things explain this shape. First, the sign is again a convention: the meaningful content is that single reversal, that is, the short end pulling one way while the long end pulls the other, that makes PC2 a slope, tilt, factor, with its one sign change. Second, and more informative, the loading is enormous at the 1-year and tiny everywhere else: the action of the slope factor is concentrated almost entirely at the front end. When the slope moves in 2022–2024, it is essentially the 1–2-year moving against a nearly static long end. Finally, PC3, the grey line, bends. It starts near zero at 1 year, dips to -0.65 in the 2–4-year belly, crosses zero around 5–6 years and rises to $+0.36$ at the long end. The middle of the curve moves opposite to the wings, a butterfly or hump, which is the curvature factor and its two reversals give the required 2 sign changes. It is the smallest factor, accounting for $\approx 3\%$ of variance, a third-order adjustment of the curve bend.

These three shapes also visually satisfy the $[0, 1, 2]$ rule by construction, which is the exact reason why the panel is certified: the level, slope and curvature labels are not assumed, they are clearly visible. The loadings decode how the curve inverted. Most day-to-day variation was still in the level factor, the whole curve rising in parallel as inflation and tightening lifted all yields, with PC1 remaining at $\approx 73\%$. But the factor that grew, the slope with a PC2 of $\approx 24\%$, is overwhelmingly a front-end phenomenon: its loading is a spike at the 1-year against an inert long end. This is the econometric fingerprint of a monetary-policy-driven slope: in 2022–2024 the policy-sensitive short rate was the active margin, hiked violently

by the ECB, while the long end, anchored by long-run expectations and by the compression from APP and PEPP asset stocks, barely moved in comparison. This reinforces once again the honest reading of Section 4.3, signalling the inversion was not a clean recession forecast [41], [26]. The curvature factor, finally, is a minor residual: a slightly belly-versus-wings twist as the market re-priced the intermediate maturities along the expected cutting path. In short, the loadings show that the rise of the slope to prominence was the ECB's policy rate acting on the front of the curve.

4.5 ROLLING PC2 SHARE AND LORD TEST FAILURE


To see the slope importance evolve continuously rather than in discrete regimes, the study runs a one-year (252 trading day) rolling change-PCA, recording at each date the PC1, PC2 and PC3 shares and a rolling Lord-validity flag. Two honest findings emerge: first, the rolling slope share is lower and noisier than Figure 18 shows, averaging $\approx 9.9\%$ across 2022–2024, so well below the $\approx 23.6\%$ obtained on the full three-year window, precisely because a one-year window contains too little slope variation to resolve the factor, the short-window unreliability that Holtes (2023) [42] warns about³. Second, for the same reason again, the rolling decomposition fails the Lord test on 16% of the sample: in those windows the level, slope and curvature reading is not valid and the rolling PC2 line must be read with caution, as seen in Figure 18. The rolling view therefore does not overturn the regime result, it contextualizes it: the trustworthy statement is the regime level $\approx 23.6\%$ over a window long enough to be reliable, not any single short rolling window.

³ Holtes (2023) [42] is a practitioner article published on a non-peer-reviewed platform. It is cited here for the methodological point it states clearly, namely that a principal-component decomposition is governed by the length of elapsed time captured rather than by the number of observations. The point is not taken on the authority of the source: it is verified independently on the data of this thesis by the rolling-window exercise of this section, in which the one-year decomposition both dilutes the 2022–2024 slope share and fails the Lord and Pelsser validity test on 16% of the sample, and by the sampling-frequency and shrinkage robustness checks of Section 4.6.

Figure 18 lets the factor shares evolve continuously in time. PC1 and PC2 are near-mirror images because they must sum, with the tiny curvature term, to one. Overlaid on top is the study's crucial safeguard: a grey shading of every date at which the rolling decomposition fails the sign-change test and its level, slope and curvature labels are therefore not to be trusted [23], [42].

PC2, the red line, stays near 10% for two decades and then spikes to $\sim 40\%$ in late 2024–2025, suggesting that the slope factor became dominant in 2025. However, this spike lies almost entirely inside a grey band, signalling that the whole of 2025 is Lord-invalid, with 100% of windows failing, so the apparent 40% slope share there is a statistical artefact of an unreliable short-window decomposition, not a real economic signal. The same is true for the beginning of the sample, the early euro-era years, between 1999 and 2002. In total, 16% of the sample fails the Lord test and it fails precisely where the rolling slope share appears highest. Conversely, over the genuine 2022–2024 regime, the rolling PC2 rises with a mean of $\approx 10\%$ on the one-year window, but far below the 24% obtained on the full three-year regime window.

Figure 18



Source: author's elaboration in Python

Figure 18 is not a confirmation of the hypothesis; it is a stress test of the study's method and its reason for honesty. It operationalizes the central lesson of Holtes (2023) [42]: a PCA decomposition is governed by the length of elapsed time captured, not the number of observations, so a one-year window is simply too short to resolve the slope factor reliably; it dilutes the true 2022–2024 effect down to $\approx 10\%$ and where it does report a large slope share, the decomposition is invalid [42].

Two conclusions follow, both essential to read the study correctly. First, the rise of the slope to prominence is a phenomenon of the regime, the multi-year build-up and unwind of the inflation and tightening episode, not of any single rolling year; it can only be measured on a window long enough to contain that whole arc, which is exactly the three-year regime block on which the 24% figure rests. Second, Figure 18 justifies the regime-by-regime design: by transparently shading 16% of the sample where its own rolling estimate is untrustworthy, it demonstrates why one must not read the slope off short windows and why the certified, long-

period regime decomposition is the right object. In a study whose headline is a claim about a factor variance share, this is the most important robustness statement.

Where the one-year rolling PCA is deliberately cautious about short windows, the structural-break toolkit of Section 5.7, applied to the slope on a longer, weekly window from 2005, supplies an independent, peer-reviewed and dated certification of the 2022–2024 regime, complementing the Lord test.

4.6 ROBUSTNESS CHECKS

The study closes with two robustness checks drawn directly from Holtes (2023) [42], whose central lesson is that a PCA decomposition is governed by the length of elapsed time, not the raw number of observations.

As regards sampling frequency, the PC2-on-changes share is re-estimated on weekly rather than daily data over the same regimes. If a result is real, it should survive the coarser sampling. It does, in direction: the 2022–2024 slope share falls from 23.6% on daily estimation to 12.25% on weekly, however remaining the highest of any regime; weekly shares range from 5% to 9% elsewhere. The magnitude is frequency-dependent, but the ranking is robust, signalling the 2022–2024 period as the slope-dominant regime. Consistent with Holtes (2023) [42], the module also flags any regime shorter than ≈ 2 years as untrustworthy regardless of its observation count: the 2025 normalization window, accounting for ≈ 1.4 years, is explicitly flagged and its shares are not to be trusted.

Analyzing the covariance estimator, Holtes (2023) [42] shows that the choice of it can distort loadings on short windows, so the study re-estimates the 2022–2024 decomposition with a constant-correlation Ledoit-Wolf shrinkage estimator, which Holtes (2023) [42] recommends. The optimal shrinkage intensity is tiny, with a $\delta=0.006$ and the slope share is essentially unchanged, going from 23.57% to 23.49%; this confirms that the result is not an artefact of the sample covariance matrix. All the data are reported in Table 12.

Table 12

Sample	PC2 daily	PC2 weekly	Number of days	Number of weeks	Years	Reliability
Convergence (1999–2007)	7,32%	8,55%	2.346	470	9	True
GFC & sovereign crisis (2008–2012)	7,37%	9,36%	1.304	261	5	True
Low or negative rates (2013–2021)	7,95%	5,43%	2.348	469	9	True
Inflation & ECB tightening (2022–2024)	23,57%	12,25%	781	156	3	True
Normalization (2025–today)	10,44%	14,33%	372	74	1,4	False

Source: author's elaboration in Python

4.7 CONCLUSION AND ECB INTERPRETATION

Stated honestly and with all caveats attached, the study's conclusion is threefold. Firstly, the hypothesis is confirmed, but only on the correct object: the slope factor variance share more than tripled to $\approx 23.6\%$ in 2022–2024, against $\approx 7\text{--}8\%$ elsewhere, on changes, while remaining $\approx 2\%$ on levels. The result is real but invisible unless changes are decomposed [37]. Secondly, it is certified and self-justifying: the Lord and Pelsser test validates the level, slope and curvature reading in the 2022–2024 regime, as well as in 2013–2021, while rejecting the pooled 1999–2026 decomposition, which is exactly what makes the regime-by-regime design correct [23]. Finally, it is robust but bounded: largest at daily frequency, still the highest of any regime at weekly frequency, unmoved by shrinkage and appropriately silent about the too-short 2025 window [42].

The economic interpretation must be equally careful and here the ECB and the Banque de France are explicit: the 2022–2024 slope is not a clean read of conventional-policy expectations or an unqualified recession signal. The negative nominal slope was driven mainly by the real curve, not by inflation compensation and, decisively for this thesis, the ECB asset-purchase stocks, namely APP and its pandemic counterpart PEPP, compressed long-term rates, so the measured slope embeds a large central-bank balance-sheet component; absent that QE the risk-free slope would have inverted less and later [41] and the inversion consequently overstates the true recession risk [26]. In other words, saying that the slope factor grew is a statement about the shape of the term structure jointly produced by a

monetary tightening and the legacy of QE, not a self-contained macroeconomic forecast; this reading is reinforced by the evidence that ECB policy and communication are first-order drivers of euro-area yield curve movements [43].

That last point is exactly what makes the final case study so important: the recognition that the ECB balance sheet, and the PEPP in particular, is a primary force shaping the OAT curve. Case Study 3 examines the mirror-image episode: not how PEPP compressed the slope over 2022–2024, but how its launch in March 2020 restored pricing efficiency during the pandemic market dislocation.

An additional and independent line of evidence removes the residual doubt attached to the short-window PCA: the structural-break analysis of Section 5.7 dates and certifies the 2022–2024 flat and inverted slope as a genuine regime, with a Markov-switching probability of 0.83, a Bai-Perron break in November 2022 and a CUSUM rejection of stability, so the slope result no longer rests on the principal-component decomposition alone.

5. CASE STUDY 3 – MARKET DISLOCATION

The original paper [1] introduces two diagnostics of market functioning: the MAE and the HPW noise measure and reads both as proxies for illiquidity and the availability or not of arbitrage capital, documenting that they spike during the two great dislocations of its sample: the September-2008 Lehman failure and the 2011–2012 euro-area sovereign debt crisis. A dislocation occurs when government bond prices and yields suddenly detach from fundamental economic indicators due to extreme stress, illiquidity, or panic. This creates abnormal mispricing, breaks historical risk-return relationships, and often triggers severe margin calls for leveraged investors. The theoretical warrant for this reading is the noise measure itself: in normal times abundant arbitrage capital keeps every bond glued to a smooth curve and pricing deviations are small, whereas during a crisis the shortage of arbitrage capital lets yields move freely relative to the curve, so the “noise” in prices rises and becomes informative about market-wide liquidity [51].

The extension of the sample to 2026 makes available the single largest liquidity event since the GFC, namely the March-2020 COVID-19 dislocation, which could not be observed by the paper [1]. The pandemic “dash for cash” abruptly drained liquidity even from core-government bond markets [52], [53] and the policy response was of unprecedented speed and scale [54]; in the euro area it took the form of the Pandemic Emergency Purchase Programme (PEPP), announced on March 18th, 2020, a temporary asset-purchase programme whose envelope was expanded from an initial €750 billion to €1,850 billion and, critically, whose purchases were conducted flexibly across time, asset classes and jurisdictions rather than strictly by the capital key [55].

Against this backdrop, the case study tries to answer two precise questions, namely whether the March-2020 dislocation was the most severe one of the French bond market and, if so, in which part of the curve, and whether the PEPP restored pricing efficiency and, if so, how quickly.

5.1 THE METHODOLOGICAL PREMISE: INVALIDITY OF COMPARISON

A central methodological commitment of this study is the recognition that a raw comparison of the overall MAE or noise across crises is not valid. The overall measure is an average across whatever bonds happen to be quoted and the cross-section differs materially from one crisis to the next in both the number and the maturity mix of securities. The problem is acute here because the OAT cross-section is small, of the order of 30–50 bonds on a given day, against the more than one hundred US Treasuries in the original HPW panel, so its very composition can sway the aggregate RMSE independently of any change in true liquidity [51]. Moreover, HPW themselves exclude securities with under one year to maturity and note that the short end is structurally noisier, so the aggregate is not a clean, composition-invariant liquidity gauge.

The study therefore reframes the comparison to be principled instead of ad hoc: crises are compared within a fixed maturity bin so that the number-and-maturity-mix confound is held constant and only genuine, like-for-like changes in pricing dispersion remain. To answer the second question, each episode is anchored at a fixed onset date and characterized by its peak dislocation and its recovery half-life, representing the number of trading days after the peak before the metric falls halfway back to its pre-crisis baseline, over a window long enough, generally ~one year, not to mistake the separate 2023–2025 structural floor for a failure to recover; the episodes are then overlaid in event time to compare magnitude and recovery speed directly. The half-life framing follows the recovery-dynamics literature, which formalizes liquidity half-life and time-to-normalization as the natural metrics for how quickly markets heal after a volatility shock [56].

A final robustness concern is addressed at the outset. Because the OAT cross-section is small, a bin's apparent dislocation could, in principle, be driven by a single squeezed security instead of by a broad-based stress, a live possibility in March 2020, when the collapse of the tradable “free float” made specific, scarce lines dislocate more than others. The study therefore also recomputes the belly

MAE on the peak days with and without the single worst-fit bond, a leave-one-out check, to distinguish a broad dislocation from a one-issue artefact.

5.2 RELATED LITERATURE

The study is positioned against and its results are compared with four strands of the recent literature. First, the noise-and-illiquidity framework: here the interpretive backbone is [51], whose noise-as-illiquidity measure is precisely the object this study tracks for the OAT market and the reason the per-bin refinement, due to small cross-section and excluding the short end, matters. Second, the COVID dislocation and the response of central banks: the March-2020 “dash for cash” and the Federal Reserve response are documented by [52]; the broader, cross-country central-bank reaction – fast, forceful and centred on quantitative easing – is derived from [54]; finally, the euro-area-specific stress in core government-bond markets is documented by [53]. The PEPP itself is described in the programme documentation [55] and its effect on euro-area sovereigns is quantified by [57], who show that the PEPP compressed euro-area sovereign bond risk premia, while the transmission channels of ECB policy to euro-area sovereign yields are surveyed in [58]. Third, the French free-float, collateral-scarcity mechanism is derived directly on the French market, with [59] documenting that the tradable free float of French government debt fell from 51.1% to 34.8% between 2015 and 2021, chiefly through Eurosystem holdings, that securities with a lower free float lost liquidity more sharply during the COVID crisis, that the Eurosystem securities-lending facility mitigated the resulting scarcity and that during the March-2020 stress liquidity deteriorated first in the repo market, then in the cash market and last in the futures market, while concluding that French public debt nonetheless remains one of the most liquid in Europe. This supplies the concrete mechanism behind the leave-one-out check and a third confound – the relative scarcity of specific paper – alongside cross-section size and genuine arbitrage-capital scarcity. Fourth and last, the operational metrics used to answer the “how fast” question, the liquidity half-life and time-to-normalization, follow

[56]; although their evidence is drawn from an emerging equity market, their duration-based recovery framework is the methodological template for measuring the speed at which the PEPP restored OAT pricing efficiency.

Taken together, these define the study's contribution: to bring the OATmeals fitting-error and noise diagnostics to bear on the March-2020 episode, to compare it with 2008 and 2011 in a cross-section-honest, per-bin manner and to measure, through the recovery half-life, whether and how fast the PEPP repaired the pricing efficiency of the French market.

5.3 PER-BIN SEVERITY TABLES

The severity results of this section are carried by two tables that report the same statistic at two levels of granularity. Both express, for each crisis and each maturity bin, the peak per-bin fitting error (MAE) in basis points, meaning the worst average deviation of the bonds in that segment from the fitted Svensson curve during the episode and both are constructed identically: from the daily maturity-specific fitting error [51], smoothed with a 5-day moving average and maximized over the ~one-year window that follows each crisis's fixed onset date.

They differ only in how finely the curve is partitioned and each view serves a distinct purpose: Table 13 uses three central bins, namely 2–5, 5–10 and 10–20 years, representing the core, most liquid part of the curve. Table 14 carries the identical statistic across all six bins, meaning 0–2, 2–5, 5–10, 10–20, 20–30 and 30–50 years, adding the short end and the ultra-long segment so that the full maturity signature of each crisis becomes visible.

What both tables define is a cross-section-normalized map of where and how severely each crisis dislocated pricing. Because the maturity segment is held fixed, a difference read down any column is a genuine difference in pricing dispersion rather than an artefact of the differing number and maturity mix of bonds quoted in each episode. This compact table answers the first research question directly,

while the extended table explains why the answer is nuanced, revealing that the four crises stressed entirely different parts of the curve.

Table 13

Peak per-bin fitting error (MAE, bp)			
Crisis	2–5 yr	5–10 yr	10–20 yr
GFC (2008)	5,47	5,05	5,11
Euro sovereign (2011)	3,12	3,56	6,96
COVID-19 (2020)	6,38	10,94	1,89
Inflation shock (2022)	8,58	7,82	9,57

Source: author's elaboration in Python

Table 14

Extended view (full curve – across all six bins)						
Crisis	0–2 yr	2–5 yr	5–10 yr	10–20 yr	20–30 yr	30–50 yr
GFC (2008)	8,60	5,47	5,05	5,11	8,27	10,45
Euro sovereign (2011)	8,61	3,12	3,56	6,96	6,41	3,77
COVID-19 (2020)	5,13	6,38	10,94	1,89	2,23	3,37
Inflation shock (2022)	9,15	8,58	7,82	9,57	9,24	22,95

Source: author's elaboration in Python

Each cell is the peak within-bin fitting error, in basis points, reached in a given crisis and maturity segment, meaning the worst average deviation of the bonds in that segment from the fitted Svensson curve at the height of that episode. Because the fitting error is the study's proxy for illiquidity and the scarcity of arbitrage capital, the table is effectively a map of where and how badly each crisis dislocated the French curve: read across a row, it shows how one crisis was distributed along the maturity spectrum; read down a column, it ranks the crises within a fixed segment.


Each of the four crises is anchored at a fixed, well-known onset date, namely the Lehman bankruptcy, on September 15th, 2008, the euro crisis reaching the core, on August 1st, 2011, the pandemic sell-off, on February 24th, 2020 and the outbreak of the Russo-Ukrainian war with the consequent inflation surge, on February 24th,

2022, instead of at a data-driven peak, so the measurement is not contaminated by later, unrelated humps. For every maturity bin the daily per-bin MAE, the mean absolute observed-minus-fitted yield error among the bonds in that bin, is smoothed with a 5-day moving average, so that it reflects the episode, not a single noisy day, and its maximum over the ~ 1 -year window following the onset is taken. That peak is the table entry. These two tables are the analytical heart of the study’s honesty, for two connected reasons: first, they neutralize the composition problem that makes a raw comparison invalid. The overall MAE or noise blends together bins containing different numbers and maturities of bonds and the OAT cross-section is small; a crisis could therefore appear “more severe” simply because more, or structurally noisier, securities happened to be quoted not because liquidity was genuinely worse⁴ [51]. By fixing the maturity bin, the tables compare like with like: any difference down a column is a genuine difference in pricing dispersion, not an artefact of the cross-section. Second, and this is the reason it is indispensable for understanding how and how much each event bit, the tables reveal that the four crises had entirely distinct maturity signatures; none of them hit the curve uniformly. COVID-19 is sharply belly-specific, for example: it peaks at 10.94 bp in the 5–10-year segment, roughly twice any other crisis there, yet is the mildest of all episodes at 10–20 years, with a small 1.89 bp and negligible beyond, the fingerprint of a pure “dash for cash” concentrated in the most heavily traded part of the curve [52], [53]. The Lehman bankruptcy is, by contrast, broad and roughly uniform, with almost ≈ 5 bp across 2–20 years and additional stress at the short end, 8.6 bp, and the ultra-long end, 10.5 bp, signalling a system-wide dislocation. The euro sovereign crisis, instead, concentrates at the short end, 8.6 bp and the 10–20-year segment, 6.96 bp, highlighting the maturities most exposed to redenomination and credit risk. Finally, the 2022 inflation shock is a long-end event, peaking at 22.95 bp in the 30–50-year bin, underscoring the ultra-long illiquidity already flagged in Chapter 2.

⁴ Recall that HPW exclude sub-1-year bonds and warn that the short end is intrinsically noisier.

These differing signatures are not a curiosity: they are the substance of the study. They allow the first research question to be answered precisely: March 2020 was the most severe dislocation only in the 5–10-year belly, not across the whole curve, a conclusion that a single aggregate number would have hidden entirely. And because the location of a dislocation encodes its mechanism, a belly dislocation signals a liquidity dash in the core trading segment, an ultra-long one signals duration or scarcity stress, a short-and-intermediate one signals sovereign-credit risk, the per-bin table is precisely what lets this study distinguish how each event stressed the market, not merely that it did.

Figure 19



Source: author's elaboration in Python

Figure 19 renders the compact per-bin severity Table 13 in visual form. Because the maturity bin is held fixed within every cluster, bar heights are directly comparable across crises, so the figure makes the cross-section-normalized comparison legible at a glance.

5.4 OVERALL NOISE: PEAK AND RECOVERY (ONE-YEAR WINDOW)

Having established where and how severely each crisis struck the curve, this section turns to the second research question, namely how large the dislocation was overall and, above all, how quickly the market healed and so to the diagnostic that speaks most directly to the PEPP. The instrument is the crisis-summary Table 15, which characterizes each episode's overall HPW noise by four quantities: the pre-crisis baseline, the peak, the peak-to-baseline multiple and the recovery half-life, which is the number of trading days after the peak before the noise falls halfway back to baseline. Two reading rules must be stated at the outset. First, the overall noise is computed on the whole, small OAT cross-section and is therefore not comparable across crises in level, so the peaks should be read through the per-bin lens of Section 5.3 and it is the recovery speed, a within-crisis time-series property largely insulated from the cross-section mix, that carries the weight here. Second, with a ~one-year window a missing half-life means not recovered within 250 trading days, not never.

Table 15

Overall noise – basis points					
Crisis	Baseline	Peak	Peak date	Peak / Baseline	Recovery half-life days
GFC (2008)	2,52	5,99	26/03/2009	2,4	17
Euro sovereign (2011)	1,66	5,91	16/11/2011	3,6	55
COVID-19 (2020)	1,63	11,04	05/08/2020	6,8	17
Inflation shock (2022)	2,60	12,86	25/10/2022	4,9	-

Source: author's elaboration in Python

Each row summarizes one episode's noise trajectory. COVID-19 produced the largest relative dislocation of the four, with a peak of 11.0 bp, 6.8 times its pre-crisis baseline, yet it recovered in just 17 trading days, exactly as fast as the Lehman shock and far faster than the euro sovereign crisis, whose noise took 55 days to halve. The 2022 inflation shock reached the highest absolute peak, ~12.9 bp, but has no defined recovery half-life.

The recovery-speed column is the most informative here. That the COVID-19 crisis healed as quickly as Lehman despite being the single most severe relative dislocation is precisely the signature of a forceful, decisive backstop: the ECB's PEPP, announced within days of the sell-off and deployed flexibly across jurisdictions [55], [57], compressed the noise back to normal at a pace no slower than the 2008 episode. The contrast with the euro sovereign crisis, three times slower to heal, is equally telling: that dislocation festered until a credible commitment, namely Draghi's July-2012 famous "whatever it takes" speech, arrived, whereas in 2020 the backstop was immediate. The detailed timeline of Section 5.4.1 formalizes this, however Table 15 already establishes the qualitative result: the PEPP restored the pricing efficiency of the French market rapidly, on a par with the fastest historical recovery.


The single absent value in Table 15 deserves careful reading, because it is the most instructive entry and it does not mean that 2022 was the worst or a permanently broken market. Mechanically, the recovery half-life is undefined because the noise never fell halfway back within the window: with a baseline of 2.60 bp and a peak of 12.86 bp, the half-recovery threshold is 7.73 bp, yet in the entire post-peak segment, from October 25th, 2022, to February 9th, 2023, the smoothed noise bottomed at only 9.32 bp, signalling it never even approached the threshold. Substantively, the reason is that the 2022 episode was not a transitory dislocation that spikes and reverts, but the onset of a structural regime change in the noise floor itself. Where the other three crises returned toward a stable, low normal, the noise after October 2022 settled at a permanently higher plateau, with a mean of ≈ 10.3 bp through 2023–2024, easing only to ≈ 8.4 bp in 2025 and did not cross back below the 7.73 threshold until June 2025, nearly two and a half years after the peak. The old baseline had simply become obsolete: the market had migrated to a new, structurally noisier level, consistent with the higher-rate, higher-volatility, larger-and-more-heterogeneous OAT universe documented in Chapter 2. The half-life concept, which presupposes reversion to a fixed pre-crisis baseline,

is therefore the wrong tool for a level shift and its absence for 2022 is not a failure of recovery but can be seen as a signature of a permanent regime change. This distinction between a transitory dislocation healed by policy, as happened during COVID-19 and a structural rise of the noise floor, as happened from 2022 onward, is exactly the separation that the March-2020 detail Figure 20 is built to make explicit, so that the PEPP's genuine, fast repair of the pandemic episode is not confused with the later, unrelated structural drift.

5.4.1 DISLOCATION AND RECOVERY IN EVENT TIME

The summary Table 15 reduced each episode to four numbers; this section restores the shape behind them. By re-aligning all four crises on a common horizontal axis, which shows the trading days measured from each onset, Figure 20 overlays their complete noise trajectories, so that the profile of dislocation and recovery, not merely its peak and half-life, can be compared at a glance. As in the preceding section, the vertical levels remain cross-section-dependent and are not the object of comparison; what the alignment makes legible is the geometry of each episode, so how fast it arrived and how it left.

Figure 20



Source: author's elaboration in Python

The overlay separates the episodes into four recognizably different shapes: the euro sovereign crisis, in grey, traces a textbook profile of a transitory dislocation: it rises within days of onset, holds an elevated plateau and then completes a clean round-trip back to its pre-crisis baseline by the end of the year, marking the only episode of the four to fully close the loop within the window. Lehman, in blue, is a low, broad and persistent hump: it never spikes sharply, drifting up to a moderate plateau and remaining mildly elevated for the entire year, the trace of the financial crisis. The 2022 inflation shock, in green, is the mirror image of the sovereign crisis: a slow, relentless, monotone climb that never turns down, the dislocation is seen as a one-way regime shift rather than a spike to be recovered from. COVID-19 is by far the most violent onset: from its pre-crisis level the noise rises almost sixfold within roughly three trading weeks, making it the steepest, fastest arrival of any of the four, and signalling the instantaneous “dash for cash” [52], [53]. It then oscillates on the highest, most turbulent plateau of the sample; its acute component visibly heals, which is the sharp mid-episode decline, but the trajectory re-elevates in the tail and does not return to its old baseline within the year, because the structural floor is already rising beneath it.

Figure 20’s contribution is to render visible a distinction the summary numbers could only imply: it cleanly separates transitory dislocations that revert from structural drifts that do not. The sovereign crisis and, in its acute component, the COVID-19 episode belong to the first family, tracing out a rise-and-return; the 2022 episode and the COVID-19 persistent tail belong to the second, drifting to a new higher level and staying there. Two readings follow. First, the onset geometry justifies the premium the study places on the speed of the policy response: the COVID-19 dislocation arrived faster and steeper than any other, so the value of a backstop deployed within days rather than months is precisely what a market with this onset profile required. Second, the overlay shows why a single trajectory cannot be read off its endpoint: the elevated one-year tail of COVID-19 is not evidence that the pandemic dislocation went unrepaired, but the superposition of a

healed acute episode and a separate, rising structural floor, two mechanisms the aggregate curve entangles.

5.4.2 THE ANATOMY OF THE TWO RECENT DISLOCATIONS

The graphical review of Figure 20 exposed two trajectories that the aggregate noise, on its own, cannot explain: the double-peaked COVID profile and the monotone non-reverting climbing of 2022. This study wants to go beneath the two curves, drawing on the recent collateral and intermediary-balance-sheet literature to give a mechanism-based and deliberately honest account of why each assumed its shape. Throughout, these are interpretations of the observed OAT noise grounded in the cited theory, not direct measurements of repo or dealer variables in the data.

First, COVID-19: an exogenous first peak, a policy-induced second. The pandemic trajectory is best read as two peaks of different origins separated by a policy intervention. The first peak is exogenous: the March-2020 “dash for cash”, in which investors liquidated even the safest assets for liquidity, arbitrage capital withdrew and the fitting dispersion spiked, the classic HPW signature of a collapse in market functioning principally caused by the absence of arbitrage capital [51], [52], [53], [60]. The PEPP, announced on March 18th, 2020, and deployed flexibly, then compressed this dislocation, restoring pricing efficiency and compressing euro-area risk premia [55], [57], the sharp mid-episode healing visible in Figure 20. The second peak, however, is endogenous, a by-product of the cure itself. By absorbing large quantities of specific eurozone government bonds, the ECB’s asset purchases removed high-quality liquid assets (HQLA) from circulation and drove collateralized borrowing rates to record lows, generating repo specialness, the wedge by which a scarce, sought-after security lends in the special-collateral (SC) segment below the general-collateral (GC) rate [61], [62]. This matters for this noise measure because, in the preferred-habitat theory of [61], bonds are simultaneously investment assets and collateral: arbitrageurs, who must borrow specific bonds in the repo market to short them, present a price-inelastic demand

for collateral, so when purchased bonds are withheld from repo their specialness feeds directly into the market price of interest-rate risk and the pricing of the entire yield curve, precisely the cross-sectional dispersion the fitting error captures. The mitigant, and the reason the second peak did not spiral, was the Eurosystem's securities-lending facilities, which lent the purchased bonds back into the repo market and relieved the scarcity [63], [64], the same facilities whose easing of low-free-float scarcity the Banque de France documents for the French market specifically [59]. The COVID-19 double peak is thus not one dislocation that failed to heal, but an exogenous liquidity shock the PEPP repaired, followed by an endogenous collateral-scarcity shock the PEPP created and the securities-lending facilities contained.


Second, the 2022 inflation shock was a structural regime change, not a dislocation. The 2022 profile, a slow, relentless climb that never turns down, is of a categorically different nature: it is the onset of a macro-financial regime change rather than an acute dislocation. [65] frames the period as a set of simultaneous regime shifts relative to the post-GFC world: supply-side shocks and inflation supersede chronic demand weakness, and the end of abundant and cheap central-bank liquidity gives way to higher rates and liquidity squeezes. The mechanism that translates this into a permanently higher noise floor is the balance-sheet constraint of the dealers who intermediate the market. [66] document a regime change in dealer positioning around the GFC: dealers moved from a net-short to a net-long government-bond position as bond supply grew five-fold while dealer balance sheets stagnated, so that post-GFC the market operates against a binding leverage constraint, the supplementary leverage ratio (SLR), whose shadow cost is visible in negative swap spreads and sizeable covered-interest-parity (CIP) deviations [66]. In 2022 this constraint bit hard: as the ECB shifted from purchases to Quantitative Tightening (QT), the flow reversed, bonds returned to a dealer sector that is net long and balance-sheet-constrained, so intermediation capacity, not a transient panic, became the binding scarcity. Because a leverage constraint is

a level feature of the post-GFC financial system rather than a shock that dissipates, the resulting pricing dispersion does not revert to the old baseline, exactly the absent recovery half-life of Figure 20. [67] and market strategists frame the same shift as structural: sovereign markets increasingly intermediated by leveraged, funding-dependent non-banks, in which “liquidity can be ample for prolonged periods yet dry up quickly” and 2026 as the entry into a regime “shaped less by cyclical recovery and more by structural change” [67], [68].

5.5 THE COVID EPISODE AND THE PEPP

Having established the mechanism behind the two recent shapes, the study now pins down the chronology of the pandemic episode and, in doing so, guards against a specific and tempting misreading. A glance at the noise series shows it elevated more or less continuously from 2020 to 2026, from which one might wrongly conclude that the French market never recovered from COVID, or that the PEPP failed. The purpose here is to separate two phenomena that the aggregate series superimposes, the transitory pandemic dislocation that the PEPP repaired and the structural rise of the noise floor that followed years later for unrelated reasons and thereby to answer the study’s second question cleanly: whether and how quickly the PEPP restored pricing efficiency.

Figure 21



Source: author's elaboration in Python

Figure 21 is deliberately built in two panels. The left panel zooms into the pandemic episode itself, from February 2020 to June 2021 and plots the daily MAE and noise together, annotated with the two dates that matter, namely the PEPP announcement, which took place on March 18th, 2020, and the noise peak. Plotting both diagnostics side by side is not redundant: the divergence between the MAE and the RMSE is itself informative, as noted below. The right panel, instead, widens the lens to the full 2018–2026 noise series, marking the pre-COVID median, in the period between 2015 and 2019, as a horizontal reference and shading the 2023–2025 window, so that the pandemic spike and the later structural floor can be seen as the distinct objects they are.

The left panel resolves the episode into a sharp acute phase and a later, milder endogenous one. The noise spikes to ≈ 10 bp on March 11th, 2020, the so-called “dash for cash” and is already easing, ≈ 5 bp, by the time the PEPP is announced on March 18th; it then rises again through the summer to its true maximum, ≈ 10.8 bp on August 5th, 2020, before falling away and normalizing completely through 2021: the half-yearly mean noise declines from 7.8 bp in 2020 to 6.6 bp in 2021 and to just 2.3 bp in the following period, so below the pre-COVID median of 2.38 bp, reaching a low of 1.9 bp in August 2021. A second feature is visible in the gap

between the two lines: at the August peak the noise reaches ≈ 11 bp while the MAE never exceeds ≈ 4 bp, a divergence far wider than in the broad March phase. The right panel then shows the sequel: after the 2021 return to normal, the noise climbs again from 2022 and settles, across the shaded 2023–2025 window, at a mean of ≈ 9.7 bp, a new plateau roughly three and a half times the pre-COVID median that persists to the end of the sample.

Three conclusions follow and together they constitute the “correction”. First, on the verdict for the study’s second question, it is shown that the PEPP restored the pricing efficiency of the French market both fully and fast. The acute March dislocation and the endogenous August peak both healed and by late 2021 the noise had returned beneath its pre-pandemic level, a complete round-trip, consistent with the euro-area evidence that the programme compressed sovereign risk premia and repaired market functioning [55], [57]. Second, the MAE-versus-noise divergence at the August peak is diagnostic of the mechanism: because the RMS statistic weights large deviations far more heavily than the mean absolute error, a peak in which the noise dwarfs the MAE signals that the dislocation was carried by a handful of securities with large individual errors rather than by a broad repricing, precisely the fingerprint of the specific-collateral scarcity induced by the ECB’s own purchases [61], [63] and the motivation for the leave-one-out test that follows. Third and most importantly, the two panels refute the false inference: the elevated noise of 2023–2025 is not persistence of the pandemic dislocation nor evidence against the PEPP, because the COVID episode had already fully normalized by 2021; it is instead the distinct, later structural regime shift documented mechanistically in the sections before, namely the interaction of QT with a balance-sheet-constrained, net-long dealer sector, in a larger and more heterogeneous post-pandemic sovereign market [66], [65], [67]. Conflating the two would commit a double error, understating the speed and completeness of the PEPP repair and misattributing a structural intermediation shift to a transient health

shock. The corrected timeline keeps them apart: a policy that worked, on an acute episode that healed, followed by a structural drift that is a separate story altogether.

This separation is placed on a formal, dated footing in Section 5.7 and following: a Bai-Perron analysis dates the COVID break on March 6th, 2020, a return to a calm ≈ 2.9 bp segment on May 28th, 2021 and a fresh structural break into the ≈ 9.6 bp floor on July 1st, 2022, while the Markov classification shades the two crises as distinct blocks separated by the intervening calm.

5.6 ROBUSTNESS: LEAVE-ONE-OUT

The per-bin result of Section 5.3, represented in Figure 19, located the severity of COVID-19 in the 5–10-year belly. Given the small OAT cross-section and the collapse of the tradable free float, a natural objection is that this belly peak might be an artefact of a single low-free-float security that became impossible to source in repo, rather than a broad market dislocation. This section tests the objection directly.

The method used is the following: on each crisis belly-peak day, the 5–10-year MAE is recomputed with and without the single worst-fitting bond; the difference, the worst-bond contribution, measures how much of the bin dislocation rests on one security. The measure is conservative by construction: [51] already exclude bonds more than four standard deviations from the curve, so one or two securities cannot drive the entire measure; the daily fit of this study applies a stronger robust screen by dropping any bond with an error above 75 bp, so any surviving single-issue component lies within the economically plausible range [51].

Table 161

Crisis dislocation Leave-one-out					
Crisis	Bonds in bin	MAE, full (bp)	MAE, drop worst (bp)	Worst-bond contribution	Share
COVID acute (23 Mar 2020)	14	5,13	2,37	2,76	54%
COVID peak (5 Aug 2020)	14	10,56	7,34	3,22	30%
GFC (13 Nov 2008)	10	3,11	2,77	0,34	11%
Inflation shock (25 Oct 2022)	11	6,53	5,83	0,70	11%

Source: author's elaboration in Python

The answer is both and is reported in Table 16: the dislocation is genuine and broad, but uniquely for COVID it is also scarcity-amplified. On one hand, the COVID belly dislocation is not an artefact of one bond: even after removing the single worst security, the August-peak belly MAE remains 7.34 bp, still above the full MAE of any other crisis, so the stress is real and shared across the segment. On the other hand, COVID carries a sizeable single-issue component that the other crises do not: the worst bond accounts for 54% of the bin MAE in the acute March phase and 30% at the August peak, against only ~11%, below 1 bp, for both GFC and the 2022 inflation shock, which are broad-based. Most tellingly, the security responsible is the same specific line on both COVID dates, a ~5-year OAT whose individual error reached 41 bp in March and 52 bp in August, exactly the concrete fingerprint of a single, scarce, low-free-float bond that could not be sourced in the repo market. This is the microstructural counterpart of the mechanism set out in Section 5.4.2: the Eurosystem purchases withdrew specific high-quality collateral, that scarcity manifested as repo specialness and one squeezed line dislocated far from the curve [61]. It matches the Banque de France's finding that the falling free float, from 51.1% to 34.8% since 2015, via Eurosystem holdings, made low-free-float securities dislocate more in March 2020, mitigated by the securities-lending facilities [59] and the FSB euro-area "dash for collateral" [53].

Two honest caveats complete the reading. First, the specific line driving the COVID single-issue component is one that later exhibited data-quality problems in the Bloomberg feed, removed elsewhere by the plausibility screen of Chapter 2;

its COVID-date errors are within the plausible range and are therefore retained, but part of the single-issue component could, in principle, be an early data artefact rather than pure scarcity, which is precisely why the leave-one-out matters, since the COVID belly dislocation is confirmed genuine and highest independently of that security. Second and as a counterweight to the free-float narrative, the FSB notes that French primary markets stayed resilient in March 2020, with large OAT and BTF issuance absorbed smoothly, so the long-run free-float decline is not the whole story for the acute shock [53]. The robust conclusion is therefore layered: COVID was the most severe belly dislocation of the sample, genuinely and broadly, but it uniquely carried a specific-collateral-scarcity component.

The structural, microstructure reading of the post-2022 floor is independently corroborated in the next sections: the noise regime is essentially uncorrelated with the ECB CISS, with a $\rho \approx -0.06$ and an agreement of 46%, so the floor is not a systemic-stress episode but the intermediation and collateral phenomenon identified here [66], [61].

5.7 FORMAL STRUCTURAL-BREAK ANALYSIS

Case Study 2 and Case Study 3 established descriptively that the OAT pricing-noise floor stepped up after 2022 and that the curve inverted in 2022–2024; this section puts dates and statistical confidence on those shifts, using peer-reviewed structural-change methods applied to the estimated price and yield series of this thesis. The same toolkit is deployed on three weekly series so that the three post-2018 episodes are dated by one consistent, certified procedure: the HPW noise, the 10-year-minus-2-year slope and the idiosyncratic OAT-Bund premium. All series are sampled weekly from 2005, which represents the modern-regime window, the volatile 1999–2003 early-euro period being a separate transition; no curve is re-estimated. Three complementary methods are therefore used.

5.7.1 THE ECONOMETRIC METHODS

The first tool, namely the Markov-switching model (Hamilton, 1989), treats each series as generated by one of a small number of unobserved regimes, each with its own mean and variance, between which the process switches according to a Markov transition matrix [48]. Fitting the model by maximum likelihood, with switching mean and variance and multiple random restarts to guard against local optima, recovers three objects of interest: the mean and volatility of each latent regime, the transition probabilities and hence each regime's expected duration and, most usefully, the smoothed probability that every individual week belongs to each regime, computed conditional on the whole sample. The number of regimes is chosen by the Bayesian Information Criterion (BIC), which selects a three-state structure, namely calm, elevated and crisis, for each series; a two-state model is rejected. This converts the qualitative claim of a rise in noise into a quantitative statement: a latent crisis regime with an estimated mean and duration, to which the model assigns specific weeks with a computed probability.

Where the Markov model classifies states, the Bai-Perron multiple structural break procedure (Bai and Perron, 2003) dates the shifts. This second tool detects an unknown number of level shifts, mean breaks, in a series and locates them optimally, partitioning the sample into homogeneous segments that minimize the within-segment sum of squares; the exact solution is obtained by dynamic programming and the number of breaks is selected by BIC [49]. Applied to each weekly series, it returns the calendar dates of the regime shifts and the mean of each segment, so that, for example, the step-up of the noise floor acquires a specific break date and a quantified segment level. The computation uses [69].

The third tool is a formal hypothesis test rather than an estimator and it provides an independent safeguard against reading breaks into a series that is in fact stable. The CUSUM-of-OLS-residual stability test (Ploberger and Krämer, 1992) regresses the series on a constant, cumulates the residuals and rejects the null of structural stability when that cumulative sum wanders implausibly far; a small p-

value therefore confirms, model-free, that a genuine structural break exists [50], corroborating the Markov and Bai-Perron findings without assuming their parametric structure.

Finally, to avoid confirming a regime with a single method on a single series, the dated regimes are cross-checked against two independent information sources, an approach adopted in the spirit of Yi, Mehra, Chen & Cartlidge (2026), who augment regime detection on the US Treasury market with an extra data source [70]. Their specific pipeline, text-based, on US data and a non-peer-reviewed preprint, is deliberately not reproduced; instead, the study adopts only its idea, using signals it already possesses: the ECB's CISS, a market-derived stress index whose own high-stress regime, a two-state Markov model, is compared with the noise regime and the hand-coded political-event calendar, against which the premium regime is aligned. Whether the price-based regimes agree or disagree with these independent signals is itself informative; an agreement confirms the reading, while a pointed disagreement discriminates between a structural and microstructural interpretation and a systemic-risk one.

Taken together, the three methods are used to answer distinct questions about each series, namely whether there is a break (CUSUM), if so when (Bai-Perron) and what the regimes and their probabilities and durations are (Markov-switching). The corroboration step tests each dated regime against evidence outside the price series itself. This is the honest, reproducible foundation on which the next section will place dates and confidence on the noise, slope and premium shifts that the descriptive analysis of the three case studies identified.

5.7.2 DATED RESULTS AND CORROBORATION

Applying the toolkit of Section 5.7.1 the models return the following results, reported series by series and then cross-checked against independent evidence.

Concerning Case Study 3, the Markov-switching model, with the regime count selected by BIC, which decisively prefers three states, with BIC 3085 at $k=3$ against 3873 at $k=2$, identifies a calm, ≈ 1.6 bp, an elevated ≈ 3.6 bp and a crisis ≈ 9.1 bp regime, the last being the most persistent, with an expected duration of ≈ 55 weeks and a 26% unconditional share of the 2005–2026 sample [48]. The Bai-Perron segmentation then dates the level shifts and its verdict is the quantitative capstone of Section 5.5: it places a break on March 6th, 2020, the COVID segment having a mean ≈ 7.8 bp, a return to a calm 2.9 bp segment from May 28th, 2021 and a separate break on July 1st, 2022, into the ≈ 9.6 bp structural floor segment that persists to the end of the sample [49]. Results are reported in Table 17.

Table 17

Regime	Mean (bp)	Volatility (bp)	Average duration (weeks)	Time share
Calm	1,64	0,54	41,9	40,7%
Elevated	3,56	0,83	25,3	33,4%
Crisis	9,06	1,64	54,6	25,9%

Source: author's elaboration in Python

Table 17 summarizes the three latent regimes that the Markov-switching model extracts from the weekly HPW noise series over 2005–2026, the model characterization of the states the OAT market moves between, as distinct from the timing of those states. Each row is one regime, ordered from the lowest- to the highest-noise state and the four columns report, respectively, the regime average noise level, its internal volatility, how long it typically lasts once entered and how often the market is found in it. Read together they say that the French market's pricing efficiency is not a single condition but a small set of well-separated states: a calm regime near 1.6 bp, an elevated one near 3.6 bp and a crisis one near 9.1 bp.

The four quantities are all by-products of the same maximum-likelihood fit [48]: the mean is the probability-weighted average noise over the weeks the model assigns to the regime, the typical level of cross-sectional pricing dispersion that

defines it; the volatility is the within-regime standard deviation of the noise. Because the model was estimated with switching variance and not only with switching mean, each regime has its own dispersion and this column measures how tightly the noise clusters around the regime mean; the average duration is the expected persistence of the regime once entered, computed from the estimated Markov transition matrix as $D_i = \frac{1}{1-p_{ii}}$, where p_{ii} is the probability the regime transitions into itself and D_i is the expected duration, the persistence of regime i , measured in weeks, since the series is sampled weekly. It is useful because it quantifies the number of weeks the market typically remains in a state once it is entered. Lastly, time share is the unconditional fraction of the sample spent in the regime, the average of its smoothed probability, meaning how common the state is over the whole 2005–2026 window.

Because the noise is a proxy for the availability of arbitrage capital and the tightness of pricing [51], the three regimes are three states of market functioning and each column carries a powerful economic message.

First, on amplitude: the crisis regime mean, 9.1 bp, is roughly five and a half times the calm regime, 1.6 bp. The calm state is the frictionless world the Svensson benchmark presumes, abundant arbitrage capital keeps every OAT glued to a single discount curve; the crisis state is its breakdown: capital scarce, individual lines drifting rich or cheap, the market pricing itself far less coherently. The elevated middle state is a regime of moderate frictions.

Second, and most striking, is persistence. The crisis regime is not the most fleeting but the most durable of the three: an expected ≈ 55 weeks, longer even than the calm regime ≈ 42 weeks, while the elevated regime is the shortest lived ≈ 25 weeks. This has a clear macro-financial reading: the noise process is effectively bimodal, with the calm and crisis states each self-sustaining and the elevated state serving as the transition the market passes through on its way between them. Economically, it means that once the pricing efficiency of the OAT


market breaks down, it stays broken for more than a year on average, meaning that dislocations in this market are sticky, not transient blips. That stickiness is exactly what one expects when a dislocation has a structural rather than a purely cyclical cause, and it is the regime-model-independent corroboration of the persistence that defined the post-2022 floor: a high-noise state that, once entered in mid-2022, was self-perpetuating.

Third, the volatility column reveals heteroskedasticity: internal volatility rises with the level, starting from 0.5, then 0.8 and peaking at 1.6 bp, so the crisis regime is not only worse-priced on average but three times more erratic than the calm one. When arbitrage capital is scarce, the mispricing is both larger and more variable, where idiosyncratic squeezes and collateral-scarcity episodes come and go, which is precisely why a switching-variance model, rather than a switching-mean one, is required to describe the series faithfully.

Finally, the time shares deliver a sobering, quantified nuance to the conventional portrait of the OAT market as exceptionally liquid. Over the modern 2005–2026 sample the market sits in the fully efficient calm regime only $\approx 41\%$ of the time, in the elevated regime $\approx 33\%$ and in the crisis regime $\approx 26\%$, so frictions of some kind prevail in roughly six weeks out of ten. The French market is highly efficient, but principally in tranquil conditions, which are a minority of a modern era dominated by the GFC, the euro sovereign crisis, the negative-rate decade, COVID-19, the post-2022 structural floor and political instability over the last few years. The crisis regime is, in effect, the statistical abstraction of those episodes and its ≈ 9 bp mean coincides almost exactly with the standalone level of the 2023–2025 floor documented descriptively in Section 5.5, the same phenomenon, now recovered as a formally estimated latent state.

Figure 22 makes this unmistakable: the Markov crisis regime shading forms two distinct blocks, COVID-19 and the post-2022 floor, separated by an un-shaded return to calm.

Figure 22



Source: author's elaboration in Python

Figure 22 overlays, on the single weekly noise series, the output of the two independent dating methods of Section 5.7.1. The shading is the Markov-switching classification: every week whose smoothed probability of belonging to the high-noise crisis regime exceeds 0.5 is coloured, so the red bands are the periods the probabilistic model defines as crisis. The dashed vertical lines are the Bai-Perron structural breaks, the six level-shift dates⁵ the deterministic, BIC-selected segmentation places in the series. The visual story is dominated by two sustained crisis blocks, the COVID episode and the post-2022 floor, from mid-2022 onward, separated by a conspicuous un-shaded interval in which the noise falls back to ≈ 2 –3 bp; a handful of brief crisis flags appear in the pre-COVID years, but no sustained crisis block precedes 2020.

The two overlays answer complementary questions and, crucially, corroborate each other: the shades give a probabilistic assignment to the calm, elevated or crisis regime of Table 17, while the dashed lines signal the deterministic dates chosen to minimize within-segment dispersion. That a probabilistic state-classification and a

⁵ The six break dates are 08-02-2008, 13-11-2009, 17-06-2016, 06-03-2020, 28-05-2021 and 01-07-2022; the seven segment means they delimit are, in order, 1.1, 3.8, 2.4, 3.6, 7.8, 2.9 and 9.6 bp.

mean-break detector, resting on entirely different statistical machinery, agree on the same turning points, a break and the onset of crisis shading at the COVID sell-off, the end of crisis shading and a break at the 2021 recovery and a fresh break and the resumption of crisis shading in mid-2022, is itself evidence of the robustness of the result: the regimes are not an artefact of the model. The Bai-Perron segment means make the dating quantitative: the series steps from a pre-GFC ≈ 1.1 bp, through the GFC ≈ 3.8 bp and a long post-crisis $\approx 2.4\text{--}3.6$ bp, to the COVID ≈ 7.8 bp segment, back down to a ≈ 2.9 bp recovery segment and finally up to the ≈ 9.6 bp structural-floor segment that runs to the end of the sample.

Figure 22 is the formal, dated certification of the corrected timeline of Section 5.5. What was there argued from the shape of the series, that is the pandemic dislocation healed and that a separate structural floor rose later, is here proven by two independent methods: the un-shaded 2021–2022 calm interval, book-ended by a recovery break in May 2021 and a fresh crisis break in July 2022, is the graphical and statistical demonstration that the post-2022 noise floor is a distinct regime, not COVID persistence.

The identity of the crisis regime is itself informative, and honestly so. The Markov crisis state is a noise state, defined by the ≈ 9 bp pricing-dispersion level and the OAT market reached it only in COVID and after mid-2022, not during the 2008–2012 GFC and sovereign crises, whose noise, though elevated, peaked near 6 bp and is classified as merely elevated. This does not mean 2008 was economically milder; it means the cross-sectional pricing dispersion of the French market reached its historical extremes only in the recent period, which is precisely what the structural interpretation predicts: a larger, more heterogeneous post-pandemic OAT universe intermediated by a balance-sheet-constrained, net-long dealer sector disperses more under stress than the smaller, less-constrained market of 2008 did [66].

Finally, Figure 22 visualizes the persistence that the regime table quantified: the 2022 crisis block, once entered, does not revert within the sample, unlike the COVID block, which closes. The structural floor is the “sticky”, self-sustaining high-noise regime whose ≈ 55 -week expected duration Table 17 reported; here it appears as an open-ended shaded band with no exit, the graphical signature of a dislocation that is structural rather than cyclical. In one image, then, Figure 22 certifies that the modern French market has moved through two genuinely distinct high-dispersion episodes, a transitory one the ECB repaired and a persistent one it did not cause and has not reversed, each dated, each agreed upon by two methods and each consistent with the mechanism the case studies identified.

5.7.3 COMPARISON WITH THE OTHER CASE STUDIES


Concerning Case Study 2, the same procedure on 10y–2y slope recovers a flat and inverted regime, ≈ 41 bp mean and genuinely negative at its trough, a moderate ≈ 119 bp and a steep ≈ 212 bp regime. The smoothed probability of the flat or inverted regime rises to 0.83 across 2022–2024 against 0.39 before 2022; the 2022–2024 inversion is thus a distinct, high-probability regime rather than a passing fluctuation and Bai-Perron brackets it with breaks at $\approx$ November 2022 and $\approx$ September 2024. This places a formal, dated and certified footing under Section 4.3.

On the idiosyncratic OAT-Bund premium the model finds a negative ≈ -12 bp, a near-zero $\approx \pm 1$ bp and a high-premium, $\approx +30$ bp, regime; the high-premium regime probability switches from 0 in 2019 to 0.90 across 2024–2026. Remarkably, the BIC-selected Bai-Perron break falls on June 7th, 2024, the week of the National Assembly dissolution, independently re-confirming, from a structural-break perspective, the event-study finding of Case Study 1 that the dissolution was the repricing event, the dominant CASC of $+24.6$ bp and a $t \approx 12$.

Lest these regimes be an artefact of the parametric machinery, the Ploberger-Krämer CUSUM test rejects the null of structural stability for all three series,

considering also the Case Study 3 itself, at $p < 0.0001$, statistics of 12.0, 7.7 and 8.7 for noise, slope and premium respectively, an independent, distribution-free confirmation that genuine breaks exist [50].

Figure 23



Source: author's elaboration in Python

Figure 23 transports the Markov-switching classification of Section 5.7.1 from the noise onto the two other series that carry the thesis's remaining regime claims. The top panel shades the weeks the model assigns to the low-slope state of the 10y–2y slope; the bottom panel shades the weeks it assigns to the high-premium state of the idiosyncratic OAT-Bund spread and overlays the tier-1 political-event calendar. Aside from the two post-2022 episodes already dated in Section 5.7.2, Figure 23 exposes the full recurrence structure of each series. The flat or inverted regime is not unique to 2022–2024: the model also flags the pre-GFC flattening of 2005–2008 and the 2019–2022 approach to zero, yet the shaded 2023–2024 stretch is visibly the only one in which the slope crosses genuinely below zero. The high-

premium regime, likewise, recurs: brief shaded blocks are present at the 2008–2009 GFC, the 2011–2012 sovereign crisis and the 2017 election, but only the post-2024 block sits directly on the domestic political-event lines.

Both series resolve into three well-separated regimes with distinctive persistence [48]: regarding the slope, a steep regime, ≈ 212 bp, a moderate one ≈ 119 bp and a flat or inverted one ≈ 41 bp, encompassing the genuinely negative 2022–2024 trough as its extreme. The steep regime is remarkably durable, expected ≈ 240 weeks or roughly four and a half years, the moderate one the most fleeting. Concerning the premium, a negative regime, ≈ -12 bp, a near-zero one, $\approx +1$ bp and a high-premium one $\approx +30$ bp. The high-premium regime is by far the rarest of all the states in this study, occupying only $\approx 15\%$ of the sample and lasting about a year at a time.

Where Section 5.7.2 established when the post-2022 regimes switched on and how tightly each aligns with an independent signal, Figure 23 makes visible what kind of states they are instances of. The slope flat or inverted shading is the recurring statistical signature of tightenings and downturns, a leading indicator whose 2022–2024 instance the ECB and Banque de France document as the steepest euro-area inversion in decades [41], [26]. The high-premium shading is the recurring signature of the repricing that pulls France away from a pure core asset, the episodes in which credit and political risk, rather than the term structure, move the OAT.

The decisive insight comes from reading the two panels together and it is that the OAT curve shape and its level relative to the Bund are governed by different forces, acting in different regimes, on different dates. The slope regime is a monetary object: its 2022–2024 flat or inverted episode is the front-end-driven inversion of Case Study 2 and its multi-year steep regime is the mirror of the QE decade that pinned the short end near zero. The premium regime is a political-and-credit object: its post-2024 episode is the dissolution-triggered repricing of Case

Study 1, coinciding not with the monetary cycle but with the domestic event calendar. That the two regimes are sequential rather than coincident is the quantitative demonstration that these are independent determinants of the French curve, not two faces of one shock.

Two further readings sharpen the picture. First, the recurrence with escalation: both the flat and inverted slope and the high-premium spread are recurring states, however their recent instances are the most extreme of their kind, the only genuinely negative slope and, as Case Study 1 showed, a persistently higher premium. The episodes differ across recurrences in a way the panel makes legible: the 2011–2012 block was the euro-area sovereign crisis, in which France was buffeted by external flight dynamics, whereas the post-2024 block is idiosyncratically French and domestic-politics-driven, the very distinction between flight-to-quality across the periphery and country-specific repricing that the euro-area literature draws [7], [32]. Second, the rarity and cost of the high-premium state: that France sits at or below a zero premium some 84% of the time and in the high-premium regime only $\approx 15\%$ quantifies, at the regime level, the core-with-occasional-repricing character of the OAT documented event by event in Case Study 1: France is usually a core euro asset and the political-premium regime is the rare, roughly annual exception into which the 2024–2025 instability pushed it. In sum, Figure 23 certifies, with the identical dated procedure applied to the noise, that the two remaining pillars of the thesis, namely the monetary inversion and the political premium, are distinct, recurring, independently timed regimes, whose most recent occurrences are also their most acute.

5.8 CONCLUSIONS

Case Study 3 set out to bring the fitting-error and noise diagnostics of [1] to bear on the March-2020 dislocation, to compare it honestly with 2008 and 2011 and to measure the PEPP's repair of the French market. Its two questions now have precise and evidence-based answers.

March 2020 was the most severe dislocation, although only in the 5–10-year belly and not across the curve. The cross-section-normalized, per-bin comparison, which is the study’s central methodological commitment, adopted precisely because the small and shifting OAT cross-section makes a raw aggregate comparison invalid, shows COVID as the clear outlier in the belly, ~ 11 bp, roughly twice any other crisis there, middle-of-pack at 2–5 years and the mildest of all crises at 10–20 years, while 2008 was broad, 2011 short-and-intermediate and 2022 an ultra-long event. COVID was thus a belly-specific dislocation and, uniquely, one carrying a specific-collateral-scarcity component traceable to a single squeezed line.

The PEPP restored pricing efficiency fully and fast. The acute March dislocation and the endogenous August peak both healed and the noise returned below its pre-pandemic level by late 2021, a complete round-trip whose speed matched the fastest historical recovery. Crucially, the corrected timeline keeps this success distinct from the structural rise of the noise floor in 2023–2025, which is not COVID persistence, but a separate regime shift born of QT and post-GFC intermediation constraints.

Beyond these answers, the case study makes a methodological point that carries the honesty of the whole thesis: a dislocation cannot be read off an aggregate number when the cross-section is small and heterogeneous. The per-bin comparison, the recovery half-life, the leave-one-out test and a formal regime-dating layer, composed of the Markov-switching, Bai-Perron and CUSUM models, together form a principled toolkit that distinguishes where a crisis struck, how fast it healed and how genuine it was and that, applied here, replaces the original hypothesis “March 2020 was the worst crisis” with the far more accurate and honest “March 2020 was the most severe belly dislocation, partly scarcity-driven, that the PEPP repaired within a year”.

Finally, the study closes the arc of the three case studies. Each isolates a different force acting on the French curve: political and fiscal risk repricing the safety premium, the monetary regime inverting the curve and lifting the slope factor and liquidity, collateral and central-bank intervention dislocating and then repairing pricing efficiency. Yet a single thread runs through all three: the ECB balance sheet, which compresses the slope in one study and rescues market functioning in another, has become a first-order determinant of the French yield curve. Carried to June 2026, the methodology of [1] thus reveals a French sovereign market that is structurally noisier, materially higher-yielding, politically repriced and more dependent than ever on the central bank at its centre.

6. CONCLUSIONS

This thesis set out to do two main things: to replicate, on independently assembled data, the Svensson yield-curve engine that [1] built for the French sovereign market and to carry it eight years beyond the original 1987–2018 sample, into a window containing everything the original sample lacked: negative policy rates, the fastest tightening in the euro’s history, a pandemic market dislocation and a domestic political crisis without precedent in the market era of the Fifth Republic. The replication succeeded on the terms of [1], with fit metrics validated against it and the fitted curves cross-checked against the Banque de France’s official zero-coupon estimates [24]; the extension turned out to be not an appendix to the replication but its justification. Every one of the three forces that the closing diagnostics of Chapter 2 flagged, namely the repricing of level and slope, the 2022–2024 inversion and the noise and fitting-error spikes, became a case study and each one returned an answer that the pooled 1987–2018 evidence could not have anticipated.

The three answers can be stated honestly and briefly, because the chapters behind them carry the evidence. France’s safety premium has been politically repriced: the OAT-Bund spread roughly doubled around the June 2024 dissolution and its widening is idiosyncratically French, where common euro-area factors, in the decomposition spirit of [35], explain, in fact, at most a fifth of its daily variation while the country-specific premium ratcheted from a two-decade norm of zero to some 40 bp, priced at long fiscal horizons rather than as rollover risk [30], [31], [33], externally consistent with the official INSEE measure to within a basis point and transmitted, at roughly two-thirds pass-through, into the borrowing costs of French firms [36], [32], [34]. The shape of the curve has changed regime: the slope factor share of daily curve variation more than tripled to $\approx 23.6\%$ during the 2022–2024 inversion, a result invisible in levels and legible only in changes [37], certified regime-by-regime by the Lord and Pelsser criterion [23] and correctly read not as a clean recession signal but as the joint product of the tightening and

the inherited stock of asset purchases [41], [26], [43]. Market pricing efficiency has a history of dislocation and repair: March 2020 was the most severe dislocation on record only in the 5–10-year belly, roughly twice any predecessor there once the comparison is properly normalized by the cross-section; the PEPP repaired it completely within a year and the elevated noise floor of 2023–2025 is not pandemic persistence but a distinct regime born of quantitative tightening and post-crisis constraints on intermediation capacity [8], [9]. None of these three claims rests on narrative alone: the structural-break layer of Section 5.7, composed of Markov-switching, Bai-Perron and CUSUM applied identically to noise, slope and premium, dates each regime and, in the premium case, places the break in the very week of the dissolution.

Read together, the case studies yield two conclusions that none of them could deliver alone. The first is the thesis's unifying thread: the ECB balance sheet has become a first-order determinant of the French yield curve in every dimension this work measured. It compresses the slope and thereby colours the meaning of the inversion [41]; its expansion in March 2020 is the identified repair of the belly dislocation; its contraction raises the structural noise floor and, as the Stage-1 and Stage-2 regressions showed, slows the transmission of policy easing to the real economy [36]. A curve that [1] could still model as the interaction of fiscal supply and private demand now has a central bank inside its every moment. The second is the independence of the forces at work: the formal regime evidence of Figure 23 shows the monetary regime, namely the flat-and-inverted slope, switched on in November 2022 by the tightening and the political regime, the high-premium state, switched on in June 2024 by the dissolution, as sequential, not coincident, meaning different forces, different dates and different objects. The French curve of 2022–2026 was not hit by one crisis with two symptoms; it was hit twice, by a monetary cycle and by a political rupture and the thesis's main contribution here is to have measured each with the other held distinctly in view.

At the very end, the verdict this work supports is deliberately double-edged. France remains, in every mechanical sense, a core sovereign: the OAT co-moves almost one-for-one with the Bund, its spread curve steepened rather than inverted, which can be seen as the market's way of saying that near-term repayment is not in question, its market remains among the deepest and most liquid in Europe and the high-premium state has historically occupied only about 15% of the sample. But the small and stable premium of the 1987–2018 era [1] is now recognizably a regime, not a property: since 2024 France pays a persistent, country-specific, long-horizon premium that two years of political calm have not eroded, that constitutes the majority of what it pays over Germany and that has migrated into the marginal cost of corporate credit at about +0.2 percentage points [36]. The premium is, in this work's view, best read as the price of consolidation capacity: the market is not doubting the Treasury's next auction but charging, at exactly the horizons where fiscal adjustment would have to occur, for a political system that has demonstrated difficulty delivering it, the configuration in which the political-risk-to-debt-sustainability loop binds hardest [33]. The risk this evidence points to is therefore not an acute crisis but a quiet chronicity: a structurally noisier, materially higher-yielding, politically repriced market, more dependent than ever on the central bank at its centre.

The limitations of this work are those it declared as it went and they bound the verdict instead of undermining it. Every design in the case studies is associational – event windows, reduced-form regressions, projections – and no structural causal parameter has been claimed; the idiosyncratic premium is a residual and therefore a bundle, in which politics coexists with liquidity and estimation noise; the transmission block omits the US long-yield driver whose footprint its own diagnostics detected [36]; the politically driven subsamples are short and the 2025–2026 window in particular will support sharper inference only as it lengthens. Upstream of all this stand the known limits of the instrument itself: a quote-based dataset on a small and heterogeneous cross-section and a parametric functional

form whose rigidities were reviewed honestly in Chapter 1 [10], [11], [27], limits that the thesis managed with filters, robustness checks and normalized comparisons rather than pretending them away.

Each limitation marks a direction for future work. The natural next step on the political premium is causal: the local-projection instrumental-variable framework of [32], instrumenting spread changes with their movements on political dates, would convert the associations of this thesis into dynamic responses. On the modelling side, the fitted curves produced here are a ready input for the dynamic term-structure models reviewed in Chapter 1, namely an arbitrage-free Nelson-Siegel or shadow-rate extension [15], [17] would price the premium rather than measure it and the machine-learning frontier [19], [20], [21] offers a complementary route where parametric rigidity binds. Finally, the pipeline itself is a monitoring instrument: the regime classification of Section 5.7 can be re-run as new data arrive and the single most consequential open question this thesis leaves, whether the high-premium state of the French curve decays, as its 2017 ancestor did, or consolidates into the new normal that its 2024–2026 persistence suggests, is one the toolkit built here can answer in real time.

REFERENCES

- [1] Grishchenko, O. V., F. Moraux and O. Pakulyak. (2020). Fuel up with OATmeals! The case of the French nominal yield curve. *The Journal of Finance and Data Science*, 6(1), 49–85.
- [2] Ahado, F., A. Barnava, L. Chantrelle, A. Chevallier and L. Desseaux. (2025). In 2024, insurers and pension funds maintained diversified portfolios in an uncertain economic environment. *Bulletin de la Banque de France*, 260/2.
- [3] Ejsing, J. W. and J. Sihvonen. (2009). Liquidity premia in German government bonds. *ECB Working Paper Series*, No 1081.
- [4] Christensen, J. H. E., J. A. Lopez and P. Shultz. (2017). Is there an on-the-run premium in TIPS? *Federal Reserve Bank of San Francisco Working Paper Series*, 2017-10.
- [5] Copin, P. and J. Dalbard. (2022). France’s Sovereign Debt Issuance Strategy. *Trésor-Economics*, No 297.
- [6] Monfort, A. and J.-P. Renne. (2011). Credit and liquidity risks in euro-area sovereign yield curves. *Banque de France Working Paper Series*, No 352.
- [7] Beber, A., M. W. Brandt and K. A. Kavajecz. (2009). Flight-to-Quality or Flight-to-Liquidity? Evidence from the Euro-Area Bond Market. *The Review of Financial Studies*, 22(3), 925–957.
- [8] ICMA. (2024). Liquidity and resilience in the core European sovereign bond markets. *ICMA Bond Market Liquidity Taskforce report*, March 2024.
- [9] Duffie, D., M. Fleming, F. Keane, C. Nelson, O. Shachar and P. Van Tassel. (2023). Dealer Capacity and U.S. Treasury Market Functionality. *Federal Reserve Bank of New York Staff Reports*, No 1070.

- [10] Ricart, R., P. Sicsic and E. Jondeau. (2005). Estimating the term structure of interest rates from French data. In: Zero-coupon yield curves: technical documentation. *BIS Papers*, No 25, 7–8.
- [11] Anderson, N. and J. Sleath. (1999). New estimates of the UK real and nominal yield curves. *Bank of England Quarterly Bulletin*, November 1999, 384–392.
- [12] Bolder, D. J. and S. Gusba. (2002). Exponentials, Polynomials, and Fourier Series: More Yield Curve Modelling at the Bank of Canada. *Bank of Canada Working Paper*, 2002-29.
- [13] Bolder, D., S. Gusba and D. Strélski. (2005). A technical note on the Merrill Lynch Exponential Spline model as applied to the Canadian term structure. In: Zero-coupon yield curves: technical documentation. *BIS Papers*, No 25, 3–5.
- [14] Diebold, F. X. and C. Li. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2), 337–364.
- [15] Christensen, J. H. E., F. X. Diebold and G. D. Rudebusch. (2011). The affine arbitrage-free class of Nelson-Siegel term structure models. *Journal of Econometrics*, 164(1), 4–20.
- [16] Andreasen, M. M. and A. Meldrum. (2013). Dynamic term structure models: The best way to enforce the zero lower bound. *Aarhus University (CREATES) and Bank of England working paper*, August 2013.
- [17] Opschoor, D. and M. van der Wel. (2024). A Smooth Shadow-Rate Dynamic Nelson-Siegel Model for Yields at the Zero Lower Bound. *Journal of Business & Economic Statistics*, 43(2), 298–311.
- [18] Bauer, M. D., G. D. Rudebusch and J. C. Wu. (2012). Correcting Estimation Bias in Dynamic Term Structure Models. *Journal of Business & Economic Statistics*, 30(3), 454–467.

- [19] Gümbel, S. and T. Schmidt. (2020). Machine Learning for Multiple Yield Curve Markets: Fast Calibration in the Gaussian Affine Framework. *Risks*, 8(2), 50.
- [20] Akiyama, S. and N. Matsuyama. (2024). Yield curve extrapolation with machine learning. *ASTIN Bulletin*, 55(1), 76–96.
- [21] Kauffmann, P. C., H. H. Takada, A. T. Terada and J. M. Stern. (2022). Learning Forecast-Efficient Yield Curve Factor Decompositions with Neural Networks. *Econometrics*, 10(2), 15.
- [22] Gürkaynak, R. S., B. Sack and J. H. Wright. (2007). The U.S. Treasury Yield Curve: 1961 to the Present. *Journal of Monetary Economics*, 54(8), 2291–2304.
- [23] Lord, R. and A. Pelsser. (2007). Level-Slope-Curvature – Fact or Artefact? *Applied Mathematical Finance*, 14(2), 105–130.
- [24] Banque de France. (2026). Zero-coupon yield curves on French government securities. *Webstat, Banque de France statistical portal*, accessed June 2026.
- [25] Jondeau, E. and R. Ricart. (1999). The information content of the French and German government bond yield curves: why such differences? *Banque de France, Notes d'Études et de Recherche*, No 61.
- [26] Bussière, M. and S. Lhuissier. (2024). What does an inversion of the yield curve tell us? *Bulletin de la Banque de France*, 250/3.
- [27] Nymand-Andersen, P. (2018). Yield curve modelling and a conceptual framework for estimating yield curves: evidence from the European Central Bank's yield curves. *ECB Statistics Paper Series*, No 27.
- [28] Vaaler, P. M., B. N. Schrage and S. A. Block. (2005). Counting the investor vote: political business cycle effects on sovereign bond spreads in developing countries. *Journal of International Business Studies*, 36(1), 62–88.

- [29] Cotoc, I., A. Johri and C. Sosa-Padilla. (2023). Sovereign Spreads and the Political Leaning of Nations. *NBER Working Paper Series*, No 29197.
- [30] Leippold, M. and F. Matthys. (2022). Economic Policy Uncertainty and the Yield Curve. *Review of Finance*, 26(4), 751–797.
- [31] De Wet, M. C. (2023). Geopolitical Risks and Yield Dynamics in the Australian Sovereign Bond Market. *Journal of Risk and Financial Management*, 16(3), 144.
- [32] Balduzzi, P., E. Brancati, M. Brianti and F. Schiantarelli. (2023). Political Risk, Populism and the Economy. *The Economic Journal*, 133(653), 1677–1704.
- [33] Ajovalasit, S., A. Consiglio, G. Pagliardi and S. A. Zenios. (2025). Are bad governments a threat to sovereign defaults? The effects of political risk on debt sustainability. *Bruegel Working Paper*, 01/2025.
- [34] Kadiric, S. (2022). The determinants of sovereign risk premiums in the UK and the European government bond market: the impact of Brexit. *International Economics and Economic Policy*, 19(2), 267–298.
- [35] Dewachter, H., L. Iania, M. Lyrio and M. de Sola Perea. (2014). A macro-financial analysis of the euro area sovereign bond market. *National Bank of Belgium Working Paper*, No 259.
- [36] INSEE. (2025). World in turmoil, growth losing steam. *Note de conjoncture – Economic outlook*, Focus of 18 March 2025.
- [37] Litterman, R. and J. Scheinkman. (1991). Common Factors Affecting Bond Returns. *The Journal of Fixed Income*, 1(1), 54–61.
- [38] Lubet, S. (2024). PCA on Inflation Swaps. *Medium*, 7 August 2024.

- [39] Karlberg, T. and N. Åkesson. (2025). Transmission of Monetary Policy Shocks in a Small Open Economy: A Sector-Level PCA-VAR Approach. *Master thesis, Lund University School of Economics and Management*.
- [40] Ang, A., M. Piazzesi and M. Wei. (2006). What does the yield curve tell us about GDP growth? *Journal of Econometrics*, 131(1-2), 359–403.
- [41] Fonseca, L., P. McQuade, I. Van Robays and A. L. Vladu. (2023). The inversion of the yield curve and its information content in the euro area and the United States. *ECB Economic Bulletin*, Issue 7/2023.
- [42] Holtes, G. (2023). Decoding Yield Curves: The Covariance Matrix and its Ripple Effect on PCA-based Decomposition. *Medium*, 26 November 2023.
- [43] Istrefi, K., F. Odendahl and G. Sestieri. (2025). ECB Communication and its Impact on Financial Markets. *Banque de France Working Paper*, No 859.
- [44] Fitch Ratings. (2025). Fitch downgrades France to ‘A+’; outlook stable. *Rating Action Commentary*, Fitch Ratings, 12 September 2025.
- [45] Eurostat. (2026). Government deficit and government debt: annual government finance statistics. *Eurostat online database*, accessed June 2026.
- [46] Joslin, S., K. J. Singleton and H. Zhu. (2011). A New Perspective on Gaussian Dynamic Term Structure Models. *The Review of Financial Studies*, 24(3), 926–970.
- [47] Joslin, S., M. Pribsch and K. J. Singleton. (2014). Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks. *The Journal of Finance*, 69(3), 1197–1233.
- [48] Hamilton, J. D. (1989). A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle. *Econometrica*, 57(2), 357–384.

- [49] Bai, J. and P. Perron. (2003). Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics*, 18(1), 1–22.
- [50] Ploberger, W. and W. Krämer. (1992). The CUSUM Test with OLS Residuals. *Econometrica*, 60(2), 271–285.
- [51] Hu, G. X., J. Pan and J. Wang. (2013). Noise as Information for Illiquidity. *The Journal of Finance*, 68(6), 2341–2382.
- [52] Milstein, E. and D. Wessel. (2024). What did the Fed do in response to the COVID-19 crisis? *Brookings Institution*, Hutchins Center on Fiscal and Monetary Policy, updated 2024.
- [53] Financial Stability Board. (2022). Liquidity in Core Government Bond Markets. *FSB Report*, 20 October 2022.
- [54] Billmeier, A. (2021). Global Central Bank Responses to the COVID-19 Pandemic. *Western Asset Management Company*, Topic of Focus, March 2021.
- [55] European Central Bank. (2026). Pandemic emergency purchase programme (PEPP). *European Central Bank*, programme documentation, accessed June 2026.
- [56] Panigrahi, A. K., A. Sharma and V. Sarda. (2026). Liquidity Recovery Dynamics Following Volatility Shocks: Evidence from an Emerging Equity Market. *International Journal of Financial Studies*, 14, 111.
- [57] Corradin, S., N. Grimm and B. Schwaab. (2021). Euro area sovereign bond risk premia during the Covid-19 pandemic. *ECB Working Paper Series*, No 2561.
- [58] Pio, L. Y. (2022). ECB policies’ impact on the euro area sovereign bond yields and channels. *Master thesis, LUISS Guido Carli, Department of Economics and Finance*, Academic Year 2021/2022.

- [59] Rossi, A., E. Lecomte, T. Legrand and B. Nguyen. (2023). French sovereign debt liquidity: main factors, recent developments and resilience during the Covid crisis. *Bulletin de la Banque de France*, 246/1.
- [60] Dils Stedman, K. (2020). Unconventional Monetary Policy, (A)Synchronicity and the Yield Curve. *Federal Reserve Bank of Kansas City Research Working Paper*, RWP 19-09.
- [61] Jappelli, R., L. Pelizzon and M. G. Subrahmanyam. (2026). Monetary Policy, the Yield Curve, and the Repo Market. *Leibniz Institute for Financial Research SAFE*, working paper.
- [62] Schaffner, P., A. Ranaldo and K. Tsatsaronis. (2019). Euro repo market functioning: collateral is king. *BIS Quarterly Review*, December 2019, 95–108.
- [63] Greppmair, S. and S. Jank. (2023). Collateral scarcity and market functioning: insights from the Eurosystem securities lending facilities. *Deutsche Bundesbank Discussion Paper*, No 31/2023.
- [64] Carrera de Souza, T. and T. Hudepohl. (2022). The Eurosystem’s bond market share at an all-time high: what does it mean for repo markets? *De Nederlandsche Bank Working Paper*, No 745.
- [65] Canuto, O. (2023). Macroeconomic Policy Regime Change in Advanced Economies. *Policy Center for the New South, Policy Brief*, PB-02/23.
- [66] Du, W., B. Hébert and W. Li. (2022). Intermediary Balance Sheets and the Treasury Yield Curve. *Federal Reserve Bank of New York Staff Reports*, No 1023.
- [67] Bank for International Settlements. (2026). High public debt and shifting financial markets: challenges for central banks. *BIS Annual Economic Report 2026*, Chapter II.

- [68] Regnat, C., J. Briggs, T. Cuillière, B. Gerard, E. Tetard and N. Naam. (2026). 2026: Entering a New Market Regime. *Natixis CIB Research*, 22 January 2026.
- [69] Truong, C., L. Oudre and N. Vayatis. (2020). Selective review of offline change point detection methods. *Signal Processing*, 167, 107299.
- [70] Yi, M., V. Mehra, J. Chen and J. Cartlidge. (2026). Enhancing Regime Shift Detection Using Unstructured Data: A Study on the Treasury Market. *arXiv preprint*, arXiv:2605.30363.