

Fixed Income Securities

Dynamic Geometric Regimes in the Yield Curve ***A Surface-Based Framework for Fixed Income Strategy Design***

Post Graduation in Data Science for Finance

2025/2026

Author:

Cristóvão Matos 20251194

2026-02-01

Table of Contents

1	Introduction	1
2	Theoretical Framework: Yield Curve Geometry and Regimes	2
2.1	The Yield Surface as a Continuous Object.....	4
2.2	Local Shape Descriptors in Maturity Space	5
2.3	Yield Curve Dynamics: Temporal Geometry	7
2.4	Local Second-Order Geometry and the Hessian.....	8
2.5	Regime Descriptor Vector	10
2.6	Yield Curve Regime Classification.....	11
2.7	Strategy Taxonomy in Geometric Terms	13
2.8	Regime-Dependent Strategy Selection	16
3	Methodological Approach and Backtesting Design	18
3.1	Data Sources and Preprocessing	18
3.2	Synthetic Instrument Construction and Limitations.....	18
3.3	Strategy Definitions and Trading Protocol.....	18
3.4	Global Backtest Design (2006-2025)	19
3.5	Randomized Out-of-Sample Backtesting Framework.....	19
3.6	Performance and Risk Measures	19
4	Empirical Results and Performance Analysis	20
4.1	Regime Selection: Rules and Empirical Illustration.....	20
4.2	Global Backtest Results (2006-2025)	21
4.3	Randomized Out-of-Sample Backtesting Results	23
5	Discussion and Conclusions.....	26
5.1	Summary of Findings	26
5.2	Interpretation and Economic Intuition	27
5.3	Directions for Further Research	27
	Bibliography	28

1 Introduction

The yield curve is a central object in fixed income markets, encapsulating information about monetary policy expectations, term premia, and macroeconomic risk across maturities. A large body of academic and practitioner research has focused on modelling the term structure of interest rates, primarily with the objective of ensuring arbitrage-free pricing, forecasting yield movements, or decomposing yields into economically interpretable factors. These models – ranging from short-rate and affine specifications to parametric curve representations and infinite-dimensional approaches (Brigo and Mercurio, 2006) – have become foundational tools for pricing, hedging, and risk management.

From a trading perspective, however, the yield curve is not a static object but a deformable structure evolving under continuously changing forces. Over some intervals, its shape evolves smoothly, preserving familiar geometric features such as slope or curvature. At other times, these features reorganise abruptly: humped curves flatten, inversions emerge or unwind, and previously stable cross-maturity relationships become unstable. Such episodes resemble structural transitions in physical systems, in which small changes in external conditions lead to qualitative changes in the configuration of the system.

Standard term-structure models are not designed to address a question that is central in this setting: under which structural conditions is a given class of fixed income strategies well-defined? In practice, strategies such as duration trades, steepeners, or curvature (butterfly) positions implicitly assume that the yield curve admits a coherent geometric description over the holding horizon. Shape-based strategies, in particular, rely on the premise that quantities such as slope or curvature behave as stable state variables. During periods of rapid deformation or regime transition, however, these quantities may lose their structural meaning, even if they appear attractive in a static snapshot of the curve.

This observation motivates a complementary perspective. Rather than introducing a new generative or forecasting model of interest rate dynamics, we develop a Dynamic Geometric Yield Surface framework that treats the observed yield curve as a continuous surface evolving jointly in maturity and calendar time. Within this representation, yield curve configurations – such as steepness, humps, or inversion – are interpreted as geometric features of the surface, while market dynamics correspond to motion, acceleration, and coupling between temporal and cross-sectional directions. This formulation enables the use of differential tools to analyse not only the instantaneous shape of the curve, but also the local stability of that shape as the system evolves.

The central premise of this study is that many fixed income strategy breakdowns are better understood as regime failures rather than signal failures. In particular, shape-based strategies are vulnerable in environments where the yield surface undergoes rapid reconfiguration, analogous to a system near a structural or phase transition, even if curvature or slope appears favourable at a given instant. By contrast, when the surface evolves smoothly and its dominant deformation modes remain well separated, shape-based strategies are structurally supported.

To formalise this distinction, we introduce a set of local geometric and dynamical descriptors derived from first- and second-order derivatives of the yield surface. These quantities characterise both the curvature of the surface and the interaction between time and maturity dimensions, providing a local measure of structural stability. They are combined into a regime classification framework that acts as a geometric filter, determining whether directional exposure, shape-based exposure, or exposure suppression is admissible at a given point in time. The framework is deliberately descriptive rather than predictive: it does not seek to forecast yield movements, but to condition strategy deployment on the compatibility between yield curve geometry and strategy design.

The empirical analysis applies this framework to U.S. Treasury data spanning 2006-2025, a period encompassing multiple monetary regimes and stress episodes. Static and regime-conditioned strategies are evaluated using synthetic instruments under identical execution and rebalancing rules. Performance is assessed both over the full sample and across randomized, paired out-of-sample windows to ensure robustness across market environments.

The contribution of this work is threefold. First, it introduces a geometric representation of yield curve dynamics that unifies shape and temporal evolution within a single analytical framework. Second, it formalises the notion of strategy admissibility based on local structural stability of the yield surface rather than unconditional deployment. Third, it provides empirical evidence that conditioning exposure on yield curve regimes can improve downside risk characteristics and risk-adjusted performance relative to static curvature strategies.

The remainder of the report is organised as follows:

- Chapter 2 develops the theoretical framework, introducing the yield surface representation, geometric descriptors, and regime classification.
- Chapter 3 describes the empirical methodology and backtesting design.
- Chapter 4 presents and discusses the empirical results.
- Chapter 5 concludes and outlines directions for future research.

2 Theoretical Framework: Yield Curve Geometry and Regimes

The theory of interest rate modelling has evolved through a sequence of increasingly general frameworks, each addressing limitations of its predecessors while preserving the central objective of arbitrage-free pricing. A unified and rigorous treatment of these developments is provided in Filipović (2009).

Short-Rate Models

The earliest formal models of the term structure describe interest rate dynamics through the evolution of a single state variable, the instantaneous short rate. In these models, the entire yield curve is derived from assumptions about the stochastic process governing the short rate under the risk-neutral measure.

Canonical examples include the Vasicek and Cox-Ingersoll-Ross (CIR) models. Their defining features are analytical tractability and closed-form bond pricing formulas, which made them attractive for early applications in pricing and risk management. Conceptually, these models impose a strong restriction: all movements of the yield curve are driven by a single factor, implying perfectly correlated changes across maturities.

While short-rate models establish the core no-arbitrage link between dynamics and pricing, their limited ability to reproduce realistic yield curve shapes and cross-maturity behaviour motivated the development of richer specifications.

Affine and Factor-Based Term-Structure Models

To overcome the limitations of single-factor dynamics, subsequent models introduce multiple latent factors, allowing the yield curve to respond differently across maturities. In affine term-structure models, yields are expressed as affine functions of a vector of state variables, preserving analytical tractability while substantially increasing flexibility.

This class encompasses both continuous-time and discrete-time formulations and includes many widely used models in empirical finance. The Nelson-Siegel and Nelson-Siegel-Svensson (NSS) specifications, while not arbitrage-free in their original form, can be interpreted as reduced-form approximations within this broader factor-based tradition. These parametric models provide an intuitive decomposition of the yield curve into level, slope, and curvature components and are particularly effective for curve fitting and descriptive analysis. In practice, they are widely used by central banks and institutions; for instance, the European Central Bank employs the NSS model to estimate euro-area yield curves and publicly releases the corresponding model coefficients, underscoring their operational relevance despite their reduced-form nature.

From a conceptual standpoint, factor models shift the emphasis from modelling a single rate to modelling the cross-sectional structure of yields, but they still rely on strong assumptions about the functional form of the yield curve and its dynamics.

Heath-Jarrow-Morton Framework

A major conceptual advance is introduced by the Heath-Jarrow-Morton (HJM) framework, which reverses the modelling perspective entirely. Rather than specifying the dynamics of a finite-dimensional state vector, HJM models define the evolution of the entire forward rate curve directly.

In this setting, the no-arbitrage condition imposes explicit restrictions on the drift of forward rates given their volatility structure. This formulation provides maximal generality: any arbitrage-free term structure model can be represented within the HJM framework under suitable regularity conditions.

The HJM approach makes explicit that the yield curve is, in principle, an infinite-dimensional object. However, this generality comes at a cost. Practical implementations typically require dimensionality reduction, parametric volatility structures, or factor representations to regain tractability.

Unifying Perspective and Limitations

Across these modelling traditions, a common objective emerges: to ensure arbitrage-free pricing while capturing the observed dynamics of yields across maturities. Short-rate models emphasise tractability, factor models emphasise empirical fit, and HJM emphasises theoretical generality.

Despite their differences, these models share an implicit focus on pricing and forecasting, rather than on strategy admissibility. They aim to describe how yields evolve under probabilistic assumptions, not to determine whether a given yield curve configuration is structurally stable for a particular trading strategy. In particular, they do not explicitly address whether the geometric features of the yield curve – such as slope or curvature – are persistent objects or transient artefacts of regime transitions.

This observation motivates the alternative perspective developed in the remainder of this chapter. Rather than proposing another generative or forecasting model of yield dynamics, our model develops a Dynamic Geometric Yield Surface framework motivated by a practical trading question: under what conditions is it structurally meaningful to trade the yield curve at all, and which type of exposure – directional, shape-based, or none – is admissible given the prevailing market geometry?

Yield Curve Dynamics as Surface Motion

Figure 1 provides a concrete visual anchor for the perspective adopted in our work. Rather than viewing the yield curve as a sequence of unrelated cross-sections, the figure represents it as a continuous surface evolving jointly in maturity (τ) and calendar time (t). Each vertical slice corresponds to the familiar yield curve at a given date, while the surface as a whole encodes how those shapes deform, propagate, and sometimes fracture over time.




Figure 1 - U.S. Treasury Yield Surface (2006-2025)

From this viewpoint, canonical yield curve configurations – upward-sloping, flat, inverted, or humped – correspond to local geometric features of the surface. A humped curve, for example, appears as a ridge or dome in maturity space; a flat curve corresponds to a locally planar region; an inverted curve manifests as a reversal of slope along the maturity axis. These are not merely descriptive shapes: they are geometric objects whose persistence or instability matters for trading.

Crucially, the surface representation makes clear that shape alone is insufficient. What matters is not only what the curve looks like, but how that shape is moving. In physical terms, maturity derivatives describe spatial geometry, while time derivatives describe motion (the dynamics). A humped curve that evolves slowly is a fundamentally different object from a humped curve undergoing rapid deformation.

This distinction becomes evident when examining periods of market stress. During relatively stable monetary environments – such as much of the post-2010 low-rate period – the yield surface evolves smoothly. Curvature in the belly persists, slopes adjust gradually, and neighbouring maturities move coherently. In geometric terms,

the surface is locally smooth, with moderate curvature and low temporal acceleration. These are precisely the environments in which shape-based strategies, such as butterflies or roll-down trades, are structurally meaningful.

By contrast, episodes such as 2008, 2020, and 2022 appear in the figure as regions of violent surface deformation. In 2008, yields across maturities collapse rapidly, producing a steep downward motion of the entire surface accompanied by sharp distortions in curvature. In 2020, a similar but even more abrupt repricing occurs, with near-vertical temporal gradients reflecting the sudden re-anchoring of the entire term structure. In 2022, the surface tilts aggressively upward, with strong time–maturity interaction as the front end reprices faster than the long end.

From a geometric perspective, these episodes are characterised by high temporal velocity and acceleration, often combined with strong coupling between time and maturity directions. The surface locally resembles a saddle or a rapidly twisting sheet rather than a stable curved manifold. In such regimes, geometric features such as slope or curvature exist only instantaneously; they are not persistent objects that can be reliably exploited.

This observation is central to the motivation of the present framework. A static butterfly strategy is constructed to be duration- and level-neutral, implicitly assuming that curvature is a well-defined and slowly evolving feature of the yield curve. The surface view reveals why this assumption fails during regime transitions: curvature is overwhelmed by temporal dynamics, and the geometric object the strategy seeks to monetise is no longer stable.

The analogy with physics is direct. In classical mechanics, equilibrium analysis is meaningful only when accelerations are small. Similarly, in this framework, shape-based trading is meaningful only when the yield surface evolves quasi-statically. When accelerations dominate, the correct action is not to refine the shape signal, but to recognise that the system has entered a transition regime.

This motivates treating the yield curve as a geometric surface or manifold and using local differential analysis as a diagnostic tool. Rather than forecasting where the surface will move, the objective is to assess its local geometry and motion:

- Is the surface locally smooth or saddle-like?
- Is curvature persistent or transient?
- Is motion dominated by maturity structure or by temporal acceleration?

The answers to these questions determine whether directional exposure, shape-based exposure, or exposure suppression is structurally justified. The remainder of this chapter formalises this intuition by introducing differential descriptors of the yield surface and showing how they can be used as a systematic geometric filter for strategy selection.

2.1 The Yield Surface as a Continuous Object

Let $t \in \mathbb{R}^+$ denote calendar time and $\tau \in \mathbb{R}^+$ denote time-to-maturity. We define the yield surface as a mapping

$$y: (t, \tau) \mapsto y(t, \tau),$$

where $y(t, \tau)$ represents the continuously compounded zero-coupon yield at time t for maturity τ (see, e.g., Capinski and Zastawniak, 2011).

Equivalently, if $P(t, \tau)$ denotes the price at time t of a zero-coupon bond maturing at $t + \tau$, the yield is defined as

$$y(t, \tau) := -\frac{1}{\tau} \log P(t, \tau), \tau > 0.$$

This representation treats yields as a smooth function of two variables, rather than as a finite-dimensional vector indexed by maturities. Throughout this study, it is assumed that the surface $y(t, \tau)$ is sufficiently smooth to admit the partial derivatives required in subsequent sections. In practice, this smoothness is ensured through interpolation or parametric fitting of observed market data.

Yield Curve Cross-Sections

For a fixed time t , the yield surface induces a one-dimensional curve in maturity space,

$$\gamma_t(\tau) := y(t, \tau),$$

which corresponds to the standard yield curve observed in financial. From this perspective, the yield curve is a cross-section of the surface at a given time. Conversely, for a fixed maturity τ , the mapping $t \mapsto y(t, \tau)$ represents the time series of a constant-maturity yield.

This distinction is important:

- variations along τ describe shape (term structure geometry),
- variations along t describe dynamics (market evolution).

Traditional fixed income analysis often treats these two dimensions separately. By contrast, the surface representation allows them to be analysed jointly.

Interpretation and Scope

Viewing the yield curve as a surface provides a natural geometric interpretation of common fixed income concepts. Local steepness, humps, and inversions correspond to first- and second-order derivatives in maturity space. Sudden repricing episodes, such as those observed during monetary regime shifts, correspond to large temporal derivatives at fixed maturities. More complex phenomena – such as the destabilisation of curvature trades during tightening cycles – emerge from interactions between time and maturity dimensions.

Importantly, the framework does not impose arbitrage-free constraints, nor does it assume a particular stochastic model for yields. Its role is descriptive rather than generative: it provides a language for classifying observed yield curve configurations and their evolution, which can then be linked to trading strategy performance in a disciplined and systematic manner.

In the following sections, we build on this representation by introducing local geometric descriptors that quantify the shape and dynamics of the yield surface, ultimately leading to a regime-based framework for strategy selection.

2.2 Local Shape Descriptors in Maturity Space

The geometric shape of the yield curve at a given point in time is encoded in the behaviour of the yield function $\gamma_t(\tau) = y(t, \tau)$ as maturity varies. To characterise this shape in a precise and model-independent manner, we introduce local differential descriptors based on derivatives with respect to maturity τ .

These descriptors provide a continuous analogue to the intuitive notions of steepness, curvature, and inversion that are commonly invoked in fixed income analysis.

Let t be fixed throughout this section.

Local Slope

The local slope of the yield curve at maturity τ is defined as the first derivative of the yield with respect to maturity,

$$S(t, \tau) := \frac{d}{d\tau} y(t, \tau).$$

The quantity $S(t, \tau)$ measures the instantaneous rate of change of yields across adjacent maturities. Positive values correspond to an upward-sloping curve at maturity τ , while negative values indicate local inversion.

Economically, the slope captures the market's compensation for extending maturity marginally, reflecting expectations about future short rates, term premia, or both. In practice, common slope measures such as the 2Y-10Y or 5Y-30Y spreads correspond to discrete approximations of this derivative over finite intervals.

From a geometric perspective, the slope determines the first-order shape of the curve and governs the local roll-down effect experienced by carry trades.

Local Curvature

The local curvature of the yield curve is defined as the second derivative with respect to maturity,

$$C(t, \tau) := \frac{d^2}{d\tau^2} y(t, \tau).$$

Curvature measures how the slope itself changes across maturities. Positive curvature corresponds to a locally convex curve segment (a “hump”), while negative curvature corresponds to local concavity.

Curvature plays a central role in relative-value strategies such as butterfly trades, which are designed to be neutral to both level and slope while loading explicitly on second-order shape effects. Discrete butterfly constructions, such as the canonical 2Y-5Y-10Y fly, can be interpreted as finite-difference approximations to $C(t, \tau)$ evaluated around the belly of the curve.

Geometrically, curvature captures the second-order structure of the curve and identifies regions where intermediate maturities are rich or cheap relative to neighbouring points.

Global Shape Factors

While slope and curvature are defined locally, trading strategies typically operate over maturity segments rather than at infinitesimal points. To bridge local geometry with practical strategy design, we introduce time-dependent global shape factors obtained by aggregating local descriptors over maturity.

Let $w_0(\tau), w_1(\tau), w_2(\tau)$ be fixed weighting functions supported on the maturity domain of interest. We define:

- Level:

$$l(t) := \int w_0(\tau) y(t, \tau) d\tau,$$

- Slope:

$$s(t) := \int w_1(\tau) \frac{d}{d\tau} y(t, \tau) d\tau,$$

- Curvature

$$k(t) := \int w_2(\tau) \frac{d^2}{d\tau^2} y(t, \tau) d\tau.$$

These quantities provide a low-dimensional summary of the yield curve’s shape at time t . In empirical implementations, they correspond to weighted averages over maturity buckets (e.g. front end, belly, long end), ensuring stability and robustness to noise.

Importantly, the decomposition into level, slope, and curvature mirrors the well-known empirical factor structure of yield curves, but here it arises naturally from the geometry of the surface rather than from statistical estimation or model assumptions.

Interpretation and Limitations

Local shape descriptors provide an instantaneous and static characterisation of the yield curve. They describe what the curve looks like at a given time, but not how that shape is evolving. In particular, two curves with identical slope and curvature profiles may correspond to very different market environments if one is stable and the other is undergoing rapid transition.

This limitation motivates the extension of the framework to temporal dynamics. In the next section, we therefore introduce derivatives with respect to calendar time, allowing us to quantify the velocity and acceleration of yield movements and to study the interaction between shape and dynamics.

Relation to Parametric Term Structure Models

The decomposition of the yield curve into level, slope, and curvature components is closely related to parametric representations of the term structure, in which yields are expressed as linear combinations of predefined basis functions of maturity (see, e.g., Brigo and Mercurio, 2006; Filipović, 2009). In such representations, the yield curve at time t is written as

$$y(t, \tau) = \sum_{i=0}^N \beta_i(t) \phi_i(\tau),$$

where the coefficients $\beta_i(t)$ summarise the dominant modes of cross-maturity variation. From a geometric perspective, such models impose a low-dimensional parametric manifold on the space of admissible yield curves. The resulting factors provide a convenient and interpretable summary of curve shape and have been shown empirically to capture a large fraction of cross-maturity variation.

The widely used Nelson-Siegel-Svensson (NSS) family corresponds to a particular choice of such basis functions. For this choice, the leading coefficients are commonly interpreted as level, slope, curvature, and an additional long-term curvature component.

The framework adopted in this study differs in two important respects:

- First, the shape descriptors introduced above – local slope $\partial_\tau y(t, \tau)$ and curvature $\partial_{\tau\tau} y(t, \tau)$ – are defined directly on the observed yield surface, without imposing a specific functional form. As a result, the analysis remains valid even when the curve exhibits features that are poorly captured by a fixed parametric basis, such as local distortions, segmented markets, or abrupt regime changes.
- Second, while Nelson-Siegel-Svensson models focus primarily on the cross-sectional shape of the curve at each point in time, the present framework explicitly integrates temporal dynamics and time-maturity interactions through higher-order derivatives. This allows the distinction between stable geometric configurations and transitional regimes, a distinction that is central to the strategy-selection problem studied in this report.

In this sense, parametric term structure models can be viewed as global approximations to the geometric quantities introduced here. The local differential descriptors provide a more granular and flexible representation, which is better suited for regime classification and for analysing the stability of shape-dependent trading strategies.

2.3 Yield Curve Dynamics: Temporal Geometry

While the shape descriptors introduced in the previous section characterise the cross-sectional geometry of the yield curve at a fixed point in time, they do not convey how that geometry evolves. For the purpose of trading strategy selection, this distinction is crucial: strategies that are profitable in a given shape regime often rely implicitly on the stability of that shape over time. Capturing such stability requires extending the analysis from maturity space to calendar time.

To this end, we study the temporal derivatives of the yield surface $y(t, \tau)$, holding maturity fixed and allowing time to vary.

Temporal Velocity

The temporal velocity of the yield at maturity τ is defined as

$$V(t, \tau) := \frac{d}{dt} y(t, \tau).$$

This quantity measures the instantaneous rate at which yields are repriced by the market. Large absolute values of $V(t, \tau)$ indicate periods of rapid adjustment, typically associated with macroeconomic shocks, policy surprises, or abrupt changes in expectations.

From an economic perspective, velocity captures the speed of information incorporation into prices. For trading strategies that rely on carry or roll-down, high velocity environments are often unfavourable, as the mark-to-market impact of repricing can dominate the intended structural exposure.

Temporal Acceleration

Beyond first-order dynamics, we define the temporal acceleration as the second derivative with respect to time,

$$A(t, \tau) := \frac{d^2}{dt^2} y(t, \tau).$$

Acceleration measures the rate of change of velocity and provides a natural indicator of regime stress or transition. Periods of elevated acceleration correspond to environments in which the market response itself is evolving rapidly, such as during the onset of monetary tightening cycles or crisis-driven repricing episodes.

Empirically, large acceleration has been observed during episodes such as the post-pandemic inflation shock, where yields not only increased sharply but did so with rapidly changing intensity across time. In such environments, strategies calibrated to historical relationships tend to break down.

Time-Maturity Interaction

The joint dependence of yields on time and maturity introduces an additional source of dynamics through the mixed partial derivative,

$$\frac{d^2}{dt d\tau} y(t, \tau).$$

This term measures how the slope of the yield curve evolves over time. In other words, it captures whether the curve is steepening or flattening dynamically, rather than statically.

From a strategic standpoint, this interaction is particularly relevant for slope-based trades. A curve that is steep at a given point in time may nevertheless be an unfavourable environment for steepeners if the slope is decaying rapidly. Conversely, modest static slopes may offer attractive opportunities if the time-maturity interaction indicates a persistent steepening trend.

Geometrically, the presence of a significant mixed derivative implies that the yield surface cannot be decomposed into independent time and maturity components. Instead, shape and dynamics are coupled, signalling a departure from stable regimes.

Dynamic Stability and Strategy Implications

The temporal derivatives introduced above allow the distinction between static shape and dynamic regime. Two yield curves may exhibit identical level, slope, and curvature profiles at a given time t , yet correspond to fundamentally different environments if their temporal velocity and acceleration differ.

This observation motivates the interpretation of yield curve regimes not solely in terms of shape, but as combinations of shape and dynamics. Stable regimes are characterised by moderate velocity and low acceleration, allowing shape-based strategies to exploit persistent geometric features. Transitional regimes, by contrast, are marked by elevated acceleration and strong time-maturity coupling, conditions under which relative-value strategies become fragile.

In the following section, these dynamic considerations are formalised through the introduction of the Hessian matrix of the yield surface, which provides a unified second-order description of both temporal and maturity-based curvature and serves as a local indicator of regime stability.

2.4 Local Second-Order Geometry and the Hessian

To unify cross-sectional and temporal aspects of yield curve behaviour, we adopt a compact second-order representation of the yield surface formulated in terms of differential operators. Henceforth, operator notation is used throughout to simplify expressions and to emphasise the geometric nature of the analysis. Let

$$\partial_t := \frac{d}{dt}, \quad \partial_\tau := \frac{d}{d\tau}$$

and assume that the yield surface $y(t, \tau)$ is twice continuously differentiable in both arguments.

The Hessian Operator

The Hessian of the yield surface at a point (t, τ) is defined as the symmetric second-order differential operator

$$H_y(t, \tau) = \begin{pmatrix} \partial_{tt}y(t, \tau) & \partial_{t\tau}y(t, \tau) \\ \partial_{t\tau}y(t, \tau) & \partial_{\tau\tau}y(t, \tau) \end{pmatrix},$$

where

$$\partial_{tt} := \partial_t^2, \quad \partial_{\tau\tau} := \partial_\tau^2, \quad \partial_{t\tau} := \partial_t \partial_\tau.$$

The Hessian aggregates three distinct sources of second-order behaviour:

- The temporal acceleration $\partial_{tt}y$,
- The maturity curvature $\partial_{\tau\tau}y$,
- and the time-maturity interaction $\partial_{t\tau}y$.

Taken together, these quantities provide a local quadratic approximation to the yield surface and describe how yields respond to simultaneous perturbations in time and maturity.

Determinant and Local Phase Classification

A key scalar invariant of the Hessian is its determinant,

$$D(t, \tau) := \det H_y(t, \tau) = (\partial_{tt}y)(\partial_{\tau\tau}y) - (\partial_{t\tau}y)^2.$$

The sign of $D(t, \tau)$ classifies the local geometry of the yield surface:

- $D(t, \tau) > 0$: the surface is locally convex or concave, with curvature of the same sign along both principal directions. This corresponds to a locally stable regime, in which second-order effects reinforce each other.
- $D(t, \tau) < 0$: the surface exhibits a saddle configuration, indicating a transitional regime in which curvature is positive in one direction and negative in another.
- $D(t, \tau) \approx 0$: the surface is locally degenerate, with one dominant curvature direction and one nearly linear direction. Such points often mark boundaries between regimes.

This classification mirrors standard results in differential geometry, but here it acquires a direct economic interpretation: stable geometric configurations support persistent curve-based relationships, while saddle-like configurations signal fragility of shape-dependent strategies.

Eigenstructure and Directionality

Let $\lambda_1(t, \tau) \geq \lambda_2(t, \tau)$ denote the eigenvalues of $H_y(t, \tau)$, with corresponding orthonormal eigenvectors

$$e_1(t, \tau), e_2(t, \tau) \in \mathbb{R}^2.$$

The eigenvalues quantify the magnitude of curvature along the principal directions of the surface, while the eigenvectors identify the directions in the (t, τ) -plane along which second-order effects are strongest.

The orientation of the dominant eigenvector $e_1(t, \tau)$ is particularly informative:

- alignment with the τ -axis indicates that second-order behaviour is dominated by maturity curvature, favouring shape-based strategies such as butterflies;
- alignment with the t -axis indicates dominance of temporal acceleration, a setting in which static shape assumptions are unreliable;

- intermediate orientations signal strong coupling between time and maturity, indicating unstable environments for both slope- and curvature-based trades.

Thus, the eigenstructure of the Hessian provides a local phase portrait for admissible fixed income strategies.

Time-Maturity Coupling Metric

To quantify the strength of interaction between temporal and cross-sectional curvature, we introduce the dimensionless coupling metric

$$\rho(t, \tau) := \frac{|\partial_{t\tau} y(t, \tau)|}{\sqrt{|\partial_{tt} y(t, \tau)| |\partial_{\tau\tau} y(t, \tau)|}}.$$

The interpretation of $\rho(t, \tau)$ is immediate:

- $\rho(t, \tau) \ll 1$: weak coupling – shape and dynamics evolve largely independently, supporting shape-based strategies;
- $\rho(t, \tau) \approx 1$: strong coupling – the surface approaches a transitional configuration;
- $\rho(t, \tau) > 1$: mixed curvature dominates, and the local geometry is effectively saddle-like.

The metric ρ therefore complements the determinant and eigenstructure by providing a continuous measure of regime instability.

Implications for Strategy Stability

The Hessian formalism provides a unified diagnostic of whether the local geometry of the yield surface is compatible with the assumptions embedded in a given trading strategy. In regions where curvature is dominated by maturity effects and coupling is weak, relative-value strategies based on slope or curvature are structurally supported. In contrast, regions characterised by strong time-maturity interaction or temporal acceleration correspond to environments in which such strategies are exposed to regime transition risk.

Importantly, this framework is not predictive in nature. Rather than forecasting yields, it evaluates the local consistency between yield curve geometry and strategy design. This distinction underlies the regime-based strategy selection mechanism developed in the subsequent sections.

2.5 Regime Descriptor Vector

The preceding sections introduced a collection of geometric and dynamical quantities that characterise the yield surface locally in time and maturity. Taken individually, these quantities provide partial information about shape, dynamics, or stability. For the purpose of strategy selection, however, it is necessary to combine them into a unified representation of the state of the yield curve.

In this section, we formalise this representation through the definition of a regime descriptor vector, which serves as the state variable of the system and underpins the regime classification and strategy-selection mechanisms developed later in the chapter.

Definition of the Regime Descriptor

Let $y(t, \tau)$ denote the yield surface, and let τ^* denote a representative maturity (or maturity band) of interest, such as the belly of the curve. We define the regime descriptor at time t as the vector

$$x_t := (l(t), s(t), k(t), H_y(t, \tau^*), \rho(t, \tau^*)).$$

The components of x_t correspond to quantities defined in the preceding sections: the global level factor $l(t)$, slope factor $s(t)$, and curvature factor $k(t)$, together with the local Hessian matrix $H_y(t, \tau^*)$ of the yield surface and the associated time-maturity coupling metric $\rho(t, \tau^*)$, both evaluated at the anchor maturity τ^* .

Collectively, these elements provide a compact state representation that integrates cross-sectional shape information with local second-order stability diagnostics, and form the basis for the regime classification and strategy-selection framework developed below.

Choice of Maturity Anchor

The inclusion of the Hessian and coupling metric requires selecting a maturity (or maturity range) at which local second-order behaviour is evaluated. This choice is not innocuous, as different regions of the curve may exhibit distinct dynamics.

In this study, τ^* is chosen to correspond to the central maturity segment of the curve, commonly referred to as the belly. This region is particularly relevant for shape-based strategies, such as butterflies, and is often the locus of curvature concentration. Evaluating second-order geometry at τ^* therefore provides a meaningful proxy for the stability of shape-dependent trades.

The framework can be readily extended to multiple anchors or to maturity-averaged Hessian measures, but such extensions are left for future research.

Interpretation as a State Variable

The regime descriptor x_t defines a point in a finite-dimensional state space $\mathcal{X}$. As time evolves, the yield surface traces a trajectory $\{x_t\}_{t \geq 0}$ through this space.

This representation allows a clear separation between:

- observation: extracting geometric and dynamical features from the yield surface;
- classification: mapping the state x_t to a discrete regime;
- decision: selecting a trading strategy conditional on the regime.

Importantly, the descriptor x_t is constructed exclusively from observable quantities. It does not rely on latent variables, estimated risk premia, or parametric assumptions about the data-generating process. As such, it is well suited for empirical evaluation and backtesting.

Role in the Overall Framework

The regime descriptor vector serves two distinct but related roles in the analysis:

- First, it provides the input to the regime classification function introduced in the next section, which maps continuous geometric information into a finite set of economically interpretable regimes.
- Second, it acts as the conditioning variable for the strategy-selection rule. In this sense, x_t plays a role analogous to a state variable in a Markov decision process, although no probabilistic transition structure is imposed at this stage.

By explicitly separating the construction of x_t from both classification and strategy choice, the framework remains modular and transparent. This modularity is particularly important for the empirical comparison between static and adaptive strategies undertaken later in the report.

2.6 Yield Curve Regime Classification

Having defined the regime descriptor vector x_t , we now formalise the notion of a yield curve regime as a discrete abstraction of the continuous geometric and dynamical state of the yield surface. The purpose of regime classification is not to forecast future yield movements, but rather to partition the state space into regions within which similar trading strategies are expected to perform in a structurally consistent manner, thereby supporting algorithmic trading decisions.

Regime Space and Classification Map

Let $\mathcal{X} \subset \mathbb{R}^d$ denote the state space of regime descriptors, where each point $x_t \in \mathcal{X}$ summarises the level, shape, and local second-order dynamics of the yield curve at time t . We define a finite regime set

$$\mathcal{R} = \{r_1, r_2, \dots, r_K\},$$

where each regime r_k corresponds to a qualitatively distinct configuration of yield curve geometry and stability.

A regime classification is then defined as a deterministic mapping

$$\Phi: \mathcal{X} \rightarrow \mathcal{R}, r_t := \Phi(x_t).$$

The map Φ induces a partition of the continuous state space $\mathcal{X}$ into K disjoint regions, each associated with a regime label.

Classification Principles

The construction of Φ follows three guiding principles:

- First, regime boundaries are defined in terms of geometric and dynamical stability, rather than economic narratives. In particular, second-order quantities derived from the Hessian – such as the determinant $D(t, \tau^*)$ and the coupling metric $\rho(t, \tau^*)$ – play a central role in distinguishing stable from transitional regimes.
- Second, the classification is coarse by design. The objective is not to capture every nuance of yield curve behaviour, but to identify broad environments in which the assumptions underlying specific trading strategies are either supported or violated.
- Third, the classification rule is time-invariant. Once specified, the map Φ is applied uniformly across the entire sample period, ensuring that regime definitions do not depend on future information.

Shape-Stability Regimes

Conceptually, regimes can be understood as combinations of two dimensions:

- Shape, summarised by the level, slope, and curvature components $(l(t), s(t), k(t))$;
- Stability, summarised by the Hessian-derived quantities $D(t, \tau^*)$ and $\rho(t, \tau^*)$.

This perspective leads naturally to regimes such as:

- stable steep curves,
- stable curved (humped) curves,
- flat or weakly curved curves,
- transitional or unstable curves characterised by saddle-like second-order geometry.

While the precise number and definition of regimes are application-specific, the unifying principle is that regimes differ in the compatibility between yield curve geometry and strategy design, rather than in the absolute level of interest rates.

Regime Persistence and Transitions

Once classified, the sequence $\{r_t\}_{t \geq 0}$ defines a discrete-time regime process. No probabilistic structure is imposed on regime transitions: regimes are not treated as latent states evolving according to a Markov process, nor are transition probabilities specified or estimated. Instead, regime persistence and transitions arise deterministically from the observed evolution of the yield surface.

Periods in which x_t remains within the same region of $\mathcal{X}$ correspond to stable regimes, while crossings of regime boundaries correspond to transitions. This viewpoint is consistent with the geometric interpretation developed earlier: transitions are associated with changes in the local curvature of the yield surface, often signalled by rising acceleration or increased time-maturity coupling.

Role in the Empirical Analysis

The regime classification step provides the critical link between the theoretical framework and the empirical backtesting exercise. By assigning each observation to a regime, it becomes possible to:

- evaluate the performance of static strategies within each regime;
- define adaptive strategies that condition explicitly on r_t ;
- compare unconditional and regime-dependent performance in a controlled manner.

The next section builds on this classification by formalising a taxonomy of fixed income trading strategies in geometric terms, clarifying which strategies are expected to be viable within each regime.

2.7 Strategy Taxonomy in Geometric Terms

With the yield curve regimes formally defined, we now introduce a taxonomy of fixed income trading strategies expressed in geometric terms. The purpose of this taxonomy is to clarify the implicit assumptions embedded in different classes of strategies and to relate these assumptions to the shape and stability properties of the yield surface.

The classification adopted here distinguishes strategies according to the order of geometric deformation of the yield curve to which they are exposed.

Directional Strategies

A directional strategy is defined as a portfolio whose payoff depends primarily on first-order movements of the yield surface in the normal (vertical) direction. Formally, such strategies exhibit non-zero exposure to the level component of the yield curve.

Let $w(\tau)$ denote a maturity-weighting function representing the portfolio exposure. The first-order change in portfolio value is approximated by

$$\Delta V_t \approx -\int w(\tau) \Delta y(t, \tau) d\tau.$$

Directional strategies satisfy

$$\int w(\tau) d\tau \neq 0,$$

and therefore load directly on parallel shifts of the yield curve.

Examples include buy-and-hold positions in fixed-rate bonds or duration-driven trades. The performance of such strategies is dominated by expectations about the level of interest rates and is largely insensitive to the finer geometric structure of the curve.

From the perspective of the yield surface, directional strategies are insensitive to slope and curvature, but highly exposed to temporal velocity and acceleration. As a result, their performance is strongly regime-dependent, particularly in environments characterised by rapid repricing.

Shape-Based Strategies

In contrast, shape-based strategies are designed to be neutral to level shifts and, in many cases, to linear tilts of the yield curve. Their payoff depends on relative movements across maturities and therefore on the geometric shape of the curve.

Formally, a shape-based strategy satisfies

$$\int w(\tau) d\tau = 0,$$

and often

$$\int \tau w(\tau) d\tau = 0,$$

ensuring neutrality with respect to level and slope.

From a differential perspective, shape-based strategies correspond to discrete approximations of derivative operators acting on $y(t, \tau)$:

- Slope strategies (e.g. steepeners and flatteners) load on first-order maturity derivatives;
- Curvature strategies (e.g. butterflies) load on second-order maturity derivatives.

In particular, a butterfly strategy can be interpreted as a finite-difference approximation to the curvature operator $\partial_{\tau\tau} y(t, \tau)$, typically centred on the belly of the curve.

Duration, DV01, and Convexity as First- and Second-Order Operators

The geometric framework developed in this chapter naturally encompasses the classical duration–convexity approximation that underpins much of fixed income risk management. To make this connection explicit, consider a default-free coupon bond with price $B(t)$ at time t . Let y denote the yield used to discount its cash flows, and assume that B is a twice continuously differentiable function of y .

A second-order Taylor expansion of $B(y)$ around a reference yield y_0 yields

$$B(y_0 + \Delta y) \approx B(y_0) + B'(y_0)\Delta y + \frac{1}{2}B''(y_0)(\Delta y)^2.$$

Dividing through by $B(y_0)$, and defining the modified duration MD and convexity C as

$$MD(y_0) := -\frac{1}{B(y_0)} B'(y_0), \quad C(y_0) := \frac{1}{B(y_0)} B''(y_0),$$

one obtains the standard duration–convexity approximation:

$$\frac{\Delta B}{B} \approx -MD \Delta y + \frac{1}{2}C (\Delta y)^2.$$

In practical fixed income applications, first-order interest rate risk is often expressed in terms of DV01 (Dollar Value of a Basis Point), defined as

$$DV01 := -\frac{\partial B}{\partial y} \times 10^{-4}.$$

DV01 measures the absolute change in bond price resulting from a one basis point increase in yield and is therefore a scaled version of modified duration. Portfolio-level DV01 is obtained by summing instrument-level DV01s weighted by position sizes.

Both modified duration and DV01 quantify first-order sensitivity to yield changes, while convexity captures second-order sensitivity.

Mapping Duration and Convexity to Global Shape Factors

The classical notions of duration and convexity can be interpreted directly within the geometric framework developed in this chapter. In particular, they correspond to specific projections of yield curve deformations onto the global shape factors $l(t)$, $s(t)$, and $k(t)$.

To make this connection explicit, consider a fixed-income instrument with price $B(t)$, whose value depends on the yield curve $y(t, \tau)$ through discounting of cash flows. A first-order variation of the price induced by a small perturbation $\delta y(t, \tau)$ of the yield curve can be written, to leading order, as

$$\delta B(t) \approx -\int \mathcal{D}(\tau) \delta y(t, \tau) d\tau,$$

where $\mathcal{D}(\tau)$ denotes the instrument's distribution of discounted cash-flow weights.

This expression highlights that duration-based sensitivity is fundamentally a functional of the yield curve deformation, not merely a scalar sensitivity to a single yield.

Duration as Projection onto the Level Direction

A parallel shift of the yield curve corresponds to a deformation of the form

$$\delta y(t, \tau) = \delta l,$$

constant across maturities. Substituting into the expression above yields

$$\delta B(t) \approx -\delta l \int \mathcal{D}(\tau) d\tau.$$

Thus, modified duration and DV01 correspond to the projection of the yield curve deformation onto the level direction $l(t)$. In this sense,

$$\text{Duration} \leftrightarrow \frac{\partial B}{\partial l}.$$

Within the present framework, duration is therefore a first-order sensitivity to the zeroth-order shape factor of the yield curve.

Classical Static Assumption

The classical duration-convexity framework implicitly assumes a static environment in which calendar time t is held fixed. Yield changes Δy are treated as exogenous shocks applied to a single yield curve snapshot, and the analysis abstracts from the temporal evolution of the curve shape.

Formally, traditional theory operates on the mapping

$$y \mapsto B(y),$$

with no explicit dependence on time or maturity structure beyond the scalar yield parameter. As a consequence, classical measures do not distinguish between yield changes occurring in stable regimes and those occurring during periods of rapid repricing or structural transition.

Dynamic Geometric Generalisation

The framework developed in this report generalises classical fixed income risk measures by embedding them in a dynamic yield curve geometry. Instead of analysing price sensitivity with respect to a single yield variable at fixed t , we consider the full yield surface $y(t, \tau)$ and its evolution in both time and maturity.

In this Dynamic Yield Curve Geometry framework:

- duration and DV01 correspond to first-order sensitivity to vertical displacements of the surface;
- convexity corresponds to second-order sensitivity along the yield dimension;
- shape-based strategies correspond to higher-order derivatives in the maturity direction;
- regime stability is assessed through second-order derivatives in both time and maturity, captured by the Hessian operator.

This extension allows classical concepts such as duration neutrality and convexity exposure to be interpreted conditionally on the local geometry and dynamics of the yield surface. In particular, it clarifies why duration-neutral or curvature-based strategies may perform well in some regimes and fail in others: their effectiveness depends not only on cross-sectional shape, but also on the temporal stability of that shape.

Butterfly as a Duration-Neutral Shape Trade

A canonical example of this generalisation is the butterfly strategy, typically implemented using three maturities $\tau_1 < \tau_2 < \tau_3$. In practice, the weights of the butterfly are chosen to ensure that the portfolio's aggregate DV01 is approximately zero, rendering the position duration-neutral with respect to parallel shifts of the curve.

Let $R(t, \tau)$ denote the traded rate at maturity τ . A stylised butterfly payoff can be written as

$$F_t = w_1 R(t, \tau_1) + w_2 R(t, \tau_2) + w_3 R(t, \tau_3),$$

with weights chosen such that

$$\sum_{i=1}^3 w_i \text{DV01}(\tau_i) = 0.$$

Under this constraint, the leading exposure of the butterfly is not to level shifts, but to relative deformations of the curve. In geometric terms, the butterfly approximates the curvature operator $\partial_{\tau\tau} y(t, \tau)$ evaluated near τ_2 .

It is therefore a pure shape trade, monetising second-order structure in maturity space while neutralising first-order directional risk.

Within the dynamic geometric framework, the profitability of such trades is expected to depend critically on regime stability, as characterised by the Hessian-based diagnostics introduced earlier. Butterflies are structurally favoured in regimes where curvature is persistent and time–maturity coupling is weak, and structurally fragile in transitional regimes dominated by temporal acceleration.

2.8 Regime-Dependent Strategy Selection

The final element of the theoretical framework is the definition of regime-dependent strategy selection. The central premise is that fixed income strategies correspond to exposures to specific geometric deformations of the yield surface, and that such exposures are admissible only in regimes where the local geometry and dynamics of the surface are compatible with the assumptions embedded in the strategy.

Accordingly, strategy selection is formulated not at the level of individual trades or instruments, but at the level of strategy classes, defined in geometric terms and activated conditionally on the prevailing yield curve regime.

Strategy Classes as Geometric Operators

Let $y(t, \tau)$ denote the yield surface. We define strategy classes as operators acting on y , characterised by the order and direction of the yield surface deformation to which they are exposed.

The following classes are considered:

- **Directional (Level) Strategies:** These strategies load on uniform vertical displacements of the yield surface and are sensitive primarily to the level factor $l(t)$. Formally, they correspond to first-order exposure to deformations proportional to the constant direction in maturity space.
- **Slope (First-Order Shape) Strategies:** These strategies load on linear tilts of the yield curve and are sensitive to the slope factor $s(t)$. Geometrically, they correspond to first-order derivatives of the yield surface with respect to maturity.
- **Curvature (Second-Order Shape) Strategies:** These strategies load on bending of the yield curve and are sensitive to the curvature factor $k(t)$. They correspond to second-order derivatives with respect to maturity and include duration-neutral constructions that monetise relative deformations across maturities.
- **Carry and Roll-Down Strategies:** These strategies exploit the joint effect of time evolution and maturity structure, and are sensitive to combined temporal and maturity derivatives of the yield surface. Their effectiveness depends critically on the stability of the local surface geometry.
- **Exposure Suppression (Risk-Off):** In regimes where yield curve geometry is unstable, the admissible strategy class is the null operator, corresponding to the deliberate suppression of shape-based and directional exposures.

At this stage, these classes are defined abstractly. Their implementation using specific instruments, maturities, and risk constraints is deferred to the methodological chapter.

Hessian-Based Regime–Strategy Rule Set

Strategy admissibility is governed by the local second-order geometry of the yield surface.

Fix an anchor maturity τ_0 (or a narrow maturity band). At each time t , compute the second-order derivatives:

$$y_{tt} := \partial_{tt}y(t, \tau_0), y_{\tau\tau} := \partial_{\tau\tau}y(t, \tau_0), y_{t\tau} := \partial_{t\tau}y(t, \tau_0),$$

and define the Hessian determinant

$$D := y_{tt} y_{\tau\tau} - y_{t\tau}^2.$$

Define the time–maturity coupling ratio

$$\rho_H := \frac{|y_{\tau\tau}|}{\sqrt{|y_{tt} y_{\tau\tau}|} + \varepsilon},$$

where $\varepsilon > 0$ ensures numerical stability.

Rule 1: Shape-Admissible versus Transition Regimes

- If $D > 0$ and ρ_H is small, the yield surface is locally elliptic, indicating a stable geometric regime. In such environments, shape-dependent strategies (curvature, slope, carry, roll-down) are admissible.
- If $D < 0$ or ρ_H is large, the surface exhibits saddle-like geometry, indicating a transition regime. In these environments, static shape assumptions are unreliable; admissible strategies are restricted to transition-aware positioning or exposure suppression.

Rule 2: Curvature versus Slope Emphasis

Within shape-admissible regimes, the relative magnitude of second-order derivatives determines the dominant deformation mode:

- **Curvature-dominated regimes:** If $|y_{\tau\tau}|$ dominates and $|y_{tt}|$ is small, curvature is locally persistent. Second-order shape strategies (curvature extraction) are favoured.
- **Slope-transition regimes:** If $|y_{\tau\tau}|$ dominates, the slope of the yield curve is changing rapidly over time. In such environments, slope-oriented strategies are preferred over static curvature trades.
- **Acceleration-dominated regimes:** If $|y_{\tau\tau}|$ dominates, temporal acceleration overwhelms cross-sectional structure. Both slope and curvature strategies are downweighted.

Rule 3: Risk Scaling via Transition Intensity

To modulate overall exposure, define a transition intensity scalar

$$\Xi(t, \tau_0) := y_{tt}^2 + 2y_{t\tau}^2 + y_{\tau\tau}^2.$$

Position size is scaled inversely with transition intensity:

$$\text{Exposure}(t) \propto \frac{1}{\Xi(t, \tau_0)}.$$

High values of Ξ correspond to unstable regimes with elevated second-order curvature, leading to systematic risk reduction. Low values indicate stable regimes in which shape-based strategies can be deployed with higher confidence.

Static versus Adaptive Strategy Deployment

A static strategy corresponds to the unconditional application of a single strategy class across all regimes, independent of the local Hessian geometry. By contrast, an adaptive strategy applies the Hessian-based rule set dynamically, selecting both the admissible strategy class and its exposure level conditional on the prevailing regime.

The empirical analysis in subsequent chapters compares static and adaptive deployments under identical market conditions, providing a direct test of the hypothesis that conditioning strategy exposure on yield surface geometry improves risk-adjusted performance.

Scope and Link to Empirical Implementation

This section specifies which geometric exposures are admissible and how they are conditioned on regime, but not how such exposures are implemented in practice. No maturity selection, instrument choice, or parameter optimisation is specified at this stage.

Concrete implementations – such as steepeners, butterflies, or buy-and-hold portfolios – are introduced in Chapter 3 as discrete instantiations of the strategy classes and rule set defined here. This separation ensures that empirical results can be interpreted as tests of regime consistency rather than artefacts of specific trade construction.

3 Methodological Approach and Backtesting Design

This chapter describes the empirical methodology used to evaluate the central hypothesis of this study: that conditioning fixed-income trading strategies on yield-curve regimes improves risk-adjusted performance relative to static deployment. The analysis is conducted through a combination of a full-sample backtest and a randomized, paired out-of-sample evaluation, using synthetic instruments derived directly from observed yield data.

The methodological design prioritises robustness, transparency, and comparability. All strategies are evaluated on identical data, instruments, and rebalancing schedules, ensuring that observed performance differences can be attributed to regime conditioning rather than implementation artefacts.

3.1 Data Sources and Preprocessing

The empirical analysis is based on U.S. Treasury constant-maturity yields obtained from the Federal Reserve Economic Data (FRED) database. Daily observations spanning January 2006 to December 2025 are used, covering a wide range of market environments, including the Global Financial Crisis, the post-crisis low-rate regime, the COVID shock, and the subsequent tightening cycle.

A cross-section of maturities from three months to thirty years is employed to estimate the yield curve on each trading day. Missing observations are forward-filled, and yields are expressed in decimal form. No filtering or de-meaning is applied beyond this preprocessing.

To obtain a smooth and differentiable representation of the yield curve, a Nelson-Siegel-Svensson (NSS) specification is fitted independently for each trading day. All maturity-domain derivatives required for regime diagnostics are computed analytically from the fitted curve rather than from raw yield observations. This smoothing step is essential to ensure numerical stability and to avoid spurious regime classifications driven by measurement noise.

3.2 Synthetic Instrument Construction and Limitations

Rather than using traded securities, the analysis employs synthetic fixed-income instruments constructed directly from yield changes. This approach ensures full control over exposures and avoids confounding effects related to coupon structure, roll-down, or security-specific liquidity.

Two instruments are defined:

- a directional instrument, approximating exposure to the five-year maturity using a duration-based price change plus carry;
- a curvature (butterfly) instrument, constructed as a DV01-neutral combination of the 2-year, 5-year, and 10-year maturities.

Instrument returns are computed using first-order duration approximations combined with daily carry. Convexity effects beyond first order, transaction costs, financing frictions, and roll-down are deliberately excluded. While this abstraction limits direct tradability, it allows the analysis to isolate the structural contribution of regime-dependent exposure selection.

Butterfly weights are rebalanced on a fixed schedule and normalised to constant gross exposure, ensuring comparability across time and strategies.

3.3 Strategy Definitions and Trading Protocol

The analysis focuses on two strategies, evaluated under identical execution rules:

- **The Static Butterfly** strategy maintains continuous exposure to the curvature instrument throughout the sample. It represents unconditional shape exposure and serves as the primary benchmark.
- **The Adaptive Geo-Regime** strategy dynamically allocates exposure based on the local geometric properties of the yield surface. At each decision date, second-order diagnostics derived from the fitted yield surface are used to classify the prevailing environment as either shape-admissible or transition-

dominated. Within shape-admissible regimes, the dominant deformation mode determines whether curvature or directional exposure is selected; in transition-dominated regimes, the strategy remains flat.

All strategy decisions are made at time t using information available up to that date and are executed from $t + 1$, eliminating look-ahead bias. Positions are held constant between rebalancing dates.

A fixed rebalancing frequency of 21 trading days is imposed across all strategies and instruments. This choice reflects a medium-term investment horizon and avoids excessive sensitivity to high-frequency fluctuations. No parameter optimisation or forecasting is performed at any stage.

3.4 Global Backtest Design (2006-2025)

The first evaluation consists of a global backtest spanning the full 2006-2025 period. Both strategies are applied continuously over the sample using the same instruments, rebalancing schedule, and return construction.

This analysis provides a baseline assessment of long-run performance, drawdown behaviour, and regime occupancy, and serves primarily as a descriptive illustration of how regime-dependent exposure alters the temporal structure of returns.

3.5 Randomized Out-of-Sample Backtesting Framework

To assess robustness and reduce dependence on any single historical path, a second evaluation is conducted using randomized, rolling out-of-sample windows.

The sample is partitioned into a training period (used exclusively to fix regime-related thresholds) and a test period over which all performance evaluation is conducted. From the test sample, multiple four-year windows are drawn using randomised starting dates subject to a minimum spacing constraint, ensuring that windows are not excessively overlapping.

Each window consists of a fixed-length sample of daily returns. An initial burn-in period is excluded from evaluation to ensure that regime classification, rebalancing logic, and position holding are fully initialised. Performance metrics are computed only on the remaining observations.

All windows are evaluated in a paired manner, meaning that both strategies are assessed on identical time intervals. This design allows direct dominance comparisons across return, risk, and tail-risk measures.

3.6 Performance and Risk Measures

Strategy performance is evaluated using a consistent set of return and risk metrics computed on daily returns and annualised where appropriate. These include annualised return and volatility, maximum drawdown, Sharpe and Sortino ratios, Sterling and Omega ratios, and Value-at-Risk and Conditional Value-at-Risk at the 95% level.

This metric set is designed to capture not only average performance, but also downside risk, tail behaviour, and drawdown severity, which are central to evaluating regime-dependent strategies.

Scope and Limitations

The methodology is intentionally conservative. The adaptive strategy does not optimise maturities, rebalance dynamically within regimes, or exploit short-term predictive signals. Its sole purpose is to test whether conditioning exposure on yield-curve geometry improves structural performance.

The use of synthetic instruments and the exclusion of transaction costs and convexity effects limit direct real-world applicability, but they allow for a clean, internally consistent comparison. Extensions incorporating tradable instruments, roll-down, or probabilistic regime transitions are left for future research.

4 Empirical Results and Performance Analysis

This chapter presents the empirical results of the backtesting exercise described in Chapter 3. The purpose of the analysis is not to optimise trading rules or to propose a deployable strategy, but rather to assess whether conditioning fixed-income exposures on yield-surface geometry leads to systematic and economically meaningful differences in performance relative to static strategy implementations.

4.1 Regime Selection: Rules and Empirical Illustration

Before analysing performance, we illustrate how the regime-selection mechanism operates in practice. Figure 2 shows the evolution of U.S. Treasury yields at the 2-year, 5-year, and 10-year maturities together with the regime classifications generated by the geometric diagnostics introduced in Chapter 3.

The framework classifies regimes in two steps. First, it distinguishes between shape-admissible and transition-dominated regimes based on the strength of time–maturity coupling and the sign of local curvature of the yield surface. Shape-admissible regimes correspond to periods in which the yield surface evolves in a sufficiently regular manner to support systematic exposure, while transition-dominated regimes reflect unstable or rapidly reconfiguring yield surfaces and result in no exposure.

Second, within shape-admissible regimes, the dominant deformation mode determines the type of exposure. Curvature-dominated regimes favour curvature-based strategies implemented via a DV01-neutral butterfly, whereas slope-dominated regimes favour directional exposure. Regime labels are computed at a daily frequency as diagnostics, while trading decisions are sampled only on fixed rebalancing dates and held over time.

Figure 2 illustrates how these rules translate into observed regime classifications. Green-shaded regions indicate curvature-dominated regimes, orange-shaded regions indicate slope-dominated regimes, and unshaded intervals correspond to transition-dominated regimes.



Figure 2 - U.S. Treasury yields (2Y, 5Y, 10Y) with geometric regime classification, 2020–2025

Extended curvature-dominated regimes coincide with periods of gradual yield curve evolution, such as the post-Global Financial Crisis low-rate environment and parts of the post-pandemic disinflation phase, during which curvature patterns are relatively stable. Slope-dominated regimes occur more episodically, most notably during sustained monetary tightening, such as the period beginning in 2022, when yields across maturities move more uniformly. Transition-dominated regimes are concentrated around abrupt repricing episodes, where yields adjust rapidly and synchronously across maturities.

Shape-admissible regimes are more prevalent than slope-dominated regimes throughout the sample. This reflects a well-documented characteristic of fixed income markets: outside of pronounced tightening or easing cycles, yield curve adjustments are dominated by changes in curvature and relative maturity structure rather than persistent parallel shifts (see, e.g., Martellini et al., 2003). Directional regimes therefore arise less frequently and are primarily associated with extended phases of coordinated level movements driven by monetary policy.

4.2 Global Backtest Results (2006-2025)

The global backtest over the 2006-2025 period serves two main objectives. First, it provides a baseline comparison between a static curvature strategy and the proposed adaptive, regime-conditioned strategy when applied over a long and economically diverse sample. Second, it offers an aggregate view of how regime conditioning affects long-run performance, risk, and drawdown characteristics before introducing more granular, window-based robustness tests in subsequent sections.

Aggregate Performance Comparison

Table 1 reports standard performance and risk metrics for the static butterfly strategy and the adaptive geometric regime strategy over the full sample.

	Static Butterfly	Adaptive Geo Regime
Annual return	0.061%	0.179%
Annual volatility	0.613%	0.904%
Max drawdown	2.103%	2.789%
Sharpe ratio	0.103	0.203
Sortino ratio	0.142	0.160
Omega ratio	1.018	1.097
Sterling ratio	0.029	0.064
VaR (0.95)	0.060%	0.036%
CVaR (0.95)	0.090%	0.114%
Global score	33.3%	66.7%

Table 1 – Aggregate performance comparison of the two strategies for the global test (2006-2025)

The adaptive geometric regime strategy delivers a higher annualised return than the static butterfly strategy (0.179% versus 0.061%) and exhibits an improvement in risk-adjusted performance. This is reflected in higher Sharpe and Sortino ratios, indicating that the additional return more than compensates for the increase in volatility on both a total and downside-risk basis. Omega and Sterling ratios also favour the adaptive strategy, suggesting a more favourable balance between gains and losses and improved performance relative to drawdowns.

These improvements are achieved despite the adaptive strategy being inactive for a substantial fraction of the sample. This reflects the conditional nature of the exposure framework: capital is deployed selectively, only when the local yield surface geometry is deemed compatible with the underlying strategy design. The resulting performance therefore arises from timing and regime selection rather than continuous risk-taking.

At the same time, the adaptive strategy experiences a larger maximum drawdown over the full sample (2.789% versus 2.103%). This indicates that regime conditioning does not uniformly reduce all measures of risk in an unconditional sense. In particular, while exposure is suppressed during periods of geometric instability, adverse outcomes may still occur during extended shape-admissible regimes. This reinforces the importance of analysing drawdown behaviour across subsamples and stress periods, rather than relying solely on a single global drawdown statistic.

Tail risk measures present a nuanced picture. The adaptive strategy exhibits a lower 95% Value-at-Risk, indicating a reduction in the frequency of moderate adverse outcomes. However, Conditional Value-at-Risk at the same confidence level is higher than for the static strategy, suggesting that when extreme losses do occur,

they tend to be more severe. This pattern is consistent with a strategy that avoids many small losses through regime filtering but remains exposed to occasional larger shocks within active regimes.

Overall, the global performance metrics point to modest but consistent improvements in risk-adjusted outcomes for the adaptive strategy. While these results do not indicate overwhelming or uniform dominance across all risk dimensions, they support the view that conditioning strategy deployment on yield surface geometry can enhance performance efficiency relative to static curvature trades.

Regime Occupancy and Trading Activity

Table 2 summarises the distribution of time spent in each regime. Transition-dominated regimes, during which no exposure is taken, account for approximately half of the sample. Curvature-dominated regimes represent a large share of the remaining time, while slope-dominated (directional) regimes are relatively rare.

Regime	Days	Share (%)
Transition (no trading)	2,855	54.7%
Butterfly	2,165	41.5%
Directional	197	3.8%

Table 2 - Trading activity and exposure statistics for the adaptive strategy

This regime composition is consistent with the yield curve dynamics discussed in Section 4.1. Outside of sustained tightening or easing cycles, yield curves tend to evolve through changes in curvature rather than persistent parallel shifts, resulting in a predominance of shape-based regimes. The limited time spent in directional regimes also explains why directional exposure contributes episodically to performance rather than acting as a continuous return driver.

The relatively high share of transition regimes implies that the adaptive strategy is deliberately inactive for long periods. This design choice materially affects both return accumulation and drawdown dynamics and should be interpreted as a structural feature of the framework rather than a weakness of implementation.

Equity Curve Behaviour

Figure 3 displays the equity curves of the static butterfly and adaptive regime-based strategies over the full 2006-2025 sample, together with the 2-, 5-, and 10-year Treasury yields (rescaled for visual comparability). Shaded regions indicate periods in which the adaptive strategy is actively deployed versus held out of the market.




Figure 3 - Equity curves of static and adaptive strategies with yield level overlay and regime activity shading

A decomposition of returns indicates that the Global Financial Crisis (2008–2009) accounts for a disproportionate share of the adaptive strategy's lifetime performance. While the static butterfly derives approximately 48% of its total return from this period, the adaptive strategy earns more than its full-sample cumulative return during the crisis alone. This outcome reflects a deliberate concentration of exposure during periods of large and coherent yield curve adjustment. Further breakdown shows that the majority of crisis-period gains originate from directional exposure, consistent with the sharp and persistent decline in yields across maturities. By contrast, the static butterfly distributes crisis gains over continuous exposure throughout the sample, resulting in a substantially smaller cumulative impact. These results suggest that the adaptive framework behaves as a selective allocator to rare, high-impact term structure events rather than as a steady carry strategy.

An analysis of worst-case rolling performance windows further clarifies the behaviour of both strategies. The adaptive strategy's weakest two-year period ends in late 2014, coinciding with the post-crisis normalization and tapering phase, characterised by frequent yield curve reconfiguration but limited sustained trends. In such environments, the strategy experiences short-lived exposures and modest entry–exit losses, resulting in underperformance despite reduced overall risk. By contrast, the static butterfly's worst two-year window occurs during the 2022-2023 tightening cycle, when persistent yield increases and curve inversion generate sustained adverse carry and convexity effects. These results highlight that the adaptive strategy's underperformance is associated with structurally ambiguous regimes rather than prolonged directional stress, reinforcing its interpretation as an event-driven allocator rather than a continuous exposure strategy.

Outside crisis episodes, movements in the adaptive equity curve are closely aligned with periods in which the yield curve evolves smoothly and coherently. During extended tightening or easing phases – visible in the figure as sustained directional shifts in the rescaled yield curves – the strategy remains selectively exposed and generates incremental gains. In contrast, periods characterised by rapid changes in slope, curvature, or relative maturity behaviour are frequently classified as transition regimes, during which exposure is reduced or eliminated, resulting in flat equity segments.

Downward adjustments in the adaptive equity curve are primarily associated with regime transitions rather than prolonged adverse exposure. These typically occur when yield curve dynamics shift abruptly, such as during changes in monetary policy stance or heightened macroeconomic uncertainty, leading to short-lived losses around entry or exit points. The static butterfly strategy, by contrast, remains continuously exposed to yield curve fluctuations, producing a smoother equity profile but with smaller cumulative gains over the sample.

Overall, the equity curve behaviour supports the conceptual motivation of the adaptive approach: returns are driven by selective participation in geometrically stable yield curve configurations, while capital preservation is prioritised during structurally unstable phases. The visual correspondence between equity dynamics and yield curve movements provides qualitative evidence that the regime-selection mechanism captures economically meaningful features of term structure evolution.

4.3 Randomized Out-of-Sample Backtesting Results

This section evaluates the empirical performance of the proposed Adaptive strategy relative to a static butterfly benchmark using randomized, paired out-of-sample backtests. Following the methodology introduced in Chapter 3, performance is assessed over fixed-length rolling windows with spaced starting dates to mitigate dependence induced by overlapping samples. All results reported below are computed on identical windows for both strategies, ensuring a fair comparison. The analysis focuses on multi-year horizons, which are more consistent with the medium-term nature of yield curve regime dynamics.

Overall backtesting results

Table 3 summarises the aggregate performance metrics computed over the full out-of-sample period. The results highlight a clear contrast between the two strategies in terms of both return generation and risk characteristics.

	Static Butterfly	Adaptive Geo Regime
Annual return	-0.026%	0.028%
Annual volatility	0.520%	0.610%
Max drawdown	0.946%	0.560%
Sharpe ratio	-0.051	0.045
Sortino ratio	-0.075	0.029
Omega ratio	0.991	1.020
Sterling ratio	-0.029	0.032
VaR (0.95)	0.049%	0.025%
CVaR (0.95)	0.078%	0.079%
Global score	22.2%	77.8%

Table 3 - Aggregate performance (median) comparison of the two strategies under randomized backtesting

The randomized backtesting results reinforce and refine the conclusions drawn from the global 2006-2025 analysis. While the full-sample test reflects performance along a single realised historical path, the randomized experiment evaluates strategy behaviour across a broad set of overlapping subperiods, thereby abstracting from path dependence and isolating structural performance characteristics.

In the randomized setting, the Static Butterfly strategy continues to exhibit weak performance. Median annualised returns across windows are slightly negative, volatility remains moderate, and downside characteristics are unfavourable. Maximum drawdown is close to 1% on average, while both VaR and CVaR at the 95% level remain elevated relative to the adaptive alternative. Consistent with these outcomes, all risk-adjusted metrics are negative or below economically meaningful thresholds: Sharpe, Sortino, and Sterling ratios indicate that persistent curvature exposure fails to deliver adequate compensation once adverse regimes are sampled repeatedly across different market environments.

By contrast, the Adaptive Geo Regime strategy demonstrates a clear improvement in performance consistency across randomized windows. Median annualised returns turn positive, albeit modest in absolute terms, and this improvement is achieved with only a moderate increase in volatility. Importantly, downside behaviour improves materially relative to the static benchmark. Maximum drawdown is significantly lower across windows, and the left tail of the return distribution is compressed, as reflected in substantially lower VaR. CVaR remains broadly comparable between the two strategies, indicating that while extreme losses are not eliminated, their frequency is reduced under the adaptive framework. As a result, all risk-adjusted performance measures favour the Adaptive strategy, including Sharpe, Sortino, Omega, and Sterling ratios. This dominance is summarised by the global score, which assigns nearly 80% of the aggregate ranking weight to the adaptive approach.

When viewed alongside the global 2006–2025 results, an important distinction emerges. In the full-sample test, the adaptive strategy delivered higher returns but also experienced larger peak-to-trough drawdowns, reflecting the influence of long realised regimes and extended exposure phases along a single historical trajectory. In contrast, the randomized backtesting shows that, across typical subperiods, the adaptive framework systematically improves downside control and risk-adjusted efficiency. This suggests that the larger drawdowns observed in the global test are not a structural feature of the strategy, but rather the consequence of specific realised regime sequences.

Return distributions

The return distributions across four-year windows show substantial overlap between the two strategies, but the Adaptive Geo Regime is modestly shifted to the right. In the boxplot this appears as a slightly higher median and upper quartile, while in the histogram the adaptive distribution places a bit more mass around small positive outcomes and near-zero returns. By contrast, the Static Butterfly exhibits a relatively larger concentration of mildly negative realizations, consistent with continuous exposure through episodes where curve geometry is not supportive.




Figure 4 - Distribution of out-of-sample annual returns across randomized 4-year windows

Maximum drawdown distributions

The drawdown distributions separate more clearly than returns (Figure 5). The Adaptive Geo Regime has a higher (less negative) median drawdown and a visibly thinner left tail, whereas the Static Butterfly shows more frequent and deeper drawdowns. In the histogram this shows up as heavier mass in the “worse drawdown” region for the static strategy, indicating that extreme drawdown realizations occur more often when the portfolio remains permanently exposed rather than scaling down or disengaging during unstable phases.




Figure 5 - Distribution of out-of-sample maximum drawdowns across randomized 4-year

Randomized Paired Window Dominance Analysis

Table 4 reports a paired dominance analysis comparing the Adaptive strategy to the static butterfly across 300 out-of-sample, one-year rolling windows with spaced starting dates. Each row corresponds to a performance metric, and all comparisons are conducted on identical windows for both strategies.

Column definitions:

- **P(Adaptive > Static):** Fraction of windows in which the Adaptive strategy outperforms the static butterfly for the given metric. Values above 0.5 indicate more frequent outperformance.
- **median (Adaptive – Static) and mean(Adaptive – Static):** Median and mean differences in the metric between the two strategies across windows. Positive values favour Adaptive.
- **p05(Adaptive – Static) and p95(Adaptive – Static):** 5th and 95th percentiles of the paired differences, providing a view of downside and upside dispersion.

Metric	P (Adaptive > Static)	Median (Adaptive-Static)	Mean (Adaptive-Static)	p05 (Adaptive-Static)	p95 (Adaptive-Static)
Annual return	0.5550	0.00023	0.00024	-0.00129	0.0019

Metric	P (Adaptive > Static)	Median (Adaptive-Static)	Mean (Adaptive-Static)	p05 (Adaptive-Static)	p95 (Adaptive-Static)
Max drawdown	0.6050	0.00422	0.00197	-0.00930	0.0106
Sharpe ratio	0.5950	0.06942	0.02977	-0.28985	0.3151
Sortino ratio	0.5700	0.06966	0.05351	-0.34746	0.3902
Omega ratio	0.4850	-0.00340	0.00977	-0.11091	0.1051
Sterling ratio	0.6200	0.05515	0.04205	-0.20019	0.2671
VaR (0.95)	1.0000	0.00024	0.00028	0.00007	0.0005
CVaR (0.95)	0.4750	-0.00001	0.00003	-0.00041	0.0005

Table 4 - Dominance statistics from randomized out-of-sample backtesting (Adaptive – Static)

Across four-year rolling windows, the Adaptive strategy shows a moderate but systematic advantage over the static butterfly. For annualised returns, the Adaptive strategy outperforms the static benchmark in around 56% of the sampled windows, with both the median and mean return differentials slightly positive. This indicates that, over multi-year horizons, the Adaptive allocation more frequently delivers incremental performance gains, though the magnitude of these gains remains modest and the overlap between strategies is substantial.

Risk-adjusted performance metrics broadly reinforce this pattern. Sharpe and Sortino ratios favour the Adaptive strategy in approximately 60% and 57% of windows, respectively, with positive median and mean differences. Sterling ratios show a similar dominance frequency, reflecting improved drawdown-adjusted performance over extended horizons. Omega ratios, by contrast, display near-symmetry between the two strategies, with dominance probabilities close to 50% and a slightly negative median difference, suggesting that gain–loss asymmetry is not a primary channel of improvement.

The clearest and most consistent advantage appears in downside risk measures. Maximum drawdown is lower for the Adaptive strategy in roughly 61% of the windows, with a positive median reduction in drawdown magnitude. Most notably, Value-at-Risk at the 95% level is lower for the Adaptive strategy in all sampled windows, indicating uniform improvement in moderate tail risk across rolling four-year periods. Conditional Value-at-Risk also slightly favours the Adaptive strategy on average, though dominance frequencies are closer to balanced, reflecting similar behaviour in extreme tail outcomes.

Taken together, the dominance patterns suggest that the Adaptive strategy's primary benefit over the static butterfly lies in improved downside control rather than large or pervasive return outperformance. Over economically relevant four-year horizons, the Adaptive framework more often delivers marginally higher returns while consistently reducing drawdown severity and tail risk exposure, supporting its use as a regime-aware alternative to continuous curvature positioning.

5 Discussion and Conclusions

5.1 Summary of Findings

This study set out to examine whether conditioning fixed income strategy deployment on the local geometry and dynamics of the yield curve improves structural performance relative to static strategy application. To this end, the yield curve was treated as a continuous surface evolving jointly in maturity and calendar time, and local differential descriptors were used to characterise its shape, motion, and stability.

The central contribution of the framework is conceptual rather than predictive. By shifting the focus from forecasting yield movements to assessing strategy admissibility, the Dynamic Geometric Yield Surface approach reframes fixed income trading as a problem of regime consistency. Strategies are interpreted as exposures to specific geometric deformations of the yield surface, and their deployment is conditioned on whether those deformations correspond to stable and persistent surface features.

Empirically, the results support this perspective. Across the 2006-2025 sample and a wide range of randomized, paired out-of-sample windows, the adaptive regime-conditioned strategy exhibits systematically improved downside characteristics relative to a static curvature (butterfly) benchmark. While absolute return improvements are modest – as expected in low-risk government bond relative-value strategies – the adaptive approach consistently reduces drawdown severity and tail risk. These improvements arise despite the strategy remaining inactive for a substantial fraction of the sample, highlighting that selective participation rather than continuous exposure is a key driver of performance differences.

Importantly, the adaptive framework does not dominate unconditionally across all metrics or horizons. Periods of underperformance are observed, particularly in environments characterised by frequent but shallow regime transitions, where repeated entry and exit generate frictional losses. This behaviour underscores that the framework is not designed to smooth returns or maximise hit rates, but to avoid structural exposure in unstable regimes.

Taken together, the findings suggest that yield curve geometry contains economically meaningful information beyond static shape measures, and that exploiting this information as a regime filter can improve the robustness of fixed income strategies.

5.2 Interpretation and Economic Intuition

From an economic standpoint, the results align with the intuition that fixed income markets alternate between quasi-static regimes, in which cross-maturity relationships are coherent, and transition regimes, in which rapid repricing overwhelms relative-value structure. Traditional static strategies implicitly assume the former, even when the market is clearly in the latter.

The geometric surface representation makes this distinction explicit. Periods such as the Global Financial Crisis, the COVID shock, and the post-pandemic tightening cycle correspond to regions of elevated temporal acceleration and strong time-maturity coupling. In these environments, slope and curvature exist only instantaneously and cannot be relied upon as stable trading objects. The adaptive strategy's tendency to disengage during such phases explains both its reduced tail risk and its episodic concentration of returns during high-impact events.

Equally important is what the framework does not do. It does not attempt to forecast yield levels or regime transitions in a probabilistic sense, nor does it generate return predictions in isolation. Instead, it provides a structural diagnostic of the local geometry of the yield surface, which can be used to identify admissible exposure types, orient directional or shape-based positions, and govern the timing of entries and exits. This distinction clarifies why the framework complements, rather than competes with, traditional term structure models and signal-based strategies: it conditions when and how exposures are taken, rather than predicting what the market will do.

5.3 Directions for Further Research

The results of this study point to several promising directions for future work.

- **Noise reduction and regime smoothing:** A natural extension is to reduce sensitivity to high-frequency noise in regime diagnostics. Rather than classifying regimes based on instantaneous geometric quantities, one could apply temporal aggregation or state estimation techniques. Kalman filtering or related state-space models offer a principled way to infer latent regime states from noisy derivative estimates, potentially reducing spurious entries and exits while preserving responsiveness to genuine transitions.
- **Optimisation of entry and exit mechanics:** The current implementation treats regime classification as a hard filter with immediate effect at rebalancing dates. More sophisticated mechanisms – such as hysteresis bands, persistence requirements, or regime confidence scores – could reduce whipsaw

behaviour. These enhancements would be particularly relevant during borderline regimes where the surface geometry oscillates between stability and transition.

- **Optimisation of maturity selection for slope and curvature trades:** The empirical analysis focuses on fixed maturity sets (e.g. a 2–5–10 butterfly and a 5-year directional proxy). A natural extension is to endogenise maturity selection. Rather than imposing a fixed configuration, the framework could select the maturity triplet (for curvature trades) or maturity pair (for slope trades) that maximises sensitivity to the relevant geometric mode at each point in time. Such optimisation could be based on local curvature intensity, slope magnitude, or DV01 efficiency, subject to liquidity and transaction-cost constraints. This would allow the strategy to adapt not only when to trade, but also where along the curve to express the intended exposure.
- **Integration with forecasting signals:** The framework is deliberately non-predictive, but it need not remain so. An important avenue for future research is to combine geometric regime filtering with traditional forecasting or signal-generation models. In such a hybrid setup, forecasting models would operate only within shape-admissible regimes, while regime diagnostics would determine when forecasts are trusted. This separation between when to trade and what direction to trade is conceptually appealing and may lead to superior risk-adjusted outcomes.
- **Hyperparameter tuning and learning:** Thresholds governing regime classification, coupling strength, and exposure scaling are currently fixed ex ante. These parameters could be tuned using cross-validation, Bayesian optimisation, or reinforcement learning techniques, subject to appropriate regularisation. Care would be required to avoid overfitting, particularly given the low signal-to-noise ratio typical of fixed income returns.
- **Extension to additional strategy classes and markets:** The geometric framework is not specific to butterflies or even to government bond markets. It could be extended to slope trades, multi-curve settings, inflation-linked bonds, or swap curves. More broadly, the methodology may be applicable to other markets where term structures arise, such as credit curves, volatility surfaces, or futures curves.
- **From local to global geometry:** Finally, the present analysis focuses on local differential properties of the yield surface. A deeper exploration of global geometric features – such as topological changes, regime persistence, or surface segmentation – may yield additional insights into long-term structural behaviour.

Bibliography

Brigo, D. and Mercurio, F. (2006), *Interest Rate Models: Theory and Practice*. Springer.

Capinski, M. and Zastawniak, T. (2011), *Mathematics for Finance: An Introduction to Financial Engineering*. Springer.

Filipović, D. (2009), *Term-Structure Models: A Graduate Course*. Springer.

Martellini, L., Priaulet, P. and Priaulet, S. (2003), *Fixed-Income Securities: Valuation, Risk Management and Portfolio Strategies*. Wiley.