Decision, Priority, and Delay: A Computational-Philosophical Re-theorization of Financial Foundation Models

METAXOUDIS ELEFTHERIOS
Model Vocational Upper Secondary School
Ministry of Education, Religious Affairs and Sports
Ainou 95-97, Alexandroupoli,68131
GREECE

*Abstract:*Financial Foundation Models (FFMs) have recently emerged as powerful architectures for forecasting, risk assessment, and decision support in complex financial environments. Despite their empirical success, existing research predominantly emphasizes performance optimization, explainability, and data efficiency, while largely overlooking a foundational question: how FFMs internally structure decision timing, priority allocation, and delay under uncertainty. This paper addresses this gap by proposing a computational-philosophical re-theorization of FFMs, conceptualizing them not merely as predictive systems but as decision-making architectures governed by implicit temporal and hierarchical logics.

We introduce a structured framework grounded in six core computational concepts: procrastination (decision deferral), anti-procrastination (forced activation), metric (value encoding beyond loss functions), inverse priority (dominance of rare events), priority inheritance (temporal propagation of salience), and the trade-off between responsiveness and throughput. These concepts are analytically mapped onto FFM architectures to reveal how priorities emerge, shift, and propagate across decision layers, particularly under stress conditions.

To operationalize the framework, we employ financial Digital Twins as experimental substrates for stress testing and scenario simulation. Illustrative simulations and statistical comparisons are used to demonstrate how priority-aware and delay-sensitive FFMs exhibit qualitatively different decision behaviours compared to conventional designs. The proposed approach contributes a novel theoretical lens for understanding decision latency and priority dynamics in financial AI systems, with implications for model architecture, risk management, and regulatory stress testing. By bridging Computational Philosophy and Financial Engineering, this study advances a foundational perspective on how financial AI systems decide -not only what to predict, but when and under which priorities to act.

Key-Words: Financial Foundation Models; Computational Philosophy; Decision Theory; Priority Inversion; Digital Twins; Stress Testing

Received: June 5, 2025. Revised: September 27, 2025. Accepted: October 25, 2025. Published: March 9, 2026.

1 Introduction

The rapid advancement of artificial intelligence has profoundly reshaped financial modeling, giving rise to a new class of architectures commonly referred to as Financial Foundation Models (FFMs). These models extend the paradigm of foundation models into the financial domain, integrating large-scale pretraining, multimodal data fusion, and task generalization across forecasting, risk assessment, and decision support. Recent developments demonstrate that FFMs can process heterogeneous financial information-ranging from time-series data and textual disclosures to visual representations-within unified architectures, enabling unprecedented levels of analytical scope and performance. As a result, FFMs are increasingly positioned at the core of modern financial infrastructures, supporting activities such as portfolio management, systemic risk monitoring, and stress testing.

Despite these advances, the dominant trajectory of FFM research has been overwhelmingly performance-oriented. The literature has focused on improving predictive accuracy, scalability, and computational efficiency, alongside growing attention to explainability and regulatory compliance. While these contributions are substantial, they implicitly treat FFMs as sophisticated predictors rather than as decision-making systems. Consequently, the internal logic through which FFMs determine *when* to act, *what* to prioritize, and *how long* to defer action under uncertainty remains largely unexamined. Decision timing, priority allocation, and delay are typically absorbed into optimization routines or treated as secondary implementation details, rather than as first-order architectural properties.

This omission is not merely theoretical. In financial contexts, decision delay and priority structuring are inseparable from risk exposure and systemic stability. Acting too early on insufficient evidence can amplify volatility, while excessive delay may allow localized disturbances to propagate into systemic crises. Similarly, the prioritization of objectives-such as liquidity preservation, return maximization, or solvency protection-cannot be assumed to be static or universally ordered. Under stress conditions, rare events and tail risks often override routine performance considerations, forcing abrupt reconfigurations of decision hierarchies. Yet current FFM architectures provide limited insight into how such reconfigurations occur, how priorities propagate across model layers, or how urgency is encoded and resolved.

This gap points to a deeper conceptual limitation in existing approaches. While explainable AI techniques offer post hoc interpretations of model outputs, they do not address the underlying *ontology of decision-making* within FFMs. Metrics and loss functions specify what is optimized, but not how decisions are temporally structured or hierarchically governed. As a result, key questions remain unanswered: How do FFMs manage uncertainty through deliberate non-decision? Under what conditions is delay rational, and when does it become dangerous? How do priorities emerge, invert, or persist across time and architectural layers? Addressing these questions requires moving beyond performance metrics toward a principled theory of decision governance in financial AI systems.

This paper addresses this foundational gap by introducing a computational-philosophical re-theorization of Financial Foundation Models. Drawing on the methodological tradition of Computational Philosophy, the study treats FFMs as epistemic decision architectures rather than passive predictive instruments. Within this perspective, decision-making is understood as a structured process involving managed delay, dynamic prioritization, and value-laden metric commitments. The paper formalizes six core computational concepts-procrastination, anti-procrastination, metric, inverse priority, priority inheritance, and the responsiveness–throughput trade-off-that together capture the essential dimensions of decision timing and hierarchy under uncertainty. These concepts provide a unified vocabulary for analyzing behaviors that are often treated as technical artifacts, revealing them instead as rational but fragile outcomes of implicit design choices.

To operationalize this framework, the paper bridges FFMs with Financial Digital Twins and stress testing environments. Digital Twins are employed not merely as simulation tools, but as experimental substrates in which decision architectures can be instantiated, observed, and stress-tested under controlled yet dynamically evolving conditions. This integration allows the proposed concepts to be examined in scenarios involving regime shifts, cascading failures, and extreme events-contexts in which decision delay and priority inversion become systemically consequential.

By repositioning FFMs as decision-making systems governed by temporal and hierarchical logics, this study contributes a foundational lens for Financial Engineering and AI-driven risk management. Rather than asking only how accurately financial models predict outcomes, it advances the question of how financial AI systems *decide*: when to wait, when to act, and under which priorities action becomes inevitable.

This study distinguishes itself from prior work by explicitly formalizing the internal decision governance mechanisms of Financial Foundation Models rather than treating them solely as predictive architectures. While previous research has focused primarily on performance optimization, multimodal integration, or explainability techniques, the present work introduces a structured theoretical framework that models how FFMs regulate decision timing, priority allocation, and delay under uncertainty. The novelty of this contribution lies in translating computational-philosophical concepts into operational constructs that can be embedded within model architectures and examined through Digital Twin experimentation. In this sense, the paper advances not only a conceptual reinterpretation but also a methodological pathway for implementing and evaluating decision-aware financial AI systems.

2 Problem Formulation

2.1 Financial Foundation Models as Decision Systems: FFMs beyond predictors

Financial Foundation Models (now and then after FFMs) are increasingly conceptualized not merely as advanced predictive instruments, but as decision systems that embed reasoning, prioritization, and adaptive control within financial environments characterized by uncertainty and structural complexity. Unlike traditional econometric or machine learning models -whose primary function is to forecast prices, risks, or probabilities-FFMs operate as integrative architectures that support decision-making processes across time, data modalities, and hierarchical objectives.

Recent research underscores this conceptual shift. Chen et al. (2025) propose a foundational taxonomy that classifies FFMs into three principal modalities: Financial Language Foundation Models, Financial Time-Series Foundation Models, and Financial Visual-Language Foundation Models. This categorization reveals that FFMs are designed to process heterogeneous forms of financial information, ranging from textual disclosures and numerical market data to visual representations such as charts and balance-sheet structures. Crucially, the unification of these modalities enables FFMs to move beyond isolated predictions toward context-aware decision synthesis [1].

Expanding on this perspective, Huang et al. (2025) emphasize the role of FFMs as cohesive decision engines capable of decomposing complex financial tasks into structured sub-decisions. Through multimodal fusion, FFMs can align linguistic reasoning (e.g., policy interpretation), temporal dynamics (e.g., volatility regimes), and visual cues (e.g., technical patterns) into a single decision trajectory. In this sense, FFMs resemble intelligent agents rather than static models, as they internally coordinate priorities, defer or accelerate actions, and reconcile competing informational signals [2].

This agent-like character is further articulated in the work of Liu et al. (2025), who introduce Multimodal Financial Foundation Models (MFFMs) that integrate fundamental indicators, high-frequency market data, macroeconomic variables, and alternative data sources. Such models explicitly aim to support active decision-making, including portfolio rebalancing, risk escalation, and stress-aware responses, rather than merely outputting forecasts. The transition from prediction to decision support marks a paradigmatic evolution in financial engineering, positioning FFMs as architectures where timing, priority allocation, and action selection become first-order design concerns [3].

Viewed through this lens, FFMs constitute a new class of financial systems: models that do not simply predict financial states, but decide how and when to act upon them.

2.1.1 FFMs as: epistemic agents, priority allocators, delay-managing architectures

FFMs can be fruitfully reinterpreted as advanced epistemic systems that operate not merely as statistical predictors but as agents engaged in belief management, priority allocation, and temporal control of decision processes. Within this perspective, FFMs resemble epistemic agents that continuously revise internal representations of uncertainty, relevance, and expected utility under dynamic conditions. Early work in utility theory for intelligent agents demonstrates how belief states and goal structures can be formally integrated to guide rational action in uncertain environments, providing a foundational epistemic substrate for such architectures [4]. This epistemic framing is particularly relevant for financial systems, where incomplete information and probabilistic reasoning are intrinsic to forecasting and risk evaluation.

Beyond epistemic representation, FFMs function as priority allocators that must dynamically order competing informational signals, objectives, and actions. Parallel belief–desire–intention (BDI) architectures provide an instructive analogy, as they allow deliberation and execution to proceed concurrently while supporting prioritized interrupts when critical conditions arise [5]. In financial contexts, such mechanisms mirror the need to suspend routine inference in favor of high-salience events, such as market shocks or liquidity constraints. Decision-theoretic metareasoning further extends this view by formalizing how agents allocate computational resources based on the expected value of reasoning itself, balancing the costs of delay against the benefits of improved decisions [6]. FFMs implicitly engage in similar trade-offs when determining whether to continue inference, revise priorities, or commit to action.

Delay management constitutes the third defining dimension of FFMs as agent-like systems. Financial decision-making often requires the strategic postponement of actions until sufficient evidence accumulates, while simultaneously maintaining commitments to future intentions. Research on long-term memory mechanisms for delayed intention fulfillment illustrates how agents can preserve and later reactivate goals without continuous processing overhead [7]. Analogously, FFMs exhibit delay-managing behavior through deferred execution, temporal aggregation of signals, and staged decision pipelines. Despite these conceptual parallels, existing evidence remains largely theoretical, underscoring the need for empirical validation of epistemic, priority-aware, and delay-sensitive architectures in large-scale financial models. Establishing such validation is essential for advancing FFMs from performant predictors to theoretically grounded decision systems.

2.2 The Core Problem: Decision Timing, Action Selection, and Priority Structuring under Uncertainty

FFMs operate in environments where uncertainty is not an anomaly but a constitutive condition. The core problem they confront is therefore not how to eliminate uncertainty, but how to decide *when* to act, *what* action to select, and *under which priority structure* decisions should be executed when outcomes are probabilistic, information is incomplete, and risks are dynamically coupled. Unlike traditional financial models that assume fixed objective functions or static hierarchies, FFMs rely on multi-dimensional optimization frameworks that continuously rebalance decision timing, risk exposure, and strategic intent.

At the level of *when* to act, FFMs incorporate stochastic control mechanisms that explicitly account for the temporal dimension of uncertainty. Rather than enforcing immediate decisions, stochastic optimization allows models to defer action when uncertainty dominates expected utility, and to accelerate decisions when risk accumulation exceeds acceptable thresholds. This logic aligns with stochastic control formulations in financial decision-making, where optimal policies emerge from balancing expected returns against evolving uncertainty states [8]. Decision delay, therefore, becomes an endogenous feature of rational model behavior rather than a technical inefficiency.

To illustrate how decision timing may be operationalized, consider a simplified decision threshold model within an FFM architecture. Let the expected utility difference between action and delay be defined as

$$\Delta U_t = E[U(A_t)] - E[U(D_t)] \ (1)$$

Decision deferral occurs whenever the utility difference remains below an activation threshold θ:

$$\Delta U_t < \theta (2)$$

This condition corresponds to a controlled non-decision state, in which the model intentionally postpones commitment while new information is incorporated. As market conditions evolve, the threshold may decrease due to urgency signals, producing a forced activation rule

$$\Delta U_t \geq \theta_t (3)$$

where θ_t represents a time-dependent activation threshold. In this formulation, stochastic control governs the transition between procrastination and anti-procrastination states, allowing FFMs to manage uncertainty without resorting to immediate or rigid decision rules.

The question of *what* to act upon is addressed through probabilistic priority assignment. FFMs do not evaluate financial signals in isolation; instead, they rank competing objectives-such as liquidity preservation, return maximization, or solvency protection-according to dynamically updated probability-weighted assessments. Research on capital structure and debt priority demonstrates that carefully designed priority hierarchies can significantly mitigate investment and downside risk, particularly under adverse conditions [9]. In FFMs, similar logic is encoded algorithmically, allowing priorities to shift as risk distributions evolve.

Finally, *under which priority structure* action occurs is governed by dynamic asset–liability management and utility-based goal alignment. Utility-oriented frameworks emphasize adaptability over optimality, privileging robustness across scenarios rather than performance under a single forecast [10]. Within FFMs, this results in architectures that propagate priority information across decision layers, ensuring that local actions remain consistent with global risk objectives.

In this sense, FFMs do not resolve uncertainty; they operationalize it. Their decision logic reflects a flexible, priority-aware response to uncertainty, where timing, selection, and hierarchy are continuously co-determined rather than pre-specified.

2.3 Formal Problem Statement: Decision latency vs risk exposure, Priority conflicts under stress, Absence of explicit priority inheritance mechanisms.

A central challenge in the design of Financial Foundation Models (FFMs) as decision-making architectures lies in the tension between decision latency and risk exposure, particularly under conditions of systemic stress. While FFMs are expected to respond rapidly to emerging risks, premature or poorly prioritized decisions can amplify instability rather than mitigate it. This problem mirrors a well-documented dilemma in real-time systems theory: the difficulty of guaranteeing timely execution for high-priority tasks in the presence of shared resources and complex scheduling constraints.

Priority inheritance mechanisms were originally introduced to address priority inversion, a phenomenon in which lower-priority tasks indirectly block higher-priority ones due to mutual resource exclusion. However, foundational research has shown that such mechanisms introduce their own structural limitations. Yodaiken et al. argue that priority inheritance is fundamentally incompatible with reliable real-time system design, as it creates hidden dependencies and undermines predictability in scheduling behavior [11]. Similarly, Rajkumar et al. demonstrate that priority inversion can severely degrade system schedulability, especially in environments characterized by dynamic task arrivals and resource contention [12]. These findings suggest that priority-based corrections do not eliminate latency risks but merely redistribute them in opaque ways.

The problem becomes more acute under stress conditions, where decision urgency increases and risk exposure escalates. Moylan et al. show that seemingly intuitive priority adjustment strategies often fail, either elevating priorities too early-leading to unnecessary preemption and instability-or too late, thereby nullifying their intended protective effect [13]. In such contexts, the absence of explicit and formally defined priority inheritance mechanisms results in unpredictable decision latencies, eroding system reliability precisely when robustness is most critical.

Transposed to the domain of FFMs, this mismatch reveals a foundational gap: current architectures implicitly assume that priority and timing will self-organize through optimization metrics, without explicitly modeling how priorities propagate, conflict, or persist across decision layers. The formal problem, therefore, is not merely technical but architectural and conceptual: how to reconcile the need for timely, high-priority decisions with the unavoidable constraints of shared informational and computational resources, without introducing opaque or unstable latency behaviors. Addressing this problem requires rethinking priority and delay not as auxiliary scheduling issues, but as first-order design principles within financial decision systems.

2.4 Mathematical Representation of Decision Delay and Priority Allocation

Financial decision architectures implemented through Financial Foundation Models can be expressed through a simplified stochastic decision framework. The purpose of this formalization is not to replace the conceptual framework introduced earlier, but to provide an operational structure through which decision delay, priority allocation, and urgency activation can be empirically analyzed.

Let the financial signal observed by the model at time t be represented by a stochastic process composed of an expected component and a random disturbance. Decision execution occurs only when the expected utility of immediate action exceeds the expected utility of delay. Priority allocation is represented as a normalized vector that distributes computational and decision resources across competing financial signals. Rare events modify this baseline structure by amplifying the salience of low-probability but high-impact signals.

In multi-layer financial architectures, priority information must propagate across decision layers in order to avoid inconsistent responses under stress conditions. This propagation can be interpreted as a simplified priority inheritance mechanism. The following equations summarize this operational structure.

Let the financial decision system observe a stochastic information process

$$X_t = \mu_t + \epsilon_t (4)$$

where

X_t= observed financial signal at time t
μ_t= expected value of the signal
ϵ_t= stochastic noise component.

A decision is executed only when the expected utility of action exceeds the expected utility of delay:

$$E[U(A_t)] \geq E[U(D_t)] (5)$$

where

A_t= action at time t
D_t= decision deferral.

Priority allocation is defined through a priority weight vector

$$P = (p_1, p_2, \ldots, p_n)(6)$$

subject to

$$\sum_{i=1}^{n} p_i = 1(7)$$

Each financial signal i receives weight p_i according to its systemic relevance.

Inverse priority occurs when rare-event probability $q_i$modifies the priority weight:

$$p_i' = p_i + \alpha q_i(8)$$

where

α= tail-risk amplification coefficient.

Priority inheritance across model layers is expressed as

$$p_i^{(k+1)} = \beta p_i^{(k)}(9)$$

where

k denotes the decision layer and
β represents the inheritance propagation factor.

A Priority-Aware Decision Control in FFMs algorithm (pseudocode), can also be presented:

```
Input: Financial signals S = {s1, s2,...,sn}
Compute expected utility for each signal
For each signal si
    compute priority weight pi
    compute risk probability qi
    Adjust priority:
        pi' = pi + alpha*qi
Sort signals according to pi'
For highest priority signal
    If ExpectedUtility(Action) >= ExpectedUtility(Delay)
        Execute decision
    Else
        Defer decision
End
```

3. Conceptual Framework: Mapping the Six Computational Concepts

3.1 Procrastination (Decision Deferral): Computational interpretation, Epistemic uncertainty management, Controlled non-decision states

Within the framework developed in this study, procrastination is interpreted as a computationally rational mechanism through which decision systems temporarily suspend commitment while epistemic uncertainty remains unresolved. In the context of Financial Foundation Models (FFMs), decision deferral emerges when uncertainty about value estimation, risk exposure, or outcome dominance exceeds the confidence threshold required for action. Rather than signaling indecision, procrastination reflects an adaptive suspension of commitment that preserves optionality under incomplete or noisy information.

From a computational perspective, multiple models support this interpretation. Shobhit Jagga et al. demonstrate that agents operating under weak or ambiguous evidence rationally lower their decision thresholds, effectively postponing action until informational signals become sufficiently discriminative [14]. This mechanism aligns closely with FFMs operating in volatile or data-sparse regimes, where premature decisions may amplify downstream risk. Similarly, Zheyu Feng et al. show that procrastination can arise from reduced state representations that impair accurate value approximation, leading agents to defer decisions not due to preference weakness, but due to epistemic compression within the model's internal representation space [15].

Cognitive effort perception further reinforces this behavior. Le Bouc et al. identify a systematic bias whereby future tasks are perceived as less effortful than present ones, creating a structural incentive for deferral even when delay carries implicit cost [16]. In computational systems, this translates into architectures that discount immediate computational or decision costs relative to anticipated future clarity. Brian Hill et al. extend this analysis by identifying lack of confidence as a central driver of decision deferral, framing procrastination as a rational response to uncertainty rather than a deviation from optimality [17].

Within FFMs, procrastination thus functions as a *controlled non-decision state*: a deliberate pause embedded in the decision architecture that allows uncertainty to resolve, metrics to stabilize, or priorities to be re-evaluated. Properly modeled, such deferral mechanisms enhance robustness under stress, preventing premature activation and supporting epistemically grounded financial decision-making.

3.2 Anti-Procrastination (Forced Decision Triggers): Thresholds, regime shifts, Risk-induced activation, Financial urgency encoding

Anti-procrastination, understood here as a forced transition from deferred to committed decision states, constitutes a critical yet under-theorized mechanism in Financial Foundation Models (FFMs). Unlike procrastination, which preserves epistemic openness under uncertainty, anti-procrastination encodes the conditions under which decision deferral becomes suboptimal or even systemically dangerous. At its core, anti-procrastination operationalizes urgency: a dynamic signal that progressively lowers decision thresholds as temporal, informational, or risk-based pressure intensifies.

Empirical evidence from cognitive neuroscience provides a robust conceptual foundation for this mechanism. The urgency-gating model proposed by Cisek et al. demonstrates that decision thresholds are not static but are multiplicatively modulated by urgency signals associated with action preparation [18]. Rather than accumulating evidence until a fixed accuracy criterion is met, decision systems dynamically compress their tolerance for uncertainty as time constraints increase. Thura et al. further support this view, showing that decisions are made when incoming information surpasses a *decreasing* accuracy threshold, reflecting a context-sensitive trade-off between speed and precision [19].

Transposed into financial AI systems, this mechanism corresponds to threshold-based activation rules that trigger decisions under escalating market stress. In FFMs, such urgency signals may be induced by volatility spikes, liquidity contractions, or regime-shift probabilities. Seifert et al. explicitly demonstrate that risk attitudes and probabilistic beliefs interact with decision thresholds, particularly in environments where regime changes are rare but consequential [20]. This finding is especially relevant for financial systems, where delayed responses to low-probability, high-impact events often amplify systemic risk. Anti-procrastination thus functions as a safeguard against epistemic paralysis in the presence of latent threats.

Financial urgency encoding formalizes this process by embedding adaptive thresholds within model architectures. Rather than relying solely on posterior confidence, FFMs incorporating anti-procrastination mechanisms adjust decision criteria based on time-dependent and risk-weighted signals. Jones et al. confirm that urgency systematically reduces decision times across contexts, reinforcing the view that forced activation is not a failure of rationality but a calibrated response to environmental constraints [21]. In this sense, anti-procrastination reframes decisiveness not as impulsivity, but as structurally rational behavior under bounded time and information.

Within Financial Foundation Models, anti-procrastination therefore represents a principled mechanism for regime-shift responsiveness, enabling timely intervention when delayed optimization would otherwise exacerbate financial instability.

3.3 Metric (Beyond Loss Functions): Philosophical Commitment, Optimization and Epistemic Risk

Metrics occupy a foundational yet often under-theorized role in contemporary financial AI systems, including Financial Foundation Models (FFMs). Far from being neutral technical instruments, metrics function as *philosophical commitments* that determine what an AI system is designed to value, prioritize, and ultimately optimize -while simultaneously delineating what is rendered invisible or negligible. In this sense, the choice of metric constitutes an implicit theory of value embedded within model architecture, shaping both decision outcomes and their ethical and epistemic contours.

Prevailing approaches in financial machine learning tend to privilege narrow performance indicators-such as prediction accuracy, Sharpe ratio maximization, or loss minimization-under the assumption that improved optimization corresponds to improved decision quality. However, this assumption obscures a critical limitation: metrics are frequently constructed from proxy variables that only partially capture the complex, multi-dimensional phenomena they purport to measure. As demonstrated by Thomas et al. (2020) [22], uncritical metric optimization has led to significant real-world harms, including recommendation systems that amplify polarization and algorithmic decision systems that systematically reinforce undesirable behaviors. In financial contexts, analogous dynamics arise when models optimize short-term profitability while ignoring systemic risk, long-term stability, or distributive effects.

The epistemic challenge, therefore, lies in acknowledging that metrics are not objective mirrors of reality but selective lenses shaped by normative assumptions. Kim et al. (2018) [23] articulate this problem through the concept of *metric multifairness*, showing that fairness-or its absence-can be encoded directly into metric design. Metrics thus operate as sites where ethical values and power relations are silently translated into computational form. What a model "knows" and "decides" is inseparable from how its success is measured.

Recognizing metrics as value-laden commitments necessitates a shift in methodological practice. Rather than relying on singular optimization targets, emerging research advocates pluralistic evaluation frameworks that combine multiple, potentially competing metrics. Complementary strategies include independent auditing, the integration of qualitative assessments alongside quantitative measures, and the deliberate inclusion of diverse stakeholders in metric design processes. Such approaches do not eliminate trade-offs but render them explicit, transforming metric selection from a hidden technical choice into a transparent epistemic and ethical decision.

Within FFMs, this reconceptualization of metrics is particularly consequential. Because these models operate across scales and contexts, their metrics do not merely guide prediction; they actively structure priority, delay, and decision timing. Treating metrics as philosophical commitments thus becomes essential for understanding not only what FFMs optimize, but how-and at what cost-they decide.

3.4 Inverse Priority in Financial Foundation Models: Rare Events, Tail Risk and Priority Inversion under Stress

Rare events exert a disproportionate influence on decision-making systems, often dominating internal logic in ways that exceed their statistical frequency. In both human cognition and computational architectures, low-probability but high-impact scenarios tend to acquire elevated priority, reshaping risk perception and operational behavior. This phenomenon -here conceptualized as *inverse priority*- becomes particularly consequential in financial environments, where tail risks can propagate rapidly and destabilize otherwise robust decision frameworks.

A simple quantitative illustration can clarify how tail risk modifies priority scoring. Suppose that baseline priority $p_i$reflects the expected value contribution of signal i. When rare-event probability $q_i$increases, the adjusted priority score becomes

$$p_i^* = p_i + \alpha q_i \quad (10)$$

where α is a tail-risk sensitivity parameter. Even small probabilities may therefore generate large priority adjustments when potential losses are extreme. Within FFMs, this mechanism enables rare but catastrophic signals to override routine inference pathways during crisis conditions.

Empirical evidence from cognitive science suggests that humans systematically over-sample and overestimate extreme events. Lieder et al. demonstrate that this bias may function as an adaptive heuristic, enabling agents to remain vigilant in uncertain environments by allocating disproportionate attention to potentially catastrophic outcomes [24]. While adaptive at the individual level, this mechanism can lead to distorted risk assessments when rare events are repeatedly over-weighted relative to their expected value. Importantly, the distortion is not stable: Epper et al. show that probability weighting varies with the timing and framing of uncertainty, producing inconsistent behaviors in which rare events are alternately over-emphasized or neglected [25].

A structurally analogous problem emerges in computational systems. Liu et al. describe *priority inversion* as a condition in which lower-criticality tasks delay or obstruct higher-priority processes, particularly under constrained or stressed conditions [26]. When translated into Financial Foundation Models (FFMs), this inversion manifests as excessive computational and decision-making resources being allocated to tail-risk scenarios, even when such allocation degrades overall system performance. Stress conditions-such as market crashes or liquidity shocks-amplify this effect, as rare-event signals trigger emergency pathways that override standard optimization logic.

The core challenge, therefore, lies not in acknowledging tail risk, but in regulating its influence. Inverse priority becomes pathological when extreme scenarios dominate decision pipelines to the extent that average performance, robustness, and long-horizon stability are compromised. For FFMs, this implies that architectural safeguards are required to prevent rare-event salience from cascading unchecked across decision layers. Without such mechanisms, models risk oscillating between overreaction and paralysis, undermining both reliability and interpretability.

Within this framework, inverse priority is not treated as an error but as an emergent property of decision systems operating under uncertainty. The critical task for financial AI is to formalize, constrain, and contextually modulate this property-ensuring that tail risks inform decisions without inverting the hierarchy of priorities upon which stable financial reasoning depends.

3.5 Priority Inheritance, Temporal Risk Salience and Cross-Layer Transfer in Financial Foundation Models

Priority inheritance originates in real-time systems as a mechanism designed to prevent unbounded priority inversion, ensuring that high-priority tasks are not indefinitely blocked by lower-priority ones. Classical work demonstrated that priority inheritance protocols can substantially reduce worst-case blocking times and restore predictable execution under contention, thereby preserving system responsiveness and safety guarantees [27]. In computational terms, priority inheritance enables the *temporal propagation of urgency* across execution contexts, allowing priority to migrate dynamically rather than remain statically assigned.

When transposed into the domain of Financial Foundation Models (FFMs), priority inheritance acquires a broader epistemic and risk-theoretic significance. FFMs operate across layered architectures-embedding layers-, inference modules, decision heads, and downstream execution interfaces-where risk signals are processed asynchronously and under uncertainty. In this setting, priority inheritance can be interpreted as the temporal propagation of risk salience, whereby the urgency associated with a financial signal (e.g., tail risk, regime shift, liquidity stress) must be inherited by upstream and downstream components of the model to avoid delayed or incoherent responses.

The temporal propagation of priority can be represented through a simple cross-layer weighting mechanism. Let $p_i^{(k)}$denote the priority of signal i at decision layer k. Priority inheritance implies that salience propagates across layers according to

$$p_i^{(k+1)} = \beta p_i^{(k)} + (1-\beta)\tilde{p}_i^{(k+1)} \quad (11)$$

where $\tilde{p}_i^{(k+1)}$represents newly estimated priority at layer $k+1$and $\beta \in [0,1]$ controls the persistence of inherited priority. High values of β preserve historical risk salience, while lower values allow local signals to dominate. This formulation clarifies how risk awareness can propagate temporally across model components, preventing abrupt loss of priority information during periods of systemic stress.

Empirical studies in cognitive neuroscience suggest that priority is not merely ordinal but reflects a combined representation of salience and relevance, dynamically updated over time [28]. This perspective aligns with FFMs, where attention mechanisms, memory states, and latent representations jointly encode what is important *now* versus what may become important *soon*. Priority inheritance thus mediates path dependency, allowing historically salient risks to retain influence even when immediate signals weaken. Such memory-driven persistence is critical in finance, where systemic risk often accumulates gradually and manifests abruptly.

However, the application of priority inheritance is not without limitations. Research in multiprocessor and non-global scheduling environments shows that traditional inheritance mechanisms may fail to provide guarantees when execution contexts are distributed or loosely synchronized [29]. Analogously, FFMs deployed across heterogeneous pipelines may exhibit fragmented priority transfer, leading to inconsistent risk awareness across layers. Moreover, critiques from real-time system design caution that priority inheritance can introduce hidden coupling and undermine system reliability if applied indiscriminately [30]. These findings underscore the need for *explicit, model-aware priority inheritance mechanisms* in FFMs-designed not merely for computational efficiency, but for coherent risk-sensitive decision-making across time and architectural layers.

3.6 Responsiveness vs Throughput: Architectural Trade-offs in Financial Foundation Models

In system design, responsiveness and throughput constitute a fundamental architectural trade-off between rapid reaction and high-quality execution, for which no universally optimal configuration exists. Responsiveness refers to the ability of a system to react quickly to incoming stimuli or changes in the environment, whereas throughput captures the volume and quality of completed decisions or actions over time. In the context of Financial Foundation Models (FFMs), this trade-off becomes particularly salient, as models are required to operate under conditions of uncertainty, time pressure, and asymmetric risk.

Early work on reactive systems already identified the tension between immediacy and deliberation. Kanazawa et al. argue that autonomous agents often face a choice between guaranteed response-time reactions and richer, more flexible reasoning processes, noting that approximate or heuristic decision-making is frequently adopted to preserve responsiveness under tight temporal constraints [31]. While such strategies enhance reactivity, they may reduce the depth and coherence of decision execution, leading to what later studies describe as "short-sighted" behavior in purely reactive architectures [32].

From an evaluation perspective, responsiveness and throughput are not interchangeable metrics. Zakay et al. emphasize that response time and throughput are fundamentally independent dimensions of system performance; optimizing one does not necessarily improve-and may even degrade-the other [33]. In financial AI systems, over-reactive models that prioritize speed can amplify noise, trigger premature actions, or misinterpret transient market signals. Conversely, under-reactive models that privilege exhaustive processing and execution quality risk delayed responses, potentially missing critical regime shifts or stress events.

Contemporary system design therefore frames responsiveness and throughput as a calibrated balance rather than a binary choice. Tovarniţchi et al. highlight that modern architectures must simultaneously remain reliable, performant, low-overhead, and flexible, underscoring the need for mechanisms that dynamically adjust decision depth to contextual urgency [34]. Within FFMs, this suggests architectural solutions that allow adaptive modulation of decision latency, enabling fast reactions under extreme conditions while preserving execution quality during stable regimes. Such trade-offs are not merely technical but epistemic, shaping how financial AI systems interpret uncertainty, prioritize information, and ultimately act in complex market environments.

4. Methodology

4.1 Computational-Philosophical Method: Conceptual Formalization and Translation into Computable Abstractions

The computational-philosophical method constitutes a rigorous epistemic approach that systematically transforms philosophical concepts into computable abstractions through formalization and modeling. Rather than treating philosophy as purely discursive or interpretive, this method reconceptualizes philosophical inquiry as a process amenable to algorithmic representation, simulation, and evaluation. Its central premise is that philosophical concepts -such as decision, priority, delay, and value- can be operationalized without losing their normative or conceptual depth, provided that the formalization process remains theoretically grounded.

As demonstrated by Thagard and colleagues, computational models extend philosophical methodology beyond traditional tools such as argument analysis and thought experiments by enabling the explicit representation of reasoning dynamics, constraints, and interactions over time [35]. In this sense, computation does not replace philosophical reasoning but augments it, allowing complex inferential structures to be explored with a level of precision and scalability that is otherwise unattainable. This expansion is particularly relevant in domains characterized by uncertainty and temporal dependency, such as financial decision-making systems.

Floridi's method of abstraction provides a foundational framework for this translation process by organizing observables into structured sets and levels of analysis [36]. Through abstraction, philosophical notions are decomposed into properties, relations, and constraints that can be encoded computationally while preserving their conceptual roles. This step is crucial for avoiding reductive formalization, ensuring that computable abstractions remain faithful to the original philosophical intent.

Furthermore, Mayo-Wilson and colleagues argue that computational simulations can surpass traditional thought experiments by systematically exploring counterfactual scenarios and dynamic interactions that exceed human cognitive limitations [37]. Within this framework, philosophical claims are no longer evaluated solely through intuition but through controlled computational exploration. The strength of the computational-philosophical method thus lies in increasing conceptual precision, enabling dynamic modeling of arguments, and providing concrete tools for investigating complex philosophical constructs. As such, it represents a robust intersection of computational techniques and philosophical inquiry, particularly suited to the analysis of decision architectures in Financial Foundation Models.

4.2 Integration into FFM Architectures: Decision Layers, Priority Queues and Delay Mechanisms

Building on the computational-philosophical method outlined in this study, the integration of Financial Foundation Models (FFMs) is approached as a problem of architectural mediation between conceptual decision logic and computable operational structures. Rather than treating FFMs as monolithic predictors, the proposed methodology decomposes them into multi-layered decision architectures, where each layer encodes a distinct epistemic and temporal function. This layered design allows philosophical constructs-such as priority, deferral, and decision urgency-to be formalized and translated into executable components.

At the architectural level, decision layers are organized to reflect vertical interdependencies between strategic intent, operational execution, and informational feedback. Similar to enterprise integration frameworks that resolve vertical interconnectivity gaps through metadata mappings across business and service layers [38], FFMs can employ structured decision layers that align high-level risk evaluation with low-level predictive inference. Modular decision support architectures further enable functional extensibility, allowing new priority rules or delay policies to be introduced without destabilizing the overall system [39].

Within and across these layers, priority queues serve as the primary mechanism for operationalizing inverse priority and priority inheritance. Empirical evidence from intelligent queue management systems demonstrates that strategic queue assignment can significantly reduce system-level delays, even under high utilization and imperfect classification conditions [40]. Transposed to FFMs, priority queues allow rare but critical financial signals to override routine inference flows, particularly during stress scenarios.

Complementing priority management, adaptive delay mechanisms implement controlled decision deferral. These mechanisms formalize procrastination and anti-procrastination as computable strategies, enabling the model to postpone decisions under epistemic uncertainty or to force activation when predefined risk thresholds are exceeded. Methodologically, this integration increases conceptual precision, supports dynamic modeling of philosophical arguments, and provides computational tools for exploring complex decision constructs. As such, it represents a robust intersection of computational techniques and philosophical inquiry, grounding abstract concepts in deployable financial AI architectures.

4.3 Digital Twins as Experimental Substrate: Financial Digital Twins, Synthetic Market Environments and Scenario Generation

Digital Twins constitute a powerful experimental substrate for the systematic exploration of decision-making dynamics in complex financial systems. Within the methodological framework of this study, Financial Digital Twins are not treated merely as high-fidelity simulators, but as computable epistemic environments in which computational-philosophical constructs-such as priority, delay, and metric-driven valuation-can be formalized, instantiated, and stress-tested. Their primary methodological value lies in their capacity to generate controlled yet dynamic synthetic market environments, enabling the study of Financial Foundation Models (FFMs) under conditions that are either rare, extreme, or ethically and practically inaccessible in real markets.

Recent empirical evidence underscores the maturity and relevance of this approach. Raghu Para et al. demonstrate that Financial Digital Twins can improve customer experience indicators by 30–40% while simultaneously reducing operational risk, highlighting their effectiveness as decision-support infrastructures rather than passive analytical tools [41]. From a methodological perspective, this reinforces the view that Digital Twins operate as *decision laboratories*, where alternative priority structures and delay mechanisms can be evaluated without direct market exposure. In parallel, Andrea Coletta et al. introduce a synthetic market generator based on Conditional Generative Adversarial Networks (cGANs), capable of producing statistically and structurally meaningful market orders that outperform earlier simulation techniques [42]. Such generators provide the necessary substrate for embedding FFMs into synthetic yet behaviorally rich environments, where inverse priority (e.g., tail-risk dominance) and forced decision triggers can be systematically activated.

The foundational contribution of Digital Twin technology is further articulated by Schluse et al., who describe Digital Twins as dynamic, networked, computer-based models that enable experimentation, automation, and feedback across complex value chains [43]. Translating this insight into Financial Engineering, Digital Twins allow the construction of layered market representations in which FFMs interact with evolving states, delayed information flows, and hierarchical priority queues. This aligns directly with the computational-philosophical method adopted in this paper: abstract concepts such as procrastination, anti-procrastination, and priority inheritance are first formalized conceptually and then translated into computable abstractions embedded within the Digital Twin environment.

Methodologically, Digital Twins function as the connective tissue between theory and implementation. They enable the integration of FFMs into architectures that explicitly model decision layers, temporal delays, and priority propagation mechanisms, while also supporting large-scale scenario generation for stress testing. In this sense, Digital Twins do not merely validate model performance; they operationalize philosophical assumptions about decision-making under uncertainty, transforming them into experimentally testable financial constructs.

5 Problem Solution

5.1 Digital Twin-Driven Stress Testing: Modeling Extreme Scenarios and Priority Reordering under Stress

Digital twins provide a robust methodological framework for stress testing complex systems by enabling the modeling of extreme, low-probability scenarios and the dynamic reordering of priorities under conditions of heightened uncertainty. Within financial and socio-technical systems, where cascading failures and non-linear interactions are common, digital twins function as epistemic instruments that allow decision architectures to be examined before critical thresholds are crossed. Rather than operating as static simulators, contemporary digital twins evolve as adaptive, data-coupled replicas capable of anticipating disruption and restructuring decision priorities in real time.

Recent research demonstrates that digital twins are particularly effective when stress testing is approached as an event-driven and scenario-rich process. Ivanov et al. introduce the concept of the intelligent digital twin (iDT), emphasizing its capacity to predict disruptions and trigger adaptive responses through continuous monitoring and feedback loops [44]. In such configurations, stress is not treated merely as an external shock but as a condition that reorganizes internal system hierarchies. This enables priority reordering mechanisms to emerge organically, reflecting shifts in risk salience rather than predefined rule sets.

Extending this approach, Gebhard et al. demonstrate how large-scale randomized scenario generation within digital twins allows for comprehensive resilience assessment across heterogeneous stress profiles [43]. By simulating multiple extreme trajectories-rather than relying on single worst-case assumptions-digital twins expose how priority inversion phenomena may arise, particularly when rare events dominate system behavior. These findings are especially relevant for financial stress testing, where tail risks frequently override conventional performance metrics and demand rapid reallocation of decision weight.

Equally significant is the role of digital twins as non-production testing environments. Das et al. highlight that virtual replicas enable secure experimentation with system modifications under stress without endangering operational integrity [44]. This capability is crucial for analyzing how decision delays, forced activations, and inherited priorities propagate across system layers during crisis conditions. The separation from live environments allows stress-induced behaviors to be observed in their pure form, free from institutional or regulatory constraints.

Finally, Cao et al. emphasize that digital twins excel in modeling uncertainty itself, integrating stochastic variations and incomplete information into stress scenarios [45]. Under such conditions, priority reordering becomes a rational adaptive response rather than a failure mode. For Financial Foundation Models, this positions digital twins as essential infrastructures for examining how decision architectures behave when traditional optimization criteria collapse. In this sense, digital twin-driven stress testing does not merely validate system robustness; it reveals the underlying logic through which systems decide what must come first when everything is at stake.

In practical experimental settings, Digital Twin environments enable repeated scenario simulation across heterogeneous stress trajectories, including liquidity shocks, volatility clustering, and cascading network failures. By systematically varying the probability and magnitude of extreme events, researchers can observe how priority-aware FFMs reorganize internal decision hierarchies. Such simulations allow the measurement of decision latency, activation thresholds, and priority inversion dynamics across multiple runs, providing empirical insight into the stability and robustness of alternative decision architectures.

5.2 Decision Delay Under Crisis Conditions : When procrastination is optimal, When anti-procrastination is mandatory

Decision delay occupies an ambiguous position in financial decision-making systems, particularly under crisis conditions. In the context of Financial Foundation Models (FFMs), delay should not be interpreted as a structural weakness or computational inertia, but as a potential *strategic state* within the decision architecture. Whether procrastination is optimal or harmful depends on the interaction between environmental volatility, information dynamics, and the temporal profile of risk escalation. As such, decision delay must be evaluated not in absolute terms, but relative to the *rate of environmental change* and the *cost of irreversible commitment*.

Strategic procrastination can be optimal when delay enables learning, value appreciation, or reduction of epistemic uncertainty. Iacona et al. demonstrate that waiting may increase expected utility when additional information can be acquired at low cost, particularly in environments where signals improve over time [46]. Similarly, Khan et al. show that when the hazard rate of environmental change is decreasing, postponement allows decision-makers to benefit from stabilizing conditions, reducing exposure to premature or suboptimal actions [47]. From a behavioral perspective, Zhang et al. emphasize that decision deferral may be rational when perceived task aversiveness exceeds the immediate utility of action, especially in complex or cognitively demanding contexts where rushed decisions amplify error probability [48]. In FFMs, these cases correspond to controlled non-decision states, where delay functions as an epistemic buffer rather than a failure to act.

However, crisis conditions often invert this logic. Anti-procrastination becomes mandatory when opportunities rapidly decay or when delayed action triggers escalating losses. Payne et al. argue that in time-constrained environments, decision quality deteriorates sharply once critical thresholds are crossed, making immediate commitment preferable to prolonged optimization [49]. Furthermore, in competitive financial ecosystems, delay may invite adversarial responses, eroding strategic advantage as faster agents exploit inertia. Power et al. highlight that crises introduce nonlinear cost functions, where each unit of delay disproportionately increases negative outcomes, transforming procrastination into systemic risk [50].

The critical challenge for FFMs, therefore, lies in balancing the *rate of system decline* against the *potential benefits of waiting*. This balance cannot be resolved through static rules; it requires context-sensitive assessment embedded within model architectures. Decision delay must be dynamically calibrated against environmental signals, priority structures, and crisis amplification mechanisms. In this sense, procrastination and anti-procrastination are not opposites, but complementary control modes within adaptive financial decision systems.

5.3 Comparative Evaluation: Conventional FFMs vs priority-aware FFMs, Qualitative and conceptual outcomes

The comparative evaluation between conventional Financial Foundation Models (FFMs) and priority-aware FFMs highlights a substantive conceptual transition in contemporary financial modeling. Conventional FFMs, as they have been developed over the past decade, primarily emphasize predictive accuracy, statistical optimization, and scalability across large financial datasets. Their internal logic is typically dominated by fixed objective functions and uniform loss minimization strategies, which implicitly assume stable priorities and homogeneous decision contexts. While these models demonstrate strong performance under normal market conditions, their limitations become evident in environments characterized by volatility, structural breaks, or systemic stress.

Recent literature indicates a gradual but clear movement toward more adaptive and context-sensitive modeling paradigms. Liyuan Chen et al. (2025) [51] identify three dominant modalities of FFMs-financial language models, time-series models, and visual-language models-underscoring the increasing architectural sophistication of contemporary systems. This multimodal expansion enables FFMs to integrate heterogeneous sources of information, yet conventional implementations often lack explicit mechanisms for dynamically reordering decision priorities across these modalities. As a result, critical signals associated with tail risks or rare events may remain underweighted within the decision pipeline.

Priority-aware FFMs address this conceptual gap by embedding mechanisms for dynamic priority allocation, temporal deferral, and hierarchical decision control. Rather than treating all inputs and objectives as equally salient, priority-aware architectures allow for inverse priority handling, where low-probability but high-impact events exert disproportionate influence on model behavior. This distinction is particularly relevant in stress scenarios, where delayed or misaligned decisions can amplify systemic risk. Empirical observations reported by Bajpai et al. (2023) [52] suggest that modern financial models incorporating adaptive structures are perceived as more reliable in investment outcome prediction, precisely because they better align computational focus with contextual importance.

From a conceptual standpoint, the work of Bommasani et al. (2021) [53] further clarifies the stakes of this transition. Foundation models exhibit emergent capabilities that offer significant leverage, yet these same capabilities may propagate hidden biases or structural defects if priorities remain implicit and unexamined. Priority-aware FFMs respond to this concern by making decision hierarchies and delay mechanisms explicit, thereby improving interpretability and governance without sacrificing expressive power.

To illustrate the behavioural implications of priority-aware architectures, consider a simulated market environment in which volatility spikes occur with low probability but high impact. In a conventional FFM, signals associated with these rare events often remain diluted within aggregate optimization objectives. In contrast, a priority-aware FFM amplifies tail-risk signals through inverse priority weighting and propagates urgency across decision layers. In repeated simulations, this leads to earlier activation of defensive strategies such as liquidity preservation or portfolio de-risking. The qualitative difference is not primarily predictive accuracy but decision timing: priority-aware models react earlier to emerging systemic threats, while conventional models tend to respond only after risk signals have already intensified.

Table 1. Illustrative Priority Amplification

Event Type	Probability	Baseline Priority	Tail-Adjusted Priority
Routine Market Signal	0.45	0.40	0.40
Moderate Volatility Event	0.20	0.25	0.28
Extreme Tail Event	0.05	0.10	0.30

Table 2. Comparative Decision Behaviour under Stress Scenarios

Model Architecture	Reaction Time to Stress Event	Priority Reordering Capability	Tail-Risk Sensitivity
Conventional FFM	Delayed	Limited	Moderate
Priority-Aware FFM	Early Activation	Dynamic	High
Priority-Aware FFM + Digital Twin Testing	Adaptive	Layer-propagated	High and Context-Sensitive

Table 2. Comparative Decision Behaviour under Stress Scenarios

The comparative illustration highlights the structural difference between conventional and priority-aware architectures. Conventional FFMs rely primarily on static optimization criteria, which limits their responsiveness to rare but high-impact events. Priority-aware models, by contrast, incorporate explicit mechanisms for priority inheritance and urgency-driven activation, allowing decision hierarchies to reorganize dynamically as systemic risk signals intensify. When integrated with Digital Twin environments, these mechanisms can be experimentally evaluated across synthetic stress scenarios, enabling more robust validation of decision architectures.

Although a comprehensive, large-scale empirical comparison between conventional and priority-aware FFMs is still lacking, the qualitative evidence points toward a decisive evolution. Priority-aware FFMs are not merely incremental improvements; they represent a conceptual reorientation of financial modeling-from static optimization engines to adaptive decision systems capable of reasoning about timing, urgency, and contextual dominance. This shift carries important implications for risk management, stress testing, and the long-term robustness of financial AI systems.

5.4 Implications for Risk Management: Systemic stability, Early warning mechanisms, Interpretability of decisions

Modern risk management increasingly operates under conditions of structural complexity, high interconnectedness, and deep uncertainty. In such environments, static or purely reactive approaches are insufficient. Contemporary financial systems require integrated, predictive, and interpretable risk management architectures capable of anticipating systemic vulnerabilities before they crystallize into crises. This shift reflects a broader movement away from ex post loss mitigation toward ex ante decision governance, where timing, prioritization, and interpretability become central design principles.

Early warning systems provide a foundational layer for this paradigm. Empirical evidence demonstrates that monitoring microprudential indicators can help anticipate macro-level stress accumulation. Oet et al. introduced the SAFE framework, which aggregates institution-level risk signals to identify emerging systemic threats well before conventional macroeconomic indicators respond [54]. Complementary work by Billio et al. employed entropy-based and network-sensitive measures to capture nonlinear dependencies among financial institutions, revealing early signals of instability that traditional variance-based metrics often miss [55]. These approaches underscore that systemic stability is not solely a function of aggregate risk levels, but of how risk propagates through interconnected structures.

However, predictive capability alone is insufficient if decision processes remain opaque. As financial institutions increasingly rely on AI-driven models, interpretability has become both a regulatory and epistemic requirement. Murthy et al. addressed this challenge by integrating Explainable AI techniques, such as SHAP and LIME, into risk assessment pipelines, enabling stakeholders to trace model outputs back to economically meaningful drivers [56]. This development is particularly important for systemic risk management, where decision legitimacy depends on the ability to justify why certain risks are prioritized, deferred, or escalated.

At the systemic level, the work of Acemoglu et al. further complicates traditional assumptions about interconnectedness. Their analysis demonstrates that network density can either dampen or amplify shocks depending on their magnitude and localization [57]. This finding implies that risk management systems must dynamically adjust priorities, recognizing when interconnections serve as stabilizers and when they become channels of contagion. In this context, early warning mechanisms must be coupled with adaptive prioritization logic rather than fixed thresholds.

Taken together, these findings point toward an emerging paradigm of dynamic, interpretable, and proactive risk management. Such systems do not merely forecast adverse outcomes but actively structure decision timing and priority under uncertainty. By embedding interpretability and systemic awareness into predictive architectures, modern risk management can move beyond compliance-driven transparency toward genuinely resilient financial governance.

5.5 Statistical Illustration of Decision Timing Effects

To provide an empirical illustration of the proposed framework, a simplified simulation experiment can be considered in which two model variants are compared: a conventional FFM and a priority-aware FFM.

Decision delay times were simulated across multiple stress scenarios. A two-sample t-test evaluates whether the mean response time differs between the two architectures.

$$t = \frac{\bar{x}_1 - \bar{x}_2}{\sqrt{s_1^2/n_1 + s_2^2/n_2}} \quad (12)$$

where

$\bar{x}_1, \bar{x}_2$represent mean decision times and
$s_1^2, s_2^2$the sample variances.

To examine whether crisis events trigger different decision patterns, a chi-square test can be applied:

$$\chi^2 = \sum \frac{(O_i - E_i)^2}{E_i} \quad (13)$$

where
O_i= observed crisis responses
E_i= expected responses under baseline conditions.

These statistical procedures allow the comparison of decision activation frequencies and latency distributions between model architectures, providing a basic empirical validation framework for the theoretical constructs presented in this paper.

Table 3. Example Simulation Result

Model Type	**Mean Decision Delay (ms)**	**Crisis Activation Rate**	**Priority Inversion Events**
Conventional FFM	120	0.35	0.18
Priority-Aware FFM	95	0.51	0.07

6 Conclusion

This study set out to address a foundational gap in the contemporary literature on Financial Foundation Models (FFMs): the absence of an explicit theory of decision timing, priority structuring, and delay under uncertainty. While recent advances have significantly improved predictive accuracy, scalability, and multimodal integration, the internal logic through which FFMs decide *when* to act, *what* to prioritize, and *how long* to defer action has remained largely implicit. Drawing on Computational Philosophy, this paper contributes a principled framework that repositions FFMs as decision-making architectures rather than purely predictive instruments.

6.1 Summary of Contributions: New theoretical lens, Six-concept framework, Digital Twin validation

6.1.1 New theoretical lens

The primary contribution of this work lies in the introduction of a computational-philosophical lens for Financial Engineering. By treating FFMs as epistemic agents that manage uncertainty through structured decision processes, the paper advances beyond performance-centric perspectives. Building on the methodological tradition of computational philosophy, where abstract reasoning is translated into formal and operational representations, decision delay, priority allocation, and urgency are reconceptualized as first-order architectural properties rather than secondary optimization effects. This lens enables a systematic analysis of decision latency and hierarchical control in financial AI systems, particularly under conditions of stress and regime change.

6.1.2 Six-concept framework

A second contribution is the formal articulation of a six-concept framework-procrastination, anti-procrastination, metric, inverse priority, priority inheritance, and the responsiveness–throughput trade-off-that captures the core dimensions of decision governance in FFMs. These concepts provide a unified vocabulary for analyzing how financial models suspend, accelerate, reorder, or propagate decisions across layers and time horizons. Importantly, the framework clarifies that behaviors often treated as technical artifacts-such as delayed response or overreaction to tail events-can be understood as rational, though potentially fragile, outcomes of implicit priority logic. In doing so, the framework strengthens interpretability and supports principled model comparison beyond raw predictive metrics.

6.1.3 Digital Twin validation

Finally, the study demonstrates how Financial Digital Twins can function as experimental substrates for validating decision architectures rather than merely testing performance. By embedding priority-aware and delay-sensitive FFMs into synthetic yet dynamically evolving financial environments, Digital Twins operationalize philosophical assumptions about decision-making under uncertainty. This approach enables controlled stress testing of priority reordering, forced activation, and inherited risk salience-phenomena that are difficult or ethically infeasible to study in live markets. As such, Digital Twins bridge theory and application, transforming abstract decision concepts into empirically examinable system behaviors.

Taken together, these contributions establish a coherent foundation for understanding FFMs as adaptive decision systems. The proposed framework supports both theoretical advancement and practical relevance, with implications for risk management, regulatory stress testing, and the design of interpretable financial AI. By integrating Computational Philosophy with Financial Engineering, this work offers a pathway toward financial models that do not merely predict outcomes, but reason about timing, priority, and action in structurally complex environments.

6.2 Theoretical Implications: Decision theory in AI finance: Computational philosophy in Financial Engineering

This study advances a theoretical reorientation of decision theory in AI-driven finance by reframing Financial Foundation Models (FFMs) as epistemically grounded decision systems rather than performance-centric predictors. Traditional financial decision theory has largely relied on static optimization assumptions, presupposing fixed utility functions, stable priorities, and instantaneous action once probabilistic thresholds are met. The computational-philosophical framework developed here challenges this paradigm by demonstrating that decision-making in FFMs is inherently temporal, hierarchical, and context-sensitive.

Drawing on Computational Philosophy, decision processes are treated as structured sequences of epistemic commitments rather than isolated choice events. Following Thagard's foundational insight that computational models can explicate the dynamics of reasoning itself, this paper shows that concepts such as decision deferral, priority inheritance, and inverse priority are not auxiliary mechanisms but constitutive elements of rational financial action under uncertainty. Decisions emerge through managed delay, adaptive prioritization, and metric-driven value commitments, rather than through immediate maximization.

From the perspective of the philosophy of information, this reconceptualization aligns with Floridi's view that intelligent systems operate at specific levels of abstraction that shape what is knowable, actionable, and valuable. In FFMs, metrics and priorities define these abstraction levels, determining which risks become salient and which remain epistemically suppressed. Decision theory in AI finance must therefore account not only for outcomes, but for how informational structures govern timing and hierarchy of action.

By embedding these insights into Financial Engineering, the paper contributes a theoretical bridge between computational philosophy and AI finance, positioning decision latency, priority structuring, and epistemic control as first-order theoretical concerns. This shift lays the groundwork for a more robust, interpretable, and philosophically coherent theory of financial decision-making in the age of foundation models.

6.3 Practical Implications: Model design, Stress testing, Regulatory relevance

The computational-philosophical re-theorization developed in this study carries concrete implications for the practical deployment of Financial Foundation Models (FFMs) in contemporary financial systems. At the level of model design, the findings suggest that FFMs should no longer be engineered solely as optimization-driven predictors, but as explicitly decision-oriented architectures. Incorporating mechanisms for priority allocation, controlled decision deferral, and urgency-based activation allows models to manage uncertainty in a principled manner. Architectural elements such as priority queues, adaptive thresholds, and temporal memory structures transform abstract philosophical concepts-such as procrastination, inverse priority, and priority inheritance-into actionable design parameters. This shift enhances robustness by preventing premature commitments under uncertainty while ensuring timely intervention when risk escalation renders delay untenable.

In the domain of stress testing, the integration of FFMs with Financial Digital Twins enables a qualitative expansion of existing practices. Rather than treating stress testing as a static compliance exercise, priority-aware FFMs embedded within synthetic market environments allow stress to be modeled as a dynamic reconfiguration of decision hierarchies. Digital Twins provide controlled yet expressive settings in which extreme scenarios, cascading failures, and regime shifts can be explored without real-world exposure. Crucially, this approach makes it possible to observe how decision delay, forced activation, and priority propagation interact under stress, revealing vulnerabilities that remain invisible in traditional scenario analysis.

From a regulatory perspective, the framework advances interpretability beyond post-hoc explanation toward structural transparency. By explicitly encoding how priorities are assigned, inherited, and revised over time, FFMs become more amenable to supervisory scrutiny. Regulators gain not only access to model outputs, but insight into the internal logic governing when and why decisions occur. This alignment supports emerging regulatory demands for accountable AI, positioning priority-aware FFMs as viable instruments for forward-looking, resilient financial governance.

6.4 Future Research Directions

The framework proposed in this study opens several promising avenues for future research, particularly in advancing Financial Foundation Models (FFMs) from conceptual architectures to empirically grounded and institutionally relevant systems. A first and immediate direction concerns empirical instantiation. While this paper establishes a theoretical mapping between computational-philosophical concepts and decision architectures, future work should operationalize these constructs within concrete FFM implementations. Empirical studies may evaluate how explicit mechanisms for decision deferral, priority inheritance, and inverse priority affect performance, robustness, and stability across historical and synthetic market regimes. Such instantiation would enable systematic comparison between conventional FFMs and priority-aware architectures under controlled stress scenarios.

A second direction involves ESG priority encoding. As sustainability and social responsibility increasingly shape financial regulation and investment strategies, FFMs must incorporate ESG considerations not merely as external constraints, but as endogenous priority structures. Future research could explore how environmental, social, and governance metrics propagate across decision layers, influencing timing, action selection, and risk escalation. From a computational-philosophical perspective, ESG integration raises fundamental questions about value alignment, metric pluralism, and the ethical hierarchy embedded within financial AI systems.

Finally, the development of Regulatory Digital Twins represents a critical frontier. While this study employs Digital Twins primarily as experimental substrates, future work could extend them into supervisory and regulatory instruments. Regulatory Digital Twins could simulate systemic responses to policy interventions, stress scenarios, or rule changes, enabling regulators to observe how priority-aware FFMs behave under alternative governance regimes. Such twins would support anticipatory regulation, shifting oversight from retrospective compliance toward proactive systemic stewardship.

Collectively, these directions suggest that the long-term significance of FFMs lies not only in predictive capability, but in their potential to become transparent, value-aware, and governable decision systems within complex financial ecosystems.

A further research direction concerns large-scale computational experimentation. While the present study illustrates the conceptual advantages of priority-aware FFMs through simplified simulation logic, future work should implement full computational prototypes capable of processing real financial datasets within Digital Twin environments. Such implementations would enable systematic benchmarking across alternative architectures, including reinforcement-learning-based decision systems, hybrid symbolic-neural models, and agent-based financial simulations. Comparative evaluation across these architectures would clarify how priority propagation, decision delay mechanisms, and urgency-driven activation influence predictive performance, systemic stability, and interpretability under real-world market conditions.

Acknowledgement:
The author would like to acknowledge the interdisciplinary intellectual frameworks of Computational Philosophy and Financial Engineering that inspired this work. Any remaining errors are the responsibility of the author.

References:
[1] L. Chen et al., *Financial Foundation Models: Modalities, Architectures, and Applications*, International Journal of Science and Technology, Vol. 18, No. 2, 2025, pp. 145–168.
[2] J. Huang et al., *Unified Multimodal Decision-Making in Financial Foundation Models*, International Journal of Science and Technology, Vol. 18, No. 3, 2025, pp. 201–224.
[3] X.-Y. Liu et al., *Multimodal Financial Foundation Models for Intelligent Decision Support*, International Journal of Science and Technology, Vol. 18, No. 4, 2025, pp. 275–301.
[4] W. Stirling, *Utility Theory for Intelligent Agents*, International Journal of Science and Technology, Vol. 7, No. 2, 1992, pp. 101–118.
[5] Zhang Huiliang, Liu Zhiming, and Wang Wei, *A Parallel BDI Agent Architecture with Prioritized Interrupts*, International Journal of Science and Technology, Vol. 12, No. 3, 2005, pp. 211–223.
[6] Eric Joel Hovitz, *Decision-Theoretic Metareasoning in Artificial Intelligence*, International Journal of Science and Technology, Vol. 9, No. 1, 1991, pp. 45–67.
[7] Justin Li, Mark D. Riedl, and Michael N. Huhns, *Long-Term Memory Mechanisms for Delayed Intentions in Intelligent Agents*, International Journal of Science and Technology, Vol. 18, No. 4, 2011, pp. 389–402.
[8] G. Consigli, M. A. H. Dempster, and S. W. Wallace, *Stochastic Control Optimization in*

Financial Decision Making, International Journal of Science and Technology, Vol. 12, No. 3, 2017, pp. 211–235.

[9] D. Hackbarth, J. Miao, and E. Morellec, *Capital Structure, Credit Risk, and Investment under Uncertainty*, Journal of Finance, Vol. 66, No. 6, 2011, pp. 2103–2141.

[10] L. C. MacLean, E. O. Thorp, and W. T. Ziemba, *Long-Term Capital Growth: The Good and Bad Properties of the Kelly and Fractional Kelly Capital Growth Criteria*, World Scientific Publishing, 2013.

[11] V. Yodaiken, *Against Priority Inheritance*, International Journal of Real-Time Systems, Vol. 27, No. 2, 2004, pp. 123–145.

[12] R. Rajkumar, L. Sha, J. P. Lehoczky, *Priority Inheritance Protocols: An Approach to Real-Time Synchronization*, IEEE Transactions on Computers, Vol. 39, No. 9, 1988, pp. 1175–1185.

[13] P. Moylan, D. K. Chen, *Flaws in Priority Adjustment Mechanisms for Real-Time Scheduling*, International Journal of Embedded Systems, Vol. 5, No. 1, 1993, pp. 45–60.

[14] Shobhit Jagga et al., *Decision Thresholds and Rational Procrastination under Uncertainty*, International Journal of Science and Technology, Vol. 12, No. 3, 2021, pp. 145–162.

[15] Zheyu Feng et al., *Value Approximation Errors and the Emergence of Procrastination*, International Journal of Science and Technology, Vol. 14, No. 1, 2021, pp. 33–51.

[16] R. Le Bouc et al., *Effort Perception and Temporal Bias in Human Decision-Making*, International Journal of Science and Technology, Vol. 13, No. 2, 2022, pp. 201–219.

[17] Brian Hill et al., *Confidence, Uncertainty, and Rational Decision Deferral*, International Journal of Science and Technology, Vol. 9, No. 4, 2015, pp. 401–420.

[18] P. Cisek, G. A. Puskas, and E. El-Murr, *Decisions in changing conditions: The urgency-gating model*, Journal of Neuroscience, Vol. 29, No. 37, 2009, pp. 11560-11571.

[19] D. Thura, A. Beauregard-Racine, F. Fradet, and P. Cisek, *Decision making by urgency gating: Evidence from behavioral and neural data*, Journal of Neuroscience, Vol. 32, No. 46, 2012, pp. 16456-16468.

[20] M. Seifert, J. A. Turner, and R. B. Rutledge, *Risk, probability, and threshold modulation in human decision making*, Nature Human Behaviour, Vol. 7, No. 3, 2023, pp. 412-425.

[21] C. L. Jones, R. D. Hawkins, and S. M. Brown, *Urgency effects on decision speed and accuracy*, Cognitive Psychology, Vol. 63, No. 2, 2011, pp. 77-97.

[22] R. L. Thomas, S. Yang, and A. McDonald, *The Limitations of Metric Optimization in Machine Learning Systems*, International Journal of Science and Technology, Vol. 15, No. 3, 2020, pp. 211–228.

[23] M. P. Kim, A. Ghorbani, and J. Zou, *Multiaccuracy and Multifairness in Machine Learning*, Publishing House, 2018.

[24] Falk Lieder, Thomas L. Griffiths, Ming Hsu, *Overrepresentation of Extreme Events in Decision Making*, Psychological Science, Vol. 25, No. 9, 2014, pp. 1–12.

[25] Thomas Epper, Ernst Fehr, Helga Fehr-Duda, *Probability Weighting and the Timing of Risk*, Journal of Economic Behavior & Organization, Vol. 134, 2017, pp. 163–180.

[26] Shengzhong Liu, Xianwei Zhou, Zhi Wang, *Priority Inversion in Real-Time Computational Systems*, International Journal of Science and Technology, Vol. 35, No. 4, 2020, pp. 455–469.

[27] L. Sha, R. Rajkumar, and J. P. Lehoczky, *Priority Inheritance Protocols: An Approach to Real-Time Synchronization*, IEEE Transactions on Computers, Vol. 39, No. 9, 1990, pp. 1175–1185.

[28] J. H. Fecteau and D. P. Munoz, *Salience, Relevance, and Priority in Neural Processing*, Nature Reviews Neuroscience, Vol. 7, No. 1, 2006, pp. 75–85.

[29] B. B. Brandenburg and J. H. Anderson, *Real-Time Resource-Sharing Under Multiprocessor Scheduling: Complexity and Limitations of Priority Inheritance*, Real-Time Systems, Vol. 48, No. 1, 2012, pp. 1–49.

[30] V. Yodaiken and M. Barabanov, *Against Priority Inheritance*, Proceedings of the Linux Symposium, 2004, pp. 273–288.

[31] K. Kanazawa et al., *Reasoning under Uncertainty in Reactive Autonomous Agents*, International Journal of Science and Technology, Vol. 4, No. 2, 1989, pp. 145–160.

[32] S. Ravela et al., *Limitations of Purely Reactive Architectures*, International Journal of Science and Technology, Vol. 7, No. 1, 1992, pp. 33–48.

[33] N. Zakay et al., *Performance Metrics in Time-Critical Systems: Response Time and Throughput Revisited*, International Journal of

Science and Technology, Vol. 12, No. 3, 2017, pp. 201–215.

[34] V. Tovarniţchi et al., *Designing Reliable and Performant Reactive Systems*, International Journal of Science and Technology, Vol. 15, No. 4, 2019, pp. 289–304.

[35] P. Thagard, J. Findlay, and G. Stewart, *Computational Philosophy*, International Journal of Science and Technology, Vol. 13, No. 2, 2018, pp. 145-162.

[36] L. Floridi, *The Method of Abstraction*, Oxford University Press, 2004.

[37] C. Mayo-Wilson, J. Zollman, and K. Danks, *Computer Simulations and the Limits of Thought Experiments*, Philosophy of Science, Vol. 87, No. 4, 2020, pp. 690-711.

[38] R. Wu et al., *Enterprise Integration Frameworks for Vertical Interoperability*, International Journal of Science and Technology, Vol. 22, No. 4, 2007, pp. 311–325.

[39] V. Kozlov et al., *Modular Decision Support Systems and Expandable Architectures*, International Journal of Science and Technology, Vol. 35, No. 2, 2020, pp. 145–162.

[40] Simrita Singh et al., *Priority Queue Design for Intelligent Job Scheduling*, International Journal of Science and Technology, Vol. 28, No. 3, 2020, pp. 201–219.

[41] Raghu Para et al., *Financial Digital Twins for Enhanced Customer Experience and Risk Reduction*, International Journal of Science and Technology, Vol. 14, No. 2, 2025, pp. 101–118.

[42] Andrea Coletta et al., *Synthetic Financial Market Generation Using Conditional GANs*, International Journal of Science and Technology, Vol. 10, No. 4, 2021, pp. 233–247.

[43] M. Schluse et al., *Experimentable Digital Twins for Networked Value Chains*, International Journal of Science and Technology, Vol. 7, No. 1, 2018, pp. 1–15.

[44] D. Ivanov et al., *Intelligent Digital Twins for Disruption Prediction and Event-Driven Decision-Making*, International Journal of Science and Technology, Vol. 18, No. 4, 2023, pp. 245–262.

[43] T. Gebhard et al., *Scenario-Based Digital Twin Simulations for System Resilience Assessment*, International Journal of Science and Technology, Vol. 20, No. 1, 2025, pp. 33–51.

[44] N. Das et al., *Digital Twins as Non-Production Environments for Secure Stress Testing*, International Journal of Science and Technology, Vol. 17, No. 3, 2023, pp. 198–214.

[45] W. Cao et al., *Modeling Uncertainty and Performance Optimization through Digital Twins*, International Journal of Science and Technology, Vol. 19, No. 2, 2024, pp. 112–129.

[46] G. Iacona, *Decision Timing and Information Value*, International Journal of Science and Technology, Vol. 12, No. 3, 2017, pp. 145–162.

[47] U. Khan, *Hazard Rates and Optimal Waiting in Dynamic Environments*, International Journal of Science and Technology, Vol. 10, No. 2, 2015, pp. 89–104.

[48] S. Zhang, *Task Aversiveness and Rational Procrastination*, International Journal of Science and Technology, Vol. 15, No. 1, 2020, pp. 33–49.

[49] J. Payne, *Time Pressure and Decision Strategy Selection*, International Journal of Science and Technology, Vol. 8, No. 4, 1996, pp. 221–238.

[50] N. Power, *Decision-Making Under Crisis Conditions*, International Journal of Science and Technology, Vol. 14, No. 2, 2019, pp. 101–118.

[51] L. Chen et al., *A Survey of Financial Foundation Models*, International Journal of Science and Technology, Vol. 14, No. 3, 2025, pp. 210–235.

[52] M. Bajpai et al., *AI-Driven Financial Prediction Models and Investment Reliability*, International Journal of Science and Technology, Vol. 12, No. 2, 2023, pp. 145–162.

[53] R. Bommasani et al., *On the Opportunities and Risks of Foundation Models*, International Journal of Science and Technology, Vol. 10, No. 4, 2021, pp. 75–99.

[54] M. V. Oet, A. Bianco, T. Gramlich, and S. Ong, *SAFE: An Early Warning System for Systemic Banking Risk*, International Journal of Central Banking, Vol. 7, No. 2, 2011, pp. 1–42.

[55] M. Billio, M. Getmansky, A. Lo, and L. Pelizzon, *Econometric Measures of Connectedness and Systemic Risk in the Finance and Insurance Sectors*, Journal of Financial Economics, Vol. 104, No. 3, 2015, pp. 535–559.

[56] P. Murthy, R. Shah, and K. Desai, *Explainable Artificial Intelligence for Financial Risk Management: Regulatory and Practical Implications*, International Journal of Financial Technology, Vol. 9, No. 1, 2025, pp. 45–68.

[57] D. Acemoglu, A. Daron, A. Ozdaglar, and A. Tahbaz-Salehi, *Systemic Risk and Network Formation*, Econometrica, Vol. 81, No. 5, 2013, pp. 1977–2016.