SUMTYME.AI, ADE JINADU

Proposal and Comment Information

Title: Regulatory Capital Rules: Regulatory Capital and Standardized Approach for Risk-weighted Assets, R-1888

Comment ID: FR-2026-0008-01-C15

Submitter Information

Organization Name: sumtyme.ai

Organization Type: Company

Name: Ade Jinadu

Submitted Date: 04/18/2026

Subject: Comment Letter on Regulatory Capital Rules: Standardised Approach [Docket No. R-1888]

Pursuant to the Administrative Procedure Act (APA), 5 U.S.C. § 553, we submit the attached comment to identify a significant Less Restrictive Alternative (LRA) to the proposed modifications of the standardised approach for Risk-Weighted Assets (RWA).

The "Complexity" Fallacy: In Section III.A.2, the Agencies suggest that operational burden prevents the adoption of more sensitive risk measures. We challenge this premise. The Causal Inception Layer (CIL) provides high-fidelity risk detection with the operational simplicity of a standardised approach. As a parameter-free, zero-training methodology, the CIL eliminates the dependency on historical data or periodic retraining that often leads to model drift and regulatory uncertainty.

Key Technical Interventions:

- Deterministic Risk Modelling: Unlike Expected Shortfall (ES), which acts as a reactive stochastic indicator, the CIL models volatility as a deterministic trajectory from the point of causal inception.

- Lead-Time Advantage: Our empirical "Propagation Proofs" establish a significant lead-time advantage. By documenting a 35-day lead time during the COVID-19 pandemic, these proofs demonstrate the capacity for proactive risk mitigation before systemic triggers are breached.

- Phase-In Waiver: A Risk-Inception Framework (RIF) provides immediate auditability and structural certainty; we propose that the 5-year phase-in period (Section VI.F) is unnecessary for portfolios governed by a RIF, supporting the Board's 2026 Economic Efficiency objectives.

- Regulatory Recommendations: We urge the Board to authorise a RIF as a "Supplemental Overlay" to justify risk-sensitive downward adjustments of the 95% corporate risk-weight floor. Furthermore, we suggest an exemption from the "Operational Risk Add-On" for portfolios governed by a verified RIF, as deterministic logs effectively mitigate model-error risk.

The CIL addresses the systemic misallocation of Tier 1 capital currently held in lagging buffers to compensate for model latency. Failure to consider this technically viable and economically significant alternative would constitute a failure to consider a critical aspect of the regulatory problem.

Respectfully submitted,

Ade Jinadu

Research Lead, sumtyme.ai

(Supporting empirical data available in the attached Appendix)

April 18th 2026

Subject: Comment Letter on Regulatory Capital Rules: Regulatory Capital and Standardised Approach for Risk-weighted Assets [Docket No. R-1888]

1. Introduction & Statement of Less Restrictive Alternative (LRA)

Pursuant to the Administrative Procedure Act (APA), 5 U.S.C. § 553, we submit this comment to identify a significant Less Restrictive Alternative (LRA) to the proposed modifications of the standardised approach for Risk-Weighted Assets (RWA). We support the effort to "de-duplicate" and simplify the capital framework into a single risk-based stack. However, we believe that the current reliance on Expected Shortfall (ES) and static risk weights, while representing a modernisation, remains a structural limitation.

We propose a Risk-Inception Framework (RIF) called the Causal Inception Layer (CIL) as a supplemental enhancement. Unlike current measures that rely on lagging indicators, the CIL identifies risk at inception. This allows for a more granular regulatory framework that avoids the restrictive capital buffers traditionally required to compensate for model uncertainty and risk detection delay, thereby optimising the balance between systemic stability and market utility.

2. Technical Critique: The "Complexity" Fallacy

In Section III.A.2, the Agencies suggest that "increased complexity and operational burden" prevent the adoption of more sensitive risk measures for standardised banks. We challenge this premise. The CIL provides the high-fidelity detection of an advanced framework with the operational simplicity of a standardised approach.

As a parameter-free, zero-training methodology, the CIL eliminates dependency on historical data, market context, or periodic retraining. Unlike the proposed Market Risk rules that rely on Expected Shortfall (ES), a reactive stochastic indicator that is blind to the causal drivers of risk, the CIL models volatility as a deterministic trajectory from the point of causal inception.

Furthermore, the proposed 5 year phase-in period (Section VI.F) is an implicit admission of the transitional instability inherent in the current proposal. As a RIF provides immediate auditability and structural certainty, we propose that this phase-in period is unnecessary for portfolios governed by a RIF, allowing for Day One implementation without sacrificing safety or soundness.

3. Risk-Inception Framework: Causal Inception Layer (CIL)

The CIL is a Risk-Inception Framework designed to establish that market volatility is deterministic. It is the visible acceleration of a causal shift that can be mapped with precision from its point of inception. The CIL identifies these shifts and maps the directional trajectory of any publicly traded asset without reliance on historical training data or market context. By leveraging a continuous data stream of publicly traded prices, the framework eliminates periodic retraining and maintains constant structural alignment with market reality. This approach transforms risk from a probabilistic estimate into a deterministic calculation.

The CIL functions as an independent observer of global markets, leveraging bidirectional temporal feedback loops to overcome the paradox of reflexivity. It provides banking organisations with the optionality to hedge incrementally, bypassing the synchronised selling that regulators aim to prevent. Every directional change is captured through these feedback loops and is auditable from the tick level, with the additional option of initiating observations at later stages of the evolution.

The CIL is governed by a core axiom: all structural changes in financial markets originate at the tick level before propagating across expanding timeframes. The framework operates through four distinct pillars:

- **Causal Identification:** The CIL detects the initial signal at the tick level to identify the precise moment equilibrium shifts. While traditional statistical approaches are forced to wait for "statistical significance" and subsequently categorise these signals as noise, the CIL captures the transition at its point of inception.
- **Structural Propagation:** A signal’s validity is contingent on successfully propagating through an ascending temporal hierarchy (Tick → Milliseconds → Seconds → Minutes → Hours → Days).
- **Minimum Threshold:** The CIL requires a minimum propagation threshold of one. If a shift fails to propagate to the adjacent temporal layer, it is identified as a non-causal shift and is discarded as it lacks the structural maturity to define a path.
- **Trajectory Projection:** While traditional models produce a widening ‘probabilistic cone’ of uncertainty, the CIL generates a non-linear price trajectory based on confirmed propagations.

Every propagation can be logged from the tick-level inception point, allowing risk officers to provide a deterministic rationale for every action. A hedge that is subsequently unwound is not a ‘false positive’ but a strategic insurance premium. It is a cost-efficient measure to avoid the destructive feedback loops of synchronised market panic. For example, a bank may initiate an Inception Phase Hedge (e.g. 10% exposure) once a signal achieves 5 minute temporal maturity, followed by a Propagation Phase Hedge at 30 minutes.

4. Comprehensive Recommendations for the Final Rule

To support the Board’s dual mandate of safety and economic growth, we recommend the following:

1. **Authorisation of a RIF as a "Supplemental Overlay":** Allow institutions to apply a RIF (such as the CIL) as a validation layer to justify a risk-sensitive downward adjustment of the 95% corporate risk-weight floor (Section III.A.2), where an inception-level hedging programme is demonstrated to mitigate risk before it manifests in broader market volatility.
2. **Elimination of the Operational Risk Add-On:** The proposed "nominal add-on to account for operational risk" is a buffer against model uncertainty. As the CIL provides deterministic logs which are explainable, auditable and traceable, it effectively mitigates model-error risk. We suggest an exemption from this add-on for portfolios governed by a verified RIF.
3. **Recognition of Temporal Maturity:** Amend Market Risk rules to acknowledge that capital charges should be calibrated based on the dynamic temporal maturity and propagation of a risk signal, rather than relying solely on static historical correlations.
4. **Immediate Implementation Option (Phase-In Waiver):** Allow banks to bypass the 5-year phase-in (Section VI.F) for portfolios governed by a verified RIF. This supports the Board’s 2026 Economic Efficiency objectives by unlocking Tier 1 capital immediately without sacrificing safety and soundness.
5. **AOCI Volatility Mitigation Credit:** For Category III and IV banks (Section II.B), provide a capital credit where the institution can demonstrate, via deterministic RIF logging, that it identified the causal inception of interest rate shifts and initiated hedges prior to systemic propagation.

5. Conclusion

The Causal Inception Layer is explainable, auditable and traceable. It addresses the systemic misallocation of billions in Tier 1 capital currently held in lagging buffers to compensate for model latency. By integrating a RIF-based approach, the Agencies can finalise a rule that is not only simpler but structurally prepared for the high-frequency realities of the modern financial landscape.

A Risk-Inception Framework represents a Pareto improvement over the current proposal; it enhances systemic safety by eliminating detection lag while simultaneously reducing the suppressed capital that currently constrains credit availability. Failure to consider such a technically viable and economically significant alternative would constitute a failure to consider a critical aspect of the problem [1] and would overlook a vital path toward optimising the resilience of the U.S. banking system.

Respectfully submitted,

Ade Jinadu

Research Lead at sumtyme.ai

ade.jinadu@sumtyme.ai

Technical documentation and supporting empirical data are available in the Appendix, with further granular documentation available upon request for the Board's quantitative staff.

Footnotes:

[1] See 5 U.S.C. § 706(2)(A): https://www.law.cornell.edu/uscode/text/5/706

APPENDIX

Propagation Proofs: CIL Lead-Time vs. Expected Shortfall (ES)

These proofs quantify the structural detection latency gap between the CIL and traditional 10 day Expected Shortfall (ES) models currently utilised by the Agencies. The data establishes that by the time ES-based triggers breach regulatory thresholds, the CIL has already confirmed structural propagation across the temporal hierarchy.

During high-volatility events, the CIL identified the precise moment of causal inception weeks before traditional measures reacted. In the case of the COVID-19 pandemic, the framework provided a 35 day lead time, while the Q4 2025 Crypto Decline was identified with a 37 day lead time. This structural advantage enables non-procyclical risk mitigation long before traditional regulatory triggers are breached. Furthermore, during the Liberation Day Tariff Announcement, ES remained directionally blind despite a 124.6% price move, whereas the CIL mapped a 121.9% causal trajectory from detection.

Deterministic Capital Efficiency

Unlike ES, which consistently overstates risk in positive market conditions, the CIL prevents the unnecessary lock-up of Tier 1 capital by identifying shifts that lack structural maturity. This allows for an Accelerated Resilience Track where capital efficiency is realised on Day One without sacrificing systemic safety.

Event	Ticker	Drawdown	CIL Mitigation
April 2025 Tariff Announcement	VIXY	124.6%	98%
Yen Carry Trade Unwind	USD/JPY	83.0%	99%
COVID-19 Pandemic	MGM	13.7%	99%
Q4 2025 Crypto Decline	BTCUSDT	36.1%	98%
2026 Iran War	UAL	26.9%	97%

Note: *We have started the observation of directional trajectories from the 1 second level for these examples, however the true inception point is at the tick level.*

Event	Liberation Day Tariff Announcement
Ticker	ProShares VIX Short-Term Futures ETF
Change Duration	48 Days
Change (%)	124.6%
% of Change Explained by the CIL	121.9%
Change Identified by the CIL	98%
Observation	ES is directionally blind and overstates risk.

The following data establishes the deterministic causal trajectory of VIXY's time series. It demonstrates the structural propagation of change from the 1 second observation through the temporal hierarchy, transforming market volatility into a traceable and deterministic calculation.

Propagation Level	Datetime	Price	% Change From 1s	Propagation Path
0	2025-02-21 14:05:38	40.7	0.00	Observed from 1s
1	2025-02-21 14:30:30	40.855	0.38	1s → 5s
2	2025-02-21 14:31:30	40.855	0.38	5s → 15s
3	2025-02-21 14:31:30	40.8813	0.45	15s → 30s
4	2025-02-21 14:32:00	40.95	0.61	30s → 1m
5	2025-02-21 18:00:00	42.33	4.00	1m → 5m
6	2025-02-21 18:10:00	42.72	4.96	5m → 15m
7	2025-02-25 10:30:00	44.45	9.21	15m → 30m
8	2025-03-10 04:00:00	51.16	25.70	30m → 60m
9	2025-04-03 10:00:00	54.96	35.04	60m → 120m
10	2025-04-03 20:00:00	61.16	50.27	120m → 240m
11	2025-04-04 00:00:00	60.7	49.14	240m → 360m
12	2025-04-07 00:00:00	73.16	79.75	360m → 720m
13	2025-04-09 12:01:45	90.33	121.94	720m → Trough

The following chart quantifies the Detection Latency Gap by contrasting the CIL's performance with the Expected Shortfall (ES) model for VIXY's time series. It establishes that while ES remained directionally blind to the 124.6% price move amplified by the Tariff Announcement, the CIL confirmed structural propagation to explain 121.9% of the causal trajectory. This demonstrates the CIL's capacity to provide a deterministic rationale for risk mitigation even when traditional statistical measures fail to trigger.

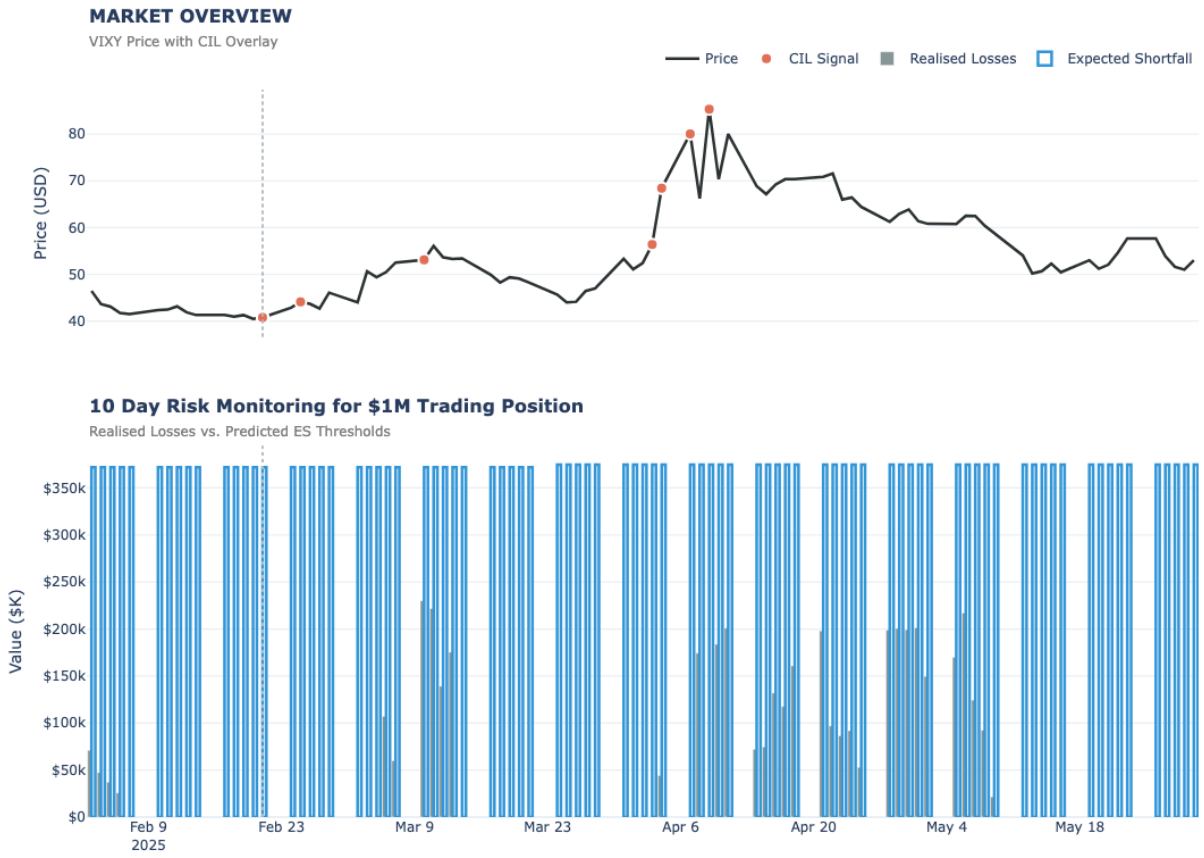

Event	COVID-19 Market Volatility
Ticker	MGM Resorts
Drawdown Duration	60 Days
Drawdown (%)	82.95%
% of Change Explained by the CIL	82.46%
Drawdown Mitigated by the CIL	99%
Detection Lead Time over Expected Shortfall	35 Days

The following data establishes the deterministic causal trajectory of MGM's time series. It demonstrates the structural propagation of risk from the 1 second observation through the temporal hierarchy, transforming market volatility into a traceable and deterministic calculation.

Propagation Level	Datetime	Price	% Change From 1s	Propagation Path
0	2020-01-21 13:46:31	33.65	0.00	Observed from 1s
1	2020-01-21 13:54:40	33.65	0.00	1s → 5s
2	2020-01-21 13:54:45	33.65	0.00	5s → 15s
3	2020-01-21 13:55:00	33.65	0.00	15s → 30s
4	2020-01-21 13:55:00	33.65	0.00	30s → 1m
5	2020-01-21 14:40:00	33.075	-1.71	1m → 5m
6	2020-01-21 19:00:00	32.715	-2.78	5m → 15m
7	2020-01-22 16:00:00	32.21	-4.28	15m → 30m
8	2020-01-27 11:00:00	29.92	-11.08	30m → 60m
9	2020-02-25 16:00:00	29.18	-13.28	60m → 120m
10	2020-02-27 12:00:00	26.37	-21.63	120m → 240m
11	2020-02-28 12:00:00	25	-25.71	240m → 360m
12	2020-03-06 00:00:00	20.59	-38.81	360m → 720m
13	2020-03-18 17:30:23	5.90	-82.46	720m → Trough

The following chart quantifies the Detection Latency Gap by contrasting the CIL's performance with the Expected Shortfall (ES) model for MGM's time series. It establishes that the CIL confirmed structural propagation 35 days before traditional regulatory triggers reacted, providing an immediate and auditable signal for risk mitigation while ES remained directionally blind to the inception of the COVID-19 market volatility.

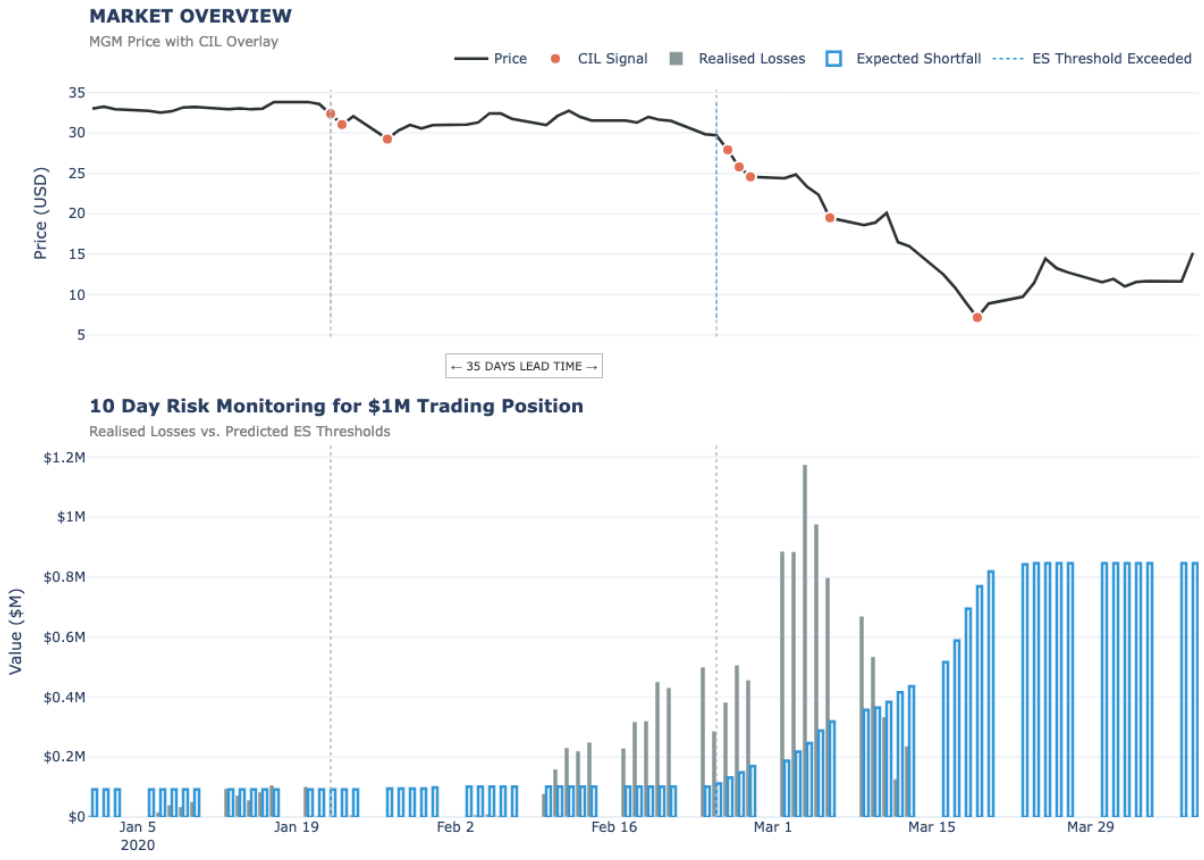

Event	Yen Carry Trade Unwind
Ticker	USD/JPY
Drawdown Duration	74 Days
Drawdown (%)	13.73%
% of Change Explained by the CIL	13.57%
Drawdown Mitigated by the CIL	99%
Detection Lead Time over Expected Shortfall	7 Days

The following data establishes the deterministic causal trajectory of USD/JPY's time series. It demonstrates the structural propagation of risk from the 1 second observation through the temporal hierarchy, transforming market volatility into a traceable and deterministic calculation.

Propagation Level	Datetime	Price	% Change from 1s	Propagation Path
0	2024-07-11 07:16:07	161.707	0.00	Observed from 1s
1	2024-07-11 07:20:40	161.677	-0.02%	1s → 5s
2	2024-07-11 07:23:20	161.655	-0.03%	5s → 10s
3	2024-07-11 07:24:45	161.636	-0.04%	10s → 15s
4	2024-07-11 07:36:30	161.574	-0.08%	15s → 30s
5	2024-07-11 12:31:00	160.864	-0.52%	30s → 1m
6	2024-07-11 12:45:00	158.752	-1.83%	1m → 5m
7	2024-07-11 12:50:00	158.503	-1.98%	5m → 10m
8	2024-07-11 13:00:00	158.642	-1.90%	10m → 15m
9	2024-07-11 15:00:00	158.55	-1.95%	15m → 30m
10	2024-07-23 18:00:00	155.684	-3.72%	30m → 60m
11	2024-08-02 14:00:00	147.055	-9.06%	60m → 120m
12	2024-08-05 04:00:00	144.356	-10.73%	120m → 240m
13	2024-09-13 00:00:00	141.512	-12.49%	240m → 360m
14	2024-09-16 09:00:00	139.75	-13.57	360m → Trough

The following chart quantifies the Detection Latency Gap by contrasting the CIL's performance with the Expected Shortfall (ES) model for USD/JPY's time series. While ES remains directionally blind to the inception of the carry trade unwind, the CIL confirms structural propagation to provide an immediate, auditable signal for non-procyclical risk mitigation.

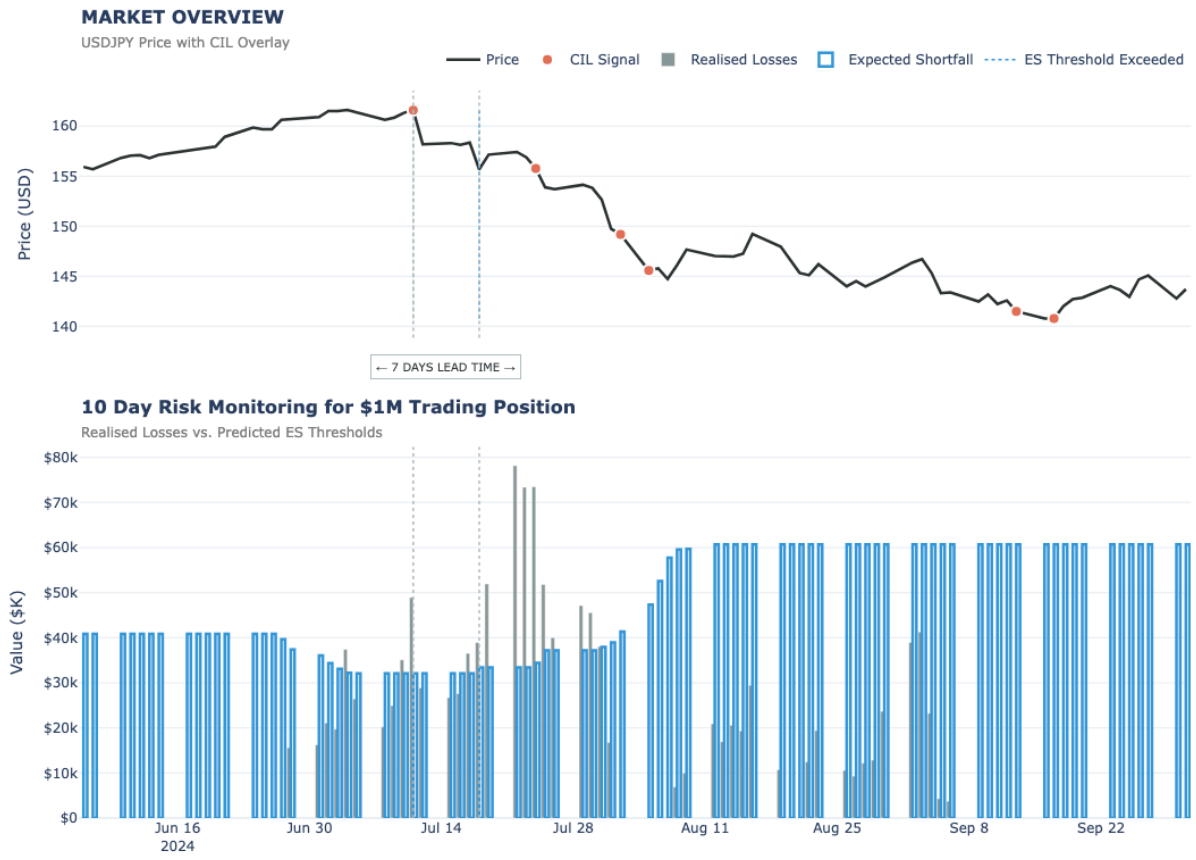

Event	Q4 2025 Crypto Decline
Ticker	BTCUSDT
Drawdown Duration	46 Days
Drawdown (%)	36.13%
% of Change Explained by the CIL	35.34%
Drawdown Mitigated by the CIL	98%
Detection Lead Time over Expected Shortfall	37 Days

The following data establishes the deterministic causal trajectory of BTCUSDT's time series. It demonstrates the structural propagation of risk from the 1 second observation through the temporal hierarchy, transforming market volatility into a traceable and deterministic calculation.

Propagation Level	Datetime	Price	% Change from 1s	Propagation Path
0	2025-10-07 13:31:15	124664.09	0	Observed from 1s
1	2025-10-07 13:39:35	124211.71	-0.36	1s → 5s
2	2025-10-07 13:42:30	124100	-0.45	5s → 15s
3	2025-10-07 13:43:00	124036.69	-0.5	15s → 30s
4	2025-10-07 13:44:00	123838.8	-0.66	30s → 1m
5	2025-10-07 14:10:00	123088.36	-1.26	1m → 2m
6	2025-10-07 23:35:00	121537.99	-2.5	2m → 5m
7	2025-10-10 17:00:00	118205.61	-5.18	5m → 10m
8	2025-10-16 14:30:00	110343.42	-11.48	10m → 30m
9	2025-10-16 21:00:00	107838	-13.49	30m → 60m
10	2025-11-04 18:00:00	106457.24	-14.6	60m → 120m
11	2025-11-21 12:29:25	80600	-35.34	120m - Trough

The following chart quantifies the Detection Latency Gap by contrasting the CIL's performance with the Expected Shortfall (ES) model for BTCUSDT's time series. It establishes that the CIL confirmed structural propagation 37 days before traditional regulatory triggers reacted, providing an auditable signal for proactive mitigation while ES remained directionally blind to the accelerating drawdown.

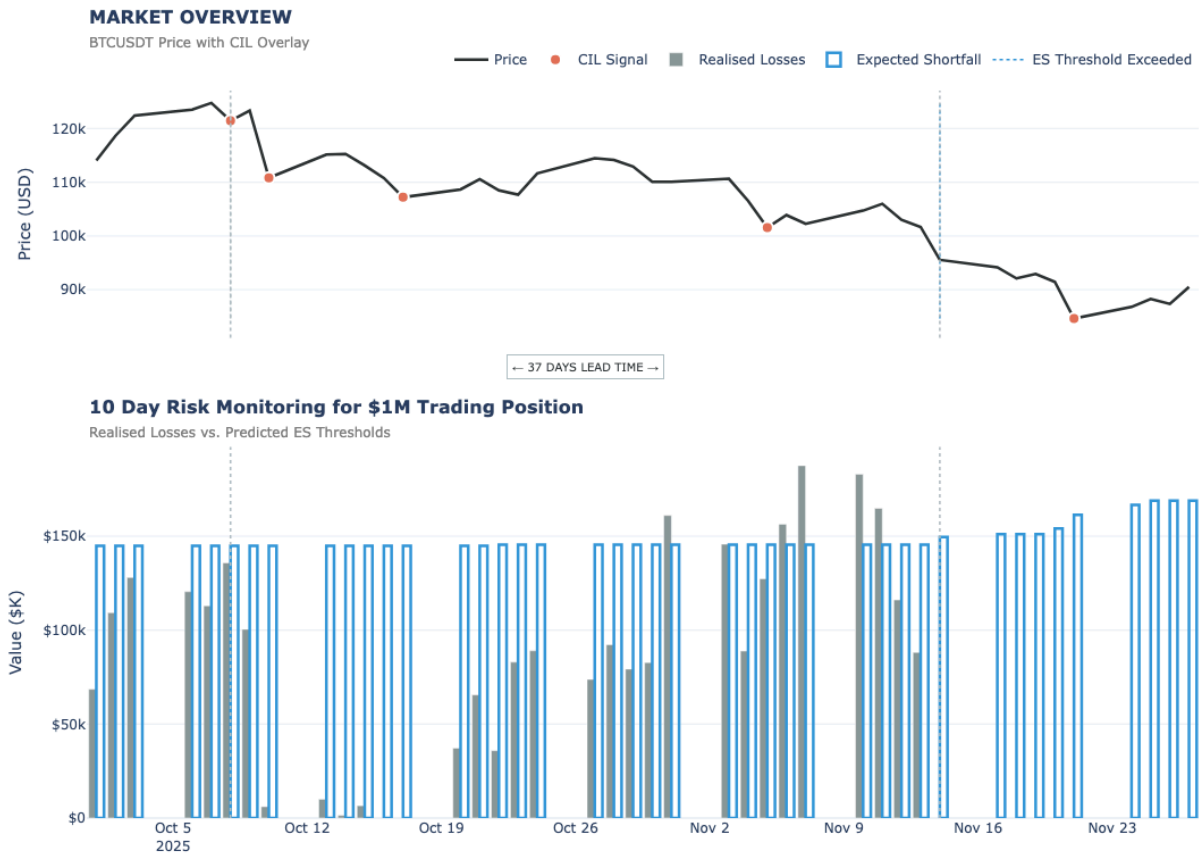

Event	2026 Iran War
Ticker	United Airlines
Drawdown Duration	8 Days
Drawdown (%)	26.88%
% of Change Explained by the CIL	26.14%
Drawdown Mitigated by the CIL	97%
Detection Lead Time over Expected Shortfall	14 Days

The following data establishes the deterministic causal trajectory of UAL's time series. It demonstrates the structural propagation of risk from the 1 second observation through the temporal hierarchy, transforming market volatility into a traceable and deterministic calculation.

Propagation Level	Datetime	Price	% Change from 1s	Propagation Path
0	2026-02-26 20:48:11	116.415	0	Observed from 1s
1	2026-02-26 20:50:40	116.17	-0.21	1s → 5s
2	2026-02-27 11:42:20	113.52	-2.48	5s → 10s
3	2026-02-27 14:30:30	110.85	-4.78	10s → 30s
4	2026-02-27 14:32:00	110.475	-5.102	30s → 1m
5	2026-02-27 17:05:00	106.36	-8.63	1m → 5m
6	2026-03-02 09:10:00	100.71	-13.49	5m → 10m
7	2026-03-02 09:30:00	101.07	-13.18	10m → 30m
8	2026-03-05 16:00:00	94.9	-18.48	30m → 60m
9	2026-03-06 14:40:04	85.98	-26.14	60m → Trough

The following chart quantifies the Detection Latency Gap by contrasting the CIL's performance with the Expected Shortfall (ES) model for UAL's time series. It demonstrates the CIL's capacity to confirm structural propagation 14 days before traditional regulatory triggers reacted, providing an immediate, auditable signal for risk mitigation while ES remained directionally blind to the inception of the 2026 Iran War drawdown.

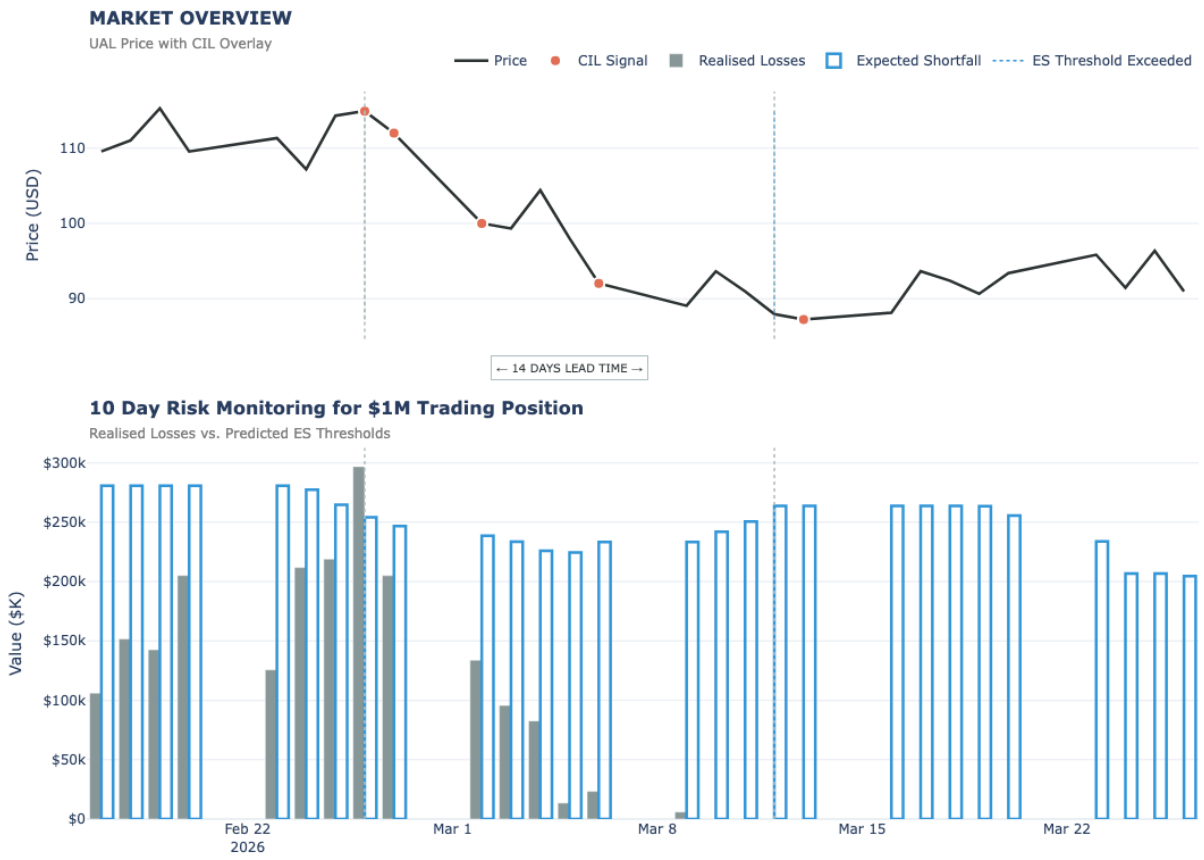

Technical Overview

Core Assumption: High-Frequency Inception

The CIL is built on the assumption that all multiscale changes originate at the highest observable time frequency before evolving into macro-change. By monitoring the micro-scale in real-time, the CIL identifies a causal directional shift at its point of inception.

- **Lead-Time Advantage:** Identifying micro-signals provides the foresight needed to stay ahead of macro-scale effects.
- **Early Detection:** The CIL captures the causal signal before the transition is recognised by statistical approaches or manifests at macro-scale resolutions (days or weeks), enabling intervention before the system reaches a critical transition.

Autonomous Causal Discovery (Category Theory)

The CIL leverages Category Theory to autonomously map the underlying causal structure of a non-stationary system using unlabelled, independent data streams.

- **Structural Linking:** Rather than seeking statistical correlations, the CIL identifies the directional morphisms that mathematically link an initial micro-signal to a macro-effect.
- **Global Function:** This process generates a unified global function that reconstructs the system's state space. This ensures the model remains consistent with the rules of the system, providing a stable framework for decision-making regardless of market regime or price volatility.

Topologically Lossless Causal Chains

The CIL maintains 100% causal accuracy across scale transitions. While standard models typically "average out" micro-data to find macro-trends, we reject this lossy approach. The CIL employs Topologically Lossless Causal Chains to preserve the specific micro-interactions foundational to macro-change.

- **Causal Signal:** Micro-activity that successfully maps to a morphism across scales, maintaining its structural identity as it matures.
- **Non-Causal Data:** High-entropy micro-activity that fails to form a structural link is systematically filtered without losing causal information.

This allows the CIL to extract deterministic causal signals from high-entropy environments, providing an auditable bridge from micro-inception to macro-reality.

Autonomous World Modelling (Adaptive Takens' Theorem)

To meet the demands of real-time detection, the CIL advances Takens' Theorem to function in non-stationary, high-entropy environments.

- **Dynamic State Projection:** The CIL reconstructs a multi-dimensional state space from a single time series using a learned global function; unlike models that fail due to numerical drift or fixed attractors, this approach maps the morphisms (rules of transformation) to achieve Structural Determinism by identifying the inevitable pathway of a transition even when the data stream appears chaotic.
- **Continuous Regime Adaptation:** The CIL uses continuous data streams and bidirectional feedback loops to constantly update its understanding of an environment's underlying structure. This means structural events typically labelled as "Black Swans" are not a surprise, but an observable, trackable transformation of the system's structural evolution.

Identifying Causal Initiation and Filtration

Using the autonomous world model, the CIL determines whether a micro-fluctuation signifies the initiation of a new causal path or the persistence of an established trajectory.

- **Lossless Filtration:** By understanding the system's underlying structure, the CIL identifies which high-frequency data is non-causal and which represents an important causal transition.
- **Solving the Computational Bottleneck:** High-frequency monitoring is traditionally resource-heavy. The CIL maintains performance speed by filtering high-frequency data once its causal transition is complete. This allows for a perfect mathematical bridge across scales without the burden of storing or processing non-causal data.