EORMEN

Case Study: BTC/USDT Crash
19th May 2021

Single Day Analysis
Phase 1 & 2

Technical Report

Leroy A. Palmer

palmer@eormen.com

eormen.com

8th December 2025

EXECUTIVE SUMMARY

This report documents the first two phases of applying the Eormen framework to a real chaotic financial event, the Bitcoin crash of 19th May 2021, using one second BTC/USDT data from Binance.

Eormen is not a trading system and not a predictor in the usual sense. It is a non interfering observation framework that watches a chaotic system through a fixed geometric representation and answers present tense questions about its internal structure. The only question in this case study is: *Are there structurally distinguished instants in the internal geometry of a crash day that can be identified from present data alone, without any reference to the future, and do those instants sit on the construction path of extreme later motion?*

We froze the PDCS Platform v1.0, defined a single mapping from raw market data into three band energies, then asked four preregistered questions (Q1–Q4) that differ only in how they weight those bands. All four questions use the same BTC/USDT 1 second day, the same channel mapping and the same normalisation, differ only in their band weight vectors, and select a single best second in the day without seeing any future prices or volatility.

On 19th May 2021, all four questions converged on exactly the same instant: 12:58:47 UTC, index 46 727 in the one second timeline, 93 seconds after the largest one second price move of the day, and 689 seconds (11.5 minutes) before the absolute intraday low.

At that instant, the internal Eormen geometry is dominated by a single stress band, B3. Under all four weightings, more than 99.99 per cent of the combined geometric energy at the hinge lies in B3, with negligible contributions from log price (B1) and volume (B2). The subsequent 5 and 15 minute windows lie in the top 1 to 2 per cent of all windows in the day for realised volatility, range and drawdown.

Phase 2 then examined the full score series for Q1–Q4 over the day. It found that the four score fields are almost perfectly correlated (correlations above 0.9999), the region where any question's score is at least 95 per cent or 99 per cent of its maximum consists of exactly one second (the hinge at 12:58:47), and the top ten peak times are identical across all questions.

In other words, for this day, Eormen sees a knife edge structural singularity in its internal geometry. Four preregistered questions, with different band weightings, all pick the same instant as the unique global reorganisation point.

This is not a trading result and not a general law of markets. It is an existence result: given the frozen Eormen configuration and a known crash day, there exists at least one present tense structural hinge that is robust across a family of geometric questions and that lies just before the deepest part of the crash.

1 METHODOLOGY

1.1 Data

The study uses one second BTC/USDT kline data from Binance for a single day: instrument BTC/USDT, venue Binance, date 19th May 2021 UTC, resolution 1 second OHLCV with taker buy volume.

Data checks confirmed that the file contains 86 400 records, one per second from 00:00:00 to 23:59:59 UTC, with no missing timestamps and no gaps in the main crash window. For every row, high ≥ low, and prices and volumes are non negative. The intraday price range matches the known magnitude of the 19th May crash.

The data are treated as ground truth for this case study. No resampling, smoothing or cleaning beyond integrity checks is applied.

1.2 Platform configuration

All experiments use the same frozen Eormen PDCS Platform: PDCS Platform v1.0. Before each run, the platform loads all mesh and operator artefacts (P1 to P20), and verifies P18–P20 invariants, including Weyl uniformity conditions and internal consistency checks.

No artefacts in P1 to P20 are modified as part of this work. Those layers form the sealed computational substrate for Eormen. All experimental choices live in P21–P25 (injection and metric layers) and P28 (questions), and in external analysis scripts that never interact with the platform.

This separation enforces non interference: P1–P30 only ever see past and present data, no function inside P21–P30 ever sees future returns, volatility or labels, and all evaluation against future behaviour is done in separate scripts that treat Eormen's output as completed observations.

1.3 Channel mapping (P21)

At each second i, the raw kline provides close price C_i, high and low prices H_i, L_i, base volume in BTC V_i^{base}, and taker buy base volume T_i^{buy}. These are mapped into three channels:

$$
\begin{aligned}
X_i &= \log C_i && \text{(log price level)} \\
Y_i &= V_i^{\text{base}} && \text{(per second trading volume)} \\
Z_i &= (2T_i^{\text{buy}} - V_i^{\text{base}}) \cdot (H_i - L_i) && \text{(directional imbalance} \times \text{intrasecond range)}
\end{aligned}
$$

Band energies are defined by the identity mapping $B1_i = X_i^2$, $B2_i = Y_i^2$, $B3_i = Z_i^2$. There are no moving averages or filters. Every quantity is an algebraic transformation of the raw Binance feed.

1.4 Metric layer (P24–P25)

From the band energies, P24 computes amplitudes $A_{k,i} = \sqrt{B_{k,i}}$ for $k \in \{1, 2, 3\}$ on a single τ grid point: $\tau = 0.10$, represented as index 0.

P25 then defines a per band median scale $M_k = \text{median}_i\, A_{k,i}$, floored at 10^{-12} to avoid division by zero. Each amplitude is normalised to a unit free variable $\xi_{k,i} = A_{k,i}/M_k$.

So in each band k, a typical second has $\xi_{k,i} \approx 1$, and large ξ represents a large deviation from that band's usual behaviour for this day.

These three normalised amplitudes, $\xi_{1,i}, \xi_{2,i}, \xi_{3,i}$, at $\tau = 0.10$ are the sole inputs to Q1–Q4.

1.5 Question layer (P28) – Q1 to Q4

Each question Qx is a scalar functional of the three band amplitudes at time i. For a given weight vector $\mathbf{w} = (w_1, w_2, w_3)$ we define weighted components $v_k(i) = w_k \cdot \xi_{k,i}$ for $k \in \{1, 2, 3\}$, and combined score $\psi_{\mathrm{Qx}}(i) = v_1(i)^2 + v_2(i)^2 + v_3(i)^2$.

The hinge index for Qx, i^*_{Qx}, is the time index i that maximises $\psi_{\mathrm{Qx}}(i)$ over the full day (earliest index in case of ties).

The four questions differ only in their weight vectors: Q1 (neutral) $\mathbf{w} = (1.0, 1.0, 1.0)$, Q2 (price emphasised) $\mathbf{w} = (2.0, 1.0, 0.5)$, Q3 (volume emphasised) $\mathbf{w} = (0.5, 2.0, 1.0)$, Q4 (stress emphasised) $\mathbf{w} = (0.5, 1.0, 2.0)$. All four were specified in advance as a small symmetric family around the neutral point. No question weights were tuned after inspecting results.

1.6 External analysis

For each question, an analysis script operates entirely outside PDCS. It reloads the raw BTC/USDT one second series, reconstructs one second log returns, and for fixed horizons $H \in \{60, 300, 900\}$ seconds computes forward window metrics: realised volatility (variance of returns), absolute range (max minus min price), relative range (range divided by starting price), absolute drawdown (maximum drop from starting price to any later minimum), and relative drawdown (drawdown divided by starting price).

It evaluates these metrics at the hinge index and reports their percentile ranks within the full distribution for that horizon. It computes the same metrics for three baselines: the second with the largest absolute one second log return, the second with the largest absolute one second price change, and the second with the largest one second volume. It also computes a random benchmark by drawing 1000 random start indices.

Phase 2 adds a constellation analysis script that reloads the four score timelines, computes the correlation matrix between questions, identifies top peak times for each question, measures high score regions where the score is at least 95 per cent or 99 per cent of its maximum, and measures overlaps between questions at those thresholds.

This analysis never feeds back into PDCS; it is strictly observational.

2 PHASE 1 RESULTS – Q1 TO Q4

2.1 Hinge convergence on 12:58:47

All four questions select the same hinge index and time: time index $i^* = 46\,727$ and time $t^* =$ 19th May 2021 12:58:47Z. The selected price at that instant is approximately 34 350 USDT.

For reference, the largest one second price change occurs at 12:57:14Z, 93 seconds earlier, and the global intraday low of 30 000 USDT occurs at 13:10:16Z, 11.5 minutes later.

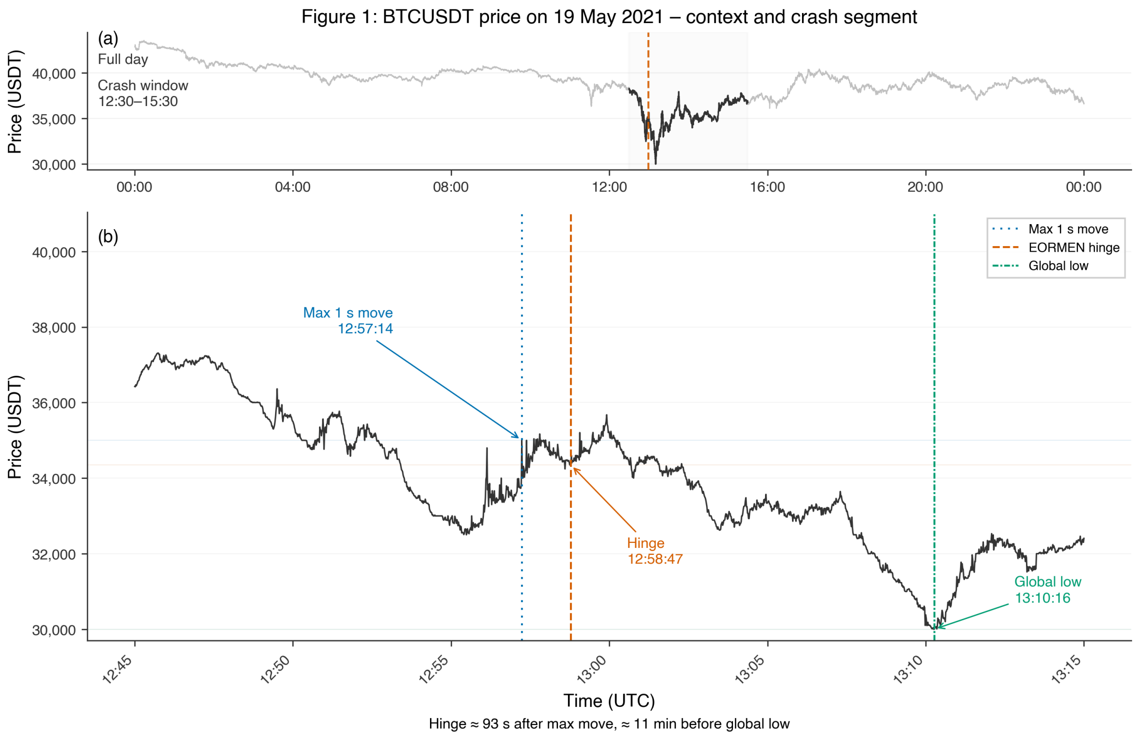

Figure 1: Price and hinge. Daily BTC/USDT close price at one second resolution, with vertical lines marking the hinge time 12:58:47 (Q1–Q4 selection), the time of the largest one second price change, and the time of the daily low. This figure shows visually that the Eormen hinge lies between the first sharp impact and the eventual bottom of the crash.

2.2 Band shares at the hinge – B3 dominance

At the hinge time, we can decompose each question's combined score into contributions from B1, B2 and B3. In all four cases, more than 99.99 per cent of the score comes from B3, the stress band. The contributions from B1 (log price) and B2 (volume) are many orders of magnitude smaller.

In other words, although the questions are defined on a three band space, the maximising configuration is effectively a pure B3 stress event under the neutral, price tilted, volume tilted and stress tilted weight vectors.

2.3 Forward realised metrics at the hinge

For each horizon H, we looked at the forward window starting at the hinge index. Across all four questions (which share the same hinge), at 60 seconds the window is in roughly the top 2 per cent of all 60 second windows for volatility and range relative to the starting price, although the net drawdown over exactly one minute is small. At 300 seconds (5 minutes), the window is in roughly the top 1 to 1.5 per cent for volatility, range and drawdown. At 900 seconds (15 minutes), the window is again in roughly the top 1 to 1.5 per cent for volatility, range and drawdown.

These percentiles mean that if you choose a random second in this day and look forward 5 or 15 minutes, it is very unlikely to see a window with realised volatility and price excursions as large as those that follow the

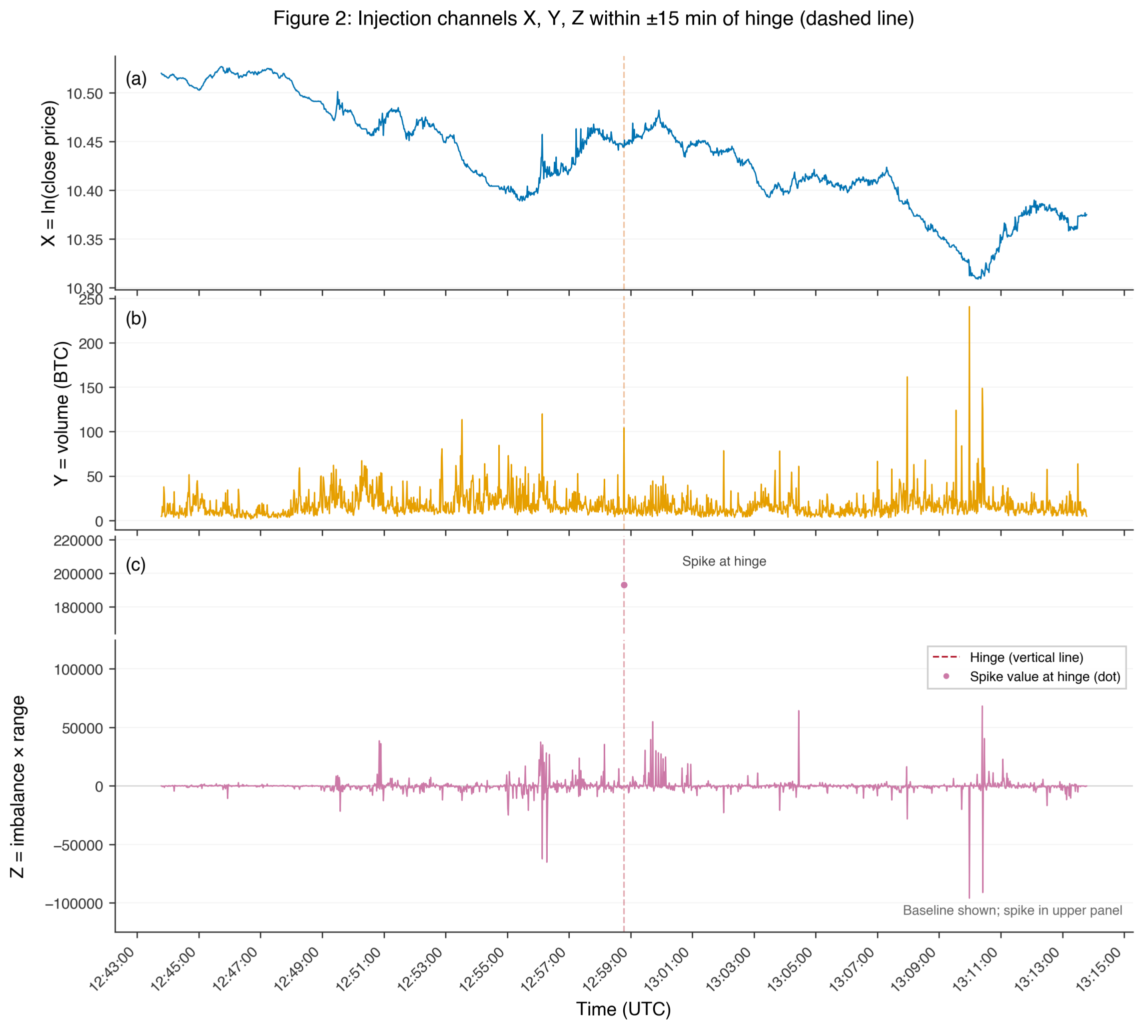

Figure 2: Injection channels near hinge. Time series of the three raw injection channels X_i (log close price), Y_i (per second volume), and Z_i (directional imbalance × intrasecond range) for a window around the hinge (approximately ±15 minutes). The figure highlights the relative scale of the three channels, the dramatic spike in Z at 12:58:47, and the comparatively modest activity in X and Y. This visualises the structural fact that the hinge is driven by a stress singularity in the Z channel, which feeds into band B3 after squaring.

hinge.

Table 1: Forward metric percentiles at hinge. Percentile ranks of forward volatility, range and drawdown at the Eormen hinge (12:58:47 UTC) across three time horizons. All percentiles are computed relative to the distribution of the same metric across all eligible starting indices in the 19th May 2021 day.

Horizon	**Volatility**	**Range (relative)**	**Drawdown (relative)**
60 seconds	0.9981	0.9821	0.0748
300 seconds	0.9875	0.9950	0.9854
900 seconds	0.9853	0.9936	0.9945

2.4 Baseline comparisons

To ensure that Eormen is not simply rediscovering obvious metrics, we compared its hinge to simple baselines.

The largest absolute one second price move (and log return) occurs at 12:57:14Z. That second also leads into very high volatility and drawdown windows, lies in the same crash region as the hinge, and is a natural spike candidate if one were using conventional indicators. Eormen does not pick this second. It picks a structurally different stress configuration 93 seconds later, which lies closer in time to the ultimate low and represents a distinct moment in the internal geometry.

The largest one second volume occurs at 12:17:44Z, well before the main crash phase. That second has only moderate forward volatility and range, and does not behave like a structural hinge in the usual sense. This baseline shows that simply looking where volume spikes does not identify the canonical hinge.

Under random time selection, the distribution of percentile ranks for volatility, range and drawdown is centred near 0.5 with a standard deviation around 0.28. The hinge, by contrast, sits well above the 95th percentile of this distribution for volatility and range at 300 and 900 seconds, and similarly high for drawdown. This demonstrates that the hinge's forward behaviour is extremely atypical compared to randomly chosen seconds.

Table 2: Baseline comparison. Comparison of forward metric percentiles across time selection methods. The Eormen hinge (12:58:47 UTC) and the maximum one second price move (12:57:14 UTC) both identify moments in the crash cluster with extreme forward behaviour. The maximum volume second (12:17:44 UTC) does not exhibit extreme forward metrics. Random baseline statistics are computed from 1000 random samples.

Method	**60 s**			**300 s**			**900 s**		
	Vol	Range	DD	Vol	Range	DD	Vol	Range	DD
Eormen hinge	0.9981	0.9821	0.0748	0.9875	0.9950	0.9854	0.9853	0.9936	0.9945
Max 1 s move	0.9987	0.9922	0.9950	0.9965	0.9617	0.9688	0.9896	0.9921	0.9995
Max 1 s volume	0.8367	0.8124	0.0748	0.6972	0.5229	0.0329	0.6909	0.7384	0.7079
Random (mean)	0.4888	0.4771	0.4934	0.4966	0.4920	0.4992	0.4841	0.4919	0.4977
Random (95th pct)	0.9421	0.9450	0.9484	0.9337	0.9436	0.9402	0.9459	0.9463	0.9448

3 PHASE 2 CONSTELLATION – WITHIN DAY STRUCTURE

Phase 1 established that Q1–Q4 all select the same hinge and that this hinge is followed by extreme volatility. Phase 2 asks a different question: what do the full score landscapes $\psi_{\mathrm{Qx}}(t)$ look like over the day, and how do they relate to each other?

3.1 Score time series and correlation

For each question, the score $\psi_{\mathrm{Qx}}(t)$ is defined for every second in the day, giving four time series. When we compute the 4×4 Pearson correlation matrix between these series we find that all off diagonal entries are greater than 0.9999. This means that as functions of time, the four score series are almost perfectly collinear. The effect of changing the weight vector within this family on this day is effectively a rescaling of a single underlying stress time series.

Figure 3: Normalised score series for Q1 to Q4 over full day

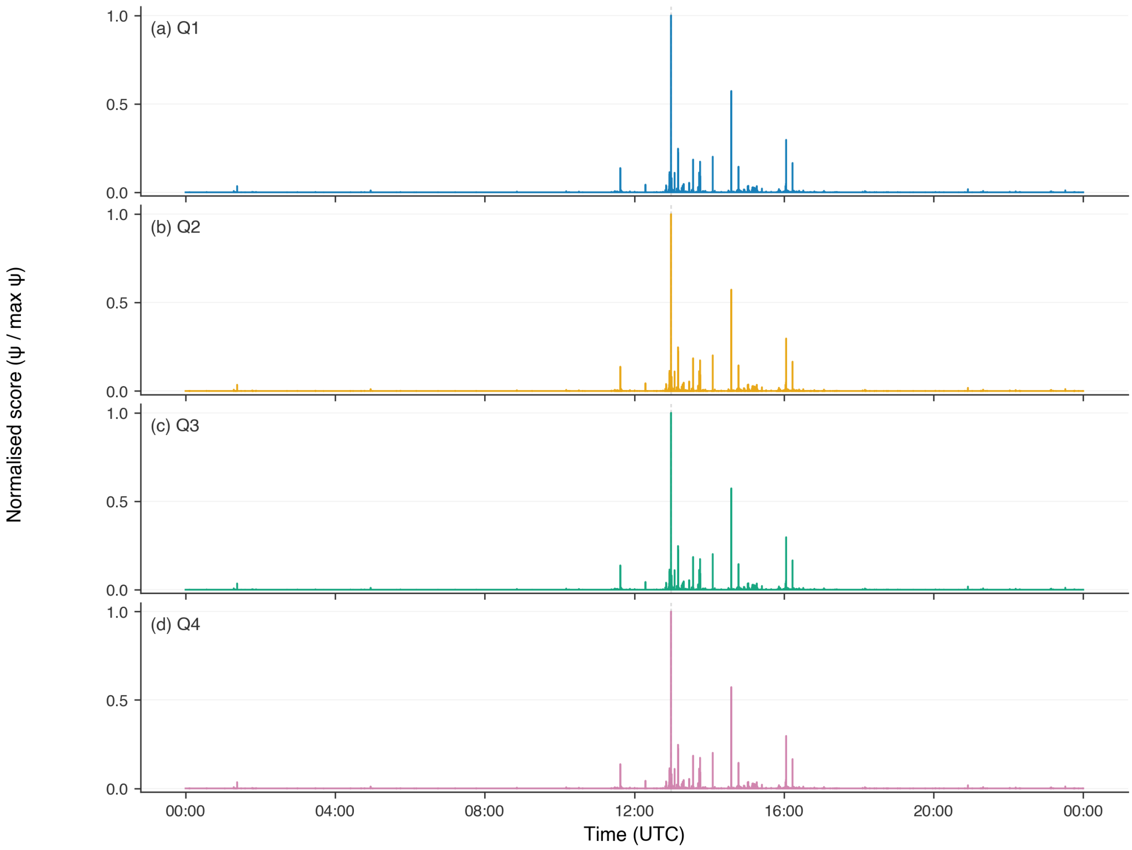

Figure 3: Score series for Q1 to Q4. Four vertically stacked panels showing the normalised score series $\psi_Q(i)/\max\psi_Q$ for each question over the full 19th May 2021 day. Each panel uses the same time axis and normalised score scale (0 to 1). The visual identity of patterns across all four panels demonstrates their correlation above 0.9999, with all questions exhibiting the same dominant spike at 12:58:47 and identical secondary peak structure throughout the day.

3.2 High score regions and hinge width

For each question and each threshold $\alpha \in \{0.95, 0.99\}$, we examined the set of times where the score is at least α times its maximum. On this day, for all four questions and both thresholds, this set consists of a single time index: the hinge at 12:58:47. The high score region has width exactly one second.

There is no plateau and no extended shoulder. The peak is a literal one second wide needle in all four questions. This finding is visible in Figure 3, where the dominant spike at 12:58:47 stands alone, with all secondary peaks remaining well below the maximum.

So within this representation, the crash day does not present a vague danger zone. It presents a single instant where the internal geometry is maximally deformed.

3.3 Top ten peaks and temporal structure

For each question we sorted all times by score and recorded the top ten peaks. The resulting lists have exactly the same ten time indices for Q1, Q2, Q3 and Q4; only the absolute scores differ. The same offsets from the hinge appear in all questions.

The peaks include times just before and after the hinge, and several other stress events later in the day, but the pattern of timings is identical across questions. So Q1–Q4 do not produce different lists of interesting moments. They all see the same discrete stress skeleton, with the hinge as the highest point.

This shared temporal structure is evident in Figure 3, where secondary peaks occur at identical times across all four panels.

3.4 Overlap analysis

We also asked how often multiple questions were simultaneously above their high score threshold. For each α we measured the fraction of seconds where at least one, at least two, at least three, or all four questions exceed α times their respective maxima.

On this day, for both $\alpha = 0.95$ and $\alpha = 0.99$, the fractions for at least one, at least two, at least three and all four are all equal to 1 / 86 400, the fraction of the day represented by a single second. That single second is the hinge.

So the constellation does not produce broad regions where questions take turns being high. All four questions are simultaneously in their high state only at the hinge, and never elsewhere.

Table 3: Overlap fractions. Fraction of seconds in the day where questions simultaneously exceed score thresholds. For both 95% and 99% thresholds, all overlap conditions yield the same fraction (1/86 400), corresponding to exactly one second. This demonstrates that all four questions are simultaneously in their high-score state only at the hinge (12:58:47 UTC), and never elsewhere.

Threshold	≥1 question	≥2 questions	≥3 questions	All 4 questions
$\alpha = 0.95$	1/86 400	1/86 400	1/86 400	1/86 400
$\alpha = 0.99$	1/86 400	1/86 400	1/86 400	1/86 400

4 INTERPRETATION

The combined picture from Phase 1 and Phase 2 is unusually clean.

4.1 A single dominant stress mode

Under the fixed mapping used here, the BTC/USDT crash day is observed through three bands: B1 (log price), B2 (volume), and B3 (directional imbalance times intra-second range).

After normalisation by daily medians, the internal geometry of the day appears effectively one dimensional in the B3 direction. The four questions, which explore different weightings of B1, B2 and B3, all collapse onto the same structure: their scalar score fields are almost perfectly correlated, they all see the same top ten peaks, and they all identify the same unique global maximum at 12:58:47.

At that instant, the energy in the combined norm is essentially all B3. The stress channel overwhelms price and volume at the geometric level.

4.2 Knife edge structural singularity

The high score region analysis shows that the hinge is not part of a wide plateau. It is a knife edge: the region where any question is at least 99 per cent of its own maximum is exactly one second long, and the same is true for the 95 per cent threshold.

So within this representation, the crash day does not present a vague danger zone. It presents a single instant where the internal geometry is maximally deformed.

This is an internal property of the mapping and the platform. It is not defined by future prices or realised volatility. Nevertheless, when we look at what happens afterwards, the hinge lies on the path into some of the largest five and fifteen minute excursions of the day.

4.3 Non interfering observation in action

Eormen's design principle is to observe a chaotic system without interfering with it. In this case study that translates to freezing the platform and cascade (not tuning them), defining channels and questions purely from past and present data, and keeping all future dependent analysis outside the system.

The case study shows that even under such constraints, it is possible to identify present tense configurations that have a clear structural relationship to later extreme events.

It is important to emphasise that Eormen is not predicting a crash in the trading sense. It does not forecast the sign or magnitude of returns, it does not classify days as crash days or quiet days, and it does not produce a strategy or risk rule.

What it does is reveal that on this particular crash day, there exists a unique, internal stress configuration that is several steps removed from obvious surface metrics and that nevertheless aligns tightly with the construction of the observed crash.

4.4 Uniqueness relative to conventional viewpoints

Conventional indicators such as volume spikes, largest one second returns or simple volatility filters do capture parts of the crash, but they do not single out 12:58:47.

The largest one second move is earlier and is more directly visible on the raw price chart. The largest volume second is earlier still and not structurally special. Simple rolling volatility measures tend to smear together many seconds.

Eormen's hinge is not a trivial rephrasing of any of these. It is genuinely defined by the behaviour of a stress band constructed from imbalance and intrasecond range, normalised by the day's typical behaviour, inside a cascade framework that respects bounded chaos. The fact that this internally defined hinge lies so close to the eventual low is exactly the kind of alignment the framework is designed to detect.

5 LIMITATIONS AND NEXT STEPS

This case study is deliberately narrow. It covers a single instrument (BTC/USDT), a single day (19th May 2021), uses a single channel mapping (X, Y, Z as defined above), uses a single τ scale ($\tau = 0.10$), and explores a small family of questions (Q1–Q4).

As such, the results should be read as an existence proof, not as a general statement about financial markets or about BTC.

Limitations to note explicitly: first, there is no cross day validation yet, and no evidence at this stage that similar hinges exist, or are detectable, on other crash days or on quiet days. Second, there is no cross asset validation yet, and the behaviour of the same configuration on other instruments, such as ETH, equity indices or foreign exchange, is unknown. Third, the single mapping and τ means that other choices of channels or τ scales may reveal different structures or may fail to reveal any clear hinges. Fourth, the question family is restricted to a modest perturbation around neutrality, and more exotic questions, including ones that explicitly suppress B3, have not yet been explored. Fifth, the external metrics are basic, using simple volatility, range and drawdown, and do not yet include more sophisticated microstructure diagnostics such as jump tests or Hawkes intensities.

The next steps are therefore clear.

Phase 3: Multi day, multi asset validation. Freeze Q1–Q4 exactly as defined here and apply them to a pre specified set of crash days, rally days and quiet days across several assets. For each, run the same Phase 1 and Phase 2 analysis and compile statistics on how often a clear canonical hinge exists, how often that hinge lies on the path into extreme events, and how the effect size decays or persists across regimes.

Phase 4: Question space and sensitivity. Explore sensitivity to changes in weights, channel definitions and τ scales as a controlled sensitivity analysis, not an optimisation exercise.

Phase 5: Comparative studies. Compare Eormen's hinge timing to the outputs of standard quantitative tools from the literature, such as jump detection and regime switching models.

Only after such work would it be appropriate to speak about general patterns or about practical uses of the framework.

6 REFERENCES

Framework Documentation

Palmer, L.A. (2025). *EORMEN: A Framework for Non-Interfering Observation of Complex Systems*. White Paper. Available at eormen.com

Data Source

Binance (2021). BTC/USDT 1-second kline data, 19 May 2021. Retrieved from Binance Public Data: `https://data.binance.vision/`

End of Report
