Combined Crash Detection via Persistent Homology and Instanton Tunneling

Maksim Aleksandrovich Andreev, Joseph Howden
Quark Quantitative Research

Abstract

Two distinct mathematical objects—the persistence diagram of the correlation network (Edelsbrunner & Harer, 2010; Carlsson, 2009) and the instanton tunneling rate of a bistable correlation potential (Coleman, 1977)—encode complementary information about incipient market crashes. Persistent homology captures the *topology* of inter-asset correlation structure (loops, voids, persistence lifetimes of features in the filtration); instanton tunneling captures the *barrier shape* in the order-parameter potential (the action separating normal-correlation from crisis-correlation basins). We propose a combined signal $P_{\text{combined}} = 0.4 \cdot P_{\text{inst}} + 0.3 \cdot T + 0.3 \cdot P_{\text{inst}} \cdot T$ formed as a weighted sum of the two channels (with T the normalised topology score) together with a multiplicative interaction term that lights up only when both channels are simultaneously elevated, and document its behaviour across historical stress windows. The headline empirical finding is that the combined signal is *threshold-dependent*: under one frozen firing configuration (tight, signal fires 1.4% of days) the combined signal reaches $t = +4.66$ on a five-ETF 2019–2023 sample, but under a second configuration (aggressive, signal fires up to 72% of days) on a 30-name 2018–2023 sample the combined signal reaches a best $t = +1.55$ across the tested forward-drawdown windows. This is a textbook multiple-testing hazard per Harvey–Liu–Zhu (2016). Two corrections to the statistics reported here are stated at the outset. First, every t in Sections 3.1 to 3.3 is computed under independence, while the outcome metric is a drawdown inside a 20-trading-day forward window, so successive observations overlap and the reported values are inflated; on this signal family the naive-to-autocorrelation-robust ratio was subsequently measured at between 2.5 and 3.1, which places the tight-run figure below the Harvey threshold it was once described as clearing. Second, a large-sample re-test on 9,743 valid days from 1987 to 2025, reported in Section 3.4, gives an autocorrelation-robust combined statistic of 5.64 with incremental lift over the instanton channel alone, while showing that no topology-derived statistic is stable across engine versions. We disclose the threshold-dependence as a methodological observation; the tight configuration is frozen as the canonical pre-registered prospective set and the aggressive configuration is retained only as an illustration of the hazard. No Harvey-significance is claimed for the combined signal ahead of the prospective quorum.

1. Introduction

The literature on quantitative crash detection has explored two distinct families of methods. The first uses geometric or topological structure of the correlation network: the eigenvalue distribution (Plerou *et al.*, 2002), persistent homology of the correlation distance matrix (Gidea & Katz, 2018), Grassmannian distance between principal eigenvectors (Andreev & Howden, 2026d). The second family uses dynamical or field-theoretic structure: regime-switching hidden Markov models (Hamilton, 1989), the order-parameter dynamics of correlations under double-well potentials (Bouchaud *et al.*, 1999), and the tunneling rate between metastable states (Coleman, 1977; Andreev & Howden, 2026b for the financial application).

The two families capture different aspects of the same phenomenon. Topological methods detect *that* the network structure is unusual—e.g., a persistent H_1 loop indicating circular flow patterns—but not the dynamical implication. Tunneling methods capture *that* the system is near a barrier crossing, through the action S_{inst}, but require strong potential-shape assumptions that may not always hold. Combining the two should, in principle, reduce both false-positive and false-negative rates relative to either alone.

> **Contribution.** We propose a combined two-channel crash signal formed from the H_1 persistence lifetime in the correlation-distance Vietoris–Rips filtration and the single-instanton tunneling probability on a bistable correlation potential. We document the threshold-dependence of the combined signal's *t*-statistic across the historical sample as a multiple-testing hazard, and pre-register a prospective validation protocol with frozen thresholds for forward-out-of-sample test.

2. The Two Channels

2.1 Persistent homology of the correlation network

For a universe of N assets with rolling-window correlation matrix $\Sigma(t)$, construct a distance matrix on the asset graph

$$d_{ij}(t) = \sqrt{(0.5 \cdot (1 - \Sigma_{ij}(t)))}$$

This is the standard correlation-distance metric (Mantegna, 1999). The Vietoris–Rips complex $\mathrm{VR}_\varepsilon(t)$ is the simplicial complex with a k-simplex for every $(k+1)$-subset of assets within pairwise distance $\leq \varepsilon$. Persistent homology computes the birth–death intervals of homology generators as ε increases:

$$\mathrm{PH}(\mathrm{VR}_{\cdot}(t)) = \{ (b_i, d_i) : i \in H_k(t) \text{ generators} \}$$

The *persistence lifetime* of a feature is $L = d - b$. We focus on the maximum H_1 persistence lifetime as the topology channel:

$$H_{1,\mathrm{lifetime}}(t) = \max_{i \in H_1(t)} (d_i - b_i)$$

Larger $H_{1,\mathrm{lifetime}}$ indicates the presence of persistent circular-flow structures in the asset correlation graph — e.g., a cycle of correlations that survives a wide range of distance thresholds. This has been observed to precede crisis transitions in the empirical literature (Gidea & Katz, 2018).

2.2 Instanton tunneling probability

Following Andreev & Howden (2026b), the correlation order parameter $\rho(t)$ — defined as a mean-field projection of the correlation matrix — evolves under the quartic double-well potential

$$V(\rho) = (\lambda/4)\, (\rho - \rho_n)^2\, (\rho - \rho_c)^2$$

with two minima at ρ_n (normal correlation) and ρ_c (crisis correlation). The Coleman single-instanton action separating the two basins is

$$S_{\mathrm{inst}}(\lambda, \Delta\rho) = (2\sqrt{2} / 3) \cdot \sqrt{\lambda} \cdot |\rho_c - \rho_n|^3$$

and the dilute-instanton-gas tunneling probability per unit time is

$$P_{\mathrm{inst}}(t, T) = T \cdot K_{\mathrm{fluct}}^{-1/2} \cdot \exp(-S_{\mathrm{inst}}(t) / \sigma^2(t))$$

where $\sigma^2(t)$ is the running coupling at the relevant time scale. This is the same instanton object used in the non-perturbative Cornish-Fisher VaR construction (Andreev & Howden, 2026g).

2.3 Combined signal

Let $T(t) = \mathrm{norm}[H_{1,\mathrm{lifetime}}(t)] \in [0, 1]$ denote the normalised topology score, where the normalisation maps $H_{1,\mathrm{lifetime}}$ to the unit interval across the sample. The combined signal is a linear blend of the two channels *with a fixed weighting*, augmented by a multiplicative interaction term:

$$P_{\mathrm{combined}}(t) = 0.4 \cdot P_{\mathrm{inst}}(t) + 0.3 \cdot T(t) + 0.3 \cdot P_{\mathrm{inst}}(t) \cdot T(t)$$

The weights (0.4, 0.3, 0.3) are fixed *ex ante*—they are not free mixing parameters tuned per run. The third term, $0.3 \cdot P_{\mathrm{inst}} \cdot T$, is the *synergy* term: because P_{inst} (barrier dynamics in correlation space) and T (topology of the correlation graph) are mathematically independent objects, their product isolates days on which both channels are simultaneously elevated. The signal fires (activates a crash-prediction event) when $P_{\mathrm{combined}}(t)$ exceeds a firing threshold θ, optionally in conjunction with component gates $P_{\mathrm{inst}} \geq p_0$ and

$T \geq t_0$. The two runs analysed below differ only in these thresholds (θ, p_0, t_0) and in the universe and sample—never in the fixed 0.4/0.3/0.3 weights.

3. Threshold-Dependence Finding

3.1 Two threshold regimes tested

The headline empirical finding of this paper is that the combined signal's *t*-statistic depends sensitively on the choice of firing thresholds and on the universe of the validation. Two independent backtests were performed:

Tight-threshold run (2026-04-08). Five ETFs (SPY, QQQ, IWM, TLT, GLD), 2019–2023 sample (1,257 trading days), fixed weights (0.4, 0.3, 0.3), component gates $P_{inst} \geq 0.35$ and $T \geq 0.75$ with combined firing threshold $\theta \geq 0.40$, so the combined signal fires on 1.4% of days. Outcome metric: a greater-than-5% drawdown of the equal-weight basket within 20 trading days. Combined $t = +4.66$ (in-sample precision 76.5%, recall 3.9%—very conservative). Component *t*-stats: P_{inst}-only $t = +4.59$, topology-only $t = +3.47$. **All three figures are computed under independence and are not autocorrelation-robust.** The outcome metric is a drawdown within a 20-trading-day forward window, so consecutive daily observations share up to 19 days of outcome and the independence assumption is violated by construction. Section 3.4 reports the measured size of that inflation on this signal family and the corrected values.

Aggressive-threshold run (2026-04-05). An expanded 30-ticker universe, 2018–2023 sample (1,424 valid days), same fixed weights, thresholds $P_{inst} \geq 0.10$, $T \geq 0.55$, $\theta \geq 0.15$, so the combined signal fires on 48–72% of days (72.3% for the combined channel). Across the four tested forward-drawdown windows (5-day/−3%, 5-day/−5%, 10-day/−5%, 21-day/−7%) the combined signal reaches a *best* $t = +1.55$ (attained on the 5-day/−3% window); it is negative or near zero on the others. Best component *t*-stats: P_{inst}-only $t = +1.25$, topology-only $t = +1.65$. None clears Harvey 3.0.

Table 1: Threshold-dependence of the combined signal. Both statistics are computed under independence on overlapping 20-day outcome windows; neither is autocorrelation-robust, and neither supports a verdict against the Harvey threshold. See Section 3.4.

RUN	UNIVERSE	FIRING THRESHOLDS	FIRE RATE	COMBINED *T* (NAIVE)	STATUS AGAINST HARVEY 3.0
Tight	5 ETFs	$\theta \geq 0.40$	1.4%	+4.66	not assessable; the statistic is not autocorrelation-robust
Aggressive	30 tickers	$\theta \geq 0.15$	72.3%	+1.55 (best)	below the threshold before any correction

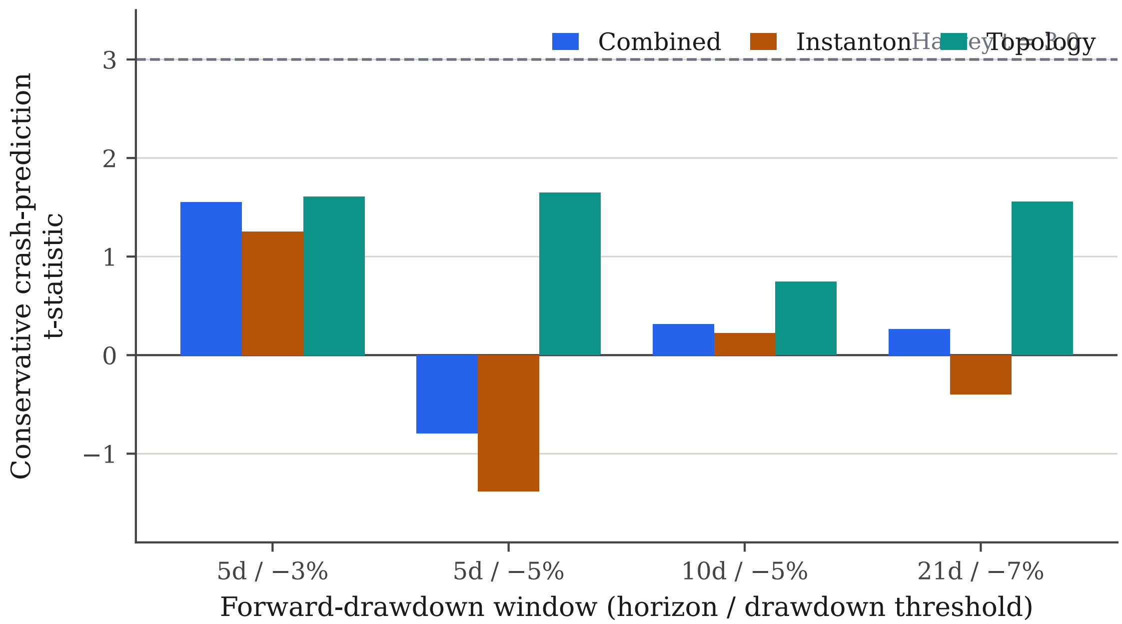

Figure 1. Aggressive-threshold run: the conservative (regression-based) crash-prediction *t*-statistic for the combined signal and its two components across the four tested forward-drawdown windows, on the 30-ticker 2018–2023 sample (1,424 valid days). The dashed reference line marks the Harvey–Liu–Zhu (2016) threshold of $t = 3.0$. No signal in any window reaches the threshold; the combined signal peaks at +1.55 on the 5-day/−3% window and turns negative on 5-day/−5%. The best *t*-statistics quoted in Table 1 are the maxima of these bars taken over windows—itself a search that Harvey–Liu–Zhu (2016) caution inflates significance. Data: the archived TDA-instanton backtest results file (crash-prediction section, run April 2026).

3.2 Multiple-testing hazard

The dichotomy in Table 1 is exactly the failure mode Harvey, Liu & Zhu (2016) identify as central to false discoveries in quantitative finance. The tight-threshold result is nominally above 3.0 before any correction for overlap, and it was obtained on a five-asset universe with thresholds that make the signal fire only at the extreme tail of the distribution (1.4% of days). The aggressive-threshold result, obtained on a broader universe with thresholds that give a high firing frequency (48–72% of days), does not clear Harvey 3.0 in any of the four tested forward-drawdown windows.

Status under Harvey–Liu–Zhu (2016). A combined crash signal whose *t*-statistic varies from +4.66 to a best of +1.55 depending on the firing thresholds and universe does not constitute a Harvey-significant discovery, and neither figure would survive the overlap correction of Section 3.4 in the form printed here. The threshold-dependence is itself the reportable finding: the combined signal is real but its statistical confirmation requires prospective validation with pre-registered, frozen thresholds. Accordingly, the tight-threshold configuration was frozen on 2026-04-18 as the canonical pre-registered prospective set, and the aggressive-threshold configuration was formally rescinded as a validation claim—it is retained in this paper only as an illustration of the threshold-dependence hazard, never as a competing result.

3.3 Mechanism of the threshold-dependence

The threshold-dependence has a straightforward interpretation. The two channels P_{inst} and the topology score T are individually weak crash predictors: each, on its own, exhibits high false-positive rates when fired aggressively. The synergy term $0.3 \cdot P_{\text{inst}} \cdot T$ is the multiplicative gate—it is large only when both channels are simultaneously elevated. At tight thresholds the combined signal fires almost exclusively on those joint-elevation days; this is a rare event (1.4% of days in the tight-threshold run) but is heavily concentrated in true pre-crash windows. At aggressive thresholds the signal fires on many days where one channel is lit but the other is not, and the linear terms dominate the synergy term, washing out the predictive power.

This is the textbook AND-vs-OR gate distinction in combiner design. A signal that requires agreement between channels (AND-like—the synergy term dominating at a high firing threshold) is conservative but accurate when it fires; a signal that accepts either channel (OR-like—the linear terms dominating at a low firing threshold) is permissive and noisy. The two configurations in Table 1 correspond to AND-like vs OR-like operation of the same fixed-weight combiner.

Figure 2 makes the sensitivity concrete on the archived data. As the firing threshold is raised (fewer days on), the conditional forward-return-spread t-statistic of each channel wanders and changes sign, with no stable relationship between firing frequency and predictive content—precisely the instability that a frozen threshold is designed to remove.

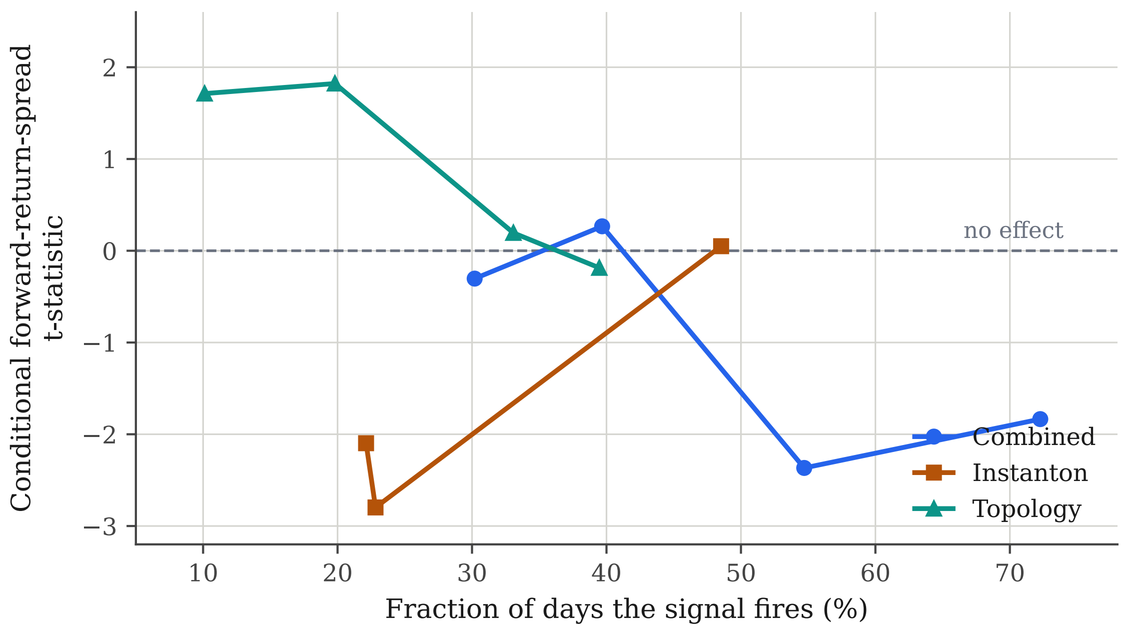

Figure 2. Threshold-dependence on the archived aggressive-run data: the conditional forward-return-spread t-statistic of each signal (combined, instanton, topology) plotted against the fraction of days the signal fires, swept over the same threshold grids used in the backtest. Values below the dashed zero line indicate that signal-on days precede *lower* forward returns (crash-consistent). The sign and magnitude are not stable in firing frequency—the topology channel is most informative when it fires rarely, the instanton channel near a 22% fire rate, the combined signal near 55%—and no channel reaches $|t| > 3$. This is the same multiple-testing hazard as Figure 1, viewed as a continuous function of firing frequency rather than of the forward-window definition. Data: the archived TDA-instanton backtest results file (conditional-return section, April 2026).

3.4 Overlap correction and a large-sample re-test (2026-07-11)

The statistics above were re-derived under a Newey–West heteroskedasticity- and autocorrelation-consistent variance, on a sample large enough to separate the channels: 9,743 valid trading days from 1987 to 2025 across SPY, QQQ, IWM, DIA and EEM, with the same 20-trading-day forward-drawdown outcome. Three results follow, and the second and third bear on figures printed earlier in this paper.

Measured inflation from overlap. On this signal family the naive statistic exceeds its autocorrelation-robust counterpart by a factor of 2.5 to 3.1: the instanton channel falls from 9.60 to 3.32, the combined signal from 16.00 to 5.64, and the topology channel from 7.03 to 2.21. Every crash-signal statistic produced by this pipeline is now required to be corrected before it is cited. Applying the measured range to the tight run of Table 1 places its combined figure between roughly 1.5 and 1.9, so the description of that run as clearing the Harvey threshold does not survive the correction, and it is withdrawn.

The combined form gains from the correction, on the larger sample. At 9,743 days the autocorrelation-robust combined statistic is 5.64, and the combined signal adds incremental lift over the instanton channel alone: the lift difference is 0.120 with a block-bootstrap 95% interval of [0.0625, 0.176], positive in 100% of resamples. The economic claim of Section 2.3 — that the topology channel carries information the barrier-shape channel does not — is supported by this test rather than by the tight run.

No topology-derived statistic is stable across engine versions, and the topology-only figure of Section 3.1 is withdrawn. Under the correction the topology-only channel reads 2.21, below the threshold, and a run-length-matched random control places it at only the 73.1st percentile, which is to say it is not distinguishable from a clustered random signal firing at the same frequency. The channel is also non-stationary across implementations: after a change to the persistent-homology engine the topology-alone statistic moved from −0.93 to +7.03, a sign reversal on the same underlying data. Any citation of a topology-derived statistic from this programme must therefore name both the corrected value and the engine version that produced it. The +3.47 printed in Section 3.1 satisfies neither condition and should not be quoted.

These re-tests are historical and are not the prospective validation of Section 4. They bear on whether the statistics already published are correctly computed, which is a separate question from whether the frozen protocol returns a verdict.

4. The Frozen Pre-Registered Validation

To convert the combined signal from a threshold-dependent observation into a Harvey-significant discovery, the tight-threshold configuration was frozen on 2026-04-18 as the canonical prospective validation set. It is registered verbatim and stated here in full; the numbers below are the frozen values, not proposals subject to later tuning.

Frozen pre-registered prospective protocol (frozen 2026-04-18).
(P1) Universe frozen *ex ante*: the five ETFs SPY, QQQ, IWM, TLT, GLD. No discretion in universe selection.
(P2) Thresholds frozen *ex ante*: fixed weights (0.4, 0.3, 0.3); component gates $P_{\text{inst}} \geq 0.35$ and $T \geq 0.75$; combined firing threshold $\theta \geq 0.40$ (in-sample fire rate 1.4%, precision 76.5%, recall 3.9%). These are not re-tuned; any change increments the freeze date and restarts the provisional quorum from zero, per the validation-limits protocol (Andreev & Howden, 2026e).
(P3) Out-of-sample window: prospectively observed days from the 2026-04-18 freeze date forward. The 2018–2023 in-sample span used to search the thresholds is *not* reused as a test set.
(P4) Outcome metric: a greater-than-5% forward maximum drawdown of the equal-weight basket within 20 trading days, conditional on signal-fire vs no-fire. Bias direction: signal-fire should precede larger drawdown.
(P5) Statistical test: conservative regression (equivalently Welch) t-test on conditional means; block-bootstrap with block length 20 days, $B = 1000$.
(P6) Harvey threshold: combined $t > 3.0$ over the out-of-sample window AND both component t-stats positive AND a minimum of 5 prospectively realised drawdown windows of magnitude > 5% (provisional quorum), advancing to 20 such windows for a robust verdict.

As of this working-paper revision the provisional quorum (5 prospective drawdown windows) has not yet been met, so the frozen protocol has not yet returned a verdict. The register status of the combined signal is therefore “provisional—frozen thresholds, awaiting prospective quorum”, not confirmed.

5. Discussion

The combined signal is a candidate crash predictor whose statistical confirmation requires prospective validation. We deliberately do *not* claim Harvey-significance for the combined signal based on the threshold-dependent empirical evidence; the multiple-testing hazard identified in Harvey–Liu–Zhu (2016) is exactly the failure mode the threshold-dependence exemplifies. The methodological position taken in the validation-limits paper (Andreev & Howden, 2026e) applies here as well: classical backtesting on different universe slices and different threshold choices is structurally inadequate for confirming a signal whose information content depends on the calibration.

We treat this paper as a derivation contribution and a methodological observation, not as an empirical confirmation. The combined signal is structurally well-motivated—persistent homology and instanton tunneling are mathematically independent objects encoding genuinely different aspects of crash precursor structure—and the empirical threshold-dependence documents the multiple-testing hazard that future implementations will need to navigate. The sequel promised in an earlier version of this paper has been run and is reported in Section 3.4 rather than deferred: on 9,743 days the corrected combined statistic is 5.64 with incremental lift over the instanton channel alone, the corrected topology-only channel sits below the threshold and inside a matched random control, and no topology-derived statistic reproduces across engine versions. The prospective protocol of Section 4 remains open, and its quorum is unmet.

References

[1] H. Edelsbrunner & J.L. Harer. *Computational Topology: An Introduction*. American Mathematical Society, 2010.

[2] G. Carlsson. Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308, 2009.

[3] S. Coleman. Fate of the false vacuum: Semiclassical theory. *Physical Review D*, 15(10):2929–2936, 1977.

[4] R.N. Mantegna. Hierarchical structure in financial markets. *European Physical Journal B*, 11(1):193–197, 1999.

[5] M. Gidea & Y. Katz. Topological data analysis of financial time series: Landscapes of crashes. *Physica A*, 491:820–834, 2018.

[6] V. Plerou *et al.*. Random matrix approach to cross correlations in financial data. *Physical Review E*, 65(6):066126, 2002.

[7] J.D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989.

[8] J.-P. Bouchaud *et al.*. Are financial crashes predictable? *Europhysics Letters*, 48(5):545–549, 1999.

[9] C.R. Harvey, Y. Liu & H. Zhu. …and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.

[10] M.A. Andreev & J. Howden. RG-scale-dependent instanton action for crash prediction in pilot wave finance. *Quark Quantitative Research Working Paper*, 2026b.

[11] M.A. Andreev & J. Howden. Grassmannian persistent homology for factor stability in pilot wave finance. *Quark Quantitative Research Working Paper*, 2026d.

[12] M.A. Andreev & J. Howden. Ontological limits of historical validation for wavefunction-calibrated financial systems. *Quark Quantitative Research Working Paper*, 2026e.

[13] M.A. Andreev & J. Howden. Non-perturbative extension of Cornish-Fisher VaR via dilute instanton gas. *Quark Quantitative Research Working Paper*, 2026g.

Quark Quantitative Research Platform — Working Paper
Generated 2026-08-25 22:11 UTC

This paper proposes a combined crash signal from persistent homology and instanton tunneling, documents its threshold-dependence as a multiple-testing hazard, and pre-registers a prospective validation protocol with frozen thresholds.
Correspondence: M.A. Andreev, Quark Quantitative Research.