

Market State Entropy: Order-Book Organization, Transition Uncertainty, and Evidence from Event-Driven Trading

AUTHORS: William Robert McConnell

SUMMARY: This paper extends the framework developed in Paper IX, focusing on the concept of market state entropy to analyze the organization of observable order-book activity and the uncertainty about the book's next state. It introduces a taxonomy of market-state entropy, covering various aspects such as event composition, queue allocation, and directional outcomes.

METHOD: The research employs a documentary case series assembled from author-supplied order-book displays and broker transaction images from a manual event-driven trading study. Six instrument examples are examined, including timed sell-then-buy sequences and long trades. The empirical evidence is derived from specific trading records and does not include a complete signal census or estimates of entropy's predictive effect.

KEY FINDINGS:

- Market state entropy provides a representation for testing, rather than a universal measure of fragility.
- The distributional summaries must be interpreted jointly with the geometry and support of the order book.
- Equal spatial entropy can accompany very different depths at the touch, and low transition entropy can describe a nearly certain adverse move.
- The framework offers a mathematically consistent public measurement system using standard finite differences and forecast-scoring comparisons.

IMPLICATION FOR TRADING: Traders can use this framework to better understand the state-dependent response of markets and the uncertainty surrounding future price movements. By analyzing the entropy of the order book, traders can make more informed decisions regarding their trading strategies, particularly in event-driven scenarios. This can help in assessing the resilience of the order book and predicting potential market transitions.

Market State Entropy

Order-Book Organization, Transition Uncertainty, and Evidence from Event-Driven Trading

MARKET PHYSICS RESEARCH SERIES | PAPER X | MPRS-010

William Robert McConnell

Acquired Market Physics LLC

Working paper | 17 September 2026

Foundational proposition

The response to pressure depends on the organization of the book that receives it. Entropy describes a specified distribution within that state; it does not replace the state, its geometry, or its capacity to resist displacement.

Continuation of Paper IX

Market Susceptibility and Structural Stiffness: A State-Based Theory of Nonlinear Price Response in Electronic Markets (July 2026).

This public research edition presents standard information-theoretic definitions and a documentary case series. Proprietary Behavioral Jitter implementation and execution logic are outside its scope.

Abstract

Electronic markets can display substantial resting quantity while remaining vulnerable to abrupt changes in executable prices. Paper IX developed market susceptibility and structural stiffness as state-dependent response constructs. This paper extends that framework by distinguishing the organization of observable order-book activity from uncertainty about the book's next state. It develops a taxonomy of market-state entropy covering event composition, queue allocation, commitment descriptions, spatial allocation, transition distributions, directional outcomes, and cross-market states. Each quantity requires an explicit alphabet, information cutoff, and sampling convention.

The central proposition is that distributional summaries must be interpreted jointly with geometry and support. Equal spatial entropy can accompany very different depth at the touch; low transition entropy can describe a nearly certain adverse move. Entropy therefore provides a representation to test, rather than a universal measure of fragility. Standard finite differences and forecast-scoring comparisons supply a mathematically consistent public measurement framework.

The empirical contribution is a documentary case series assembled from author-supplied book displays and broker transaction images associated with a manual event-driven trading study. Six dated instrument examples are examined, including a timed sell-then-buy FN sequence and several long trades. Later experimental trades are excluded from the first-trade-per-ticker selection rule. The records establish displayed book configurations and reconciled transaction outcomes; they do not constitute a complete signal census or estimates of entropy's predictive effect. A supplemental SPCX example is kept separate because its receipt date is cropped.

The paper supplies a falsifiable bridge from state-conditioned response to distributional measurement, with a prospective evaluation design that separates advance decisions, realized pressure, and subsequent outcomes. Proprietary production methods are not needed to reproduce the public mathematical examples and are not disclosed.

Keywords: order-book kinematics; market microstructure; entropy; liquidity geometry; state dependence; market susceptibility; event-driven trading.

JEL classifications: C58; G12; G14; G17.

Reader guide

Part I. The public research framework

Sections 1-5 connect this paper to Paper IX, define the observational unit, and distinguish deterministic book updates from uncertainty about future messages.

Sections 6-11 develop the entropy taxonomy, spatial counterexample, transition interpretation, temporal differences, and history-sensitive forecast evaluation.

Section 12 states the hypotheses. Section 13 establishes the information timeline.

Part II. Documentary evidence

Sections 14-15 define the case-series selection and transaction register. Sections 16-21 examine APLD, HOOD, FN, STX, RDDT, PLTR, and the supplemental SPCX record. Source plates accompany the three principal illustrated cases.

Section 22 states the combined empirical finding. Sections 23-25 define the next tests and robustness requirements.

Part III. Interpretation and reporting

Sections 26-28 discuss information, applications, disclosure, and conclusions. References and Appendix A complete the manuscript.

How to read the evidence

“Displayed” identifies content visible in a supplied source. “Reconciled” identifies arithmetic checked against that source. “Author-reported” identifies the observer’s account of the decision process. “Hypothesis” identifies a relationship proposed for further testing. These terms allow a documented result to stand without assigning it a stronger evidentiary role than the record supports.

All empirical amounts are USD. Source images are observations supplied by the author, not independently authenticated broker or exchange extracts. Tables transcribe relevant fields; original account identifiers and unrelated transactions are not reproduced.

1. Introduction

A limit order book is a changing set of executable commitments. At a given instant, displayed quantity has a location, a side, and a place in the matching process. Subsequent submissions, executions, cancellations, and amendments alter that configuration. The same aggressive demand can therefore encounter very different resistance at different times.

Paper IX made this state dependence central to the interpretation of price response. It distinguished market susceptibility from static thinness: a shallow book can replenish rapidly, whereas a visually deep book can lose support. Paper X asks how the distribution of activity, allocation, and possible successor states can be represented within that framework.

Entropy is useful because it makes the object being summarized explicit. An event-type distribution concerns the composition of observed messages. A spatial distribution concerns allocation across defined coordinates. A transition distribution concerns probabilities assigned to future states. Their entropies answer different questions. Treating them as one interchangeable measure of market disorder loses the distinction that makes them informative.

The proposed contribution has three parts. First, a disciplined taxonomy prevents confusion between diversity, location, and transition risk. Second, a set of public mathematical examples identifies claims that can be falsified without access to a proprietary signal engine. Third, documentary trading cases place the research questions in observed market settings, while keeping transaction outcomes separate from tests of the proposed entropy features.

The empirical motivation comes from the author's manual observation of bilateral asymmetry around scheduled events. The author reports that the earnings study extended prior macro-event work. The cases are presented as records of that application, not as a new randomized experiment or a completed estimate of strategy performance.

2. The bridge from Paper IX

Paper IX already treats response as directional, horizon-specific, and conditional on market state. It also discusses conditional distributions and recovery history. This paper does not claim to introduce those ideas for the first time. Its contribution is a particular information-theoretic representation and a clearer separation of the objects being measured.

The simplest bridge is a pair of questions. Susceptibility asks how much displacement accompanies specified pressure in a given state. Entropy asks how dispersed a specified distribution is. A market can be highly susceptible while its next move is relatively predictable, or resilient while its message stream contains a diverse mixture of activity.

Research continuity

The receiving state remains the central object. Distributional summaries are added to that state description and tested against conventional depth, spread, activity, and order-flow controls. They do not replace those controls or create information beyond the inputs from which they were computed.

Paper IX foundation	Paper X development
State-conditioned response	Explicit distributional summaries within the receiving state
Directional susceptibility	Separate probability, direction, and severity of future outcomes
Persistent support and replenishment	Distinguish quantity allocation from durability and event composition
Recovery and history dependence	Compare snapshot and history-based forecasts using held-out outcomes
Falsifiable empirical response theory	Test the added value of the public representation against strong baselines

3. Order-book kinematics and the observable state

Kinematics is used here to describe changes in recorded book coordinates and quantities over time. It does not imply a conserved liquidity mass, a universal market force, or an immutable law of future price movement. An observed state depends on venue coverage, feed semantics, book reconstruction, and the clock on which it is sampled.

At a minimum, a study must distinguish displayed bid and ask quantities at specified price coordinates, the best quotes, and the events that alter them. A market-by-price feed supplies aggregate depth; it does not necessarily supply persistent individual order identities. A consolidated display can combine sources whose clocks and update rules differ. These distinctions matter before any statistical transformation is applied.

Observed persistence is equally specific. A horizontal heatmap band can establish persistence of aggregate displayed liquidity at approximately the same coordinate. It cannot by itself establish that the same orders remained, that no cancellations occurred, or that the displayed quantity would absorb unlimited pressure. The terms “Iron” and “Sand,” where used informally in the research program, describe contrasting apparent persistence; they are not claims of guaranteed support or participant intention.

The public framework therefore separates three records: the book state visible before a decision, the subsequent sequence of observed updates, and the execution record. Each can be valuable even when the others have different temporal resolution. Joining them requires time alignment; it does not require treating every visual feature as an order-level measurement.


4. Determined updates and uncertain paths

A matching rule determines how a valid incoming instruction changes a sufficiently specified book. That conditional determinism does not specify the next instruction, its arrival time, hidden liquidity, or the behavior of other venues. It therefore does not make the entire future price path deterministic.

A useful mathematical abstraction is a state update $B(\text{next}) = F(B(\text{now}), \text{message}, \text{venue rules})$. Reconstruction asks whether the recorded sequence and the applicable rules reproduce the observed book. Forecasting asks which future messages and states will occur. Those are distinct tasks.

The distinction matters when a best quote moves across several levels. The event can reflect executions, cancellations, replacements, or a combination. A screenshot of the resulting staircase is evidence of the displayed path. Attributing each step to a particular event class requires the corresponding messages. An apparently empty displayed level also does not establish the absence of non-displayed interest.

Figure 1 makes the information boundary explicit. The update operator can be studied without claiming that the incoming sequence was known before the event.



The next message is not known merely because the update rule is known.

Figure 1. Book update state machine.

5. What “market-state entropy” means

For a discrete probability vector p over a fixed alphabet of K categories, Shannon entropy is defined in bits by Equation (1). The convention is $0 \log_2(0) = 0$. The probability vector must be nonnegative and sum to one. Entropy is a property of that vector, not a free-standing property of an asset (Shannon, 1948).

$$H(p) = - \sum_{k=1}^K p_k \log_2 p_k \quad (1)$$

$$H_{\text{norm}}(p) = \frac{H(p)}{\log_2 K}, \quad K > 1 \quad (2)$$

Normalization facilitates comparison only when the category construction is itself comparable. Altering the number of bins, the price range, or the event definition can change the result without a corresponding change in economic conditions. A no-event window has no empirical event-share vector unless a separate convention is declared; it should not automatically receive the same interpretation as a concentrated active window.

Three interpretations must remain separate. Marginal entropy describes diversity within an observed distribution. Conditional entropy describes residual uncertainty given stated information. Estimated forecast entropy describes dispersion in a fitted probability forecast, whose accuracy must be evaluated separately.

The term “disorder” is consequently restricted to an economic interpretation supported by additional evidence. High entropy can accompany orderly adaptation. Low entropy can accompany one-sided collapse. Neither value identifies a trading direction without information about what the categories mean and where support lies.

6. Taxonomy of distributional objects

The taxonomy preserves the breadth of the original framework while separating mathematical definitions from any production classifier. The categories below are research objects, not components of a disclosed trading score. Each requires its own observation rule and validation.

No universal aggregation is proposed. Combining these entropies into one number would require choices about units, weights, missing values, and economic objectives. Those choices would need an independent justification. Reporting the components separately is more informative at the present stage.

Family	Distribution being summarized	Interpretive boundary
Event	Shares of defined message types	Composition is not message intensity
Queue	Allocation across observable queue categories	Aggregate depth cannot reveal missing order identities
Commitment	Shares of explicitly defined research persistence categories	Commitment is not directly observed intention
Spatial	Displayed size shares across fixed price bins	Permutation invariant; location must be reported separately
Transition	Probabilities of defined successor states	Forecast concentration is not forecast correctness
Directional	Probabilities of up, flat and down outcomes	Entropy alone does not identify the favored direction
Cross-market	Distribution over synchronized market states	Common shocks and stale observations can create apparent dependence

7. Event composition and queue allocation

Event entropy summarizes the mixture of recorded activities over a stated window. A public example uses mutually exclusive categories for additions, executions, and cancellations, with amendments treated according to the feed's own semantics. If $n(k)$ is the number of category- k observations and N is their sum, the empirical shares are $n(k)/N$ for $N > 0$. Counts and entropy should be reported together.

A window with three equally common event types can have the same event entropy whether it contains thirty messages or thirty thousand. Activity intensity is therefore a separate covariate. A nearly homogeneous cancellation stream has low event entropy even though the associated withdrawal can be economically important. Event diversity is not a monotone measure of danger.

Queue entropy poses a different observation problem. If the feed supplies individual displayed order quantities $q(j)$, their shares can describe concentration among those orders. If only aggregate price-level quantities are available, the same construction describes level concentration, not order-level queue composition. Hidden orders and feed-specific aggregation limit the interpretation in either case.

The prospective implementation should report the denominator, eligible categories, empty-window policy, and treatment of feed resets. These are necessary to reproduce a public measure. They do not require disclosure of proprietary behavioral state assignments, economic weights, or live decision rules.

The literature motivates this separation: Cont, Kukanov, and Stoikov (2014) examine price changes in relation to order-book events, while Eisler, Bouchaud, and Kockelkoren (2012) distinguish immediate effects from changes in subsequent event rates and gaps. An entropy transform is a candidate summary of such recorded activity, not a substitute for the underlying event history.

8. Spatial entropy needs coordinates

Let $q(j)$ denote displayed quantity in a fixed set of price bins and let $s(j) = q(j)/\sum q(j)$, when total quantity is positive. Applying Equation (1) to s measures allocation dispersion. It does not identify which bin is nearest the touch.

The constructed example below is deliberately simple. The bins are ordered from nearest to farthest from a fixed reference. Both states contain 100 units and have identical Shannon entropy. Yet one places 90 units in the nearest bin and the other places none there. Their immediate displayed resistance differs even though their entropy does not.

Constructed counterexample, not an empirical estimate

Geometry must accompany entropy: the reference price, ordered depth vector, empty levels, spread, and price-bin width. A mean distance can help summarize location but is not sufficient to distinguish every pattern either. A small set of summaries should not be confused with the complete ordered book.

Quantity / statistic	State A	State B
Nearest bin, distance 1	90	0
Middle bin, distance 2	10	10
Farthest bin, distance 3	0	90
Total displayed quantity	100	100
Shannon entropy (bits)	0.4690	0.4690
Size-weighted mean distance	1.1	2.9

9. Transition uncertainty and directional risk

For a stated horizon, a public outcome alphabet can distinguish upward displacement, no qualifying displacement, and downward displacement. The price reference and displacement rule must be fixed before outcomes are evaluated. Other alphabets can distinguish spread widening, repair, or repeated dislocation, but their definitions must not depend on the proposed predictor.

Let $\hat{p}(k | I)$ be a forecast based on information I available at the decision cutoff. Its entropy summarizes the forecast's concentration. A concentrated forecast can be wrong. Forecast entropy therefore belongs alongside calibration and proper scoring, rather than serving as evidence of forecast quality by itself.

$$H_{\text{forecast}}(I) = - \sum_k \hat{p}_k(I) \log_2 \hat{p}_k(I) \quad (3)$$

Consider two forecasts over up, flat and down: (0.98, 0.01, 0.01) and (0.01, 0.01, 0.98). They have identical entropy and opposite directional implications. Similarly, a nearly certain damaging event has low entropy. A risk estimate must retain the probability and magnitude of the adverse outcome; entropy alone does not supply either.

Marginal state occupancy and conditional transition uncertainty can also differ sharply. An alternating two-state sequence can spend equal time in each state while the next state is predictable from the current state. Equal occupancy is compatible with high marginal entropy and low conditional entropy. This standard distinction is sufficient for the public argument; it does not require publishing a proprietary switching estimator.

10. Commitment, direction and cross-market context

Commitment entropy remains a conceptual category in this edition. A reproducible empirical version would require publicly specified, observable categories and a stated observation window. A proprietary classification cannot be treated as a fully reproducible public measure merely by publishing its final entropy. Nor does a persistent band reveal the economic intent of the participants whose orders contribute to it.

Directional entropy can be computed from a forecast over signed outcomes or from a clearly defined historical outcome distribution. Those constructions should not be mixed. A historical distribution estimated using future outcomes cannot be used as though it were available at the earlier decision cutoff. A directional probability vector should be reported alongside its entropy because the entropy is unchanged when up and down labels are exchanged.

Cross-market analysis introduces another distinction. The diversity of state labels across assets at one instant is not the joint entropy of a multivariate time series. Neither quantity alone establishes transmission between markets. Different trading hours, reporting delays, missing symbols, and common news can change the apparent relationship.

A public cross-market study should specify the synchronized observation set, category definitions, weighting convention, and missing-data policy. A lead-lag claim requires clocks and information availability consistent with the claimed lead. In the present documentary examples, no cross-market entropy estimate is reported.

These families remain in the theory because they ask useful questions. Their empirical status is explicit: definitions and proposed research uses are retained, while unavailable estimates and production implementations are not supplied.

11. Temporal change and history

An entropy series requires two separate time choices: the estimation window h and the displacement between evaluations Δ . Confusing them produces incorrect derivative notation. With a fixed construction of $H(t; h)$, backward finite differences provide causal descriptive changes at time t .

$$v_H(t) = \frac{H(t; h) - H(t - \Delta; h)}{\Delta} \quad (4)$$

$$a_H(t) = \frac{H(t; h) - 2H(t - \Delta; h) + H(t - 2\Delta; h)}{\Delta^2} \quad (5)$$

If physical time is measured in seconds, the units are bits per second and bits per second squared. With an event-index clock, the units change accordingly. Neither expression is presented as a proprietary signal. Both are standard differences whose behavior depends on sampling, estimation noise, and overlapping windows.

For history, the population information identity concerns conditional entropy averaged over the conditioning variables. With future outcome Y , current state S , and earlier history R :

$$I(Y; R | S) = H(Y | S) - H(Y | S, R) \geq 0 \quad (6)$$

This identity does not establish that useful history is present in a particular dataset. Estimated, pointwise forecast entropies can move in either direction. The empirical question is whether including eligible history improves predictions on observations not used to fit the model.

$$\Delta L = \frac{1}{m} \sum_{i=1}^m \log_2 \frac{\hat{p}_i^{\text{history}}(y_i)}{\hat{p}_i^{\text{snapshot}}(y_i)} \quad (7)$$

A positive held-out ΔL indicates lower average log loss for the history model on that evaluation set. Here m is the number of evaluated outcomes, y_i is the realized category, and the two estimated probabilities are indexed by observation i . It is not automatically a causal effect or a universally positive “hysteresis gap.” Forecast probabilities, treatment of zeros, and the fitting protocol must be specified before testing (Gneiting and Raftery, 2007).

12. Falsifiable propositions

The framework is useful only if its representations survive comparisons that could reject them. The following propositions are distinct from the documentary results later in the paper.

H1. Representation value

At a fixed forecast horizon, a public entropy representation improves held-out probability scores relative to a specified reduced baseline. A stronger comparison includes the underlying share vector. If the entropy summary adds no value, the proposed representation benefit is unsupported. A deterministic transform cannot add information beyond all of its inputs.

H2. Geometry beyond dispersion

Ordered depth and gap information distinguish outcomes among observations with similar total depth and similar spatial entropy. Failure to do so would weaken the practical importance of the spatial distinction in that setting, although the mathematical permutation counterexample remains true.

H3. State-conditioned response

Within supported comparisons of directional pressure, response differs across pre-pressure states. The test must add value beyond Paper IX's conventional controls. Conditioning retrospectively on realized pressure is not the same task as forecasting before that pressure exists.

H4. History beyond snapshot recovery

Eligible prior history improves held-out forecasts after displayed depth has recovered. If snapshot-only forecasts perform equally well within a predeclared practical tolerance, the history claim is unsupported for that sample.

H5. Generalization

Any observed improvement persists across a chronological holdout and survives reasonable variation in sampling conventions. A result concentrated in one selected event is a case finding, not a general law.

These propositions do not presuppose that entropy must succeed. A clear negative result would identify where the receiving-state representation should be simplified or replaced.


13. The information timeline

The relevant forecast cutoff is the time at which information could have been used, not simply the timestamp printed on the underlying message. Exchange event time, local receipt time, display time, and decision or order time can differ. A study must choose the clock appropriate to its claim.

For manual trading, an executed position before a documented release can establish advance directional commitment without a separately written directional call. The accompanying book record supports the explanation of the observed state. A contemporaneous written call is another useful record, but it is not a universal prerequisite.

The earnings release time and the time the observer's news system received that release are also different quantities. "Before the packet" requires a definition of the packet, its source, and the relevant receipt clock. A calendar time near an announcement is not, by itself, a measured matching-engine response interval.

The timeline in Figure 2 separates observations available for a decision from the later pressure and outcomes used to evaluate it. A later screenshot may preserve a historical point; the capture time and historical cursor time must then be kept distinct.



Historical cursor time, local receipt time and execution time are distinct.

Figure 2. Observation, decision cutoff and subsequent outcomes.

14. Documentary case-series design

The case material comes from author-supplied Bookmap displays, transaction-history images, and research records associated with an event-driven trading study. The author describes the process as observing bilateral asymmetry before earnings information arrived and entering when the asymmetry was judged sufficiently strong. This description provides the study context; exact entry-to-release timing is reported only where the records establish it.

The selection rule for this revision is the first trade per ticker within the study interval, July 27 through August 18, 2026. Later experimental trades are excluded from the primary case register. This rule was clarified during retrospective assembly; it is not represented as a prospectively registered sampling plan. When first-trade order is inferred from a display rather than execution times, that inference is identified.

Six dated instrument examples are retained in the main register. APLD's first-trade assignment is provisional. The series is a selected documentary collection, not a census of all eligible signals, all trades, or all account activity. Consequently the paper reports individual outcomes, not a cohort win rate, a strategy return, or a significance test based on the number of profitable rows.

The evidence has three roles. Images document what was displayed. Broker fields document the recorded transaction amounts and, where visible, times. The author's account identifies the intended interpretation and study process. These roles are joined case by case rather than collapsed into an unsupported claim of fully reconstructed exchange causality.

No original entropy estimates are extracted from the still images. The cases motivate and illustrate the state-based questions. They do not test the predictive contribution of entropy or certify the proprietary method.

15. Dated transaction register

The table transcribes one candidate first round trip per instrument. Amounts are differences in the supplied cash fields, including displayed transaction fees. They are not account returns, risk-adjusted returns, or an estimate of expected performance. Unshown borrowing, financing and tax costs are not estimated. FN is a sell-then-buy sequence; the other listed sequences are buy-then-sell.

APLD is retained as an apparent first trade based on descending history order. STX's smaller 100-share round trip is retained because the author identifies the 40-share trade as a follow-up. PLTR's later 500-share trade is excluded. The selection therefore does not retain the largest profit merely because it is larger.

**APLD ordering is inferred, not time-verified. Prices may be rounded on the source display; cash fields govern reconciliation. Sources: D1-D3 in Appendix A. There is no aggregate total because these are selected cases with incomplete first-trade and signal-denominator verification.*

Date / ticker	Sequence and shares	Displayed prices	Net cash difference
27 Jul / APLD*	Buy, sell / 800	\$26.395 → \$28.0576	\$1,329.46
28 Jul / STX	Buy, sell / 100	\$744.28 → \$748.00	\$370.44
29 Jul / HOOD	Buy, sell / 500	\$89.585 → \$92.50	\$1,456.45
30 Jul / RDDT	Buy, sell / 400	\$178.64 → \$182.00	\$1,342.42
03 Aug / PLTR	Buy, sell / 1,400	\$125.95 → \$129.00	\$4,266.01
17 Aug / FN	Sell, buy / 40	\$640.54 → \$576.55	\$2,559.11

16. APLD: the receiving state

The two APLD frames provide a useful example of state observation before a reported trade. The earlier display shows \$26.41 and a latest visible Time & Sales entry at 15:55:35.955. The later display shows \$26.42 and a latest visible entry at 15:55:39.871. The top Order Book gauge changes from +00% to +19%, while Volume changes from +36% to +37%. These are platform-displayed summaries, not independently reconstructed order-flow measures.

Both frames show prominent displayed depth around \$26 and other lower coordinates, together with overhead clusters. This is a spatially structured book. The bright bands support a description of aggregate persistence over the displayed interval. They do not establish that individual orders never changed, and the annotation beside the last-price line is not treated here as an identified ask queue.


The apparent first transaction is an 800-share purchase at \$26.395 followed by a sale at displayed average price \$28.0576. The supplied cash entries are \$21,116.00 paid and \$22,445.46 received, giving \$1,329.46. The separate 500-share round trip is not added to the primary result.

The empirical content is concrete: a documented book configuration, a reported directional interpretation, and a profitable recorded outcome. The present images do not measure a cancellation fraction or transition frequency. They also do not identify the fill times. Those narrower limitations leave the documented state and cash outcome intact.


Source: E1 and D1. The calendar date is supplied by the transaction record and the author's case association; it is not visible in the two book frames.

Exhibit E1. APLD book-state sequence

Original author-supplied frames, scaled proportionally. Earlier and later refer to the latest visible tape times, not independently established capture times. No proprietary calculation is overlaid.



E1a. Displayed price \$26.41; Order Book +00%; Volume +36%.



E1b. Displayed price \$26.42; Order Book +19%; Volume +37%.

17. HOOD: persistence and discrete displacement

The HOOD pre-event image shows a displayed price near \$89.58, prominent lower bands near \$88, \$86 and \$85, and variable overhead depth. The appropriate description is macroscopic persistence at approximately fixed coordinates. There are visible intermediate overhead clusters, so the upper region is not characterized as a completely empty corridor.

The historical pulse image presents a stepped quote path from roughly \$88.40 through \$90.34 and above. Its tooltip identifies a historical point on July 29 at 16:04:58.015267051 EDT. The display therefore preserves a fine timestamp and an abrupt path. Timestamp digits alone do not establish effective measurement accuracy, exchange sequencing, or the latency of a news-to-engine response. The contemporaneous tape pane in that capture is not treated as belonging automatically to the historical cursor point.

The broker transaction image shows a 500-share purchase at \$89.585 and sale at \$92.50. Purchase cash of \$44,792.50 and sale proceeds of \$46,248.95 reconcile to \$1,456.45. The \$1.05 sale fee reconciles the gross price gain of \$1,457.50 with those proceeds.

HOOD thus supplies a useful combination of book geometry, a displayed discrete path, and a matched cash outcome. It does not, in this edition, isolate how much of the displayed path arose from cancellations rather than executions. Nor is the historical point assigned to a release packet solely because it lies near a nominal announcement time.


Source: E2 and D1. The earlier \$89.58 state and the later approximately \$88.40 baseline are separate observations; intervening repricing is not filled in by an invented continuous record.

Exhibit E2. HOOD state and historical pulse

Both panels reproduce supplied Bookmap displays. The pulse panel's historical tooltip is the relevant time reference. Other visible panes can reflect a different display time.



E2a. Displayed price \$89.58; persistent lower bands near \$88, \$86 and \$85.



E2b. Historical tooltip: 29 July 2026, 16:04:58.015267051 EDT.

18. FN: a timed short-side outcome

FN supplies the clearest execution-time sequence in the selected transaction material. The broker image identifies August 17, 2026, a sale of 40 shares at displayed price \$640.54 at 16:15:44, and a purchase of 40 shares at \$576.55 at 16:25:07. The difference between the displayed times is 9 minutes 23 seconds. The timezone is not shown on the receipt.

The cash reconciliation uses the amounts actually printed on the source: \$25,621.65 for the sale, \$23,062.00 for the purchase, and \$0.54 commission. Their difference is \$2,559.11, matching the displayed realized P&L. Multiplying the rounded sale price by 40 gives a value five cents below the displayed sale amount; the underlying unrounded fill detail is not supplied, so the cash fields are retained without inventing a correction.


The receipt supports a sell-then-buy outcome consistent with the author's reported short trade. It demonstrates that the supplied results are not exclusively long-side examples. That observation does not by itself establish the distribution of all long and short signals or remove selection bias.

A subsequent FN book display shows a decline toward the mid-\$560s and a later displayed price of \$570, with tape entries around 16:28:55. This is useful outcome context, not a pre-entry forecast exhibit. Another FN image in the private collection explicitly carries a delayed-data notice; that image is not used as evidence of contemporaneous signal availability.

FN strengthens the documentary outcome record. It is not presented as proof that the short entry preceded a particular earnings packet, because that comparison requires a source-defined release and aligned clock.

Exhibit D3. FN transaction sequence


Supplied broker receipt. Trade date, amounts, commission and execution times are visible. The displayed timezone is unspecified. The existing blue marks are part of the supplied image.



40 shares sold, then bought back. Displayed realized P&L: \$2,559.11.

Exhibit E3. FN outcome context

Original supplied book display, after the execution times shown in D3 if the display and receipt clocks are aligned. This exhibit is used only as outcome context, not proof of a pre-entry signal.



FN:NQ@BMD. Visible tape entries near 16:28:55; displayed price \$570.

19. STX: preserving the first-trade rule

The STX example illustrates why selection must be independent of which row produces the most impressive gain. The supplied transaction record contains a 100-share round trip and a later 40-share round trip. The author identifies the latter as a follow-up. The primary case therefore uses the 100-share purchase at \$744.28 and sale at \$748.00.

The purchase amount is \$74,428.00. The sale proceeds are \$74,798.44 after the displayed \$1.56 fee. Their difference is \$370.44. This is the STX outcome retained in the main register, even though the follow-up has a much larger percentage move.

The 40-share follow-up is excluded from the primary case series. Its source cash arithmetic is \$32,326.56 less \$29,940.40, or \$2,386.16. A secondary summary in the source collection reported a different amount; the underlying receipt governs. The correction does not alter the retained 100-share case.

This choice is scientifically important. A retrospective collection should not silently substitute a later successful experiment for the first decision under study. The same rule applies to smaller gains and losses if they occur in the eligible record. First-trade selection reduces one form of discretionary substitution, but it does not turn an incomplete archive into a complete prospective sample.

STX is retained as transaction evidence. This edition does not assign a measured entropy value to its setup or claim an independently established lead time from the supplied date-only transaction rows.

Source: D1. The first-trade assignment combines the supplied transaction record with the author's identification of the follow-up.

20. RDDT and PLTR: additional dated outcomes

RDDT

The July 30 RDDT record shows a 400-share purchase at \$178.64 and sale at \$182.00. Cash paid of \$71,456.00 and proceeds of \$72,798.42 give a net difference of \$1,342.42. The private archive also contains a directional research narrative associated with this event.

The narrative is relevant to the author's interpretation, but its operational ratios and state thresholds are not reproduced. Their omission does not erase the trade outcome. It keeps the empirical statement focused on what can be checked in the public record: symbol, transaction date, quantity, displayed prices, and cash reconciliation. No entropy-performance estimate is derived from the narrative.

PLTR

The August 3 PLTR record shows a 1,400-share purchase at \$125.95 and sale at \$129.00. Cash paid of \$176,330.00 and proceeds of \$180,596.01 reconcile to \$4,266.01. A later 500-share PLTR round trip in the collection is excluded under the first-trade rule.

The size of the PLTR cash outcome reflects both position size and the change in price. It is not directly comparable to another instrument's percentage return or to return on account equity. No capital-allocation rule is inferred from the amount alone.

Role in this paper

These records broaden the documentary examples beyond one ticker. Their inclusion does not make the selected rows independent statistical replications, establish an exhaustive signal denominator, or validate the omitted proprietary method. They provide additional dated outcomes to which a fuller event chronology can be attached without changing the public theoretical framework.

Sources: D1-D2. Execution times are not visible on these receipt images.

21. Supplemental SPCX record

SPCX is retained as a supplemental example outside the dated main register. The author-supplied receipt shows 1,000 shares purchased at \$125.45 and sold at \$128.00, with purchase cash of \$125,450.00 and net sale proceeds of \$127,997.16. The difference is \$2,547.16. The relevant receipt crop does not show the transaction date.

The associated book frames show displayed prices of \$124.92 and \$124.44. Their latest visible tape entries are 15:37:02.479 and 15:40:36.447, respectively. The Order Book gauge changes from +12% to +37%, while Volume changes from +11% to +8%. The observed price change between these setup frames is downward, not the later profitable upward outcome.

The closer view includes resting clusters around \$126, \$127 and \$128. The wider view uses a different vertical scale and includes more distant liquidity. These are not treated as directly comparable individual queues without accounting for display aggregation. The term variable-density corridor is appropriate; an empty path to the eventual exit is not established.

The author supplies a NORTH interpretation linking the setup and the trade. The current narrative contains the realized outcome as well as the interpretation. It can therefore serve as a descriptive case narrative, while the timing of the original decision remains a separate item in the research record.

The missing calendar date determines the supplemental status. This keeps interval eligibility consistent while retaining the separately documented transaction outcome.

Sources: E4 and D4. The instrument label is transcribed as supplied; this edition does not independently establish its security-master mapping.

22. What the case series establishes

The six dated examples and the supplemental SPCX record provide documentary evidence of displayed book structures and recorded trading outcomes. The FN receipt additionally supplies an explicit execution interval. The source images make the cases more concrete than a purely hypothetical discussion of liquidity states.

The principal empirical lesson is that the observational objects must remain distinct. A bright persistent band describes aggregate displayed support over an interval. A best-quote staircase describes a displayed path. A broker cash difference describes a transaction outcome. An advance directional commitment is established by appropriately timed decision or entry records. None of these objects should be silently substituted for another.

The first-trade rule prevents later experiments from inflating the primary register. APLD remains an inferred first-trade assignment; supplemental undated records remain outside the dated set. The collection does not supply the full opportunity set, including decisions not to trade. For this reason, no “all signals correct” statement or probability of success is calculated from the selected profitable examples.

The cases are consistent with the research motivation that book state matters to event-driven price response. They do not identify an incremental entropy effect, distinguish all information and liquidity channels, or demonstrate a universal deterministic path. Those claims require the designs in the next sections.

This is a positive empirical contribution with a defined scope: it assembles observable states, transaction records, and explicit source boundaries into a case series that can be extended. The contribution does not depend on disclosing the author's live classifier or portraying the present sample as a completed strategy validation.

23. A prospective test of the public framework

The next empirical stage should begin with an eligible event list fixed before outcomes are observed. For each event, record the instrument, session, relevant announcement source, information cutoff, and whether a trade was taken. No-trade decisions belong in the opportunity record. Failed signals and interrupted data sessions should remain visible under declared handling rules.

The public test can use ordinary feed-observable variables: spread, depth at fixed distances, signed order-flow summaries, event counts, and the entropy examples defined here. It need not reproduce the proprietary production engine. This preserves a meaningful research test while keeping the operational architecture separate.

Freeze feature construction and outcome definitions before the chronological evaluation interval. If parameters are fitted, use earlier data for fitting and validation, then hold the later interval aside. Event windows that overlap in time must not be treated as independent simply because they contain different snapshots.

Primary outcomes should answer a small number of questions: the probability of a defined price move, its signed magnitude, or the distribution of transition states. Separate these from any execution-performance outcome. A statistically useful state forecast can be uneconomic after costs; a profitable discretionary trade does not alone establish forecast calibration.

The original case series remains valuable as motivation and as a source of hypotheses. It should not become both the design set and the supposedly untouched test set. The prospective study succeeds scientifically if it reports the prespecified results and rejection conditions, including the possibility that a simpler representation performs as well.

24. Baselines and evaluation

The benchmark should be capable enough to make improvement informative. A baseline that omits spread, depth, activity, or ordinary order-flow information can give an entropy transformation undeserved credit for variables it merely summarizes. The same observations and forecast cutoffs must be used for every comparison.

Compare at least three representations: conventional covariates; conventional covariates plus the public entropy summaries; and, where feasible, conventional covariates plus the underlying share vectors and geometry. The third comparison distinguishes useful compression from a claim of new information.

For probabilistic forecasts, use a proper score with a stated sign convention. Log loss is lower when the model assigns more probability to the realized outcome. Calibration checks whether events assigned similar probabilities occur at corresponding frequencies. Sharpness is concentration of the forecast; it is desirable only in conjunction with calibration. Report uncertainty in score differences rather than the best isolated event.

For state-conditioned response, distinguish a forecast made before future flow from a retrospective response analysis that holds realized flow approximately comparable. Realized pressure is a future variable at the earlier forecast cutoff. Using it for an ex-ante forecast would change the task.

Matched-pressure comparisons can support conditional associations but do not by themselves remove all confounding. Information releases, volatility changes, and market-wide conditions can jointly affect pressure and support. The model and interpretation should reflect that limitation instead of assigning all displacement to one mechanism.

The general scoring approach follows Gneiting and Raftery (2007). The receiving-state comparison follows the empirical question posed in Paper IX.

25. Robustness, nulls and rejection

A useful robustness exercise changes an assumption that could matter. Vary the public sampling window within a declared range, check whether price-bin definitions alter the result, and test sensitivity to feed gaps or session boundaries. Repeating the same favorable chart with a new label is not an independent test.

Null models must correspond to the statistic being tested. Permuting events within a fixed window while preserving category counts leaves marginal event entropy unchanged. Such a permutation can test temporal ordering or transition dependence, but it is not a meaningful null for the unchanged marginal statistic. Likewise, permuting price-level quantities preserves spatial entropy while deliberately altering geometry.

Dependence requires attention at the event and session level. Thousands of overlapping windows from one earnings episode do not supply thousands of independent earnings events. Confidence intervals and resampling units should match the dependence structure, and the number of hypotheses should be limited before testing.

The public framework should be narrowed when improvements disappear under strong baselines, when forecast probabilities are systematically miscalibrated, or when performance depends on a single undocumented data correction. Boundary conditions are findings: a representation may be useful in one venue, horizon, or event class and unhelpful elsewhere.

Case-level accounting and model-level validation answer different questions. Reconciled cash outcomes remain valid observations even when a proposed explanatory model fails. Conversely, accurate accounting does not substitute for a test of the model. This separation allows both positive and negative research findings to be retained.

26. Information and mechanical response

Information and order-book mechanics are compatible explanations at different levels. New information can change valuations and quoting decisions; those decisions enter a market through submissions, cancellations, replacements, and executions. The realized path then depends on the receiving book and the arriving sequence.

This paper therefore does not infer that discrete quotes refute informational efficiency. It asks whether a richer state description explains or predicts variation that a specified reduced model misses. Evidence of state dependence would be economically meaningful without requiring the stronger claim that information is irrelevant.

Nor do anonymous book updates establish balance-sheet withdrawal, institutional identity, or intent. A decrease in displayed supply is observable under appropriate data. Its economic motivation is generally not identified by the book alone. Figure 3 sets out plausible channels without declaring any one of them estimated in the present case series.




Figure 3. Candidate information and mechanical channels; no causal effect is estimated here.

27. Applications and research disclosure

The public framework can support research on event monitoring, liquidity recovery, and state-conditioned execution costs. These are possible applications, not deployment recommendations or validated trading instructions. A production rule requires its own calibration, monitoring, and evaluation.

Behavioral Jitter is retained as the name of a separate proprietary research component. This public edition does not disclose its estimator, state classifier, weights, decision thresholds, window schedule, calibration procedure, or execution logic. The paper's reproducible mathematical content consists of standard entropy definitions, illustrative counterexamples, and a public evaluation design. The documentary cases do not depend on a reader reconstructing the withheld component.

Data and reproducibility

The author retains original screenshots, transaction images, and a source manifest. This edition includes selected book displays, a relevant FN receipt, and transcribed cash fields. It does not distribute a raw exchange-message dataset or executable production system. The public case evidence can be inspected at the documentary level; message-level reconstruction and prospective forecasting tests remain distinct research stages.

Author contribution and interests

The author developed the research program, supplied the observations and transaction records, and directed the manuscript. Artificial-intelligence tools assisted with drafting, mathematical checks, source organization, and editorial revision. The author is responsible for the final claims and submission.

The author is affiliated with Acquired Market Physics LLC and reports commercial and patent-pending interests related to the research program. No patent specification, application identifier, private conversation, or claim of granted patent protection is included. This statement is a research-interest disclosure, not an assessment of patent scope.

28. Conclusion

Paper IX established the receiving market state as the object through which directional pressure produces displacement. Paper X extends that argument by distinguishing several distributional descriptions of the state and its possible successors. Its central result is conceptual but exact: dispersion is not location, uncertainty is not direction, and conditional book mechanics are not knowledge of the future message sequence.

The spatial counterexample shows why entropy cannot replace geometry. The transition examples show why low entropy can describe either favorable stability or a nearly certain adverse move. The history formulation supplies a valid forecast comparison without assigning an impossible sign to population conditional entropy. These distinctions make the framework testable rather than merely metaphorical.

The documentary cases add observed book configurations and reconciled trading outcomes to the series. APLD and HOOD illustrate pre-event state records; FN contributes a timed short-side outcome; STX, RDDT and PLTR provide additional dated transaction examples. SPCX remains a supplemental record because the supplied receipt crop does not establish its date. First-trade selection prevents later experiments from being substituted for the decisions under study.

The evidence supports publishing a substantive observational case series. It does not yet estimate entropy's independent predictive contribution or establish complete strategy performance. The next step is a fixed opportunity set, aligned event records, and held-out comparison against strong conventional baselines.

The momentum from Paper IX is therefore preserved in a precise way: from how state conditions response to how state organization and transition uncertainty can be measured. The public contribution stands on its definitions, observations, and rejection criteria, while proprietary implementation remains outside the manuscript.

References

Cont, R., Kukanov, A., and Stoikov, S. (2014). The price impact of order book events. *Journal of Financial Econometrics*, 12(1), 47-88. Author manuscript: <https://arxiv.org/abs/1011.6402>

Eisler, Z., Bouchaud, J.-P., and Kockelkoren, J. (2012). The price impact of order book events: Market orders, limit orders and cancellations. *Quantitative Finance*, 12(9), 1395-1419. Author manuscript: <https://arxiv.org/abs/0904.0900>

Gneiting, T., and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359-378. <https://doi.org/10.1198/016214506000001437>

McConnell, W. R. (2026). Market Susceptibility and Structural Stiffness: A State-Based Theory of Nonlinear Price Response in Electronic Markets. Market Physics Research Series, Paper IX, MPRS-009. Acquired Market Physics LLC. July working-paper version supplied by the author as the SSRN edition.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379-423 and 623-656. Publisher archive: <https://www.nokia.com/bell-labs/publications-and-media/publications/a-mathematical-theory-of-communication/>

Documentary sources

The empirical exhibits are author-supplied primary observations rather than published literature. Appendix A identifies the source objects and their evidentiary role. Their inclusion does not imply endorsement by the broker, charting platform, exchange, or data provider.

Bibliographic continuity is anchored to the supplied July Paper IX. A separate manuscript with a different title is not substituted as the predecessor.

Appendix A. Source map and reporting standard

The source labels below identify the retained private originals without publishing personal filesystem paths. Images have been reproduced proportionally, without adding claimed measurements. Transcriptions preserve the relevant cash fields. Source identity and cash arithmetic are separate from independent authentication.

Minimum standard for subsequent estimates

Report venue/feed coverage and clocks; the observation and event population; category definitions and sampling; the forecast cutoff; outcome definitions; handling of missing data; fitting and holdout periods; baseline comparisons; dependence-aware uncertainty; exclusions and failed tests. Operational information withheld from this edition must not be portrayed as publicly reproducible.

Notation

H: discrete entropy in bits. K: number of categories. h: estimation-window length. Δ : evaluation step. I: information available at cutoff. Y: future outcome. S: current state. R: eligible prior history. ΔL : held-out log-loss improvement from adding history.

Label	Retained source	Use in this edition
D1	July 31 field-note transaction exhibit; original Schwab history captures	APLD, STX, HOOD and RDDT cash fields
D2	kb3 .jpg	Dated PLTR and RDDT transaction rows
D3	ikbr fn.jpg	FN date, quantities, times, amounts and commission
D4	kb 1.jpg / duplicate kb 5.jpg	Supplemental SPCX transaction; date cropped
E1	apld.png; apld2.png	APLD displayed book-state sequence
E2	hood4.png; Screenshot 2026-07-30 113847.png	HOOD state and historical pulse display
E3	fn 11.png	FN subsequent book-path context
E4	spcx.png; spcx2.png	Supplemental SPCX setup observations