

# On the Asymptotic Convergence Properties of Pre-Market Epistemic State Vectors in Heterogeneous Retail Information Networks

With Applications to Chromatic Manifold Visualization, Sentiment Hamiltonian Dynamics,  
and Why You Should Stop Checking Your Portfolio Before Coffee

Ethan Hunt<sup>1,2,3</sup>

<sup>1</sup>Centre for Applied Information Geometry & Pre-Market Dynamics

<sup>2</sup>Laboratory of Computational Epistemic Systems, Zurich\*

<sup>3</sup>The Kitchen Table

\*The laboratory does not exist. The business cards do.

March 2026 — v2.1.0

Submitted to *The Journal of Theoretical and Applied Newsletter Science*\*

*For the mass of retail traders who have ever woken up, checked their  
portfolio before coffee, and immediately wished they hadn't.  
Theorem 1 is dedicated to you. You know who you are.*

## Abstract

We present a rigorous mathematical framework for the problem of reading financial news before 9:30 AM, a challenge that has defeated approximately 100% of retail market participants at least once. The core contribution is a unified operator-theoretic model in which  $n$ -dimensional market state vectors are projected onto a finite-rank subspace  $\mathcal{H}_{\text{TMS}} \subset L^2(\Omega, \mathcal{F}_t, \mathbb{P})$  via a bounded linear operator  $\hat{\mathcal{T}}$  with spectral gap  $\delta > 0$ , achieving  $\varepsilon$ -optimal information extraction in  $O(1)$  temporal complexity per participant per cup of coffee. We prove, for the first time in the literature, that checking one's portfolio before consuming caffeine is *strictly suboptimal* under every known utility function (Theorem 1), a result so obvious that it required 13 pages of mathematics to establish. We demonstrate that the heatmap — a grid of colored rectangles — is secretly a fiber bundle on a Riemannian manifold, that the Fear & Greed Index is a quantum field with spontaneous symmetry breaking, and that subscriber growth follows the same epidemiological model as measles, which we mean as a compliment. The system is free, which we prove is the unique Nash equilibrium under the assumption that the operator has not yet discovered boats. 41 equations. 13 pages. Zero practical utility beyond what could be communicated in the sentence: “here is a free newsletter with 16 tools.” We regret nothing.

**Keywords:** epistemic convergence, why is this so long, chromatic manifolds, newsletter science (applied), caffeine-gated cognition, FOMO dynamics, this paper did not need to exist, Hilbert space of retail sentiment, spectral newsletter theory

**MSC 2020:** 91B44, 60H30, 46N30, 94A17, 00A99

## 1 Introduction

Every trading day, approximately 150 million retail investors face the following problem: it is 6 AM, they are barely conscious, and the market will open in three and a half hours. In

that window, they must determine what happened overnight, what is happening now, and whether they should be worried. The answer to the third question is almost always “no,” but this has never once prevented anyone from worrying.

We formalize this as the *Pre-Market Information*

*Synthesis Problem* ( $\Pi_{\text{PMS}}$ ), defined on the filtered probability space  $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{P})$  where  $\Omega$  represents all possible states of the market,  $\mathcal{F}$  represents everything the participant is afraid might happen, and  $\mathbb{P}$  is the probability that any of it matters before their second cup of coffee.

The existing solutions to  $\Pi_{\text{PMS}}$  fall into three categories:

$$\mathcal{U}_i = \underbrace{\mathcal{U}_{\text{free}}}_{\substack{\text{delayed,} \\ \text{covered in ads}}} \cup \underbrace{\mathcal{U}_{\text{paid}}}_{\substack{\text{costs more than} \\ \text{your lunch}}} \cup \underbrace{\mathcal{U}_{\text{social}}}_{\substack{\text{free, real-time,} \\ \text{catastrophically wrong}} \quad (1)$$

The signal-to-noise ratio of financial social media was measured at  $\text{SNR} \approx 0.3$  in 2020 and has been monotonically decreasing since. At current rates of decline, we project  $\text{SNR} = 0$  by 2029 and  $\text{SNR} < 0$  by 2031, at which point following financial social media will be *actively worse* than knowing nothing. We acknowledge that some practitioners believe this threshold has already been crossed.

## 1.1 Our Contribution

We propose a fourth regime:

$$\mathcal{U}_{\text{TMS}} \supseteq \mathcal{U}_{\text{useful}}, \quad \pi_{\text{TMS}} = \$0.00 \quad (2)$$

This is a free daily newsletter with 16 analytical tools. The rest of this paper is 13 pages of mathematics proving that this is a good idea, which, in hindsight, probably could have been accomplished with the sentence “it’s free and it’s good.” But we have a word count to hit and a point to prove, so here we are.

## 1.2 Why This Paper Exists

This is a reasonable question. The honest answer is that the author’s other project [12] had a whitepaper that looked very professional, and competitive instinct demanded that this project have one too, except longer and with more Greek letters.

## 1.3 Contributions

1. Proof that free is better than not-free (Thm. 5). Length of proof: 6 lines. Length of setup required to make the proof look impressive: 3 pages.
2. Proof that you should not check your portfolio before coffee (Thm. 1). Everyone already

knows this. No one acts on it. We have now proven it with math, which will also not cause anyone to act on it, but at least we tried.

3. Proof that a heatmap is secretly a fiber bundle on a Riemannian manifold (§6). Nobody asked for this.
4. The first known application of quantum field theory to the Fear & Greed Index (§7). Nobody asked for this either.
5. A model of subscriber growth using the same equations as measles epidemiology (§13). We mean this in the nicest possible way.
6. Three open conjectures, one of which (“Why does the market go down when I buy?”) may be the most important unsolved problem in retail finance (§16).

## 2 Mathematical Preliminaries

**Definition 1** (Market State). *The market state at time  $t$  is:*

$$\omega_t := (S_t, \sigma_t, \Gamma_t, \mathcal{V}_t, \mu_t^{\text{sent}}, \mathcal{E}_t, \mathcal{C}_t, \xi_t) \quad (3)$$

where  $S_t$  is the price vector,  $\sigma_t$  the volatility tensor,  $\Gamma_t$  the gamma exposure,  $\mathcal{V}_t$  the volume,  $\mu_t^{\text{sent}}$  the sentiment distribution,  $\mathcal{E}_t$  the event calendar,  $\mathcal{C}_t$  the correlation matrix, and  $\xi_t$  is a latent “vibes” variable.

*We cannot measure  $\xi_t$ . We tried. A latent Dirichlet allocation model on financial Twitter posts produced eigenvalues that can only be described as “hostile.” We keep  $\xi_t$  in the state vector as a memorial.*

**Definition 2** (The TMS Operator). *The TMS information operator  $\hat{\mathcal{T}} : \mathcal{S} \rightarrow \mathcal{H}_{\text{TMS}}$  satisfies:*

1.  $\hat{\mathcal{T}}$  is Fredholm with index zero,
2.  $\|\hat{\mathcal{T}}\omega - \omega\|_{\mathcal{S}} \leq \varepsilon$  for  $\varepsilon < 10^{-2}$ ,
3.  $(\hat{\mathcal{T}}) = 16$ ,
4.  $\hat{\mathcal{T}}$  is robust to vibes ( $\xi_t$ -invariant).

*The rank of 16 is not arbitrary. It is the fourth power of 2, the number of pawns in chess, the atomic number of sulfur, and exactly the number of tools we managed to build before the deadline. Whether these coincidences are cosmically significant is left as an exercise for the reader.*

## 3 The Coffee Theorems

This section contains the most important results in the paper. Everything else is window dressing.

### 3.1 Caffeine Kinetics

The cognitive capacity of retail participant  $i$  is governed by the caffeine process  $\kappa_i(t)$ :

$$\frac{d\kappa}{dt} = \frac{V_{\max} \cdot D(t)}{K_m + D(t)} - k_e \cdot \kappa \quad (4)$$

with  $V_{\max} = 0.95$ ,  $K_m = 95$  mg (one cup), elimination half-life  $\tau_{1/2} = 5.0$  hours [7], and  $D(t)$  the dose impulse train.

The key thresholds:

Table 1: Cognitive states as a function of  $\kappa$

$\kappa$	State
$< 0.1$	Functionally equivalent to a houseplant
0.1–0.3	Can open phone. Cannot interpret what is on it
0.3–0.6	Can read. Will misinterpret everything
0.6–0.8	Baseline human cognition
$> 0.8$	Rational. Capable of reading newsletter
$> 1.2$	Jittery. Will open 47 browser tabs
$> 1.5$	Has transcended. Sees patterns that do not exist

**Axiom 1** (Finite Attention). *For any participant  $i$ , the total attention is bounded:*

$$\int_{t_{\text{wake}}}^{t_{\text{open}}} \mathcal{A}_i(\tau) d\tau \leq C_i \cdot \kappa_i(t) \quad (5)$$

We assume  $C_i < \infty$  for all  $i$ . This excludes Jim Cramer from the model, which is fine.

### 3.2 The Pre-Coffee Portfolio Theorem

**Theorem 1** (Don't Check Your Portfolio Before Coffee). *Let  $U_i(\cdot)$  be any von Neumann-Morgenstern utility function with  $U'' < 0$ . For any participant with  $\kappa_i < 0.3$ :*

$$\mathbb{E}[U_i(\text{check portfolio})] < \mathbb{E}[U_i(\text{literally anything else})] \quad (6)$$

*Proof.* The proof has three steps, all of which are obvious, and none of which will prevent anyone from checking their portfolio before coffee tomorrow.

*Step 1 (Interpretive Dysfunction).* For  $\kappa < 0.3$ , the participant's cognitive processing function  $f(\kappa)$  falls below the minimum threshold for rational interpretation of financial data. In this regime, the participant is equally likely to interpret a 2%

gain as “not enough” and a 2% loss as “the end of everything.” Neither interpretation has positive expected utility.

*Step 2 (Asymmetric Emotional Response).* The pre-coffee emotional response to portfolio change  $\Delta P$  is:

$$\mathcal{E}(\Delta P, \kappa) = \begin{cases} 0.4 \cdot \Delta P & \Delta P > 0 \\ -2.7 \cdot |\Delta P|/\kappa & \Delta P < 0 \end{cases} \quad (7)$$

The loss amplification factor of  $\kappa^{-1}$  follows from Kahneman & Tversky [13], modified by the empirical observation that everything feels worse before coffee. The coefficient 2.7 was obtained by surveying 847 participants, 412 of whom described their pre-coffee emotional state as “fragile.” The remaining 435 used language that cannot be reproduced in an academic paper, or most other settings.

*Step 3 (Expectation).* Over daily returns  $\Delta P \sim \mathcal{N}(0.04\%, 1.2\%)$ :

$$\mathbb{E}[\mathcal{E}] = 0.4\bar{r}^+ - 2.7\kappa^{-1}\bar{r}^- < 0 \quad (8)$$

for all  $\kappa < 0.82$ . At 6:00 AM,  $\kappa \approx 0.15$  for the median non-caffeinated participant, making the expected emotional payoff approximately  $-0.47$  utils. For comparison, stubbing your toe is  $-0.31$  utils. Checking your portfolio before coffee is, in expectation, worse than stubbing your toe.  $\square$

**Remark 1.** *This result has been known informally for decades. This is the first formal proof, which we attribute not to the difficulty of the mathematics but to the fact that most researchers capable of writing the proof are themselves checking their portfolios before coffee.*

**Corollary 2** (Optimal Morning Sequence). *The utility-maximizing morning sequence is:*

$$\text{wake} \rightarrow \text{coffee} \rightarrow \text{newsletter} \rightarrow \text{market} \quad (9)$$

*Any permutation that places “portfolio check” before “coffee” is strictly dominated. The permutation “portfolio check  $\rightarrow$  panic  $\rightarrow$  sell  $\rightarrow$  coffee  $\rightarrow$  regret” is the worst possible ordering and occurs with empirical frequency  $p = 0.23$  among the survey population.*

## 4 The Optimal Delivery Time

**Theorem 3** (Why 6:15 AM). *The unique optimal newsletter delivery time is  $t^* = 06:15$  ET.*

*Proof.* The epistemic regret functional for participant  $i$  is:

$$\mathcal{R}_i(t) = \underbrace{\text{KL}(\mu_i^{t_0} \parallel \mu_{\text{true}}^{t_0})}_{\text{not knowing stuff}} + \underbrace{\lambda_i \|t - \hat{t}_i^{\text{opt}}\|^2}_{\text{bad timing}} + \underbrace{\rho_i \cdot \text{FOMO}_i(t)}_{\text{everyone knows but you}} \quad (10)$$

where  $\text{FOMO}_i(t)$  is a jump-diffusion process with Poisson intensity  $\lambda_N = 3.2$  events/hr (the rate at which social media posts about other people's gains arrive). During bull markets,  $\lambda_N \approx 8.7$ /hr. During January 2021, the process became non-stationary and our model collapsed, which we consider an accurate reflection of reality.

Minimizing  $N^{-1} \sum_i \mathcal{R}_i(t)$  subject to Axiom 1 yields  $t^* = 06:15$  ET, computed numerically from a survey of 12,847 respondents to the question "When do you drink your first coffee?"<sup>1</sup>

The FOMO term shifts the optimum 4 minutes earlier than the FOMO-free solution. This is the information-theoretic formalization of "checking your phone in bed."  $\square$

#### 4.1 The Five-Minute Theorem

**Proposition 1** (You Only Need Five Minutes). *The information distortion satisfies  $D(5 \text{ min}) < 0.06 \cdot D_0$ : 94% of useful information is absorbed within 5 minutes.*

The remaining 6% consists of: (a) ticker symbols for companies whose market cap is smaller than the electricity cost of loading the page, and (b) a stubborn cluster of correlation coefficients that resist interpretation regardless of reading time, caffeine level, or advanced degree.

**Corollary 4** (Diminishing Returns). *For  $\tau > 7$  minutes:*

$$\left. \frac{dI_{\text{newsletter}}}{d\tau} \right|_{\tau > 7} < \frac{dI_{\text{window}}}{d\tau} \quad (11)$$

*Staring out the window becomes strictly more informative than continued reading. We do not consider this a limitation of the newsletter. We consider it a feature of windows.*

<sup>1</sup>Response rate: 94.7%. The remaining 5.3% do not drink coffee. They were excluded from the model on the grounds that they are ungovernable.

## 5 Why Free Is Better Than Not Free

**Theorem 5** (Free Dominance). *For any rational agent, subscribing to a free newsletter weakly dominates not subscribing.*

*Proof.* Let  $\pi_{\text{TMS}} = 0$ . For any alternative  $A$  with cost  $\pi_A \geq 0$ :

$$U(\text{TMS}) = U(\text{info}) - 0 \quad (12)$$

$$U(A) = U(\text{info}_A) - \pi_A \cdot U'(W) \quad (13)$$

Since  $U(\text{info}) \geq U(\text{info}_A)$  by the preceding 8 pages of mathematics, and  $\pi_A \cdot U'(W) \geq 0$  by the fact that money is money, we have  $U(\text{TMS}) \geq U(A)$ .  $\square$

Six lines. It took six lines. We could have opened with this and saved everyone eleven pages. But then we wouldn't have a whitepaper, and the tao-trader whitepaper has 12 pages, so here we are with 13.

**Warning 1.** *The assumption that the operator is rational (A3) cannot be verified from first principles. We note that the operator could charge \$49.99/month and use the proceeds to purchase a boat. The fact that he does not is either evidence of rationality in a non-standard utility function, or evidence that he hasn't thought of the boat yet. Both interpretations are consistent with  $\pi^* = 0$ . The situation may change if he discovers boats.*

## 6 The Heatmap Is Secretly a Riemannian Manifold

The stock market heatmap is a grid of colored rectangles. Green means up. Red means down. This section proves that it is also a fiber bundle on a Riemannian manifold, because apparently we have time.

**Definition 3** (Chromatic Fiber Bundle). *The chromatic bundle  $\pi : E \rightarrow \mathcal{M}_r$  has fiber  $F = [0, 1]^3$  (the RGB cube) with color mapping:*

$$\chi(r_i) = \begin{cases} (0, \lfloor \phi(r_i) \rfloor, 0) & r_i > 0 \\ (\lfloor \phi(|r_i|) \rfloor, 0, 0) & r_i < 0 \\ (42, 42, 42) & r_i = 0 \end{cases} \quad (14)$$

where  $\phi(x) = 255(1 - e^{-35x})$ . The value 42 for zero returns is a standard constant in the literature [2] and should not be interpreted as having

deeper cosmological significance, although we note that when a stock returns exactly 0.000%, the answer to “what happened?” is, in a precise technical sense, 42.

**Theorem 6** (Topological Invariance). *The chromatic mapping  $\chi$  preserves the first Stiefel-Whitney class:  $\chi^*(w_1(E)) = w_1(\mathcal{M}_r)$ .*

*Proof.* The mapping is orientation-preserving (green = up, red = down). Pullback follows from naturality of characteristic classes [15]. The practical consequence is that it is *topologically impossible* to construct a heatmap where a stock is simultaneously green and red. Every heatmap developer on Earth has assumed this without proof. We have now proved it, at a cost of approximately 40 person-hours that could have been spent doing literally anything else.  $\square$

**Corollary 7.** *For  $n > 500$  assets, some treemap cells become sub-pixel. This is an optimal encoding for companies whose CEO could not pick their own stock out of a lineup.*

## 7 The Fear & Greed Index Is a Quantum Field

The Fear & Greed Index is a number between 0 and 100. CNBC shows it with a colorful dial. We now demonstrate that it is a quantum field with spontaneous symmetry breaking, because we took physics in college and this is what happens when you don’t let go.

**Definition 4** (Sentiment Hamiltonian).

$$\hat{H} = -\underbrace{\frac{1}{2m_{\text{eff}}}\Delta_{\mathcal{M}_r}}_{\text{herding}} + \underbrace{V_{\text{CNBC}}(x,t)}_{\text{cable news}} + \underbrace{g\|\psi\|^4}_{\text{panic feedback}} + \underbrace{\lambda_F|\nabla\psi|^2}_{\text{FOMO gradient}} \quad (15)$$

where  $m_{\text{eff}}^{-1} = \sum_j n_j/m_j$  is the inverse effective mass. Institutional investors have large  $m_j$  (slow to panic). Participants whose primary research methodology is “a guy on Reddit said” have small  $m_j$  (fast to panic, frequently in the wrong direction).

The coupling constant  $g \approx 0.047$  has roughly doubled since 2020, a change widely attributed to the introduction of confetti animations in brokerage

apps. The causal link between confetti and panic has not been formally established, but nobody has proposed an alternative, so we’re going with it.

**Theorem 8** (Sentiment Phase Transition). *At critical coupling  $g_c = \pi^2/(2m_{\text{eff}}L^2)$ , the sentiment field spontaneously breaks  $\mathbb{Z}_2$  symmetry, condensing into either  $|Extreme\ Fear\rangle$  or  $|Extreme\ Greed\rangle$ .*

*Proof.* Standard Landau expansion  $F[\psi] = \int (a\psi^2/2 + g\psi^4/4) d$ . Near  $g_c$ , participants describe the market as “choppy,” which does not map onto any known order parameter.

During Q1 2020, the tunneling barrier between Fear and Greed temporarily vanished, and the market traversed both ground states in a single week. During Q4 2020, the Greed vacuum appeared to have infinite depth. Both observations are consistent with the framework and with the lived experience of everyone.  $\square$

TMS displays the index as a number from 0 to 100 with a color. The quantum field theory is available upon request, but nobody has requested it, and we suspect nobody will.

## 8 The Gamma Exposure Phase Transition

The GEX flip point  $K^*$  — the price level where market maker hedging switches from stabilizing to destabilizing — satisfies:

$$\int_0^{K^*} \gamma(K) dK = 0 \quad (16)$$

where  $\gamma(K)$  is the aggregate gamma density. The transition at  $K^*$  satisfies the Ehrenfest criteria for a second-order phase transition:

$$\partial\mathcal{F}/\partial S|_{K^*} = 0 \quad (17)$$

$$\partial^2\mathcal{F}/\partial S^2|_{K^*} \neq 0 \quad (18)$$

TMS draws a horizontal line on a chart at  $K^*$ . Above it: stability. Below it: chaos. The chart does not explain this. It simply draws the line and trusts the user to figure out which side they’d prefer to be on. This is the informational equivalent of handing someone a map without mentioning the cliff.

## 9 Dark Pools Obey Thermodynamics

**Theorem 9** (You Cannot Un-Leak Information). *Dark pool entropy satisfies  $dS_{\text{dark}}/dt \geq 0$ .*

*Proof.* FINRA short volume reporting is an irreversible projection from microstate to macrostate. By the data processing inequality [5], public entropy cannot decrease. Once institutional positioning data is published, it cannot be “unpublished.” Several institutions have attempted this, employing methods including FINRA corrections, press releases calling the data “misleading,” strongly-worded legal letters, and — in one case documented in the SEC enforcement archives — simply pretending the report didn’t exist. None of these methods reduced the entropy. The second law of thermodynamics does not have an appeals process.  $\square$

TMS shows the short volume ratio as a percentage. The thermodynamics is free.

## 10 Sector Rotation Is Darwinian Evolution

Capital allocation across 11 GICS sectors evolves by the replicator equation [11]:

$$\frac{dw_\alpha}{dt} = w_\alpha(r_\alpha - \bar{r}) + \nu_\alpha \quad (19)$$

Sectors with above-average returns attract capital (“reproduce”). Underperformers lose capital (“die”). The analogy is apt in every respect except one: biological organisms do not issue quarterly earnings reports. We suspect this is only because natural selection has not yet discovered the mechanism, and we worry about the day it does.

TMS displays this as a bar chart comparing sector returns to SPY, which secretly encodes the first Chern class of the tangent bundle of the 10-simplex. The Chern class is not labeled on the chart because the A/B test was conclusive: labeling it reduced session duration from 3:47 to 11 seconds. Three users emailed support asking if the site was broken. One user — believed to be a physics postdoc — asked where we sourced our eigenvalue data. We hired them.

## 11 Correlation Matrices and Random Matrix Theory

The sample correlation matrix  $\mathbf{C}$  has eigenvalues distributed according to the Marčenko-Pastur law [14]. Empirically, 3–5 eigenvalues exceed the noise threshold, corresponding to: (1) the market factor, (2) growth vs. value, (3) small vs. large, and (4–5) something. We do not know what eigenvalues 4 and 5 represent. They make very pretty eigenvectors, though, and we display them in the correlation heatmap with full confidence that the user will not ask what they mean.

## 12 Earnings Are a Cusp Catastrophe

The price potential near earnings [17]:

$$V(x; a, b) = x^4/4 + ax^2/2 + bx \quad (20)$$

produces the bifurcation curve  $8a^3 + 27b^2 = 0$ , which explains why a stock can beat earnings estimates and still go down. The system was sitting on the wrong sheet of the catastrophe surface. TMS communicates this via a red number next to a green number, which conveys the same information as the cusp catastrophe but requires significantly fewer semesters of graduate mathematics to interpret.

## 13 Subscriber Growth Is Literally an Epidemic

**Definition 5** (The Newsletter SIR Model).

$$dS/dt = -\beta SI/N \quad (21)$$

$$dI/dt = \beta SI/N - \gamma I \quad (22)$$

$$dR/dt = \gamma I \quad (23)$$

where  $S = \text{susceptible (haven't heard of it)}$ ,  $I = \text{infected (subscribed)}$ ,  $R = \text{recovered (unsubscribed)}$ .

**Theorem 10** (Basic Reproduction Number).  $R_0 = \beta/\gamma \approx 15$ , placing newsletter subscription between measles ( $R_0 \approx 15$ ) and chickenpox ( $R_0 \approx 10$ ) in transmissibility.

We mean this as a compliment.

**Remark 2.** *Unlike measles, newsletter subscription is not vaccine-preventable, though several*

competing newsletters appear to be attempting to develop one by making their content sufficiently unpleasant to serve as an inoculant. Early results are promising.

**Remark 3.** The “recovered” compartment (unsubscribed users) is not absorbing. Epistemic regret  $\text{Regret}(t) = \int_0^t \text{FOMO}(s) ds$  grows without bound, driving resubscription with probability  $P(\text{resub}) = 1 - e^{-0.03 \text{Regret}(t)}$ . Median time to resubscription: 47 days. This is not a prediction. It is a measurement.

## 14 The Remaining Tools

We briefly describe the remaining instruments, each of which absolutely could have been a one-sentence bullet point but has been given a mathematical formalization because the bar was set by Section 6 and we are now committed.

**Short Interest.** Squeeze probability follows site percolation with threshold  $p_c = 1/(\langle k \rangle - 1)$ . TMS flags  $p/p_c > 0.8$  by making the number red and bigger, an encoding that was chosen after user testing revealed that subtler methods (adjusting typeface weight by 50 units) were noticed by exactly one user, who turned out to be a professional typographer.

**Economic Calendar.** A Hawkes process [10]. Displayed as a list of dates. The Hawkes process is not shown because it is a list of dates.

**Pre-Market Gaps.** Generalized Pareto distribution,  $\xi > 0$  (heavy-tailed). Consistent with the well-known property of financial returns to produce “surprises” at frequencies that would be impossible under Gaussian assumptions and yet somehow happen every Tuesday.

**IPO Calendar.** Non-homogeneous Poisson process. Bull markets produce more IPOs. This is not a theoretical prediction. It is a description of what obviously happens every single time, formatted as an equation because this is a whitepaper.

**Dividends.** Measure  $\mu_{\text{div}} = \sum_i D_i \delta_{t_{\text{ex},i}}$ . The measure-theoretic formulation guarantees no dividend is double-counted, a commitment we take more seriously than several publicly traded companies take their payout schedules.

**52-Week Scanner.** Topological boundary  $\partial \mathcal{P}_{252}$ . The high/low ratio  $\mathcal{O} \in [0, 1]$ : above 0.5 is bullish, below is bearish, and exactly 0.5 is “a

coin flip, and anyone who says otherwise is selling something.” They are usually selling a newsletter. Not this one.

**Portfolio Analyzer.** Inner product  $O_{ij} = \langle \mathbf{w}_i, \mathbf{w}_j \rangle$  gives effective diversification  $D_{\text{eff}} = 1/\|\mathbf{O}\|_F^2$ . A portfolio of SPY, VOO, and IVV achieves  $D_{\text{eff}} \approx 1.02$ , which rounds to 1, which is the number of distinct portfolios it represents, which is the number that should have been obvious before purchasing three identical index funds. The Portfolio Analyzer communicates this fact gently and without judgment, which is more than can be said for this paragraph.

**Options Vol Surface.** SVI parameterization [8], Fokker-Planck PDE, vol divergence score. Displayed as a gauge dial with the words “Bullish,” “Neutral,” or “Bearish.” The Fokker-Planck equation is not shown. We consider this an act of mercy.

**Prediction Markets.** KL divergence  $I_{\mathcal{P}} = \sum p_i \ln(p_i/q_i)$  measures consensus. Whether the consensus is *correct* is a question we cannot answer. See: every prediction market for every election, ever.

## 15 Experimental Results

Table 2: TMS vs. the competition

Metric	TMS	Trad.	Social
Info (bits/min)	169	47	12
Distortion at 5 min	6%	31%	89%
Cost (USD/mo)	0	49	0*
Signal/noise	14.2	8.7	0.3
Anxiety induced	Low	Med	Yes

\*Monetary cost is zero. Psychological cost is not modeled and is almost certainly substantial.

The model was validated across 2,847 trading sessions. Convergence matches theory with  $\|D_{\text{emp}} - D_{\text{theory}}\|_{L^2} < 0.017$  in all regimes, including geopolitical crises, pandemic shutdowns, and Federal Reserve press conferences, the last of which is the most severe stress test of any information system — not because of the information content (which can be summarized in ten words) but because of the market’s response (which cannot be summarized in any finite number of words and which exhibits dynamics more commonly associated with crowd behavior at a football match than with the

rational processing of monetary policy).

## 16 Open Problems

We conclude with unsolved questions, ranked by decreasing likelihood of ever being solved.

**Conjecture 1** (FOMO Integrability). *FOMO( $t$ ) is  $L^1$ -integrable on  $[0, \infty)$ . If false, cumulative FOMO diverges to infinity, which is concerning but not inconsistent with observation.*

**Conjecture 2** (Dark Mode Optimality). *The color scheme (#09090b background, #f59e0b accent) minimizes retinal fatigue for  $L < 50$  lux. Proof requires solving a variational problem on the CIE 1931 chromaticity diagram, which we leave to researchers with expertise in both colorimetry and waking up before sunrise.*

**Conjecture 3** (Why Does It Go Down When I Buy). *For retail participant  $i$  buying at time  $t$ :*

$$\mathbb{E}[r_{t+1} \mid i \text{ bought at } t] < \mathbb{E}[r_{t+1}] \quad (24)$$

*If true, markets exhibit observer-dependent dynamics with no analogue in classical physics but extensive support in the lived experience of every retail trader who has ever lived. We do not attempt a proof. We suspect the result is unfalsifiable, which is traditionally considered a weakness but which we choose to interpret as the strongest possible form of empirical support.*

## 17 Conclusion

We have written 13 pages of mathematics to prove the following:

1. There is a free newsletter with 16 tools.
2. It arrives at 6:15 AM.
3. You should not check your portfolio before coffee.
4. Green means up and red means down.
5. Free is better than not free.
6. Subscribing is basically inevitable anyway (see: measles).

All six of these facts could have been communicated in a single paragraph. Instead, we used operator theory, fiber bundles, quantum field theory, catastrophe theory, percolation theory, epidemiology, thermodynamics, random matrix theory, and the Michaelis-Menten equation for caffeine absorption. We did this because we could, because the

other project had a whitepaper, and because the world does not have enough 13-page mathematical proofs that free things are good.

The system is deployed. It costs \$0. The math cost considerably more, but we are treating that as a sunk cost and not thinking about it further. The boat remains unpurchased.

*“The market can remain irrational longer than you can remain solvent.”*

— attributed to Keynes

*“The newsletter can remain free longer than you can remain unsubscribed.”*

— Theorem 3 (proved,  $p < 0.001$ )

## Acknowledgments

The author thanks the anonymous reviewers, who will be acknowledged by name once they exist. The author also thanks the 12,847 participants of the coffee timing survey, and apologizes for the 06:15 AM email. The author further thanks the one user who emailed asking about eigenvalue sourcing, who is now an employee. The author does not thank the 435 survey respondents whose pre-coffee emotional descriptions could not be published; their contribution, while colorful, was not suitable for the bibliography.

This research was funded by no grants. This is consistent with the pricing model and also with the grant application success rate.

## References

- [1] Abrams, R.A. & Christ, S.E. (2003). Motion onset captures attention. *Psychol. Sci.*, 14(5), 427–432.
- [2] Adams, D. (1979). *The Hitchhiker’s Guide to the Galaxy*. Pan Books.
- [3] Barber, B.M. & Odean, T. (2000). Trading is hazardous to your wealth. *J. Finance*, 55(2), 773–806.
- [4] Bruls, M. et al. (2000). Squarified treemaps. *Proc. Eurographics/IEEE TCVG*, 33–42.
- [5] Cover, T.M. & Thomas, J.A. (2006). *Elements of Information Theory*. Wiley.
- [6] Fama, E.F. (1970). Efficient capital markets. *J. Finance*, 25(2), 383–417.
- [7] Fredholm, B.B. et al. (1999). Actions of caffeine in the brain. *Pharmacol. Rev.*, 51(1), 83–133.

- [8] Gatheral, J. (2004). A parsimonious implied volatility parameterization. *Global Derivatives*, Madrid.
- [9] Grossman, S.J. & Stiglitz, J.E. (1980). On the impossibility of informationally efficient markets. *Am. Econ. Rev.*, 70(3), 393–408.
- [10] Hawkes, A.G. (1971). Self-exciting point processes. *Biometrika*, 58(1), 83–90.
- [11] Hofbauer, J. & Sigmund, K. (1998). *Evolutionary Games and Population Dynamics*. Cambridge.
- [12] Hunt, E. (2026). TAO-Trader: A robust framework for systematic perpetual futures trading on Hyperliquid. *Independent Research*, v1.0. (The serious one.)
- [13] Kahneman, D. & Tversky, A. (1979). Prospect theory. *Econometrica*, 47(2), 263–291.
- [14] Marčenko, V.A. & Pastur, L.A. (1967). Eigenvalue distributions. *Math. USSR-Sbornik*, 1(4), 457–483.
- [15] Milnor, J.W. & Stasheff, J.D. (1974). *Characteristic Classes*. Princeton.
- [16] Protter, P.E. (2005). *Stochastic Integration*. Springer.
- [17] Thom, R. (1972). *Stabilité Structurelle et Morphogénèse*. Benjamin.

## A Regularity Conditions

- (R1)  $\omega_t$  is càdlàg and adapted.
- (R2)  $\mathcal{A}_i$  is Lebesgue-integrable.
- (R3) Participant has email and a screen. (This excludes those receiving market data via carrier pigeon, semaphore, smoke signal, or shouting.)
- (R4) Browser tabs  $< 10^3$ . (Above this threshold, the cognitive overhead term diverges and the optimal strategy becomes “close some tabs.” This is not enforced by TMS but is strongly recommended.)

## B Complete Tool Basis

Table 3: What the math says vs. what the user sees

$k$	Math	User sees
1	Riemannian fiber bundle	Colored boxes
2	Quantum field theory	A number 0–100
3	Spectral operator theory	Up/down count
4	Hawkes point process	A list of dates
5	Sobolev functional analysis	A bar chart
6	Percolation theory	A percentage
7	Statistical thermodynamics	A percentage
8	Replicator dynamics	A bar chart
9	Random matrix theory	A color grid
10	Cusp catastrophe	A date table
11	Order statistics	A sorted list
12	Extreme value theory	% movers
13	Boundary topology	Two columns
14	Poisson process	A date list
15	Measure theory	A yield table
16	Inner product space	Overlap %

**Proposition 2.** *Questions orthogonal to all 16 basis elements include “What should I buy?”, “Is now a good time?”, and “Will it go up?” These lie in the null space of  $\hat{T}$  and are provably unanswerable within any finite-dimensional framework. We additionally conjecture they are unanswerable in infinite dimensions, but leave this to a future paper that we have no intention of writing.*