

The Wisdom of Clocks

Crowds calibrate to the percentage point wherever a bell closes the book; the market that never settles overpays half its listings fivefold — and the exchanges are unbolting the bell.



RIKO KARDAMOW

AUG 10, 2026

It is seven o'clock on a Monday evening in Dublin and I am finalising this piece with the tabs still open — Kalshi's Fed markets repricing the September decision after Friday's ugly July jobs print, and an AP wire from last week reporting that a federal judge has let Utah enforce its anti-gambling laws against the very same exchange, a ruling that landed on 5 August, the same day Morgan Stanley published the report this piece is about.

Mauboussin and Callahan's *The Wisdom of Crowds in Markets* ate my weekend — a mug of tea went cold twice on Saturday morning — and I want to be honest about why: somewhere around page twenty-three it hands you, almost in passing, the thing I have been circling since I wrote in February that prediction markets are truth compressors.

What's Truth Worth?

RIKO KARDAMOW • FEB 25



"Not everything that counts can be counted, and not everything that can be counted counts." — attributed to William Bruce Cameron

[Read full story](#)

The report's framework is Surowiecki's, tidied for practitioners: crowds are wise when three conditions hold — cognitive diversity, an aggregation mechanism, and incentives — and mad when one of them breaks.


My reading, after a weekend of sitting with it, is that the report proves a fourth condition on every page while never listing it in the trio, and the fourth condition is not a property of the crowd at all. It is a clock — a scheduled moment at which the market's opinion stops mattering and the world takes over, the bell the racecourse calls post time — and everything the report measures divides cleanly into markets that have one and the single market that does not.

The fourth condition is a clock

The founding experiment already contained it.

Galton's 1906 Plymouth crowd — seven hundred and eighty-seven legible tickets at sixpence each — guessed the dressed weight of an ox at an average of 1,197 pounds against an actual weight of 1,197, and the guesses did not even distribute normally around the truth they collectively found

Exhibit 1: The Guesses of the Ox's Weight Do Not Follow a Normal Distribution



Source: Kenneth F. Wallis, "Revisiting Francis Galton's Forecasting Competition," *Statistical Science*, Vol. 29, No. 3, August 2014, 420-424.

Run through Page's diversity prediction theorem, the average individual error of 5,408.2 and the prediction diversity of 5,408.1 net to a collective error of 0.1 — the crowd beats essentially everyone in it.


But notice what the ox supplies that no condition in the trio names: a carcass, weighed, on a date. The competition worked because it ended — and that ending, not the butchers in the crowd, is what the rest of the report is secretly about.

The race runs whatever the odds say

Every market in the report that owns a post time is calibrated to a degree that should embarrass the people who still call betting irrational.

Across more than seventy-two million trades on Kalshi, the contract price and the realised win rate lie almost exactly on the forty-five-degree line — a contract at seventy cents wins about seventy per cent of the time


Exhibit 2: Actual Win Rate versus Contract Price on Kalshi



Source: Counterpoint Global based on Jonathan Becker, "The Microstructure of Wealth Transfer in Prediction Markets," Independent Research, January 18, 2026.

The Iowa Electronic Markets, across the 1988–2000 presidential cycles, predicted vote shares a week out with an average absolute error of 1.5 percentage points against the final Gallup poll's 2.1, the error grinding down as the clock runs toward election day


Exhibit 3: Information Revelation Through Time on Iowa Electronic Markets



Source: Counterpoint Global based on Iowa Electronic Markets and Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz, "The Promise of Prediction Markets," *Science*, Vol. 320, No. 5878, May 16, 2008, 877-878.
Note: Dotted blue line is a fitted trend line.

Roughly four hundred thousand football matches across six thousand leagues produce the same picture, and thirty thousand horses across two and a half thousand races produce it again on the tote.


Exhibit 8: Probability Implied by Odds and Win Rate for Horseracing



Source: Counterpoint Global based on Constancedongg, "Horse-Racing," GitHub, <https://github.com/constancedongg/Horse-Racing>, Last modified July 6, 2019.

Four mechanisms, one geometry — and the geometry is only measurable because in each case the game was played, the race was run, the ballot was counted.

Exhibit 7: Probability Implied by Odds and Win Rate for Soccer Matches



Source: Counterpoint Global based on Cuamckuu, "Domain-driven-soccer-probs," GitHub, <https://github.com/cuamckuu/domain-driven-soccer-probs>, Last modified November 1, 2025 and Artur Karimov, Aleksandr Koshkin, Dmitrii Kaplun, and Denis Butusov, "Domain-Driven Identification of Football Probabilities," Mathematics, Vol. 13, No. 24, Article 3976, 2025.

The deviation that does exist proves the point rather than breaking it.

The favourite-longshot bias — favourites paying slightly better than their odds imply, longshots slightly worse — shows up on Kalshi, in fixed-odds football, on the racetrack, and Snowberg and Wolfers' work attributes it to misperceived probabilities rather than a taste for lottery tickets. It persists for a mechanical reason the report is precise about: the profit in the bias is smaller than the toll for collecting it, a vigorish of around four to five per cent on straight sports bets and a parimutuel takeout of fifteen to twenty-four per cent depending on the pool.

Work the report's own tote example — a ten-thousand-dollar pool, four thousand on Horse A, a twenty per cent takeout — and a horse with a forty per cent chance of winning pays out one dollar sixty against a

two-dollar ticket; the bettor needs the true probability above fifty per cent before the wager breaks even

Exhibit 9: Simple Example of Implied Probability of Winning in a Parimutuel System

Horse	Bet	Winning Chance
A	\$4,000	40.0%
B	3,000	30.0
C	2,000	20.0
D	750	7.5
E	250	2.5

Source: Counterpoint Global.

A mispricing that cannot fund its own correction is not an inefficiency in any sense a desk can use; it is the entry fee for the answer key, and the crowd that pays it is buying, whether it knows it or not, the single thing the stock market cannot sell — a marked exam.

*The deeper discipline is that in these markets the outcome is causally independent of the price. The teams play the game whatever the line says; **the horse does not read the tote board.***


A political market can be manipulated with the intent of swaying voters — the one channel where the posted probability reaches back toward the outcome — and the report shows even that channel failing against the clock, which is where the story turns from calibration to something closer to justice.

The bell disciplines even the cheats

The manipulation record reads like a sequence of expensive lessons in the same rule.

The speculative attack on the TradeSports contract for George W. Bush on 13 September 2004 knocked the price down about thirteen points on heavy volume — and the price recovered within hours.

Exhibit 4: Temporary Impact of Attempted Price Manipulation of a TradeSports Contract




Source: Counterpoint Global based on Paul W. Rhode and Koleman Strumpf, "Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data," Working Paper, February 2006.
Note: Greenwich Mean Time.

Rothschild and Sethi's forensic work on Intrade in 2012 found a single trader carrying a third of all the money spent on Romney in the final fortnight, apparently defending a floor under one contract and a ceiling over another — the positions were eventually pulled, the Intrade-Betfair gap snapped shut, and the trader lost an estimated seven million dollars.

Camerer ran the experiment properly on the racetrack, placing and cancelling large bets to see whether he could lure the pool into a false read, and found the odds moved and then reverted before post time, the informed money cancelling the noise.

Exhibit 10: Temporary Impact of Manipulation of a Market for Horse Racing



Source: Counterpoint Global based on Colin F. Camerer, "Can Asset Markets Be Manipulated? A Field Experiment With Racetrack Betting," *Journal of Political Economy*, Vol. 106, No. 3, June 1998, 457-481.
Note: Post time is when the race begins and betting is closed.

A market with a settlement date converts every manufactured belief into a bill, presented on schedule, to the account of whoever manufactured it.

Even the crowd's honest enthusiasms get invoiced.

Twenty-four horses entered the Belmont positioned to win the Triple Crown between 1950 and this year, carrying an average implied win probability of sixty-one per cent; twenty-one per cent of them won.

Exhibit 11: Implied Win Probability and Success Rate for Triple Crown Contenders, 1950-2026


Year	Horse	Implied Probability	Won Belmont / Triple Crown?	Odds (Fractional)
1958	Tim Tam	87.0%	No	3-20
1961	Carry Back	71.4%	No	2-5
1964	Northern Dancer	55.6%	No	4-5
1966	Kauai King	62.5%	No	3-5
1968	Forward Pass	50.0%	No	1-1
1969	Majestic Prince	43.5%	No	13-10
1971	Canonero II	58.8%	No	7-10
1973	Secretariat	90.9%	Yes	1-10
1977	Seattle Slew	71.4%	Yes	2-5
1978	Affirmed	62.5%	Yes	3-5
1979	Spectacular Bid	83.3%	No	1-5
1981	Pleasant Colony	55.6%	No	4-5
1987	Alysheba	55.6%	No	4-5
1989	Sunday Silence	55.6%	No	4-5
1997	Silver Charm	50.0%	No	1-1
1998	Real Quiet	55.6%	No	4-5
1999	Charismatic	40.0%	No	3-2
2002	War Emblem	44.4%	No	5-4
2003	Funny Cide	50.0%	No	1-1
2004	Smarty Jones	74.1%	No	7-20
2008	Big Brown	80.0%	No	1-4
2014	California Chrome	54.1%	No	17-20
2015	American Pharoah	57.1%	Yes	3-4
2018	Justify	55.6%	Yes	4-5
Average Success rate		61.0%	20.8%	

Source: Counterpoint Global and Wikipedia.

Crist, the handicapper who made a quarter of a year's salary betting against Smarty Jones in 2004, called the phenomenon plainly — “the best friend that horseplayers have are the big race days” — because a wave of sentimental money is only exploitable in a market where the race settles the argument.

And the settlement ledger shows exactly who ends up holding it: across nearly six hundred million Polymarket trades from late 2022 through this March, the top one per cent of winners captured seventy-seven per cent of the profits, the top ten per cent captured ninety-six, and sixty-nine per cent of 2.4 million users lost money outright.


Exhibit 5: Concentration of Gains on Polymarket



Source: Counterpoint Global based on Pat Akey, Vincent Grégoire, Nicolas Harvie, and Charles Martineau, "Who Wins and Who Loses In Prediction Markets? Evidence from Polymarket," Working Paper, June 12, 2026.

On Kalshi, the same transfer is visible at the level of market structure — makers earn an average excess return of positive 112 basis points, and takers pay the mirror-image negative 112.

Exhibit 6: Average Excess Returns of Makers versus Takers on Kalshi



Source: Counterpoint Global based on Jonathan Becker, "The Microstructure of Wealth Transfer in Prediction Markets," Independent Research, January 18, 2026.


The bell does not make the crowd rich; it makes the crowd honest, by settling every error into someone's account before the story can be revised.

The paper nobody marks

Now cross to the one market in the report with no post time, and watch every measurement instrument fail at the border.

The stock market is the only positive-sum venue in the set — ninety-one trillion dollars of wealth created over the century to 2025 against the 5.8 trillion, in 2025 dollars, that American gamblers have lost since 1929 — and it is simultaneously the only one whose accuracy cannot be scored, because a share's fair value is never revealed short of takeover or bankruptcy.

Exhibit 12: Annual Losses from Gambling in the U.S., 1929 to 2025



Source: Counterpoint Global; U.S. Bureau of Economic Analysis, Personal consumption expenditures: Services: Gambling [DGAMRC1A027NBEA], retrieved from FRED, Federal Reserve Bank of St. Louis; U.S. Bureau of Economic Analysis, Gross domestic product (implicit price deflator) [A191RD3A086NBEA], retrieved from FRED, Federal Reserve Bank of St. Louis.

The report is candid that testing efficiency here collapses into the joint hypothesis problem: any test of the market is also a test of your pricing model, so a failed test convicts neither.

And when academics built a workaround — thirteen thousand eight hundred US listings from 1975 to 2019, the present value of each firm's remaining lifetime earnings measured against its first traded price — the answer split down the middle of the question: the average ratio of lifetime earnings to price was 0.9, price within ten per cent of value, a

portrait of efficiency; the *median* was 0.2, half of all firms delivering a fifth or less of what their debut price implied, a portrait of systematic overpayment rescued in aggregate by a handful of enormous winners.

This is where Tony Lawson's 2013 critique of the reflexivity framework earns its place in a markets piece thirteen years after publication.

Lawson's correction — aimed at the theory's own author — was that the load-bearing distinction is not between natural and social phenomena, or science and non-science, but between closed systems, where event regularities hold and prediction is possible, and open ones, where they do not; the sin of modern economics is that it "treats open systems as if they are actually closed".

Exhibit 13: Regulatory Topics for Prediction, Sports Betting, Parimutuel, and Stock Markets

Market	Main U.S. Regulators	What They Regulate	Main Overlap
Prediction Markets / Event Contracts	Commodity Futures Trading Commission (CFTC) for event contracts listed on CFTC-regulated designated contract markets (DCMs). National Futures Association may apply to derivatives intermediaries. Kalshi and Polymarket US are currently DCMs as designated by the CFTC.	Event contracts are derivatives tied to a specified event, occurrence, or value. CFTC rules and the Commodity Exchange Act give the CFTC authority to review and prohibit event contracts tied to activities such as terrorism, assassination, war, gaming, or conduct unlawful under federal or state law.	Largest overlap is with sports betting, including game outcomes, proposition bets, spreads, championships, or combinations that are similar to parlays. Platforms and the CFTC have argued that federal oversight preempts state jurisdiction. Several states are challenging this interpretation.
Sports Betting	Mostly state gaming commissions, lotteries, casino control boards, and state attorneys general. Tribal gaming involves tribal regulators and the National Indian Gaming Commission. Agreements between the state and tribes are reviewed by the Bureau of Indian Affairs, which operates under the U.S. Department of the Interior.	Licensing of sportsbooks, permitted bet types and wagering menus, geolocation, age limits, advertising, taxes, responsible-gambling rules, integrity monitoring, and illegal-operator enforcement. Federal overlays include payment restrictions under the Unlawful Internet Gambling Enforcement Act and enforcement for federal gambling and criminal law.	Overlaps with prediction markets when an event contract regulated by the CFTC is functionally a sports wager. Also overlaps with parimutuel regulation where the same state agency regulates sports wagering and horse racing.
Parimutuel Betting	Primarily state racing or gaming commissions. For Thoroughbred racing, Horseracing Integrity and Safety Authority operates under Federal Trade Commission oversight for racetrack safety, anti-doping, and medication rules. States may retain parimutuel and other racing authority.	Pooled wagering: bettors wager with each other through a pool, not against a bookmaker. The Interstate Horseracing Act defines parimutuel wagering, interstate off-track wagers, consent requirements, and takeout rules for interstate off-track wagering.	Overlaps with sports betting at the state-agency level. Structurally different from sportsbook odds because the operator extracts a pool takeout rather than act as bookmaker. Less direct CFTC overlap, except where platforms emphasize peer-to-peer market structure.
Stock Market / Securities Markets	The Securities and Exchange Commission (SEC) is the primary federal securities regulator. The Financial Industry Regulatory Authority regulates broker-dealers under SEC oversight. National securities exchanges and clearing agencies are SEC-regulated Self-Regulatory Organizations. State securities regulators also enforce state securities laws.	Issuers, public-company disclosures, broker-dealers, exchanges, securities trading, clearing, market manipulation, insider trading, investor protection, and security-based swap markets.	Overlaps with the CFTC for hybrid derivatives: SEC regulates security-based swaps, CFTC regulates swaps, and both jointly regulate mixed swaps; security futures are also jointly regulated.

Source: Counterpoint Global.

Read the report's two great summary tables through that lens and they stop being comparisons and become a single boundary drawn twice: Exhibit 13 shows four regulators policing products that increasingly overlap, and Exhibit 14 shows the deeper split those regulators are groping toward — three markets whose diversity breakdowns are “limited” or “do not persist” because outcomes arrive on schedule, and

one whose correlated beliefs “can dominate” because nothing ever arrives at all.

Betting markets are the one place in economic life where the closed-system assumption is approximately true, because the institution manufactures the closure — a defined event, a binary payoff, a forced settlement — and calibration charts are simply what a closed system looks like when you photograph it.

I should admit the personal stake in that reading.

I spent years of my working life on the settlement side of fixed income — the unglamorous end, where a trade is not an opinion but an obligation with a value date, and where every mismatch between what the front office believed and what actually existed surfaced in the fails queue within a day.

Exhibit 14: Comparison of Prediction, Sports Betting, Parimutuel, and Stock Markets

Market	Accuracy	Aggregation Mechanism	Role of Intermediary	Diversity Breakdowns	Cost to Participate	Where the Wisdom Comes From
Prediction Market	Very well calibrated. Observed frequency of outcomes close to probabilities implied by prices. Accuracy can be compromised by thin markets.	Continuous double auction. For binary contracts, price approximately reflects probability before costs.	Platforms match orders, handle settlement, and charge fees. They generally do not take directional risk.	Partisan participants can create short-term movement.	Some charge fees for trades or withdrawals. Customers may not receive interest on cash balances.	Broad, diverse participation with skin in the game; arbitrageurs pull price toward truth.
Sports Betting Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot, home-team, and recency biases, but they vary by sport.	Bookmaker line with small and infrequent revisions throughout the week. Odds approximately reflect probability before costs.	Bookmaker takes the other side, manages risk, and may balance exposure or shift prices to exploit bettor biases.	Limited because bookmakers set and adjust the odds and outcomes are independent of implied probabilities.	Vigorish varies by sport, market, and type of wager, but it is around 4-5% for straight bets and higher for parlays.	Bookmakers who take into consideration bets by smart money.
Parimutuel Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot bias.	Gross bet distribution determines implied probabilities. Payout odds reflect the takeout.	None. Track is an operator and not a counterparty. Payouts are after a takeout.	Manipulations move odds in the short term but do not persist. Periodically, large money comes in on a favorite such as a Triple Crown candidate.	Takeout rate 15-24% depending on pool. Low end for win, place, and show pools and higher for exotic pools.	Aggregation of all bets, with smart money tending to come in late if advantageous.
Stock Market	Efficiency difficult to quantify. Hard to generate excess returns, but markets veer to booms and busts from time to time.	Continuous double auction. Most liquid of markets.	Market makers, specialists, and high-frequency traders. Most liquidity providers seek to maintain no net exposure.	Correlated beliefs, narratives can dominate, short-term investors fail to discount future cash flows, and diversity breakdown effects are often non-linear.	Cost for institutions is commission and market impact. Often zero explicit commissions for retail investors. Relatively low.	Diverse investors (time horizon, sophistication, decision rules). Payoffs for producing costly information. Arbitrageurs push price toward fundamentals.

Source: Counterpoint Global.

The wisdom-of-crowds literature spent twenty years asking what makes a crowd smart, and the honest answer sitting in that table is that we can only *know* a crowd is smart where its claims come due — everywhere else, wisdom is an unmarked exam and the grade is a rumour.

The loop completes only after the bell fails

The reflexive mechanism that runs wild in the unsettled market is buildable in four steps from the report's own evidence, and it is worth building explicitly, because the missing bell is load-bearing at exactly one of them.

- Step one, a belief: the price is the probability, the market knows.
- Step two, the belief shapes behaviour: investors conform — the CFA Institute's survey of seven hundred investors found peer influence the single most damaging bias to their decisions — and short-horizon traders extrapolate the price rather than discounting the dividend, precisely what Hirota and Sunder produced in the laboratory, where long-horizon sessions converged on fundamental value and short-horizon sessions, paid on future prices rather than the terminal payout, bubbled.
- Step three, the behaviour alters the fundamental: in a perpetual market the price *is* an input to the thing it prices — it sets the firm's cost of capital, feeds the narratives that Shiller models as epidemics and Mangel shows investors invent to metabolise novelty, and censors the bears, since the report notes a wide dispersion of analyst views fails to reach the price when pessimists cannot or will not participate.
- Step four, the altered fundamental feeds the belief back — confirming it, valuation by valuation, until the correlated position is large enough that a small further loss of diversity moves the price non-linearly. In sports the loop dies at step three because the odds cannot make the horse faster; in equities step three is the business model, and the inflection — the moment self-reinforcing turns self-defeating — arrives unscheduled, which is why it arrives as a crash.

The strongest counter deserves its full weight: calibration is not wisdom.

A sportsbook line is set by a bookmaker, not discovered by a crowd — Levitt showed the books deliberately shade spreads into bettor biases — so the tidy forty-five-degree lines may certify the oddsmaker rather than the mob; superforecasters working in teams beat prediction markets outright, and sophisticated aggregations of ordinary polls match them, so the market mechanism is not even uniquely good at its own game.

All granted — and all of it, I think, sharpens the thesis rather than blunting it.

The claim was never that a bell makes crowds wise; the claim is that a bell makes wisdom checkable, and checkability is the entire difference between a statistic and a story.

Lawson's point lands here with full force: prediction in an open system is impossible except by chance, so the stock market's wisdom is not merely undemonstrated but undemonstrable from the inside — the joint hypothesis problem is what an open system says when you ask it for a grade.

Unbolting the clock

Which makes the regulatory news of this summer read like a controlled experiment someone decided to run on the thesis.

On 29 May the CFTC — down to a single confirmed commissioner — approved Kalshi's application to list a perpetual futures contract referencing spot bitcoin, in one day, and issued a policy statement inviting every designated contract market to self-certify similar contracts; on 18 June, CME sued the agency and its chairman in federal

court in Washington, arguing the contracts are swaps under Dodd-Frank and the approval process was unlawful.

Strip the litigation of its tax and turf dimensions and what remains is precisely the fourth condition changing hands: the prediction-market industry, which earned its calibration record — the entire forty-five-degree case for its own existence — from contracts that settle at a predetermined time, is importing the one instrument defined by never settling.

A perpetual contract has no post time; it is marked against a funding rate forever, which is to say it is the stock market's epistemology sold on a betting exchange.

The same platforms are simultaneously listing contracts on corporate key performance indicators, and the incentive to reach into the outcome is already operational — a weather station near Charles de Gaulle apparently tampered with to win a Polymarket temperature bet in April, half a million Spotify streams of one artist scrubbed in July after suspicious Kalshi wagers on the month's most-played song. The bell disciplined the cheats; the cheats have begun working on the bell.

So, now that the whole ledger is on the table, here is what it amounts to: the wisdom of crowds was never a property of crowds. It was a property of settlement — of the institutional decision to make beliefs come due — and the three conditions of the literature are merely what determines *how much* wisdom the settlement date can extract.

My verdict is that the report, the best single document yet written on collective intelligence in markets, still files its most important finding under table formatting: the column that separates Exhibit 14's calibrated markets from its bubbling one is a clock, and the industry with the best clock in finance is at this moment dismantling it for volume.

Exhibit 14: Comparison of Prediction, Sports Betting, Parimutuel, and Stock Markets

Market	Accuracy	Aggregation Mechanism	Role of Intermediary	Diversity Breakdowns	Cost to Participate	Where the Wisdom Comes From
Prediction Market	Very well calibrated. Observed frequency of outcomes close to probabilities implied by prices. Accuracy can be compromised by thin markets.	Continuous double auction. For binary contracts, price approximately reflects probability before costs.	Platforms match orders, handle settlement, and charge fees. They generally do not take directional risk.	Partisan participants can create short-term movement.	Some charge fees for trades or withdrawals. Customers may not receive interest on cash balances.	Broad, diverse participation with skin in the game; arbitrageurs pull price toward truth.
Sports Betting Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot, home-team, and recency biases, but they vary by sport.	Bookmaker line with small and infrequent revisions throughout the week. Odds approximately reflect probability before costs.	Bookmaker takes the other side, manages risk, and may balance exposure or shift prices to exploit bettor biases.	Limited because bookmakers set and adjust the odds and outcomes are independent of implied probabilities.	Vigorish varies by sport, market, and type of wager, but it is around 4-5% for straight bets and higher for parlays.	Bookmakers who take into consideration bets by smart money.
Parimutuel Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot bias.	Gross bet distribution determines implied probabilities. Payout odds reflect the takeout.	None. Track is an operator and not a counterparty. Payouts are after a takeout.	Manipulations move odds in the short term but do not persist. Periodically, large money comes in on a favorite such as a Triple Crown candidate.	Takeout rate 15-24% depending on pool. Low end for win, place, and show pools and higher for exotic pools.	Aggregation of all bets, with smart money tending to come in late if advantageous.
Stock Market	Efficiency difficult to quantify. Hard to generate excess returns, but markets veer to booms and busts from time to time.	Continuous double auction. Most liquid of markets.	Market makers, specialists, and high-frequency traders. Most liquidity providers seek to maintain no net exposure.	Correlated beliefs, narratives can dominate, short-term investors fail to discount future cash flows, and diversity breakdown effects are often non-linear.	Cost for institutions is commission and market impact. Often zero explicit commissions for retail investors. Relatively low.	Diverse investors (time horizon, sophistication, decision rules). Payoffs for producing costly information. Arbitrageurs push price toward fundamentals.

Source: Counterpoint Global.

The payout runs through contract form, and the molecule is where it lands.

Physical delivery is the energy market's post time — the expiring benchmark that must converge on a cargo, a pipeline nomination, a half-hourly balancing settlement — and it is the reason commodity

curves stay tethered to storage economics while equity-adjacent energy paper can run on narrative for years.

As perpetual structures migrate from crypto toward event and commodity exposure — Hyperliquid's oil perps were the early tell — the desk-level move is that the basis between perpetual paper and the expiring physical contract becomes a live gauge of detachment: the funding rate on a perp is what belief costs per day when nothing forces it to be tested, and a widening, persistent basis against the physical front month is the market quietly telling you which of its two clocks it trusts.

The structural implication is a premium migrating to physically-settled benchmarks — the contract form that gains share in a perpetualising complex is the one that still delivers, because it is the only one whose price retains an audit.

And if you take nothing else to the desk tomorrow morning, take three rules the report's own numbers hand you.

- In a calibrated market, the only durable edge is structural, not informational. Kalshi's makers earn positive 112 basis points and its takers pay negative 112 — a 224-basis-point round trip that is, in a market already sitting on the forty-five-degree line, essentially the whole distribution of returns — so if you are using event contracts the way that New York bar owner used the Knicks contract, as a hedge, rest your orders and be the maker, because the seventy-seven per cent of profits captured by the top one per cent is not mostly clairvoyance; it is the patient side of the book collecting the toll from the impatient side, sixty-nine per cent of whom lose.
- Fade sentiment only where a settlement date forces convergence. Crist's big-race-day rule works — the sixty-one per cent implied against twenty-one per cent realised on Triple Crown contenders is genuinely collectable — because the Belmont runs on schedule and the pool must pay; the equity analogue of the sentimental

favourite, the hot new listing whose median delivers a fifth of its debut price, is *not* collectable on the same terms, because nothing forces convergence before your financing, your patience, or your risk committee expires. In settled markets, size for being right; in unsettled ones, size for how long you can be right without being paid.

- Use the checkable market to mark the uncheckable one.

The report notes prediction markets now list contracts on corporate key performance indicators, and the Fed's own researchers have started treating Kalshi's macro contracts as benchmarks — so when a calibrated, settling probability diverges from what an equity price or an OIS curve implies for the same event, that divergence is information with an audit trail on one side of it, and it is currently free.


One cost footnote that is worth actual money before anyone builds a systematic strategy on any of this: the wagering-loss provision that took effect for tax years after December 2025 caps deductible losses at ninety per cent of losses, so a bettor who wins ten thousand dollars and loses ten thousand owes tax on a thousand dollars of income that never existed — and the report is blunt that the treatment of prediction-market gains and losses is not yet clearly specified at all.

Stack the tolls honestly — four to five per cent vigorish on straight sports bets, fifteen to twenty-four on the tote, the 224-basis-point taker round trip on event contracts, tax on phantom gains for high-turnover books — and the price list of the answer key is steep enough that most strategies which look beatable pre-cost are the favourite-longshot bias in a different coat: real, measurable, and smaller than the fee for collecting it.

So here is what I will actually be watching, and I would encourage you to watch it with me.

- The first gauge sits in a Washington courtroom: if the judge voids the 29 May approval, the clock gets bolted back on and this piece becomes a warning shot that landed early; if the approval stands, every exchange in the country gets an invitation to list contracts that never settle — and the lovely forty-five-degree line of Exhibit 2, the whole calibration case for prediction markets, becomes something the new contracts can never earn, because you cannot plot a market against outcomes it has promised never to produce.

Exhibit 2: Actual Win Rate versus Contract Price on Kalshi



Source: Counterpoint Global based on Jonathan Becker, "The Microstructure of Wealth Transfer in Prediction Markets," Independent Research, January 18, 2026.

- The second gauge is Kalshi's own bitcoin perp, and I will be honest that this is the one I find genuinely fascinating — watch its funding basis to spot through the first real volatility event, because if the basis holds tight, the exchanges will have pulled off something clever, borrowing settlement discipline from a reference price they do not control, and if it gaps and stays gapped, well, Exhibit 14 told you first.

Exhibit 14: Comparison of Prediction, Sports Betting, Parimutuel, and Stock Markets

Market	Accuracy	Aggregation Mechanism	Role of Intermediary	Diversity Breakdowns	Cost to Participate	Where the Wisdom Comes From
Prediction Market	Very well calibrated. Observed frequency of outcomes close to probabilities implied by prices. Accuracy can be compromised by thin markets.	Continuous double auction. For binary contracts, price approximately reflects probability before costs.	Platforms match orders, handle settlement, and charge fees. They generally do not take directional risk.	Partisan participants can create short-term movement.	Some charge fees for trades or withdrawals. Customers may not receive interest on cash balances.	Broad, diverse participation with skin in the game; arbitrageurs pull price toward truth.
Sports Betting Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot, home-team, and recency biases, but they vary by sport.	Bookmaker line with small and infrequent revisions throughout the week. Odds approximately reflect probability before costs.	Bookmaker takes the other side, manages risk, and may balance exposure or shift prices to exploit bettor biases.	Limited because bookmakers set and adjust the odds and outcomes are independent of implied probabilities.	Vigorish varies by sport, market, and type of wager, but it is around 4-5% for straight bets and higher for parlays.	Bookmakers who take into consideration bets by smart money.
Parimutuel Market	Very well calibrated, with high correlation between probability implied by odds and win rate. Evidence of favorite-longshot bias.	Gross bet distribution determines implied probabilities. Payout odds reflect the takeout.	None. Track is an operator and not a counterparty. Payouts are after a takeout.	Manipulations move odds in the short term but do not persist. Periodically, large money comes in on a favorite such as a Triple Crown candidate.	Takeout rate 15-24% depending on pool. Low end for win, place, and show pools and higher for exotic pools.	Aggregation of all bets, with smart money tending to come in late if advantageous.
Stock Market	Efficiency difficult to quantify. Hard to generate excess returns, but markets veer to booms and busts from time to time.	Continuous double auction. Most liquid of markets.	Market makers, specialists, and high-frequency traders. Most liquidity providers seek to maintain no net exposure.	Correlated beliefs, narratives can dominate, short-term investors fail to discount future cash flows, and diversity breakdown effects are often non-linear.	Cost for institutions is commission and market impact. Often zero explicit commissions for retail investors. Relatively low.	Diverse investors (time horizon, sophistication, decision rules). Payoffs for producing costly information. Arbitrageurs push price toward fundamentals.

Source: Counterpoint Global.

What gets me, sitting with all of this at seven on a Monday evening in Dublin, is that the industry doing the unbolting is the one industry that should know better — it spent two decades proving crowds are wise exactly as often as their bets come due, and it is now dismantling the machinery of coming due, for volume.

Post time was never the constraint on the wisdom.

It was the wisdom.



✓ Subscribed