PREDICTION MARKETS · THE NEXT PHASE

PROSPECTORS

From Competing for Trading Access to Competing for the Power to Define Outcomes

Traders of rule discrepancies earn a spread; those who define outcomes share in the future.

TWO PROBABILITIES · ONE CONTRACT

THE GAP IS THE ALPHA →

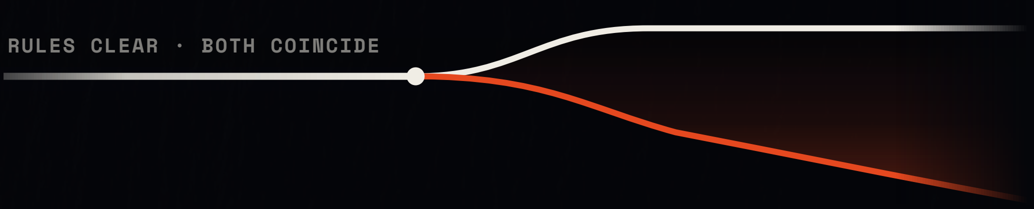

PROBABILITY THE EVENT OCCURS
100%

PROBABILITY THE RULES RECOGNIZE IT
0%

STRATEGY · 2026-05 · FACT UNDISPUTED · $972M RESOLVED NO

Prediction Markets · Series No. 03
Analysts Merida Lee & Kiko · August 2026

Question everything.

AT A GLANCE

TL;DR

Finding the Next Wave of Industry Alpha, from Settlement Friction to Cross-Platform Infrastructure

- **The prediction-market front end is getting crowded. The next wave of growth may come from back-end infrastructure.** Platform competition has converged on traffic, licenses, and liquidity. Yet every new market still requires someone to interpret the rules, verify the evidence, confirm the outcome, and trigger payment. These repeatedly rebuilt capabilities are becoming suitable for standardization and reuse across platforms.

- **Uncertainty in rules and settlement is creating alpha.** HR's regression model shows that the 0.487% of markets that entered dispute accounted for 8.64% of trading volume, while similar contracts across platforms continue to exhibit persistent price gaps. Shallow order books and execution costs, however, constrain deployable capital. The more scalable opportunity is to turn rule comparison, evidence verification, outcome confirmation, and payment triggering into products that can be called repeatedly.

- **The outcome layer has entered commercialization.** Currently observable annual revenue is approximately $15 million-$37 million. Covering the entire existing market would imply roughly $64 million-$161 million, while monetization at the level of mature projects would put the industry's upper annual-revenue bound at about $456 million.

- **An industry opportunity is taking shape, and the outcome layer remains at the stage best suited to early research and positioning.** Platform equity already prices in substantial front-end growth, while HYPE and ICE still provide limited capture of outcome-layer value. The real signal to track is which early-stage projects begin winning cross-platform calls and converting those calls into attributable, renewable revenue. The ideal investable asset has not fully formed—which also means this layer of industry value has not yet been fully priced.

> Traders of rule discrepancies earn a spread; those who define outcomes share in the future.

ABOUT HR

Heretic Research is an independent research firm focused on the intersection of crypto and AI. We look past the noise to uncover the underlying data, then investigate it in depth. We remain independent of every market narrative. All views expressed are solely those of Heretic Research.

Question everything. Alpha hides in doubt.

Analysts **Merida Lee & Kiko**

PREFACE

Prediction Markets' Overlooked Second Probability: Whoever Adjudicates the Outcome Rewrites the Price

Prediction markets usually treat the price of a Yes contract as the probability that an event will occur. That interpretation misses their most valuable variable: after a fact occurs, whether the rules recognize it ultimately determines whether the contract pays out.

A contract therefore contains two probabilities at once:

- the probability that the event occurs;
- the probability that the event is recognized as Yes under the rules.

When the rules are clear and the evidence is consistent, the two probabilities converge. When deadlines, information sources, evidentiary standards, or final adjudication diverge, they separate. That gap is not merely a technical bug. It changes cash flows and can itself be treated as an asset.

This report focuses on that still-underappreciated "invisible hand" of prediction markets and examines it through three definitions. Interpretive authority determines whether the rules recognize a fact. A **rule discrepancy** is the trading alpha created when that recognition gap enters the price. The **outcome layer** packages rule comparison, evidence verification, outcome confirmation, and payment triggering into cross-platform interfaces.

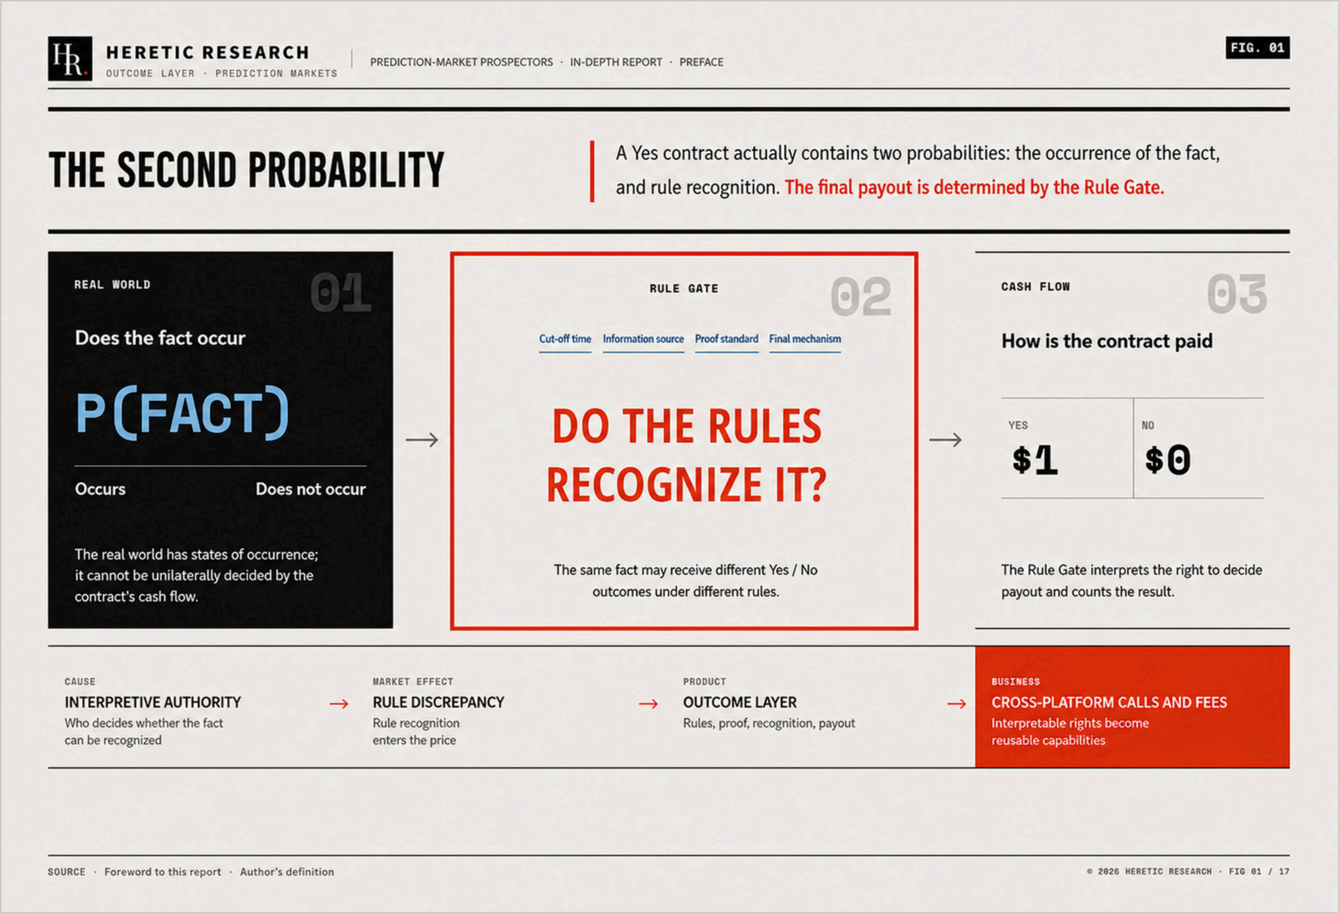

FIG. 01 — THE SECOND PROBABILITY

The three form a causal chain: interpretive authority is the cause, rule discrepancy is the effect, and the outcome layer is the commercial form of interpretive authority.

Whoever decides whether prediction-market rules recognize a fact therefore controls the overlooked second probability. In trading, it appears as rule discrepancy. At the industry level, it drives rule comparison, evidence verification, outcome confirmation, and payment triggering to be packaged as shared interfaces.

The same rule discrepancy thus creates two pools of profit. The portion captured by the outcome layer is the second profit pool named in this report's title. It is an industry-level estimate of profit distribution; no individual entity currently discloses it separately.

Combining the data evidence with industry research, this report identifies a growth opportunity that the market has not yet fully priced: **the outcome layer that can replicate across platforms and generate recurring revenue will become the main source of incremental industry growth in the next phase of prediction markets.**

SECTION 01

0.487% of Markets Entered Dispute but Accounted for 8.64% of Volume: Large Capital Behind a Small Probability

Measured by market count, disputes are a low-frequency operating issue. That denominator systematically understates the economic value of rule discrepancies. Their incidence may be small, but their economic significance is determined by the amount of capital exposed to them.

As of July 15, 2026, this report identified 282,191 markets for dispute-rate and regression analysis among closed markets with final notional volume of at least $1,000 and retained UMA status histories. UMA is the onchain dispute-resolution mechanism used by Polymarket. Within this sample, only **0.487%** of markets entered dispute, yet they accounted for **8.64%** of final notional volume.

Rule discrepancies concentrate in the most expensive markets.

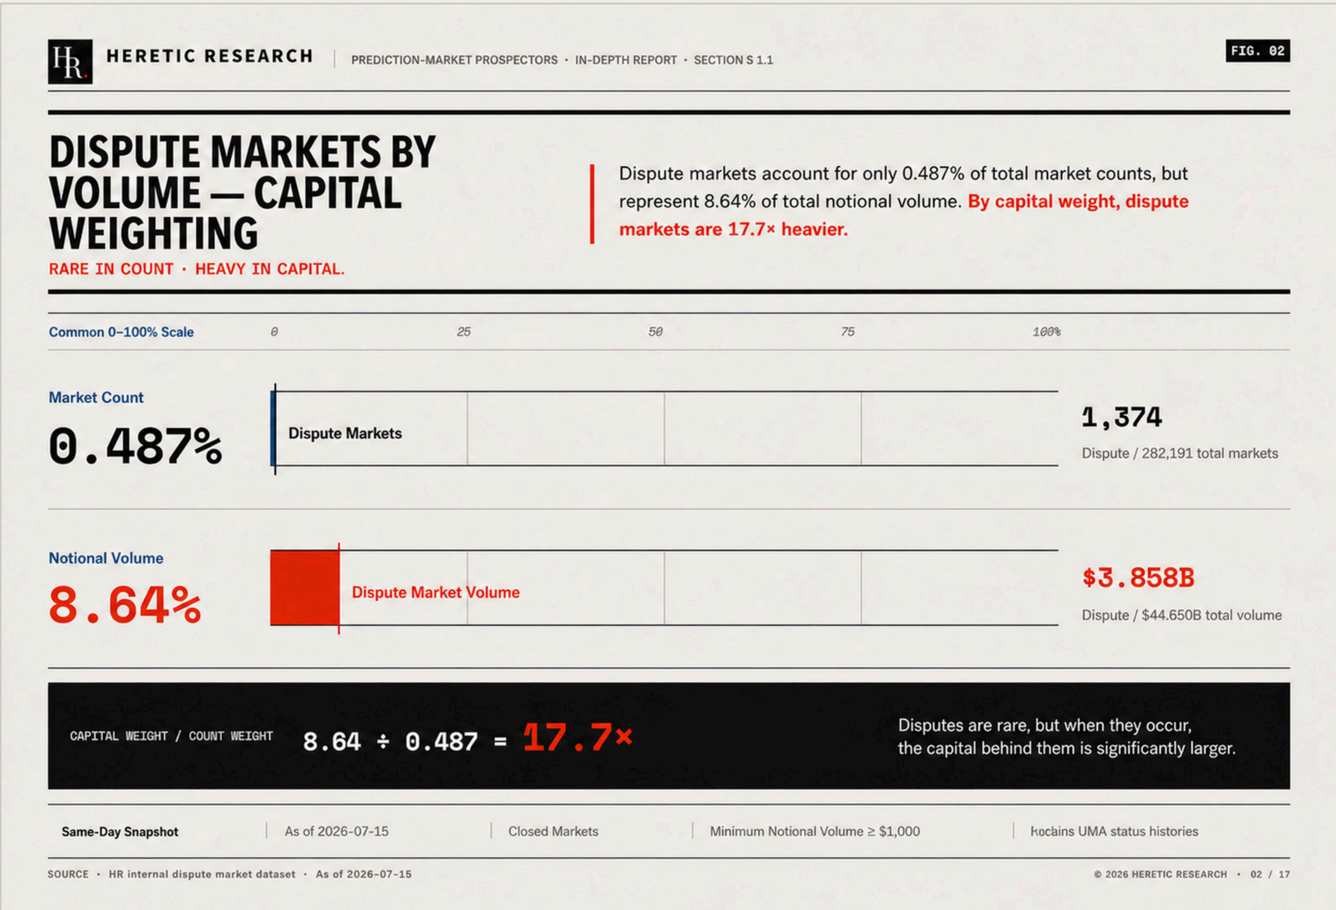

FIG. 02 — DISPUTES CONCENTRATE TRADING VOLUME

1.1 The Greater the Trading Volume, the More Concentrated the Dispute Risk

Across tiers of final notional volume, the sample dispute rate rises monotonically, from **0.26%** in the lowest tier to **30.14%** in the highest.

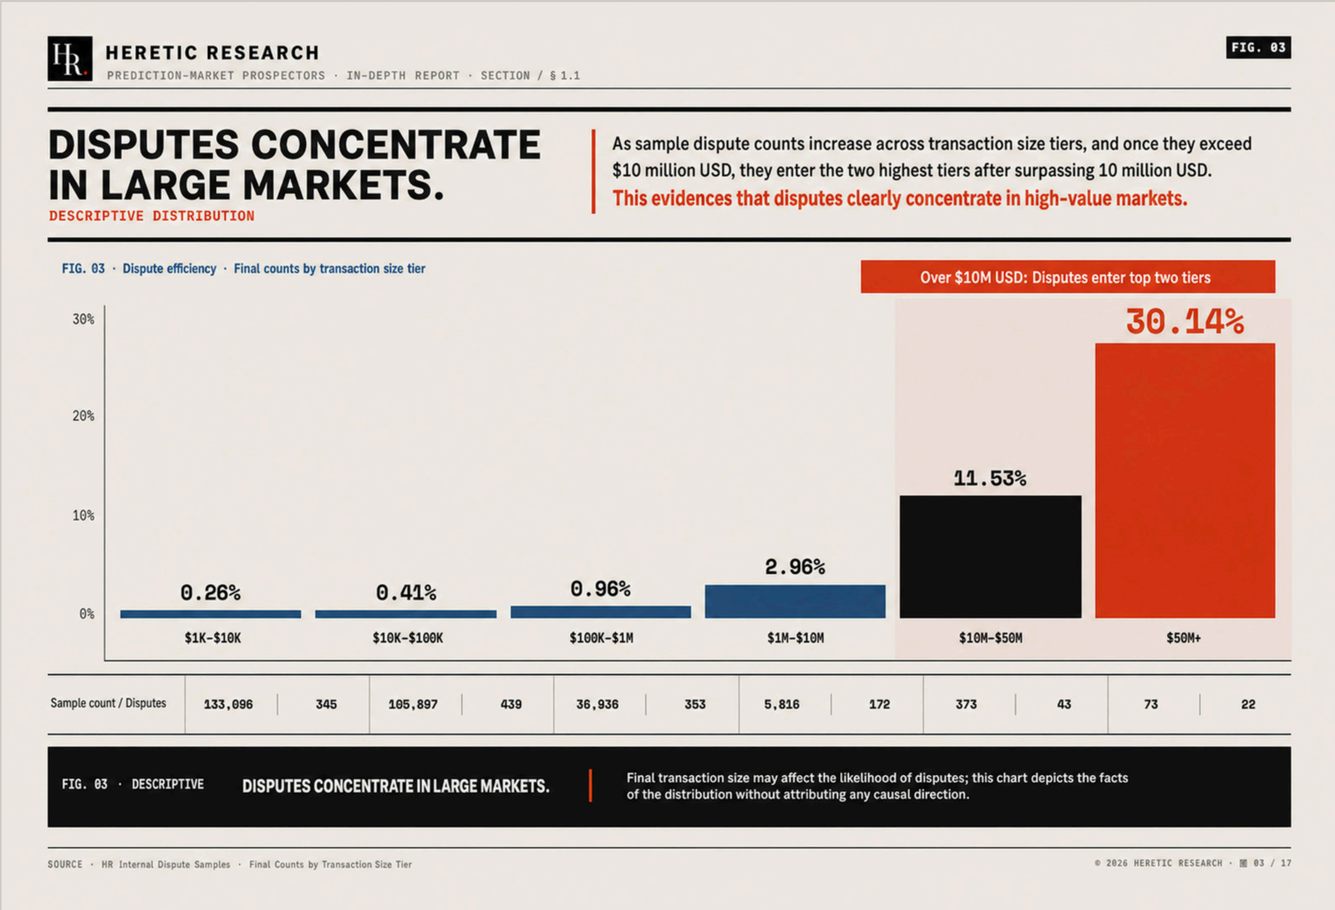

FIG. 03 — DISPUTE RATE BY VOLUME TIER

Regression modeling shows that after controlling for market year, category, duration, and other factors, volume remains **significantly positively associated** with disputes: for each tenfold increase in single-market volume, the relative propensity for that market to enter dispute rises to approximately 2.27 times its previous level. This shows that disputes are clearly concentrated in high-value markets. Because measured volume may include trading after a dispute began, however, the result establishes association, not whether high volume causes disputes or disputes generate more volume. Either way, it has become a recurring cost borne by large pools of capital.

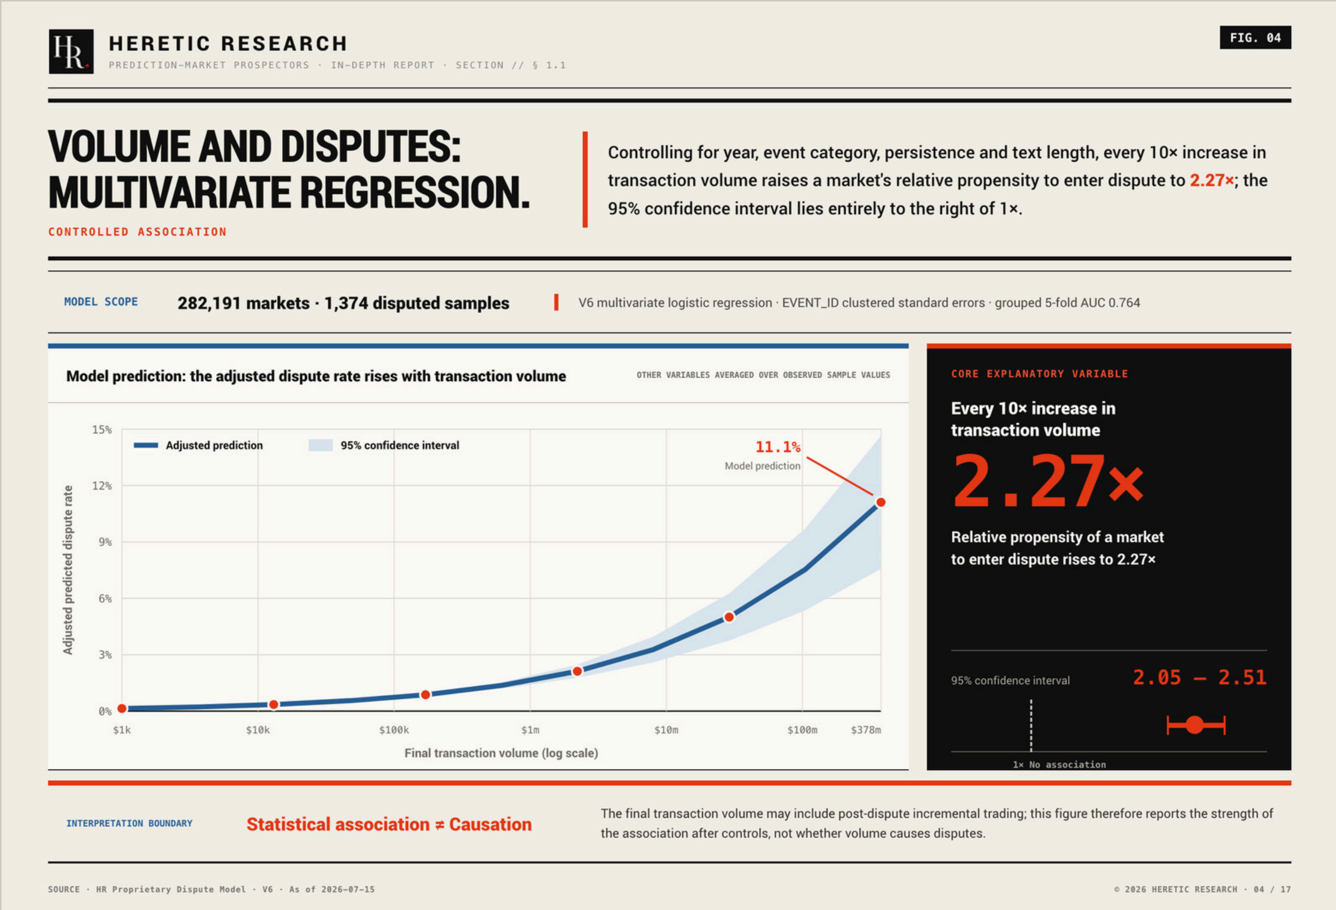

FIG. 04 — REGRESSION RESULT

1.2 High-Value Trading Brings Interpretive Authority from the Back Office to the Front

When a single event determines only a few thousand dollars, ambiguous rules can be handled ad hoc by operations staff. When it determines tens of millions of dollars, ambiguity simultaneously changes payouts, market-maker positions, and platform credibility.

High-value trading therefore requires four things to be visible before an order is placed: whether rules can be compared, evidence verified, outcomes confirmed, and payments triggered. When any one is missing, market makers charge for uncertainty through wider spreads, smaller positions, or higher capital requirements. Those costs are ultimately passed to traders as higher execution costs, lower executable size, and poorer capital efficiency.

The 0.487% dispute rate matters precisely because it immobilizes far more capital than its share of market count suggests. When rule discrepancies concentrate in the most expensive markets, the central question becomes: who has the power to change the payout?

SECTION 02

Concentrated Interpretive Authority: A Handful of Wallets Control the Final Vote on Disputes

The original terms of a prediction market cannot cover every fact. Where the terms run out, outcome confirmation shifts from rule execution to interpretive authority: whoever makes the final decision can rewrite the cash flow. This is the first source of rule discrepancy—**the allocation of interpretive authority**, or who has the final say within a platform. The more concentrated that authority becomes, the harder it is to predict the probability of rule recognition.

2.1 The Strategy Case: A Fact Can Occur Without the Contract Recognizing It

The relevant Polymarket contract stated that the market should resolve Yes if Strategy sold any bitcoin before 11:59 p.m. ET on May 31, 2026. Strategy later disclosed that it had sold **32 BTC** between May 26 and May 31.

The market ultimately resolved No. The dispute was not about whether a sale occurred, but whether it had to occur before the deadline or be publicly confirmed before the deadline. The original rule was written around the time of the event. A clarification issued after June 1 incorporated the timing of public confirmation. Galaxy Research's review of the rule text and clarification shows that traders faced two different standards for rule recognition.

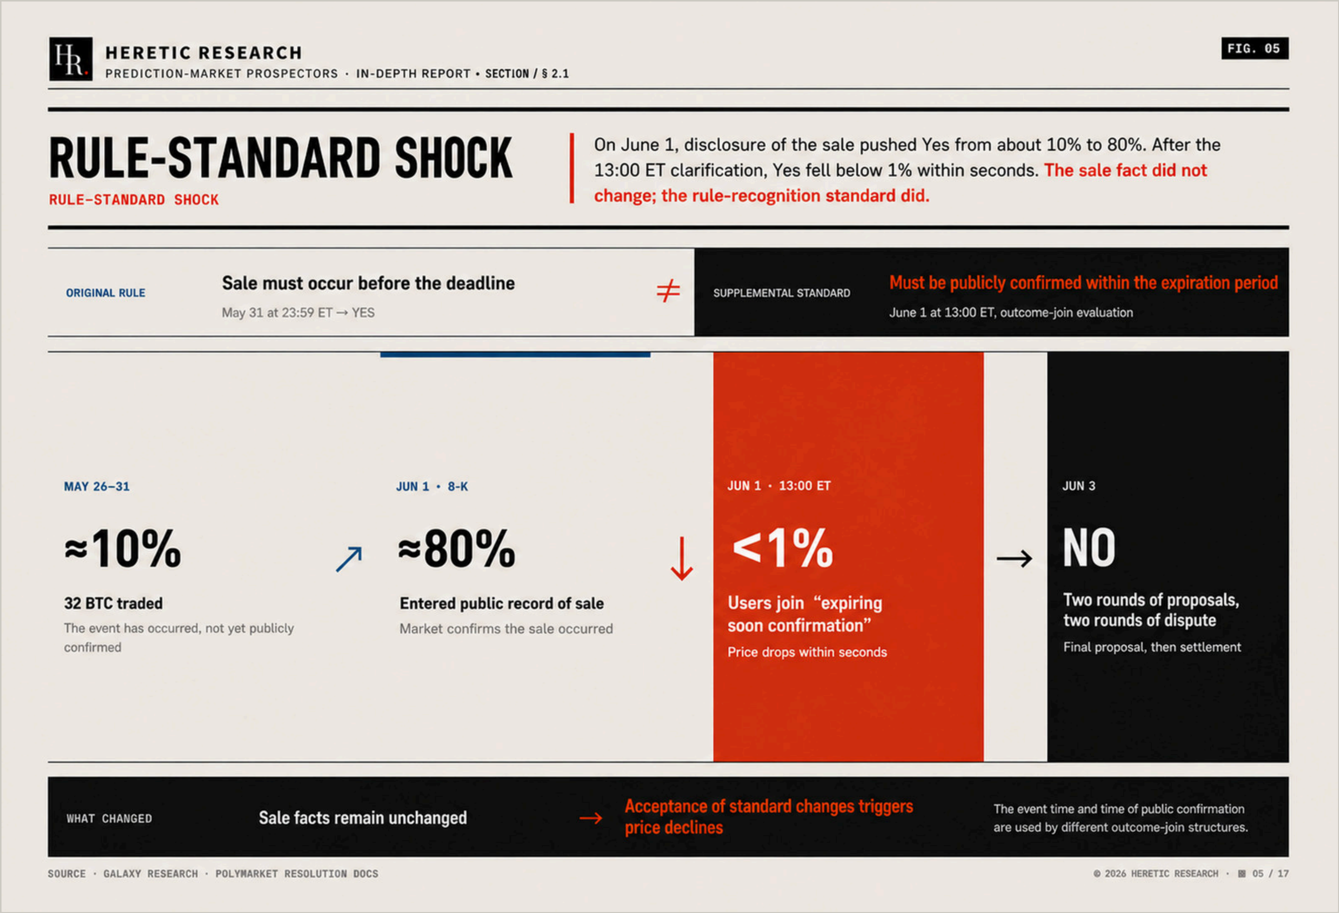

FIG. 05 — THE STRATEGY CASE

This was not simply "a dispute caused price volatility." It was an identifiable rule-standard shock. Disclosure of the fact pushed Yes upward; a clarification then drove it back down. The fact did not change. What changed was the probability that the rules would recognize it. Interpretive authority directly determined the cash flow—here through the clarification and the voting process that followed.

2.2 Open Participation Has Not Eliminated Concentrated Power

The voting process is open to everyone. Polymarket disputes proceed to a UMA vote, and any tokenholder can theoretically participate. In practice, decision power is allocated by token holdings.

Bloomberg's investigation of onchain voting found that among more than 6,400 addresses that participated in dispute votes over three years, just **nine wallets** supplied roughly half of the voting power. A separate Wall Street Journal investigation, cited by other media, found that in more than 300 disputes at least one voter also held a position in the relevant market.

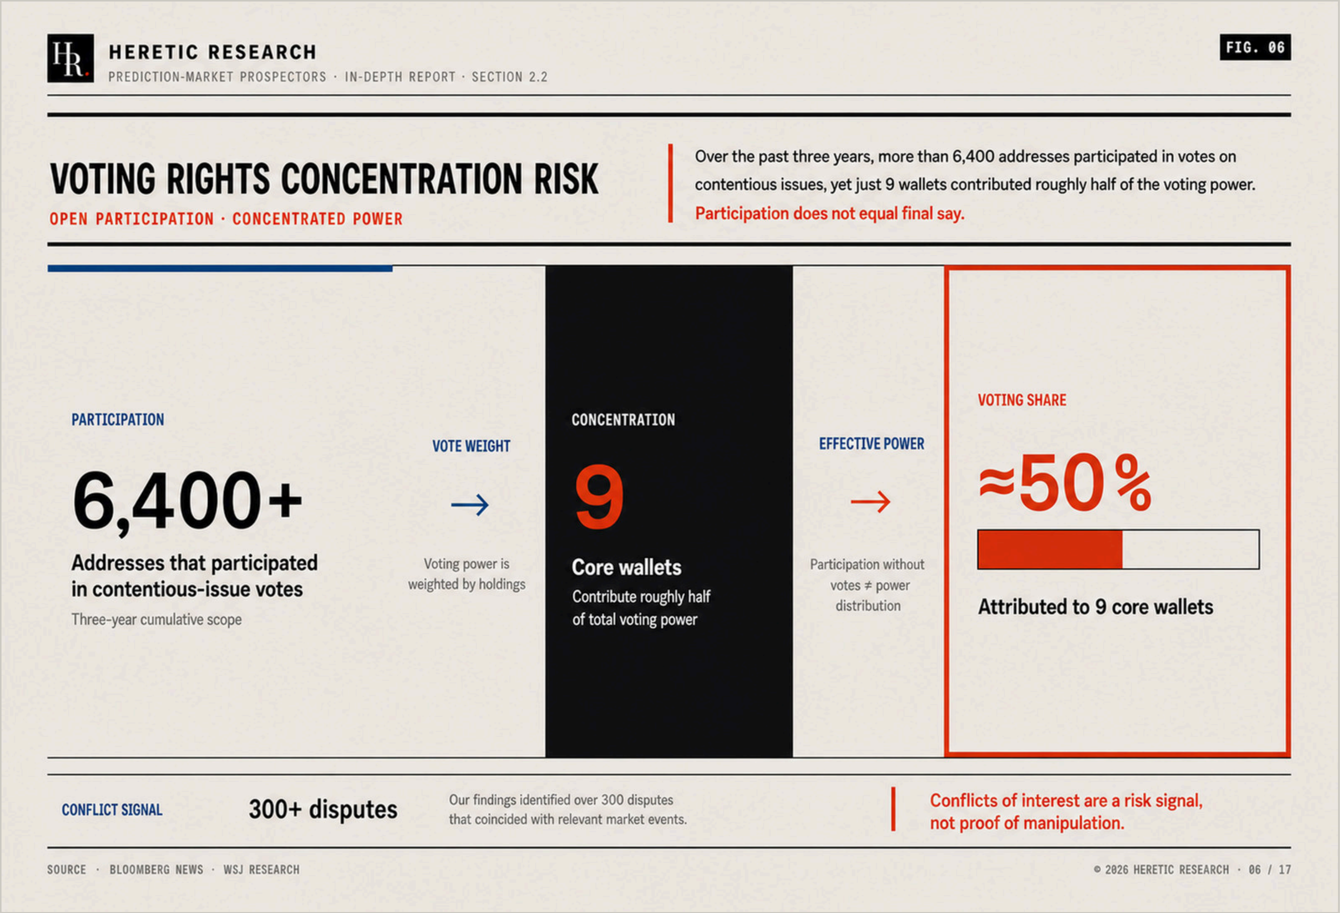

FIG. 06 — CONCENTRATION OF DISPUTE VOTING POWER

There is no need to prove manipulation for the conflict of interest itself to constitute a risk. A small number of wallets may both hold relevant positions and help determine the final settlement. Once trading interests and outcome adjudication are concentrated among the same participants, the market must price in that potential influence. Assessing its magnitude starts with the cost and threshold for entering the dispute process.

2.3 Large Positions Increase the Economic Incentive to Enter the Dispute Process

The threshold for initiating a dispute is not high. Polymarket's dispute documentation shows that a challenger generally needs only to match the proposer's bond, commonly $750. For a trader with a large position, that cost may be far smaller than the capital affected by the outcome.

Assume a trader holds a Yes position worth $100,000. If the market is proposed to resolve No, the position could lose nearly all its value. Even if the trader believes the dispute has only a 10% chance of succeeding, paying $750 to send the case into the dispute process may still be rational.

Ignore research, gas, and other additional costs for the moment. If the dispute succeeds, the trader preserves the $100,000 position and receives half of the losing side's bond, or $375. If it fails, the trader loses the $750 bond. Expected value is approximately:

$$10\% \times (\$100{,}000 + \$375) - 90\% \times \$750 = \$9{,}362.50$$

Even with a low probability of success, the gap between the amount potentially preserved and the bond creates a strong incentive to dispute. The larger the position, the lower the minimum success probability needed to justify entering the process.

Expected value of a dispute is therefore:

$$\text{Dispute EV} = q \times (V + 0.5B) - (1 - q) \times B - C$$

Here, V is the position cash flow preserved if the dispute succeeds, B is the dispute bond, C represents research, gas, and opportunity costs, and q is the estimated probability that the favorable interpretation will ultimately be accepted.

The break-even condition for initiating a dispute is:

$$q > (B + C) \div (V + 1.5B)$$

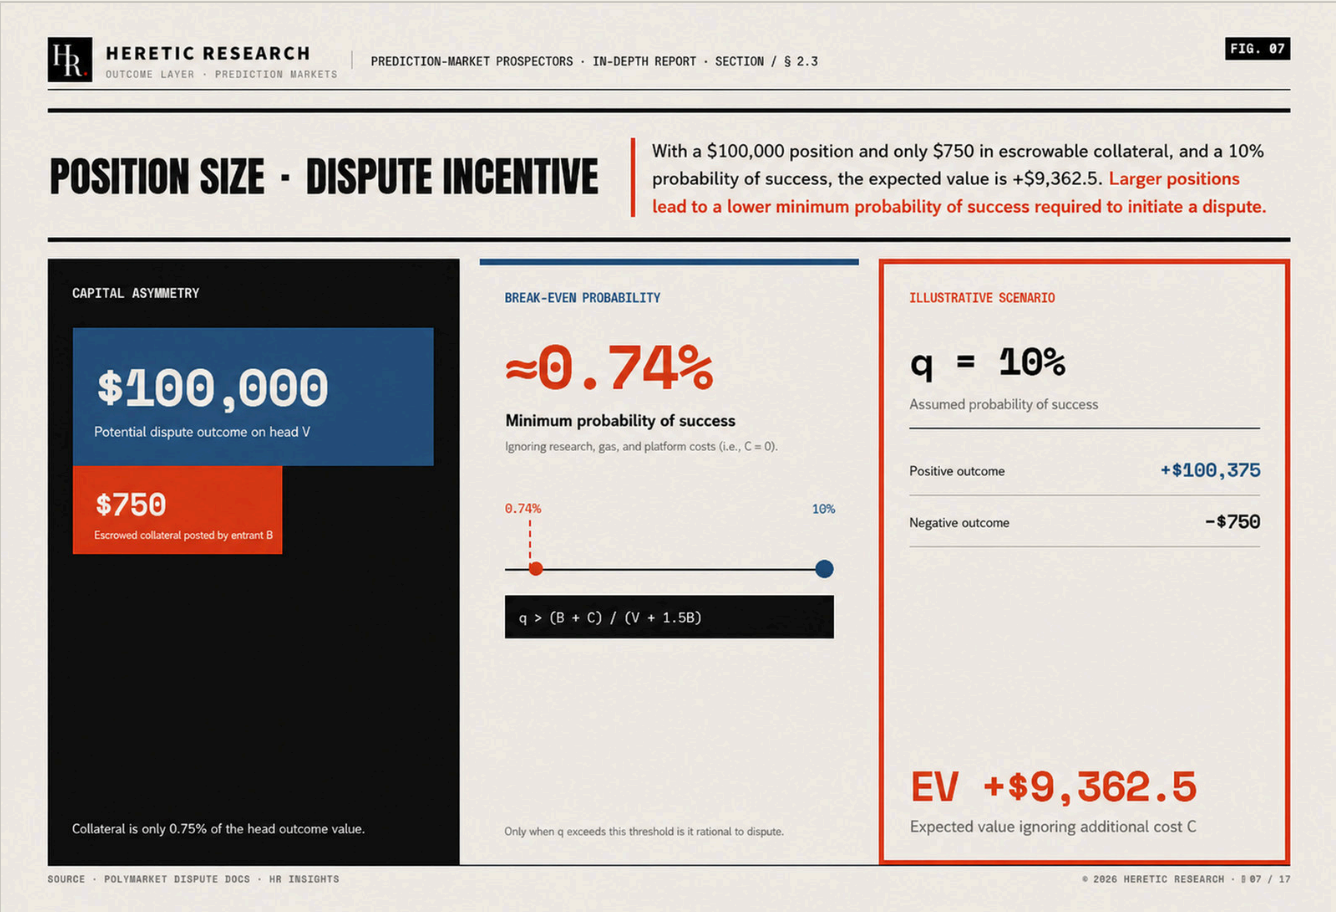

FIG. 07 — DISPUTE INCENTIVE THRESHOLD

The greater the position value, the lower the minimum success probability needed to initiate a dispute. A $750 bond cannot determine the final vote, but it is enough to put a favorable interpretation into the process that determines cash flow.

SECTION 03

Rule Discrepancies Are Already in the Price, but Strategy Capacity Stops at Order-Book Depth

The preceding section identified who holds interpretive authority. We now turn to the fact that platforms also apply different interpretive standards. This is the second source of rule discrepancy: **differences in interpretation standards**. For the same second probability, concentrated authority makes it unpredictable, while divergent standards make it unequal across platforms.

Rule discrepancy has already become a tradable pricing divergence. Beyond judging whether an event will occur, prediction-market traders must first determine whether two apparently identical contracts use the same deadline, information source, and settlement standard—and whether they pay out under the same circumstances. A cross-platform price gap can be an arbitrage only when the settlement conditions are identical. **When the rules differ, traders are really betting that the market has mispriced those rule differences.**

3.1 Cross-Platform Pricing Evidence: Semantically Equivalent, Yet Still Differently Priced

A 2026 cross-platform study covering ten major platforms and more than 100,000 events manually validated natural-language descriptions, outcome-confirmation semantics, and time frames. It found that some events were listed on multiple platforms, yet semantically equivalent markets still exhibited persistent, executable price divergences averaging **2%-4%**. The study calls this “semantic non-fungibility”: without standardized event identifiers and outcome-confirmation standards, prices cannot fully converge.

A more current order-book observation points in the same direction. From July 1 to July 7, 2026, among BTC year-end $120,000 contracts traded on both venues, Kalshi's Yes price averaged approximately 2.08 percentage points above Polymarket's for the week. The matched Polymarket-Kalshi contract observation turns the idea that “rules and platform structure enter the price” into a visible spread.

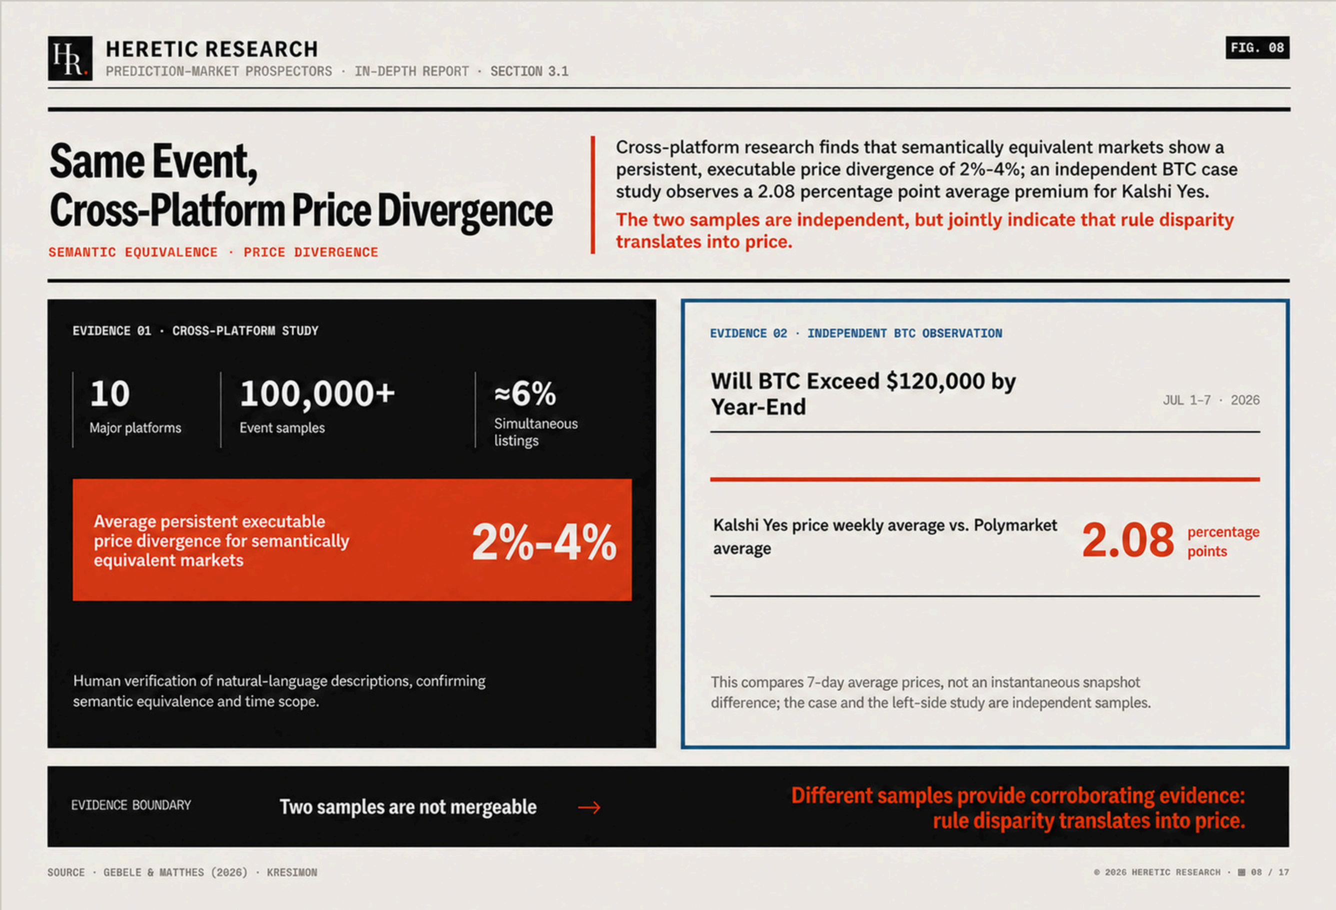

FIG. 08 — CROSS-PLATFORM PRICING DIVERGENCE

The market is not failing to price the second probability; it simply lacks a common language for doing so. For traders, that is precisely the source of alpha.

3.2 Identical Settlement Conditions Earn a Spread; Different Conditions Earn a Judgment Premium

Cross-platform arbitrage is a classic prediction-market strategy. But before trading, one must first determine whether two apparently identical contracts really pay out under the same circumstances.

Assume Platform A and Platform B both list a market on whether bitcoin will exceed $120,000 at year-end. A Yes contract on Platform A costs $0.54, while a No contract on Platform B costs $0.43, for a combined cost of $0.97.

If the contracts use exactly the same deadline, price source, and settlement standard, one side will pay $1 regardless of the outcome. After fees, slippage, and funding costs, the trader can lock in the spread so long as total cost remains below $1.

If Platform A instead uses an index price at 23:59 UTC on December 31 while Platform B uses the Eastern Time closing price, the contracts may settle differently. Buying both sides for $0.97 is no longer risk-free arbitrage. The trader is now judging whether the market underestimates the probability that the two rule sets produce different outcomes.

Cross-platform trades therefore fall into two categories:

- **Identical settlement conditions:** the trader earns a pricing dislocation and can lock in a return whenever all-in cost is below the final payout.
- **Different settlement conditions:** the trader earns a judgment premium, with returns determined by whether their understanding of deadlines, information sources, and settlement standards is better than the market's.

For the second category, the relevant estimate is not whether the event itself occurs, but the probability that the contract ultimately resolves Yes under its rules.

Suppose a Yes contract costs $0.54 and total fees, slippage, and funding cost amount to $0.015. The trader must assign at least a 55.5% probability that it ultimately resolves Yes for the trade to have positive expected value.

> Break-even probability of rule recognition = market price + all-in cost per contract

If the trader estimates a 60% probability of a Yes settlement—above the 55.5% break-even line—expected profit is about $0.045 per contract. If the estimate is only 53%, the trade is not worthwhile even if the trader believes the underlying event itself is likely to occur.

Two contracts referring to the same predicted event are therefore not automatically arbitrageable. Traders must first determine whether their settlement conditions match. If they do, the opportunity is a price spread. If they do not, the opportunity is superior understanding of the rules.

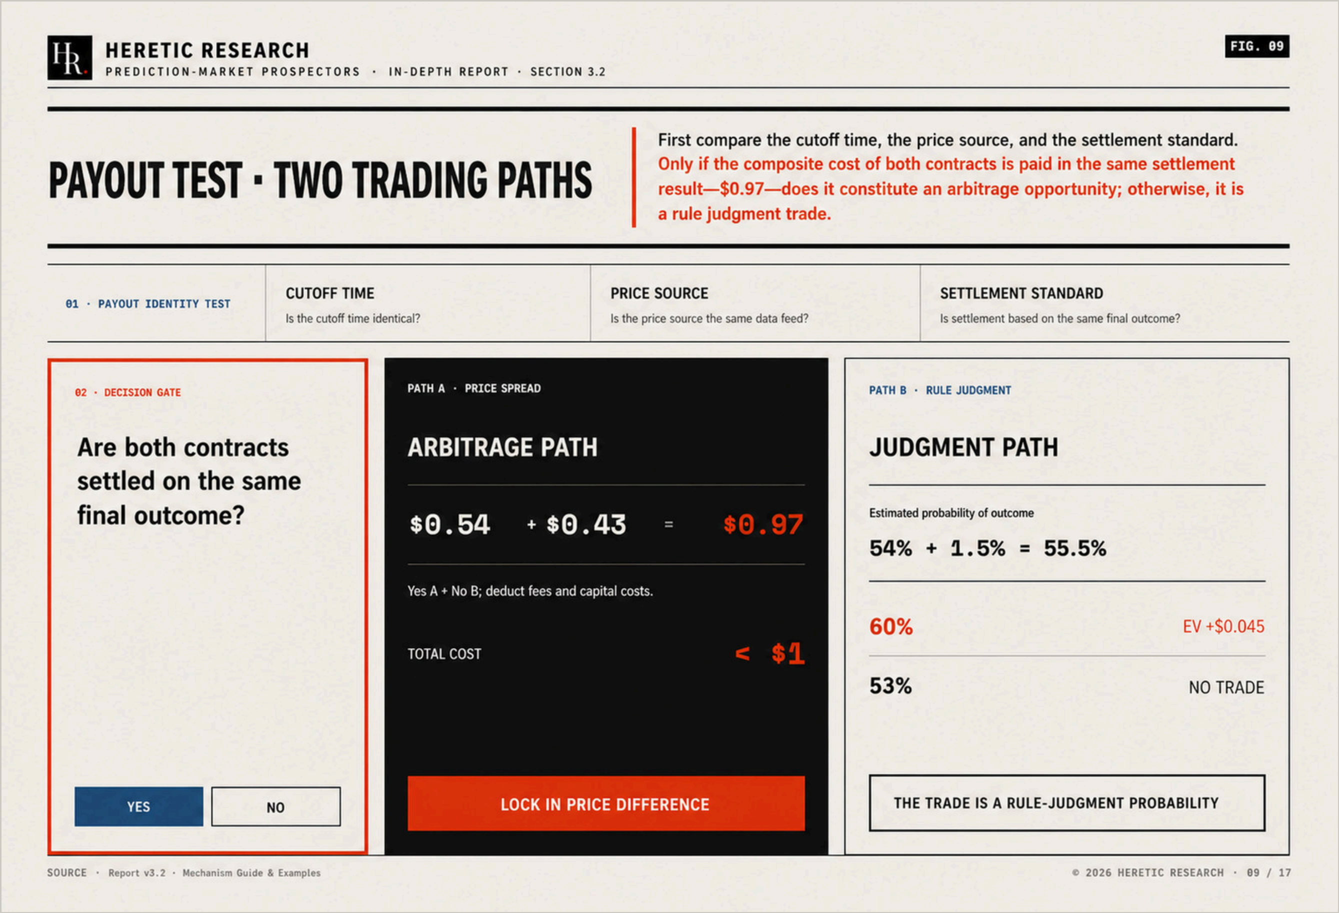

FIG. 09 — PRICE SPREAD VERSUS RULE JUDGMENT

3.3 The Spread Is Real and Arbitrageable, but Capacity Is Constrained by Order-Book Depth

The preceding section explained how to identify and construct a rule-discrepancy trade. Whether such a strategy can absorb significant capital, however, depends on two practical questions: can the spreads actually be realized, and how much capital can each opportunity accommodate?

First, these spreads are not merely theoretical. They have already formed a strategy market. Onchain research into markets settled between April 2024 and April 2025 found that within-market rebalancing and cross-market combination trades generated approximately **$40 million in realized arbitrage profit**. Prediction-market mispricing can therefore be identified systematically and converted into actual returns.

Second, executable scale remains extremely limited. A study covering 173 NBA games reconstructed more than 75 million order-book snapshots and identified 290 combination-arbitrage episodes. Their median gross order-book return per opportunity was only about 1.01%; 76.9% were constrained by depth, and depth-constrained opportunities allowed an average execution of just 14.8 contracts. Executable single-market anomalies were rarer still, with a median duration of only 3.6 seconds.

If an arbitrage opportunity offers a gross return of roughly 1%, each $100 theoretically produces about $1. But if the book can fill only about 15 contracts at nearly $1 each, **only about $15 can actually be deployed, for perhaps $0.15 in gross profit**. The return percentage looks attractive, but deployable capital is tiny and the opportunity may vanish within seconds.

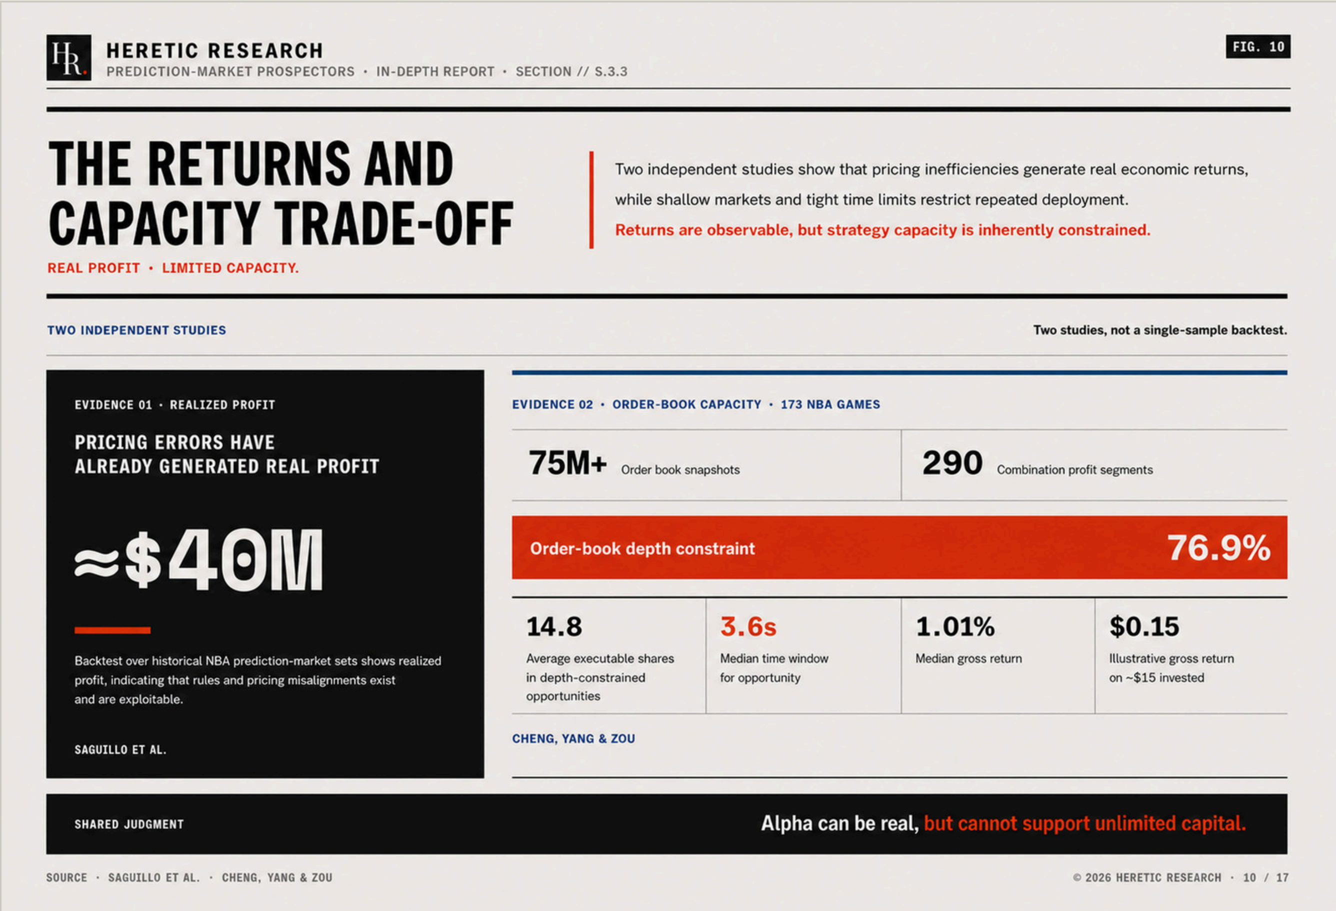

FIG. 10 — ARBITRAGE CAPACITY AND ORDER-BOOK DEPTH

Rule-discrepancy trading is therefore a business built on research and execution, not an infinitely scalable market windfall. Contract research, semantic matching, genuine order-book identification, and multi-leg execution determine whether the opportunity can be captured. Depth and opportunity duration determine how much capital the strategy can absorb.

Traders earn a one-off mispricing and must find and execute the next one from scratch. The business that can scale with aggregate prediction-market volume may not be trading these frictions directly, but charging continuously for infrastructure that reduces rule uncertainty, improves outcome confirmation, and limits manual discretion.

SECTION 04

The Bigger Alpha Is Not in the Order Book: Turning Interpretive Authority into a Product

Why are prediction-market order books generally shallow? One important reason is uncertainty in settlement rules. Market makers cannot know with precision how a platform will ultimately recognize a fact, so they widen spreads and reduce quoted size to compensate for the risk.

That uncertainty produces two different businesses. Traders identify and exploit mispricing caused by rule discrepancies, but their returns are limited by individual-market depth and opportunity duration. The outcome layer instead seeks to reduce uncertainty in rules and settlement by turning rule comparison, evidence verification, outcome confirmation, and payment triggering into repeatable services sold to multiple markets and platforms.

Both arise from the same friction: traders make money because the friction exists; the outcome layer makes money by reducing it.

To do so, the outcome layer must first unbundle how outcomes are confirmed. Different events involve different levels of difficulty. Facts such as prices, times, and onchain states can be read directly by machines. Events with scattered evidence but relatively clear rules can use AI for retrieval and organization. Ambiguous rules, large financial stakes, or conflicting evidence still require a human final judgment.

Productization begins with this division of labor. Each confirmation task that becomes standardized adds another capability that platforms can procure and call.

4.1 Deterministic Outcomes Are Being Standardized

The outcome layer begins productization with the clearest, easiest-to-verify events: the bitcoin price at a specified time, whether an onchain address completed a transfer, or whether a state appeared before a deadline. If the data source, read time, and decision condition are specified in advance, the system can read the data, determine Yes or No, and trigger payment automatically.

For a market asking whether bitcoin exceeds a specified price at a specified time, the market need only define the price source, observation time, and threshold for Yes. At the deadline, the system reads the data and settles automatically, without ad hoc rule interpretation.

Hyperliquid's HIP-4 turns this process into a standardized product. It integrates binary outcome contracts into a unified trading system so that rule setting, trading, outcome confirmation, and payment occur within the same infrastructure—and the template can be reused for the next set of events. Demand has validated this standardized product: short-duration Polymarket crypto markets powered by Chainlink have accumulated more than **$3.4 billion in volume**.

Regulated platforms are adopting the same approach. Some Kalshi markets in gold, crude oil, and agricultural products use external price data from Pyth as the basis for settlement. Outcome confirmation can thus move from an internal platform process to a service supplied by a specialist data provider and reused across markets.

Deterministic outcomes matter not only because they are easier to automate. They are the first proof that outcome confirmation can be decomposed into a standardized interface and sold as a foundational outcome-layer product.

4.2 AI's Value Is to Compress the Cost of Evidence

Deterministic outcomes allow direct data reads, but many prediction markets face a different problem: the rules are written, while the relevant evidence is dispersed across announcements, news reports, regulatory filings, and conflicting sources. The costly part is often not the final judgment, but finding the materials, filtering valid evidence, and organizing the basis for a decision.

AI is best suited to this work. It can accelerate retrieval, synthesize sources, identify conflicts, and generate an evidence package for final review. Whether it can replace outcome adjudication, however, depends on observed performance.

A study of UMA dispute markets found that AI achieved **89.58%** agreement in post-dispute review, but only **33.88%** recall when identifying disputes in advance. The former suggests that AI can often reproduce the eventual decision once the materials are available; it does not show that AI independently discovered the “true answer.” The latter shows that AI still fails to identify most high-risk markets before a dispute.

AI therefore belongs in the middle of the decision process, not at its end. Its current advantage is reducing the cost of evidence processing before judgment.

A more workable division of labor is: machines read unambiguous data; AI retrieves, filters, and organizes evidence; humans make the final call where rules are ambiguous, stakes are large, or evidence conflicts.

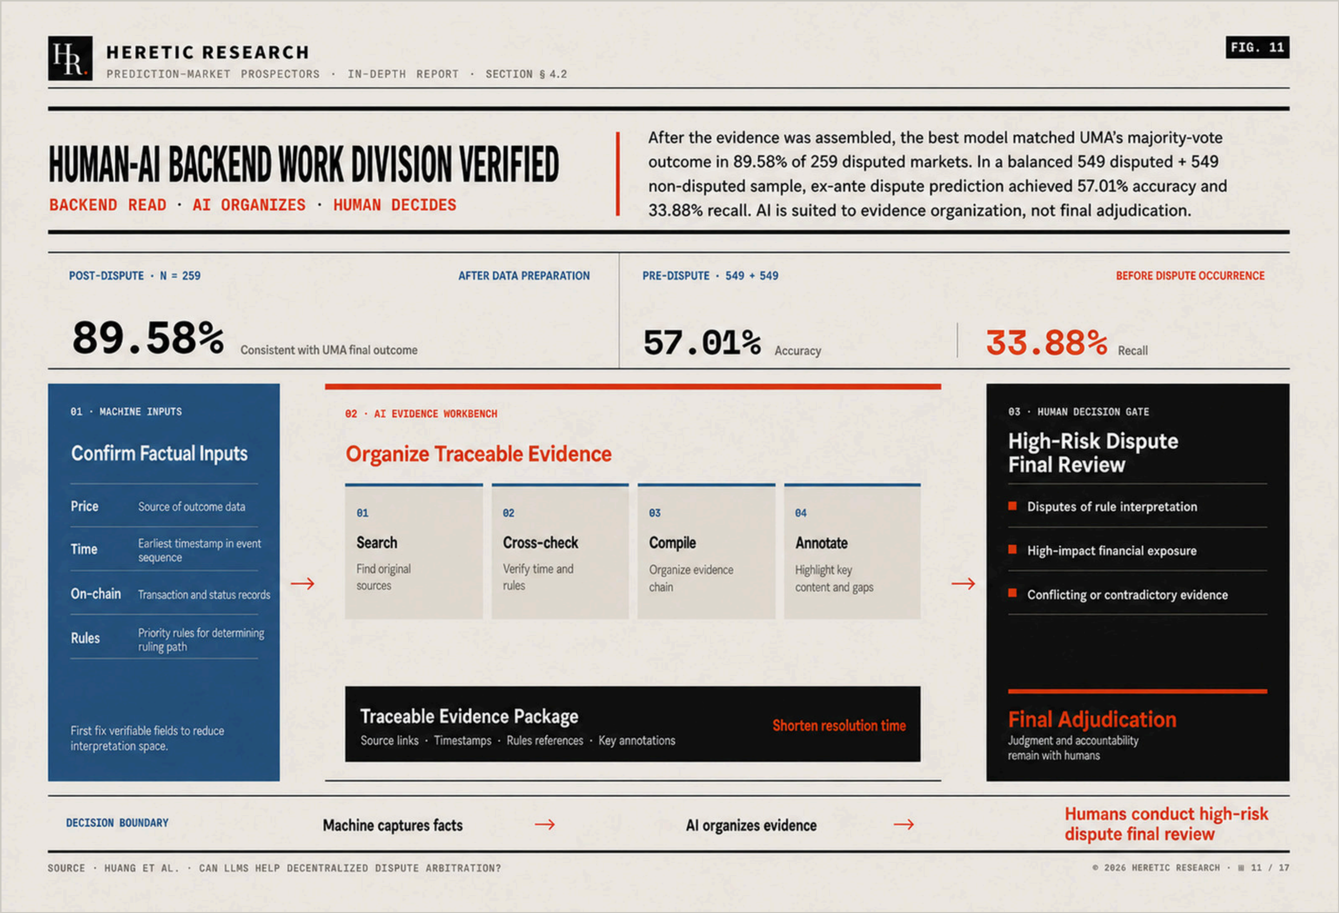

FIG. 11 — MACHINE, AI, AND HUMAN DIVISION OF LABOR

Once standardized, outcome confirmation is no longer merely an internal operating process. Rule comparison, evidence verification, outcome confirmation, and payment triggering can be exposed as separate interfaces and repeatedly called by platforms, market makers, institutions, and agents.

That is the starting point of the outcome layer: turning rules, evidence, and settlement capabilities scattered inside individual platforms into infrastructure that can be purchased across platforms. The next chapter asks how repeated calls to these capabilities become revenue.

SECTION 05

Charging for Interpretive Authority: The Outcome Layer Is Forming a Second Profit Pool

Once rule comparison, evidence verification, outcome confirmation, and payment triggering become reusable cross-platform interfaces, revenue comes not only from final settlement but from repeated calls throughout the trading chain. Interpretive authority is no longer merely governance power; it begins to generate a second profit pool.

Two conditions are required for independent outcome-layer revenue: the same capability must be reusable across markets rather than rebuilt for each event, and platforms, institutions, or developers must be willing to pay for those calls on an ongoing basis.

HIP-4 first demonstrates the former: outcome confirmation can be automated and replicated across different kinds of markets.

5.1 HIP-4 Completes the Outcome-Asset Product Loop, but Volume Remains Unstable

HIP-4 has already shown that deterministic outcomes such as prices can be machine-read and settled automatically. Its further commercial significance is that the same trading and confirmation process need not serve only one event.

When a new outcome market launches, the platform does not need to rebuild the entire system. Market creation, rule encoding, order distribution, liquidity integration, outcome confirmation, and payment can all reuse the established process:

> Event becomes a rule → rule creates an outcome asset → market trades continuously → machine confirms outcome → funds settle automatically → template is reused for the next events

Real trading has validated that reusability. Less than three months after launch, HIP-4 had generated approximately **$277 million** in cumulative volume and **$182 million** over the latest 30 days. Users are willing to trade machine-confirmed, automatically settled outcome assets—not merely accept them in theory.

The trading mix also shows that the same infrastructure can serve different demand. Sports markets contributed approximately **82%** of this volume. After major-event peaks subsided, however, activity shifted mainly back to crypto markets. Sports provide event-driven peaks; crypto provides more frequent and stable everyday demand.

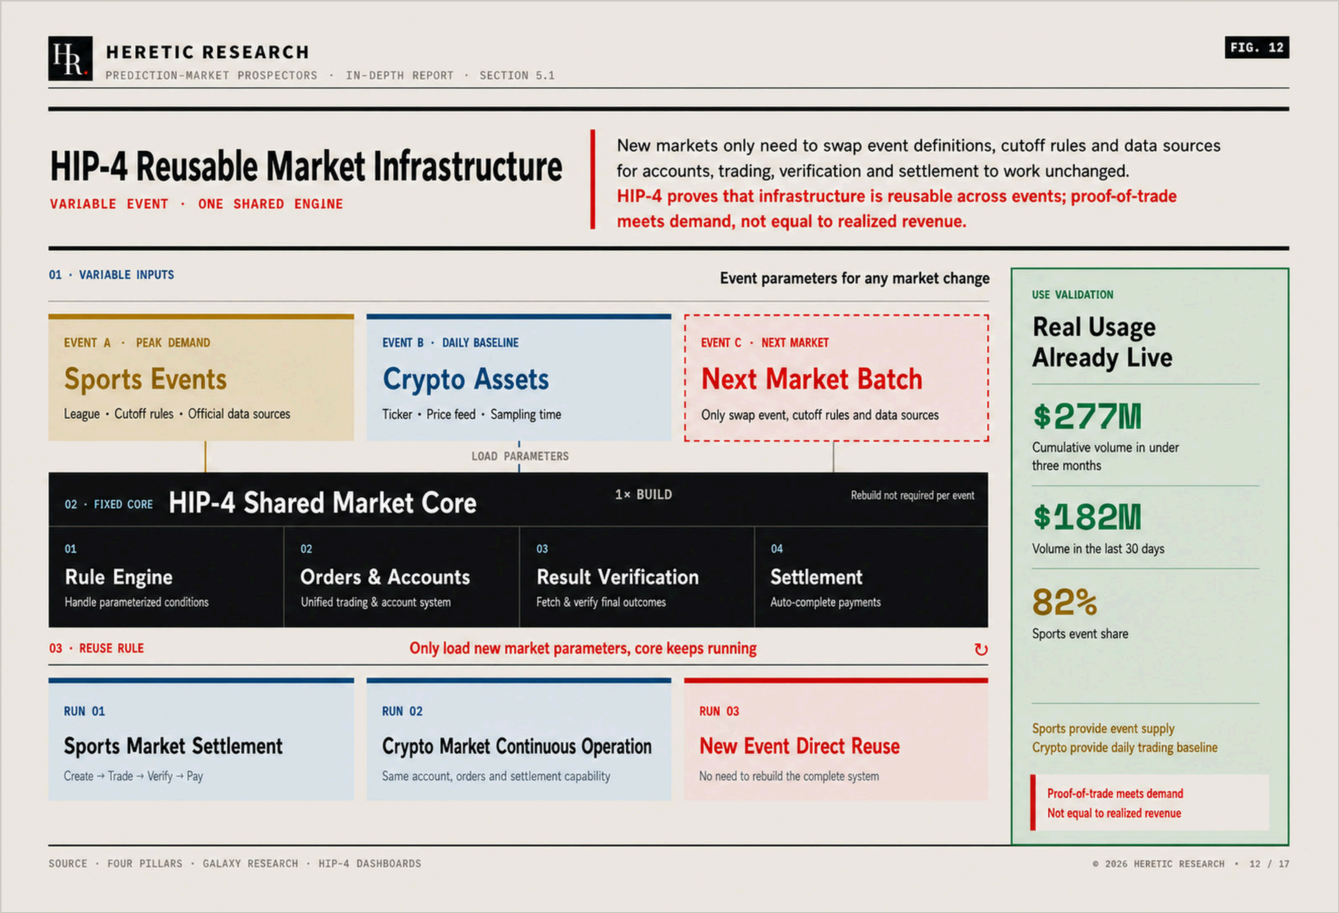

FIG. 12 — HIP-4 PRODUCT LOOP AND TRADING MIX

This gives the outcome layer the growth characteristics of software and data infrastructure. New markets no longer build settlement from scratch; the four established capabilities can be replicated across events. More categories produce more calls, raising the marginal value of the standardized interface. Whether those calls become independent revenue still depends on who pays and how.

5.2 Calls, Workflows, and Data Licensing: Three Revenue Paths and a Shift in Buyers

The outcome layer can charge along three paths: outcome calls and payment triggers; institutional data and audit workflows; and outcome-data licensing and distribution. Each corresponds to a different buyer, reason for purchase, and unit of charge.

FIG. 13 — THREE OUTCOME-LAYER REVENUE PATHS

The paths differ in significance. Institutional contracts can stabilize revenue. Data subscriptions and licensing can increase recurring revenue and improve valuation quality. Scale, however, still depends primarily on charges for outcome calls and settlement.

Four emerging suppliers show buyers expanding from onchain protocols to regulated platforms and institutional budgets:

• **Chainlink proves demand for high-frequency outcome-interface calls.** Chainlink-powered short-duration Polymarket crypto markets have exceeded **$3.4 billion** in cumulative volume. When standardized data can automatically confirm outcomes, the same interface can be called repeatedly by large numbers of markets.

• **Azuro proves that infrastructure calls can become actual revenue.** Azuro generated approximately **$4.5 million** in cumulative protocol revenue from roughly **$530 million** in cumulative volume. It provides not only final outcomes but infrastructure for market creation, data provision, and settlement. Prediction-market back-end capabilities can therefore become independent revenue through protocol charges and revenue sharing rather than remain only internal platform costs.

• **Pyth proves that regulated platforms will procure external outcome data.** Kalshi uses Pyth data in some gold, crude-oil, and agricultural markets. Unlike per-call onchain protocol charges, this resembles an exchange procuring external market data and settlement inputs. The purchasing rationale shifts from onchain traffic to compliance, reliability, and division of responsibility.

- **ICE proves that prediction-market data can enter institutional information products.** ICE distributes Polymarket probability and sentiment data to institutional clients through real-time feeds and historical databases. What is sold is no longer only a market's final outcome, but the data continuously produced by prediction markets. Charging accordingly shifts toward subscriptions, licensing, and revenue sharing.

Together, the cases prove revenue formation at different points: Chainlink demonstrates call demand, Azuro demonstrates protocol charges, Pyth demonstrates regulated procurement of external inputs, and ICE demonstrates institutional distribution of outcome data.

Combined with HIP-4, they reveal a relatively complete commercialization chain:

> Outcome products are traded → interfaces are called repeatedly → platforms procure external capabilities → outcome data enters institutional workflows

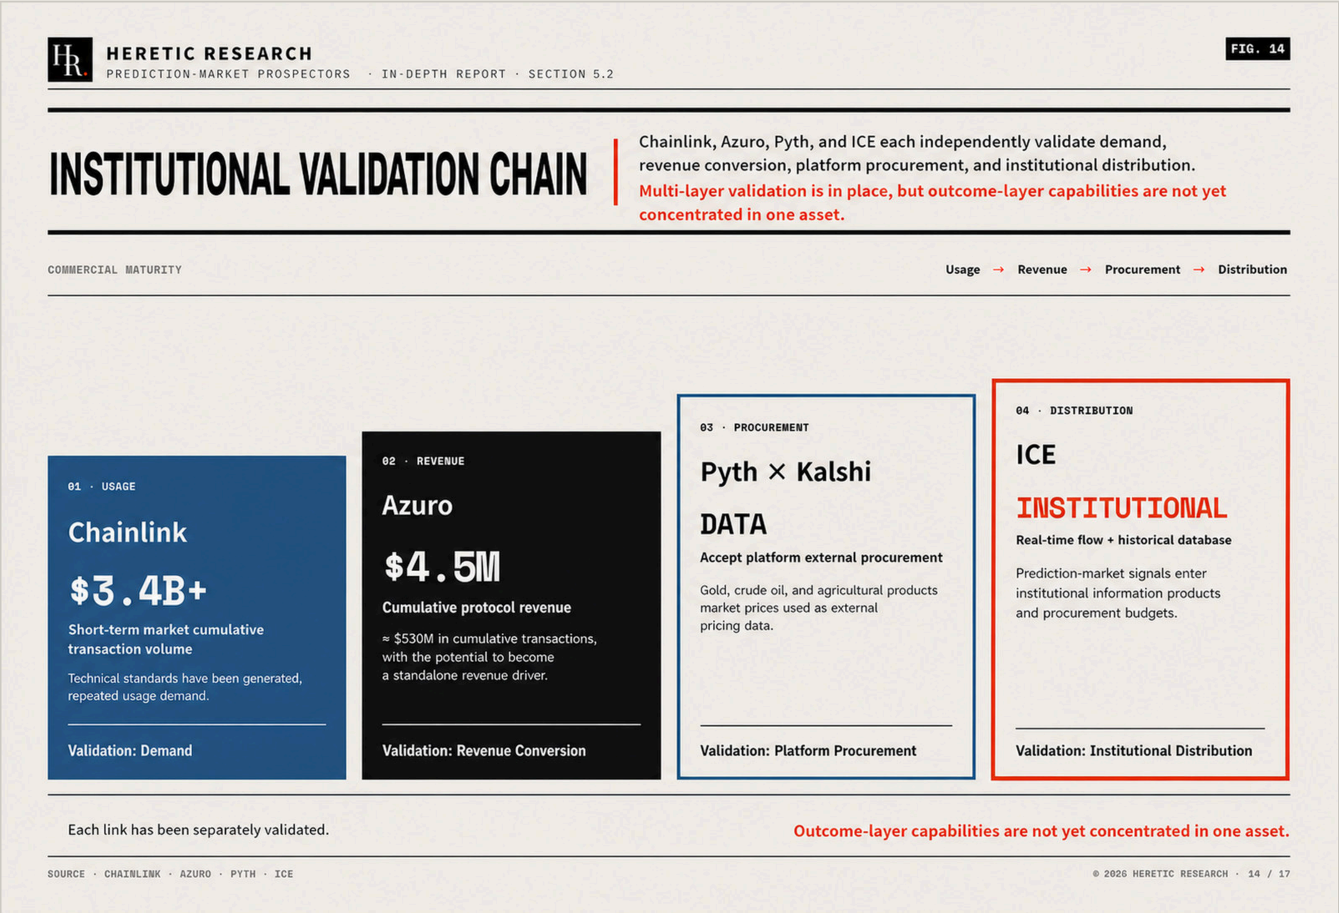

FIG. 14 — OUTCOME-LAYER COMMERCIALIZATION CHAIN

Market structure also determines the importance of each revenue path. In June 2026, combined volume across Kalshi, Polymarket US, and international Polymarket reached **$44.8 billion**, implying a simple annualized run rate of **$537.6 billion**. Regulated markets accounted for **77.1%**, and onchain markets for **22.9%**.

The two segments need different things from the outcome layer. Onchain markets typically unbundle confirmation among oracles, dispute mechanisms, and smart-contract payments, making them suited to per-call, per-settlement, or protocol-revenue-share pricing. Regulated platforms retain final settlement responsibility but may still procure external data, audit tools, and institutional distribution services, favoring annual contracts, API subscriptions, and data licensing.

Because regulated markets already account for most volume, the larger future revenue opportunity may come from compliance data, audit tools, and institutional workflows rather than simply replicating onchain dispute fees. This conclusion has an important condition: independent suppliers must provide those capabilities externally. If exchanges build and keep them in-house, they remain platform costs and cannot be extrapolated across total market volume as external outcome-layer revenue.

5.3 From a Tens-of-Millions Base to a Hundreds-of-Millions Second Profit Pool

The preceding sections establish genuine demand and emerging charges through protocol revenue sharing, external data procurement, and institutional licensing. The industry does not yet disclose “outcome-layer revenue” separately, so a complete market-size figure cannot be observed directly.

This section therefore estimates it in three steps:

1. begin with protocol revenue already generated to estimate currently observable outcome-layer revenue;

2. assume the same monetization capability covers the existing prediction market to estimate scalable revenue capacity;

3. use a project with an established commercial model to estimate a mature-state upper bound.

These are not three independent forecasts, but three stages along one commercialization path. Potential revenue rises at each stage, as does reliance on assumptions.

Currently Observable Commercial Base: $15 Million-$37 Million

As of July 28, 2026, the 33 onchain prediction-market protocols tracked by DefiLlama generated approximately $146 million in revenue over the prior year and about $30.26 million over the latest 30 days.

Not all of that revenue belongs to the outcome layer; it also includes platform, liquidity, and other protocol components. To estimate the outcome-layer share, this report uses Azuro's actually implemented allocation rule: **10%** of pool profit is assigned to Data Providers.

At that proportion:

- prior-year revenue implies about $14.61 million for the outcome layer;

- annualizing the latest 30 days implies about $36.82 million.

Rounded, public revenue and an implemented allocation rule jointly support an observable outcome-layer revenue range of approximately $15 million-$37 million.

This is not the industry's total revenue. It is the portion currently confirmable from public data—commercialization that has begun, not a forecast.

Covering the Existing Market: $64 Million-$161 Million

Using the June volume denominator above, onchain markets imply approximately **$123.1 billion** in annualized volume. Dividing observable outcome-layer revenue by that volume produces the currently observable monetization rate:

> $15 million-$37 million ÷ $123.1 billion = 1.19-2.99 bp

At 1.19-2.99 bp, each $10,000 of prediction-market volume currently generates about **$1.19-$2.99** of outcome-layer revenue. If that monetization rate expanded beyond onchain markets to the whole existing prediction market, annual industry revenue would be approximately **$64 million-$161 million**.

This case does not require continued rapid growth in prediction markets. It assumes only that more existing volume uses external outcome-layer services. It still requires two conditions: outcome confirmation, data provision, and payment triggering must cover existing volume; and regulated platforms must continue outsourcing data, audit, and outcome workflows, producing observable contracts, invoices, and renewals.

Mature Productization Upper Bound: Approximately $456 Million

The second stage uses the industry's current, relatively low monetization rate. A more mature case can be anchored to Azuro's established charge structure.

Azuro's cumulative protocol revenue equals approximately **84.9 bp** of cumulative volume. With roughly 10% allocated to Data Providers, the implied outcome-layer monetization rate is:

> 84.9 bp × 10% = 8.49 bp

That equals approximately **$8.49** of outcome-layer revenue per $10,000 of volume. It is not a theoretical fee but a level achieved by an operating project. If the entire prediction market reached the same rate, annual outcome-layer revenue would be approximately **$456 million**.

Reaching it would require the outcome layer not only to cover more markets but also to capture a larger share of rule comparison, evidence verification, outcome confirmation, and payment triggering. The $456 million figure is therefore better treated as a mature-business-model reference ceiling than as a near-term forecast.

The 10-18× range comes from transactions involving data and financial-infrastructure businesses including IHS Markit, Black Knight, and Adenza, and is used to value recurring, highly visible industry revenue.

Once total outcome-layer annual revenue exceeds **$100 million**, a **$1.0 billion-$1.8 billion** operating value would move the industry from early product validation into the range of institutional allocation.

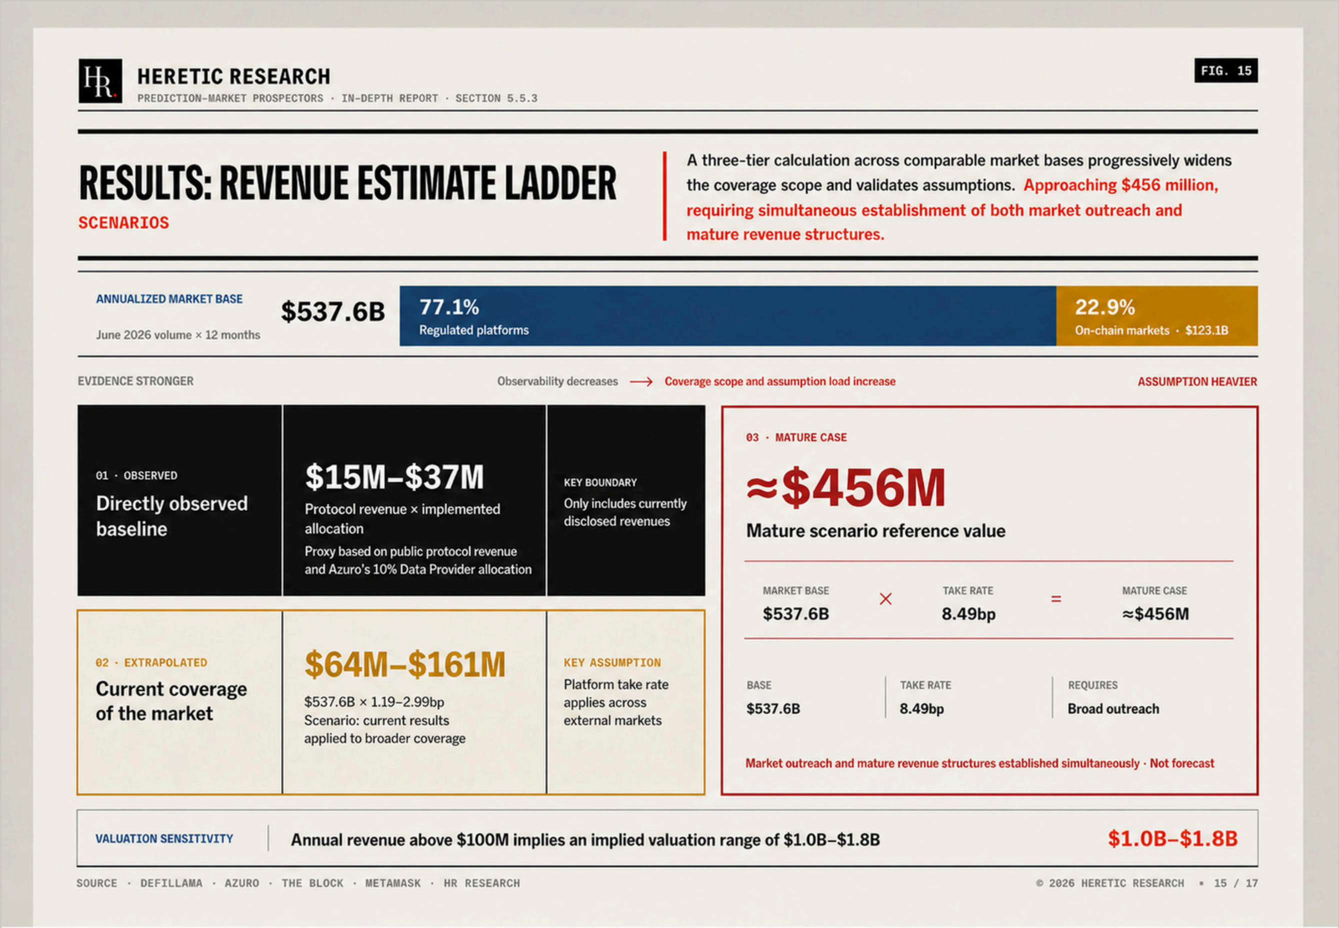

FIG. 15 — OUTCOME-LAYER REVENUE SCENARIOS

Demand, charging, and distribution already exist, but remain dispersed across platforms and suppliers. HIP-4 generates calls; Azuro produces revenue sharing; Pyth, Chainlink, and ICE occupy data and distribution interfaces. Outcome-layer industry value can now be estimated. Investment returns will depend on who can convert those calls into recurring revenue.

SECTION 06

Mapping the Industry Trend to Investable Assets

The first five chapters established why the market needs an outcome layer, how it becomes a product, and how much revenue it could generate. The next question follows naturally: even if the outcome layer has industry value, can investors capture it through existing assets?

6.1 Prediction-Platform Equity Includes Outcome-Layer Capabilities, but Is Not the Most Direct Instrument

The most obvious route is platform equity. Yet prediction-platform valuations include much more than the outcome layer: trading traffic, licenses, liquidity, distribution, brand, and expectations for prediction-market growth.

Two major valuation anchors are observable. Kalshi completed a new financing round in May 2026 at a $22 billion valuation. Polymarket was valued at approximately **$8 billion pre-money** when ICE invested in October 2025, followed by another **$600 million** direct investment from ICE in March 2026.

Platform equity provides exposure to prediction-market growth, but not a pure outcome-layer instrument. At current valuations, investors are paying primarily for trading growth, licenses, and network effects—not for a separately validated outcome-layer asset.

Our conclusion is therefore explicit: **do not chase platform equity at current prices.**

6.2 Investment Threshold: Cross-Platform Calls and Attributable Revenue

Four positions have emerged around the outcome layer. Ranked from highest to lowest by the directness of revenue attribution, they are:

- **deterministic data and machine confirmation** – Pyth and Chainlink;
- **subjective outcomes and dispute handling** – dispute mechanisms such as UMA and Kleros;
- **application distribution and workflow entry points** – distribution interfaces such as DFlow, agent interfaces, and specialist rule tools, referred to in this report as Builder applications;
- **institutional data standardization and distribution** – ICE, discussed separately in Section 6.4.

The first two directly hold outcome-layer revenue attribution. The third depends on growth for monetization. The fourth already appears in a public company's financial statements.

Private-market research should prioritize two types of project: those callable across platforms rather than dependent on one exchange, and those able to convert calls into ARR, traffic fees, or data revenue sharing. The final step in establishing an asset is revenue attribution. The easiest place to test that standard is HIP-4, which has volume, a fee schedule, and a token.

6.3 HYPE: Exposure to Platform Growth, but HIP-4 Does Not Yet Make a Material Fee Contribution

HYPE is among the most tradable related assets, but represents the growth of Hyperliquid as a whole, not pure outcome-layer revenue. The key question is not HIP-4 volume itself, but the share of total Hyperliquid fees that HIP-4 generates.

> HIP-4 fee contribution = HIP-4 fees ÷ total Hyperliquid fees

At **$182 million** in latest-30-day volume and an assumed **4-7 bp** fee rate, HIP-4 would generate approximately **$72,800-$127,400** in monthly fees. Hyperliquid's total fees over the same period were approximately **$57.5 million**, putting HIP-4's contribution at just **0.13%-0.22%**.

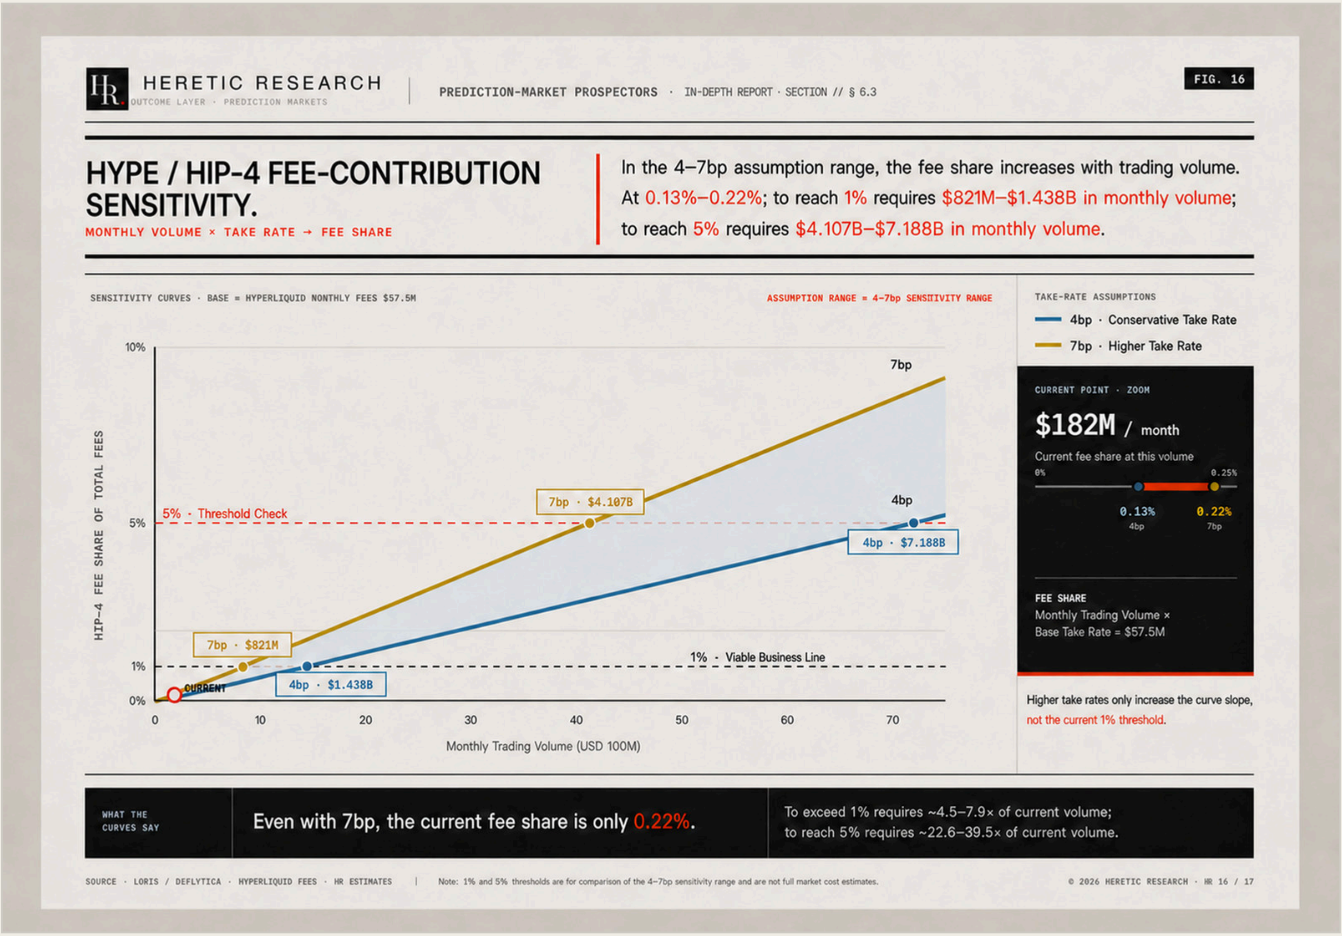

FIG. 16 — HIP-4 FEE CONTRIBUTION

That share shows HIP-4 has validated product demand but remains too small to create an independent valuation increment for HYPE. Below 1%, it is closer to a product option; sustained above 1%, it begins to form a visible business line; above 5%, it could materially affect the HYPE valuation case.

HIP-4 matters because the fee pathway already exists. As outcome markets expand and volume shifts from event peaks toward steadier everyday demand, that revenue may gradually enter protocol fees.

6.4 ICE: Financially Transparent, but Outcome-Layer Exposure Is Too Small

ICE is the easiest public prediction-market asset to verify. It owns Polymarket equity and has integrated Polymarket data into institutional feeds and historical databases, allowing both equity value and data-distribution value to reach public financial statements.

As of March 31, 2026, ICE's Polymarket Series D/E preferred shares had a carrying value of approximately **$2 billion**, representing about **23%** of issued shares and **14%** on a fully diluted basis. In the first quarter of 2026, ICE recognized roughly **$389 million** in non-cash fair-value gains from observable price changes.

ICE offers clear verification. Relative to ICE's overall business, however, this exposure is still too small to change earnings or valuation materially. ICE is better treated as a window into whether outcome-layer commercialization enters institutional budgets and public statements than as a high-beta investment vehicle.

6.5 Early-Stage Projects and Infrastructure Protocols: Closest to the Outcome Layer, but Revenue Attribution Remains Unclear

Compared with HYPE and ICE, oracles, dispute-resolution systems, distribution interfaces, and specialist Builders sit closer to the outcome layer itself. Pyth and Chainlink provide deterministic data and machine confirmation; UMA and Kleros address subjective outcomes and disputes; projects such as DFlow connect platforms, agents, and application workflows.

These projects are more focused and more likely to capture cross-platform calls directly. Yet most face the same problem: product usage does not mean revenue accrues reliably to equity or tokenholders. The number of partner platforms, supported markets, and calls proves demand, but not investment value by itself.

Two indicators therefore determine whether such an asset is investable: whether it can be called repeatedly by multiple markets without depending on one platform, and whether those calls convert into ARR, traffic fees, or data revenue sharing attributable to the relevant asset.

The opportunities closest to the outcome layer remain concentrated among these early projects, but the revenue loop is incomplete. Whoever first converts cross-platform calls into attributable recurring revenue could become one of the first true outcome-layer assets.

6.6 Current Conclusion: The Opportunity Is Forming, but the Ideal Asset Has Not Yet Appeared

The outcome-layer industry thesis is progressively taking shape, but investable assets remain early. Platform equity is expensive; HIP-4 does not yet make a material fee contribution; ICE lacks meaningful upside sensitivity; and projects closest to the outcome layer still await proof of cross-platform calls and revenue attribution.

The priority is therefore not to force a "pure-play outcome-layer asset," but to track which project first meets both conditions: providing rules, evidence, and outcome-confirmation capabilities to multiple platforms; and converting those calls into recurring revenue attributable to equity or tokenholders.

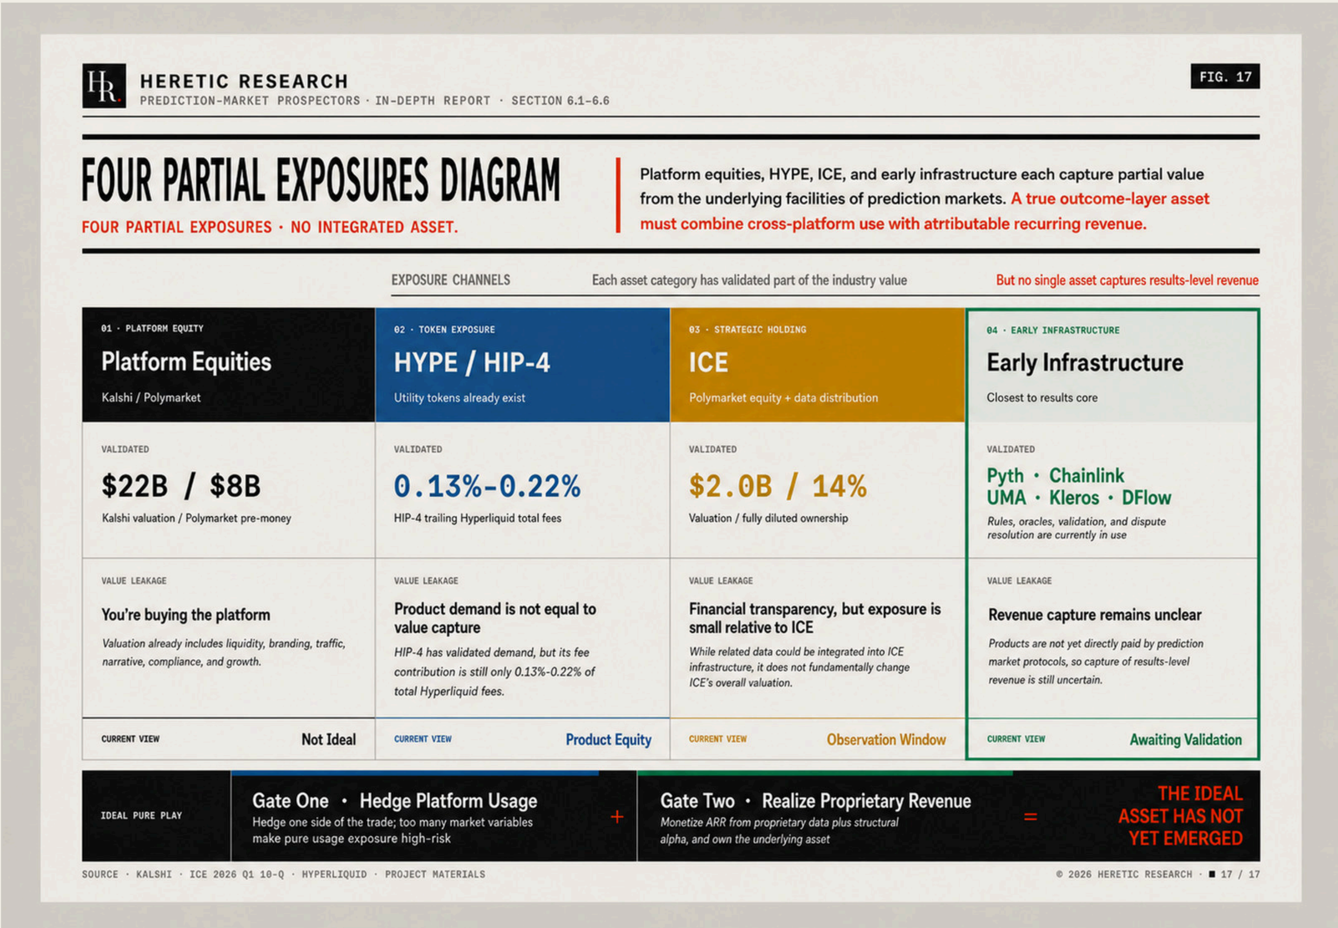

FIG. 17 — OUTCOME-LAYER INVESTMENT MAP

SECTION 07

Conclusion: The Earliest Alpha Often Appears Before the Asset Takes Shape

Prediction markets' rapid growth has so far occurred at the front end: more platforms, more categories, and competition centered on traffic, licenses, and liquidity.

As the market expands, however, the costs of rule interpretation, evidence verification, and outcome confirmation rise with it. Building a separate process at every platform duplicates investment and produces incompatible standards. Capabilities once scattered inside platforms are therefore becoming suitable for standardization and shared use across markets.

Different platforms must read the same facts, process similar evidence, interpret comparable rules, and bear the same risk of erroneous settlement. Today these capabilities remain dispersed among oracles, dispute mechanisms, data providers, and platform operations teams. Once standardized, they could become infrastructure procured by many markets, like market data, clearing, and payments.

That is why the outcome layer matters. It need not compete for every trader anew. As the number of markets, trading categories, and settlements grows, it serves demand that already exists. Front-end competition determines where traffic goes; the outcome layer becomes a service that traffic must call regardless of the platform it passes through.

The market has no clear leader yet, and revenue does not reliably accrue to any single asset class. Yet the most valuable stage of an industry opportunity often comes before the landscape is settled—when demand begins appearing across platforms, product boundaries are taking shape, and revenue attribution remains underpriced.

The signals to track are therefore who begins serving multiple platforms, who turns rules, evidence, and outcome confirmation into standardized products, and who first converts those calls into recurring revenue.

The first phase of prediction markets belonged to the expansion of trading access. In the next phase, more value will flow to infrastructure that does not own the user but determines how the market completes its final payment.

Heretic Research will continue tracking the migration of rule, evidence, and outcome-confirmation capabilities from internal platform functions to industry infrastructure, with particular attention to cross-platform adoption, charging models, and revenue attribution. We will update this research as a clearer industry structure and investable assets begin to emerge.