

Prediction Laundering: The Sociotechnical Production of Truth on Polymarket

Yasaman Rohanifar¹, Syed Ishtiaque Ahmed¹, Sharifa Sultana²

¹Computer Science, University of Toronto

²University of Illinois Urbana-Champaign

yasamanro@cs.toronto.edu, ishtiaque@cs.toronto.edu, sharifas@illinois.edu

Abstract

The growing reliance on decentralized prediction markets as epistemic infrastructures has positioned platforms like Polymarket as providers of objective, real-time probabilistic truth. However, the authoritative signals produced by these platforms often obscure the underlying uncertainty, strategic manipulation, and capital asymmetries inherent in their production, leading to an inappropriate epistemic trust by external observers. This paper presents a qualitative sociotechnical audit of Polymarket ($N = 27$), combining digital ethnography with interpretive walkthroughs and semi-structured interviews to examine how probabilistic authority is enacted and contested. We introduce the concept of “*Prediction Laundering*” to describe the sociotechnical process by which prediction markets sanitize subjective bets, strategic hedges, and capital-heavy moves into a single, authoritative probabilistic signal, obscuring the capital asymmetries and contested governance behind its production. We trace this through a four-stage laundering lifecycle: (1) Structural Sanitization, where a centralized ontology scripts the bet-able future; (2) Probabilistic Flattening, where the interface collapses heterogeneous motives into a singular percentage; (3) Architectural Masking, which utilizes selective opacity to hide capital-driven influence; and (4) Epistemic Hardening, where messy governance disputes are erased to produce an objective historical record. Our findings reveal that this process induces a state of Epistemic Stratification, where technical elites audit underlying data while the general public consumes a sanitized, capital-driven signal. We challenge the narrative of frictionless collective intelligence and conclude by arguing for a shift toward Friction-Positive Design, advocating for interventions that expose rather than erase the social and financial friction inherent in the production of synthetic truth.

Introduction

Prediction markets have gained renewed global attention as computational systems capable of aggregating dispersed beliefs into probabilistic forecasts (Hitchcock 2025; Massa and Avan-Nomayo 2025; Cross 2026). From high-stakes geopolitical conflicts to public health and economic instability, market prices are increasingly cited in media, policy discussions, and online publics as reliable, real-time indicators of future events (Miah 2025; Tapscott 2024). Platforms such as

Polymarket (Polymarket 2026b), operating via blockchain infrastructure and pseudonymous participation, are framed as uniquely credible due to their perceived transparency and decentralized financial incentives (Yahoo Finance Editorial 2025). This growing reliance positions decentralized speculation not merely as a venue for wagering, but as a primary infrastructure for public truth.

However, this growing reliance raises critical questions for sociotechnical scholarship: how do these markets reconfigure epistemic authority, and what systemic risks are hidden behind their probabilistic outputs? This question takes on urgent significance as market data moves beyond human media consumption and directly into the automated ecosystems of artificial intelligence. Major financial networks have already integrated live prediction market data into their terminal frameworks (TradingView 2024). At the same time, this data has increasingly begun appearing in everyday LLM search results and AI-driven conversational assistants, where it is presented to users as objective ground truth regarding world events (Britton 2026). When global media, policymakers, and conversational AI tools treat these speculative prices as flawless mirrors of reality, they risk institutionalizing a form of manufactured truth that launders raw capital concentration into an unearned public consensus. While prior research has focused on forecasting accuracy and market efficiency, far less attention has been paid to prediction markets as infrastructures that actively shape knowledge, belief, and responsibility. Treating probabilistic prices as neutral or objective obscures the ways in which platform design, governance rules, and social dynamics influence whose beliefs count, which futures are legible, and how uncertainty is morally framed. By centering these dynamics, we highlight how the commodification of uncertainty through such infrastructures can undermine epistemic justice and the sustainability of public discourse.

In this paper, we argue that prediction markets function through a process we term “*Prediction Laundering*”. Drawing on MacFarlane’s work on knowledge transmission (MacFarlane 2005, 2014), we define prediction laundering as the selective transformation of heterogeneous, strategic, and power-laden wagers into a single authoritative probabilistic signal. Crucially, this process retains the appearance of epistemic objectivity while obscuring socially relevant context about how that signal was produced. This includes

the erasure of strategic hedging, unequal influence of capital, and the contested processes through which outcomes are resolved.

Importantly, prediction laundering does not operate through outright concealment. While transaction data may remain publicly accessible, the laundering occurs through contextual displacement: the interface foregrounds a clean probabilistic output while backgrounding the interpretive, financial, and governance conditions that shape it. As a result, users and external observers encounter a signal that appears to represent collective belief, despite being shaped by asymmetrical power and strategic behavior. While such dynamics are not unique to markets, their consequences differ when prediction market outputs are treated as epistemic infrastructure. What distinguishes prediction markets in this context is not the presence of these dynamics, but their packaging into a single, highly legible probabilistic signal that is easily circulated and interpreted outside the market itself. These outputs act as widely consumed signals used in media, public discourse, and policy reasoning. In these contexts, the loss of context is not merely an analytical limitation but a source of potential misinterpretation and distortion.

This paper contributes to the sociotechnical AI research community by (1) conceptualizing how prediction market platforms function as epistemic infrastructures rather than purely financial tools, (2) identifying the risks introduced when financially mediated signals are interpreted as neutral indicators of reality, and (3) offering design and policy directions for improving transparency and accountability in such systems.

We investigate these dynamics through a qualitative sociotechnical audit of Polymarket ($N = 27$). Combining digital ethnography with interpretive walkthroughs, design probes, and semi-structured interviews, we sampled both domain experts and generalist users across a spectrum of information literacies. We examine how probabilistic authority is produced and contested across platform design, market mechanics, community practices, and off-platform coordination. Our research is guided by three core questions:

- **RQ1:** How do prediction market design and pricing mechanisms transform the users' contested beliefs and value-laden judgments into seemingly neutral probabilistic signals?
- **RQ2:** How do users interpret, rely on, and contest these probabilistic signals, particularly when moral discomfort, uncertainty, or exclusion arise?
- **RQ3:** What kind of risks and challenges occur for users when reality, rumors, and probabilities as numbers convolute together?

We show that Polymarket curates a narrow ontology of the future, privileging capital-driven influence over diverse belief aggregation while offloading intensive interpretive labor onto participants. Our findings trace this through a four-stage laundering lifecycle that reveals systemic failures in epistemic governance:

1. **Structural Sanitization:** The platform utilizes a centralized ontology to script which versions of the future are bet-able, pre-formatting reality for speculation.

2. **Probabilistic Flattening:** The interface collapses heterogeneous motives, such as strategic hedging and interpretive oscillation, into a singular numerical truth.
3. **Architectural Masking:** Selective opacity hides capital concentration and whale influence, creating a mirage of democratic consensus.
4. **Epistemic Hardening:** Messy governance disputes and ethical tensions are scrubbed from the final record, leaving only a de-contextualized, authoritative outcome.

By identifying the epistemic stratification and accountability gaps created by this process, this work challenges optimistic narratives of collective intelligence. We call for greater scrutiny of the computational systems mediating how synthetic truths are imagined, priced, and trusted.

Related Work

To understand the sociotechnical mechanisms of prediction markets, we must engage with literature that spans epistemology, economic sociology, and algorithmic accountability. This section establishes the theoretical foundations for our audit by examining three core areas. First, we explore the philosophy of “Epistemic Transmission” and the concept of “laundering” to understand how information is transformed as it moves between contexts. Second, we analyze the “Wisdom of Crowds” as an idealized framework for collective intelligence and contrast it with the realities of market stratification. Finally, we situate this study within the field of algorithmic accountability, drawing on sociotechnical auditing methods to investigate the governance and labor required to maintain these informational systems.

Knowledge Laundering and Epistemic Transmission

The primary theoretical anchor for this paper is the concept of “Knowledge Laundering”. This term was coined by John MacFarlane to describe a fundamental failure in the transmission of information between distinct contexts. In the framework established by MacFarlane (MacFarlane 2014, 2005), knowledge is not a fixed or invariant property. Instead, it is often assessment-sensitive, meaning its validity depends on the epistemic stakes and the specific context of the person evaluating a claim. The process of laundering occurs when a piece of information moves from a high-stakes context, where a subject might be cautious or uncertain, to a low-stakes context, where an external observer accepts the claim as a settled fact. This transition happens without the underlying uncertainty being transmitted along with the information itself.

This theory challenges the “Transmission Principle” in classical epistemology which traditionally assumes that if Person A knows p and tells Person B, then Person B also knows p (Lackey 2008). MacFarlane argues that this principle breaks down when the pedigree of the knowledge is scrubbed. If Person A only knows p relative to a very low standard of evidence but Person B accepts p as an objective and high standard fact, the knowledge has been laundered. This creates a loophole where the public can end up trusting

a claim that no single individual in the chain actually felt certain about.

In the context of Algorithmic Accountability, this framework provides a powerful lens for examining how computational systems mediate belief. While MacFarlane originally applied this to human testimony, we argue that prediction markets function as automated infrastructures for epistemic laundering. It is critical to distinguish this process from mere concealment. While blockchain infrastructures offer radical transactional transparency, they simultaneously facilitate a form of contextual opacity. The laundering occurs not by hiding the raw transaction data, but by transforming the subjective, value-laden pedigree of a wager into a de-contextualized numerical fact. By the time a high-uncertainty bet is aggregated into a market price, the strategic motives and capital-driven pressures behind it are discarded, leaving only an authoritative signal for public consumption.

Scholarship in sociotechnical systems research regarding inappropriate reliance (Metaxa et al. 2021) and algorithmic opacity (Burrell 2016) has traditionally focused on how users over-trust model outputs. By integrating MacFarlane's work, we shift the focus to the sociotechnical transformation of the signal itself. We contend that prediction markets do not just display information. They perform a sociotechnical scrubbing that converts strategic and capital driven bets into de-contextualized probabilistic truth. This process enables what we define as "Prediction Laundering", which refers to the systemic erasure of the social and financial context, including hedges and power asymmetries, from the final public signal.

Information Markets and the "Wisdom of Crowds"

The conceptual foundation for prediction markets rests on the Hayekian view of markets as a unique solution to the problem of dispersed knowledge (Hayek 1933). Hayek argued that the price mechanism allows society to aggregate fragmented bits of information that no single central planner could possess. In this view, prices function as signals that coordinate human action by summarizing vast amounts of local and specialized data. Modern proponents of prediction markets extend this logic by suggesting that these platforms are the most efficient tools for "discovery" in uncertain environments. Wolfers and Zitzewitz formalize this in their analysis of prediction markets in theory and practice, stating that the financial incentives of the market force participants to reveal their true beliefs (Wolfers and Zitzewitz 2004). They argue that this makes market prices more reliable than polls or expert opinions because the "skin in the game" requirement filters out insincere or low-quality information.

Standard economic interpretations of prediction markets are primarily rooted in the "Wisdom of Crowds" hypothesis. This framework suggests that a large group of independent actors will collectively produce a more accurate forecast than any individual expert (Surowiecki 2005). The core assumption is that a market mechanism can effectively filter out individual biases and noise by providing financial incentives for accuracy. Wolfers and Zitzewitz argue that this process makes prediction markets uniquely efficient at infor-

mation aggregation because participants are forced to back their beliefs with capital (Wolfers and Zitzewitz 2004). In this idealized model, the market price is viewed as a neutral and mathematical consensus that reflects the total sum of available public and private information.

Beyond the aggregation of existing information, we must account for the performativity of markets as theorized by scholars such as Callon (Callon 1998) and MacKenzie (MacKenzie 2008). Rather than acting as a camera that passively records a pre-existing reality, a prediction market functions as an engine that actively shapes the world it predicts. By generating a specific probabilistic signal for a geopolitical or social event, the platform influences public sentiment, media narratives, and the strategic behavior of the actors involved. In this sense, the truth produced by Polymarket is not merely discovered. It is sociotechnically constructed through a process that can alter the very outcomes it claims to forecast.

Sociological critiques of finance suggest that markets are better understood as specialized micro-structures rather than broad and egalitarian crowds. For example, Knorr Cetina notes that global financial markets are deeply stratified environments where information is mediated by social networks and technical expertise (Knorr Cetina and Bruegger 2002). In these settings, a small subset of participants with high capital and specialized platform knowledge exerts a disproportionate influence over the resulting signal, creating a state of specialized competition where the price does not reflect a collective belief but rather the strategic moves of a few dominant actors. We build on these critiques by examining how the "crowd" in decentralized platforms like Polymarket is structured. We focus on the gap between the platform's presentation as an open and egalitarian marketplace and the reality of its capital-driven hierarchy.

Algorithmic Governance and the Erasure of Context

Scholarship in critical data studies suggests that data is never raw but is always cooked through specific technical and social choices (Gitelman 2013). We build on this by examining how the curation of binary contracts acts as a form of epistemic enclosure. By pre-formatting complex social realities into bet-able outcomes, platforms narrow the scope of what is knowable to only what is financially tradable. This process aligns with broader trends in algorithmic governance where computational systems prioritize legibility and scale over local nuance and situated knowledge.

This governance is maintained through intensive human labor that is often rendered invisible by the platform interface. Drawing on the concept of "Broken World Thinking" developed by Steven Jackson, we argue that the credibility of prediction markets often depends on constant repair work performed in off-platform and community-led spaces such as Discord and Telegram (Jackson 2014). While these platforms are marketed as automated, their governance is characterized by resource dependencies and asymmetrical power relations between platform owners and users. The platform captures the value of a "clean" probabilistic signal while

offloading the “dirty” work of dispute resolution and fact-checking onto a pseudonymous public.

To investigate these dynamics, we follow the tradition of auditing from the outside (Sandvig et al. 2014). This approach uses qualitative user experiences to reverse-engineer the social consequences of opaque systems, acknowledging that knowledge is inherently situated in people and locations. It recognizes that the wisdom produced by a platform is inseparable from the social dynamics of the community that generates it. By focusing on the lived experiences of participants, we can observe the emotional and cognitive negotiations that are typically rendered invisible by raw transaction data and market charts.

Finally, we address the risk of inappropriate reliance where external observers or policymakers treat market outputs as de-contextualized facts (Metaxa et al. 2021). This reliance is encouraged by a platform design that prioritizes cognitive ease and social proof at the expense of epistemic rigor. As Ananny and Crawford argue, transparency alone is insufficient for accountability if the underlying logic of the system remains opaque (Ananny and Crawford 2018). By identifying the mechanisms of Prediction Laundering, we bridge the gap between platform studies and financial psychology. We argue that the lack of accountability in these systems is not a technical bug but a structural feature that allows speculative infrastructures to operate independently of traditional status markers or ethical constraints.

Background

Markets have long been understood in sociotechnical research as complex assemblages embedded in cultural, social, and material practices rather than mere arenas of economic exchange. Traditional and informal markets unfold through situated practices of negotiation, trust, and testing that constitute them as lived social worlds rather than mere price-setting mechanisms. Recently, decentralized prediction platforms like Polymarket have emerged, financializing future uncertainty into tradable outcome contracts. To analyze the production of truth on Polymarket, we must first understand the technical and institutional mechanisms that govern the platform. Unlike traditional centralized betting sites, Polymarket utilizes a hybrid architecture that combines off-chain user interfaces with on-chain settlement. This section details the market mechanics, the hybrid order book, and the decentralized oracle system that collectively enable the transformation of subjective wagers into authoritative probabilistic facts.

Prediction Market Fundamentals

Polymarket achieves information aggregation through the trading of event contracts or outcome tokens. Each contract pays 1.00 USD if a specified event occurs and 0.00 USD otherwise. While many markets are structured as simple Yes or No outcomes, others allow participants to trade across a range of mutually exclusive categorical results. The trading price of a share fluctuates between 0.00 and 1.00 USD and acts as a collective assessment of probability. For instance, a contract trading at 0.75 USD signals a 75 percent market-estimated likelihood of the event. This financial incentive

structure is designed to reward accuracy and penalize misinformation. When the event is resolved, shares in the correct outcome are redeemable for 1.00 USD while all other shares become worthless.

Hybrid Architecture and the CLOB

Polymarket operates through a hybrid-decentralized model to ensure technical efficiency. While the final execution of trades and the custody of funds occur on the Polygon blockchain, the platform utilizes an off-chain Central Limit Order Book, or CLOB, to match buy and sell orders. This design allows for a frictionless trading experience by removing the need for users to pay blockchain gas fees for every individual order placement or cancellation. However, this architectural choice also creates a significant sociotechnical gap. The front-end interface presents a singular price to the user, yet this number is the result of an underlying flow of high-frequency capital and strategic hedging that remains largely obscured from the casual observer.

The Resolution Lifecycle and the UMA Oracle

The final stage of any market is its resolution, which is managed through a decentralized governance mechanism known as the UMA Optimistic Oracle. When a market reaches its expiration date, a human proposer submits what they believe to be the correct outcome along with a financial bond. This starts a challenge period during which any other participant can dispute the proposal by posting their own bond. If no dispute occurs within the allotted time, the outcome is finalized and the market settles on the blockchain. If a dispute arises, the governance process moves to a vote by UMA token holders. This resolution lifecycle relies on a combination of automated smart contracts and human consensus to settle outcomes and distribute payments.

Methods

We conducted a qualitative sociotechnical audit of Polymarket to examine how probabilistic authority is enacted, interpreted, and contested through market rules and community practices. Polymarket was selected as the primary site of study because it is the leading global, blockchain-based prediction platform. Given the platform’s complexity and pseudonymous culture, we relied on guided user sessions to capture participant experiences without requiring actual monetary trades.

Multi-Modal Qualitative Approach

We conducted a four-month qualitative study to examine both the technical features and social dynamics of the platform. We structured our study to connect macro-level platform observations (such as trading trends and community discussions) with micro-level participant interviews about trust, power, and market influence. The study began with phase 1, which focused on a Digital Ethnography where we conducted systematic observation of platform dynamics across the primary interface and its peripheral social ecosystems. This phase consisted of three sub-components including a technical audit of Polymarket affordances, the monitor-

Total: 27 (Female: 11, Male: 16)				
Age Range	Education		Crypto Familiarity	Expertise
20-29: 8	Highschool: 1		Novice: 10	Domain Expert: 10
30-39: 9	Undergraduate: 12		Exchange-Reliant: 11	Generalist User: 17
40-49: 5	Masters: 9		Crypto-Native: 6	
50+: 5	Doctorate: 5			

Table 1: Demographics of Participants

ing of social shadow infrastructures on X, Reddit, and Discord, and the tracking of high-capital whale movements via block explorers. In this stage, we adopted a non-participant observer role to archive field notes, screenshots, and transaction data without intervening in the social dynamics.

Following the ethnography, we conducted one-on-one sessions with the full sample of participants ($N = 27$) in phase 2. We structured each session around three complementary methods to thoroughly evaluate participant interactions and reasoning. The engagement began with interpretive walkthroughs using think-aloud protocols while navigating real market interfaces. This stage focused on immediate and visceral reactions while capturing the real-time reasoning process as users encountered specific event contracts and price fluctuations.

The session then moved into semi-structured interviews investigating epistemic entry points. Rather than focusing solely on historical trust, these interviews explored the specific technical, social, and aesthetic cues participants use to evaluate the legitimacy of decentralized information. We documented the criteria users employ to transition from passive observers to active interpreters of speculative data, focusing on how they bridge existing knowledge with the probabilistic outputs generated by the system.

The session concluded with design probes, where participants reviewed and responded to specific artifacts captured during the digital ethnography. We presented participants with archived comment threads, data visualizations of whale positions, and ethical vignettes depicting real-world trading scenarios. Using these authentic platform artifacts, we encouraged participants to reflect on how they navigate social and technical friction when evaluating market information. This two-phase approach ensured that findings from our macro observations directly informed our individual participant sessions.

Participants and Sampling

We employed a hybrid purposive and snowball sampling strategy to recruit twenty-seven participants ($N = 27$) representing a diverse range of backgrounds and technical proficiencies. To ensure the audit captured a wide spectrum of information literacies, we seeded our recruitment through professional and academic networks, targeting individuals with varying degrees of exposure to financial forecasting and decentralized technologies. We categorized these participants into two primary groups based on their professional and academic backgrounds. “Domain experts” ($n = 10$) included individuals with background in computer science, economics, or data science, allowing them to technical audit

the platform’s blockchain mechanics and algorithmic logic. Conversely, “generalist users” ($n = 17$) included individuals working in non-technical fields like education, health-care, and the arts, capturing how everyday consumers interact with the platform without specialized background in finance or crypto.

We also accounted for participants’ prior crypto experience, as the ability to navigate blockchain tools significantly influenced how much control and insight a user had on the platform. We define “crypto-native” participants as those who regularly interact with self-custodial wallets and decentralized protocols, possessing the literacy to audit transactions via block explorers. “Exchange reliant” users are those familiar with centralized cryptocurrency exchanges who understand the value of digital assets but lack experience with on-chain governance or smart contract interaction. Finally, “novice” participants include individuals with little to no prior exposure to cryptocurrency, for whom the platform’s blockchain-based settlement remains an opaque black box. By organizing our sample across these dimensions, we observed how power shifts between technical insiders who can verify raw data and casual users who must trust the main interface.

Following these initial sessions, we utilized snowball sampling (Goodman 1961; Noy 2008) to reach participants who occupy specific outsider niches, such as individuals with strong moral or religious objections to speculative markets. This approach allowed us to move beyond a convenience sample and intentionally include voices that are systematically excluded or marginalized by the platform’s technical jargon and complex crypto infrastructure. By including these epistemic outsiders, we were able to document the barriers to entry and the moral friction that characterize the experience of those for whom the market’s logic is not inherently intuitive. This inclusive sampling strategy ensures that our audit accounts for the lived experiences of marginalization and exclusion that often remain invisible in purely quantitative analyses of market efficiency.

Analysis Procedure

All sessions were recorded, transcribed, and analyzed through iterative thematic coding, moving from initial open codes to more structured conceptual categories. Our analytical framework was primarily abductive, meaning we treated moments of breakdown, irony, and contestation as critical entry points into the ways participants and traders make sense of uncertainty. Rather than seeking a unified user experience, we intentionally focused on the friction between the platform’s technical affordances and the users’ situated




Figure 1: Methodological workflow and data integration diagram illustrating our two-phase research design moving from macro-level platform ethnography to micro-level participant sessions. The ethnographic data archive from phase 1 directly informs the design probes in phase 2 to enable a confrontation between users and real-world platform artifacts.

knowledge.

This approach allowed us to identify the specific patterns of sociotechnical mediation that facilitate the translation of subjective belief into probabilistic signals. By centering power relations, we interrogated how the platform’s interface and governance structures privilege certain forms of expertise while marginalizing others. This method of “critical listening” ensured that our findings reflect not just the stated functionality of the platform, but the lived experiences of those navigating its speculative infrastructure. Ultimately, our analysis sought to connect individual moments of confusion or skepticism to broader systemic accountability gaps, providing a grounded basis for our later theorization of the platform’s epistemic outcomes.

Findings: The Life Cycle of a Laundered Prediction

Our sociotechnical audit reveals that Polymarket’s primary sociotechnical effect is the systematic “laundering” of subjective, value-laden, and often adversarial human inputs into objective-looking probabilistic signals. While the platform’s interface presents a veneer of mathematical transparency, this clarity is achieved through a series of structural and social filters that strip away the “dirty” context of individual wagers, including financial hedging, coordinated manipulation, and capital-driven bias. We trace this process through a “Laundering Lifecycle” and demonstrate how the platform’s mechanics consistently privilege capital and technical literacy over democratic information aggregation. By examining the transition from market creation to final settlement, we identify the specific points where the system collapses complex social disagreement into a sanitized numerical truth, effectively insulating the resulting data from accountability or

ethical interrogation.

Stage 1: Structural Sanitization (The Curated Ontology)

The first stage of prediction laundering occurs before any wager is placed, at the point where a real-world event is translated into a tradable contract. Our audit reveals that Polymarket does not function as a decentralized marketplace for all possible futures, but instead enforces a rigid, centralized ontology. While the platform encourages “market suggestions” via Discord and X, users cannot directly create or list their own markets (Polymarket 2026a). Instead, a dedicated “Markets Team” acts as a primary epistemic filter, vetting proposals based on their “news value”, trading demand, and technical resolve-ability. This structural gate-keeping functions as a form of epistemic enclosure. By pre-formatting reality into a curated set of tradable contracts, the platform decides which versions of the future are permitted to be quantified. One expert participant noted that this suggestion-based model creates a facade of community participation while maintaining strict institutional control:

“I think people forget that Polymarket’s design actually shapes what you’re even allowed to bet on. On paper, “anyone” can suggest creating a Polymarket! But in reality, the moderators have to approve the titles and how it’s all going to be settled. If it’s too vague, like ‘Will the new iPhone be a success?’ they’ll reject it. Same for subjects that are too weird or [morally] bad... like markets on someone getting hurt. The result is this curated set of Polymarkets that they allowed for us... basically just the stuff that are measurable, legible, and fit their brand!” (P3, Male, Crypto-native, Domain Expert)

By filtering out the ambiguous or the socially unacceptable, the platform creates an epistemic boundary that privileges data-centric outcomes over complex, non-binary social realities. This curation suggests that the “wisdom” produced by the market is not a reflection of universal human speculation, but a subset of possibilities that have been preformatted for blockchain settlement. In the forthcoming discussion, we will further argue that the laundering process begins with this initial curation, where the platform decides which versions of reality are “fit” to be quantified.

Stage 2: Probabilistic Flattening (The Collapsing of Belief into Price)

Once a market is live, the platform utilizes a pricing mechanism that acts as an “epistemic solvent” by dissolving the specific context and motivations behind individual wagers. While participants engaged with the platform for very different reasons, ranging from systematic data analysis to personal bias, the interface boiled these distinct motivations down into a single percentage. This mechanism collapses disagreement and moral contestation into a single signal that circulates as objective knowledge, effectively masking the messy reality of how that number was reached. One participant highlighted the shift in perception that happens when personal opinions are turned into clean percentages:

“It’s funny because once I see a number attached to a possible outcome, it stops feeling like an opinion and more like... real truth. It doesn’t matter who is betting or why. You just see 72 percent and it feels like that is the truth, even though I know it is coming from... just a mess of guesses and incentives!” (P9, Female, Exchange-Reliant, Generalist User)

This flattening process often leads to a state of “interpretive oscillation” in which users fluctuate between treating the market as a credible tool and dismissing it as a gamified environment. Despite the interface’s polished presentation, participants frequently questioned the underlying validity of the numbers, wondering if price movements stemmed from real information or speculative hype. One user expressed this confusion when reconciling the probabilistic framing with the underlying market mechanics:

“I’m looking at all these percentages and I know it’s supposed to be the probability... but there is this other language. The percentages don’t add up. How much fee is taken by the platform? Is this number even realistic or just betting hype? What triggers the number to change? Like... does someone input it? I feel like I’m at a casino and everything is staged to make me bet!” (P12, Male, Novice, Generalist User)

The solvent effect is further complicated by “the hedge problem” or the platform’s inability to distinguish user intent behind capital deployment. Participants who bet purely as financial insurance appear indistinguishable from those expressing strong ideological belief. As a result, strategic hedging is laundered into apparent consensus, masking a messy mix of risk-management tactics behind a clean probabilistic figure. Another participant described this gap be-

tween personal belief and financial strategy when hedging against their own preferred outcome:

“If I were to bet on this election outcome, I’d bet against my own candidate just to soften the blow... if they lose, at least I get a payout! The problem is that other users see that I’m predicting a loss, when I’m actually just buying peace of mind... I wish there was a way to show which bets come from real belief and which are just financial hedges. These markets don’t show if a politician is actually popular, but the media uses these numbers as the absolute truth!” (P4, Male, Exchange-Reliant, Domain Expert)

The transition from a mess of guesses to a singular percentage represents a fundamental loss of information. The platform strips away the “why” of the bet, leaving only the “what”, which effectively misleads the observer into believing the price reflects a consensus of belief rather than a cluster of financial strategies. In our discussion, we explore this through the lens of “epistemic dilution”, where the market’s need for a singular signal inevitably produces a distorted representation of public sentiment. Ultimately, the platform erases all traces of hedging and doubt, presenting users on Polymarket with nothing more than a single percentage.

Stage 3: Architectural Masking (The Illusion of Crowdsourced Truth)

Following the consolidation of price, the third stage of laundering involves obscuring the specific actors responsible for market movements. Although prediction markets are frequently framed as “democratic aggregators”, our audit reveals how participants perceive and navigate the ways epistemic influence is disproportionately asserted through concentrated capital. The platform’s pseudonymity hides the identity of whales or large-scale traders whose massive positions can single-handedly shift the market, allowing their private financial interests to be laundered into the appearance of a broad, neutral consensus. This structural invisibility makes every signal look like a broad consensus, regardless of whether it reflects many participants or just a few large wallets. One participant noted how the weight of capital often overrides the wisdom of the general public:

“Obviously some people move the market way more than others. Just like real life, if you have money... or if you really care about a topic to lobby on it, or if you have insider information... your belief counts for more. Everyone else is kind of just reacting to that, not really contributing. The whale moves the price and then the crowd just follows along with that momentum... even if they don’t know why it happened.” (P17, Female, Exchange-Reliant, Generalist User)

This asymmetry is maintained through a condition of “selective opacity”. While the primary Polymarket interface is clean and simple, it hides the complex on-chain data that reveals who is actually moving the money. We observed a significant literacy gap between crypto-native participants, who utilized third-party block explorers to audit whale behavior, and novice users, who were forced to take the front-

end price at face value. This gap creates a two-tiered epistemic environment where the reality of market manipulation is legible only to technical insiders. A crypto-native participant described the necessity of looking under the hood to find the truth behind the price:

“To gain the most, you have to follow the whales. I’ll just go to Polymarket leaderboards, Etherscan or Arkham... [third-party tools] to see who is actually dumping millions into a vote. If you just look at the Polymarket homepage, it looks like many people are agreeing... but it’s only when you dig deeper that you realize it is just one massive wallet!” (P21, Male, Crypto-Native, Domain Expert)

This selective opacity allows capital-driven signals to be laundered into the appearance of democratic agreement. By hiding the architects of the price movement, the platform prevents the average observer from questioning the motives behind a shift in odds. This dynamic produces power distortion, where concentrated financial influence is laundered into the appearance of distributed consensus. In our discussion, we connect this to the theory of “Epistemic Stratification”, where the ability to see through the laundry is a privilege reserved for those with specific technical and financial literacy, while the general public is left to consume the sanitized, consolidated output as fact.

Stage 4: Epistemic Hardening (The Oracle Fallacy)

The final stage of prediction laundering occurs during market resolution, where the “hardening” of a result erases the history of social, moral, and financial friction. Ultimately, this sanitization process strips the moral gravity from real-world events, framing them as plain topics for estimation. Converting human tragedies (like war or political instability) into market odds allows participants to view real-world suffering through a detached, analytical lens. This objectification of harm normalizes speculation while insulating both the platform and the trader from ethical accountability. This hardening contributes to an accountability gap, where contested and strategic resolution processes are erased from the final record, obscuring responsibility for the outcome. One participant described the discomfort of profiting from disaster and how the platform’s design mitigates that guilt:

“There is a weird discomfort with betting on people dying in a war... but when you look at the chart, it just looks like any other stock market chart! The way all of the information is framed makes it feel like you are solving a puzzle rather than profiting from a disaster!” (P2, Female, Exchange Reliant - Generalist User)

By transforming a moral crisis into a speculative puzzle, the platform performs a form of moral laundering. It takes the weight of a sensitive ethical choice and places it on the mechanical neutrality of the market. This hardening ensures that once the oracle clears, the messy human cost of the wager is completely obscured. As another participant noted, the platform effectively absorbs the user’s moral burden: “it’s like the platform will take the responsibility of being a bad person from you and place it on itself. You just do the betting!”

This moral distancing is paired with a tactical struggle to control the platform’s resolution mechanisms. While the platform markets itself as a decentralized system, the actual process of resolving the truth relies on a stratified financial hierarchy. When a market outcome is ambiguous, the messy human labor of dispute resolution takes place in off-platform social silos like Discord, often involving tactical deception and succumbing to exit strategies. Once the market is settled by the oracle, however, this friction is scrubbed from the user interface, presenting the final result as an objective fact. One domain expert described the fundamental dishonesty that often precedes the hardening of a market signal:

“Polymarket is just a zero-sum game... and because of that, the comment section is probably one of the most intellectually dishonest places on the internet!! When you see someone being helpful or sharing a link, they almost always have a financial incentive to make you believe them. They want you to take the other side of their trade so they can get exit liquidity... they’ll spam the record with news clips just to influence how the oracle eventually settles the truth.” (P1, Female, Crypto-Native, Domain Expert)

This tactical labor is eventually buried beneath the platform’s formal resolution mechanics, which are shielded from the average user by high financial barriers. A participant with a technical background expressed concern over how this hardening process is governed:

“On paper, markets get resolved in a democratic way... but you have to stake a 750 USD bond just to propose an outcome. Most people don’t have that... you need capital just to have a say in how the truth is settled. Then, if you want to challenge, it’s another 750 USD. People with the UMA tokens vote on it... so if you are wealthy, you can basically buy enough tokens to outvote everyone and make sure the market settles in your favor!” (P16, Female, Exchange-Reliant, Domain Expert)

The final “cleansing” of the record occurs at the moment of settlement. Another participant noted how the interface removes all traces of this chaotic behind-the-scenes labor once a result is reached:

“If a resolution is confusing... everyone runs to Discord to argue with the mods. It’s total chaos behind the scenes. But then the Oracle clears... and all that mess is gone. On the main page, all that is left is a ‘Yes’ or ‘No’... like there was never any doubt at all.” (P25, Male, Crypto-Native, Domain Expert)

The “Oracle Fallacy” lies in the public perception of resolution as a neutral, data-driven event. In reality, it is an epistemic settlement achieved through a contest of capital and strategic noise. The interface erases the chaotic behind-the-scenes labor once a result is reached, presenting the final output as a de-contextualized fact.

Discussion

Our findings demonstrate that Polymarket’s epistemic authority is not a byproduct of neutral information aggrega-

tion, but the result of a deliberate sociotechnical scrubbing process we have termed “*Prediction Laundering*”. It is important to note that statistical accuracy is not the same as transparency. A market can accurately predict an outcome while disguising the motives and capital behind the price. By tracing the lifecycle of a wager, we have shown how the platform systematically strips away the “dirty” context of individual motives such as financial hedging, capital-driven manipulation, to produce a sanitized, authoritative signal. This transformation creates a profound tension between the platform’s front-end veneer of “crowdsourced truth” and the back-end reality of capital concentration and technical gatekeeping. In the following sections, we discuss how this process induces a state of “Epistemic Stratification” and creates significant accountability gaps, ultimately challenging the idealized “wisdom of crowds” narrative that currently dominates the discourse on prediction markets.

The Asymmetric Transparency of Synthetic Truth

One of the cores promise of prediction markets is to solve the problem of “expert failure” by democratizing information aggregation. However, our findings suggest that Polymarket’s architecture does not distribute epistemic power but merely shifts it from traditional institutional experts to a new, invisible class of capital-rich technical elites. This results in what we call “*Epistemic Stratification*”, where the platform produces a synthetic truth that appears democratic on the front-end but is technically and financially gated on the back-end.

This stratification is a direct consequence of architectural masking. By design, the interface offers a flattened, highly accessible view of reality for everyday users, while masking the audit trail of capital influence beneath the complex layers of the blockchain. As we observed in Stage 3, the transparency of the system is asymmetric. While the blockchain is technically open for anyone to see, the literacy required to decipher “whale” movements or identify “wash trading” creates a functional barrier to entry. This mirrors what sociotechnical scholars describe as “professionalization”, where a supposedly public resource (in this case, “the truth” about the future) is captured by specialists who possess the tools to see through the “laundered” signals.

This capture fundamentally undermines the “Wisdom of the Crowd” theory. If the majority of users are not contributing independent beliefs but are instead reacting to the momentum of a few significant wallets, the market is no longer an aggregator of diverse perspectives. Instead, it becomes a feedback loop of capital. In this environment, the “consensus” is not discovered, but manufactured by those with the financial resources to move the price and the pseudonymity to hide their motives. By laundering this capital-driven influence into a clean, probabilistic percentage, the platform grants “whale” interests the unearned authority of a democratic mandate.

Consequently, the resulting data is not a neutral reflection of public sentiment, but a laundered signal that has been stripped of its financial origins. When policymakers or news organizations cite these percentages as “objective” forecasts, they are inadvertently participating in the final stage of the

laundering process. They are granting institutional legitimacy to a signal that was forged in a state of extreme capital asymmetry, effectively allowing wealth of a few to masquerade as collective intelligence.

The Accountability Gap: Laundering Moral and Institutional Responsibility

A critical consequence of the laundering life cycle is the creation of an accountability gap. By the time a prediction reaches its final “hardened” state, the platform’s architecture has successfully decoupled the resulting data from the human and ethical costs associated with its production. This process mirrors the final stage of financial laundering integration, where the “cleaned” product is merged into the legitimate economy. In the context of prediction markets, this integration allows both the platform and its users to evade moral and institutional responsibility by attributing controversial outcomes to the “objective” mechanics of the market.

This gap is primarily sustained through “moral displacement”, a phenomenon we observed in Stage 4. When ethically sensitive events, such as military conflicts or political instability, are re-framed as speculative puzzles, the interface performs a form of ethical sanitization. The “laundered prediction” allows the user to profit from tragedy without the social stigma of being a disaster speculator. As our participants noted, the platform “takes the responsibility of being a bad person” away from the individual, effectively commodifying human suffering into a neutral data point.

Further, the “oracle fallacy” provides the platform with institutional deniability. By offloading the “dirty labor” of truth-resolution to decentralized oracles like UMA, the platform can claim it has no hand in the final result. However, our audit shows that this “decentralization” is often a mirage. The high financial barriers, such as the \$750 bonds, ensure that only a specific class of actors can influence the final settlement. When a dispute arises, the platform can point to the “code” or the “oracle” as a neutral arbiter, even if the resolution was determined by capital-weighted voting. This creates a state of institutional evasion, in which the platform takes credit for the “accuracy” of its signals while distancing itself from the contested and often biased methods used to produce them.

This lack of accountability has implications for how these signals are consumed by the public. When news organizations treat a “hardened” Polymarket percentage as a factual forecast, they are interacting with a signal that has been scrubbed of its moral and tactical friction. The laundering is complete when the “dirty” history of Discord disputes or exit-liquidity tactics is erased, leaving only a “clean” number that appears to be beyond human manipulation. This results in a dangerous epistemic loop: the market produces a signal through capital-driven bias, the platform hardens it into “truth”, and society consumes it as a neutral fact, with no one held responsible for the underlying distortion.

Algorithmic Gatekeeping and the Curated Future

The laundering process begins with an act of ontological power which is the ability to define the boundaries of what

is considered “bet-able” and, by extension, what is “knowable”. While prediction markets are often theorized as open, permission-less mirrors of reality, our findings in Stage 1 reveal a rigorous system of algorithmic gatekeeping. By centralizing the creation of contracts within a “markets team”, the platform does not merely aggregate existing information, but actively scripts the narrative of the future. This curation ensures that the only versions of reality allowed to reach the flattening stage are those that are measurable, binary, and compatible with blockchain settlement.

This gatekeeping functions as a form of epistemic enclosure. When a platform decides which questions are valid (e.g., “Who will win the election?”) and which are “too vague” (e.g., “Will the community be better off?”), it effectively prioritizes data-centric reality over complex social lived experiences. As our participants observed, this creates a curated set of futures that fit a specific brand. This curation is the first step in the laundry: it scrubs away the messy, non-binary complexities of social life, such as “why” an event matters, and replaces them with a clean, tradable ontology. The platform essentially manufactures a reality that is preformatted for speculation, ensuring that the “wisdom” produced is always limited by the platform’s own interests.

This curation also creates a false neutrality. Because the front-end interface is clean and automated, users often perceive the available markets as a neutral reflection of global interest. However, our audit shows that the “markets team” acts as a primary editorial board, deciding which geopolitical tensions or social crises are worthy of being turned into financial instruments. This is not a passive reflection of the world, but an active commodification of uncertainty. By selecting specific events to be “hardened” into data, the platform exercises a subtle but immense influence over public attention and political discourse.

Ultimately, this gatekeeping ensures that the “laundering” is successful from the very start. By the time a trader enters the market, the complex “dirty” reality of a social event has already been converted into a “clean” binary contract. The platform has already decided the terms of the debate, leaving the user to simply provide the capital that will eventually be laundered into “truth”. While this qualitative audit prioritizes interpretive depth over statistical generalizability, our final section explores how to break this cycle by introducing design interventions that restore the very friction the laundering process seeks to erase.

Design Recommendations: Toward Friction-Positive Interfaces

The laundering process we have described relies on the “smoothing” of information by stripping away context, motives, and disputes to produce a clean, authoritative signal. To counter this, we propose a shift toward “friction-positive design”. “Positive friction” refers to deliberate design choices that slow the user down, encouraging critical reflection rather than passive consumption (Chen and Schmidt 2024). For prediction markets, this means re-inserting the very “noise” that the platform currently seeks to launder. We suggest four primary design interventions to restore epistemic integrity:

Exposing Capital Asymmetry To break the “Architectural Masking” of Stage 3, platforms should implement “whale alerts” and “concentration metrics” directly on the front-end. If a 70% probability is being driven by a few large wallets, the interface should visually flag this concentration. By moving this data from the specialist layer of block explorers to the public UI, the platform would dismantle the two-tiered hierarchy of Epistemic Stratification.

Visualizing Interpretive Context To counter the “probabilistic flattening” of Stage 2, interfaces should allow for “multi-signal disclosure”. Instead of a single percentage, the platform could display a “sentiment map” that distinguishes between conviction bets, financial hedges, and tactical noise. Disclosing the intent behind the capital, even through optional user tagging, would prevent the “epistemic dilution” that occurs when a hedge is mistaken for a prediction.

Downstream Integration Standards for AI and Media Because platforms benefit from low-friction signals, interventions must also target external consumers. Search engines, LLMs, and media outlets integrating live market data should implement automated provenance tags. Rather than rendering a raw probability percentage (e.g., ‘72% likelihood’), AI assistants should append structural contextual cues (e.g., ‘driven by low order-book depth’ or ‘high whale concentration’) to prevent downstream epistemic inflation.

Audit-Ready Resolution Trails To address the “accountability gap” and the “oracle fallacy” of Stage 4, the history of a market’s resolution must remain visible. A friction-positive design would archive and link this Oracle votes and dispute logs directly to the final result. This ensures that the “hardened” fact remains connected to its contested origin, preventing the platform from evading institutional responsibility for controversial settlements.

By embracing these “programmed inefficiencies”, platforms can move from being machines of “laundering” to tools of epistemic accountability. These interventions do not seek to make the market faster or more efficient; rather, they seek to make it more honest. By exposing the friction behind the signal, we can ensure that prediction markets serve as a transparent mirror of collective intelligence rather than a sanitized megaphone for capital-driven influence.

Policy Implications Beyond interface design, these findings raise important policy considerations. As prediction market outputs are increasingly integrated into media reporting and decision-making contexts, there is a need for standards around interpretability and disclosure. These concerns are particularly salient in contexts where such signals may influence public decision-making without sufficient scrutiny of their underlying assumptions. This may include requirements to surface concentration metrics, clarify the role of financial incentives in shaping probabilities, and provide traceable records of dispute and resolution processes. Without such interventions, there is a risk that financially mediated signals are treated as neutral indicators of reality in contexts where their underlying assumptions and limitations are not understood.

Ethical Statement

This study was conducted with the approval of the authors' Research Ethics Board (REB) to ensure the protection of participant anonymity and data security. To mitigate risks associated with pseudonymous and volatile prediction platforms, we designed our protocol so participants were never required or incentivized to incur real-world financial risk. Furthermore, we recognize our role as researchers in a sociotechnical audit as one of non-participant observation. We maintained a deliberate outsider stance during the digital ethnography phase to avoid influencing market sentiment or community discourse on X and Discord. This positionality allowed us to document the platform's social and technical friction as it naturally occurs, providing a neutral baseline for our analysis of user interpretation.

Declaration on the Use of AI

The authors used Google Gemini and ChatGPT to assist in the preparation of this manuscript. Specifically, these generative AI models were used for refining the linguistic clarity and academic tone of the paper. All core concepts including the "Prediction Laundering" framework, the four stages of the laundering lifecycle, and the design recommendations were conceptualized by the human authors based on original empirical research. The authors reviewed, edited, and verified all outputs produced by the generative AI to ensure accuracy and adherence to academic standards. The authors take full responsibility for the content of this publication.

References

- Ananny, M.; and Crawford, K. 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society*, 20(3): 973–989.
- Britton, D. 2026. Perplexity Upgrades Finance Capabilities. *Yahoo Finance*. Accessed: 2026-05-03.
- Burrell, J. 2016. How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms. *Big Data & Society*.
- Callon, M. 1998. *The Laws of the Markets*. Oxford: Blackwell Publishers. ISBN 9780631206088.
- Chen, Z.; and Schmidt, R. 2024. Exploring a behavioral model of "positive friction" in human-AI interaction. In *International Conference on Human-Computer Interaction*, 3–22. Springer.
- Cross, V. 2026. Prediction Markets Record All-Time High \$5.3B Volume to End 2025. *DeFi Rate*. Accessed: January 5, 2026.
- Gitelman, L. 2013. *Raw data is an oxymoron*. MIT press.
- Goodman, L. A. 1961. Snowball Sampling. *The Annals of Mathematical Statistics*, 32(1): 148–170.
- Hayek, F. A. 2013. The use of knowledge in society. In *Modern understandings of liberty and property*, 27–38. Routledge.
- Hitchcock, L. 2025. Google Finance Integrates Polymarket, Kalshi Prediction Market Data. *Yahoo Finance*. Accessed: 2026-05-20.
- Jackson, S. J. 2014. Rethinking repair. In *Media technologies: Essays on communication, materiality, and society*, 221–240. MIT Press.
- Knorr Cetina, K.; and Bruegger, U. 2002. Global microstructures: The virtual societies of financial markets. *American journal of Sociology*, 107(4): 905–950.
- Lackey, J. 2008. *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.
- MacFarlane, J. 2005. The Making of Relativism. *Oxford Studies in Epistemology*, 1: 141–174.
- MacFarlane, J. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford University Press.
- MacKenzie, D. 2008. *An engine, not a camera: How financial models shape markets*. MIT Press.
- Massa, A.; and Avan-Nomayo, O. 2025. Kalshi and Polymarket Battle Gambling Companies in Prediction Market Frenzy. *Bloomberg*. Accessed: January 5, 2026.
- Metaxa, D.; et al. 2021. Auditing Algorithms: Understanding Algorithmic Systems from the Outside In. *Foundations and Trends® in Human-Computer Interaction*.
- Miah, S. 2025. 7 Media Companies Embracing Prediction Markets. <https://www.webopedia.com/crypto/learn/media-companies-embracing-prediction-markets/>. Accessed: 2026-01-05.
- Noy, C. 2008. Sampling Knowledge: The Hermeneutics of Snowball Sampling in Qualitative Research. *International Journal of Social Research Methodology*, 11(4): 327–344.
- Polymarket. 2026a. How Are Markets Created? Polymarket Documentation. Accessed: January 13, 2026.
- Polymarket. 2026b. Polymarket — The World's Largest Prediction Market. Accessed: January 5, 2026.
- Sandvig, C.; Hamilton, K.; Karahalios, K.; and Langbort, C. 2014. Auditing algorithms: Research methods for detecting discrimination on internet platforms. In *Data and discrimination: converting implicit bias into explicit design*.
- Surowiecki, J. 2005. The wisdom of crowds. *Surowiecki, J.*
- Tapscott, A. 2024. Forget the pundits and polls—internet prediction markets anticipated Biden's withdrawal weeks ago. *Fortune*. Accessed: January 5, 2026.
- TradingView. 2024. Bloomberg adding Polymarket election data to Terminal. *TradingView*. Accessed: April 2026.
- Wolfers, J.; and Zitzewitz, E. 2004. Prediction markets. *Journal of economic perspectives*, 18(2): 107–126.
- Yahoo Finance Editorial. 2025. Polymarket Achieves 94% Accuracy Predicting Real-World Outcomes. *Yahoo Finance*. Accessed: January 5, 2026.