Aalto-yliopisto
Aalto-universitetet
Aalto University

Bachelor's Programme in Industrial Engineering and Management

Mechanisms, Efficiency and Probabilistic Accuracy of Scalable Prediction Markets

Herman Bergring
Aalto University, School of Science

Bachelor's thesis 2025

Author Herman Bergring

Title of thesis Mechanisms, Efficiency and Probabilistic Accuracy of Scalable Prediction Markets

Programme Bachelor's Programme in Science and Technology

Major Industrial Engineering and Management

Thesis supervisor Risto Rajala

Thesis advisor Ruth Kaila

Date 14.12.2025 **Number of pages** 45 **Language** English

Abstract

Prediction markets offer the potential to aggregate dispersed information into accurate, real-time probability estimates, yet their practical impact has historically been constrained by liquidity frictions and regulatory uncertainty. While foundational literature often proposes Automated Market Makers (AMMs) and scoring-rule mechanisms to address the thin market problem, modern scalable prediction markets such as Kalshi instead adopt the Continuous Double Auction (CDA) architecture familiar from traditional exchanges. This thesis validates the design principles of scalable prediction markets by comparing Kalshi's market microstructure and institutional features against theoretical benchmarks for optimal information aggregation and efficiency.

The analysis shows that for high-volume markets, CDAs outperform subsidized AMMs regarding cost efficiency and price discovery. The thesis argues that Kalshi approximates the theoretical optimum through institutional solutions, specifically CFTC-regulation and central clearing, which foster the liquidity necessary for efficiency. Empirically, Kalshi's prices are suggested to be well-calibrated and capable of outperforming standard statistical forecasting models. Furthermore, the observed Favorite-Longshot Bias is reinterpreted as a structural liquidity premium induced by tick sizes and fees, rather than a behavioural failure.

The thesis characterizes these market prices as "meta-signals", wealth-weighted signals of all available information that are largely orthogonal to other data sources. Consequently, the thesis concludes that scalable prediction markets can transform collective beliefs into institutional-grade data products, providing reliable probabilistic inputs for financial hedging, strategic risk management and governance.

Keywords Prediction markets, event derivatives, information aggregation, probabilistic forecasting, market microstructure

Författare Herman Bergring

Titel Mekanismer, effektivitet och probabilistisk precision i skalbara prognosmarknader

Utbildningsprogram Kandidatprogrammet i teknikvetenskap

Huvudämne Produktionsekonomi

Ansvarslärare Risto Rajala

Handledare Ruth Kaila

Datum 14.12.2025 **Sidantal** 45 **Språk** Engelska

Sammandrag

Prognosmarknader har potentialen att aggregera spridd information till precisa sannolikhetsestimat i realtid, men deras praktiska genomslag har historiskt begränsats av låg likviditet och regulatorisk osäkerhet. Medan grundläggande litteratur ofta föreslår automatiserade marknadsgaranter (AMM) för att lösa problemet med tunna marknader, använder moderna skalbara prognosmarknader som Kalshi istället arkitekturen för kontinuerlig dubbelauktion (CDA) och orderbok som är standard på traditionella börser. Arbetet validerar designprinciperna för skalbara prognosmarknader genom att jämföra Kalshis marknadsmikrostruktur och institutionella egenskaper med teoretiska riktmärken för optimal informationsaggregering och marknadseffektivitet.

Analysen visar att för marknader med hög volym överträffar CDA-mekanismen subventionerade AMM:er gällande kostnadseffektivitet och prisupptäckt. Arbetet argumenterar för att Kalshi approximerar ett teoretiskt optimum genom institutionell design istället för algoritmiska lösningar, specifikt genom CFTC-reglering och central clearing, som möjliggör den likviditet som krävs för effektivitet. Empiriskt visar sig Kalshis priser vara välkalibrerade och i flera fall kunna överträffa statistiska standardmodeller för prognoser. Vidare omtolkas systematiska avvikelser, såsom ”Favorite–Longshot Bias”, som strukturella likviditetspremier orsakade av tick-storlek och avgifter, snarare än beteendemässiga brister.

Avhandlingen karaktäriserar dessa marknadspriser som ”metasignaler”: ekonomiskt viktade sammanvägningar av all tillgänglig information som i stor utsträckning är ortogonala mot andra datakällor. Slutsatsen är att skalbara prognosmarknader kan omvandla kollektiva uppfattningar till dataprodukter av institutionell kvalitet, vilket erbjuder tillförlitliga sannolikhetsbaserade inmatningsvärden för finansiell hedging, strategisk riskhantering och styrning.

Nyckelord Prognosmarknad, händelsederivat, informationsaggregering, sannolikhetsprognoser, marknadsmikrostruktur

Table of contents

Notations and abbreviations

Notations

Ω	Sample space
$\mathcal{F}, \mathcal{S}$	Sigma-algebras on distinct sample spaces
$\sigma(C)$	Sigma-algebra generated by set C
$\mathcal{P}(S)$	Power set of set S
$\mathbb{P}$	Probability measure
$\mathbb{E}$	Expectation
$\mathcal{S}_1 \otimes \mathcal{S}_2$	Product sigma-algebra
$\Omega_{\mathrm{Y}} \times \Omega_{\mathrm{S}}$	Product space
$P(C\|\mathrm{A}^{\mathrm{c}}B)$	Probability of event C, given complement of event A and event B

Abbreviations

0DTE	Zero Days to Expiration
AMM	Automated Market Maker
CCP	Central Counterparty
CDA	Continuous Double Auction
CFTC	Commodity Futures Trading Commission
CLOB	Central Limit Order Book
DCO	Designated Clearing Organization
DCM	Designated Contract Market
EMH	Efficient-Market Hypothesis
FLB	Favorite-Longshot Bias
FOMC	Federal Open Market Committee
HFT	High-Frequency Trading
IEM	Iowa Electronic Markets
LMSR	Logarithmic Market Scoring Rule
(L/M/P) SR	(Logarithmic/Market/Proper) Scoring Rule
MM	Market Maker
PM	Prediction Market

1 Introduction

The concept of prediction markets (PM) has been around since 1988, when the Iowa Electronic Markets (IEM) was founded to study market dynamics and forecast outcomes such as the United States presidential election (Tziralis & Tatsiopoulos, 2007). Since then, interest in and research on PMs have grown significantly, given the originality and conceptually attractive utility. Over the years, elegant theory has developed around the mechanisms, information aggregation capabilities, efficiency, and accuracy of PMs, but applications never exceeded small-scale experiments due to regulatory and liquidity constraints (Arrow et al., 2008). Treated more as gambling than as financial infrastructure, most PMs were constrained to strict position limits and legal uncertainty. This kept robust and diverse participation out, leading to illiquid markets, wide spreads, and noisy prices. Arrow et al. (2008), a group of 22 prominent social scientists including four Nobel prize winners in economics, argued for reforming regulation around prediction markets to realize the promised potential and utility of PMs. This promise is in the process of being delivered, as in late 2020, the Commodity Futures Trading Commission (CFTC) licensed the prediction market offered by the startup firm Kalshi.

Accompanied by technological advancements, the regulatory shifts have enabled different applications of scalable prediction markets to emerge. These emerging PMs, including Kalshi and rivalling Polymarket, gained extensive adoption and media coverage during the 2024 U.S. presidential election, where the ‘Trump’ or ‘Harris’ market attracted record volumes. In the US, prediction markets saw $1 billion traded, with $3.6 billion traded in the rest of the world. The subsequent market prices, interpreted as probabilities, provided a more accurate estimate than polling in the run-up to the election (Morrow, 2024). Following this, media pronouncements have emphasized the asymmetric potential of prediction markets. Charles Schwab, the founder of the $176 billion brokerage firm bearing his name, describes Kalshi as a company that “could fundamentally change financial markets” (Park, 2025). Coinbase CEO Brian Armstrong stated in a CNBC interview that PMs “could one day serve as an alternative to The New York Times” (as cited in Park, 2025), hinting at redefining how news is both presented and consumed, basing coverage of upcoming and on-going events on real-time probabilities.

For the first time at scale, there exists publicly available, real-time, cross-domain market data with predictive properties. The result is a continuous flow of accurate information about the future, which is immediately applicable in many high stakes domains, such as finance, corporate risk management, policy and governance, where reliable estimates are scarce yet very valuable. Central to this thesis is evaluating the reliability and impact of the prices produced by scalable prediction markets and identifying the structural conditions under which they can be trusted and applied as probabilities at scale.

Much of the foundational theory on prediction markets focused on solving issues such as the *thin market problem* (Hanson, 2003; Karimi & Dimitrov, 2024; Whelan, 2025) and the *no-trade theorem* (Chen et al., 2010; Whelan, 2025), two forces that create a self-reinforcing loop that has plagued previous PMs. Literature has often proposed different Automated Market Makers (Hanson, 2003; Hanson, 2012; Othman et al., 2013) or Market Scoring Rules (Hanson, 2003; Chen & Pennock, 2010; Luckner et al., 2012) to solve this negative loop by guaranteeing liquidity and incentivizing participation. Scalable PMs like Kalshi have instead opted for implementing a Continuous Double Auction (CDA), a proven mechanism for liquid markets in traditional finance. This poses the question whether emerging event exchanges like Kalshi, that aim to foster liquid, thick markets at scale, make a compromise or the optimal choice under the influence of real-world factors and constraints.

This thesis argues that Kalshi's microstructure does not diverge from optimal prediction market design but rather achieves it by solving for scalability given its novel positioning, combining market microstructure, institutional-scale design, and CFTC regulation. Therefore, the objective of this thesis is to validate the design principles of scalable prediction markets, using Kalshi as a benchmark, demonstrating how the distinct structure and solutions create conditions for optimal performance. This validation provides the basis for transformative applications, reframing probabilities as a reliable, real-time, institutional-grade data product for finance, risk management, and governance. As such, this thesis aims to answer the following set of research questions:

RQ1: How does Kalshi's market architecture, as a regulated and scalable PM, align with theoretical ideals of liquidity and informational efficiency?

RQ2: What key design features enable Kalshi to maintain liquidity and produce well-calibrated, information-rich price signals?

RQ3: What do existing theory and evidence suggest about the potential uses of scalable prediction markets as institutional-grade inputs, and what limitations or conditions are identified for their reliable application?

By addressing these questions, this thesis provides an evidence-based assessment of the potential for widespread institutional use of scalable prediction markets and their prices, grounded in three decades of theory. The primary contribution of this thesis is the development of a benchmark framework for evaluating scalable prediction markets.

The thesis is a literature review that combines peer-reviewed articles with working papers, theses and practitioner reports on prediction markets, market microstructure and information aggregation. Due to the recency of Kalshi's emergence and its niche regulatory positioning, there exists limited peer-reviewed literature directly analyzing its market structure or performance. Consequently, the thesis supplements foundational theory with secondary sources such as regulatory filings, analyses from Kalshi's research blog and platform documentation, and third-party analysis.

As for the research methodology, the academic evidence was gathered through targeted searches in major databases and working-paper repositories. Search queries were structured around the thesis' core topics: information aggregation, informational efficiency, probabilistic forecasting, liquidity and market microstructure, informed trading and manipulation, as well as commercial prediction markets. Sources were selected based on relevance to these topics and their usefulness for establishing differences and similarities between scalable prediction markets and academic theory. The thesis integrates these inputs through a benchmark framework that translates the most critical conditions for optimality into explicit assessment criteria. Kalshi is evaluated against these criteria by (i) extracting theoretical mechanisms from literature, (ii) using empirical findings from analyses on early Kalshi market data to ground expectations, and (iii) applying the specific institutional and regulatory context to identify where Kalshi plausibly satisfies, partially satisfies, or fails key conditions.

The scope of this thesis is a comparative analysis, focusing on Kalshi as a single case study against established PM theory. The theoretical focus is restricted to the elements directly used in the benchmark assessment: proper scoring rules and automatic market maker (AMM) design, informational efficiency, interpretation and accuracy of probabilistic prices, and their translation into scalable market microstructure.

The thesis is structured as follows: Section 2 defines the concepts discussed in the thesis; Section 3 covers the core theoretical foundations of prediction markets and examines scalable prediction markets against these theoretical foundations; Section 4 synthesizes the findings from Section 3 in the context of Kalshi's distinct solutions; Section 5 discusses the implications and applications of the validated results; Section 6 synthesizes the thesis and presents key findings and limitations as well as proposals for future research.

2 Preliminaries

To effectively examine the theory and applications of prediction markets, this section establishes precise definitions of the concept itself, its foundations, and the primary applications, along with other key terms and concepts used throughout this thesis. This includes notation for key mathematical concepts such as the price, outcome and payoff of a contract, information in the context of the price, and different contract types.

2.1 Prediction markets

In essence, prediction markets are event exchanges where participants can trade contracts on uncertain future events, and the resulting prices are commonly interpreted as collective probability estimates. The event exchanges match people who place bets on an event happening with people who bet the event won't happen. Through the mechanisms of crowdsourcing and the premise of the efficient-market hypothesis, PMs are assumed to function as efficient aggregators of dispersed and private information. Thus, the prices produced aim to reflect all available information and can be interpreted as probabilistic estimates of real-world outcomes.

The simplest and most common market format is the binary option for a discrete event market, referred to as an event contract, where the contract has two options, typically 'Yes' and 'No'. In an optimal market without fees, the sum of the price of each contract always totals 1, reflecting the fact that the probability of the whole sample space is always one. After the outcome of a contract is clear, each option settles at either 0 or 1. The current price of an open market reflects the implied probability of the event occurring, meaning the price of a 'Yes' contract being 0.7$, typically shown in cents as 70c, means that the market assigns a probability of 70% to the event happening. If the outcome happens, each 'Yes' contract settles at 1 and the value of 'No' contracts settle at 0. Thus, the core idea is that PMs incentivize the public to back their knowledge and beliefs, and in averaging the activity, PMs can transform ambiguity and chaos into a single, clear signal.

For the mathematical notation of prediction market theory, basic properties are denoted below in an introductory manner. We define a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with the sample space being a product space $\Omega = \Omega_Y \times \Omega_S$ of the space of outcomes $\Omega_Y = \{0,1\}$, defined on the measurable space $(\Omega_Y, \mathcal{P}(\Omega_Y))$, and the space of information paths Ω_S on $[0, T]$, with T being the time of settlement, defined on the measurable space $(S, \mathcal{S})$. Every $\omega \in \Omega$ is a vector (y, s) where $y \in \{0,1\}$ and $s = (s_u)_{0 \leq u < T}$ is a time-indexed stream of evidence observable at time u. This produces a rich space while keeping the event binary. Thus, $\mathcal{F} = \mathcal{P}(\Omega_Y) \otimes \sigma(s_u: 0 \leq u \leq T)$ denotes the sigma-algebra of information, containing relevant events $E \in \mathcal{F}$ and complements $E^c = \Omega \backslash E \in \mathcal{F}$. The filtration $(\mathcal{F}_t)_{0 \leq t < T}$ describes the information available up to a time $t < T$, and is defined $\mathcal{F}_t = \sigma(\emptyset, \{0,1\}) \otimes \sigma(s_u : 0 \leq u \leq t)$ that satisfies $\mathcal{F}_s \subseteq \mathcal{F}_t \subseteq \mathcal{F}$ for $0 \leq s \leq t$. The price at time t is denoted $p_t = \mathbb{E}[X|\mathcal{F}_t]$, where $p_t \in [0,1]$ and the price of the complement $\bar{p}_t = 1 - p_t$. The payoff of a binary contract is denoted $X = 1_{\{Y=1\}} \in \{0,1\}$, where $Y \in \{0,1\}$ is the realized outcome.

Beyond simple binary markets, PMs also run categorical markets with $K > 2$ mutually exclusive outcomes (e.g., which candidate wins), and range markets that divide a continuous value into buckets, with prices across buckets forming a single consensus, as illustrated in an example from Kalshi in Figure 1. Thus, the payoff of a binary contract in a categorical or range market is denoted $X_k = 1_{\{Y=k\}} \in \{0,1\}$, for outcomes or buckets $k = 1,2, \ldots, K$. As outcomes in both range and categorical markets are mutually exclusive, we have that $\sum_k X_k = 1$ and thus the price $p_{tk} = \mathbb{E}[X_k|\mathcal{F}_t]$ satisfies $\sum_k p_{tk} = 1$ and $p_{tk} \in [0,1]$.

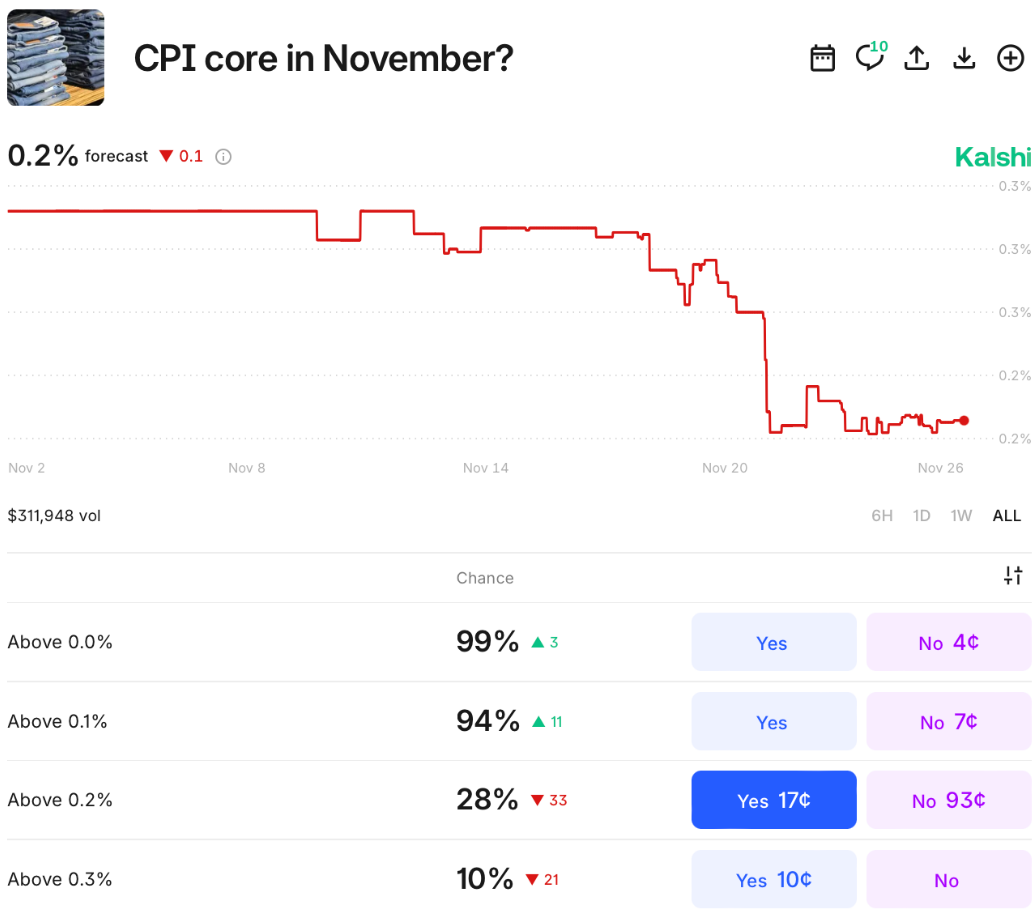

Figure 1. A sample Kalshi market: U.S. CPI core of November of 2025 (Kalshi, accessed 2025-12-01)

2.2 Kalshi

Kalshi is a U.S. based and internationally available prediction market that was approved by the Commodity Futures Trading Commission (CFTC) as a Designated Contract Market (DCM) in November 2020, allowing it to legally list and match trades (Kalshi, 2020). The DCM approval marked the first U.S.-regulated path for event trading at scale, turning prediction markets from contained experiments into institution-grade infrastructure with widespread access for all interested parties.

Kalshi is currently one of the largest PMs by notional volume and open interest and appears to lead its peers as of October 2025 (DuneData, 2025). Kalshi offers regulated event contracts across a wide variety of domains, including politics, economics, companies, technology, culture, sports, and weather. Unlike earlier PMs, Kalshi operates at scale with position limits at $7 million per contract, and as of October 2025, Kalshi's global liquidity pool spans over 140 countries, making it a viable venue for a wide range of market participants.

The traders on Kalshi's markets can be classified into distinct participant types, including market makers (MM), informed speculators, hedgers, and retail ("noise") traders (Kalshi, 2023b). MMs are usually professional firms providing liquidity by quoting bids and asks, whereas hedgers are institutions or individuals using contracts to balance other, often greater financial risks. Informed speculators are traders who possess superior information about outcomes, whereas retail traders are made up of the public, trading on public information or personal beliefs.

Kalshi's institutional focus is reflected in its market infrastructure and partnerships. Since early 2024, a global quantitative trading firm, Susquehanna International Group (SIG), has acted as a dedicated MM in Kalshi's markets, supporting tight spreads and continuous liquidity in key contracts. On the customer side, Kalshi offers a production-grade API, simplifying the integration of event contracts into existing trading systems for quantitative funds and brokers. Finally, clearing is handled by Kalshi Klear, Kalshi's own dedicated Derivatives Clearing Organization (DCO) that centrally manages collateral, margin and settlement, solidifying Kalshi's resemblance to established derivatives exchanges

2.3 Similarities to traditional markets and exchanges

The structure and payoff of Kalshi's contracts are very close to cash-settled futures that settle at 0 or 1. The only real difference is that the underlying is an event rather than an asset price. In this sense, Kalshi can be viewed as an exchange for event derivatives, applying familiar futures logic to real-world outcomes.

The trading mechanism is likewise close to major futures and equity exchanges. Kalshi uses a Central Limit Order Book (CLOB), illustrated by an example contract in Figure 2, with limit and market orders, discrete tick sizes, continuous trading, and posted fees. Order submission, matching and price formation follow the same logic as on major exchanges, so for most traders the only new element is the contract nature, not the trading process itself.

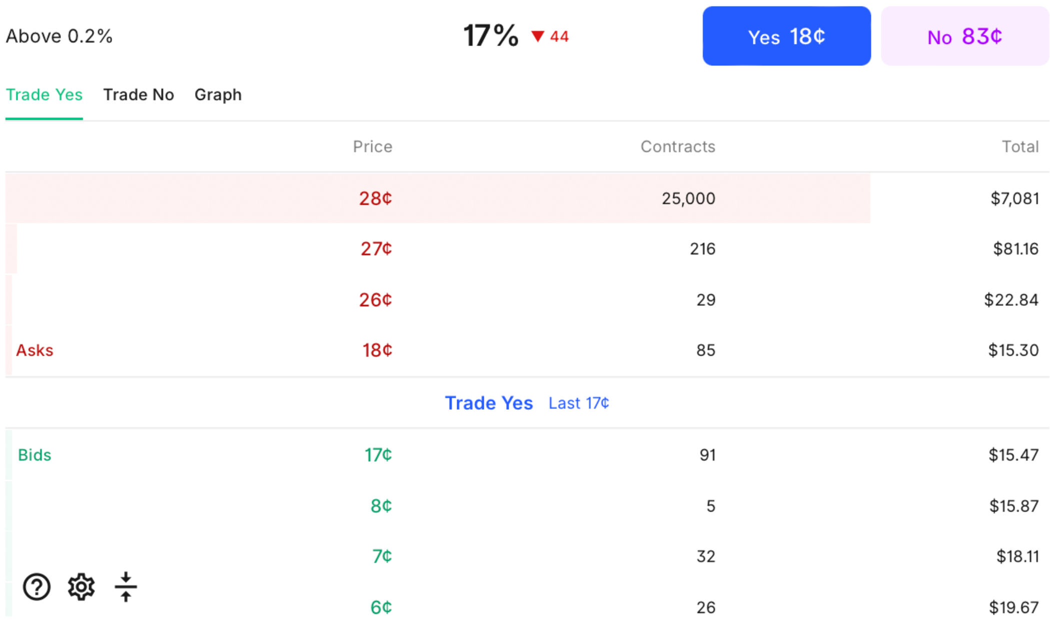

Figure 2. Options for predicting for or against the option 'above 0.2%' for U.S. CPI core of November of 2025 (Kalshi, accessed 2025-12-01)

Similarly, clearing, collateral and risk management also mimic traditional markets. Positions are fully collateralised against the bounded payoff, and Kalshi's clearing arrangements manage counterparty risk centrally. As in traditional futures markets, traders face marked-to-market profit and loss, margin requirements and position limits, which makes for straightforward integration into existing risk and reporting frameworks.

Thus, the main difference lies in the nature of the underlying and its fundamental value. Ordinary futures are written on tradeable assets that can be hedged and valued using standard models. On the contrary, event contracts are written on binary real-world outcomes that cannot be traded or hedged directly. Their fundamental value is the market's subjective probability of the event, reflected in the contract price. This preserves much of the familiar machinery of traditional exchanges while creating a distinct asset class in which what is traded is not an asset price but a state of the world.

3 Theoretical validation of the optimal prediction market

This section sets out a benchmark for what an optimal prediction market would look like in thick, frictionless conditions. This is done by drawing on the existing literature on mechanism design, market microstructure, and information aggregation. In this thesis, optimality is not synonymous with frictionless perfection, and the benchmark used is constrained optimality. For the thesis, such a benchmark is referred to as an optimal prediction market, in the sense that its design and outcomes are consistent with the main theoretical criteria proposed in prior work.

At a high level, these criteria can be grouped into three dimensions: (i) mechanism design, (ii) liquidity provision, and (iii) information aggregation and calibration. These are summarized into the causal structure referred to in this section, as presented in Figure 3. Mechanism design covers the trading rules and microstructure that support price discovery and orderly trading, even when activity is low. Liquidity provision concerns how the PM encourages a broad set of participants to supply orders and capital, so that the market can sustain large positions and relatively tight spreads. Information aggregation captures the extent to which dispersed signals are reflected in prices, and whether prices behave as probabilities over time.

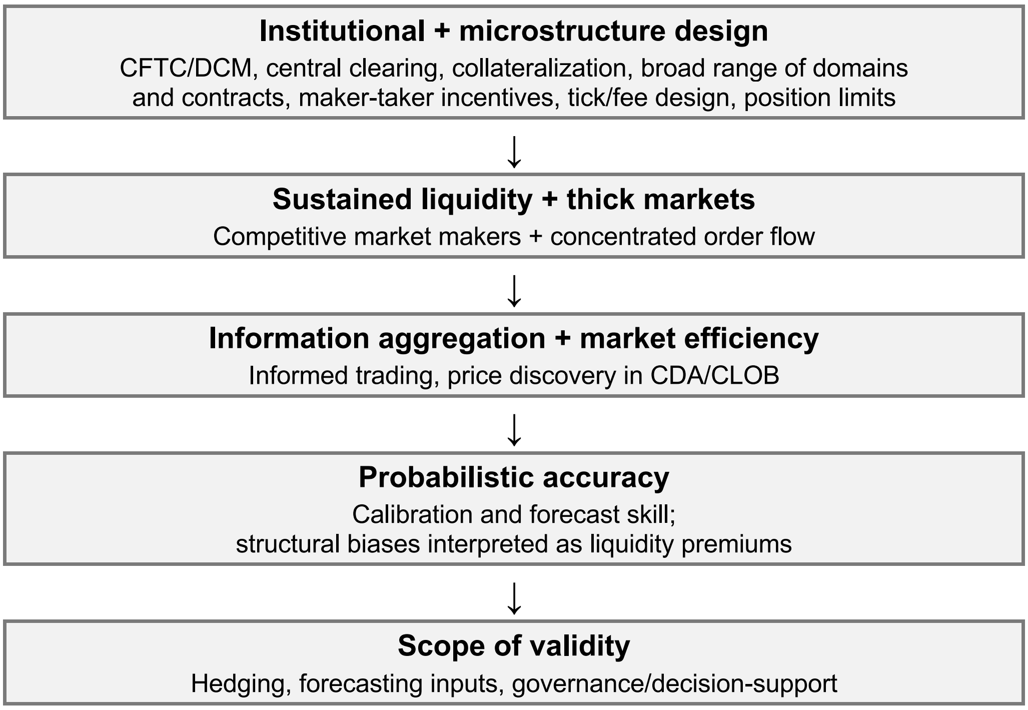

Figure 3. Benchmark model of a scalable prediction market

In addition, an optimal prediction market in this sense would be grounded in an institutional framework that makes these properties usable at scale. This includes a sufficiently broad set of contracts across many meaningful domains, position limits and margin rules that allow for substantial exposure, and a regulatory setting that permits participation by both retail and institutional users. The remainder of this section formalises these three dimensions and derives testable implications that are compared to the observed design and performance of Kalshi.

3.1 Mechanism design

Prediction market design involves choosing a market structure that achieves specific objectives, such as maximizing accuracy, managing liquidity and ensuring incentive compatibility (Chen & Pennock, 2010). This section introduces the two most common such mechanisms, namely Continuous Double Auctions (CDAs), used for CLOBs, and Market Scoring Rules (MSRs), used for Automated Market Makers (AMMs), (Luckner et al., 2012). In a CDA market potential buyers submit their bids and potential sellers submit their ask prices, and then the market establishes a price at which the market is cleared (Karimi & Dimitrov, 2024), much like in traditional markets and exchanges.

On the other hand, Proper Scoring Rules (PSRs) are payment schemes constructed so that a rational forecaster with a specific true belief optimizes their expected score by reporting probabilities equal to the belief, making truthful probabilistic reporting incentive-compatible (Waghmare & Ziegel, 2025). MSRs extend this idea into a mechanism. To avoid the thin market problem an AMM maintains a current forecast and lets any participant replace it with a newer forecast, receiving the incremental payoff between the two forecasts once the outcome is realized (Waghmare & Ziegel, 2025). In theory, this continuously provides liquidity while incorporating information through forecast updates. Thus, AMMs exchange directly with traders, rather than facilitating traders trading directly with each other (Karimi & Dimitrov, 2024).

3.1.1 Continuous Double Auctions: Makers and Takers

In contrast to AMMs, CDAs facilitate peer-to-peer trading based on participant orders rather than an algorithm. Kalshi uses a Central Limit Order Book (CLOB) infrastructure, the specific data structure that implements a CDA. A CLOB is an electronic ledger that matches limit orders based on price-time priority, matching orders firstly based on price and in ties matched on time. This means that the highest ("best") bid order and the lowest ("cheapest") ask order make for the best market for a trade, which Kalshi calls a "maker-and-taker" market structure (Bürgi et al., 2025). This mechanism is transparent, anonymous and low cost. Kalshi operates commercially and charges a per-trade fee of $F = \theta p_t(1 - p_t)$, where at present $\theta = 0.07$. This fee structure has a concave cost shape, and peaks at $p_t = 0.5$. Kalshi's approach to fees imply smaller distortions in prices than a flat fee on all contracts (Whelan, 2023) but nevertheless has implications for previous theoretical models and results in PM literature that assume no fees (Bürgi et al., 2025).

Providers of liquidity (makers) are traders who submit passive limit orders that rest on the order book, create depth and define the spread (Bürgi et al., 2025). Consumers of liquidity (takers) are traders who submit aggressive orders that execute instantly against the best available resting orders (Bürgi et al., 2025). Makers aim for better prices but face execution risk, while takers get instant execution but pay the spread (Whelan, 2025). Thus, the bid-ask spread represents the cost of instant liquidity and potential profit for makers. These spreads are influenced by market thickness, volatility and information asymmetry (Whelan, 2025).

A thick market implies high liquidity and tight spreads, whereas a thin market implies low liquidity, wide spreads and low match rates (Whelan, 2025). In a thin market the CDA prices are often prone to two main issues, volatility and inconsistency with true consensus as well as wide spreads resulting in withholding of information due not being able to profit from trading (Karimi & Dimitrov, 2024). The former issue is called the *thin market problem*, whereas the latter is classified as the *no-trade theorem*.

Liquid markets across traditional exchanges and all asset classes almost exclusively run on CLOBs, with the reason being their competition inducing nature that effectively commoditizes market function and creates better outcomes for all genuine participants.

3.1.2 Automated Market Makers: scoring rules and incentive

MSRs, which function as AMMs, are one of the primary solutions for eliciting information (Chen & Pennock, 2010). AMMs are algorithmic agents that enable participation and information elicitation in electronic markets (Othman et al., 2013). In practice, AMMs subsidize a set amount, defined by the bounded loss, to effectively aggregate the traders' beliefs on the outcome of a future event (Hanson, 2003). This requires a variable cost to be paid upfront by the market organizer every time a market opens (Karimi & Dimitrov, 2024), which is seen as the main drawback of using AMMs. However, AMMs solve the mentioned *thin market problem* of CDAs by constantly providing liquidity at consensus prices (Hanson, 2003). In practice, AMMs are subsidized counterparties that enable trading even among rational traders, resolving the *no-trade theorem* issue as well (Chen et al., 2010).

Proper Scoring Rules (PSRs) provide the standard framework for eliciting and evaluating probabilistic estimates (Waghmare & Ziegel, 2025). Formally, PSRs are defined as mechanisms designed to elicit truthful, subjective probabilities from traders by maximizing their expected score when they report honestly (Armantier & Treich, 2013; Waghmare & Ziegel, 2025). Theoretically, PSRs guarantee truthfulness through incentive compatibility under risk neutrality (Armantier & Treich, 2013), paying an agent a score $s_i(\pi, i)$ based on their reported distribution $\pi = \{\pi_i\}$ and the actual outcome *Y=i*, where *i* denotes all possible states. The scoring rule is 'proper' if a risk-neutral agent's expected score is uniquely maximized by reporting true private belief $\boldsymbol{p} = \{p_i\}$ (Hanson, 2012):

$$\boldsymbol{p} = argmax_{\boldsymbol{\pi}} \left(\sum_i p_i \, s_i(\pi, i) \right),$$

where a risk-neutral agent cares only about maximizing expected payoff, not about risk or payoff variance (Hanson, 2012). This incentive compatibility is the theoretical foundation for all MSRs. An example of a PSR is the Logarithmic Scoring Rule (LSR) proposed by Hanson (2012), that emphasizes modularity, which means that bets on a specific conditional outcome (e.g., event A given event B) updates the probability $P(A|B)$ while preserving the marginal probability of the conditioning event *P(B)*, essential for complex information structures.

The Logarithmic Market Scoring Rule (LMSR), derived from the LSR, is a specific AMM design that characterises many of the optimal properties of an optimal market design (Hanson, 2003). These include bounded loss for the market organizer (Chen et al., 2010), cost efficiency where the cost of subsidizing is related to the information gain (Hanson, 2003), and modularity. The prices of an LMSR are derived from a cost function *C*, as defined below. For a market with n outcomes, the total cost for the MM to have sold $\boldsymbol{q} = q_1, q_2, \ldots, q_n$ shares of each outcome is:

$$C(\boldsymbol{q}) = b \; log \left(\sum_Y e^{q_Y/b} \right),$$

where b is the liquidity parameter. The market price for the next share for an outcome Y is then the derivative of C, $p_Y = \partial C / \partial q_Y$. This single formula provides the LMSRs key properties: the guaranteed bounded loss for the MM equal to $b \log(\mathrm{n})$, while providing guaranteed liquidity at all prices (Hanson, 2003).

Hanson (2012) also proves that the LMSR is the unique rule with the property of preservation of conditional independence, denoted as *I*(*A,B,C*), implying that logarithmic rule bets on A given B preserve *P(B)*, and for any event *C,* preserve the conditional probabilities $P(C|AB)$, $P(C|\mathrm{A^c}B)$, and $P(C|\mathrm{B^c})$ (Hanson, 2012). Practically, this allows the market to decompose complex probability spaces and prevent combinatorial explosions of updates when traders reveal specific, superior information (Hanson, 2012).

However, the LMSR and the underlying LSR are only guaranteed to be myopically, that is shortsightedly, truthful, which creates the possibility of non-myopic, strategic manipulation (Chen et al., 2010). Thus, informed traders might bluff or strategically withhold information to maximize gains (Chen et al., 2010). In addition, the underlying PSRs have been shown to produce biased prices if traders are not risk-neutral or have external financial positions, if for instance traders are hedging on the outcome (Armantier & Treich, 2013).

As AMMs are passive, when the real-world price moves, they are picked off by arbitrageurs who move the price towards the current, true value (Othman et al., 2013). The profit, that in other asset classes would be earned by skilled MMs, is thus transferred to arbitrageurs. In microstructure terms, AMMs are constantly exposed to similar risks without the ability to adjust quotes dynamically (Linnainmaa, 2010). Therefore, the losses suffered by an AMM during volatile periods are effectively permanent subsidies, reflecting systematic exploitation by informed traders.

3.1.3 Capital efficiency and the cost of liquidity

By using MSRs, AMMs offer guaranteed liquidity at the cost of a subsidy, in theory solving the *thin market problem* and the *no-trade theorem* (Hanson, 2003). On the other hand, by using CDAs, CLOBs offer competitive price discovery through peer-to-peer order matching but require endogenous liquidity from Makers (Whelan, 2025). Karimi and Dimitrov (2024) frames it as “to subsidize or not to subsidize”, when investigating which mechanism leads to better price discovery.

Although MSRs have been proposed as the appropriate marked design for prediction markets (Hanson, 2003; Luckner et al., 2011) in both thin and thick markets, the claims lack empirical verifications (Karimi & Dimitrov, 2024). Hanson (2003) argues that markets using LMSRs match the accuracy of CDAs when the number of trades per trader increases. However, Othman et al. (2013) identifies two key issues for scalable, practical use of LMSRs: the price movement is determined by a manually set liquidity parameter and doesn’t adapt to trading volume, and the required subsidy grows with liquidity, creating a cost barrier to scaling.

In PM literature, ex-post analysis of price discovery is considered the predominant method for comparing market mechanisms (Wah et al., 2016). Price discovery is defined as the root-mean-square deviation between the market price and the ground truth (Wah et al., 2016), essentially measuring how efficient the aggregation of new information into prices is. Instead of measuring ex-post accuracy, Karimi & Dimitrov (2024) argue for using price discovery share, which measures a “market’s proportional contribution to price discovery relative to other markets”. This is measured using information share, that identifies which market moves first upon new information, and component share, that measures each market’s contribution to the common, efficient price in the long run (Karimi & Dimitrov, 2024).

The incentive structures also differ. As noted, the underlying MSRs of AMMs are only myopically incentive compatible, potentially allowing strategic manipulation. In a CLOB, Makers face the risk of their resting orders being executed by Takers with superior information, called adverse selection risk, which incentivizes makers to be well-informed (Linnainmaa, 2010). The transparency of the order book allows informed traders to compete directly and punish mispriced orders, increasing non-myopic incentive for truthful pricing in liquid markets (Linnainmaa, 2010). Simply put, an AMM is a passive MM who is always quoting, but never with the benefit of constant aversion of adverse selection. Thus, in terms of subsidizing for liquidity, it is a question of adverse selection in AMMs versus inventory risk in CDAs. Out of the options, inventory risk is preferred by professional MMs, as they allow for active position adjustments and risk management (Ho & Stoll, 1981) as opposed to depositing liquidity into a static pool.

Empirically, the choice of mechanism depends on market thickness. In thin markets, characterised by low volumes, Karimi & Dimitrov (2024) confirm that subsidized MSRs show higher price discovery shares. As for thick markets, however, they provide the first empirical evidence that CDAs outperform MSRs in price discovery share. This contrasts with the prior consensus of equal performance. They conclude that when the number of trades in a market is high (informally, n>982), a market organizer should prefer an unsubsidized CDA market, as it's cheaper and faster at aggregation information (Karimi & Dimitrov, 2024). Scalable prediction markets can be assumed to consistently and significantly surpass this threshold, as the volumes of key contracts range from hundreds of thousands to tens or hundreds of millions in dollars.

Given this, the CLOB mechanism is both theoretically and practically the optimal choice. It leverages the superior price discovery of CDAs in high volume markets, avoids the subsidy costs and scaling issues of AMMs, and fosters strong information aggregation incentive induced by transparent competition. Thus, we can conclude that there is significant evidence to support that for a scalable, regulated, institutional-grade prediction market like Kalshi, the CDA implemented by using a CLOB is the optimal mechanism choice for efficient and cost-effective information aggregation.

3.2 Efficiency and information aggregation

Having established the theoretical and practical superiority of using the Continuous Double Auction (CDA) mechanism for a scalable prediction market, this section validates the performance of this given mechanism. As mentioned, the central promise of prediction markets is their informational efficiency. In the foundational paper presenting the efficient-market hypothesis (EMH), Fama (1970) states that in an efficient market, prices fully reflect all available information, which underpins the idea of PM prices being accurate probabilistic estimates. This ability of aggregating large amounts of dispersed information into a single, consensus price distinguishes PMs as powerful predictive tools (Wolfers & Zitzewitz, 2004). To further validate Kalshi as an optimal PM, we assess whether the fundamental goal of informational efficiency is achieved.

3.2.1 Efficiency in prediction markets

By establishing informational efficiency as the theoretical benchmark for an optimal market and the meaning of informational efficiency in the context of prediction markets, this section focuses on the statistical properties of an efficient prediction market. The premise of the EMH claiming that the current price reflects all information contained in historical prices can be likened to the martingale hypothesis and thus measured statistically (Richard & Vecer, 2021). The property of this hypothesis, called the martingale property, states that the last quoted price is the best predictor of the outcome, and that all previous information, including past prices, should be statistically insignificant in predicting the next price movement, and thus the outcome (Richard & Vecer, 2021). Formally, given all information available $\mathcal{F}_t$ at time t, the expected price at time t+1 is equal to the current price:

$$\mathbb{E}[\mathrm{p}_{t+1}|\mathcal{F}_t] = \mathrm{p}_t.$$

This is also the classic, simple notation that defines Weak-Form Efficiency (Fama, 1970).

Efficiency is not a static property, but rather an active process of price discovery. As previously defined, price discovery measures how efficiently new information is aggregated into prices (Wah et al., 2016), and the speed of this process is the primary empirical measure of Semi-Strong Form Efficiency (Fama, 1970). Fama (1970) categorizes market efficiency into three forms: Weak, Semi-Strong and Strong, where Semi-Strong Form Efficiency, a superset of Weak Form Efficiency, asserts that prices adjust rapidly and accurately to all publicly available information, including historical prices as well as related data, news and announcements (Fama, 1970).

Beyond definitions, price discovery can be modeled as the equilibrium outcome of a dynamic "game" between traders (Wah et al., 2016). In this view, strategic traders form beliefs based on their private information and the current price, then decide whether to submit market or limit orders, or to opt out. The market mechanism, the CLOB in Kalshi's case, provides the incentive structure that shapes these choices and aggregates competing strategies into a single price path (Wah et al., 2016). In an equilibrium, that path defines both the information of traders' beliefs and the design of the mechanism that aggregates these beliefs.

In this equilibrium view, whether information is reflected in prices is thus dependent on both who trades and how the market is structured. Corgnet et al. (2023) outline the conditions for aggregation and demonstrate that the process is “fragile” and often fails in most asset market structures. In their paper, Corgnet et al. (2023) fail to replicate a seminal experiment on information aggregation, and in turn proposed an alternate market structure that simplifies the asset and introduces contingent claims, where the payoff of an asset is dependent on the outcome of a future event, mirroring Kalshi’s contracts specifications. For this market structure, Corgnet et al. (2023) find “robust evidence of information aggregation”, which implies that the formal and informal rules that govern a market, called institutional features, have a critical impact on market efficiency.

In Kalshi’s case, institutional features include mechanism design, the regulatory framework, contract specifications and participation rules. From this follows that the specific institutional design of prediction markets, like Kalshi’s, makes them uniquely effective at information aggregation. In particular, the simple, easily priced binary contracts, structured as contingent claims on well-defined events, appear to be a key feature in promoting information aggregation.

3.2.2 Market microstructure and equilibrium dynamics

The process of rapid price discovery requires a functioning and well-designed market, as outlined in section 3.1. A prerequisite for informational efficiency is that traders must be able to act on their beliefs instantly and at marginal cost. Thus, market efficiency is not just an outcome but is contingent on the market’s underlying liquidity, which manifests in the form of market depth, bid-ask spreads, and the absence of arbitrage opportunities (Bürgi et al., 2025; Hegarty & Whelan, 2025; Whelan, 2025).

Whelan (2025) establishes that exchanges like Kalshi can settle into multiple equilibria: thick or thin markets. Liquidity in a CLOB is endogenous and not guaranteed, depending on the willingness of makers to post orders. If the probability of a match is considered low, makers demand a wider spread to compensate for the risk of tied up capital, creating a loop that keeps markets thin (Whelan, 2025). On the opposite, in a thick equilibrium, high match rates encourage makers to post tighter spreads, further attracting takers and deepening liquidity (Whelan, 2025).

A structural challenge for liquidity in prediction markets in general, is that event contracts naturally resolve, forcing exchanges to continuously list new contracts as old ones expire. This creates a constant “cold start” problem where liquidity must be built from scratch. This feature is fundamentally unfavorable for liquidity and increases the probability of markets remaining thin, due to market makers (MMs) being hesitant to quote large sizes on *ad-hoc* contracts as the inability to train quantitative models creates unmanageable return uncertainty (Ho & Stoll, 1981).

This means that MMs prefer markets with large amounts of data from consistent, recurring events. Without confident MMs, markets remain thin and price discovery suffers.

Kalshi addresses this by listing recurring, non-ad-hoc markets like weekly Consumer Price Index (CPI) reports and daily S&P500 movements, where the underlying event is of recurring nature. In addition, 0DTE (Zero Days to Expiration) contracts on events such as Nasdaq and S&P Daily Closes, 10-Year Treasury Daily Yields, and WTI Oil prices further increase the appeal of market making (Kalshi, 2023c). These types of contracts provide the required continuous data series that MMs need to train quantitative models, price contracts quickly and to measure risk effectively, giving them confidence to quote larger sizes and maintain deep liquidity. Moreover, Kalshi onboarded the professional trading SIG in 2024 (Kalshi, 2024) and started its dedicated market maker program in 2025, explicitly targeting the creation of thick markets.

To bridge the gap between retail flow and institutional capital, Kalshi has introduced block trading via their own entity, Kalshi Klear. In standard CLOBs, large orders face significant price impact where a large buy order can sweep the entire order book, driving the price up and worsening the entry price of the order, which discourages institutional participation. Block trades, or "off-exchange trades", allow institutions to negotiate large transfers of risk privately, bypassing the public order book (Kalshi Klear LLC, 2025). This mechanism is critical for efficiency, as it allows the market to absorb massive liquidity without the immediate volatility in the CLOB, while the subsequent public reporting of the trade still contributes to price discovery.

Despite these solutions, specific design choices induce frictions that affect efficiency. Firstly, Kalshi imposes a fee structure that includes a transaction fee of $0.07p_t(1-p_t)$ (Bürgi et al., 2025). While Whelan (2023) notes that this concave fee has less effect than a flat fee, it still imposes a cost that widens the spread required for makers to break even. As fees decrease on both sides of the trade, they disproportionately impact the expected returns of takers, specifically on low-priced "longshot" contracts (Whelan, 2023).

Another possible friction is the tick size, as Kalshi's minimum tick size is 1c (1%) per contract (Bürgi et al., 2025), contrasted with other PMs, such as Polymarket, who operate with a 0.1c (0.1%) tick. This discrete tick mechanically enforces a floor of 1c for the bid-ask spread (Harris, 1994), which leads to an important trade-off in efficiency. On one hand, the 1c tick significantly weakens the ability of informed traders to express their superior information. To illustrate, a highly informed trader who prices a contract at 92.6c with a 0.1% error cannot express this by quoting a 92.5c-92.7c spread but is rather forced to quote the much wider 92c-93c spread. A trader with lower confidence, with say a 0.5% error, is also forced to quote the same spread. On the other hand, a 1c spread represents a potentially up to ten times larger profit for MMs on a single trade compared to a market with a 0.1c spread, which essentially serves as a minimum profit margin for liquidity providing MMs (Harris, 1994). A *binding* tick, meaning that the natural spread would be lower than the tick, is thus a necessary factor in ensuring wide enough spreads (Harris, 1994) to cover inventory risk faced by makers, functioning as a critical liquidity preservation mechanism.

Consequently, a potential conflict of interest arises regarding liquidity provision, given the market structure imposed by the tick size. Bürgi et al. (2025) find that makers generally earn positive returns whereas takers earn negative returns. Since Kalshi's affiliate, Kalshi Trading, participates in the market as a liquidity provider (Kalshi, 2023b), Kalshi technically benefits from the spreads and potential inefficiencies that its tick-size rules enforce. While this participation ensures liquidity and a solution to the no-trade theorem, it raises questions about whether user execution quality or operator revenue is prioritized in Kalshi's market design. In any case, the 1c tick is another institutional feature of Kalshi that breeds thicker markets, which is for the greater good of a scalable PM and its users in general.

3.2.3 Informed trading and insider activity

According to the EMH, Strong Form Efficiency asserts that prices reflect all information, both public and private ("insider") (Fama, 1970). This form implies that no one can consistently outperform and is generally considered not to hold in practice, which is why "insider trading" is prohibited in traditional financial markets. If insiders were allowed to trade on private information in traditional markets, they would essentially gain value from retail traders, effectively driving liquidity out of the market. However, in prediction markets the objective is not building capital nor fair exchange but rather generating accurate price signals. Thus, excluding those with the best information, the insiders, is theoretically counterproductive to the utility of prediction markets.

Kalshi explicitly prohibits traders with the ability to influence the outcome of an event from trading on that specific contract. However, Kalshi allows expert-based informed traders, so in the theoretical discussion that follows, the term 'insider' refers to traders with expert-based or superior information rather than to the legal notion of insider trading in securities law.

Consequently, PMs are often described as if they could achieve Strong Form Efficiency. An argument against this is illustrated by Grossman and Stiglitz (1980), as they showed that perfectly revealing prices cannot be an equilibrium when superior information exists. If prices were to be *perfect*, gathering information would be unprofitable, resulting in no reason to trade. In practice, an informationally efficient market should therefore leave small, systematic profit opportunities that compensate insiders and informed traders for their information and inventory risk.

The mechanism through which informed traders improve efficiency is distinct in contingent claims markets, that is markets where assets pay out based on specific event outcomes (Corgnet et al., 2023). Corgnet et al. (2023) find that "markets with competing insiders effectively disseminate the insiders' information to other traders". Furthermore, they note that in contingent claims markets, insiders help traders identify states of the world that should be viewed as unlikely as insiders will be selling those assets. As a result, insiders not only move prices towards the true probability, but actively shape the probability distribution for the rest of the market, showing which outcomes are effectively impossible (Corgnet et al., 2023).

However, the efficiency gained by allowing informed trading is contingent on the other participants' beliefs. One factor that Corgnet et al. (2023) don't address is the effect of ambiguity on informational efficiency, where traders have loose beliefs about complex events and avoid trading. Galanis et al. (2024) show that ambiguity can cause prices to remain at incorrect values and become prone to manipulation, as traders don't trade against a mispricing that they are unsure about. In this context, insiders are central in resolving ambiguity. By trading aggressively on superior information, insiders encourage broader market participation and allow for the market to converge towards the true probability (Galanis et al., 2024).

Despite the benefits in moving toward more accurate probabilities, the suspicion of insider activity brings about a challenge regarding trust. Choo et al. (2022) find empirically that the mere suspicion of manipulation or private, superior information can weaken trust in the market, leading to suboptimal utilization of the information revealed in prices, even if prices were to be accurate. This presents a challenge for exchanges like Kalshi, who must facilitate informed trading to ensure optimal accuracy while maintaining a market structure that appears fair, to sustain liquidity from retail traders. To do this, Kalshi enables expert-based informed trading, but strictly enforces rules around blatant insider activity.

In any case, empirical evidence suggests that PMs should in general be resilient toward manipulation and insider activity. Buckley and O'Brien (2017) conducted experiments on PM manipulation and found that while attempts to distort prices have an immediate effect, the effects are "rapidly ameliorated by rational traders". In a liquid, scalable market like Kalshi, informed traders thus act as a balancing force, effectively seeing manipulation as an arbitrage opportunity, quickly trading against distortions and returning prices to efficient equilibria (Buckley & O'Brien, 2017). Therefore, in a CDA market, aggressive trading by highly informed participants acts less like a source of unfairness and more like a mechanism for quick error correction and ambiguity resolution.

3.3 Prices as probabilities

The foundational premise of PMs is that the trading prices of event contracts can be interpreted as the market's aggregated probability of those events occurring. While prices are useful estimators, they are not probabilities in a statistical sense, but rather prices that are formed through information aggregation; essentially complex equilibrium outcomes derived from wealth, risk preferences and belief distributions (Wolfers & Zitzewitz, 2006). This is at the core of PMs, laying the foundational view that PM prices can and should be viewed as the mean subjective probability assigned by market participants.

Wolfers and Zitzewitz (2006) provide the analytical framework, showing that prices correlate strongly with the mean belief of market participants under certain conditions. The four criteria found by Wolfers and Zitzewitz (2006) to imply that prices in fact are accurate probabilities are: traders being risk neutral, having logarithmic utility, paying no transaction fees and having no limits. In theory, these criteria form a frictionless, limitless and idealized market structure where traders are indifferent to variance in returns and focus solely on maximizing expected value, while also maximizing the growth rate of their wealth rather than settling for linear profits (Wolfers & Zitzewitz, 2006). Scalable PMs like Kalshi partially satisfy the structural criteria set by Wolfers and Zitzewitz (2006) but struggle with the frictionless criteria due to real-world regulatory and operational constraints. Thus, validating Kalshi as a source of accurate predictive data requires decomposing the mechanics that translate beliefs into these equilibrium prices.

3.3.1 Theoretical foundations

This transformation of beliefs into market prices is constrained by the risk tolerance of traders. If all traders in a market have logarithmic utility, they will bet sizes proportional to their confidence and wealth (Wolfers & Zitzewitz, 2006). Wolfers and Zitzewitz (2006) proved that in such a market, the equilibrium price exactly equals the wealth-weighted mean belief of all traders. Moreover, He and Treich (2017) demonstrate that risk aversion, the opposite of risk-neutrality, also drives a wedge between the true probability and the price, reinforcing the result of Wolfers and Zitzewitz (2006) that the equilibrium price has a component of risk-adjusting in the aggregation. This means that prices are skewed toward the beliefs of wealthier or less risk-averse traders, as they are willing to take larger positions.

Consequently, as mentioned before, the prices shouldn't be seen as true probabilities, but rather risk-adjusted, wealth-weighted probabilities. Furthermore, the incentive to reveal truthful information is often distorted by external exposure. Armantier and Treich (2013) show that if traders have positions in the outcome outside the market, as in for example hedging against a recession, they will bias their trading behaviour to insure themselves rather than to maximize speculative profits. This adds the aggregation of hedging demand to the notion of equilibrium prices, which has been proved by Buckley and O'Brien (2017) to distort prices immediately, but however the effects are "rapidly ameliorated by rational traders", correcting the manipulation as such.

Despite these distortions, the efficiency of price paths can be rigorously tested using the martingale property. As defined by Richard and Vecer (2021), a market is efficient if the current price is the best predictor of the outcome. If prices follow a martingale, they effectively are rational probability estimates that update instantly to new information. As the market mechanism achieves efficiency, the resulting price p_t represents the wealth-weighted, risk-neutral and continuously updated expected value of the outcome $X \in \{0,1\}$: $p_t = \mathbb{E}[X|\mathcal{F}_t]$. For event contracts this conditional expectation equals the current implied probability for the event, so testing for the martingale property is equivalent to testing whether the market generates dynamically consistent probability estimates.

3.3.2 Empirical validation and the Favorite-Longshot Bias

Recent empirical evidence on Kalshi suggest that, despite the FLB, the prices on Kalshi would remain highly accurate. A third-party analysis of early Kalshi market data by Clay (2022) analyzed Kalshi's historical performance using the Brier Skill Score (BSS), which combines aspects of predictive ability, calibration and confidence into a value with the highest score of 1. As shown in Figure 4, Kalshi prices from 8,476 markets in their last hour before settlement are tightly clustered around the line of perfect calibration, $\Delta y / \Delta x = 1$, illustrated as the dashed line. The red dots in the plot are the observed proportions of markets that eventually resolved to 'Yes' for all markets that traded at each price (Clay, 2022). The BSS of 0.624 one hour before settlement, contrasted to a BSS of 0.361 one week before settlement, suggests that prices remain well-calibrated to outcomes and improve in accuracy as information arrives. This means that these event contracts, analyzed long before today's deep liquidity, priced at for instance 70c occurred roughly 70% of the time. While this evidence is not peer-reviewed, it offers a preliminary indication of forecasting skill.

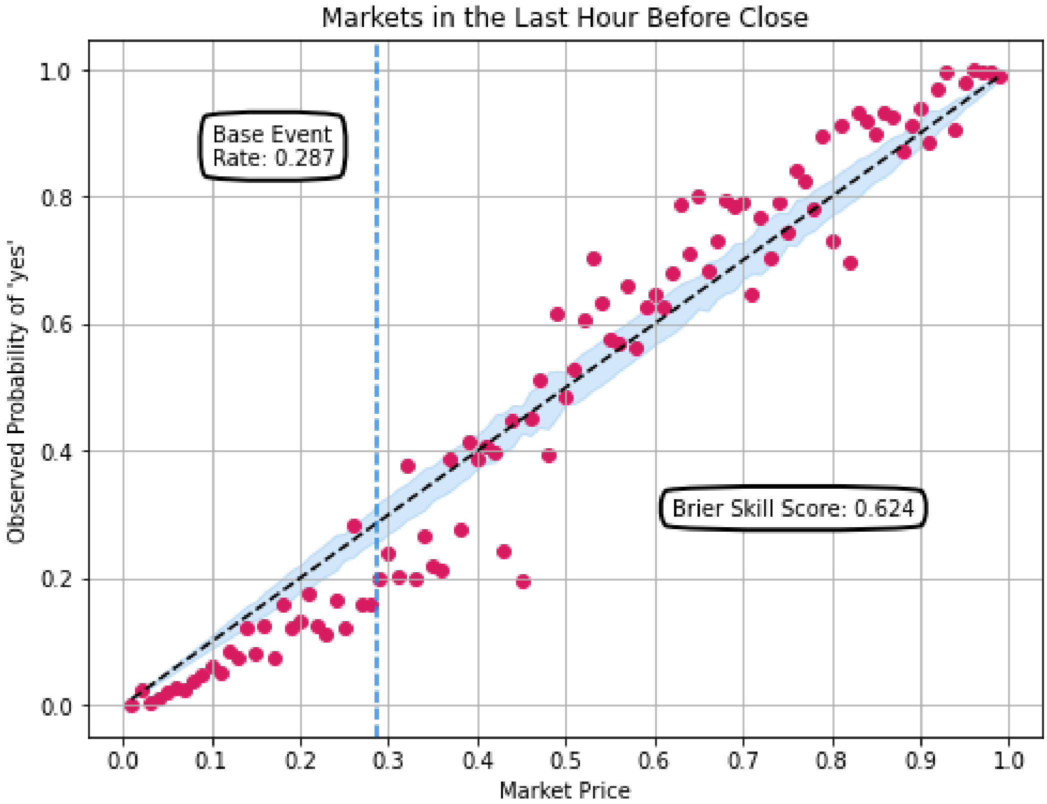

Figure 4. Calibration of Kalshi 'Yes' prices in the last hour before settlement (Clay, 2022)

In a more recent comparative analysis of Kalshi's predictions in oil and yield markets and traditional forecasting methods, Johansson (2024) reported evidence to support that market prices could produce superior forecasts. Rather than comparing direct prices to model outputs, Johansson (2024) interprets Kalshi's range and bucket contracts as discrete probability distributions over the outcomes, to derive a point forecast by computing the corresponding expectations. These forecasts are evaluated against sophisticated benchmarks like ARIMA time-series models and Long Short-Term Memory (LSTM) neural networks using standard error metrics in oil and treasury yield markets. In the study, Johansson (2024) found that Kalshi's high-volume contracts "vastly outperform" these widely used statistical and machine learning benchmarks. This suggests that the wealth-weighted, risk-adjusted aggregation mechanism effectively incorporates dispersed signals that algorithmic models fail to capture.

While theory predicts distortions due to risk aversion, empirical literature consistently identifies a specific and systematic deviation known as the Favorite-Longshot Bias (FLB) (Ottaviani & Sørensen, 2008). The FLB is the primary divergence from perfect probabilistic accuracy (Bürgi et al., 2025). The bias can be characterized through the following inequalities:

$$FLB : \begin{cases} \mathbb{E}['\text{Win}'|\text{Price} = \mathrm{p}_t] < \mathrm{p}_t \\ \mathbb{E}['\text{Win}'|\text{Price} = 1 - \mathrm{p}_t] > 1 - \mathrm{p}_t \end{cases}, for\ high\ \mathrm{p}_t.$$

Essentially illustrating that contracts trading at low prices (e.g., 10c), called longshots, win less than 10% of the time, and contracts trading at high prices (e.g., 90c), favorites, win more than 90% of the time. The FLB can also be noticed in Figure 4, where points for low-price buckets are slightly below the dashed line, and vice versa for high-price buckets, illustrating how longshots are slightly overpriced whereas favorites are slightly underpriced. This phenomenon, which represents a market inefficiency, has been widely documented in betting markets and has also been empirically confirmed to exist in Kalshi's markets by Bürgi et al. (2025).

However, in the context of a regulated, institutional exchange like Kalshi, FLB is better understood as a structural necessity than a behavioral flaw. Hegarty and Whelan (2025) argue that in markets with a Maker-Taker (CDA) microstructure, FLB arises from the market structure itself, specifically because of the interaction of fees, tick sizes and the elasticity of demand. It likely reflects the risk premium demanded by informed liquidity providers (Hegarty & Whelan, 2025), meaning that the bias observed in Kalshi’s prices is the rational cost of liquidity provision in the maker-taker structure, the mathematical result of the fee structure, rather than a failure of the crowd’s intelligence.

In this view, the FLB is simply the expression of the spread required to cover the maker’s inventory risk. Specifically, on low-probability longshot contracts, the 1c tick represents a premium that skews the implied probability higher, resulting in witnessing the FLB. Seen this way, the FLB is arguably necessary for the market to function. If prices were constantly perfectly aligned with true probabilities, there would be no risk-adjusted profit opportunities and little incentive to supply superior information or liquidity. This is the manifestation of the logic proposed by the Grossman-Stiglitz paradox, introduced in Section 3.2.3. In this sense, Kalshi appears to operate in an efficiency sweet spot: prices are accurate enough for serious applications, while remaining just inefficient enough to sustain the environment of informed trading that makes these prices possible in the first place.

Table 1. Summary of empirical accuracy studies on Kalshi markets (Clay, 2022; Johansson, 2024; Bürgi et al., 2025).

Performance Metric	**Theoretical Benchmark**	**Observed Performance on Kalshi**	**Study; data**
Calibration of prices	Prices equal outcome probabilities, events with probability p occur with probability p in the long run.	Prices are tightly clustered around the line of perfect calibration; small deviations only at very low and very high prices.	Clay (2022); 'Yes' prices from 8,476 finalized markets
Brier Skill Score	Positive BSS relative to an uninformed baseline, improving as information arrives.	BSS is consistently greater than 0 several days before settlement and increases with time and new information.	Clay (2022); 'Yes' prices from 8,476 finalized markets
Favorite-Longshot Bias	In a frictionless, risk-neutral market, no systematic under- or overpricing of favorites and longshots.	Prices are "informative and improve in accuracy". Documents a systematic FLB linked directly to the market microstructure.	Bürgi et al. (2025); 300,000+ Kalshi contracts
Comparison to statistical models	An efficient prediction market should at least match and often outperform standard time-series models.	Kalshi forecasts, especially in high-volume markets, "vastly outperform ARIMA and LSTM models".	Johansson (2024), WTI Oil Prices, 10Y Treasury Yields

In conclusion, while Kalshi's prices exhibit a small, structural FLB as predicted by Hegarty and Whelan (2025), the empirical evidence in Table 1 shows that, under thick-market conditions, prices are well-calibrated, display positive and improving forecast skill and closely match or outperform standard models and theoretical benchmarks. Taken together, these metrics indicate that Kalshi's markets come very close to theoretical benchmarks for an informationally efficient PM, with biases interpretable as necessary conditions for healthy markets. Consequently, Kalshi can be regarded as approximating the theoretical optimum as closely as is feasible under real-world constraints such as fees, regulation, and discrete tick sizes.

3.4 Synthesis of benchmark

Synthesizing the theoretical components developed in Sections 3.1, 3.2 and 3.3 and the causal logic laid out in Figure 3, this section breaks down "constrained optimality" as a standard for the evaluation of scalable prediction markets. Section 3.1 motivates mechanism design, specifying how market microstructure and incentive design shape liquidity provision and the costs of launching and sustaining markets. Section 3.2 then frames the efficiency and information-aggregation criteria that follow from thick-market conditions, clarifying what it means for prices to reflect dispersed information under realistic constraints and why imperfections to an extent may be structurally necessary to sustain informational efficiency. Finally, Section 3.3 develops the probabilistic interpretation argument, establishing when prices can be treated as reliable probability estimates and how systematic deviations should be interpreted in a frictional but competitive and incentive-compatible setting.

Table 2 decomposes the three headline dimensions in Figure 3 into nine evaluation dimensions into a structured benchmark rubric. Specifically, it synthesizes theoretical requirements and observable indicators for each dimension for an assessment of Kalshi, backed by the anchoring claims used in the thesis.

Table 2. Benchmark dimensions and assessment of Kalshi

Benchmark dimension	**Definition of "optimal"**	**Observable indicators**	**Assessment; conditions**	**Anchor claim**
1) Mechanism design choice	In high-volume markets, the trading mechanism supports low-cost liquidity and effective price discovery.	CDA/CLOB vs AMM/MSR; maker–taker logic; argument on scalability and subsidy requirements.	Meets; high-volume markets	Kalshi's CDA/CLOB structure is argued to outperform subsidized AMMs on cost-efficiency and price discovery when volume is high.
2) Liquidity formation and sustainability	Thick markets exist (or are feasibly sustainable) without permanent subsidy. Liquidity supports larger participant range and trade sizes.	Role of professional market makers; concentration or network effects; position limits; spread and depth narrative.	Meets; liquidity concentration	Regulation + institutional design are argued to attract participation and enable thick markets, but liquidity may concentrate in a small number of contracts/venues.
3) Institutional reliability and trust	Institutional setup reduces historical constraints (legal uncertainty, counterparty risk) to broaden participation and capital inflow.	CFTC/DCM framing; clearing/settlement structure; collateralization; counterparty-risk logic.	Meets	Regulatory status and clearing/settlement design are positioned as trust infrastructure that supports liquidity and institutional participation.
4) Capital efficiency and cost of liquidity	Design avoids open-ended liquidity subsidy while compensating liquidity provision via stable equilibrium frictions.	Tick/fee design; spread floor; inventory risk; "efficiency sweet spot" logic.	Partially meets	Small frictions (fees/ticks) create bounded micro-inefficiencies but are framed as necessary compensation to sustain liquidity provision and informed participation.

Benchmark dimension	Definition of "optimal"	Observable indicators	Assessment; conditions	Anchor claim
5) Information aggregation and informed trading	Structure enables information revelation by informed traders while limiting manipulation and ensuring credible resolution.	Informed trading pathways; supervision/regulation; manipulation discussion; no-trade / thin market loop framing.	Meets	Allowing expert participation under supervision is argued to strengthen aggregation while managing manipulation and resolution risk.
6) Probabilistic accuracy (calibration)	Prices can be treated as probabilities (wealth-weighted) under thick-market conditions.	Calibration evidence: "prices improve with time"; interpretation of probabilities.	Meets; cited analyses	Existing reported analyses indicate well-calibrated prices (especially late), supporting probability interpretation under liquid conditions.
7) Probabilistic accuracy (forecasting skill)	Market forecasts match or exceed standard statistical models in relevant domains (where evaluated).	Model comparison narrative (e.g., time-series baselines); macro/financial markets framing.	Meets; cited study, high-volume	Reported comparisons suggest high-volume markets can outperform standard forecasting baselines in selected domains.
8) Systematic bias and interpretation (FLB)	Deviations from frictionless probability pricing are explainable, bounded, and structurally interpretable.	FLB evidence; link to tick/fees; liquidity premium framing; "necessary imperfection" logic.	Partially meets; structural framing of bias	Observed bias is reframed as structural liquidity premium driven by microstructure (tick/fees) rather than pure behavioral failure.
9) Scope of validity	Identify where use is defensible and conditional rather than universal; specify assumptions.	Explicit thick-market qualifier; application examples (hedging/risk/governance); external validity limits (thin/niche markets).	Partially meets; validity of secondary sources, high-volume markets	Applications are argued as plausible where liquidity and institutional reliability hold, while acknowledging evidence is concentrated in early/high-volume markets and may not generalize.

4 Kalshi: tending toward the optimum

While the literature extensively debates the relative merits of different market mechanisms, the primary cause of failure for previous PMs has been the inability to sustain liquidity. The self-reinforcing loop of wide spreads discouraging trading, leading to lower information revelation, all caused by the thin market problem, has historically driven away liquidity providers. Kalshi's microstructure makes the case that the solution to this problem is not algorithmic, as in a more optimal AMM, but institutional. By restructuring the prediction market to satisfy the constraints of institutional capital, Kalshi fosters the thick-market conditions required for a Continuous Double Auction (CDA) to function optimally, underpinning the production of well-calibrated prices that can be interpreted as reliable probability estimates. Thus, Section 4 follows Table 2, treating Kalshi's institutional solutions as an independent variable, liquidity as the facilitating condition, and probabilistic accuracy as the dependent outcome.

4.1 Institutional solutions

To answer the question of which features allow Kalshi to function optimally where others have failed, as posed in Research Question 2, first "institutional" must be defined as the specific set of market mechanics that breed trust, capital efficiency and access. In terms of Table 2, this section discusses Dimensions 2, 3 and 4, covering institutional reliability, liquidity formation as well as capital efficiency and cost of liquidity.

As Kalshi is a CFTC-regulated Designated Contract Market (DCM), it is required by law to clear all trades through a CFTC-regulated clearinghouse, referred to as a Designated Clearing Organization (DCO) (Kalshi, 2023b). The foundational result of this structure is regulatory certainty, where a critical component is the elimination of *counterparty risk*, that is the risk of a counterparty defaulting before settlement, through *novation* (Kalshi, 2023b). Novation refers to the process of the regulated clearinghouse, in this case Kalshi Klear, becoming the central counterparty in a transaction, substituting itself for both the buyer and seller (Kalshi, 2023b). This effectively means that the exchange itself is the counterparty, which guarantees solvency and transforms unpredictable counterparty risk into standardized credit risk against the clearinghouse.

Furthermore, all retail trades are fully cash collateralized, making the credit risk and broader systemic risks faced by institutions negligible (Kalshi, 2023b). While institutions thrive on taking market risk, the counterparty risk of trading in an unregulated market remains unacceptable for fund managers. Thus, Kalshi's Designated Contract Market (DCM) status has solved the trust and counterparty risk problems by providing the legal certainty required for regulated funds to deploy capital. These are the core mechanism behind Table 2's assessment that Kalshi's regulatory and clearing setup functions as reliable infrastructure. Consequently, the DCM status fundamentally changes the financial structure of trades on Kalshi, converting "bets" into cleared financial instruments, and removing participation constraints that historically prevented thick-market conditions from forming.

Kalshi also reduces the capital and operational frictions that would otherwise discourage institutional participation. As mentioned, retail trades are fully collateralized and there is no leverage, but for professional MMs and institutional users, central clearing through Kalshi Klear allows exposure to be managed in a familiar way. This includes netting and risk-based margin at the clearinghouse level. This lowers the cost of providing liquidity and supports tighter spreads than would be possible on an uncleared PM. In addition, Kalshi aligns with global financial standards through connectivity. Kalshi offers a FIX (Financial Information eXchange) API, the standard communication protocol in global markets, which allows quantitative funds and HFT firms to connect to Kalshi using existing infrastructure. Together, FIX connectivity and clearing make event contracts like any other exchange-traded derivative from an institutional perspective.

In conclusion, Kalshi's regulatory status should be understood not merely as a forcibly imposed legal necessity, but as an offensive strategic moat designed to secure a thick market equilibrium. In market microstructure theory, liquidity is often fragmented across exchanges; however, by obtaining DCM status, Kalshi puts up a barrier to entry that effectively prevents the fragmentation of its most valuable source of liquidity, institutional capital. While unregulated or decentralized competitors keep capturing retail flow, they are legally blocked from servicing the regulated funds and corporate hedgers whose participation underpins the conditions for the premise of this thesis, mimicking an established financial exchange rather than a volatile betting pool.

Finally, the durability of this thick market equilibrium is reinforced by the proprietary nature of event contracts. As opposed to traditional equities, where a share of stock purchased on the NYSE can be sold on the NASDAQ, Kalshi's event contracts exist solely within its specific regulatory and clearing framework. As a result, a 'Yes' contract on Kalshi cannot be transferred to, or settled on, a competing PM like Polymarket. The absence of fungibility creates a powerful 'winner-takes-all' dynamic driven by network effects, effectively setting up conditions where liquidity converges to the PM with the deepest order book. From this assertion follows the idea that institutions will be faced with a binary choice: to trade on the 'winner' exchange to absorb their size, or not to trade at all. Thus, the market structure naturally trends toward a monopoly or duopoly where the thickest markets become the only viable options for institutional scale.

To summarize, these institutional choices address the two historical hurdles keeping PMs small and experimental. Regulation and clearing make Kalshi a trusted, legally accessible prediction market, while professional market making and concentrated order flow generate the thick markets required for CDA-based price discovery. Table 2 formalizes this as a conditional claim: if institutional trust holds and liquidity is sufficiently thick, then a CDA can realize its advantages in price discovery and enable probabilistic accuracy. In other words, Kalshi solves the fundamental constraints mentioned in the introduction by adopting the institutional microstructure of a modern derivatives exchange.

4.2 Convergence with theory

This institutional structure allows us to assess how Kalshi aligns with theoretical benchmarks, as posed in Research Question 1. Thus, framed against Table 2, this section evaluates Kalshi's convergence primarily through mechanism design choice and probabilistic accuracy. The key question is therefore conditional: whether the market achieves enough scale for CDA price formation to dominate the micro-frictions introduced by fees and discrete ticks

A central tension in this design is the trade-off between granular precision and robust aggregation. While the tick size of 1c creates micro-inefficiencies in terms of ±1% pricing errors, the institutional volume creates macro-efficiency. The wisdom of the crowd on Kalshi may be off by a percentage point at the micro level, but it captures the fundamental state of the world better than AI, algorithms, or other models, as it successfully and effectively aggregates diverse private information that otherwise could not be accessed. As shown by Johansson (2024), Kalshi's high-volume markets vastly outperform sophisticated ARIMA and LSTM models, suggesting that the market's ability to aggregate diverse, private information outweighs the micro-inefficiencies of its fees and tick size.

Furthermore, the FLB observed in Section 3.3 should be interpreted as the price of the institutional solution, not as a failure of the crowd. To get thick markets, you need MMs, and to keep MMs solvent against informed traders, you need spreads and fees (Hegarty & Whelan, 2025). Therefore, the FLB is effectively the operating cost of a sustainable, liquid market: a small, structural wedge that compensates liquidity providers for inventory and information risk. Table 2 separates the usefulness of probabilities (calibration) from frictionless pricing by allowing systematic wedges when they are stable and tied to microstructure. As mentioned in Section 3.3.2, Kalshi appears to be operating in an efficiency sweet spot, not falling short of the frictionless benchmark. Prices are accurate enough to be treated as institution-grade probabilities, yet just imperfect enough to leave a narrow, predictable return to retain informed liquidity providers, in line with the Grossman-Stiglitz paradox, which states that fully and freely informative prices cannot be an equilibrium (Grossman &

Stiglitz, 1980). By tolerating these structural costs, Kalshi avoids the subsidy pitfalls of AMMs, while solving the thin market problem through a design that is both practically sustainable and nearly theoretically optimal.

The causal chain from design to near-optimality, as presented in Figure 3 and formalized in Table 2, is therefore straightforward. Regulation generates trust; trust attracts institutional liquidity; institutional liquidity enables professional market making; professional market making creates thick markets; and thick markets allow the CDA to achieve maximum efficiency, producing well-calibrated and robust probabilities. These reliable probability estimates in turn support the high-value applications discussed in Section 5, ultimately attracting further institutional participation and liquidity, reinforcing the entire cycle.

In conclusion, Kalshi approximates the theoretical optimum of a liquid CDA market, not by being perfect but by being sufficiently robust to attract the liquidity needed for smoothing over those imperfections.

5 Implications and applications

Having shown in Sections 3 and 4 that Kalshi's design and microstructure approximates the theoretical benchmark set out in Section 3 under thick-market conditions, this section examines how such markets can be used in practice. Following the causal chain in Figure 3, the applications discussed below are established as conditional on the scope of validity presented in Table 2 (Dimension 9). This means that applications are strongest in contract categories and time windows where the prerequisites and liquidity conditions assumed in the benchmark plausibly hold. As such, this section focuses on three domains where these reliable probabilities could have direct value: asset management, strategic risk management, and governance.

5.1 Probabilities as an asset class

Given the evidence in Sections 3.3.1 and 3.3.2 that Kalshi prices are well-calibrated and represent a risk-adjusted, wealth-weighted consensus that aggregates private information better than algorithms, these prices can be treated as reliable probability estimates. Section 3.1.3 showed that in thick markets, the CDA mechanism is the cost-efficient choice, that ensures competitive price formation. As formulated in Section 2.3, event contracts can be viewed as binary futures on states of the world, where market prices are implied probabilities for those states. In this sense, Kalshi operates as an exchange for event derivatives, and makes event contracts a distinct asset class whose payoff structure is comparable to, but broader than, traditional options and futures.

This positions Kalshi as a primary source of alternative data, a niche that is growing as quantitative hedge funds constantly seek new ways to generate *alpha*, that is financial returns that are uncorrelated with the broader market. Furthermore, Kalshi's data is not just alternative, it is largely orthogonal in the sense of being weakly correlated with traditional asset returns. In comparison, most alternative data, such as satellite imagery and transaction data, is granular input that still attempts to predict correlated financial outcomes. A Kalshi price is a meta-signal, in the sense that it is a real-time, risk-adjusted, wealth-weighted output of all other available public information, as aggregated by a provably competitive market. Many of the markets listed on Kalshi are based on non-financial events such as political and economic outcomes or weather patterns, which are structurally uncorrelated with traditional asset classes, making the data a clean source of orthogonal signals for diversification and alpha generation. These properties make Kalshi's implied probabilities natural inputs for hedging and scenario analysis, as demonstrated in Sections 5.2 and 5.3. However, traders are faced with the choice of interpreting raw prices, including the small liquidity premium discussed in Section 3.3.2, or an adjusted probability estimate as the relevant input to their models.

As Kalshi offers a FIX API, quantitative funds can both trade contracts directly and incorporate event risk as a real-time data feed into trading strategies. Examples include a macro strategy, comparing distributions implied in Kalshi data to proprietary models, for instance for inflation, and trade the spread. Alternatively, a fund could set up an arbitrage strategy, profiting on the structural FLB by acting as Makers and providing liquidity at the market's extremes and profiting on the risk premium. In addition to using data in trading strategies, funds can use the probabilities as weights in portfolio models, risk management frameworks, or stress tests, subject to the usual basis-risk and liquidity constraints.

5.2 Strategic risk management

Section 3.2 showed that markets with simple binary payoffs are particularly effective at aggregating information, while Section 4.1.1 argued that Kalshi's regulatory and clearing structure enables large positions with minimal counterparty risk. Therefore, event contracts can be used as distinct hedges for non-tradeable risks. Traditional insurance and derivatives cover standard risks such as currency fluctuation, fire, theft and so forth, but they cannot cover event risk such as "Will the FTC block a merger?" or "Will a specific tariff pass?". Other uninsurable risks include risks that are binary or discrete, important for profits, or not hedgeable via standard FX, interest-rates or equity derivatives. Table 3 presents four risk categories, with corresponding examples for events, contracts on event, as well as how and why the event could be hedged against.

Table 3. Framework for event-based corporate hedging

Risk Category	Event Exposure	Illustrative Contract	Hedging Rationale
Political	A battery company's subsidies depend on election results.	"Party X to win congress."	If the unfavorable outcome occurs, contract payout offsets loss of subsidies.
Climate	An events company's revenue is lost if it rains.	"Will rain in NYC be > 0.5 inches on _?"	If it rains, the payout offsets lost revenue.
Supply Chain	A producer relies on components from a port that may close due to a strike.	"Will Port X report > 10 'down days' this month?"	If the port closes, the payout offsets the costs of sourcing more expensive alternatives.
Economic	An investment portfolio is sensitive to high inflation.	"Will the next CPI reading be > 3.5%?"	If inflation rises significantly, the payout acts as a direct, non-correlated hedge.

Hedging positions should be built in correlation with both the cost and the implied probability of the risk. In practice, these hedges are still imperfect in the sense that contract definitions don't always match exposure exactly, so hedgers face basis risk between the event and their profit and loss (P&L).

As Kalshi supports high position limits, up to $7M or more via hedging exemptions, it allows for corporations to effectively buy insurance against these specific political, economic or regulatory outcomes. In terms of Section 2.3, hedgers are effectively buying binary futures on states of the world that matter for their cash flows. This application is only possible because of the thick market solutions discussed in 3.1 and 3.2, as without deep liquidity provided by MMs, such as SIG, a corporation trying to place a $5M hedge would crash the market. Kalshi's institutional structure is what makes this specific utility function possible.

5.3 Governance and policy

Section 3.1.3 argues that CDAs incentivize truth-telling over strategic manipulation. Unlike corporate forecasts or political polling, which are biased due to for instance principal agent problems, market prices are "skin-in-the-game" estimates. Furthermore, Section 3.2.3 argued that allowing expert-based informed trading can improve informational efficiency in markets, if the markets are liquid and robust to manipulation, whereas Section 3.3.2 notes that given the calibration and markets against models comparison, these markets can be treated as serious forecasting tools.

Governments and boards can use these liquid markets as inputs in decision-making, using the forward-looking prices to get insights on the realism of macro events affecting the instances concerned. As an example, an intelligence agency could monitor markets on geopolitical events as an early warning system (Ferris & Ferris, 2025). Other use cases in governance could be a health agency tracking markets on disease spread to allocate resources in advance, or a central bank or ministry tracking inflation, unemployment and default probabilities to complement internal models and surveys. The fact that Kalshi allows hyperlocal, expert-based informed trading means that these prices often capture internal sentiment before it reaches the public or affected institutions, thus acting as a potentially valuable early-warning indicator in many domains.

6 Conclusion

This section concludes the thesis by tying the theoretical benchmarks, empirical evidence and institutional analysis together. It summarizes the main findings in the context of the research questions, reflecting on their broader implications for prediction markets and financial infrastructure, finally outlining key limitations and directions for future research.

6.1 Summary of findings

By comparing Kalshi's CDA-CLOB microstructure and institutional setup to the theoretical benchmarks in Section 3, the thesis argues that in high-volume environments, a regulated CDA is better aligned with mechanism-design and liquidity criteria than subsidized AMMs. This conclusion is bound to markets where liquidity and participation are strong, and where the benchmark conditions are clearly supported by the evidence.

Firstly (RQ1), a CLOB with competitive market makers, central clearing and a broad variety of contracts, allows Kalshi to behave much like a traditional futures exchange for event contracts, while mitigating subsidy and scalability issues of MSR or AMM based designs. Secondly (RQ2), the analysis identifies a set of structural features that are critical for approximating optimal performance. These include professional market makers competing in a maker-taker fee system, discrete ticks that define minimum spreads, fully collateralized positions and high position limits. Together, these features breed a small, structural Favorite–Longshot Bias which, in line with the Grossman–Stiglitz paradox, should be interpreted as the necessary compensation to informed traders and liquidity providers for inventory and information risk. Allowing expert-based informed trading by specialized participants, under CFTC supervision, further strengthens information aggregation while managing the risks of manipulation.

Thirdly (RQ3), existing theory and evidence suggest that, under thick-market conditions, Kalshi's prices can be used as institutional-grade probabilistic inputs for forecasting, decision-support and hedging. This is conditional on strong liquidity and participation as well as credible resolution. Under these conditions, empirical evidence on calibration, Brier Skill Score and FLB indicates that prices can be treated as reliable probability estimates. However, this inference is bound to specific scopes and should not be generalized to thin or niche markets where the benchmark conditions weaken.

Overall, these results suggest that Kalshi is more than a speculative event exchange. By combining a thick, transparent and competitive market structure with institutional solutions and a broad variety of contracts, Kalshi effectively turns collective beliefs about real-world events into a tradeable, institutional-grade data product. In this sense, Kalshi acts as a bridge between speculative markets and information markets. Prices serve both as payoffs for traders and as forward-looking probabilities that in theory could be placed alongside bond yields, interest- and currency rates on Bloomberg terminals. In practice, these points are implications of the benchmark assessment rather than direct empirical findings.

Therefore, the thesis concludes that Kalshi approximates the optimal prediction market as closely as is feasible under real-world constraints and provides a foundation on which future utilities incorporating PMs can be built.

6.2 Discussion

This thesis started from a long-standing stalemate, where theory predicted powerful prediction markets, but practice mostly delivered small, experimental PMs. The analysis suggests that what was missing was an institutional setting, where a standard CDA could perform at scale. Once regulation, clearing and market making were in place, the simple order book could produce liquid, competitive and informative markets that approximate theoretical benchmarks.

The results also show that Kalshi's prices are not merely bets, but serious alternatives to traditional models for various macro variables. These markets complement existing models in the sense that models can flag systematic mispricings, while markets offer real-time, weighted probabilities that models alone cannot produce. As a result, Kalshi's event contracts turn collective beliefs about inflation, policy, climate, or any other significant domain into something that can be traded, hedged and used as input in professional decision-making.

Finally, the thesis argues that efficiency in Kalshi's markets is best interpreted as a sweet spot. A small FLB, 'Maker-Taker' fees, and discrete ticks are not flaws, but the incentive that keeps informed traders and liquidity providers returning to participate, keeping prices accurate and meaningful. Because contracts are proprietary and backed by DCM status, liquidity is likely to concentrate on one or a few PMs, deepening liquidity and improving price discovery, although leading to greater dependence on a small number of venues. Kalshi's set-up and design therefore points to a realistic and sustainable notion of "optimal". Prices are accurate enough for serious applications, yet not so perfect that the incentives to create and trade on information disappear.

6.3 Limitations

Firstly, most of the empirical evidence is drawn from existing studies rather than from raw Kalshi data. Moreover, the findings in this thesis rely in part on documentation and analysis produced or published by Kalshi itself. While these sources provide valuable technical and institutional details, they may reflect a promotional or selective view. In addition, the empirical literature consists of a limited set of early Kalshi studies that mostly cover high-volume, U.S. markets. The conclusion that Kalshi approximates the optimal prediction market may therefore overstate how well the design works in general for thin or niche markets and say little about domains where liquidity is inherently hard to build. Secondly, the conclusions of near optimality are tied to the specific setting of a CFTC-regulated CDA with professional market making and central clearing. Thus, these conditional conclusions do not automatically extend to smaller or decentralized PMs, where different mechanisms or designs are used. However, the objective for this thesis is to show that Kalshi, rather than PMs in general, approximates optimality.

6.4 Future research

This thesis only briefly covers how different contract designs and user behaviour impacts performance in the long-term. Questions such as how to design more robust recurring or perpetual event contracts, how institutions and organizations use given probabilities in practice, and how retail perceptions of manipulation or conflicts of interest affect trust, need their own studies. In addition, if network effects seem to concentrate liquidity and attention toward the biggest exchanges in the future, research should evaluate the regulatory and systemic implications of relying on such venues as part of the core information infrastructure for the high-stakes domains mentioned in this thesis. Addressing these limitations would deepen the understanding of when prediction markets can safely be treated as core information infrastructure rather than dismissed as niche speculative tools.

References

Armantier, O., & Treich, N. (2013). Eliciting beliefs: Proper scoring rules, incentives, stakes and hedging. *European Economic Review*, *62*, 17–40. https://doi.org/10.1016/j.euroecorev.2013.03.008

Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., Levmore, S., Litan, R., Milgrom, P., Nelson, F. D., Neumann, G. R., Ottaviani, M., Schelling, T. C., Shiller, R. J., Smith, V. L., Snowberg, E., Sunstein, C. R., Tetlock, P. C., Tetlock, P. E., … Zitzewitz, E. (2008). The Promise of Prediction Markets. *Science*, *320*(5878), 877–878. https://doi.org/10.1126/science.1157679

Buckley, P., & O'Brien, F. (2017). The effect of malicious manipulations on prediction market accuracy. *Information Systems Frontiers*, *19*(3), 611–623. https://doi.org/10.1007/s10796-015-9617-7

Bürgi, C., Deng, W., & Whelan, K. (2025). Makers and Takers: The Economics of the Kalshi Prediction Market. SSRN. https://doi.org/10.2139/ssrn.5502658

Clay, R. (2022, September 18). *Calibration and skill of the Kalshi prediction markets*. Commonwealth Data Solutions. https://www.cwdatasolutions.com/post/calibration-and-skill-of-the-kalshi-prediction-markets

Chen, Y., Dimitrov, S., Sami, R., Reeves, D. M., Pennock, D. M., Hanson, R. D., Fortnow, L., & Gonen, R. (2010). Gaming Prediction Markets: Equilibrium Strategies with a Market Maker. *Algorithmica*, *58*(4), 930–969. https://doi.org/10.1007/s00453-009-9323-2

Chen, Y., & Pennock, D. M. (2010). Designing Markets for Prediction. *AI Magazine*, *31*(4), 42–52. https://doi.org/10.1609/aimag.v31i4.2313

Choo, L., Kaplan, T. R., & Zultan, R. (2022). Manipulation and (Mis)trust in Prediction Markets. *Management Science*, *68*(9), 6716–6732. https://doi.org/10.1287/mnsc.2021.4213

Corgnet, B., Deck, C., DeSantis, M., Hampton, K., & Kimbrough, E. O. (2023). When Do Security Markets Aggregate Dispersed Information? *Management Science*, *69*(6), 3697–3729. https://doi.org/10.1287/mnsc.2022.4463

DuneData. (2025). *Prediction Markets – Notional Volume Dashboard*. Dune Analytics. https://dune.com/dunedata/prediction-markets

Ferris, S. P., & Ferris, R. M. (2025, July 16). The market knows best: Using data from prediction markets to assess national security threats. *Military Intelligence Professional Bulletin*. https://mipb.ikn.army.mil/issues/jul-dec-2025/the-market-knows-best/

Fama, E. F. (1970). Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*, *25*(2), 383–417. https://doi.org/10.2307/2325486

Galanis, S., Ioannou, C. A., & Kotronis, S. (2024). Information Aggregation Under Ambiguity: Theory and Experimental Evidence. *Review of Economic Studies*, *91*(6), 3423–3467. https://doi.org/10.1093/restud/rdae009

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *The American Economic Review*, *70*(3), 393–408.

Hanson, R. (2003). Combinatorial Information Market Design. *Information Systems Frontiers*, *5*(1), 107–119. https://doi.org/10.1023/A:1022058209073

Hanson, R. (2012). Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation. *The Journal of Prediction Markets*, *1*(1), 3–15. https://doi.org/10.5750/jpm.v1i1.417

Harris, L. E. (1994). Minimum Price Variations, Discrete Bid–Ask Spreads, and Quotation Sizes. *Review of Financial Studies*, *7*(1), 149–178. https://doi.org/10.1093/rfs/7.1.149

He, X.-Z., & Treich, N. (2017). Prediction market prices under risk aversion and heterogeneous beliefs. *Journal of Mathematical Economics*, *70*, 105–114. https://doi.org/10.1016/j.jmateco.2017.02.005

Hegarty, T., & Whelan, K. (2025). Market structure and prices in online betting markets: Theory and evidence. *Oxford Economic Papers*, gpaf023. https://doi.org/10.1093/oep/gpaf023

Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, *9*(1), 47–73. https://doi.org/10.1016/0304-405X(81)90020-9

Johansson, F. (2024). Markets vs. Machines: An investigation of prediction market forecasting accuracy of economic variables compared to time series forecasting methodologies [Master's thesis, Uppsala University]. DiVA Portal. https://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-529868

Kalshi. (2020, November 4). *Kalshi wins CFTC approval for exchange to trade event contracts*. Kalshi News. https://news.kalshi.com/p/kalshi-designation

Kalshi. (2023a, May 2). *Prediction markets vs. CME Fed Watch: A new age of forecasting Federal Reserve actions.* Kalshi News. https://news.kalshi.com/p/prediction-markets-vs-cme-fed-watch-a-new-age-of-forecasting-federal

Kalshi. (2023b, May 5). *Who am I trading with on Kalshi?* Kalshi News. https://news.kalshi.com/p/who-am-i-trading-against-on-kalshi

Kalshi. (2023c, June 12). *0DTE on Kalshi*. Kalshi News. https://news.kalshi.com/p/0dte-on-kalshi

Kalshi. (2024, April 3). *Liquid prediction markets are (finally) here...* Kalshi News. https://news.kalshi.com/p/liquid-prediction-markets-are-finally-here

Kalshi Klear LLC. (2025, July 18). *Kalshi Klear rulebook* (Version 1.3). U.S. Commodity Futures Trading Commission.

Karimi, M., & Dimitrov, S. (2024). To Subsidize or Not to Subsidize: A Comparison of Market Scoring Rules and Continuous Double Auctions for Price Discovery. *Information Systems Frontiers*, *26*(2), 801–823. https://doi.org/10.1007/s10796-023-10384-8

Linnainmaa, J. T. (2010). Do Limit Orders Alter Inferences about Investor Performance and Behavior? *The Journal of Finance*, *65*(4), 1473–1506. https://doi.org/10.1111/j.1540-6261.2010.01576.x

Luckner, S., Schröder, J., Slamka, C., Franke, M., Geyer-Schulz, A., Skiera, B., Spann, M., & Weinhardt, C. (2012). *Prediction Markets: Fundamentals, Designs, and Applications*. Gabler Verlag. https://doi.org/10.1007/978-3-8349-7085-5

Morrow, A. (2024, November 8). How prediction markets saw something the polls and pundits didn't. *CNN.* https://edition.cnn.com/2024/11/08/business/polymarket-election-trump-nightcap

Othman, A., Pennock, D. M., Reeves, D. M., & Sandholm, T. (2013). A Practical Liquidity-Sensitive Automated Market Maker. *ACM Transactions on Economics and Computation*, *1*(3), 1–25. https://doi.org/10.1145/2509413.2509414

Ottaviani, M., & Sørensen, P. N. (2008). The Favorite-Longshot Bias: An Overview of the Main Explanations. In *Handbook of Sports and Lottery Markets* (pp. 83–101). Elsevier. https://doi.org/10.1016/B978-044450744-0.50009-3

Park, A. (2025, September 15). Why billionaires are piling into prediction markets. *Forbes.* https://www.forbes.com/sites/aliciapark/2025/09/15/why-billionaires-are-piling-into-prediction-markets/

Richard, M., & Vecer, J. (2021). Efficiency Testing of Prediction Markets: Martingale Approach, Likelihood Ratio and Bayes Factor Analysis. *Risks*, *9*(2), 31. https://doi.org/10.3390/risks9020031

Tziralis, G., & Tatsiopoulos, I. (2007). Prediction Markets: An Extended Literature Review. *The Journal of Prediction Markets*, *1*(1), 75–91. https://doi.org/10.5750/jpm.v1i1.421

Waghmare, K., & Ziegel, J. (2025). *Proper scoring rules for estimation and forecast evaluation* (No. arXiv:2504.01781). arXiv. https://doi.org/10.48550/arXiv.2504.01781

Wah, L. (2016). An Empirical Game-Theoretic Analysis of Price Discovery in Prediction Markets. IJCAI International Joint Conference on Artificial Intelligence, 510-516.

Whelan, K. (2023). On prices and returns in commercial prediction markets. *Quantitative Finance*, *23*(11), 1699–1712. https://doi.org/10.1080/14697688.2023.2257756

Whelan, K. (2025). Agreeing to Disagree: The Economics of Betting Exchanges (CEPR Discussion Paper No. 20633). CEPR. https://cepr.org/publications/dp20633

Wolfers, J., & Zitzewitz, E. (2004). Prediction Markets. *Journal of Economic Perspectives*, *18*(2), 107–126. https://doi.org/10.1257/0895330041371321

Wolfers, J., & Zitzewitz, E. (2006). *Interpreting Prediction Market Prices as Probabilities* (NBER Working Paper No. 12200). National Bureau of Economic Research. https://doi.org/10.3386/w12200