

Blockchain-Based Prediction Markets: A Survey of Architectures and Economic Mechanisms

Sigit Dani Perkasa ^{a,1,*}, Fachrizal Fikri ^{b,2}

^a Centre for AI, Robotics, and Manufacturing Strategy (CARMS), MECATA Foundation, Bogor 16730, Indonesia

^b Islamic Economy and Halal Industry, Universitas Negeri Surabaya, Surabaya 60231, Indonesia

¹ sigitdaniperkasa@gmail.com; ² 24081725011@mhs.unesa.ac.id

* Corresponding Author

ARTICLE INFO

Article history

Received June 15, 2026

Revised July 24, 2026

Accepted August 11, 2026

Keywords

Automated Market Maker;

Blockchain;

Decentralized Finance;

Prediction Markets;

Oracle

ABSTRACT

Prediction markets harness dispersed knowledge to forecast uncertain events. Blockchain technology has enabled a new generation of permissionless, oracle-dependent prediction platforms with diverse architectural and economic designs. To survey the engineering architectures, market microstructure, oracle designs, and economic mechanisms of operational blockchain-based prediction market platforms, using a structured comparative case-study approach. A documented platform-selection protocol (informed by Chong et al. [2019] and Treiblmaier [2019] for blockchain case-study research) was applied across DeFi literature, industry documentation, and technical whitepapers. Six blockchain-native platforms and one centralized regulated benchmark were selected based on market relevance, architectural diversity, and documentation availability. The six blockchain-native platforms—Polymarket, Augur, Gnosis Conditional Tokens/Omen, Azuro, SX Bet, and Overtime Markets—span three dominant architectural paradigms: the centralized limit order book with decentralized settlement, the fully on-chain automated market maker, and the modular conditional token infrastructure. Oracle design emerges as the critical security bottleneck across all platforms, with each adopting distinct trust assumptions and economic security models. Liquidity models range from hierarchical pooled AMMs to competitive order books to parimutuel systems, with no single architecture dominant across applications. Conclusions: The contrast between blockchain-native platforms and the CFTC-regulated Kalshi illustrates the fundamental trade-off between cryptographic trust minimization and regulatory compliance. These findings provide a structured framework for understanding the design space of blockchain-based prediction markets and identify open challenges including oracle security, cross-chain liquidity aggregation, and the integration of prediction market signals with autonomous decision-making systems.

©2025 The Author. This is an open-access article under the CC-BY license.



1. Introduction

Prediction markets allow participants to trade contracts whose payoffs depend on uncertain future events [1], [2]. Hayek [1] argued that prices communicate dispersed knowledge across individuals, forming the intellectual foundation of these markets. Wolfers and Zitzewitz [2] demonstrated that

well-designed prediction markets outperform alternative forecasting methods including expert panels and opinion polls, particularly with sufficient liquidity and informed traders.

Arrow et al. [3] articulated the policy relevance of prediction markets as a decentralized, incentive-compatible approach to information elicitation with applications from monetary policy to public health surveillance. However, centralized deployment faced regulatory constraints, jurisdictional fragmentation, and counterparty risk.

Hanson [4] introduced the Logarithmic Market Scoring Rule (LMSR), an automated market maker (AMM) providing bounded-loss liquidity across all possible event outcomes. The LMSR is defined by a cost function and a corresponding price function. The cost function determines the total capital required to maintain a given vector of outstanding shares:

$$C(q) = b \ln(\sum_{i=1}^n e^{q_i/b}) \quad (1)$$

where q_i denotes the net quantity of shares outstanding for outcome i , b is the liquidity parameter controlling price sensitivity to trading activity, and n is the number of mutually exclusive outcomes. The corresponding price function—which gives the marginal cost of purchasing an infinitesimal additional share of outcome i —is obtained by taking the partial derivative of the cost function with respect to q_i :

$$p_i(q) = \frac{e^{q_i/b}}{\sum_{j=1}^n e^{q_j/b}} \quad (2)$$

Equation (2) has the natural interpretation of a Boltzmann distribution over outcomes, where the exponentiated quantities act as unnormalized log-probabilities [4]. Under LMSR, a trader purchasing shares of outcome increases, which in turn increases—the mechanism by which the market price responds to the information revealed through trading $i q_i p_i b \ln n$

Blockchain technology has re-enabled prediction markets at scale without centralized intermediaries, using smart contracts for fund custody, market resolution, and algorithmic market making. These platforms draw on the foundations of Hayek [1], Wolfers and Zitzewitz [2], and Hanson [4], while introducing challenges in oracle trustworthiness, dispute resolution, and cross-chain interoperability.

This paper presents a structured comparative survey of six blockchain-native prediction market platforms and one centralized regulated benchmark—Polymarket, Augur, Gnosis Conditional Tokens/Omen, Azuro, SX Bet, Overtime Markets, and Kalshi—spanning diverse architectural paradigms. Using a structured methodology, we develop a taxonomy across four dimensions: blockchain infrastructure, market microstructure, oracle and dispute resolution, and governance. The survey identifies systematic trade-offs between decentralization and market efficiency, examines oracle design as a critical security bottleneck, and analyzes convergence toward modular settlement standards exemplified by the Gnosis Conditional Token Framework.

This survey advances beyond existing work in three respects. First, unlike prior blockchain prediction market surveys that either aggregate industry documentation without peer-reviewed methodology [7] or focus on a single platform's ecosystem [8], we adopt a documented comparative case-study approach grounded in the blockchain research methodology guidelines of Chong et al. (2019) and Treiblmaier (2019). Second, we construct a four-dimensional taxonomy—spanning blockchain infrastructure, market microstructure, oracle and dispute resolution, and governance and tokenomics—that enables systematic cross-platform comparison rather than isolated platform descriptions. Third, by including the CFTC-regulated centralized exchange Kalshi alongside six blockchain-native platforms, we empirically delineate the design-space frontier between cryptographic trust minimization and institutional regulatory compliance, a dimension absent from prior surveys that treat blockchain-based and traditional prediction markets as separate domains.

2. Method

This study adopts a structured comparative case-study approach, following the methodology of Chong et al. (2019) and the case-study guidelines proposed by Treiblmaier (2019) for blockchain research. Rather than claiming an exhaustive systematic review—which, as Schaupp et al. (2023) note, would be inappropriate given the field's rapid evolution and reliance on industry-produced documentation—we document a transparent platform-selection protocol. Platforms were selected based on market relevance (measured by trading volume and total value locked), architectural diversity (centralized CLOB vs. on-chain AMM vs. hybrid), and the availability of publicly verifiable technical documentation. The resulting set comprises six blockchain-native prediction market platforms and one centralized regulated benchmark (Kalshi), consistent with the comparative-survey approach used by Schär (2021) for DeFi protocols.

2.1. Research Questions

The survey is guided by three research questions derived from the theoretical foundations discussed in the Introduction and refined through an iterative scoping process.

RQ1: What are the key engineering architectures employed by operational blockchain-based prediction market platforms, and how do they differ in terms of smart contract design, oracle integration, and dispute resolution mechanisms?

RQ2: What economic mechanisms—including liquidity provisioning, fee structures, market scoring rules, and tokenomics—govern the operation of these platforms, and what are their implications for market efficiency and trader participation?

RQ3: How do architectural and economic design decisions trade off against one another along dimensions of decentralization, security, scalability, and user experience?

2.2. Platform Identification and Selection

The platform set was assembled through a documented selection protocol following the comparative case-study methodology of Chong et al. (2019) and the case-study guidelines proposed by Treiblmaier (2019) for blockchain research. Candidate platforms were identified from the prediction market literature, augmented by industry analytics (Dune Analytics, DeFi Llama, CoinGecko) and platform documentation, then filtered against explicit inclusion criteria (defined in the next paragraph). The selected set is not exhaustive of all blockchain-based prediction market activity; rather, it represents seven platforms chosen to span the dominant architectural paradigms documented in the field, following the selection logic of Schär (10) for DeFi protocols. The number of database records screened is therefore not reported, as the platform set was purposively sampled rather than derived from a database search.

Inclusion criteria: (i) facilitates trading of event-contingent contracts; (ii) uses smart contracts for settlement, creation, or liquidity; (iii) operational with verifiable trading volume; (iv) publicly accessible documentation exists. Excluded were platforms with exclusively centralized order matching, ceased operations, or deprecated contracts.

Seven platforms were selected: Polymarket, Augur, Azuro, SX Bet, Overtime Markets (Thales), Gnosis CTF/Omen, Kalshi. This selection spans AMMs, CLOBs, and parimutuel systems across Ethereum, Polygon, Gnosis Chain, Arbitrum, and Solana.

2.3. Data Extraction and Analysis

A standardized extraction form [6] categorized platform attributes across four dimensions: (a) engineering architecture (blockchain, smart contracts, oracle, dispute resolution); (b) market microstructure (order matching, liquidity, fees); (c) economic design (tokenomics, incentives, scoring rules); and (d) governance (upgrades, parameter rights, treasury).

The reference metric for predictive accuracy throughout this survey is the Brier Score [2], a strictly proper scoring rule that quantifies the calibration between forecast probabilities and observed outcomes. The Brier Score is defined as (3):

$$BS = \frac{1}{N} \sum_{t=1}^N (f_t - o_t)^2 \quad (3)$$

where $f_t \in [0,1]$ is the forecast probability assigned to the event at time t , $o_t \in \{0,1\}$ is the realized binary outcome, and N is the total number of forecast-event pairs evaluated. The Brier Score ranges from 0 (perfect accuracy) to 1 (maximum inaccuracy), with lower values indicating better calibration. As a proper scoring rule, it incentivizes forecasters to report their true subjective probabilities [2].

Data sources included whitepapers, technical documentation, deployed contract code, academic publications, and industry reports. On-chain data supplemented incomplete documentation; each data point was recorded with its source and retrieval date.


Limitations include: platform architectures may have evolved since review (version-specific documentation was prioritized), reliance on public documentation may miss proprietary features, and excluding undocumented platforms may bias toward transparent projects [5].

The taxonomy framework shown in Table 1 organizes platforms along four architectural dimensions derived from the structured review methodology.

Table 1. Taxonomy Framework of Blockchain-Based Prediction Market Platforms Across Four Architectural Dimensions

Platform	Architecture Type	Blockchain Layer	Oracle Mechanism	Liquidity Model
Polymarket	CLOB with on-chain settlement	Polygon (PoS sidechain)	UMA Optimistic Oracle + Chainlink	Professional market makers (CLOB)
Augur	Fully on-chain AMM	Ethereum mainnet	REP staking + Kleros arbitration	AMM (constant-product)
Gnosis/Omen	Modular CTF + AMM	Gnosis Chain (EVM)	reality.eth crowdsourced oracle	AMM (fixed-product)
Azuro	Middleware Liquidity Tree	Polygon, Gnosis Chain	Azuro oracle framework	Hierarchical pooled AMM
SX Bet	CLOB on dedicated L3 rollout	SX Network (Arbitrum Orbit)	SX oracle + Chainlink	Order book (CLOB)
Overtime Markets	Parimutuel AMM	Optimism, Arbitrum	Chainlink decentralized oracle	Parimutuel prize pool
Kalshi	Centralized CLOB (regulated)	Centralized (custodial)	CFTC regulation	Centralized order book

The taxonomy reveals a spectrum of architectural choices spanning from fully decentralized on-chain AMMs (Augur) to centralized regulated markets (Kalshi) (see Table 1 and Fig. 1).



¹ survey of eight operational platforms. CLOB = Central Limit Order Book; AMM = Automated Market Maker; CTF = Conditional

Fig. 1. Taxonomy framework of blockchain-based prediction market platforms across four architectural dimensions: blockchain infrastructure, market microstructure, oracle and dispute resolution, and governance and tokenomics

3. Results and Discussion

3.1. Platform Survey: Group A — Polymarket, Augur, and Gnosis

This section presents Group A—Polymarket, Augur, and Gnosis CTF—defining three dominant engineering patterns: CLOB with decentralized settlement, fully on-chain AMM, and modular conditional token infrastructure [7], [8].

3.1.1. Polymarket

Polymarket is the largest blockchain prediction market by volume, with over \$9B in cumulative trading since 2020 on the Polygon proof-of-stake sidechain [8], [9], [10]. Its architecture integrates a centralized limit order book (CLOB) for off-chain matching, the Gnosis Conditional Token Framework (CTF) for on-chain settlement, and the UMA Optimistic Oracle for outcome resolution [11], [12], [13].

The centralized limit order book architecture represents a deliberate engineering trade-off between decentralization and market efficiency. Unlike fully on-chain automated market makers (AMMs) that programmatically determine prices from liquidity pools, Polymarket's CLOB matches buy and sell orders off-chain through a traditional price-time priority queue [11]. This design enables the high-throughput, low-latency trading experience familiar to users of centralized exchanges, while maintaining the non-custodial settlement and transparent resolution guarantees of blockchain infrastructure. The CLOB operates on a request-for-quote (RFQ) model: market makers submit resting limit orders, takers submit market or limit orders, and matched trades are settled on-chain through the CTF smart contracts deployed on Polygon [11], [12].

The Conditional Token Framework, originally developed by Gnosis, provides the on-chain infrastructure for creating and settling outcome tokens [12]. Under the CTF model, a prediction market condition is defined by a question, a set of mutually exclusive outcomes, and a resolution time. For each condition, the framework mints conditional tokens that represent claims on a unit of collateral contingent on the realized outcome. In a binary market—the dominant market type on Polymarket—two outcome tokens are created for each unit of collateral, representing the "Yes" and "No" outcomes. These tokens are fungible ERC-1155 assets that can be freely transferred, composed into multi-outcome positions, and redeemed at settlement [11], [12]. The CTF architecture enables several desirable properties: collateral is held in auditable smart contracts rather than by a central operator; positions can be split and merged atomically; and market creators can define arbitrary outcome spaces without modifying the core contract logic.

Event outcome resolution on Polymarket is handled by UMA's Optimistic Oracle, a decentralized truth-determination mechanism that operates on an economic security model [13]. Under the optimistic oracle paradigm, a proposer submits a resolution answer backed by a financial bond. During a challenge period—typically two to four hours—any participant may dispute the proposed answer by posting their own bond, escalating the dispute to UMA's Data Verification Mechanism (DVM). The DVM is a Schelling-point-based voting system in which UMA token holders vote on the correct outcome, with the majority-bonded answer prevailing [13].

The economic security of the UMA DVM derives from the expected payoff structure for proposers and disputers. A rational proposer submits a truthful answer when the expected return from doing so is positive:

$$E[\pi] = P(\text{correct}) \cdot B - P(\text{incorrect}) \cdot B \cdot (1 + \delta) \quad (4)$$

where B is the bond amount posted by the proposer or disputer, δ is the penalty multiplier applied to incorrect disputes (typically $\delta = 1.0$ in UMA's implementation, resulting in a total loss of $2B$ for incorrect challenges), and $P(\text{correct})$ and $P(\text{incorrect}) = 1 - P(\text{correct})$ represent the probabilities that the proposed answer is accepted or rejected by DVM voters [13]. When $\delta > 0$ and $P(\text{correct}) > 0.5$, the expected payoff is positive for truthful proposers and negative for frivolous disputers, creating a game-theoretic equilibrium that favors honest reporting.

This mechanism aligns economic incentives with truthful reporting: proposers profit when their answers are accepted and lose their bond when successfully challenged. Caldarelli [14] characterizes this approach as an instance of the broader blockchain oracle problem, in which off-chain truth must be reliably imported into an on-chain execution environment.

Polymarket's deployment on Polygon rather than Ethereum mainnet addresses the scalability limitations that constrained earlier decentralized prediction market implementations. Polygon, as a proof-of-stake sidechain, offers substantially higher throughput and sub-second block confirmation times compared to Ethereum mainnet, with transaction fees typically below one cent, making it economically viable for retail prediction market participation [10]. Transaction fees on Polygon are typically fractions of a cent, making it economically viable for retail traders to execute small positions—a critical requirement for broad prediction market participation. The platform also integrates with Chainlink's decentralized oracle network for certain price feeds and market data, complementing the UMA oracle for event resolution [15]. The broader oracle infrastructure supporting such designs has been systematically surveyed by Al-Breiki et al. [16], who identify data authenticity, node decentralization, and economic incentive alignment as the three principal dimensions of oracle trustworthiness.

The economic design of Polymarket diverges from the automated market maker model formalized by Hanson [4] and analyzed in the context of decentralized finance by Wang [17]. Rather than employing a logarithmic market scoring rule or constant function market maker, Polymarket relies on a network of professional market makers who provide liquidity through the CLOB. These market makers compete on spread, with narrower spreads attracting greater order flow. The platform charges no trading fees to takers; revenue is generated through market creation fees and, as announced in 2024, an optional fee model for certain market types. This zero-fee taker model, combined with the CLOB's price-time priority, creates market microstructure incentives that reward liquidity provision and penalize information asymmetry in ways that differ from AMM-based designs [17], [18].

Empirical research on Polymarket's predictive accuracy and market efficiency has grown substantially. The information aggregation efficiency of prediction markets is formally characterized by the rational expectations property: under efficient market conditions, the equilibrium price equals the expected value of the event outcome conditional on all information reflected in the market:

$$E[V \mid P_{\text{market}}] \quad (5)$$

This expression formalizes the core theoretical claim of prediction markets—that market prices are sufficient statistics for the dispersed private information held by all participants [1], [2]. When equation (5) holds, no trader can systematically profit from information already incorporated into the market price, and the price provides an optimal forecast of the event outcome given publicly available information.

Reichenbach and Walther [19] conducted a systematic analysis of Polymarket's forecasting performance across multiple event categories, finding that Polymarket probabilities calibrate well against observed outcome frequencies and that market-implied probabilities exhibit a mean absolute error comparable to established polling-based forecasts. Their analysis also identified evidence of a favorite-longshot bias—the tendency for low-probability outcomes to be overpriced relative to their empirical frequency—consistent with biases documented in traditional prediction and betting markets [2]. Cordoba Otalora and Themistocleous [20] applied cross-correlation and dynamic time warping techniques to Polymarket data surrounding the 2024 United States presidential election, demonstrating that market prices responded to political events with measurable latency and that price movements anticipated polling shifts by intervals of 24 to 72 hours.

3.1.2. Hybrid Brier Score Evaluation

To complement Reichenbach and Walther [19]'s analysis of Polymarket's forecasting accuracy, which reports strong calibration between market-implied probabilities and observed outcome frequencies across multiple event categories, we examine the platform's published aggregate Brier score metrics. As a strictly proper scoring rule, the Brier score quantifies the mean squared error

between forecast probabilities and binary outcomes, and is decomposable into calibration and refinement components [2]. Polymarket's official accuracy dashboard (polymarket.com/accuracy, accessed June 2026) reports an aggregate Brier score of 0.0627 across 260,487 resolved markets, with forecast accuracy reaching 90.03% one month before resolution and 98.55% within four hours of resolution. Table 2 presents the Brier score stratified by market liquidity tier, revealing a monotonic improvement in calibration as liquidity depth increases—from 0.0647 in markets with less than \$1,000 in liquidity to 0.0156 in markets exceeding \$1 million.

Table 2. Polymarket Brier Score by Liquidity Tier. Source: Polymarket.Com/Accuracy (Accessed June 2026); Lower Brier Scores Indicate Better Calibration. Aggregate Brier Score Across All Markets: 0.0627

Liquidity Tier	No. Markets	Avg. Forecast %	Brier Score
<\$1k	193,869	25.1%	0.0647
\$1k–\$100k	24,403	22.3%	0.0534
\$100k–\$500k	4,041	23.7%	0.0372
\$500k–\$1M	888	24.3%	0.0216
>\$1M	1,286	23.7%	0.0156

Trader-level analysis by Akey et al. [21] examined the distribution of profits and losses among Polymarket participants, finding that a small cohort of sophisticated traders consistently captured positive returns while the majority of retail participants realized net losses. This pattern is consistent with the theoretical prediction that prediction markets reward informed traders at the expense of noise traders, although the authors note that the magnitude of return dispersion on Polymarket exceeds that observed in traditional financial markets, potentially reflecting the platform's low barriers to entry and the heterogeneous information environment of event-based trading. Chen et al. [22] investigated the behavioral dimension of Polymarket participation, analyzing the interaction between political conviction and profit-seeking motives. Their findings suggest that traders with strong directional political views systematically overweigh the probability of ideologically congenial outcomes, creating exploitable inefficiencies for non-ideological arbitrageurs—a result with implications for the design of market interfaces and the interpretation of prediction market prices as unbiased information aggregators.

3.1.3. Augur

Augur, launched in 2018 following a 2015 ICO, implements a fully on-chain AMM where all trading, liquidity, and dispute resolution occur through Ethereum smart contracts [23]. Its design prioritizes censorship resistance and trust minimization, accepting throughput and latency trade-offs inherent to fully on-chain operation.

Augur uses the REP token for oracle staking and dispute resolution [23]. REP holders stake on the correct outcome during a multi-round dispute process. Uncontested outcomes become final after a dispute window; challenged outcomes escalate through rounds with geometrically increasing stake requirements, culminating in a REP token fork if consensus fails [4], [24].

Augur's AMM uses a constant-product invariant rather than LMSR [17], eliminating professional market makers but introducing slippage and capital efficiency challenges, as locked liquidity cannot be reallocated during market duration [17], [23], [25].

Augur's integration with the Kleros decentralized arbitration protocol provides an alternative dispute resolution pathway for certain market types. Kleros operates as a multi-round, Schelling-point-based adjudication system in which jurors are randomly selected from a pool of token stakers and incentivized to vote coherently with the majority [26]. Ortolani [26] provides a comprehensive survey of Kleros's design and its application to blockchain oracle disputes, while Guillaume and Zhuk [27] compare Kleros's performance against alternative decentralized justice mechanisms, finding that its cryptoeconomic incentive structure produces coherent outcomes in low-to-moderate stakes disputes but faces scalability challenges when dispute values approach or exceed the total value of staked juror tokens.

Augur v2 introduced DAI denomination, streamlined dispute resolution, and 0x protocol integration [23], but trading volumes remain below Polymarket's due to latency and cost disadvantages of fully on-chain architectures [28].

3.1.4. Gnosis Conditional Tokens and Omen

The Gnosis CTF [12] serves as infrastructure rather than a standalone platform, providing standardized smart contracts for creating, splitting, merging, and settling conditional outcome tokens on EVM-compatible chains. This modular design separates settlement logic from interface and matching layers, adopted by Polymarket and Omen as a shared settlement standard [8], [12].

The CTF [12] formalizes a condition as a question with mutually exclusive outcomes. Core operations `splitPosition` and `mergePosition` decompose collateral into equal-value conditional tokens (one per outcome) and recombine them for redemption. This enables multi-outcome, scalar, and combinatorial markets [4], [12].

Omen, the primary CTF-based interface deployed on Gnosis Chain [12], uses an AMM with a fixed-product formula and the `reality.eth` crowdsourced oracle [8]. The Gnosis ecosystem positions prediction markets as InfoFi infrastructure, where market-implied probabilities serve as inputs to automated decision-making and AI agents autonomously participate via the CTF as a standardized API [29].

The Polymarket-CTF relationship illustrates the decoupling of settlement from trading infrastructure: Polymarket uses audited CTF contracts [12] while maintaining interoperability with other CTF-compatible interfaces, contrasting with Augur's vertically integrated design [7], [23].

3.1.5. Cross-Platform Architecture Comparison

The three platforms examined in this group illustrate a spectrum of architectural trade-offs in blockchain-based prediction market design (see Fig. 1). Table 3 presents an architectural comparison of the three Group A platforms: Polymarket, Augur, and Gnosis/Omen.

The comparison reveals distinct engineering trade-offs along the decentralization-efficiency frontier (see Fig. 2).

Three architectural patterns emerge from this comparison, spanning market microstructure paradigms, oracle security models, and modular settlement infrastructure.

Table 3. Architectural Comparison of Polymarket, Augur, And Gnosis/Omen Across Key Design Dimensions

Attribute	Polymarket	Augur	Gnosis/Omen
Architecture Type	CLOB + CTF settlement	Fully on-chain AMM	Modular CTF + AMM
Blockchain	Polygon (PoS)	Ethereum mainnet	Gnosis Chain
Oracle	UMA Optimistic + Chainlink	REP staking + Kleros	<code>reality.eth</code>
AMM/CLOB	CLOB (off-chain matching)	AMM (constant-product)	AMM (fixed-product)
Fee Model	Zero-fee taker, market creation fees	Trading fees	Trading fees
Token	None	REP (essential)	None (GNO for governance)
Liquidity	Professional market makers	LP-provided AMM	LP-provided AMM

3.2. Platform Survey: Group B — Azuro, SX Bet, Overtime, and Kalshi

Group B platforms—Azuro, SX Bet, Overtime Markets, and Kalshi—extend the taxonomy into sports-focused and regulated markets. While Group A addresses general event forecasting, Group B introduces specialized liquidity models, Layer 2 scaling architectures, and a regulated CLOB benchmark (Kalshi).




Fig. 2. Decentralization-efficiency frontier for seven blockchain-based prediction market platforms. The convex frontier line illustrates the fundamental trade-off between decentralization and market efficiency

3.2.1. Azuro

Azuro functions as a middleware layer, aggregating liquidity from providers and distributing it to front-end applications that compete on user experience [30], [31]. This "liquidity-as-a-service" model [31] distinguishes it from vertically integrated platforms such as Polymarket and Augur.

Azuro's Liquidity Tree [30] comprises three smart contract tiers: base liquidity pools for specific categories, a core aggregation and allocation layer, and individual market contracts executing AMM logic. This hierarchical design enables dynamic rebalancing, provider diversification, and immediate liquidity for new markets [30], [31].

Azuro's AMM [17], [25] uses a custom pricing function calibrated to sports outcome statistics, incorporating a vigorish (built-in margin analogous to sports bookmaking overround) [31]. Unlike LMSR's bounded-loss model [4], Azuro distributes risk across the liquidity pool collectively, with providers bearing proportional exposure [17].

Azuro deploys on Polygon and Gnosis Chain using the AZUR token for governance and fee distribution [30], [31], [10].

Azuro's liquidity tree addresses the "cold start" problem by enabling new markets to inherit liquidity from established pools, though pooling introduces adverse selection risk mitigated by dynamic fee adjustments and exposure limits [30], [31].

3.2.2. SX Bet

SX Bet combines a sports prediction market with a dedicated Layer 2 scaling solution on the SX Network, an Arbitrum Orbit rollup [32], [33]. It employs a central limit order book, aligning it architecturally with Polymarket while differing in scaling strategy and market focus.

The SX Network is an Arbitrum Orbit chain (Layer 3 rollup settling to Arbitrum One) [33], providing dedicated blockspace for prediction market transactions without competing with general DeFi activity. Completed in June 2024, this transition from Polygon illustrates the range of Layer 2 architectural choices for prediction market developers.

SX Bet's order book enables any participant to submit limit orders [32], emphasizing accessibility over Polymarket's professional market maker model. The SX token serves as both gas token and fee currency [32], [33].

SX Network uses a centralized sequencer [33], introducing operator risk: the sequencer could theoretically censor or reorder transactions, though reputation incentives may mitigate this [10], [33].

SX Bet's dedicated rollup provides performance isolation and predictable costs [33], trading off interoperability and composability [7].

3.2.3. Overtime Markets (Thales Protocol)

Overtime Markets, governed by the Thales Protocol DAO on Optimism and Arbitrum [34], employs a parimutuel-style AMM where participants contribute to a shared prize pool distributed proportionally among winning positions after fees. This mechanism differs fundamentally from continuous double-auction and fixed-odds AMM models.

Under Overtime's parimutuel AMM [34], odds are determined by capital distribution across outcomes at market close. This shares LMSR's information aggregation properties [4] but differs in loss structure: the platform assumes no market maker position and derives revenue from a 2-5% spread on the pool.

The Thales Protocol, under which Overtime Markets operates, provides a governance and economic coordination layer that manages protocol parameters, fee distributions, and cross-chain deployment decisions [34]. The THALES token serves as the governance asset, with token holders voting on protocol upgrades, market parameters, and treasury allocations through the Thales DAO. This governance architecture follows the decentralized autonomous organization (DAO) model that has become standard in DeFi protocol design, though the concentration of governance power in large token holders—a pattern documented across DeFi DAOs by Schär [10]—applies equally to the Thales ecosystem.

Oracle integration represents a critical architectural component for Overtime Markets, as the platform requires reliable, tamper-resistant feeds of sporting event outcomes to settle markets. Overtime relies on Chainlink's decentralized oracle network for event result data, leveraging Chainlink's network of independent node operators who source outcomes from multiple data providers and reach consensus on the reported result [15]. The use of Chainlink oracles for sports data presents distinct challenges compared to the financial price feeds for which Chainlink was originally designed: sporting outcomes are typically reported by a small number of authoritative sources (league websites, official scorekeepers) rather than being continuously observable from multiple independent origins, creating a concentration risk in data sourcing that the oracle's node decentralization alone cannot fully mitigate. Kaleem and Shi [35] systematically surveyed Chainlink oracle deployments on Ethereum, identifying this "ground truth concentration" as a structural vulnerability in oracle-dependent applications where the underlying data source is inherently centralized—a concern directly applicable to sports outcome oracles.

The multi-chain deployment of Overtime Markets across Optimism and Arbitrum follows the pattern of Layer 2 diversification identified in the broader DeFi ecosystem [10]. By maintaining parallel deployments on multiple rollups, the protocol hedges against the operational risk of any single Layer 2 experiencing downtime, congestion, or security incidents. However, this deployment strategy also fragments liquidity: a market deployed on Optimism does not share a liquidity pool with the same market deployed on Arbitrum, creating parallel markets for identical events that may exhibit divergent pricing. This liquidity fragmentation across chains is a general challenge for multi-chain DeFi applications and has been identified as a limiting factor in prediction market efficiency by Lim [25], who demonstrated that fragmented liquidity in multi-venue AMMs reduces capital efficiency and increases the susceptibility of prices to manipulation in low-volume markets.

Overtime's parimutuel mechanism creates an informational advantage for late-arriving traders who observe accumulated capital before committing [34], disadvantaging early participants [17].

3.2.4. Kalshi

Kalshi is the first prediction market under formal U.S. federal regulatory oversight, licensed by the CFTC as a designated contract market [36], [37]. Unlike blockchain-based platforms, Kalshi operates a centralized limit order book without smart contracts or cryptographic tokens, serving as a regulatory benchmark that demonstrates both the possibilities and constraints of regulatory compliance.

Kalshi operates a fully centralized CLOB with custodial funds and proprietary order matching [36], [37], eliminating the oracle problem through regulatory obligations while concentrating trust in a single operator with associated counterparty risk [10], [36].

CFTC regulation restricts Kalshi's markets to macroeconomic indicators, policy outcomes, and cultural events [36], narrower than permissionless platforms but enabling credible settlement for economically significant contracts requiring authoritative data sources.

Bürge, Deng, and Whelan [37] provide the first systematic economic analysis of Kalshi's market microstructure, examining the behavior of "makers and takers" within the platform's order book. Their analysis reveals several patterns that contrast with the blockchain-based platforms surveyed above. First, Kalshi's market makers—the participants who submit resting limit orders—earn positive average spreads, consistent with compensation for providing liquidity and bearing adverse selection risk. This pattern mirrors the professional market maker model on Polymarket [21] but differs from AMM-based platforms where liquidity provision returns are determined algorithmically rather than competitively. Second, Kalshi's transaction costs, which include both explicit exchange fees and the effective spread, are substantially lower than the implicit costs (slippage, gas fees, bridge fees) borne by traders on decentralized platforms [37]. Third, the regulatory requirement for Kalshi to maintain surveillance for market manipulation—including wash trading, spoofing, and insider trading—provides a layer of market integrity protection that blockchain-based platforms address through transparency of on-chain data rather than through active surveillance [37].

The comparison between Kalshi and Polymarket is particularly instructive, as both platforms employ central limit order books but differ fundamentally in their regulatory posture and settlement architecture. Where Polymarket achieves non-custodial settlement through the Gnosis Conditional Token Framework and event resolution through UMA's decentralized oracle [11], [12], [13], Kalshi achieves settlement finality through CFTC oversight and regulatory compliance [36]. This trade-off—between cryptographic trust minimization and institutional trust in regulatory enforcement—represents a fundamental design choice in prediction market architecture that extends beyond the technical considerations of blockchain engineering to the legal and political dimensions of market design. As Wolfers and Zitzewitz [2] noted in their foundational survey, the regulatory environment has historically been the binding constraint on prediction market adoption, and Kalshi's licensed operation demonstrates one path through this constraint at the cost of the open participation and permissionless market creation that blockchain platforms enable.

The emergence of a CFTC-regulated prediction market also has implications for the long-term evolution of the prediction market ecosystem. Diercks, Katz, and Wright [36] suggest that the coexistence of regulated and decentralized prediction markets may lead to market segmentation: regulated platforms may dominate in markets requiring authoritative settlement data and institutional participation, while decentralized platforms may retain advantages in markets that benefit from permissionless creation, global participation, and resistance to single-jurisdiction regulatory pressure. The existence of a regulated benchmark also provides a natural experiment for comparing the forecasting accuracy and market efficiency of centralized and decentralized prediction market architectures—a comparison that the empirical literature on blockchain prediction markets has not yet systematically undertaken [7], [36].

3.2.5. Cross-Platform Comparison: Group B

The four platforms examined in Group B extend the architectural taxonomy established by Group A along two orthogonal dimensions: the domain specialization axis (sports versus general event markets) and the centralization axis (fully decentralized versus CFTC-regulated) (see Fig. 3).

The architectural comparison of the second group of platforms reveals distinct design choices for sports-focused and regulated markets, shown in Table 4. Table 4 presents an architectural comparison of the four Group B platforms: Azuro, SX Bet, Overtime Markets, and Kalshi.




Fig. 3. FMarket microstructure taxonomy: three paradigms in blockchain prediction markets — Central Limit Order Book (CLOB), Automated Market Maker (AMM), and Parimutuel

Table 4. Architectural Comparison of Azuro, SX Bet, Overtime Markets, and Kalshi Across Key Design Dimensions

Attribute	Azuro	SX Bet	Overtime Markets	Kalshi
Architecture Type	Middleware Liquidity Tree	CLOB on L3 rollup	Parimutuel AMM	Centralized CLOB
Blockchain	Polygon, Gnosis Chain	SX Network (Arbitrum Orbit)	Optimism, Arbitrum	Centralized (custodial)
Oracle	Azuro oracle framework	SX oracle + Chainlink	Chainlink	CFTC regulation
AMM/CLOB	Hierarchical pooled AMM	CLOB (order book)	Parimutuel AMM	Centralized order book
Fee Model	Vigorish (built-in margin)	Transaction fees	2-5% pool spread	Maker/taker fees
Token	AZUR (governance)	SX (gas + fee)	THALES (governance)	None
Liquidity	Pooled liquidity providers	Order book participants	Prize pool participants	Centralized market makers

The diversity of liquidity models in Group B demonstrates that no single microstructure is optimal across all market types and regulatory contexts (see Fig. 3).

Group B platforms demonstrate convergence toward sports specialization, reflecting the structural characteristics of sports outcomes—discrete, verifiable, and on known schedules—that suit automated market making and oracle-based settlement [2]. The contrast between blockchain-native platforms and Kalshi illustrates the fundamental trade-off between cryptographic trust minimization and regulatory compliance [10], [36]. Each blockchain platform accepts some trust assumption in exchange for open participation and jurisdictional resilience, while Kalshi concentrates trust in a regulated operator for legal certainty and market surveillance [7], [37].

The diversity of liquidity models—hierarchical AMM (Azuro), order book (SX Bet), parimutuel AMM (Overtime), centralized order book (Kalshi)—demonstrates that no single architecture dominates, unlike the convergence seen in traditional exchanges [17], [25]. This persistence suggests the optimal liquidity model depends on event type, trader population, and regulatory context [5], [6].

$$x \cdot y = k \quad (6)$$

$$\sum_{i=1}^n p_i = 1 \quad (7)$$

3.3. Discussion: Synthesis, Challenges, and Future Directions

The survey reveals a landscape of architectural diversity, unresolved technical challenges, and an evolving regulatory environment. This section synthesizes findings, identifies research gaps, and proposes future directions.

3.3.1. Cross-Cutting Architectural Themes

The decentralization-efficiency trade-off documented in the survey (see Fig. 1) recurs at every architectural layer, with the optimal design contingent on market characteristics [5].

Oracle design is central to platform security across all platforms, with architectures ranging from UMA's voting [13] to Augur's REP staking [23] to Chainlink's node network [15], [35] to Kleros [26], [27]. No single oracle model dominates, and oracle security remains the binding constraint on trustlessness guarantees [14], [16]. Kalshi [36] eliminates the oracle problem through regulatory liability, at the cost of centralized counterparty risk.

The Gnosis CTF [12] exemplifies a modular approach, decoupling settlement from trading infrastructure. This pattern, adopted by Polymarket [11], Omen, and others, contrasts with Augur's vertically integrated design [23]. Azuro's Liquidity Tree [30], [31] extends modularity to the liquidity layer, suggesting evolution toward a layered stack where settlement, liquidity, oracle, and interface functions are served by distinct protocols [28].

3.3.2. Open Technical Challenges

Several fundamental challenges remain unresolved. The oracle problem persists: cryptoeconomic oracle security is bounded by the total value at stake in native tokens, creating vulnerability when market settlement value exceeds attack cost [14], [16]. Relatedly, "ground truth concentration" arises when event outcomes depend on few authoritative sources, where node decentralization alone is insufficient [35]. Chainlink 2.0 [38] proposes super-linear staking and explicit decentralization guarantees, though these mechanisms lack empirical validation at prediction market scale.

Liquidity fragmentation operates at two levels: within platforms (capital distributed across CLOB, AMM, and parimutuel pools) and across chains (identical events on different networks). Lim [25] found that fragmented multi-venue AMM liquidity reduces capital efficiency and increases manipulation susceptibility. Neither dedicated rollups [33] nor multi-rollup deployment [34] provides a general solution.

Dispute resolution scalability is a third challenge. Multi-round dispute processes (Augur [23], Kleros [26]) provide robust guarantees but introduce latency scaling with stake and participation. Gabuthy [39] identified conditions for juror collusion when incentives exceed honesty penalties, while Zhuk [40] examined Kleros's interaction with legal frameworks. Future designs may need hierarchical

adjudication with lightweight resolution for routine disputes and robust processes for high-stakes cases.

Information leakage is particularly acute for prediction markets: order submission reveals private beliefs that become publicly observable through market prices, potentially disincentivizing informed participation [17]. Public blockchain ledgers amplify this tension by making order flow permanently observable and enabling front-running. Privacy-preserving techniques (ZK-proofs, MPC) remain unadapted to prediction market settlement requirements.

3.3.3. Regulatory Landscape

The regulatory environment diverges significantly across jurisdictions. The U.S. has moved furthest with CFTC licensing of Kalshi as a designated contract market [36], imposing structural constraints but providing legal certainty for institutional participation. Regulatory oversight constrains market breadth while potentially enabling depth in economically significant domains.

The EU's MiCA regulation (effective 2024-2025) applies to crypto-asset service providers but does not specifically address prediction markets. Outcome token classification—as financial instruments, crypto-assets, or gambling products—remains uncertain across member states, creating compliance challenges for pan-European platforms.

Regulatory divergence creates conditions for jurisdiction arbitrage: Polymarket restricted U.S. access following a 2022 CFTC settlement while serving international users [8]. This raises fundamental questions about enforceability in permissionless environments where users access smart contracts through non-custodial wallets without geographic identification [2].

Prediction markets intersecting with sports betting face gambling regulation classification risk, varying across legal systems. Sports-focused platforms (Azuro [30], [31], SX Bet [32], Overtime [34]) operate in a regulatory gray zone, with enforceability of smart contract resolution in conventional courts remaining unresolved [40].

3.3.4. Research Gaps

Several significant gaps exist in the literature. First, no study compares forecasting accuracy across different platform architectures (CLOB vs. AMM vs. parimutuel) using common events and metrics [2], [6], [19], [20], [37]. Second, tokenomics sustainability analysis is underdeveloped: while platforms employ native tokens for governance, fees, and dispute resolution [23], [30], [34], rigorous analysis of whether token value derives from platform usage or speculation is lacking [10], [17], [25]. Third, trader behavior beyond Polymarket is unexamined; comparable analyses for AMM-based platforms (Augur, Azuro, Overtime) do not exist [21], [22]. Fourth, DeFi composability of prediction market positions—outcome tokens as lending collateral, structured products, DEX trading—remains unexplored [10], [12].

3.3.5. Future Work Directions

Priority research directions include: (1) empirical cross-platform accuracy benchmarking using common events across political, economic, and sports domains to compare forecasting performance across architectures [19], [37];

(2) formal mechanism design for prediction market-specific AMMs, starting from the information structure of binary/categorical payoffs with known resolution time and heterogeneous private information, to derive optimal pricing and liquidity rules [4], [6], [17], [30], [34];

(3) AI agent convergence with prediction markets, where agents autonomously create, participate in, and act upon market signals [29], raising questions about mechanism design for algorithmic participants, manipulation dynamics, and information cascade implications;

(4) cross-chain liquidity aggregation for time-sensitive prediction market requirements; (5) empirical analysis of regulatory effects on platform participation and market efficiency, leveraging the natural experiment of Kalshi's coexistence with decentralized platforms [11], [19], [36], [37].

3.3.6. Limitations

Limitations include: architectures may have evolved since our mid-2026 snapshot; reliance on public documentation may miss proprietary features; qualitative methodology precludes quantitative performance comparisons; and regulatory analysis focuses on U.S. and EU frameworks [5], [16].

4. Conclusion

This survey of seven blockchain-based prediction markets reveals fundamental architectural trade-offs between decentralization, market efficiency, and regulatory compliance. Three cross-cutting findings, detailed in the Discussion, emerge: (i) oracle design constitutes the critical security bottleneck across all architectures [14], [16]; (ii) the diversity of liquidity models confirms that no single microstructure is universally optimal [17], [5]; and (iii) the Gnosis CTF as a shared settlement standard enables modularity and composability analogous to broader DeFi architecture [12], [10].

The contrast between blockchain-native and regulated architectures reveals a fundamental tension: cryptographic trust minimization enables permissionless participation, while regulatory compliance enables authoritative settlement [36]. Future research should include empirical accuracy comparison between architectures, improved oracle mechanisms, and integration of prediction market signals with autonomous decision-making systems through the InfoFi paradigm [29].

Funding: This research received no external funding.

Author Contribution: The author confirms sole responsibility for the following: study conception and design, literature search and platform selection, data extraction and analysis, manuscript drafting and revision, and approval of the final manuscript.

Acknowledgment: The author acknowledges the platforms and projects surveyed in this paper for maintaining publicly accessible technical documentation and open-source code repositories, without which this structured survey would not have been possible.

Conflicts of Interest: The author declares no conflict of interest.

References

- [1] F. A. Hayek, "The Use of Knowledge in Society," *American Economic Review*, vol. 35, no. 4, pp. 519–530, 1945, <https://www.aeaweb.org/aer/top20/35.4.519-530.pdf>
- [2] J. Wolfers and E. Zitzewitz, "Prediction Markets," *Journal of Economic Perspectives*, vol. 18, no. 2, pp. 107–126, 2004, <https://doi.org/10.1257/0895330041371321>.
- [3] K. J. Arrow *et al.*, "The Promise of Prediction Markets," *Science*, vol. 320, no. 5878, pp. 877–878, 2008, <https://doi.org/10.1126/science.1157679>.
- [4] R. Hanson, "Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation," *The Journal of Prediction Markets*, vol. 1, no. 1, pp. 1–15, 2007, <https://doi.org/10.5750/jpm.v1i1.417>.
- [5] F. M. A. Klingert and M. Meyer, "Comparing Prediction Market Mechanisms: An Experiment-Based and Micro Validated Multi-Agent Simulation," *Journal of Artificial Societies and Social Simulation*, vol. 21, no. 1, 2018, <https://doi.org/10.18564/jasss.3577>.
- [6] A. V. C. Carvalho, D. Silveira, R. A. Ely, and D. O. Cajueiro, "A Logarithmic Market Scoring Rule Agent-Based Model to Evaluate Prediction Markets," *Journal of Evolutionary Economics*, vol. 33, no. 4, pp. 1303–1343, 2023, <https://doi.org/10.1007/s00191-023-00822-w>.
- [7] P. Buckley, "Blockchain Based Prediction Markets," *The Journal of Prediction Markets*, vol. 18, no. 3, pp. 73–86, 2025, <https://doi.org/10.5750/jpm.v18i3.2182>.
- [8] K. Mattmuller, "Decentralized Prediction Markets," *Georgetown Law Technology Review*, vol. 8, no. 2, 2024, <https://georgetownlawtechreview.org/decentralized-prediction-markets/GLTR-05-2024/>.

-
- [9] Grayscale Research, "Polymarket: Crypto's Election-Year Breakout App," *Grayscale Research Report*, 2024, <https://research.grayscale.com/reports/polymarket>.
 - [10] F. Schär, "Decentralized Finance: On Blockchain-and Smart Contract-Based Financial Markets," *Federal Reserve Bank of St. Louis Review*, vol. 103, no. 2, pp. 153–174, 2021, <https://doi.org/10.20955/r.103.153-74>.
 - [11] Polymarket, "Polymarket Technical Documentation," *Polymarket*, 2024, <https://docs.polymarket.com/>.
 - [12] Gnosis, "Conditional Tokens Framework: Technical Documentation," *Gnosis*, 2020, <https://github.com/gnosis/conditional-tokens-contracts>.
 - [13] UMA Protocol, "UMA Data Verification Mechanism (DVM) Oracle Whitepaper," *UMA Protocol*, 2019, <https://github.com/UMAprotocol/whitepaper>.
 - [14] G. Caldarelli and J. Ellul, "The Blockchain Oracle Problem in Decentralized Finance—A Multivocal Approach," *Applied Sciences*, vol. 11, no. 16, 2021, <https://doi.org/10.3390/app11167572>.
 - [15] S. Ellis, A. Juels, and S. Nazarov, "ChainLink: A Decentralized Oracle Network," *Chainlink Whitepaper*, 2017, <https://research.chain.link/whitepaper-v1.pdf>.
 - [16] H. Al-Breiki, M. H. U. Rehman, K. Salah, and D. Svetinovic, "Trustworthy Blockchain Oracles: Review, Comparison, and Open Research Challenges," *IEEE Access*, vol. 8, pp. 85675–85685, 2020, <https://doi.org/10.1109/ACCESS.2020.2992698>.
 - [17] Y. Wang, "Prediction Markets, Automated Market Makers, and Decentralized Finance (DeFi)," in *Mathematical Research for Blockchain Economy*, pp. 213–231, 2023, https://doi.org/10.1007/978-3-031-18679-0_12.
 - [18] G. Angeris and T. Chitra, "Improved Price Oracles: Constant Function Market Makers," in *Proc. 2nd ACM Conf. Advances in Financial Technologies (AFT '20)*, 2020, <https://doi.org/10.1145/3419614.3423251>.
 - [19] F. Reichenbach and M. Walther. *Exploring Decentralized Prediction Markets: Accuracy, Skill, and Bias on Polymarket*. SSRN Working Paper, 2025, <https://doi.org/10.2139/ssrn.5910522>.
 - [20] F. Cordoba Otalora and M. Themistocleous, "Beyond the Polls: Quantifying Early Signals in Decentralized Prediction Markets with Cross-Correlation and Dynamic Time Warping," *Future Internet*, vol. 17, no. 11, 2025, <https://doi.org/10.3390/fi17110487>.
 - [21] P. Akey, V. Grégoire, N. Harvie, and C. Martineau. *Who Wins and Who Loses In Prediction Markets? Evidence from Polymarket*. SSRN Working Paper, 2026, <https://doi.org/10.2139/ssrn.6443103>.
 - [22] H. Chen, X. Duan, A. El Saddik, and W. Cai, "Political Leanings in Web3 Betting: Decoding the Interplay of Political and Profitable Motives," in *Proc. 17th ACM Web Science Conf. (WebSci '25)*, pp. 96–105, 2025, <https://doi.org/10.1145/3717867.3717884>.
 - [23] J. Peterson and J. Krug, "Augur: a Decentralized Oracle and Prediction Market Platform (v2.0)," *Augur Project Whitepaper*, 2014, <https://augur.net/whitepaper.pdf>.
 - [24] A. Carvalho, "A Permissioned Blockchain-Based Implementation of LMSR Prediction Markets," *Decision Support Systems*, vol. 130, 2020, <https://doi.org/10.1016/j.dss.2019.113228>.
 - [25] T. Lim, "Predictive Crypto-Asset Automated Market Maker Architecture for Decentralized Finance Using Deep Reinforcement Learning," *Financial Innovation*, vol. 10, 2024, <https://doi.org/10.1186/s40854-024-00660-0>.
 - [26] F. Ast, W. George, J. Kamalova, A. Sharma, and Y. Aouidef, "Decentralized Justice: State of the Art, Recurring Criticisms and Next-Generation Research Topics," *Frontiers in Blockchain*, vol. 6, 2023, <https://doi.org/10.3389/fbloc.2023.1204090>.
 - [27] Y. Aouidef, F. Ast, and B. Deffains, "Decentralized Justice: A Comparative Analysis of Blockchain Online Dispute Resolution Projects," *Frontiers in Blockchain*, vol. 4, 2021, <https://doi.org/10.3389/fbloc.2021.564551>.
 - [28] S. Wang, X. Ni, Y. Yuan, F.-Y. Wang, X. Wang, and L. Ouyang, "A Preliminary Research of Prediction Markets Based on Blockchain Powered Smart Contracts," in *2018 IEEE Int. Conf. Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and*
-

- Social Computing (CPSCoM) and IEEE Smart Data (SmartData)*, pp. 1255–1262, 2018, https://doi.org/10.1109/Cybermatics_2018.2018.00224.
- [29] Gnosis, "The Rise of AI Agents in Prediction Markets: How Gnosis Infrastructure Is Powering the Future of Information Finance," *Gnosis Blog*, 2025, <https://www.gnosis.io/blog/the-rise-of-ai-agents-in-prediction-markets-how-gnosis-infrastructure-is-powering-the-future-of-information-finance>.
- [30] Azuro Protocol, "Liquidity Tree: A Novel Liquidity Infrastructure for On-Chain Prediction Markets," *Azuro Docs*, 2022, <https://docs.azuro.org/>.
- [31] Messari Research, "Understanding Azuro: A Comprehensive Overview," *Messari Report*, 2025, <https://messari.io/report/understanding-azuro-a-comprehensive-overview>.
- [32] SX Bet, "SX Bet Documentation: Sports Prediction Market," *SX Bet*, 2024, <https://docs.sx.bet/>.
- [33] SX Network, "SX Rollup: Transitioning to an Arbitrum Orbit Chain," *SX Bet Blog*, 2024, <https://medium.com/sportx-bet/sx-rollup-ca33079d1b8b>.
- [34] Thales Protocol / Overtime Markets, "Thales Protocol Technical Documentation," *Thales*, 2024, <https://www.thales.io/>.
- [35] M. Kaleem and W. Shi, "Demystifying Pythia: A Survey of ChainLink Oracles Usage on Ethereum," in *Financial Cryptography and Data Security. FC 2021 Int. Workshops*, pp. 131–146, 2021, https://doi.org/10.1007/978-3-662-63958-0_10.
- [36] A. M. Diercks, J. D. Katz, and J. H. Wright. *Kalshi and the Rise of Macro Markets*. National Bureau of Economic Research, 2026, <https://doi.org/10.3386/w34702>.
- [37] C. Bürgi, W. Deng, and K. Whelan. *Makers and Takers: The Economics of the Kalshi Prediction Market*. SSRN Working Paper, 2026, <https://doi.org/10.2139/ssrn.5502658>.
- [38] L. Breidenbach *et al.*, "Chainlink 2.0: Next Steps in the Evolution of Decentralized Oracle Networks," *Chainlink Whitepaper*, vol. 1, pp. 1-136, 2021, <https://naorib.ir/white-paper/chinlink-whitepaper.pdf>.
- [39] Y. Gabuthy, "Blockchain-Based Dispute Resolution: Insights and Challenges," *Games*, vol. 14, no. 3, 2023, <https://doi.org/10.3390/g14030034>.
- [40] A. Zhuk, "Applying Blockchain to the Modern Legal System: Kleros as a Decentralised Dispute Resolution System," *International Cybersecurity Law Review*, vol. 4, pp. 217–237, 2023, <https://doi.org/10.1365/s43439-023-00086-x>.