Evidence Markets

Safwan Hossain[*] Harvard University
Gabriel Andrade Gensyn AI
Chengqi Zang Gensyn AI
Yiling Chen Harvard University

Abstract

Modern prediction markets face two limitations that restrict their applicability in a range of settings: (i) they reveal what the crowd believes but not the evidence or reasoning behind those beliefs, and (ii) they require an exogenous event with an external ground truth that resolves at a known future date. We address these twin challenges by introducing evidence markets, a generalization of prediction markets that incentivizes the submission of evidence alongside beliefs and can be resolved both exogenously and endogenously from crowd-sourced evidence when external resolution is unavailable. At its core, the market uses a logarithmic market scoring rule whose liquidity parameter changes dynamically with the accumulated evidence. We prove that platform loss is bounded, evidence is rewarded in proportion to current market uncertainty, and that the mechanism can be equivalently implemented through an automated market maker. When the market resolves endogenously, we characterize how withholding evidence shifts a trader's belief about resolution and use this to prove that truthful belief and evidence reporting always constitutes an ε-subgame perfect Nash Equilibrium. For exogenous resolution, truthful belief and evidence submission is a dominant strategy equilibrium. To address operational considerations, we propose evidence verification via an LLM-as-a-Judge framework with staking, and give an asynchronous execution algorithm that is not bottle-necked by verification. Throughout the work, we use LLM evaluations—determining which model is best for a given task—as a salient and representative running example for our proposed market.

1 Introduction

Prediction markets have experienced a meteoric rise over the past few years, moving from an academic curiosity to mainstream infrastructure for forecasting and decision-making. Platforms such as Polymarket and Kalshi now host markets on elections, geopolitics, scientific milestones, and economic indicators, drawing substantial volume and attention from policymakers, journalists, and researchers. Their appeal rests on a clean theoretical foundation: under standard assumptions, the market price aggregates participants' dispersed beliefs about an uncertain event into a single, continuously updating probability estimate—one that often outperforms expert forecasts and polls. Yet despite their growing prominence, prediction markets in their current form face two fundamental limitations that constrain both the depth and the breadth of what they can offer.

The first limitation is interpretive. Prices on a prediction market summarize the weighted aggregate of participants' beliefs into a single probability, but the evidence and reasoning underlying those beliefs remain privately held by each trader. In short, the current iteration of prediction markets tell us *what* the crowd believes, but not *why*. Understanding this reasoning is crucial to have confidence in the forecast and justify any downstream actions. Decision-makers acting just on price signals are forced to infer this reasoning themselves, eroding much of the efficiency that motivated such markets in the first place. In contrast, a mechanism that incentivizes traders to disclose the evidence behind their beliefs would let the market summarize precisely this information. Future traders could then update their beliefs in light of the evidence already submitted, yielding a more informed aggregate and a richer information product overall.

The second limitation is more fundamental and pertains to which events can even be considered for prediction markets. Currently, they can only consider events which are externally resolved by time — i.e.

[*]Correspondence to `shossain@g.harvard.edu`

Table 1: Summary of the guarantees given in this work for the proposed Evidence Market under truthful submission of belief and complete evidence set.

	Endogenous Resolution	**Exogenous Resolution**
Dominant Strategy Equilibrium	✗ Exact Equilibrium	✓ Exact Equilibrium
Sub-game Perfect Nash Equilibrium	✓ ε-Exact Equilibrium	✓ Exact Equilibrium

events whose outcome will, at some specific future date, be publicly observable and verifiable independent of the market itself. This is the natural setting for elections, sporting events, asset prices, and other phenomena that nature or institutions will resolve on a known timeline. It fails, however, for the much larger class of events whose answers depend on judgment, evidence, or interpretation rather than on the passage of time: questions about scientific replication, policy effectiveness, the quality of competing artifacts, and so on. By incentivizing evidence submission, we may be able to endogenously resolve such questions based on the submitted evidence.

The evaluation of large language models (LLMs) stands out as both a pressing and structurally illustrative example of these concerns. LLMs are released at a rapid pace that has outstripped the capacity of static benchmarks to keep up. Further, benchmarks are costly to produce, saturate quickly, leak into training data, and capture narrow slices of capability that may be relevant only to a few. It is possible to capture such context-specific evaluation by framing it as an event: *Which model is better at task X?* However, there is no external, time-bound, and clear resolution of such an event; indeed resolution is itself an exercise in building an evaluation dataset (evidence aggregation). This makes LLM evaluation a natural setting for a market mechanism that crowd-sources an evidence dataset and endogenously resolves through that evidence. Ideally, participants can not only bet on which model wins, but also submit supporting evidence in the form of evaluation questions which will be used to determine the winner. This forms the running example used throughout this work and a novel application facilitated by it.

Our Contribution: We propose *evidence markets* that addresses the two limitations above in tandem. This mechanism (1) incentivizes both belief reporting and evidence submission, and (2) can be resolved externally by realized events or endogenously through the submitted evidence. Traders are free to report either belief or evidence (or both), thereby generalizing existing prediction markets as a special case of our model. Section 3 outlines the formal model used in this paper, outlines our proposed market, and formalizes the trader interactions as an extensive-form game. We observe that in endogenous resolution, traders affecting market resolution through evidence submission means there is a multitude of consistent beliefs they can hold about resolution depending on what evidence they submit. Section 4 formalizes this epistemic question and bounds how selectively submitting private evidence can affect a trader's resolution belief. Section 5 then discusses the overall payoff mechanism for this market under both types of market resolution by leveraging the formalisms of logarithmic market scoring rules (LMSRs) and automated market makers (AMMs). At a high-level, we propose a dynamic liquidity parameter that *decreases* as the evidence quality in the market *increases* and show that truthful belief and full evidence submission is a ε-subgame perfect Nash Equilibrium (SPNE) for endogenous resolution and strict dominant strategy equilibrium (DSE) for exogenous resolution. The equilibrium results are outlined in Table 1. In Section 6 we comment on the evidence verification problem and highlight the unsuitability of standard approaches like peer-prediction for this problem. Building on past related works, we propose an LLM-as-a-Judge framework and augment it with staking where relevant to address this challenge. We also give a practical algorithm for executing orders that is not blocked by evidence verification. We conclude in Section 7 with a discussion on the broader implications of our work and open questions it raises.

2 Related Work

Prediction markets and market scoring rules: A long line of empirical work has established that prediction markets aggregate dispersed information effectively and often outperform alternative forecasting methods [24, 15, 1]. The theoretical foundation for our mechanism rests on the seminal work of Hanson which proposed logarithmic market scoring rule, an incentive-compatible mechanism for truthful belief submission [7, 8]. Building on this, Chen and Pennock [3] formalizes the connection between such market scoring rules being duals of automated market makers (AMMs) with a given cost function, allowing versatility in implementation. Our construction inherits the LMSR cost-function structure but modifies the liquidity parameter to be a decreasing function of cumulative evidence quality, coupling price-formation directly to the evidence record.

Markets for unverifiable outcomes and explanations: The self-resolving prediction market of Srinivasan et al. [21] also looks to address one of the limitations we point out: existing markets require an externally verifiable ground truth to resolve. Their mechanism terminates randomly with probability α after each report and pays earlier participants based on the final reporter's prediction via negative cross-entropy, achieving truthful reporting as a perfect Bayesian equilibrium without ever observing an outcome. Our work shares the motivation but takes a structurally different approach. Rather than treating the final belief report as a proxy for ground truth, we open up a second axis, submission of evidence that can discriminate between alternatives, and resolve endogenously using the crowd-sourced evidence. This shift converts the resolution signal from a belief aggregate into an evidence aggregate, and thereby grounds resolution away from self-referential beliefs and toward externally inspectable artifacts that can be of independent interest.

The other limitation mentioned – that markets reveal beliefs but not the reasoning behind them – is taken up by Srinivasan et al. [20], which proposes directly and only eliciting rationales for events as opposed to beliefs. Their model captures the insight that traders' beliefs are drawn from overlapping information sources, so belief reports alone fail to reveal what is shared versus what is new; rationales uncover this distinction and enable more efficient aggregation. Our mechanism pursues the same goal – uncovering the *why* behind the price – but through a different channel: rather than only eliciting verbal rationales, we make evidence/explanations a market commodity that sits alongside beliefs. Importantly, this allows traders who may not be able to articulate their rationales well or have no novel information, to still trade in the market through just beliefs as in standard prediction markets. Indeed, a trader's information/explanation may already exist in the market, but they can draw different conclusions from it and thus arrive at a distinct belief. Our work allows this flexibility not found in Srinivasan et al. [20].

On the whole, these two papers independently address the resolution problem and the explanation problem that motivated this work. To our knowledge, no prior work addresses both simultaneously. This is our core contribution: a single mechanism that can reveal the motivating basis of the price as well as endogenously resolve the market through such evidence. Crucially, our design is also *flexible* in a way that neither prior mechanism is; traders may submit beliefs without evidence (recovering standard prediction markets trading), submit evidence without taking a position (acting as pure contributors to the resolution record), or do both. This flexibility lowers the participation barrier for traders whose comparative advantage lies on only one side or have different risk profiles.

Peer prediction. A separate literature elicits truthful reports without ground truth by paying agents based on the reports of peers [16, 17, 5, 12, 18]. Multi-task extensions including Correlated Agreement [19] and Determinant-based Mutual Information [11] strengthen these results under weaker assumptions. We engage directly with this literature as a possible way to verify evidence in the absence of ground truth. Unfortunately, when agents can participate as both traders and peer verifiers of evidence, peer predictions guarantees do not map cleanly. This motivates our optimistic verification with staked disputes, in which a frontier LLM serves as default verifier and agents dispute by staking, with adjudication having access to the whole history. This draws on the now well-established asymmetry between generation and verification of reasoning: process reward models [14], generative reward models [25], and verifier-augmented thinking models [10] all exploit the fact that checking a candidate solution is structurally easier than producing one. Lastly, the recent WOMAC mechanism [22] also explores adjudication structures for prediction competitions and is conceptually adjacent, though aimed at a different (competition-style) setting.

LLM evaluation. The most pressing application for our framework is the comparative evaluation of LLMs, where static benchmarks saturate quickly, suffer from contamination, and capture narrow capability slices [13, 23]. Crowd-sourced pairwise-preference platforms such as Chatbot Arena [4] and judged benchmarks such as MT-Bench [26] partially address these issues but lack incentive structure for participants and do not produce a market-style continuously-updated estimate. Our mechanism is designed for precisely this regime: the evaluation question/event can be highly targeted and the crowd-sourced evidence set can be used for resolution, alongside more detailed model evaluation and other downstream tasks.

3 Model

Agent Evidence and Beliefs: Consider an event with n possible alternatives or outcomes. Participants, agents, or *traders* arrive to the market sequentially, and each trader t's belief about the resolution of this event is denoted by a probability vector $q_t \in \mathrm{int}(\Delta(n-1))$, the interior of the $n-1$ dimensional simplex[1]. Each trader also holds some evidence set $E_t = (e_{t1}, \ldots, e_{t\ell})$ where each e_{tj} can be seen as some atomic unit of evidence/explanations and $|E_t| \leq h_{max}$[2]. It is perfectly valid for $E_t = \emptyset$, implying the traders holds no evidence, generalizing existing predictions markets. In the specific case of LLM evaluation, this atomic evidence can be seen as an evaluation question and answer pair: $e_{tj} = (\text{question}_{tj}, \text{answer}_{tj})$. In general, we consider traders as being able to selectively choose or subsample the evidence they submit to the market, denoted E'_t. Note that this does not typically include them manipulating or tampering the atomic units e_{tj}. Section 6 broadens this perspective and considers the evidence verification problem for (1) exogenous resolution and (2) endogenous resolution for LLM evaluations. Overall, for endogenously resolving based on evidence, we make this structural assumption connecting evidence to outcomes:

Definition 3.1. *Each alternative or outcome $i \in [n]$ has a binary outcome X_{ie} with respect to an atomic evidence e. $X_{ie} = 1$ denotes the evidence supporting or corroborating event alternative i, and $X_{ie} = 0$ denotes the evidence refuting this alternative.*

While the definition above is simplistic, it allows theoretical tractability and is especially pertinent in the LLM evaluation setting where e is a (question, answer) pair: X_{ie} simply maps to whether model i answered that question correctly or not. We use E_t^{total} to denote the total amount of evidence in the market at some time t, with the cumulative evidence submitted so far being publicly revealed. Key to our model, is the notion of *evidence quality* which is applied to arbitrary sets of evidence. More formally:

Definition 3.2. *For evidence $E = (e_1, \ldots, e_\ell)$, let $r(E)$ denote the quality of that evidence set. Further, let $R_t = r(E_t^{total})$ denote the quality of the cumulative evidence set that exists in the market at time t. We assume that $r(\cdot)$ is non-negative and monotone increasing: $r(E) \geq 0$ and for $E_1 \subseteq E_2$, $r(E_1) \leq r(E_2)$*

The quality function r is clearly contextual and our results do not make any assumptions beyond the stated non-negativity and monotonicity, which are fairly natural for such settings. The function r also encapsulates any verification or filtering that is pertinent for a given market and could depend on E_t^{total}: we may, for instance, only want relevant questions that are not semantically similar to evidence that is already in the market. We comment more on the specific instantiations of r and it's nuances as it relates to both endogenous and exogenous resolution in Section 6.

Market Desiderata: Our goal is to propose a market for the submission of belief and evidence which can aggregate these diverse positions and payout traders such that truthful behaviour is incentivized. More specifically, we wish to build a market that satisfies the following desiderata:

[1] In practice, it is standard to assume trader belief lie in the simplex interior; beliefs on the simplex edge can have unbounded worst-case loss for traders in LMSR style mechanisms and thus becomes impossible for platform to collect.

[2] We consider traders as innately possessing their evidence pair. While another natural model could consider some effort/cost in evidence acquisition, this is not our setting.

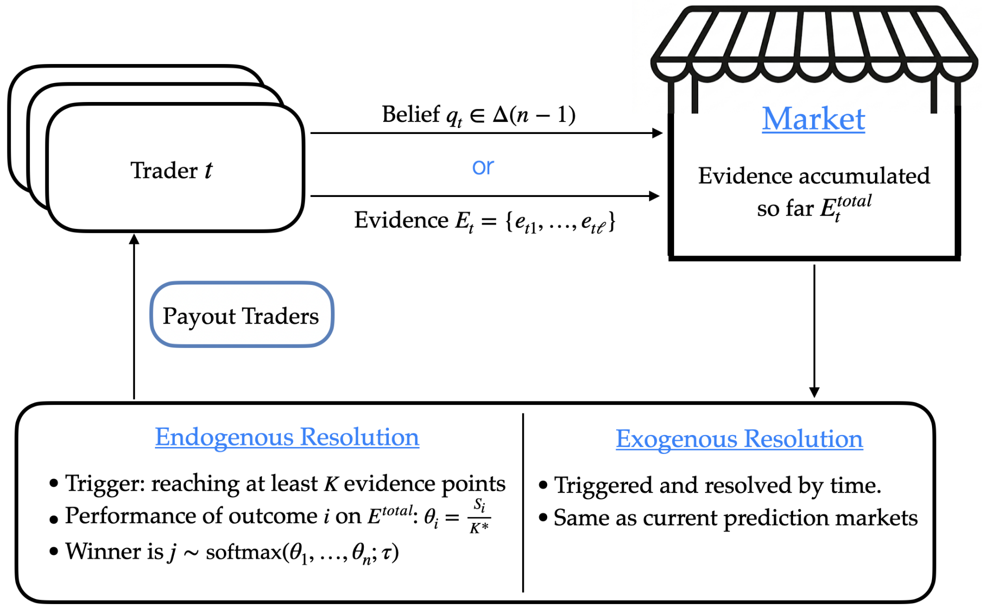

Figure 1: High-Level flow of the proposed evidence market. Section 4 studies the connection between belief and evidence under endogenous resolution in detail, and Section 5 outlines our specific payout mechanism that build upon and generalizes LMSRs.

Definition 3.3. *Evidence markets should allow traders to submit beliefs and/or evidence and satisfy:*

- *Axiom 1: Can be resolved exogenously or endogenously through the collected evidence.*
- *Axiom 2: It is optimal for each trader t to submit their full evidence set E_t.*
- *Axiom 3: It is optimal for each trader t to report/act with respect to their true belief q_t.*
- *Axiom 4: The numerical market belief at any time should be clearly interpretable.*
- *Axiom 5: The order in which an agent submits her evidence/beliefs/trades should be irrelevant.*

Market Structure: We now propose our high-level market structure, visualized in Figure 1. It consists of a finite number of traders arriving sequentially and submitting beliefs and evidence to the market. Beliefs can be submitted directly or indirectly through buying shares of resolution outcomes (see Section 5.2). If the market is externally resolved, the resolution is triggered at the associated time T, and traders receive payoff. In the case of endogenous resolution, resolution is triggered at the first time instance when at least K pieces of evidence has accumulated. If after some time T_{end}, the accumulated evidence has not reached K, the market times out and no payoff is achieved. Observe that in any resolved market, the accumulated evidence $K^* := |E_{total}| \in [K, K + h_{max} - 1]$. The choice of K is an event-creator chosen hyperparameter, and we assume that it is smaller than the cumulative amount of private evidence held by all possible traders for an event between $t \in [0, T_{end}]$. Endogenous resolution then follows as such:

Definition 3.4. *Endogenous resolution is triggered at the first time instance when at least K pieces of evidence has accumulated. For each resolution alternative i and accumulated evidence size at resolution $K^* = |E_{total}|$, let $\theta_i = \frac{1}{K^*}\sum_e X_{ie}$ denote the* score *or the fraction of E^{total} supporting i. Then the market resolves by sampling alternative $i \sim softmax(\theta_1, \ldots, \theta_n; \tau)$ where τ is the temperature of the softmax. If timeout T_{end} occurs prior to at least K pieces of evidence accumulating, the market terminates with no payoff or loss to any party.*

This market interaction can also be interpreted as an extensive form game between the traders. Formally, an instance of this game is characterized by:

- **Sequential Arrivals:** Each trader arrives to market once and at some fixed and distinct time t[3].
- **Strategy Space:** Each trader's action is to submit a belief q_t and some evidence set E'_t conditioned on the history of play $h_t = (q_1, E'_1, \ldots, q_{t-1}, E'_{t-1})$[4]. We denote the strategy function $\gamma_t : h_t \to (q_t, E'_t)$ mapping histories to actions. Traders do not have any knowledge or priors about future evidence to be submitted.
- **Payoffs:** Each trader's utility/payoff is a function of the strategies $(\gamma_1, \ldots,)$, the complete history h_T, as well as other market hyperparameters: $u(\gamma_1, \ldots, \gamma_n, \cdot)$.

We specify the exact payoff structure of this market/game in Section 5. This extensive form game allows us to consider two possible equilibrium notions: *dominant strategy equilibrium* and *subgame-perfect Nash Equilibrium*. They are formally defined in this setting as follows:

Definition 3.5 (Dominant Strategy Equilibrium)**.** *A complete strategy profile $(\gamma_1, \ldots, \gamma_n)$ is a dominant strategy equilibrium (DSE) if for each agent i, strategy function γ_i maximizes i's utility regardless of any other strategies taken by past or future agents: $\forall i, \forall \gamma'_i, \boldsymbol{\gamma}'_{-i} : u_i(\gamma_i, \boldsymbol{\gamma}'_{-i}) \geq u_i(\gamma'_i, \boldsymbol{\gamma}'_{-i})$.*

Definition 3.6 (Subgame Perfect Nash Equilibrium)**.** *A complete action profile $(\gamma_1, \ldots, \gamma_n)$ is a subgame perfect Nash Equilbrium (SPNE) if for any sub-game rooted at (t, h_t), the continuation profile $(\gamma_t, \gamma_{t+1}, \ldots, \gamma_T)$ is a Nash Equilibrium: $\forall \tau \geq t : u_\tau(\gamma_\tau, \boldsymbol{\gamma}_{-\tau} | h_t) \geq u_\tau(\gamma'_\tau, \boldsymbol{\gamma}_{-\tau} | h_t)$*

Since agent actions can be adaptive to history, dominant strategy equilibrium is a very strong desiderata. An agent's strategy needs to be optimal even if future agents are adversarial to them, herd in specific ways, or take up random strategies. SPNE, on the other hand, is weaker and a refinement of the standard Nash Equilibrium notion from simultaneous games. It requires that for a given action profile, no agent can benefit by unilaterally deviating from the given profile.

The remainder of the paper is devoted to precisely specifying each part of this market/game, characterizing the equilibria, and validating that they satisfy the desiderata outlined in Definition 3.3. We first consider the epistemic question that arises in endogenous resolution: how does selective evidence submission affect a trader's belief about resolution and what is the sensitivity of this relationship. Later, we outline the market's payoff mechanism: a modified version of an LMSR that can be equivalently implemented through an AMM.

4 Beliefs under Endogenous Resolution

In traditional prediction markets, trader actions cannot affect outcome resolution. As such, the belief q they hold about resolution is an exogenous quantity and coincides with their belief about the event as a whole. This structurally changes in endogenous resolution where evidence submitted by the traders is used to resolve the market. As such, a distinction emerges between a trader's belief about the underlying event, and their belief about market resolution, which in our market is sampled from the softmax over the terminal scores of each alternative. It is a trader's belief about this latter that we denote by q. To be precise about this, observe that we can decompose the score of alternative j from the perspective of trader t as a function of the current evidence in the market, the evidence submitted by trader t, and future evidence. Specifically, let $k(t) = |E^t_{total}|$ denote the total pieces of evidence collected at time t, with $k_j(t) = \sum^{k(t)} X_{je}$ counting how much of the current evidence supports alternative j. For trader t holding evidence set E_t, let $E'_t \subseteq E_t$ of which $E'_{t,j}$ support outcome j. Lastly, let $S_{K_{T-t}} = K^* - k(t) - |E'_t|$ denote the amount of future evidence left to arrive, with the random variable $S_{K_{T-t},j}$ denoting the number of future evidence that supports alternative j. Then for any E'_t this trader may submit:

$$\text{Trader } t\text{'s view of alternative } j\text{'s score: } \theta_{tj}(E'_t) = \frac{k_j(t) + |E'_{t,j}| + S_{K_{T-t},j}}{k(t) + |E'_t| + S_{K_{T-t}}} = \frac{k_j(t) + |E'_{t,j}| + S_{K_{T-t},j}}{K^*}.$$

[3]Due to the bijectivity of arrival time and agent identity, we index agents and their strategy by time wherever convenient.
[4]In Section 5.2 traders purchase shares as opposed to submitting beliefs.

Clearly, trader t's perception of the score for alternative j is a function of the evidence they choose to submit, and correspondingly, so is their belief about resolution[5]. We now formalize the relationship below, with $\boldsymbol{S_{K_{T-t,*}}} = (S_{K_{T-t,1}}, \ldots, S_{K_{T-t,n}})$ denoting the random variable tuple of future evidence:

Lemma 4.1. *A trader arriving at time t has the following belief about the market resolving to outcome $j \in [n]$ when they submit an evidence set $E'_t \subseteq E_t$:*

$$q_{tj}(E'_t) = \mathbb{E}_{\boldsymbol{S_{K_{T-t,*}}}} \left[\frac{\exp(\frac{1}{\tau}\theta_{tj}(E'_t))}{\sum_\ell \exp(\frac{1}{\tau}\theta_{t\ell}(E'_t))} \right] \tag{1}$$

Proof. Due to the randomness of future evidence, θ_{tj} is a random variable. As per our resolution procedure (see Definition 3.4), we take the softmax based on these scores. So the interim belief of trader t regarding resolution can be expressed as a function of the evidence they submit:

$$\begin{aligned} q_{tj}(E'_t) &= P[j \sim \text{softmax}(\theta_{t1}(E'_t), \ldots, \theta_{tn}(E'_t); \tau)] \\ &= \sum_{\boldsymbol{S_{K_{T-t,*}}}} P(\boldsymbol{S_{K_{T-t,*}}}) \, \text{softmax}(\theta_{t1}(E'_t), \ldots, \theta_{tn}(E'_t); \tau | \boldsymbol{S_{K_{T-t,*}}}) \\ &= \sum_{\boldsymbol{S_{K_{T-t,*}}}} P(\boldsymbol{S_{K_{T-t,*}}}) \frac{\exp(\frac{1}{\tau}\theta_{tj}(E'_t))}{\sum_\ell \exp(\frac{1}{\tau}\theta_{t\ell}(E'_t))} = \mathbb{E}_{\boldsymbol{S_{K_{T-t,*}}}} \left[\frac{\exp(\frac{1}{\tau}\theta_{tj}(E'_t))}{\sum_\ell \exp(\frac{1}{\tau}\theta_{t,\ell}(E'_t))} \right] \end{aligned}$$

□

The sensitivity of a trader's belief to the evidence they submit is not, in general, a desirable feature. This can lead to traders strategically sub-sampling their evidence to influence their corresponding belief for a higher payoffs. This in turn means the evidence collected by the market cannot claim to be a representative of the crowd's wisdom, diminishing its usefulness. As such, our goal as the market designer is to ensure that traders can only change their resolution belief by some arbitrarily small ε through selective evidence submission. If combined with a trader payoff that is (1) bounded and (2) incentivizes truthful belief reporting (see Section 5), we can claim ε incentive compatibility (IC) to submit the entire evidence set along with the corresponding true belief. The extent to which this is possible depends on equilibrium notions we consider and the strategy-space of future traders under that notion. We specifically consider two settings now: history-independent strategies, where trader actions are not a function of history, and history-dependent strategies, where the trader actions are.

History-Independent Strategies: In this setting, we show that the sensitivity of trader beliefs about resolution to their own evidence submission can be made arbitrarily small. This is primarily due to future evidence being a random variable that does not depend on how any other agent is acting.

Theorem 4.1. *Under history-independent strategies, consider a trader who arrives at time t with total evidence E_t. For strategies such that the market resolves before timeout even when E_t is fully withheld, the sensitivity of the trader belief between submitting E_t and E'_t is given by:*

$$||q(E_t) - q(E'_t)||_1 \leq \frac{|E_t| + h_{max-1}}{\tau K}.$$

Correspondingly, for any $\varepsilon > 0$, choosing $\tau \geq \frac{2h_{max}-1}{K\varepsilon}$ ensures that for any trader t: $||q(E_t) - q(E'_t)||_1 \leq \varepsilon$.

[5]We generally only consider the trader as deciding to strategically choose a subset of their entire evidence to submit. Manipulation in the form of fabricating or tampering evidence is part of verification which we discussed in Section 6.1.

Proof. From Lemma 4.1, we note that this difference in belief can be expressed as follows:

$$||q(E_t) - q(E_t')||_1 = \sum_{j=1}^{n} \left| \sum_{\boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}} P(\boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}) \frac{\exp(\frac{1}{\tau}\theta_{tj}(E_t))}{\sum_\ell \exp(\frac{1}{\tau}\theta_{t,\ell}(E_t))} - \frac{\exp(\frac{1}{\tau}\theta_{tj}(E_t'))}{\sum_\ell \exp(\frac{1}{\tau}\theta_{t,\ell}(E_t'))} \right|$$

$$= \sum_{\boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}} P(\boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}) \sum_{j=1}^{n} \left| \text{softmax}(\theta_{t1}, \ldots, \theta_{tn} | \boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}, E_t) - \text{softmax}(\theta_{t1}', \ldots, \theta_{tn}' | \boldsymbol{S}_{\boldsymbol{K}_{T-t,*}}, E_t') \right|$$

For brevity, let $\boldsymbol{x} = (\theta_{t1}, \ldots, \theta_{tn})$ under a realization of $S_{K_{T-t},j}$ and submission of true evidence E_t. Further, let $\sigma(\cdot) := \text{softmax}(\cdot)$. We will first determine the sensitivity of the softmax function at $\boldsymbol{x}$ for any perturbation δ. We first note the Jacobian of the softmax is the following:

$$\frac{\partial \sigma_j}{\partial \boldsymbol{x}_m} = \frac{1}{\tau}\sigma_j(\boldsymbol{x})\Big(\mathbf{1}[j = m] - \sigma_m(\boldsymbol{x})\Big) \tag{2}$$

For any direction δ, we then observe the following using the identities $(1 - \sigma_j) = \sum_{m \neq j} \sigma_m$ and $\sum_j \sigma_j = 1$:

$$\begin{aligned}
\nabla \sigma_j(\boldsymbol{x}) \cdot \delta &= \frac{1}{\tau}\sum_{m=1}^{n} \sigma_j(\boldsymbol{x})\Big(\mathbf{1}[j = m] - \sigma_m(\boldsymbol{x})\Big) = \frac{1}{\tau}\delta_j\sigma_j \sum_{j \neq m} \sigma_m - \frac{1}{\tau}\sum_{m \neq j} \delta_m \sigma_j \sigma_m \\
&= \frac{1}{\tau}\sigma_j \left(\sum_{m \neq j} \delta_j \sigma_m - \delta_m \sigma_m \right) = \frac{1}{\tau}\sigma_j \left(\delta_j(1 - \sigma_j) + \sigma_j\delta_j - \sigma_j\delta_j - \sum_{m \neq j} \sigma_m \delta_m \right) \\
&= \frac{1}{\tau}\sigma_j \left(\delta_j - \sum_{m} \sigma_m \delta_m \right) := \frac{1}{\tau}\sigma_j \left(\delta_j - \bar{\delta} \right)
\end{aligned}$$

where $\bar{\delta} = \sum_m \sigma_m \delta_m$ is simply the weighted average of the δ vector according to the weights of σ_m. By the mean value theorem, there exists a $u \in [0, 1]$ such that: $\sigma_j(\boldsymbol{x} + \delta) - \sigma(\boldsymbol{x}) = \nabla \sigma_j(\boldsymbol{x} + u\delta) \cdot \delta$. Then using the above, we have that:

$$\sum_{j=1}^{n} |\sigma_j(\boldsymbol{x} + \delta) - \sigma(\boldsymbol{x})| = \frac{1}{\tau}\sum_{j=1}^{n} |\sigma_j(\boldsymbol{x} + u\delta)(\delta_j - \bar{\delta})| \tag{3}$$

Let J^+ denote the set of indices where $\delta_j - \bar{\delta} \geq 0$ and let J^- denote the set of indices where $\bar{\delta} - \delta_j \geq 0$. Note that $\sum_{j=1}^{n} \sigma_j(\delta_j - \bar{\delta}) = 0$ by definition: softmax sums to 1 and $\bar{\delta}$ is the weighted average with respect to $\sigma_j(\boldsymbol{x} + u\delta)$. As such $V = \sum_{j \in J^+} \sigma_j(\delta_j - \bar{\delta}) = \sum_{j \in J^-} |\sigma_j(\delta_j - \bar{\delta})|$. So it suffices to bound V. Let $p = \sum_{j \in J^+} \sigma_j(\boldsymbol{x} + u\delta)$. Then the following holds:

$$\begin{gathered}
V \leq p \cdot \max_{j \in J^+} \delta_j - \bar{\delta} \quad \text{and} \quad V \leq (1 - p) \cdot \min_{j \in J^-} \bar{\delta} - \delta_j \\
\implies V^2 \leq p(1 - p)\frac{\Delta\delta^2}{4} \leq \frac{\Delta\delta^2}{16} \implies V \leq \frac{\Delta\delta}{4}
\end{gathered}$$

where we use the fact that $p(1 - p) \leq \frac{1}{4}$ and letting $\Delta\delta = \max_{ij} \delta_i - \delta_j$, we use the AM-GM inequality since $\max_{j \in J^+} \delta_j - \bar{\delta} + \min_{j \in J^-} \bar{\delta} - \delta_j = \Delta\delta$. Thus we can claim that:

$$\sum_{j=1}^{n} |\sigma_j(\boldsymbol{x} + \delta) - \sigma(\boldsymbol{x})| \leq \frac{1}{\tau}2V \leq \frac{\Delta\delta}{2\tau} \tag{4}$$

Under history-independent strategies, the future evidence stream — the sequence of arriving traders and the evidence sets they submit — is a random process that does not depend on trader t's action. We can thus fix any realization of this stream, and compare a scenario A in which trader submits complete evidence E_t, to one B where they withhold and selectively choosing what to submit, $E_t' \subseteq E_t$. It suffices to then bound the total variation between the two softmax lotteries. Let $M = |E_t|$, $M' = |E'_t|$. Being strategic is equivalent then to withholding $W = M - M'$ of evidence. Next let F_A and F_B denote the future evidence arriving. Then at resolution trigger in either scenario, we observe that:

$$K_A^* = k(t) + M + F_A \quad K_B^* = k(t) + M - W + F_B \tag{5}$$

Note that every submission included in scenario A (except the withheld pieces) is included in B and that B will occur weakly later. Indeed, B will rely on Z additional pieces that are not in A) to achieve resolution, with $|Z| \le W + h_{max} - 1$ and $|Z| = F_B - F_A = K_B^* - K_A^* + W$.

Let S_j^A denote the number of supports at resolution for alternate j in scenario A and d_j the number of supports for j in the withheld evidence W. Then $S_j^B = S_j^A - d_j + Z_j$, where Z_j is the support for j in the additional evidence. Let δ_j denote the difference in score for alternative j between the two scenarios. We have that:

$$\delta_j = \theta_j^A - \theta_j^B = \frac{S_j^A}{K_A^*} - \frac{S_j^A - d_j + Z_j}{K_B^*} = S_j^A \frac{K_A^* - K_B^*}{K_A^* K_B^*} + \frac{d_j - Z_j}{K_B^*} \tag{6}$$

Then for any parameters δ_j and δ_m, we have that:

$$\delta_j - \delta_m = \underbrace{(S_j^A - S_m^A)\frac{K_A^* - K_B^*}{K_A^* K_B^*}}_{\le \frac{h_{max}-1}{K}} + \underbrace{\frac{d_j - d_m}{K_B^*}}_{\le \frac{W}{K}} + \underbrace{\frac{Z_m - Z_j}{K_B^*}}_{\le \frac{W+h_{max}-1}{K}} \tag{7}$$

$$\le \frac{h_{max}-1}{K} + \frac{W}{K} + \frac{W + h_{max} - 1}{K} \le \frac{2W + 2h_{max} - 2}{K} \le \frac{2|E_t| + 2h_{max} - 2}{K} \tag{8}$$

Plugging this into the softmax sensitivity bound derived in Equation 4, we have that the maximum deviation in beliefs is given by:

$$|q(M) - q(M')|_1 \le \frac{\Delta\delta}{2\tau} \le \frac{|E_t| + h_{max} - 1}{K\tau}. \tag{9}$$

If we wish to ensure there is no more than ε difference between beliefs due to evidence submission across all traders, it suffices to choose the temperature τ uniformly. Since $|E_t| \le h_{max}$ for every trader,

$$\tau \;\ge\; \frac{2h_{max} - 1}{\varepsilon K} \;\ge\; \frac{|E_t| + h_{max} - 1}{\varepsilon K} \quad\implies\quad \Delta q \;\le\; \frac{|E_t| + h_{max} - 1}{\tau K} \;\le\; \varepsilon \tag{10}$$

□

We make a few observations. First the limit to which any trader can affect their belief about resolution depends on the relative size of their evidence with respect to K. Crucially, it does not depend on what that future evidence looks like (i.e. which alternatives they support). For the platform then, it suffices to estimate the largest evidence whale in the market, which is captured by h_{max}. The larger this value, the larger the temperature τ becomes to diminish how much beliefs can change through selective evidence submission.

While this approach can ensure arbitrarily small deviations, it means that market resolution becomes less sharp, especially in the presence of whales. Contrasting to exogenously resolved markets where the market belief gets sharper as one gets closer to resolution, this may not be the case here. This is less of a dilemma than it appears for two reasons. Firstly, resolution time is binding in the exogenous case and as such, it is natural that as one gets closer to it the more uncertainty disappears and the sharper the beliefs. In the endogenous case, the choice of resolution trigger K is arbitrary. Secondly, in exogenous resolution, trader beliefs reflect their opinion on the unrealized world event and are of primary importance. By contrast, in endogenous markets, the submitted beliefs are about resolution and can be viewed separately from the metric implied by the actual evidence being submitted. The event-creator may, for instance, care about the argmax of the scores rather than the softmax lottery. So long as traders are incentivized to submit their full evidence set, any desired function of the cumulative evidence can be faithfully computed, independent of the resolution mechanism's sharpness.

History-Dependent Strategies: The most general setting in extensive form games allow each trader to condition their strategy arbitrarily on the history. This is a very rich space and can lead to phenomena such as herding that make tight sensitivity guarantees as given above difficult. We illustrate this below:

Proposition 4.1. *Under history-dependent strategies, no bound on belief sensitivity that vanishes with the traders evidence share can hold. Specifically, for any K and h_{max} there exist history-dependent strategies under which selective submission shifts a trader's resolution belief by $|q(E_t) - q(E_t')|_1 = \Omega(1/\tau)$, matching, up to constants, the maximum belief shift achievable by modifying all K pieces of evidence.*

Proof. Consider a sequence of agents $a_1, a_2, \ldots$, two possible resolution alternatives o_1, o_2, and $K = 100$ data points needed to resolve the market. Suppose each agent has 2 data points, one supporting each alternative. The agents arrive to market in the given order, and all herd on whatever agent 1 does. Specifically:

- If a_1 only submits evidence for o_1, all the rest will submit only data supporting o_1.
- If a_1 only submits evidence for o_2, all the rest will submit only data supporting o_2.
- If a_1 submits evidence for both o_1 and o_2, all the rest will submit both their data points.

It is clear that by a_1 choosing to submit evidence for o_1 can collapse the market to resolving to o_1 with probability $\text{softmax}(1, 0; \tau)$; submitting for o_2 mean resolving to o_2 with this probability. Thus, even though each agent controls only 2% of the evidence needed to resolve, their resolution belief can completely change based on their evidence submission decision. □

The result above suggests that it may be difficult to ensure DSE for endogenous resolution markets since any trader strategy therein must be optimal even against such intense herding behavior of others, which can greatly affect beliefs and thereby payoffs[6].

5 Payoff Mechanism

Having explored the relationship between market beliefs in endogenous resolution and evidence, we now turn to the question of designing payoffs for our proposed market. This is of central importance in both exogenous and endogenous resolution since payoff structure determines whether traders are correctly incentivized to submit true beliefs and evidence and characterize the equilibrium behaviour. We propose a modified version of a classic payoff function and consider two possible implementations.

5.1 Direct Belief Elicitation and Market Scoring Rules

Classical Perspectives: Our starting point is market scoring rules: a family of incentive-compatible mechanisms that aggregate probabilistic forecasts from multiple participants. They are the workhorse of academic literature on prediction markets, with the most canonical version being the logarithmic market scoring rule. Traders arrive sequentially and submit their complete beliefs over the space of alternatives. At event realization, their payoff is proportional to how well they can predict the realized outcome as compared to the trader immediately before them. Specifically, the payoff considers the difference between the log probabilities assigned to the realized outcome, scaled by a *liquidity parameter* β. The formal definition is given below:

Definition 5.1 (Logarithmic Market Scoring Rule (LMSR))**.** *Consider traders arriving sequentially with trader t holding a private belief q_t. They may submit a possibly different belief $\hat{q}_t$ to the mechanism and their payoff when the event resolves to outcome ω is given by:*

$$\textit{payoff}_t(\hat{q}_t) = \beta \log \hat{q}_t^\omega - \beta \log q_{t-1}^\omega$$

[6] This herding example can be re-worked to show that truthful submission of complete evidence set cannot be a DSE.

In such markets, it is (1) dominant strategy incentive compatible (DSIC) for all traders to submit their true belief $q' = q$ and (2) the platform's loss is bounded by $\beta \log(n)$.

There are two core properties of this mechanism. First, it is strictly optimal for each trader to submit their true belief q_t and not manipulate it. Second, while the cumulative amount (sum of all the payoffs) paid out by the platform running this mechanism is non-zero – i.e. the platform can lose money – this is bounded by $\beta \log n$. To see this, observe that sum of the payoffs form a telescoping series, with the total payout being the difference between the log probabilities assigned in the first ($t = 0$) and last belief ($t = T$). In the worst case, the market starts out at uniform belief and converges to the resolved winner with probability 1. While the platform losing money can seem unsatisfactory, this loss can be economically interpreted as the cost of information aggregation with strong incentive guarantees.

Note that trader t can also lose money in the mechanism (i.e. have negative payoff) if they were worse at predicting than their predecessor. This is natural in any market with risk. Lastly and from an operational perspective, observe that traders submit complete beliefs up-front and receive or pay out money at market resolution.

Evidence-Augmented LMSR: We now modify this mechanism to our generalized prediction market that allows for the submission of evidence alongside beliefs. Recall that $r(E)$ is a metric to compute the quality of an evidence set E, and $R_t = r(E_t^{total})$ denotes the quality of the cumulative evidence in the market at time t. The core novelty we introduce is a *dynamic* liquidity parameter $\beta(\cdot)$ that adjusts based on cumulative evidence quality R_t. We formalize this below:

Definition 5.2 (Evidence-Augmented LMSR)**.** *Consider a trader arriving at time t with evidence E_t. Then their payoff in submitting a subset of evidence $E'_t \subseteq E_t$ and a belief $\hat{q}_t$ (which can be distinct from their true belief) when the resolution outcome is ω is given by[7]:*

$$\begin{aligned} \mathit{payoff}_t(\hat{q}_t, E'_t) &= \beta(r(E_{t-1}^{total} \cup E'_t)) \log \hat{q}_t^{\omega} - \beta(r(E_{t-1}^{total})) \log q_{t-1}^{\omega} \\ &= \beta(R_t) \log \hat{q}_t^{\omega} - \beta(R_{t-1}) \log q_{t-1}^{\omega} \end{aligned}$$

Observe that if a trader t provides evidence that increases the cumulative quality, their log probability will be scaled by a different β than the trader before them at $t-1$. In providing no evidence (or evidence of 0 marginal quality according to r) they face the same liquidity β as their predecessor and their payoff is similar to classical LMSR payoff. In other words, this mechanism is a strict generalization. Observe that since the trader is free to submit any belief and evidence subset to the mechanism and the platform wants to limit exposure, our goal is to ensure (1) platform loss is bounded, (2) traders are incentivized to report true belief, which can correspond to the evidence they submit for endogenous resolution, and (3) traders are incentivized to submit their whole evidence set. The first two goals mirror the guarantees provided by standard LMSRs and we start with this.

Proposition 5.1. *Suppose trader t submits evidence set $E'_t \subseteq E_t$ and holds true belief q_t or $q(E'_t)$ for exogenous or endogenous resolution respectively. Then under the Evidence Augmented LMSR, it is optimal for them to submit q_t or $q(E'_t)$ in either respective case. Further, the platform loss is bounded by $\beta(R_0) \log n := \beta_0 \log n$.*

Proof. Fix the evidence submission E'_t and for brevity, let q_t denote the true belief (so $q_t = q_t(E'_t)$ for endogenous resolution). Let $\beta_t = \beta(R_t)$ and $\beta_{t-1} = \beta(R_{t-1})$. Next, observe that trader t's objective can be simplified as such:

$$\max_{\hat{q}} \mathbb{E}_{q_t} \left[\beta_t \log \hat{q}_t - \beta_{t-1} \log q_{t-1}\right] = \max_{\hat{q}} \beta_t \cdot \mathbb{E}_{q_t}[\log \hat{q}] \tag{11}$$

Next, observe the following where the first inequality holds due to Jensen's

$$0 = \log \mathbb{E}_{q_t} \left[\frac{\hat{q}_t}{q_t}\right] \geq \mathbb{E}_{q_t} \left[\log \hat{q}_t\right] - \mathbb{E}_{q_t} \left[\log q_t\right] \tag{12}$$

[7] For exogenously resolved events, traders come to market with some true belief q_t that is unaffected by what evidence they choose to submit. For endogenous resolution, the relationship between evidence submission and corresponding true belief is denoted $q_t(E'_t)$ and is outlined formally in Lemma 4.1

As such, the trader's payoff is maximized when they report their true belief q_t. As for the platform loss, observe that this reduces to a telescoping sum:

$$\sum_{t=0}^{T} \beta(R_t) \log q_t^{\omega} - \beta(R_{t-1}) \log q_{t-1}^{\omega} = \beta(R_T) \log q_T^{\omega} - \beta(R_0) \log q_0^{\omega}$$

Since log probabilities are negative, this is maximized when final belief is 1 on the resolved outcome and the initial belief is uniform. This yields a maximum loss of $\beta(R_0) \log n$. □

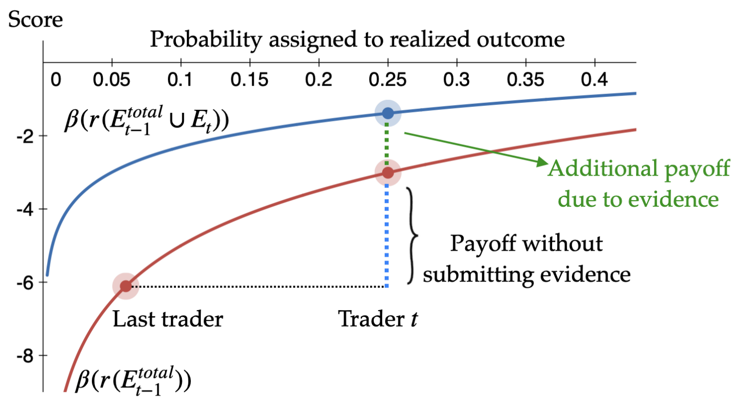

Figure 2: Log Scoring curves under different liquidity parameters. The red denotes the curve before evidence submission, and in blue, the curve after evidence submission. We visualize trader $t-1$ submitting probability 0.05 on realized outcome and trader t submitting probability 0.25 along with evidence E_t.

We now focus on the last desiderata and argue that traders are incentivized to submit their complete evidence set. This depends crucially on the shape of this liquidity curve $\beta(\cdot)$. Since log probabilities are negative, for positive $\beta(R)$ that *decreases* in R, the log probability curve shifts up and becomes closer to 0. This means that keeping the belief constant (trader t submits the same belief as $t-1$), by simply submitting quality evidence, they decrease β, resulting in non-zero payoff to this trader. When they submit a distinct belief alongside this, they obtain the standard LMSR payoff (which could be positive or negative) alongside this non-zero evidence payoff (see Figure 2 for a visual). This is the core feature that incentivizes full evidence submission and we formalize it Theorem 5.1. For exogenous resolution, the incentive to submit complete evidence hold strictly; for endogenous resolution, we appeal to Theorem 4.1 to give an ε incentive compatibility guarantee. We first state a technical lemma before proceeding to the main result.

Lemma 5.1. *Let $\beta(R)$ be positive and decreasing in R and suppose trader t submits evidence set $E_t' \subseteq E_t$. Then $x = \beta(r(E_{t-1}^{total})) - \beta(r(E_{t-1}^{total} \cup E_t'))$ is nonnegative and the expected payoff to trader t can be decomposed as:*

$$\textit{payoff}_t(E_t') = \underbrace{\beta(R_{t-1}) D_{KL}(q_t(E_t') || q_{t-1})}_{\textit{(1) Belief payoff}} + \underbrace{x H(q_t(E_t'))}_{\textit{(2) Evidence payoff}}$$

Proof. Under exogenous resolution, q_t is independent from this and under endogenous resolution, let $q_t = q_t(E')$ for brevity. Due to Proposition 5.1, the trader will always submit this belief, so the payoff is now only a function of the evidence. Since β is non-negative and decreasing and $r(E)$ is always monotone increasing, there must exist a non-negative x such that:

$$\beta(R_t) = \beta(r(E_{t-1}^{total} \cup E_t')) = \beta(r(E_{t-1}^{total})) - x = \beta(R_{t-1}) - x \tag{13}$$

Then their expected payoff can be expressed as follows, with $H(\cdot)$ denoting the Shannon Entropy function and $\mathrm{D_{KL}}$ the KL divergence:

$$\begin{aligned}
\text{payoff}_t(E'_t) &= \sum_\omega q_t^\omega \left[\beta(R_t)\log q_t^\omega - \beta(R_{t-1})\log q_{t-1}^\omega\right] \\
&= \beta(R_{t-1})\sum_\omega q_t^\omega[\log q_t^\omega - \log q_{t-1}^\omega] - x\sum_\omega q_t(\omega)\log q_t^\omega \\
&= \beta(R_{t-1})\sum_\omega q_t^\omega[\log q_t^\omega - \log q_{t-1}^\omega] + xH(q_t) \\
&= \beta(R_{t-1})\sum_\omega q_t^\omega \log\frac{q_t^\omega}{q_{t-1}^\omega} + xH(q_t) = \underbrace{\beta(R_{t-1})\mathrm{D_{KL}}(q_t||q_{t-1})}_{\text{(1) Belief payoff}} + \underbrace{xH(q_t)}_{\text{(2) Evidence payoff}}
\end{aligned}$$

□

Theorem 5.1. *Consider each trader adopting a history independent strategy of submitting their entire evidence set E_t alongside reporting their corresponding true belief q_t. Then:*

- *Under exogenous resolution, this is a dominant strategy equilibrium.*
- *Under endogenous resolution, for any $\varepsilon > 0$, there exists a temperature τ such that this is an ε-subgame perfect Nash equilibrium.*

Proof. For exogenous resolution, q_t has no relationship with E'_t. This means that the trader can only optimize the second term of their payoff decomposition outlined in Lemma 5.1: $[\beta(r(E_{t-1}^{total})) - \beta(r(E_{t-1}^{total} \cup E'_t))]H(q_t(E'_t))$. However, since Shannon entropy is always non-negative, they maximize this by maximizing the multiplier, which is maximized by submitting the complete evidence set E_t due to the monotone increasing nature of $r(E)$ and decreasing nature of $\beta(\cdot)$[8]. In other words, it is optimal to submit the complete evidence set under exogenous resolution regardless of the actions of any other traders, past or future.

For endogenous resolution, the resolution belief that traders hold depends on the evidence they submit. While a trader can deviate in both submitting partial evidence and incorrect beliefs, Proposition 5.1 ensures that it is always dominant for them to report the true belief corresponding to whatever evidence E'_t they submit. As such, we only need to consider deviation in evidence set.

Recall the resolution rule here computes a softmax over the scores for each alternative j, where score $\theta_j \in [0,1]$. This means that the resolution probability given a complete set of k evidence items E must satisfy:

$$p^\omega(E) \geq \frac{1}{\sum_\ell \exp(\frac{1}{\tau}\theta_\ell)} \geq \frac{1}{n\exp(\frac{1}{\tau})} = \frac{1}{n}\exp(-\tfrac{1}{\tau}) := p_{min}$$

Since Proposition 5.1 ensures only true beliefs are submitted, and any true belief q_{t-1} is a mixture over such beliefs (see Lemma 4.1), such beliefs must lie in the simplex interior. In other words: $q_{t-1}^\omega \geq p_{\min} > 0$ for all ω. Next, let q_1, q_2 denote beliefs that can be induced by trader t due to evidence submission $E_1 \subseteq E_t$ and any subset $E_2 \subseteq E_t$ respectively. Then we observe:

$$\text{payoff}_t(E_1) - \text{payoff}_t(E_2) = \beta(R_{t-1})\left[\mathrm{KL}(q_1\|q_{t-1}) - \mathrm{KL}(q_2\|q_{t-1})\right] + \left[x_1H(q_1) - x_2H(q_2)\right]. \quad (14)$$

Let $f(p) := \mathrm{KL}(p\,\|\,q_{t-1}) = \sum_\omega p^\omega \log p^\omega - \sum_\omega p^\omega \log q_{t-1}^\omega$ denote the KL function on the relative interior of the simplex. Observe that the gradient is given by: $\frac{\partial f}{\partial p^\omega} = \log p^\omega + 1 - \log q_{t-1}^\omega$. Let $C := \log(1/p_{\min})$. Then:

$$\|\nabla f(\xi)\|_\infty \leq 1 + |\log(p^\omega)| + |\log(q_{t-1}^\omega)| \leq 2C + 1 := L_{\mathrm{KL}}.$$

By the mean value theorem, there exists ξ on the segment between q_1 and q_2 such that $f(q_1) - f(q_2) = \nabla f(\xi)\cdot(q_1 - q_2)$. Combining this with Holder's inequality, we have (where $\delta = \|q_1 - q_2\|_1$):

$$\left|\mathrm{KL}(q_t\,\|\,q_{t-1}) - \mathrm{KL}(q_{t'}\,\|\,q_{t-1})\right| \leq \|\nabla f(\xi)\|_\infty \cdot \|q_t - q_{t'}\|_1 \leq L_{\mathrm{KL}}\,\delta.$$

[8]Note that beliefs are always in the simplex interior.

Now for the entropy term. Let $E_1 = E_t$ and E_2 any subset as before. Then due to r being monotone increasing and β decreasing and non-negative, it must be that $x_1 \geq x_2$. We can thus decompose:

$$x_1 H(q_1) - x_2 H(q_2) \ = \ (x_1 - x_2)\, H(q_1) + x_2 \left[H(q_1) - H(q_2)\right]$$

The first term is always nonnegative and thus favours full evidence submission. So we focus on the second term. Observe that for Shannon entropy $H(p)$, $\partial H/\partial p^\omega = -\log p^\omega - 1$. Over the line segment defined by q_1, q_2, the norm of this gradient is bounded by $\log(1/p_{\min}) + 1 =: L_H$. Thus by applying the mean value theorem and noting that $x_2 \leq \beta(R_{t-1}) \leq \beta_0$, we have:

$$x_1 H(q_1) - x_2 H(q_2) \geq -\beta_0 L_H \delta$$

Combining it all together, we have that the payoff from submitting the complete evidence set $E_1 = E_t$ is at most $\delta(L_{KL}\beta_0 + L_H\beta_0)$ worse than submitting any other set E_2, where $\delta := ||q_1 - q_2||_1$ is the belief sensitivity.

We now argue that each trader t adopting a history independent strategy of submitting their entire evidence set E_t is a ε-SPNE. Pick any time t and history h_t which has accumulated evidence E_t^{total}, and consider any deviating action $E_t' \subseteq E_t$ at this node. Since the continuation prescribed by the full-submission profile is constant across histories, the future evidence stream is identical under E_t and E_t'. We partition the realizations of this stream by whether the market resolves under the withholding branch E_t', and bound the deviation gain on each event separately.

On realizations where the market resolves under E_t' (and hence also under E_t, whose trigger is weakly earlier), the conditions of Theorem 4.1 apply at (t, h_t), so the belief sensitivity satisfies $||q_1 - q_2||_1 = \delta \leq \frac{|E_t| + h_{max} - 1}{\tau K}$ and the analysis above bounds the deviation gain by $\frac{\beta_0(L_{KL}+L_H)(|E_t|+h_{max}-1)}{\tau K}$.

On realizations where the market does not resolve under E_t', the deviating trader's payoff is exactly 0 due to timeout. Under full submission E_t, the market may or may not resolve on such a realization; if it times out the payoff is again 0, and if it resolves the expected payoff is $\beta(R_{t-1})\mathrm{KL}(q_1 \| q_{t-1}) + x_1 H(q_1) \geq 0$ by the decomposition of Lemma 5.1, both terms being non-negative. Hence the deviation gain is at most 0.

The expected deviation gain is a mixture of the two events and is therefore bounded by the first case. Thus:

$$\text{payoff}_t(E_t) - \text{payoff}_t(E_t') \geq -\frac{\beta_0(L_{KL} + L_H)(|E_t| + h_{max} - 1)}{\tau K} := -\varepsilon \implies \tau \geq \frac{\beta_0(L_{KL} + L_H)(2h_{max} - 1)}{\varepsilon K}$$

By selecting the temperature constant accordingly, we ensure that no deviating action can achieve more than ε additional payoff over full submission for any trader, at any history.

□

Given the result above, our market satisfies all axioms set out in Definition 3.3. It can be resolved both exogenously and endogenously (Axiom 1) and incentivizes traders to submit their complete evidence set (Axiom 2) and corresponding true belief (Axiom 3) — albeit under different equilibrium guarantees for exogenous and endogenous cases. The market belief in an LMSR is always the last belief submitted (Axiom 4) and it is clear from the decomposition above that jointly submitting evidence and belief yields the same payoff as splitting this into two consecutive trades in either order, which ameliorates any arbitrage concern (Axiom 5). This also highlights that risk-averse traders in this market can achieve a *risk-free* payout through evidence submission. Specifically, a risk-averse trader t can submit $q_t = q_{t-1}$ and their evidence set E_t to attain a payoff of $[\beta(R_{t-1}) - \beta(R_t)]H(q_{t-1}) \geq 0$. Since the last trader's belief captures the market belief, the marginal value of evidence is highest when the market is most entropic and the liquidity curve is the steepest. This is an attractive property that goes beyond the core axioms highlighted in Section 3. We state this formally below:

Corollary 1. *Consider a risk-averse trader t who will not engage with a market if any negative payoff is possible. By submitting $q_t = q_{t-1}$ and their evidence set E_t, they achieve a non-negative payoff in every realization with expected value $[\beta(R_{t-1}) - \beta(R_t)]H(q_{t-1}) \geq 0$.*

Lastly, one may have two nagging concerns about endogenous resolution given the results above. Is insufficient evidence such that the market times out part of some equilibrium strategy? Given that this is clearly undesirable, we give a positive result below ruling out this possibility. Second, we sketch an approximate SPNE for this setting, but do not comment on the possibility of DSE for this setting. Leveraging the herding example of Proposition 4.1, we show in Proposition 5.3 that truthful belief and evidence submission cannot constitute such an equilibrium in this setting, implying SPNE is the best we can hope for.

Proposition 5.2. *For endogenous resolution, any strategy profile under which the market times out with probability 1 is not a subgame-perfect Nash equilibrium.*

Proof. Consider the extensive form of our game. For a node (t, h_t) with accumulated evidence E_t^{total}, call a trader's evidence set E_{t_i} *robustly informative at* (t, h_t) if it has strictly positive marginal quality with respect to every superset of the current record: for all $E \supseteq E_t^{total}$, $r(E \cup E_{t_i}) > r(E)$. Call a node (t, h_t) *m-resolution-feasible* if there exist m traders $t \leq t_1 < t_2 < \cdots < t_m \leq T$, each holding robustly informative evidence at (t, h_t), such that $\sum_{i=1}^{m} |E_{t_i}| \geq K - k(t)$. This property is preserved under any play: evidence only accumulates, so $K - k(t)$ only shrinks along any path, and robust informativeness at (t, h_t) implies robust informativeness at any successor node since the future record is a superset of the current one. By assumption (see Section 3), the root node is m-resolution-feasible for some finite m.

We prove the following claim by strong induction on m: *in any SPNE, at every m-resolution-feasible node (t, h_t), (i) the continuation value of trader t_1 (the first trader in the feasibility certificate) at their realized history is strictly positive, and (ii) the market resolves with strictly positive probability under the equilibrium continuation from (t, h_t).*

Base case ($m = 1$). Here trader t_1 can unilaterally trigger resolution: $|E_{t_1}| \geq K - k(t)$. Consider the action available to t_1 of submitting E_{t_1} in full while reporting $\hat{q} = q_{t_1-1}$. This action triggers resolution with certainty, and by Corollary 1 its realized payoff on every resolution outcome ω is $[\beta(R_{t_1-1}) - \beta(R_{t_1})](-\log q_{t_1-1}^{\omega}) > 0$, where strict positivity follows from β strictly decreasing, the strictly positive marginal quality of E_{t_1}, and the interiority of q_{t_1-1}. Sequential rationality requires trader t_1's equilibrium continuation value to weakly exceed the value of this available action, giving (i). For (ii), observe that the timeout event yields payoff exactly 0, while payoffs conditional on resolution are bounded above (beliefs are interior and $\beta \leq \beta_0$). A strictly positive continuation value is therefore impossible if the equilibrium continuation times out with probability 1; hence the market resolves with strictly positive probability.

Inductive step. Suppose the claim holds for all $m' < m$, and let node (t_1, h_{t_1}) be m-resolution-feasible. Consider this node, and the action available to them of submitting E_{t_1} in full. If this submission can trigger resolution, we are in the base case and the argument above applies directly. Otherwise, this evidence accumulated and we move to a node (t_2, h_{t_2}) at which point:

$$K - k(t_2) \leq K - k(t_1) - |E_{t_1}| \leq \sum_{i=2}^{m} |E_{t_i}| \tag{15}$$

which is by definition $m-1$-resolution-feasible. Since an SPNE requires the continuation profile to be an equilibrium at *every* node, including those off the equilibrium path, the induction hypothesis applies at (t_2, h_{t_2}): the market resolves with strictly positive probability under the equilibrium continuation from there. The full submission action thus yields trader t_1 at the m-resolution-feasible node a strictly positive payoff. Any SPNE continuation must achieve at least this utility and a continuation that times out with probability 1 has 0 payoff. Thus, claims (i) and (ii) follow at this node exactly as in the base case: sequential rationality gives t_1 a strictly positive equilibrium continuation value, which is incompatible with the market timing out with probability 1.

The proposition follows by applying claim (ii) at the root node: under any SPNE, the market resolves with strictly positive probability, so no profile that times out with probability 1 can be an SPNE. □

Proposition 5.3. *There exists an instance of the endogenous resolution markets under which complete evidence submission is not a dominant strategy equilibrium.*

Proof. Consider the herding instance outlined in Proposition 4.1: A sequence of agents $a_1, a_2, \ldots$, two possible resolution alternatives o_1, o_2, and $K = 100$ data points needed to resolve the market. Suppose each agent has 2 data points, one supporting each alternative. The agents arrive to market in the given order, and all herd on whatever agent 1 does.

The market starts with uniform belief $q_0 = (0.5, 0.5)$ and no evidence. Consider two possible profiles: the herding profile γ_h where trader a_1 only submits the evidence for o_1 and all future traders herd on this, with everyone having belief $q_h(\tau) = \text{softmax}(1, 0; \tau)$. Let γ_t denote the truthful profile where all traders submit both their complete evidence and the true belief $q_t = (0.5, 0.5)$. Given that the market starts at uniform belief, the traders simply get paid for evidence and not belief updates in this setting:

$$\begin{aligned}\text{payoff under } \gamma_t \text{ for any agent} &= [\beta(r(E_{t-1}^{total})) - \beta(r(E_{t-1}^{total} \cup E'_t))]H([0.5; 0.5]) \\ &= [\beta(r(E_{t-1}^{total})) - \beta(r(E_{t-1}^{total} \cup E'_t))]\end{aligned}$$

Now consider the herding profile specifically for trader t_1. Their payoff here includes both a belief update component and an evidence component:

$$\text{payoff to trader } a_1 \text{ under } \gamma_h = \beta(r(\emptyset))D_{KL}(q_h(\tau); q_0) + [\beta(r(\emptyset)) - \beta(r(E'_1))]H(q_h)$$

Let $\beta(r(\emptyset)) = 1$ and the liquidity function decrease arbitrarily slow. That is $\beta(r(\emptyset)) - \beta(r(E'_1)) \le \varepsilon$, where we can choose any $\varepsilon > 0$ to satisfy this. Then the payoff to trader a_1 under γ_t is ε and under γ_h it is at least $D_{KL}(q_h(\tau); q_0)$. For any given τ, choosing the liquidity slope such that $\varepsilon < D_{KL}(q_h(\tau); q_0)$ ensures that if other traders are selecting the herding strategy, trader a_1 will prefer it over the truthful one. Thus, truthful behaviour is not DSE under endogenous resolution. □

5.2 An Equivalent Automated Market-Maker

While the evidence augmented LMSR satisfies all our core desiderata, it suffers from two operational challenges. In LMSRs, traders submit beliefs without any stake/money; it is only upon resolution that a platform must collect any amount they are owed. This can lead to collection issues in practice. Second, traders in practice are more accustomed to equity markets which allow them to buy "shares" of different outcomes and hold for payoffs. This is also the design of markets like Kalshi and Polymarket where traders can purchase shares of different alternatives.

Fortunately, LMSRs can be equivalently implemented through automated market makers (AMMs) that offer such an interaction and alleviate the twin concerns above. Each "share" of an alternative is an Arrow-Debreu contract that pays out \$1 if that alternative is realized, and \$0 if not. The core argument rests on introducing a *cost-function* that takes as input the total number of shares sold for each alternative so far. The derivative of this denotes the instantaneous share price of any alternative. A trader's execution price for a batch of shares is determined by integrating over their orders. The classic result of [8, 3] show that incentives in this market coincide exactly with LMSRs.

We now show this equivalence can be extended to our evidence-augmented LMSR. That is, we propose a modified cost function defined over both cumulative shares and evidence, whose derivative defines the marginal price of shares and evidence. Formally, let $s_{i,t}$ denote the number of shares (Arrow-Debreu contracts) that have been sold for alternative i up until time t, with $\boldsymbol{s}_t \in \mathbb{R}^n_{\ge 0}$ denoting the vector of shares sold[9]. Then we define the cost-function as follows:

Definition 5.3 (Evidence-Augmented AMM)**.** *Consider a market at time t where $\boldsymbol{s}_t$ shares have been sold and cumulative evidence of R_t collected. Then we define the evidence-augmented cost function $C(\boldsymbol{s}_t, R_t)$ and*

[9]We do not consider/allow traders to be short in the AMM. From an information aggregation perspective where trading shares is simply a conduit to expressing beliefs, any belief that can be expressed through shorting shares, can be equivalently expressed through purchasing shares due to the zero-sum nature of the prices. For binary alternatives for instance, shorting alternative A is equivalent to buying alternative B

price function $\pi(\boldsymbol{s}_t, R_t)$ *at this state as:*

$$C(\boldsymbol{s}_t; R_t) = \beta(R_t) \log \sum_i \exp\left(\frac{s_{i,t}}{\beta(R_t)}\right) \quad ; \quad \pi_i(\boldsymbol{s}_t; R_t) = \frac{\partial C}{\partial s_{i,t}} = \frac{\exp(\frac{s_{i,t}}{\beta(R_t)})}{\sum_i \exp(\frac{s_{i,t}}{\beta(R_t)})} \tag{16}$$

The partial derivative of cost, $\pi_i(\boldsymbol{s}_t; R_t)$*, is the instantaneous price of buying 1 share of alternative* i*, and* $\frac{\partial C}{\partial R}$ *can be seen as the price of providing evidence. A trader looking to buy (sell)* $\boldsymbol{z}_t \in \mathbb{R}^n$ *shares and supply evidence* E_t *then pays (receives)* [10]*:*

$$C(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, R_{t-1}) = \int_{(\boldsymbol{s}_{t-1}, R_{t-1})}^{(\boldsymbol{s}_{t-1}+\boldsymbol{z}_t, R_t)} \frac{\partial C}{\partial R} dR + \sum_{i=1}^{n} \pi_i(\boldsymbol{s}; R) d\boldsymbol{s} \tag{17}$$

Observe that the prices here are not "locked in" for a given order. Every increment of the order (shares or evidence) being executed changes the prices at which the next increment is executed. That said, since the prices are the gradient of a potential function (i.e. the cost function), the corresponding vector field is conservative and thus, any integral is path independent. This means the manner in which the order is executed does not matter, giving us the no-arbitrage behaviour we desire (Axiom 5). It is also immediate from the cost function that the sum of the prices for the n alternatives sum to 1. They can thus be interpreted as the current market belief over the outcomes (Axiom 4).

The core novelty in our cost expression as compared to the classic one proposed by Hanson [8], Chen and Pennock [3] is the dependence on the evidence quality R in the liquidity parameter $\beta(\cdot)$. Indeed, this dynamism has a natural interpretation in the AMM setting. Rational and risk-neutral traders have some true belief q_t over outcomes which can be distinct from the current market prices/belief. As their order is executed, prices shift, and if they are not budget-constrained, they trade until the prices reflect their belief. By supplying evidence and decreasing β, they make the price curve π steeper, implying less shares need to be bought to shift the market belief to their q_t. We now formalize this intuition and more generally show that for such traders their optimal payoff under this AMM market coincides exactly with the evidence-augmented LMSR, thereby inheriting all its strong incentive properties.

Theorem 5.2. *For an evidence augmented AMM at state* $(\boldsymbol{s}_{t-1}, R_{t-1})$*, the expected payoff from a trade* $(\boldsymbol{z}_t, E_t)$ *is equivalent to the expected profit from updating an evidence-augmented LMSR from belief* $q_{t-1} = \pi(\boldsymbol{s}_{t-1}; R_{t-1})$ *and liquidity* $\beta(R_{t-1})$ *to a belief* $q_t = \pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t; r(E_{t-1}^{total} \cup E_t))$ *and liquidity* $\beta(r(E_{t-1}^{total} \cup E_t))$*. Then for a trader who has true belief* q_t *and evidence* E_t*, there exists a trade* $\boldsymbol{z}_t$ *such that* $\pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t; r(E_{t-1}^{total} \cup E_t)) = q_t$ *and it is a* ε*-SPNE/DSE for the trader to submit trade* $(\boldsymbol{z}_t, E_t)$ *as opposed to any other trade under endogenous/exogenous resolution respectively.*

Proof. Suppose the realized outcome at resolution is alternative i. Let us look at the payoff under the AMM market. The trader earns $\boldsymbol{z}_{i,t}$ for the winning alternative but pays the cost of transaction:

$$\begin{aligned}
\text{AMM payoff} &= \boldsymbol{z}_{i,t} - [C(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, R_{t-1})] \\
&= \boldsymbol{z}_{i,t} - \left[\beta(r(E_{t-1}^{total} \cup E_t)) \log \sum_j \exp\left(\frac{s_{j,t-1} + z_{j,t}}{\beta(r(E_{t-1}^{total} \cup E_t))}\right) - \beta(R_{t-1}) \log \sum_j \exp\left(\frac{\boldsymbol{s}_{j,t-1}}{\beta(R_{t-1})}\right)\right]
\end{aligned}$$

Now consider the evidence-augmented LMSR where the last reported belief is $q_{t-1} = \pi(\boldsymbol{s}_{t-1}; R_{t-1})$. Consider the true belief of a trader to be $q_t = \pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t; r(E_{t-1}^{total} \cup E_t))$ and they have evidence E_t. Since the modified LMSR has truthful reporting as part of an equilibrium strategy, the trader will report this true

[10] Recall $R_{t-1} = r(E_{t-1}^{total})$.

value and submit their evidence under this. Their payoff then is:

$$\begin{aligned}
\text{LMSR payoff} &= \beta(r(E_{t-1}^{total} \cup E_t)) \log \pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t; r(E_{t-1}^{total} \cup E_t)) - \beta(R_{t-1}) \log \pi(\boldsymbol{s}_{t-1}; R_{t-1}) \\
&= \beta(r(E_{t-1}^{total} \cup E_t)) \log \left[\frac{\exp(\frac{s_{i,t-1}+\boldsymbol{z}_{i,t}}{\beta(r(E_{t-1}^{total} \cup E_t))})}{\sum_j \exp(\frac{s_{j,t-1}+\boldsymbol{z}_{j,t}}{\beta(r(E_{t-1}^{total} \cup E_t))})} \right] - \beta(R_{t-1}) \log \left[\frac{\exp(\frac{s_{i,t-1}}{\beta(R_{t-1})})}{\sum_j \exp(\frac{s_{j,t-1}}{\beta(R_{t-1})})} \right] \\
&= (\boldsymbol{s}_{i,t-1} + \boldsymbol{z}_{i,t}) - \beta(r(E_{t-1}^{total} \cup E_t)) \log \sum_j \exp(\frac{s_{j,t-1}+\boldsymbol{z}_{j,t}}{\beta(r(E_{t-1}^{total} \cup E_t))}) - \boldsymbol{s}_{i,t-1} + \beta(R_{t-1}) \log \sum_j \exp(\frac{s_{j,t-1}}{\beta(R_{t-1})}) \\
&= \boldsymbol{z}_{i,t} - \left[\beta(r(E_{t-1}^{total} \cup E_t)) \log \sum_j \exp(\frac{s_{j,t-1}+z_{j,t}}{\beta(r(E_{t-1}^{total} \cup E_t))}) - \beta(R_{t-1}) \log \sum_j \exp(\frac{s_{j,t-1}}{\beta(R_{t-1})}) \right]
\end{aligned}$$

which we observe to be the same payoff under the AMM. Next note that in the evidence-augmented LMSR, submitting their true belief q_t along with the evidence E_t is ε optimal for the trader. Next we note that since the price function is essentially a scaled softmax function, it is surjective over the simplex. So for any belief q' and evidence $E_t' \subseteq E_t$, there exists a trade $\boldsymbol{z}'$ such that $\pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}'; r(E_{t-1}^{total} \cup E_t')) = q'$. Now suppose that the trader chooses a $(\boldsymbol{z}', e')$ such that $q' \neq q$. As we show above, this would lead to the same payoff as submitting belief $q' \neq q$ under the modified LMSR, which we know leads to strictly worse payoff than submitting true beliefs (Proposition 5.1). Next, consider an trader choosing a $(\boldsymbol{z}', E_t' \subset E_t)$ such that $\pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}'; r(E_{t-1}^{total} \cup E_t') = q_t$. Again, as we show above, this would lead to the same payoff as submitting belief q_t with partial evidence $E_t' \subset E_t$ under modified LMSR. But we know from Theorem 5.1 that the trader payoff in the LMSR will be $-\varepsilon$ higher in submitting the full evidence along with the true belief. Thus, the trader only submits the trade with full evidence that updates market price to q. □

With the result above, we satisfy the incentive compatibility desiderata of Axioms 2 and 3. Further, since any trader interaction here has an analogue in the LSMR realm, we also inherit the $\beta_0 \log n$ bounded worst-case platform loss for running this mechanism. We lastly turn to understanding the value of evidence in this market. In the evidence-augmented LMSR, traders are able to directly submit belief and evidence. As such, we decomposed the payoff into the value of belief update and the value of evidence submission. The latter was shown to be proportional to the entropy of the market belief.

In the evidence-augmented AMM, there is no explicit way to submit "belief" since beliefs are endogenously defined as a function of both shares bought and evidence submitted. Even if one contributes evidence, the market belief shifts since the price curve re-adjusts to be steeper. So while the equilibrium price after a rational risk-neutral trader completes their order (buying shares and submitting evidence) $\pi(\boldsymbol{s}_{t-1} + \boldsymbol{z}, r(E_{t-1}^{total} \cup E_t))$ should match their true belief q_t there is no clean decomposition of the AMM payoff into belief update and evidence update in the style done in LMSR. That said, a more natural decomposition in the AMM setting is on the execution cost based on the shares bought $\boldsymbol{z}$ and evidence submitted E_t:

$$\begin{gathered}
\text{Execution Cost: } C(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, R_{t-1}) \\
= \underbrace{C(\boldsymbol{s}_{t-1}, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, R_{t-1})}_{\text{cost of evidence}} + \underbrace{C(\boldsymbol{s}_{t-1} + \boldsymbol{z}_t, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, r(E_{t-1}^{total} \cup E_t))}_{\text{cost of position/shares}}
\end{gathered}$$

If the cost of the first term is shown to be negative (i.e the trader received a discount for submitting evidence) then we spiritually recover the same incentives for risk averse traders: if one wishes to take on no risk (or is budget constrained), they can buy no shares and can still get full benefit for submitting evidence. It is clear from Definition 5.3 that this evidence cost is exactly equal to $\int_{R_{t-1}}^{R_t} \frac{\partial C}{\partial R}$ where $\frac{\partial C}{\partial R}$ is the marginal cost of evidence submission. We now show that for decreasing β, this is negative and proportional to the entropy of the market prices, spiritually similar to the LMSR evidence payoff result.

Theorem 5.3. *At any market state $(\boldsymbol{s}, R)$, the marginal cost of evidence is $\frac{\partial C}{\partial R} = \frac{\partial \beta}{\partial R} H(\pi(\boldsymbol{s}, R))$. As such, $C(\boldsymbol{s}_{t-1}, r(E_{t-1}^{total} \cup E_t)) - C(\boldsymbol{s}_{t-1}, R_{t-1})$ is always non-positive for any decreasing $\beta(R)$ function.*

Proof. Let $K = \sum_i \exp\left(\frac{s_i}{\beta(R)}\right)$. The price of information/evidence $\frac{\partial C}{\partial R}$ is given as follows:

$$\frac{\partial C}{\partial R} = \frac{\partial \beta}{\partial R}\log(K) + \frac{\beta(R)}{K}\frac{\partial K}{\partial R} \quad ; \quad \frac{\partial K}{\partial R} = -\sum_i \exp\left(\frac{s_i}{\beta(R)}\right)\frac{s_i}{\beta^2(R)}\frac{\partial \beta}{\partial R}$$

We therefore have that:

$$\frac{\partial C}{\partial R} = \frac{\partial \beta}{\partial R}\left[\log(K) - \frac{1}{K}\sum_i \exp\left(\frac{s_i}{\beta(R)}\right)\frac{s_i}{\beta(R)}\right] = \frac{\partial \beta}{\partial R}\left[\log(K) - \sum_i \underbrace{\frac{1}{K}\exp\left(\frac{s_i}{\beta(R)}\right)}_{\pi_i(\boldsymbol{s};R)}\frac{s_i}{\beta(R)}\right]$$

Next, observe that by definition $\log(\pi_i) = \frac{s_i}{\beta(R)} - \log(K)$. Thus, by substituting in this relationship to the expression above, we have:

$$\begin{aligned}\frac{\partial C}{\partial R} &= \frac{\partial \beta}{\partial R}\left[\log(K) - \sum_i \pi_i(\boldsymbol{s};R)(\log \pi_i(\boldsymbol{s};R) + \log(K))\right] \\ &= \frac{\partial \beta}{\partial R}\left[\log(K) + H(\pi(\boldsymbol{s};R)) - \log(K)\sum_i \pi_i(\boldsymbol{s};R)\right] = \frac{\partial \beta}{\partial R}H(\pi(\boldsymbol{s};R))\end{aligned}$$

where $H(p) = -\sum_i p_i \log p_i$ is the Shannon entropy. Since entropy is always positive, $\frac{\partial C}{\partial R}$ has the same sign as $\frac{\partial \beta}{\partial R}$, which is decreasing by definition. □

Beneath the equivalence of the two market forms lies a subtle asymmetry. For a risk-neutral trader the two representations coincide exactly, just as the classical correspondence between scoring rules and market makers predicts. The risk-averse, evidence-only trade is where they diverge. In the scoring-rule view, such a trader reports an unchanged belief, $q_t = q_{t-1}$, along with evidence E_t; the market belief is by definition left untouched; only the liquidity, and hence the payoff, responds to the new evidence. The corresponding risk-averse trade AMM trade is to buy nothing – $\boldsymbol{z} = 0$ and submit evidence E_t. This does, however, move the market belief since $\pi(\boldsymbol{s}; R)$ depends on the evidence record R as well as on the shares $\boldsymbol{s}$. So fresh evidence reprices the book even when no shares change hands. The same evidence is therefore belief-neutral in one representation but belief-moving in the other. In short, evidence-only trades differ in both their payoff and their effect on the market belief between the modified LMSR and AMM implementations of our market[11].

6 Practical Considerations

6.1 Evidence Verification

Our work so far considers one axis of strategic reasoning about evidence: withholding. However, traders can also (knowingly or not) submit dubious evidence. Verification is thus a first-class concern for any practical instantiation of an evidence market. Indeed, the entire incentive structure derived in Sections 4 and 5 rests on the assumption that the quality function $r(\cdot)$ faithfully reflects the informativeness of submitted evidence. If a trader can affect endogenous resolution or drive up R_t and harvest the entropy payoff by submitting irrelevant, fabricated, or duplicate content, then the platform loses out in both money and poor information aggregation. The health of the market thus rests on being able to filter out such submissions. The quality function r introduced in Definition 3.2 is where we model this filtering taking place. Conceptually, r should assign zero marginal contribution to evidence that is (i) *irrelevant* to the event or the alternatives being considered, (ii) *wrong* or factually/objectively incorrect, or (iii) *a semantic duplicate* of evidence already

[11] In both cases, however, the evidence only payoff is proportional to the current entropy of beliefs.

collected in E_{t-1}^{total}. The monotonicity assumption $r(E_1) \leq r(E_2)$ for $E_1 \subseteq E_2$ then ensures that such evidence has no effect on the market[12].

In any case, the choice of r depends significantly on whether resolution is exogenous or endogenous. In the exogenous case, evidence only affects payoff magnitudes and explanation informativeness; the resolution itself is determined externally. Verification thus needs to filter out content that would waste platform funds or pollute the aggregation, but the consequences of a verification error are limited to just this. In the endogenous case, by contrast, evidence directly determines resolution, and a trader injecting fabricated evidence can shift resolution probabilities in their favor, thereby increasing payoff in both the evidence and belief axes. The verification machinery must accordingly be more robust.

A first instinct in both cases may be to turn to the rich literature on peer prediction [16]. These are mechanisms that truthfully elicit opinions from agents about unverifiable signals (absence of ground truth) by rewarding them on how their reports relate to those of their peers. Unfortunately, the incentive guarantees here typically require peers to have no external incentives, which is generally not possible in our setting since a peer could also be a trader. If they have an exiting trade in the market and resolution is endogenous, they are incentivized to only approve signals that confirm their trade. If they may trade in the future, they are incentivized to not approve any evidence (in either endogenous or exogenous resolution) since it would eat up liquidity (i.e. decrease $\frac{\partial \beta}{\partial R}$) that they would otherwise access. Freeman et al. [6] show that ensuring incentive compatibility with general external factors essentially requires increasing verification payments and transaction costs. Given the external incentive in the endogenous case is prediction market payoff, this could be quite large and this approach would be generally unpalatable. As such, we consider verification through the LLM-as-a-Judge paradigm, with different implementations for the endogenous and exogenous resolution cases. In both cases, the LLM inference cost for verification is expected to come from transaction fees.

Endogenous Resolution. We start by considering what an appropriate quality function $r(\cdot)$ looks like in this regime by using the LLM-evaluation example. Recall that in the endogenous setting each atomic evidence e has a binary outcome $X_{ie} \in \{0, 1\}$ with respect to each alternative i; for LLM evaluation $X_{ie} = 1$ corresponds to model i answering question e correctly. One natural per-evidence quality metric is the *discriminative score*: the ratio of alternatives that fail on e to those that succeed on e. Evidence on which no alternative succeeds receives zero credit.

Definition 6.1 (Discriminative Score)**.** *For an atomic evidence e with binary outcomes $\{X_{ie}\}_{i=1}^{n}$, let $c(e) := \sum_i X_{ie}$ count the alternatives supporting e and $w(e) := n - c(e)$ count those refuting it. The per-evidence discriminative score is*

$$r_{\text{disc}}(e) \;:=\; \begin{cases} w(e)/c(e), & c(e) \geq 1, \\ 0, & c(e) = 0, \end{cases} \quad \textit{and} \quad r(E) := \sum_{e \in E} r_{\text{disc}}(e) \tag{18}$$

r_{disc} is non-negative and (additively) monotone by construction, satisfying Definition 3.2. The per-evidence score has a clean interpretation; it is maximized precisely when exactly one alternative is corroborated by e and all others are negated. This nicely aligns with the ideal discriminating question in an LLM-evaluation context where the goal is to find questions where a few models answer correct and others answer wrong. This score drops to zero when the evidence is non-discriminating – either all alternatives are corroborated or negated by e.

Given this concrete quality function, the verification machinery's task reduces to a single binary decision per submission: should a submitted piece of evidence e be admitted to E^{total}, with the outcomes $\{X_{ie}\}_i$ then computed and used for resolution. As discussed, verification must reject e if it is irrelevant, incorrect, or a duplicate. Given the aforementioned challenges of peer prediction, we propose an *LLM verification with staked disputes* mechanism, drawing on the proof-of-stake paradigm in distributed systems [2] and the optimistic-rollup design pattern from Ethereum scaling [9]. When evidence e is submitted, a frontier

[12]The current model does not penalize such dubious submission - it just gives zero benefit for doing so. In practice, it may be fruitful to consider additional external incentives to discourage such behavior. Staking to submit evidence may be one such approach.

LLM judge — equipped with broad world knowledge and existing market context such as E_{t-1}^{total} — issues a tentative accept/reject decision along with a rationale. A bounded dispute window then opens, during which any market participant may challenge the decision by staking capital. The right to dispute is allocated via a second-price auction over stakes, with the winning challenger required to submit a written counter-argument. The judge re-adjudicates with the full history in context (the original evidence, its initial verdict and rationale, and the disputer's counter-argument). If the re-adjudication overturns the initial decision, the disputer's stake is returned and a small reward is paid out from the platform's transaction-fee pool; otherwise the stake is slashed by the platform.

This design inherits the core security argument from proof-of-stake. Economic skin in the game replaces any assumption that participants are intrinsically honest, and slashing penalties incentivizes honest disputes. Crucially, the mechanism sidesteps the peer-prediction objection raised earlier due to the LLM judge acting as the gate-keeper. Peers are not asked to truthfully report a signal in the presence of side-channel incentives; they are instead asked to stake capital to override an LLM decision they believe to be wrong. A trader with private information that the LLM erred has positive expected value in disputing; a trader without such information faces zero expected gain and will not dispute frivolously. Staking thus provides a costly-but-available channel for traders to inject information when they deem the verifier to have erred. The frontier LLM on the other hand serves as the arbiter given multiple pieces of evidence, as opposed to a one shot oracle.

From a practical perspective, the dispute window must be bounded so verification eventually terminates. Upper and lower bounds need to be enforced for staking to deter spurious or capital rich adversaries from dominating disputes. The inference cost of running the judge must be funded by the platform through transaction fees, with slashed dispute stakes providing additional revenue. Lastly, we cannot claim the mechanism is bulletproof since the LLM judge remains a centralized point of failure. However, we believe the increasing ability of LLMs will only minimize such errors.

Exogenous Resolution. When resolution is determined by an external time-bound event, evidence plays no role in adjudicating the outcome. It only enriches the aggregated information product and scales the liquidity parameter. Verification is correspondingly less critical and we simply need to design r to ensure the platform is not paying out for and aggregating poor quality information. For this we propose a lightweight scheme inspired by the LLM-judge construction of Srinivasan et al. [20].

Concretely, consider a frontier LLM judge J equipped with broad world knowledge and the cumulative evidence record E_{t-1}^{total}. We elicit from J a belief distribution $p_J(\,\cdot\,; E_{t-1}^{total})$ over the n alternatives. When a new submission E_t arrives, we append it to the judge's context and re-elicit a posterior belief $p_J(\,\cdot\,; E_{t-1}^{total} \cup E_t)$. The marginal quality contribution of the new evidence is then defined as the divergence between these two distributions:

$$r(E_{t-1}^{total} \cup E_t) - r(E_{t-1}^{total}) := D_{KL}\big(p_J(\cdot\,; E_{t-1}^{total} \cup E_t) \,\|\, p_J(\cdot\,; E_{t-1}^{total})\big), \tag{19}$$

with R_t being the running sum of such contributions and $r(\emptyset) = 0$. This definition satisfies the non-negativity and monotonicity requirements of Definition 3.2 by construction, since KL divergence is non-negative. Indeed, KL-divergence in (19) can be replaced by any other f-divergence (total variation, Jensen-Shannon) without affecting anything qualitatively.

The key property of this construction is that a competent judge should be *unmoved* by evidence lacking genuine informational content. If E_t is irrelevant to the event in question or is already contained in E_{t-1}^{total}, the judge has no reason to revise its belief; if E_t is factually wrong, the judge's world knowledge should lead it to discount the submission. In these failure modes, the divergence in (19) collapses to zero (or near-zero), and the trader receives no evidence payoff. Conversely, evidence that is novel, correct, and pertinent will move the judge's belief and accrue credit proportional to its informational content. This aligns the incentives of evidence submission with the platform's epistemic goals: traders are paid for information that genuinely changes a well-informed observer's view of the event. The overall scheme inherits the verification asymmetry discussed in Section 2 — the judge need only *evaluate* a submitted piece of evidence rather than generate it, which is structurally easier and well-supported by recent work on LLM verifiers [14, 25, 10]. As before, while the LLM-judge does remain a central point of failure, it becomes less so as their power and capabilities improve.

6.2 Execution Algorithms

The AMM mechanism developed in Section 5 implicitly assumes that trades arrive and are executed sequentially: each order's evidence is fully verified, $r(E_t)$ computed, the cumulative quality R_t updated, the order executed before the next trader interacts with the market. In practice this is unworkable. Verification, particularly through an LLM judge of the kind described just earlier, takes meaningful time, and a market that pauses for each verification will simply be too slow for most traders. Concretely, consider a trader t who submits an order $(\boldsymbol{z}_t, E_t)$ and is followed seconds later by a trader $t+1$ with $E_{t+1} = \emptyset$. Verification of E_t may not have completed by the time $t+1$ arrives. What price should the market quote $t+1$? Should they have to wait until E_t is verified to trade?

We address this through an asynchronous execution protocol consisting of two concurrent algorithms: a *trade execution* routine that processes incoming orders immediately at a conservative price, and a *verification adjustment* routine that consumes the evidence queue in arrival order and refunds the difference between the conservative charge and the true synchronous cost. The protocol is designed to satisfy four desiderata:

- **Equivalence to synchronous execution.** Each trader's final net payment matches what they would have paid under sequential synchronous execution, preserving all incentive guarantees of Section 5.
- **Immediate execution for evidence-free orders.** A trader submitting $E_t = \emptyset$ never waits on any verification queue.
- **At-most-one charge per order.** Traders can only be charged (transfer of money from agent to platform) at-most once during the execution of a single order. Any subsequent transfer is a rebate from platform to trader.
- **Lower-bounded refund at execution.** The platform can quote the trader a guaranteed minimum refund at the time the order is executed.

The central design choice is to execute each order *pessimistically*, i.e. charge the maximum cost the trader could face once the unverified evidence resolves and then refund the difference once verification completes. Whether the pessimistic regime corresponds to higher or lower β depends on the order, since traders do not uniformly prefer high liquidity. The following lemma characterizes this dependence and gives us a computable worst-case β within any liquidity interval.

Lemma 6.1. *For fixed cumulative shares $\boldsymbol{s}$, share order $\boldsymbol{z}$, and liquidity β, let $\xi(\boldsymbol{z}; \boldsymbol{s}, \beta) := C(\boldsymbol{s}+\boldsymbol{z}; \beta) - C(\boldsymbol{s}; \beta)$ denote the execution cost of the shares. Then*

$$\frac{\partial \xi}{\partial \beta} \;=\; H\big(\pi(\boldsymbol{s}+\boldsymbol{z}; \beta)\big) - H\big(\pi(\boldsymbol{s}; \beta)\big), \tag{20}$$

So ξ is increasing in β when the trade raises market entropy and decreasing when it lowers it.

Proof. Since $C(\boldsymbol{s}; \beta) = \beta \log K$ with $K = \sum_i \exp(s_i/\beta)$, the product rule gives

$$\frac{\partial C}{\partial \beta} \;=\; \log K + \frac{\beta}{K} \cdot \frac{\partial K}{\partial \beta}, \qquad \frac{\partial K}{\partial \beta} \;=\; -\frac{1}{\beta^2} \sum_i s_i \exp\!\left(\frac{s_i}{\beta}\right).$$

Combining and recognizing $\pi_i(\boldsymbol{s}; \beta) = \exp(s_i/\beta)/K$ yields $\partial C/\partial\beta = \log K - \sum_i \pi_i (s_i/\beta)$. Substituting $s_i/\beta = \log \pi_i + \log K$ and using $\sum_i \pi_i = 1$ simplifies this to $\partial C/\partial \beta = H(\pi(\boldsymbol{s}; \beta))$. Equation (20) follows immediately. □

We can now state the two algorithms. The trade execution algorithm accepts each incoming order $(\boldsymbol{z}_t, E_t)$ at time t, appends E_t to the unverified evidence queue Q, and charges trader t the worst-case share execution cost over the possible range of cumulative quality the addition of this evidence could lead to. Let R denote the verified cumulative quality computed so far, and let r_Q denote an upper bound on the total quality contribution of evidence currently in Q.[13] When trader t's trade raises market entropy – so $\partial\xi/\partial\beta > 0$ by Lemma 6.1 – the trader prefers *low* β. So the worst case for them is the *highest* feasible β, achieved at $R_t^{\mathrm{ex}} = R$ (i.e. assume the queue contributes nothing). When the trade lowers market entropy, the worst case is instead $R_t^{\mathrm{ex}} = R + r_Q$ (assume the queue contributes its maximum). A market snapshot at the time of execution is stored for use by the verification adjustment routine.

Algorithm 1: Trade Execution

Input: Verified state so far $(\boldsymbol{s}, R)$; queue Q with quality upper bound r_Q; arriving order stream
`snapshots` $\leftarrow \{\}$; `// trader index` $\rightarrow$ `execution-time state`
while *the order stream is not empty* **do**
  pull earliest order $t : (\boldsymbol{z}_t, E_t)$
  if $H(\pi(\boldsymbol{s}+\boldsymbol{z}_t;\beta(R))) - H(\pi(\boldsymbol{s};\beta(R))) \geq 0$ **then**
    $R_t^{\mathrm{ex}} \leftarrow R$; `// trade raises entropy; pessimistic: highest` β
  else
    $R_t^{\mathrm{ex}} \leftarrow R + r_Q$; `// trade lowers entropy; pessimistic: lowest` β
  `snapshots`$[t] \leftarrow \big(\boldsymbol{s}, \boldsymbol{z}_t, R_t^{\mathrm{ex}}, \text{copy of } Q\big)$
  append (t, E_t) to Q
  charge trader t: $\xi_t^{\mathrm{ex}} := C(\boldsymbol{s}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{ex}})) - C(\boldsymbol{s};\beta(R_t^{\mathrm{ex}}))$
  update $\boldsymbol{s} \leftarrow \boldsymbol{s} + \boldsymbol{z}_t$

The verification adjustment algorithm pulls evidence from Q in arrival order, verifies it to compute $r(E_{t-1}^{total} \cup E_t)$, and refunds trader t the difference between ξ_t^{ex} and the synchronous cost. This refund decomposes naturally into two non-negative components: a *transaction refund* reflecting that R_t^{ex} executed the shares at the trader's worst-case β, and an *evidence refund* reflecting the quality contribution $r(E_t)$ of the trader's own evidence.

Algorithm 2: Verification Adjustment

Input: Queue Q from Algorithm 1; snapshots dictionary `snapshots`
`verified` $\leftarrow \{\}$; `// trader index` $\rightarrow r(E_t)$
while *Q is not empty* **do**
  pop earliest (t, E_t) from Q; verify and set `verified`$[t] \leftarrow r(E_t)$
  let $\big(\boldsymbol{s}^{\mathrm{ex}}, \boldsymbol{z}_t, R_t^{\mathrm{ex}}, Q_t\big) \leftarrow$ `snapshots`$[t]$
  $R_t^{\mathrm{true}} \leftarrow \sum_{(t', E_{t'}) \in Q_t}$ `verified`$[t']$; `// queue entries preceding` t `are now verified`
  `trans_refund`
   $\leftarrow \big[C(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{ex}})) - C(\boldsymbol{s}^{\mathrm{ex}};\beta(R_t^{\mathrm{ex}}))\big] - \big[C(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{true}})) - C(\boldsymbol{s}^{\mathrm{ex}};\beta(R_t^{\mathrm{true}}))\big]$
  `evidence_refund` $\leftarrow C(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{true}})) - C(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{true}}+r(E_t)))$
  transfer `trans_refund` + `evidence_refund` to trader t

Theorem 6.1. *Under Algorithms 1 and 2, the final net payment from each trader t matches the synchronous sequential execution cost. Each trader is charged once at execution; all subsequent transfers are non-negative refunds. Furthermore, the platform can quote each trader a deterministic lower bound on their total refund at the time of execution.*

Proof. Fix trader t with snapshot $(\boldsymbol{s}^{\mathrm{ex}}, \boldsymbol{z}_t, R_t^{\mathrm{ex}}, Q_t)$, true cumulative quality R_t^{true} at the time of their arrival, and own evidence quality $r_t := r(E_t)$. Their synchronous cost is

$$\xi_t^{\mathrm{sync}} \;:=\; C\big(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{true}}+r_t)\big) - C\big(\boldsymbol{s}^{\mathrm{ex}};\beta(R_t^{\mathrm{true}})\big)\,.$$

Adding and subtracting $C(\boldsymbol{s}^{\mathrm{ex}}+\boldsymbol{z}_t;\beta(R_t^{\mathrm{true}})) - C(\boldsymbol{s}^{\mathrm{ex}};\beta(R_t^{\mathrm{true}}))$, we obtain the identity

$$\xi_t^{\mathrm{sync}} \;=\; \xi_t^{\mathrm{ex}} \;-\; \texttt{trans\_refund} \;-\; \texttt{evidence\_refund}, \tag{21}$$

where the two refund terms are as defined in Algorithm 2. Hence the trader's net payment after refund equals ξ_t^{sync}.

[13]In practice, r_Q can be set by the platform as a sum of per-evidence quality ceilings — e.g. by capping the divergence in equation (19) at some platform-chosen value.

Both refund components are non-negative. For `trans_refund`, Lemma 6.1 together with the pessimistic choice of R_t^{ex} implies that $\xi(\boldsymbol{s}^{\text{ex}}, \boldsymbol{z}_t, \beta(R_t^{\text{ex}})) \geq \xi(\boldsymbol{s}^{\text{ex}}, \boldsymbol{z}_t, \beta(\rho))$ for every ρ in the plausible range $[R, R + r_Q]$, and in particular for $\rho = R_t^{\text{true}}$. For `evidence_refund`, Theorem 5.3 establishes $\partial C/\partial R = (\partial\beta/\partial R)H(\pi) \leq 0$, so C is non-increasing in R and the difference is non-negative. The trader thus faces a single charge at execution followed only by transfers in their favor.

Finally, the platform can lower bound the refund at execution time. Since `trans_refund` ≥ 0, the entire refund is bounded below by `evidence_refund`, which in turn can be expressed as the line integral $\int_{R_t^{\text{true}}}^{R_t^{\text{true}}+r_t} |\partial\beta/\partial R| \cdot H(\pi(\boldsymbol{s}^{\text{ex}} + \boldsymbol{z}_t; \beta(R')))dR'$. Bounding the integrand below over the plausible range $R' \in [R, R + r_Q + r_t^{\max}]$ — with $r_t^{\max}$ a platform-set ceiling on per-order quality contribution — and assuming a platform-set minimum quality $r_t^{\min}$ for evidence to be deemed valid, gives the executable lower bound

$$\texttt{evidence\_refund} \;\geq\; r_t^{\min} \cdot \min_{R' \in [R,\, R + r_Q + r_t^{\max}]} \left\{ \left| \frac{\partial\beta}{\partial R} \right|_{R'} \cdot H(\pi(\boldsymbol{s}^{\text{ex}} + \boldsymbol{z}_t; \beta(R'))) \right\},$$

which is computable from quantities known at execution time and given to the trader as a conditional guarantee (i.e. contingent on their evidence passing verification with quality at least $r_t^{\min}$). □

7 Discussion

This work introduced evidence markets, a generalization of prediction markets in which traders can convey beliefs alongside evidence, and which resolves either exogenously through a realized event or endogenously through the submitted evidence itself. This addresses the two limitations we highlighted: the opacity of the reasoning behind a market price, and the inapplicability of prediction markets to questions that admit no time-bound external resolution. Concretely, the mechanism extends the logarithmic market scoring rule with a liquidity parameter that decreases as cumulative evidence quality grows; under it, reporting both one's belief truthfully and one's evidence wholly is a dominant strategy equilibrium when the market resolves exogenously and a ε-subgame perfect Nash Equilibrium when it resolves endogenously. The same mechanism can equivalently be cast as an automated market maker. In both views, the trader's payoff can be cleanly split to capture the value of evidence so that even a risk-averse trader who takes no directional position earns a non-negative reward simply for contributing good evidence. Although LLM evaluation is our running example, the design fits any setting where the question is comparative and its answer must be assembled from a body of evidence rather than read off from nature; examples include the replicability of a scientific finding, the effectiveness of a policy, or the relative quality of competing artifacts.

Several questions remain. The most pressing being a deeper treatment of verification; our LLM-gated mechanism with staked disputes is a constructive starting point, but hardening it to become production-ready for an adversarial, real-money market is a substantial design and engineering problem. A second concerns the acquisition cost of evidence. We treat evidence as free to acquire, whereas producing a genuinely informative piece is often effortful, and modeling that cost and the resulting tension between reporting one's belief and actively gathering evidence for it would bring the framework closer to the reality we have in mind. A third concerns the trading interface. Platforms such as Kalshi and Polymarket run on central limit order books rather than cost-function based AMMs; how best to implement evidence markets in that setting with similar guarantees remains open.

Acknowledgments

We thank Sid Srinivasan, Bo Waggoner, and Raf Frongillo for their insightful comments and the many helpful discussions we have had.

References

[1] Kenneth J. Arrow, Robert Forsythe, Michael Gorham, Robert Hahn, Robin Hanson, John O. Ledyard, Saul Levmore, Robert Litan, Paul Milgrom, Forrest D. Nelson, et al. The promise of prediction markets. *Science*, 320(5878):877–878, 2008.

[2] Vitalik Buterin and Virgil Griffith. Casper the friendly finality gadget. *arXiv preprint arXiv:1710.09437*, 2017.

[3] Yiling Chen and David M Pennock. A utility framework for bounded-loss market makers. In *23rd Conference on Uncertainty in Artificial Intelligence, UAI 2007*, pages 49–56, 2007.

[4] Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios N. Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. Chatbot arena: An open platform for evaluating LLMs by human preference. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, 2024.

[5] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd International Conference on World Wide Web (WWW)*, pages 319–330, 2013.

[6] Rupert Freeman, Sébastien Lahaie, and David Pennock. Crowdsourced outcome determination in prediction markets. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, 2017.

[7] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

[8] Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.

[9] Harry Kalodner, Steven Goldfeder, Xiaoqi Chen, S Matthew Weinberg, and Edward W Felten. Arbitrum: Scalable, private smart contracts. In *27th USENIX Security Symposium*, pages 1353–1370, 2018.

[10] Muhammad Khalifa et al. ThinkPRM: Process reward models that think. *arXiv preprint arXiv:2504.16828*, 2025.

[11] Yuqing Kong. Dominantly truthful multi-task peer prediction with a constant number of tasks. pages 2398–2411, 2020.

[12] Yuqing Kong and Grant Schoenebeck. An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. In *ACM Transactions on Economics and Computation*, volume 7, pages 1–33, 2019.

[13] Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, et al. Holistic evaluation of language models. *arXiv preprint arXiv:2211.09110*, 2022.

[14] Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. Let's verify step by step. In *International Conference on Learning Representations (ICLR)*, 2024. Preprint: arXiv:2305.20050.

[15] Charles F. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429, 2006.

[16] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.

[17] Dražen Prelec. A Bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.

[18] Grant Schoenebeck and Fang-Yi Yu. Learning and strongly truthful multi-task peer prediction: A variational approach. In *Proceedings of the 12th Innovations in Theoretical Computer Science Conference (ITCS)*, 2021.

[19] Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C. Parkes. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 179–196, 2016.

[20] Siddarth Srinivasan, Ezra Karger, Michiel Bakker, and Yiling Chen. Tell me why: Incentivizing explanations. *arXiv preprint arXiv:2502.13410*, 2025.

[21] Siddarth Srinivasan, Ezra Karger, and Yiling Chen. Self-resolving prediction markets for unverifiable outcomes. In *Proceedings of the 26th ACM Conference on Economics and Computation (EC)*, pages 547–573, 2025. Preprint: arXiv:2306.04305.

[22] Siddarth Srinivasan, Tao Lin, Connacher Murphy, Anish Thilagar, Yiling Chen, and Ezra Karger. WOMAC: A mechanism for prediction competitions. *arXiv preprint arXiv:2508.17907*, 2025.

[23] Aarohi Srivastava et al. Beyond the imitation game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research*, 2023.

[24] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.

[25] Lunjun Zhang, Arian Hosseini, Hritik Bansal, Mehran Kazemi, Aviral Kumar, and Rishabh Agarwal. Generative verifiers: Reward modeling as next-token prediction. *arXiv preprint arXiv:2408.15240*, 2024.

[26] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging LLM-as-a-judge with MT-Bench and Chatbot Arena. In *Advances in Neural Information Processing Systems (NeurIPS), Datasets and Benchmarks Track*, 2023.