

Trust as a thermodynamic quantity

Contents

Mathematical Framework for Agent Value Estimation in Prediction Markets

Abstract. We present a mechanism-agnostic framework for quantifying agent value in prediction markets based on Kullback-Leibler divergence reduction, adjusted for prediction timeliness and market-specific cost structures. The system value function $V_{\text{system}}(a)$ aggregates agent contributions across markets, incorporating temporal decay and mechanism-dependent normalization factors.

1. Preliminaries

Let $\mathcal{M} = \{1, \dots, M\}$ denote a set of markets and $\mathcal{K} = \{1, \dots, K\}$ a set of agents. For market $m \in \mathcal{M}$, let \mathcal{L}_m be the outcome space and P_m^{final} the true distribution revealed at time $t_1^{(m)}$.

For each market m , let $Q_m = (q_{m,1}, \dots, q_{m,T_m})$ be the sequence of bets where $q_{m,t} \in \Delta(\mathcal{L}_m)$, with $\phi_m : \{1, \dots, T_m\} \rightarrow \mathcal{K}$ mapping time indices to agents.

2. Core Value Functions

Definition 2.1 (Base Value). The information value of bet t in market m is the reduction in Kullback-Leibler divergence from the final state:

$$v_{m,t} = D_{\text{KL}}(P_m^{\text{final}} \| P_{m,t-1}) - D_{\text{KL}}(P_m^{\text{final}} \| P_{m,t})$$

where $D_{\text{KL}}(P \| Q) = \sum_{i \in \mathcal{L}_m} P(i) \log \frac{P(i)}{Q(i)}$.

Definition 2.2 (Temporal Decay). The timeliness of prediction t is encoded via decay function $\delta : [0, t_1] \rightarrow \mathbb{R}_+$:

$$\delta(t) = \exp(-\lambda(t_1 - t)), \quad \lambda > 0$$

or alternatively $\delta(t) = 1 - t/t_1$ for linear decay.

Definition 2.3 (Normalized Value). Accounting for market frictions, the adjusted value is:

$$v_{m,t}^{\text{adj}} = \delta(t) \cdot \frac{v_{m,t}}{c_{m,t}}$$

where $c_{m,t}$ represents the mechanism-specific cost of information at time t .

3. Mechanism-Specific Cost Functions

Bayesian Markets. State represented by pseudo-counts $\alpha_t \in \mathbb{R}_+^{|\mathcal{L}_m|}$ with update $\alpha_t = \alpha_{t-1} + w \cdot q_t$. The cost of information is the prior strength:

$$c_{m,t}^{\text{Bayes}} = \|\alpha_{t-1}\|_1 = \sum_{j \in \mathcal{L}_m} \alpha_{t-1}(j)$$

LMSR Markets. State represented by log-quantities $q_t \in \mathbb{R}^{|\mathcal{L}_m|}$ with cost function $C(q) = b \log \sum_i \exp(q_i/b)$. The cost is:

$$c_{m,t}^{\text{LMSR}} = \nabla C(q_{t-1}) \cdot b_t = \sum_{i \in \mathcal{L}_m} P_{m,t-1}(i) \cdot b_t(i)$$

where b_t is the share vector purchased.

4. System Value Aggregation

Definition 4.1 (Agent System Value). The total value of agent $a \in \mathcal{K}$ is:

$$V_{\text{system}}(a) = \sum_{m \in \mathcal{M}} \sum_{\substack{t=1 \\ \phi_m(t)=a}}^{T_m} \delta(t) \cdot \frac{D_{\text{KL}}(P_m^{\text{final}} \| P_{m,t-1}) - D_{\text{KL}}(P_m^{\text{final}} \| P_{m,t})}{c_{m,t}}$$

Proposition 4.2 (Additivity). System value is additive across agents:

$$V_{\text{system}}(\mathcal{K}) = \sum_{a \in \mathcal{K}} V_{\text{system}}(a)$$

Proposition 4.3 (Counterfactual Equivalence). $V_{\text{system}}(a)$ equals the expected KL divergence between the final outcome and the counterfactual market state without agent a 's participation.

5. Incentive Alignment

Let $R(a)$ denote the reward function under mechanism \mathcal{M} . The **alignment gap** is defined as:

$$\Gamma(a) = V_{\text{system}}(a) - \mathbb{E}[R(a)]$$

Bayesian Markets. With scoring rule $S(P, \omega)$ where ω is the realized outcome:

$$R_{\text{Bayes}}(a) = \sum_{m,t:\phi_m(t)=a} S(P_{m,t}, \omega_m) - S(P_{m,t-1}, \omega_m)$$

The alignment gap depends on prior strength $\|\alpha_0\|$ and decay parameter λ .

LMSR Markets. With profit function π :

$$R_{\text{LMSR}}(a) = \sum_{m,t:\phi_m(t)=a} \pi_{m,t}$$

The alignment gap depends on liquidity parameter b and timing t .

6. Optimization Framework

The system designer solves:

$$\max_{\lambda, b, \alpha_0} \sum_{a \in \mathcal{K}} V_{\text{system}}(a; \lambda, b, \alpha_0)$$

subject to participation constraints $\mathbb{E}[R(a)] \geq u_a^{\text{outside}}$ for all $a \in \mathcal{K}$.

The Lagrangian yields optimal decay rate:

$$\lambda^* = \frac{\sum_a \frac{\partial V_{\text{system}}(a)}{\partial \lambda}}{\sum_a \frac{\partial \mathbb{E}[R(a)]}{\partial \lambda}}$$

7. Conclusion

The framework provides a mechanism-agnostic metric $V_{\text{system}}(a)$ for agent evaluation, incorporating temporal discounting $\delta(t)$ and mechanism-specific costs $c_{m,t}$. Future work concerns dynamic optimization of (λ, b) to minimize the alignment gap $\Gamma(a)$ across heterogeneous agent populations.