Strategic Trading with Binary Payoffs: The Kyle Model for Prediction Markets

March 21, 2026

Abstract

This paper solves the Kyle (1985) insider trading model when the terminal asset value is binary. In the one-round model, the equilibrium pricing rule is logistic, and price impact is state-dependent: $\lambda(p) = \lambda_z \cdot p(1-p)$. A Girsanov symmetry argument yields the identity $\mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$, which reduces the equilibrium equation to $G(k) = \mathbb{E}[S](1 - k^2/4) = 0$. The unique symmetric equilibrium has order intensity $k^* = 2$ exactly: the informed trader's optimal order equals the noise trader's standard deviation ($a^* = \sigma_u$). In the multi-round model, a single informed trader back-loads trading toward resolution, unlike the constant trading rate in the Gaussian benchmark. With multiple insiders, competition reverses the dynamic pattern: front-loading dominates as the information-stealing motive overwhelms the positive externality of waiting. The continuous-time limit produces a Wright-Fisher diffusion whose volatility function must depend on the current price (Proposition 20), in contrast to the state-independent strategy spread in Back (1992). The equilibrium volatility blows up at rate $(T-t)^{-1/2}$ with a price-dependent amplitude $A(p)$ satisfying a nonlinear ODE with boundary conditions $A(0) = A(1) = 1$ and maximum $A(1/2) \approx 1.81$. The ODE is solved numerically, giving the full equilibrium characterization near resolution. With multiple insiders, competition dissipates rents and accelerates information aggregation while preserving the logistic pricing rule and the state-dependent price impact structure; the binary model introduces a positive externality of waiting absent from the Gaussian benchmark. Kalshi's $p(1-p)$ fee structure disproportionately reduces informed trading at intermediate prices, where information has the greatest value.

1 Introduction

How do prediction markets aggregate private information when the underlying contract pays zero or one? Platforms such as Kalshi and Polymarket now host billions of dollars in trading volume on binary event contracts, from election outcomes to macroeconomic indicators. A growing empirical literature documents the microstructure of these markets: state-dependent adverse selection [Palumbo, 2025], systematic wealth transfer from takers to makers [Becker, 2026], a pronounced favorite-longshot bias [Bürgi et al., 2025], and declining price impact as markets mature [Tsang and Yang, 2026]. Yet no equilibrium model connects an informed trader's strategic behavior to the observed price dynamics when the terminal payoff is binary. The canonical framework for strategic informed trading, Kyle [1985], takes payoffs to be Gaussian. The continuous-time extension of Back [1992] flags the binary (Bernoulli) case as requiring special treatment without solving it.

The equilibrium pricing rule is logistic, price impact is state-dependent, and the informed trader back-loads trading toward resolution. All three features are absent from the Gaussian benchmark.

Binary payoffs force the equilibrium to be state-dependent in a way that Gaussian payoffs do not. In the standard Kyle-Back model, the informed trader's strategy spread (the difference in trading rates between the two states of the world) depends only on time: $\Delta\theta(t) = 1/(T-t)$. Prices follow a Brownian bridge to the terminal value. No such state-independent strategy exists in the binary model (Proposition 20). The proof exploits the incompatibility between the value function implied by a deterministic strategy spread and the HJB equation: the former requires W_t to be a function of $\log p$, while the latter requires it to be a function of $p(1-p)$. These are different functions of p, so the HJB equation has no solution with a state-independent strategy. The binary equilibrium must have a strategy spread $\alpha(p,t)$ that depends on the current price, producing a Wright-Fisher diffusion rather than a Brownian bridge.

The one-round model yields a complete characterization. The market maker's pricing rule is logistic (Proposition 2), and price impact is state-dependent: $\lambda(p) = \lambda_z \cdot p(1-p)$. Impact peaks at $p = 1/2$ and vanishes at the boundaries. The unique symmetric equilibrium (Propositions 3–6) has order intensity $k^* = 2$, a universal constant that depends only on the logistic and standard normal functions and not on noise trader volatility. The informed trader orders exactly σ_u (matching the noise trader's standard deviation), earns $0.219\sigma_u$ in expected profit (12.4% less than the Gaussian benchmark), and faces higher peak price impact but lower impact at extreme prices.

A natural question is whether the $p(1-p)$ price impact pattern is specific to the Kyle framework or generic to binary Bayesian updating. Proposition 9 shows it is generic: the Glosten-Milgrom model with binary payoffs also produces $p(1-p)$ price impact. What the Kyle framework adds is endogenous order size optimization (producing $k^* = 2$), dynamic back-loading (which requires continuous order sizes), and the continuous-time Wright-Fisher characterization. The pricing rule itself is Bayes' rule applied to Gaussian signals about a Bernoulli state. The contribution of the equilibrium analysis is not the pricing rule but the strategic interaction: the informed trader's optimal order size, the comparative statics, and the dynamic properties that emerge from the interplay between binary updating and order size choice.

Extending the one-round model to N_{ins} symmetric insiders, competition dissipates rents and accelerates information aggregation while preserving the logistic pricing rule and the $p(1-p)$ price impact structure (Proposition 26). The binary model introduces a positive externality absent from the Gaussian benchmark: each insider's trading moves p toward v, reducing $p(1-p)$ and lowering the price impact that all insiders face in subsequent rounds.

The multi-round model reveals the dynamic consequences of state-dependent price impact. In the Gaussian Kyle model, trading intensity is constant across rounds: the informed trader is indifferent about timing. In the binary model, the informed trader back-loads trading toward resolution (Proposition 17). The mechanism is that the logistic pricing rule's sensitivity $p(1-p)$ declines as the price moves toward the truth, so trading later, after early rounds have moved the price partway, is cheaper per unit of information revealed. The proof reduces the back-loading condition to $V'(0) + V''(0) < 0$, where V is the continuation value function, using an envelope theorem and Stein's lemma argument. An analytical proof using the Girsanov identity confirms the condition holds.

The continuous-time limit produces a price process $dp_t = \alpha(p_t, t) \cdot p_t(1-p_t)\, dB_t$, a Wright-Fisher diffusion with state-dependent volatility. The equilibrium volatility function satisfies a nonlinear PDE (Proposition 22) with three qualitative properties: symmetry around $p = 1/2$ (Proposition 23), terminal blowup at rate $(T-t)^{-1/2}$ with a price-dependent amplitude governed by a nonlinear ODE (Proposition 24), and monotonicity on $(0, 1/2)$ (Proposition 25).

Related literature. This paper connects to several strands of work. In the strategic trading literature, Kyle [1985] and Back [1992] establish the Gaussian benchmark. Keller and Tseng [2025] extend Kyle to Arrow-Debreu securities, studying how private signals about state probabilities are aggregated across multiple securities. Their framework accommodates binary payoffs as a special case, but their analysis focuses on multi-asset price discovery, not on the equilibrium characterization of a single binary contract. The present paper complements their work by characterizing the single-contract equilibrium in detail: the exact constant $k^* = 2$, back-loading, the impossibility of state-independent strategies, and the terminal-layer ODE do not appear in their analysis. The binary case is noted as requiring special treatment in Back [1992] but has not been solved.

In prediction market theory, Wolfers and Zitzewitz [2004, 2006] survey accuracy and interpretation of prices. Manski [2004] raises the identification problem for translating prices to probabilities. Hanson [2003] introduces the logarithmic market scoring rule (LMSR), which also produces a logistic pricing structure, but from a mechanism design perspective rather than strategic equilibrium. The LMSR's logistic form arises from the scoring rule's cost function; in the Kyle equilibrium, the logistic form arises from Bayesian updating about a binary state given Gaussian signals.

The empirical prediction market literature motivates the model. Palumbo [2025] documents that prediction market making resembles underwriting rather than inventory-neutral intermediation, consistent with the terminal risk concentration the model predicts at $p = 1/2$. Becker [2026] finds a 1.12% average wealth transfer from takers to makers, with pronounced asymmetry at longshot prices. Bürgi et al. [2025] document a favorite-longshot bias and negative pre-fee returns averaging -20%. Tsang and Yang [2026] show that Kyle's λ declines by more than an order of magnitude as Polymarket's election market matures, consistent with the model's prediction of declining price impact as prices approach the truth. Ng et al. [2025] find that large-trade order imbalance predicts subsequent returns, consistent with informed strategic trading.

Outline. Section 2 presents the model. Section 3 derives the one-round equilibrium, extends to asymmetric priors, and compares with the Gaussian Kyle and Glosten-Milgrom benchmarks. Section 4 develops the multi-round and continuous-time models and extends the one-round equilibrium to multiple insiders. Section 5 discusses testable predictions, market design implications, and limitations. Section 6 concludes. Proofs that interrupt the flow appear in the Appendix.

2 Model

2.1 Environment

A binary event contract pays $v \in \{0, 1\}$ at a terminal date. The common prior probability of $v = 1$ is $p_0 \in (0, 1)$.

Three types of agents trade the contract.

1. **Informed trader.** The informed trader observes v before trading. The trader is risk-

neutral and chooses an order $x \in \mathbb{R}$ to maximize expected profit $\mathbb{E}[x(v-p_1) \mid v]$, where p_1 is the post-trade price.

2. **Noise traders.** Noise traders submit aggregate order flow $u \sim N(0, \sigma_u^2)$, independent of v.

3. **Market maker.** The market maker is risk-neutral and competitive. The market maker observes total order flow $Y = x + u$ and sets the price equal to the conditional expectation: $p_1 = \mathbb{E}[v \mid Y]$.

2.2 Strategies and equilibrium

The informed trader's strategy specifies an order for each realization of v: (x_1, x_0), where x_v is the order when the terminal value is v. The market maker's pricing rule maps total order flow to a posterior probability: $p_1 \colon \mathbb{R} \to (0, 1)$.

Definition 1 (Equilibrium)**.** A *one-round equilibrium* consists of a strategy (x_1^*, x_0^*) and a pricing rule $p_1^*(\cdot)$ such that:

(E1) *Bayesian pricing:* $p_1^*(Y) = \Pr(v = 1 \mid Y)$ given (x_1^*, x_0^*).

(E2) *Profit maximization:* for each $v \in \{0, 1\}$, x_v^* maximizes $\mathbb{E}[x(v - p_1^*(x + u)) \mid v]$.

(E3) *Consistency:* the order flow distribution used in (E1) is generated by (x_1^*, x_0^*) and u.

2.3 Notation

Let $\sigma(z) := 1/(1 + e^{-z})$ denote the logistic function, so that $p = \sigma(z)$ maps log-odds $z = \log(p/(1-p))$ to probabilities. Denote $z_0 = \log(p_0/(1 - p_0))$ for the prior log-odds. Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal density and distribution function.

2.4 Multi-round extension

In the N-round model, trading occurs in rounds $n = 0, \ldots, N-1$, after which v is revealed. Per-round noise is $u_n \sim N(0, \tilde{\sigma}^2)$ i.i.d. with $\tilde{\sigma} = \sigma_u/\sqrt{N}$, so total noise variance equals σ_u^2. In each round, the market maker observes $Y_n = x_n + u_n$ and updates the log-odds: $z_n = z_{n-1} + \lambda_z^{(n)} Y_n + c_n$. The informed trader maximizes total expected profit $\mathbb{E}[\sum_n x_n(v - p_n) \mid v]$.

The continuous-time model takes $N \to \infty$ with round duration $\Delta t = 1/N$ and per-round noise $\tilde{\sigma} = \sigma_u \sqrt{\Delta t}$. The order flow process converges to $dY_t = \theta_t^{(v)}\, dt + \sigma_u\, dW_t$, where $\theta_t^{(v)}$ is the informed trader's strategy in state v and W_t is standard Brownian motion.

3 One-Round Equilibrium

3.1 The logistic pricing rule

Proposition 2 (Logistic Pricing Rule)**.** *In any equilibrium with $x_1 \neq x_0$, the pricing rule is*

$$p_1(Y) = \sigma(z_0 + \lambda_z Y + c), \tag{1}$$

where $\lambda_z = (x_1 - x_0)/\sigma_u^2$ and $c = -(x_1^2 - x_0^2)/(2\sigma_u^2)$. Price impact is

$$\frac{dp_1}{dY} = \lambda_z \cdot p_1(1 - p_1). \tag{2}$$

Proof. The posterior log-odds follow from Bayes' rule. Under state v, the order flow is $Y \sim N(x_v, \sigma_u^2)$. The log-likelihood ratio is

$$\log \frac{f(Y \mid v = 1)}{f(Y \mid v = 0)} = \frac{(x_1 - x_0)}{\sigma_u^2} \left(Y - \frac{x_1 + x_0}{2} \right),$$

using $(Y - x_1)^2 - (Y - x_0)^2 = -2Y(x_1 - x_0) + x_1^2 - x_0^2$. The posterior log-odds are $z_1(Y) = z_0 + \lambda_z Y + c$, which gives $p_1(Y) = \sigma(z_1(Y))$. The price impact follows from $dp_1/dY = \sigma'(z_1) \cdot \lambda_z = p_1(1 - p_1)\lambda_z$. □

The $p(1 - p)$ factor is the binary signature. Near $p = 1/2$, maximum uncertainty makes order flow maximally informative. Near $p \in \{0, 1\}$, the market maker's strong prior makes order flow nearly uninformative. In the Gaussian Kyle model, price impact λ is constant. The state-dependence here arises entirely from the bounded support of v.

3.2 Symmetric equilibrium

Proposition 3 (Symmetric Equilibrium)**.** *When $p_0 = 1/2$, a symmetric equilibrium exists with $x_1 = a^*$, $x_0 = -a^*$, where $a^* = k^*\sigma_u/2$ and k^* is a positive solution to*

$$G(k) := \mathbb{E}_Z[\sigma(S)] - \frac{k^2}{2} \mathbb{E}_Z[\sigma(S)(1 - \sigma(S))] = 0, \tag{3}$$

with $S = -k^2/2 - kZ$ and $Z \sim N(0, 1)$.

Proof. With $z_0 = 0$ and $c = 0$ (by symmetry), the trader with $v = 1$ choosing order $\tilde{a}$ faces posterior $p_1 = \sigma(\lambda_z(\tilde{a}+u))$ where $\lambda_z = 2a/\sigma_u^2$. Expected profit is $\Pi(\tilde{a}; a) = \tilde{a}\,\mathbb{E}[\sigma(-m-\tau Z)]$, where $m = 2a\tilde{a}/\sigma_u^2$, $\tau = 2a/\sigma_u$, and Z is standard normal (absorbing the constant 1 from the $v = 1$ payoff, since profit per contract is $1 - p_1$).

The first-order condition at $\tilde{a} = a$ uses $d\Pi/d\tilde{a} = \mathbb{E}[S_a] - \tilde{a}\lambda_z\mathbb{E}[S_a(1 - S_a)]$, where $S_a = \sigma(-k^2/2 - kZ)$. Setting $\tilde{a} = a$ gives $G(k) = 0$.

Existence follows from the profit function's shape: $\hat{\Pi}(k) = (k\sigma_u/2)\mathbb{E}[S]$ satisfies $\hat{\Pi}(0) = 0$, $\hat{\Pi} > 0$ for small k, and $\hat{\Pi} \to 0$ as $k \to \infty$ (by dominated convergence, since $S \to 0$ pointwise). An interior maximum exists, where $G(k^*) = 0$. □

Proposition 4 (Universal Constant)**.** *The equilibrium satisfies $k^* = 2$ exactly. The equilibrium quantities are:*

$$a^* = \sigma_u, \qquad \lambda_z = \frac{2}{\sigma_u}, \qquad \Pi^* = 0.219\sigma_u.$$

Proof. The key step is an identity: $\mathbb{E}[S(1 - S)] = \mathbb{E}[S]/2$ for all $k > 0$.

Let $W = -k^2/2 - kZ$ with $Z \sim N(0, 1)$. Since $S(1 - S) = \sigma'(W) = e^{-W}/(1 + e^{-W})^2$:

$$\mathbb{E}[S(1 - S)] = \int \frac{e^{-W}}{(1 + e^{-W})^2}\, \phi(z)\, dz.$$

Apply the substitution $z \to -z - k$. Under this map, $W = -k^2/2 - kz$ transforms to $k^2/2 + kz = -W$, and the Gaussian density picks up the factor $e^{-kz-k^2/2} = e^W$. The integrand becomes

$$\frac{e^W}{(1 + e^W)^2} \cdot e^W = \frac{e^{2W}}{(1 + e^W)^2} = S^2,$$

where the last equality uses $S = e^W/(1 + e^W)$. Therefore $\mathbb{E}[S(1 - S)] = \mathbb{E}[S^2]$. Since $\mathbb{E}[S(1 - S)] = \mathbb{E}[S] - \mathbb{E}[S^2]$, combining gives $\mathbb{E}[S^2] = \mathbb{E}[S]/2$, hence $\mathbb{E}[S(1 - S)] = \mathbb{E}[S]/2$. The full proof appears in Appendix A.1.

Substituting into $G(k) = \mathbb{E}[S] - (k^2/2)\mathbb{E}[S(1 - S)] = \mathbb{E}[S] - (k^2/4)\mathbb{E}[S] = \mathbb{E}[S](1 - k^2/4)$. Since $\mathbb{E}[S] > 0$ for all k, the equation $G(k) = 0$ reduces to $k^2/4 = 1$, giving $k^* = 2$. □

The constant $k^* = 2$ is universal: equation (3) involves only k, the standard normal distribution, and the logistic function. The noise trader volatility σ_u scales all quantities proportionally (a^*, Π^*) or inversely (λ_z), but k^* itself is parameter-free. The informed trader's optimal order matches the noise trader's standard deviation exactly: $a^* = \sigma_u$. This scale invariance carries over to the multi-round model (Lemma 13).

Remark 5 (Economic intuition for $k^* = 2$). In the Gaussian Kyle model, the informed trader's optimal order is $\beta(v - \mu) = (\sigma_u/\sigma_v)(v - \mu)$: the deviation from the prior, scaled by the signal-to-noise ratio. The proportionality constant $\beta = \sigma_u/\sigma_v$ depends on both noise and fundamental volatility.

In the binary model, the fundamental "volatility" is fixed: $v \in \{0, 1\}$, so $\sigma_v = \sqrt{p_0(1 - p_0)} = 1/2$ at the symmetric prior. The informed trader cannot adjust for σ_v because it is absorbed into the contract's binary structure. The logistic pricing rule is "softer" than the linear rule at the margin: the first dollar of order flow moves the log-odds by λ_z, but the probability only moves by $\lambda_z p(1-p) = \lambda_z/4$ at $p = 1/2$. This reduced marginal cost of trading (relative to information revealed) is exactly offset at $a = \sigma_u$, where the informed trader's order is indistinguishable in scale from the noise. The identity $\mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$ captures this balance: the expected price impact cost (proportional to $\mathbb{E}[S(1-S)]$) is exactly half the expected gross gain (proportional to $\mathbb{E}[S]$), forcing the optimal scale to be σ_u.

Proposition 6 (Uniqueness)**.** *The equation $G(k) = 0$ has exactly one positive solution.*

Proof. The identity $\mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$ (Proposition 4) reduces $G(k) = \mathbb{E}[S](1 - k^2/4)$. Since $\mathbb{E}[S] > 0$ for all $k > 0$, the equation $G(k) = 0$ is equivalent to $k^2 = 4$, which has the unique positive solution $k = 2$. Equivalently, $G'(2) = -\mathbb{E}[S] < 0$ (since $G'(k) = \mathbb{E}'[S](1-k^2/4) + \mathbb{E}[S](-k/2)$ and the first term vanishes at $k = 2$): the function crosses zero exactly once and with negative slope. □

3.3 Asymmetric priors

Proposition 7 (Asymmetric Prior)**.** *The symmetric equilibrium at $p_0 = 1/2$ extends smoothly to p_0 near $1/2$. For $p_0 > 1/2$: $x_1^*(p_0) < |x_0^*(p_0)|$. The informed trader confirming the prior trades less than the informed trader contradicting it.*

Proof. Define $F(x_1, x_0; p_0) = 0$ as the two-equation first-order condition system. At $(a^*, -a^*; 1/2)$, $F = 0$ by Proposition 3. The Jacobian $J = \partial F/\partial(x_1, x_0)$ has diagonal entries equal to the second-order conditions (strictly negative at a maximum) and off-diagonal entries capturing cross-effects through λ_z and c. By symmetry at $p_0 = 1/2$, $J = \left(\begin{smallmatrix} -a & b \\ b & -a \end{smallmatrix}\right)$ with $a > |b|$ (diagonal dominance from the direct second-order effect exceeding the indirect pricing-rule effect). The implicit function theorem gives smooth functions $x_1^*(p_0)$, $x_0^*(p_0)$ near $p_0 = 1/2$.

For the comparative static: increasing p_0 shifts the prior toward $v = 1$. The $v = 1$ trader's informational advantage shrinks (the market already leans the right way), reducing x_1^*. The $v = 0$ trader's informational advantage grows (the market leans the wrong way), increasing $|x_0^*|$. Formally, $-J^{-1}\partial F/\partial p_0$ evaluated at $p_0 = 1/2$ has the required signs because $\partial F_1/\partial p_0 < 0$ (higher prior reduces the $v = 1$ trader's edge) and $\partial F_0/\partial p_0 > 0$ (higher prior increases the $v = 0$ trader's edge). □

The asymmetry in order sizes reflects differential information value. When the prior already favors the truth, the informed trader confirming the prior has less new information to trade on. The informed trader contradicting the prior has more valuable information, reflected in a larger order.

3.4 Comparison with Gaussian Kyle

Proposition 8 (Binary vs. Gaussian). *At $p_0 = 1/2$, matching the first two moments of v:*

Quantity	*Gaussian Kyle*	*Binary Kyle*
Pricing rule	$p_1 = 1/2 + \lambda Y$	$p_1 = \sigma(\lambda_z Y)$
Price impact	$1/(4\sigma_u)$ *(constant)*	$\lambda_z p(1-p)$, *peak* $0.500/\sigma_u$
Order size $(v = 1)$	σ_u	σ_u
Expected profit	$0.250\sigma_u$	$0.219\sigma_u$

The binary informed trader earns 12.4% less than the Gaussian benchmark.

The binary trader earns less because the bounded payoff limits the maximum edge per contract to 1. In the Gaussian model, extreme realizations of v create large profits. The logistic pricing rule partially compensates (it extracts less rent per unit than the linear rule at extreme prices), but the bounded payoff dominates.

3.5 Comparison with Glosten-Milgrom

Proposition 9 (Glosten-Milgrom Comparison). *The binary Kyle and binary Glosten-Milgrom models share the $p(1-p)$ state-dependent price impact structure. They differ in three ways:*

(a) Order size: *Glosten-Milgrom constrains trades to $x \in \{-1, +1\}$; Kyle allows continuous order sizes, producing the endogenous $k^* = 2$.*

(b) Information per round: *Kyle reveals more information per unit of noise because the informed trader optimally scales order size.*

(c) Dynamics: *The back-loading result (Proposition 17) requires continuous order size optimization. Glosten-Milgrom's unit-trade framework cannot generate timing effects.*

Proof. In the binary Glosten-Milgrom model with informed trader probability μ: at prior p, the ask price satisfies $\text{ask}(p) = p(1+\mu)/[p(1+\mu) + (1-p)(1-\mu)]$. The log-odds shift per buy order is $\Delta z = \log((1+\mu)/(1-\mu))$. The price impact $\text{ask}(p) - p$ is proportional to $p(1-p)$, confirming the shared Bayesian structure.

In the Kyle model, the information content per round (expected squared log-odds shift) is $k^{*4}/4 + k^{*2} = 8$. In Glosten-Milgrom, it is $[\log((1+\mu)/(1-\mu))]^2 \approx 4\mu^2$ for small μ. For comparable noise-to-signal ratios, Kyle reveals more because the informed trader optimally exploits the continuous order size. $\square$

The Glosten-Milgrom comparison clarifies the contribution of the Kyle framework to the binary setting. The $p(1-p)$ pattern is a property of binary Bayesian updating, common to any model with Bernoulli fundamentals. The Kyle framework adds endogenous order size, dynamic optimization, and the resulting back-loading and continuous-time characterization.

3.6 Position limits

Proposition 10 (Position Limits)**.** *Impose* $|x| \le \bar{Q}$ *with* $p_0 = 1/2$*.*

(a) The constraint is slack for $\bar{Q} \ge \sigma_u$*.*

(b) The constraint binds for $\bar{Q} < \sigma_u$*; the equilibrium becomes* $x_1 = \bar{Q}$*,* $x_0 = -\bar{Q}$*.*

(c) Information aggregation, measured by the variance of the posterior, increases in $\bar{Q}$ *for* $\bar{Q} \le a^*$*.*

(d) Informed profit increases in $\bar{Q}$ *for* $\bar{Q} \le a^*$*.*

Proof. Part (a) follows from $a^* = \sigma_u$. For (b), when $\bar{Q} < a^*$, the unconstrained optimum exceeds the limit, so the trader submits the maximum allowed. Parts (c) and (d): a larger order increases λ_z, which increases both the informativeness of order flow (the signal-to-noise ratio in the Bayesian update) and the trader's expected profit (up to the unconstrained optimum). □

Position limits reduce the informed trader's ability to move prices toward the truth. The threshold $\bar{Q} = \sigma_u$ provides an exact calibration benchmark: PredictIt's $850 cap is binding whenever σ_u (in dollar terms) exceeds $850, consistent with documented inefficiencies in PredictIt prices.

3.7 Market maker viability

Proposition 11 (Market Maker Risk)**.** *A risk-averse market maker with mean-variance preferences faces expected profit of zero (by competitive pricing) and profit variance*

$$Var(\pi_{MM}) = p_0(1-p_0) \cdot \mathbb{E}[(Y - p_1(Y))^2 \mid \textit{losing side}],$$

which is maximized at $p_0 = 1/2$*. Risk-averse market makers demand wider spreads at intermediate prices.*

At $p = 1/2$, every trade has a roughly equal chance of ending on the winning or losing side at resolution. The market maker's terminal exposure is binary: the entire position pays off or becomes worthless. This "all-or-nothing" terminal risk differs qualitatively from equity market making, where the market maker faces continuous price risk that can be hedged. Palumbo [2025] documents that Kalshi market makers behave more like underwriters than inventory-neutral intermediaries, consistent with this terminal risk structure.

3.8 Fee distortion

Proposition 12 (Fee Distortion). *A proportional fee $f \cdot p_1(1-p_1)$ per contract (the Kalshi fee structure) reduces equilibrium trading intensity: $k^*(f) < k^*$ for small $f > 0$. The distortion is largest at $p = 1/2$.*

Proof. The fee adds a marginal cost $f\,\mathbb{E}[S(1-S)] > 0$ at the frictionless optimum, shifting the informed trader's best response leftward. By continuity of the equilibrium map and the envelope theorem, the comparative static holds for small f. The fee $f \cdot p(1-p)$ is maximized at $p = 1/2$, so the marginal cost of trading is highest where information has the greatest value for moving prices toward the truth. □

Welfare analysis. In the frictionless equilibrium ($f = 0$), the informed trader's expected profit $\Pi^* = \sigma_u \cdot \mathbb{E}[S] = 0.225\sigma_u$ equals the noise traders' expected loss (the market maker breaks even). Information aggregation, measured by the posterior log-odds shift ${k^*}^2/2 = 2$, represents the rate at which the price incorporates private information.

The fee f reduces trading intensity from $k^* = 2$ to $k^*(f) < 2$. The welfare cost has two components. First, the informed trader's profit falls, directly reducing information aggregation: the log-odds shift drops from 2 to $k^*(f)^2/2 < 2$. Second, noise traders benefit from reduced adverse selection: their expected loss falls from Π^* to $\Pi^*(f) < \Pi^*$. Total surplus (informed profit minus noise trader loss) is zero by the competitive market maker condition, so the welfare effect operates entirely through the information channel. A social planner who values price accuracy at rate ν per unit of log-odds shift faces the tradeoff: fee revenue $f \cdot \mathbb{E}[p_1(1-p_1)] \cdot |x_1^*|$ versus information loss $\nu \cdot (2 - k^*(f)^2/2)$. The optimal f equates the marginal revenue from fees with the marginal value of information destroyed.

4 Dynamic Extensions

4.1 The N-round model

Lemma 13 (Scale Invariance). *The equilibrium constant k^* in the one-round symmetric equilibrium does not depend on noise volatility: equation* (3) *involves only k, the standard* *normal distribution, and the logistic function. In the last round of the N-round model, the symmetric equilibrium at $z = 0$ has the same k^* but order size $a^*_{N-1} = k^*\sigma_u/(2\sqrt{N})$.*

Proposition 14 (Bellman Equation)**.** *The value function $V_n(z, v)$ for the informed trader satisfies*

$$V_n(z, v) = \max_{x_n} \mathbb{E}_{u_n}\big[x_n(v - \sigma(z_n)) + V_{n+1}(z_n, v)\big], \tag{4}$$

where $z_n = z + \lambda_n(x_n + u_n) + c_n$ and λ_n, c_n are determined by the equilibrium strategy at round n. The terminal condition is $V_N(z, v) = 0$.

Proof. The total payoff $\sum_n x_n(v - p_n)$ decomposes into the round-n payoff and the continuation. The log-odds z_n is a sufficient statistic for the round-n state because it summarizes all public information relevant to future pricing and strategies. The Bellman equation follows by standard dynamic programming. □

4.2 Information conservation and back-loading

The central dynamic question is whether the informed trader front-loads or back-loads trading. In the Gaussian Kyle model, trading intensity is constant: the informed trader is indifferent about timing. Binary payoffs break this indifference.

Proposition 15 (Value Function Properties)**.** *Define $V(z) := V_2(z, 1)$, the $v = 1$ informed trader's equilibrium profit in the one-round game at prior log-odds z with noise $\tilde{\sigma}$.*

(a) $V(z) > 0$ *for all* $z \in \mathbb{R}$.

(b) $V(z) \to 0$ *as* $z \to +\infty$.

(c) $V'(0) < 0$.

(d) $V'(0) + V''(0) < 0$.

Parts (a) and (b) are immediate: the informed trader always has an edge (the posterior is bounded away from the truth), and the edge vanishes when the prior equals the truth. Part (c) follows from the envelope theorem applied to the equilibrium system; the proof appears in Appendix A.2. Part (d) is proved analytically. The envelope theorem gives $V'(0) = a^*(k^{*2}D/2 - 2A)$ where $A = \mathbb{E}[S(1-S)]$ and $D = \mathbb{E}[S(1-S)(1-2S)]$ (Appendix A.2). A second application of the envelope theorem gives $V''(0)$. Combining, $V'(0) + V''(0) = -0.116\tilde{\sigma} < 0$ (Appendix A.3).

Proposition 16 (Information Conservation)**.** *In the two-round model with $p_0 = 1/2$ and $\tilde{\sigma} = \sigma_u/\sqrt{2}$:*

(a) *The last round is a one-round game with noise* $\tilde{\sigma}$.

(b) *The round-1 equilibrium order satisfies* $a_0^* < k^*\tilde{\sigma}/2$ *(strictly less than the single-round optimum).*

(c) *Total two-round profit exceeds the single-round profit:* $\Pi^{(2)}_{total} > \Pi^{(1)}$.

Proof. Part (a) holds because no continuation exists after round 2.

For part (b), the round-1 first-order condition equates the marginal round-1 profit with the marginal loss in continuation value. At $a_0 = a^{\text{single}} := k^*\tilde{\sigma}/2$, the round-1 marginal profit is zero (by the single-round first-order condition). The continuation value term, evaluated using Stein's lemma, is

$$\frac{d}{da_0}\mathbb{E}[V_2(z_1,1)]\bigg|_{a^{\text{single}}} = \frac{2k^*}{\tilde{\sigma}}\big(\mathbb{E}[V'(z_1)] + \mathbb{E}[V''(z_1)]\big),$$

where $z_1 = k^{*2}/2 + k^*Z$ under $v = 1$. Computing $V(z)$ on a fine grid ($\Delta z = 0.01$, 601 points) via the one-round equilibrium solver at each z, and evaluating the expectation by Gauss-Hermite quadrature over the z_1 distribution, gives $\mathbb{E}[V'(z_1) + V''(z_1)] = -0.036 < 0$ (Proposition 15(d) and Table above confirm $V' + V'' < 0$ pointwise). The continuation derivative is negative. The total marginal payoff at a^{single} is therefore negative, which means the optimum satisfies $a_0^* < a^{\text{single}}$.

Part (c): the single-round strategy (trade a^{single} in round 1, do nothing in round 2) is feasible in the two-round game and earns $\Pi^{(1)}$. The optimal two-round strategy earns strictly more because the informed trader can profitably use round 2 (retaining private information after round 1). □

The Stein's lemma step is the key technical device. Standard Stein's lemma gives $\mathbb{E}[V'(z_1)Z] = k^*\mathbb{E}[V''(z_1)]$ for $z_1 = k^{*2}/2 + k^*Z$. Combined with $dz_1/da_0 = (2/\tilde{\sigma})(k^* + Z)$, the continuation derivative factors as $\mathbb{E}[V'(z_1)+V''(z_1)]$, reducing the back-loading condition to properties of the value function at a single point.

Proposition 17 (Back-Loading)**.** *In the two-round model, the informed trader's expected round-2 profit exceeds the expected round-1 profit.*

Proof. By Proposition 16(b), $a_0^* < a^{\text{single}}$, so round-1 profit $\Pi_1 < \Pi^{(1)}$ (the trader trades less than the single-round optimum, and the single-round optimum is the unique maximizer). By Proposition 16(c), total two-round profit exceeds $\Pi^{(1)}$: $\Pi_1 + \mathbb{E}[\Pi_2] > \Pi^{(1)}$. Combining: $\mathbb{E}[\Pi_2] > \Pi^{(1)} - \Pi_1 > 0 > \Pi_1 - \Pi^{(1)}$, so $\mathbb{E}[\Pi_2] > \Pi_1$. □

Two forces drive back-loading, both unique to the binary model.

First, state-dependent price impact creates a timing advantage. After round 1, the price p_1 has moved away from 1/2 (toward the truth), so $p_1(1-p_1) < 1/4$. Price impact is lower in round 2 than in round 1: trading later is cheaper.

Second, the bounded payoff interacts with the logistic curvature. The maximum gain per contract is 1, regardless of the price. As the price approaches the truth, per-contract gain and per-unit price impact both shrink. In the Gaussian model, both scale linearly, canceling out. In the binary model, price impact (via the logistic's curvature) shrinks faster, creating a net advantage to waiting.

The back-loading result for $N = 2$ rests on the analytically verified condition $V'(0) + V''(0) = -0.116\tilde{\sigma} < 0$ (Appendix A.3). For general N, the same mechanism operates, but a formal proof requires solving the full N-round backward induction. We state the general-N result as a conjecture.

Conjecture 18 (General-N Back-Loading)**.** *In the N-round model, the expected per-round profit is increasing in the round number for all $N \geq 2$.*

The conjecture is supported by the $N = 2$ proof and the economic mechanism: in each round, trading moves the price toward the truth, reducing future price impact and making later rounds relatively more profitable.

4.3 Continuous-time price dynamics

Proposition 19 (Price Dynamics)**.** *In the continuous-time limit, the price process satisfies*

$$dp_t = \alpha(p_t, t) \cdot p_t(1 - p_t)\, dB_t, \tag{5}$$

where B_t is the innovation process (standard Brownian motion under the market maker's filtration) and $\alpha(p_t, t) = (\theta_t^{(1)} - \theta_t^{(0)})/\sigma_u$ is the normalized strategy spread.

Proof. The Kushner-Stratonovich nonlinear filtering equation for binary states with Gaussian observation noise gives

$$dp_t = p_t(1 - p_t)\frac{\theta_t^{(1)} - \theta_t^{(0)}}{\sigma_u}\, dB_t,$$

where $dB_t = (dY_t - (p_t\theta_t^{(1)} + (1 - p_t)\theta_t^{(0)})dt)/\sigma_u$ is the innovation. The diffusion coefficient $p(1 - p)$ is the binary signature: prices follow a Wright-Fisher-type diffusion with state-dependent volatility. The Wright-Fisher diffusion, originally developed in population genetics to model allele frequencies bounded in $[0,1]$, arises here because prediction market prices share the same bounded state space and variance structure proportional to $p(1-p)$. □

4.4 The impossibility of state-independent strategies

The binary model's equilibrium structure differs fundamentally from the Gaussian benchmark.

Proposition 20 (State-Dependent Strategy Spread)**.** *In the binary Kyle-Back model, the equilibrium strategy spread $\alpha(p,t) = (\theta^{(1)}(p,t) - \theta^{(0)}(p,t))/\sigma_u$ must depend on p. No state-independent strategy spread is consistent with the HJB equation.*

Proof. **Step 1: Implications of a deterministic spread.** Suppose $\alpha(t)$ depends only on time. The price process becomes $dp_t = \alpha(t)\, p_t(1-p_t)\, dB_t$, a Wright-Fisher diffusion with time-dependent coefficient.

The HJB first-order condition (from Proposition 21 below) gives

$$W_p(p,t) = -\frac{\sigma_u}{p \cdot \alpha(t)},$$

which integrates to $W(p,t) = -(\sigma_u/\alpha(t)) \log p + h(t)$ for some function $h(t)$.

Step 2: The HJB constraint. Substituting W_p and $W_{pp} = \sigma_u/(p^2\alpha(t))$ into the HJB equation at the optimum yields, after simplification,

$$W_t = -\frac{\alpha(t)\sigma_u}{2}\, p(1-p).$$

Step 3: The contradiction. From Step 1, $W_t = (\sigma_u \alpha'(t)/\alpha(t)^2) \log p + h'(t)$. Setting Step 2 equal to Step 3:

$$\frac{\sigma_u \alpha'(t)}{\alpha(t)^2} \log p + h'(t) = -\frac{\alpha(t)\sigma_u}{2}\, p(1-p).$$

The left side is affine in $\log p$; the right side is $p(1-p)$. These are linearly independent functions of p on $(0,1)$. They cannot be equal for all p unless both coefficients vanish: $\alpha'(t) = 0$ (so α is constant) and $h'(t) = -(\alpha\sigma_u/2)p(1-p)$ for all p (impossible, since $h'(t)$ does not depend on p). Contradiction.

Step 4: Why Gaussian is different. In the Gaussian model with $v \in \mathbb{R}$, the value function is $W(p,t) = (v-p)^2/(2\Sigma(t))$ (quadratic in $v-p$). The first-order condition gives a linear strategy $\theta^{(1)} = (v-p)\beta(t)$, so the spread $\Delta\theta = \beta(t)$ is deterministic. The quadratic value function is consistent with the HJB because Gaussian dynamics preserve the quadratic form. In the binary model, the value function cannot be quadratic (the payoff is bounded), so the HJB forces α to depend on p. $\square$

The proof identifies the structural source of the difference: the logarithmic value function ($W \propto \log p$) implied by a deterministic strategy is incompatible with the $p(1-p)$ structure of binary dynamics. In the Gaussian model, the quadratic value function and linear dynamics are mutually consistent. Binary payoffs break this consistency.

4.5 PDE characterization

Proposition 21 (HJB Equation)**.** *The value function $W(p,t)$ for the $v = 1$ informed trader satisfies*

$$0 = W_t + \max_{\theta^{(1)}} \left\{ \theta^{(1)}(1-p) + \mu(p,t;\theta^{(1)})W_p + \frac{1}{2}\nu(p,t)^2 W_{pp} \right\},$$

where $\mu(p,t;\theta^{(1)}) = p(1-p)\alpha_{eq}(p,t)(\theta^{(1)} - \hat{\theta}(p,t))/\sigma_u$, $\nu(p,t) = p(1-p)\alpha_{eq}(p,t)$, $\hat{\theta} = p\theta^{(1)}_{eq} + (1-p)\theta^{(0)}_{eq}$, and the first-order condition is

$$(1-p) + \frac{p(1-p)\alpha_{eq}(p,t)}{\sigma_u} W_p = 0. \tag{6}$$

Proof. Under $v = 1$, $dY_t = \theta^{(1)}dt + \sigma_u\, dW_t$. The innovation is $dB_t = (\theta^{(1)} - \hat{\theta})dt/\sigma_u + dW_t$. The price dynamics under $v = 1$ are $dp_t = p(1-p)\alpha_{eq}[(\theta^{(1)} - \hat{\theta})/\sigma_u\, dt + dW_t]$. The running payoff is $\theta^{(1)}(1-p)\, dt$. The HJB equation follows from standard stochastic control. The first-order condition differentiates the maximand with respect to $\theta^{(1)}$. $\square$

Proposition 22 (PDE for the Volatility)**.** *The equilibrium volatility $\alpha(p,t)$ satisfies the nonlinear PDE*

$$\frac{\alpha_t}{p\alpha^2} = -\frac{1}{2}\left[(1-2p)(\alpha + (1-p)\alpha_p) + p(1-p)^2\alpha_{pp}\right], \tag{7}$$

with the boundary condition $\alpha(p,t) \to \infty$ as $t \to T$ (to ensure $p_T \in \{0,1\}$ a.s.) and regularity at $p \in \{0,1\}$.

Proof. From the first-order condition (6): $W_p = -\sigma_u/(p\alpha)$. Differentiating: $W_{pp} = \sigma_u/(p^2\alpha) + \sigma_u\alpha_p/(p\alpha^2)$. Substituting into the HJB at the equilibrium and simplifying yields

$$W_t = -\frac{\sigma_u p(1-p)}{2}[\alpha + (1-p)\alpha_p].$$

From the first-order condition, $W(p,t) = -\sigma_u \int_{\bar{p}}^{p} [q\alpha(q,t)]^{-1} dq + h(t)$, so $W_t = \sigma_u \int_{\bar{p}}^{p} \alpha_t(q,t)/(q\alpha(q,t)^2)\, dq + h'(t)$. Differentiating both expressions for W_t with respect to p and equating gives (7). $\square$

4.6 Qualitative properties of the equilibrium

Proposition 23 (Symmetry). *If the continuous-time equilibrium is unique, then* $\alpha(p,t) = \alpha(1-p,t)$ *for all* $p \in (0,1)$ *and* $t \in [0,T)$.

Proof. At $p_0 = 1/2$, the map $(p,v) \mapsto (1-p, 1-v)$ is an automorphism of the game: it permutes the states, reverses the prior, and preserves all equilibrium conditions. Under this map, $\theta^{(1)}(p,t)$ becomes $\theta^{(0)}(1-p,t)$ and $\theta^{(0)}(p,t)$ becomes $\theta^{(1)}(1-p,t)$. The strategy spread transforms as $\Delta\theta(p,t) \mapsto \Delta\theta(1-p,t)$. If the equilibrium is unique, it must be invariant under this automorphism, giving $\alpha(p,t) = \alpha(1-p,t)$. □

Proposition 24 (Terminal Blowup). *As* $t \to T$*, the equilibrium volatility blows up as*

$$\alpha(p,t) \sim \frac{A(p)}{\sqrt{T-t}}, \tag{8}$$

where $A(p)$ *satisfies the nonlinear ODE*

$$\frac{1}{2pA} = -\frac{1}{2}\left[(1-2p)(A+(1-p)A') + p(1-p)^2 A''\right], \tag{9}$$

on the domain $p \in [1/2, 1]$ *(from the* $v = 1$ *trader's HJB), with boundary conditions* $A'(1/2) = 0$ *and* $A(1) = 1$*. The equilibrium symmetry* $A(p) = A(1-p)$ *extends the solution to* $[0, 1/2]$.

Proof. Substitute the ansatz $\alpha(p,t) = A(p)/\sqrt{T-t}$ into the PDE (7). The left side becomes $1/(2pA\sqrt{T-t})$. The right side is $-1/(2\sqrt{T-t})$ times the bracketed expression in (7) with A replacing α. Canceling $1/\sqrt{T-t}$ gives (9).

Boundary condition. Set $q = 1-p$ and evaluate the ODE as $q \to 0$. The second-derivative term $p(1-p)^2A'' = (1-q)q^2A''$ vanishes as q^2. The first-derivative term $(1-2p)(A+(1-p)A') = (2q-1)(A+qA') \to -A$. The forcing term $1/(pA) \to 1/A$. The dominant balance is $-A + 1/A = 0$, giving $A(1) = 1$. By symmetry, $A(0) = 1$.

Local behavior at $p = 1/2$. At $p = 1/2$: $1 - 2p = 0$ and $p(1-p)^2 = 1/8$, so the ODE gives $A''(1/2) = -16/A(1/2) < 0$.

Boundary expansion near $p = 1$. Setting $A = 1 + a_1 q + a_2 q^2 + \cdots$ with $q = 1 - p$ and matching powers of q in the ODE determines $a_1 = 3$, $a_2 = -7/2$, $a_3 = 47/10$ uniquely. There is no free parameter at the boundary: the ODE and $A(1) = 1$ determine the local expansion completely. □

The terminal blowup at rate $(T-t)^{-1/2}$ is the binary analog of Back [1992]'s result $\sigma(t) = \sigma_v/\sqrt{T-t}$ in the Gaussian case. The Gaussian model has a time-dependent amplitude that does not depend on the price. The binary model has a price-dependent amplitude $A(p)$, a direct consequence of the state-dependent strategy (Proposition 20).

Solution of the ODE. The boundary value problem $A(0) = A(1) = 1$, $A'(1/2) = 0$ has a unique smooth solution on $[0, 1]$, obtained numerically by collocation (Runge-Kutta with boundary matching). Table 1 reports $A(p)$.

p	0.00	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00
$A(p)$	1.00	1.28	1.50	1.66	1.77	1.81	1.77	1.66	1.50	1.28	1.00

Table 1: The terminal-layer amplitude $A(p)$ solving the ODE (9) with boundary conditions $A(0) = A(1) = 1$ and symmetry $A(p) = A(1-p)$. The maximum $A(1/2) = 1.81$ occurs at the symmetric prior. The N-period computational results give a rescaled amplitude $\alpha_N(0, t) \cdot \sqrt{1-t}$ that ranges from 1.61 at $t = 0$ to 1.73 at $t = 0.88$ for $N = 32$. The self-similar ansatz is asymptotically valid near $t = T$; the rescaled amplitude trends toward $A(1/2)$ as $t \to T$, with the remaining 5% gap attributable to finite-N discretization.

The solution has a natural economic interpretation. The amplitude $A(p)$ measures the intensity of equilibrium trading at price level p as resolution approaches. At $p = 1/2$ (maximum uncertainty), trading intensity is highest: both the informed trader's edge and the market maker's uncertainty are maximal. At p near 0 or 1 (the market is nearly certain), $A(p) \to 1$: trading intensity drops but remains positive because the informed trader on the wrong side of the market still has an edge worth exploiting. The fact that $A(1) = 1 > 0$ (rather than $A(1) = 0$) reflects a fundamental asymmetry of binary markets: when the market is "almost right," the minority-side informed trader trades aggressively into the confident market, generating nonzero volatility even at extreme prices.

Proposition 25 (Monotonicity)**.** *The terminal-layer amplitude satisfies $A'(p) > 0$ for $p \in (0, 1/2)$ and $A'(p) < 0$ for $p \in (1/2, 1)$. The amplitude A is maximized at $p = 1/2$.*

Proof. At $p = 1/2$: $A'(1/2) = 0$ (symmetry) and $A''(1/2) = -16/A(1/2) < 0$, so A is strictly concave and decreasing for p slightly above $1/2$. At $p = 1$: $A(1) = 1$ and the boundary expansion gives $A'(1) = -a_1 = -3 < 0$, so A is decreasing near $p = 1$. The numerical BVP solution confirms $A'(p) < 0$ for all $p \in (1/2, 1)$ with no interior critical points. Symmetry gives the result on $(0, 1/2)$. □

4.7 Computational verification

To verify the analytical predictions and demonstrate convergence to the continuous-time limit, we solve the N-round equilibrium numerically for $N \in \{1, 2, 4, 8, 16, 32\}$ by backward induction over the Bellman equation (Proposition 14). Each round's equilibrium uses 20-node Gauss-Hermite quadrature for the expectation integrals.

Per-round trading intensity. Table 2 reports the per-round trading intensity k_n (at $z = 0$) and the back-load ratio $k_{\text{last}}/k_{\text{first}}$ for each N.

N	k_{first}	k_{last}	Back-load ratio	Total profit/σ_u
1	2.000	2.000	1.00	0.219
2	1.207	2.000	1.66	0.277
4	0.866	2.000	2.31	0.321
8	0.583	2.000	3.43	0.351
16	0.420	2.000	4.76	0.372
32	0.285	2.000	7.02	0.386

Table 2: N-round equilibrium at $z = 0$. The back-load ratio $k_{\text{last}}/k_{\text{first}}$ grows monotonically with N. The $N = 1$ row uses the analytical $k^* = 2$. For $N > 1$, the last-round intensity $k_{\text{last}} = 2.0$ matches the analytical value exactly.

Three features stand out. First, the last-round intensity $k_{\text{last}} = 2$ for all N: the final round reproduces the one-round equilibrium, as Lemma 13 predicts. Second, the back-load ratio grows monotonically, confirming Conjecture 18 for all N tested. Third, total informed profit increases with N but at a decreasing rate (increments roughly halve: $+0.044, +0.043, +0.031, +0.021, +0.014$), suggesting convergence to a limit around $\Pi_\infty \approx$ 0.40–0.42.

Terminal blowup verification. Define the rescaled intensity $\alpha_N(0, t) = k_n \cdot \sqrt{N}$, where k_n is the per-round intensity at round $n = tN$. The self-similar ansatz (Proposition 24) predicts $\alpha(0, t) \approx A_0/\sqrt{1 - t}$ for some constant A_0. Table 3 reports the product $\alpha_N(0, t) \cdot \sqrt{1 - t}$ at fixed time points.

A power-law regression of $\alpha_{32}(0, t)$ on $(1 - t)$ gives exponent -0.529 (standard error 0.02), consistent with the predicted $-1/2$. The amplitude at $p = 1/2$ is $A_0 \approx 1.65$. In the Gaussian Kyle-Back model, the analogous blowup coefficient is $\sigma_v/\sqrt{T - t} = 0.5/\sqrt{T - t}$ (for $\sigma_v = 1/2$), so the binary model produces approximately 3.3 times higher trading intensity at the symmetric prior. The logistic pricing rule, which is "softer" than the linear rule at extreme order flows, permits more aggressive trading.

Summary. The computational results establish three facts. The continuous-time equilibrium exists as the limit of N-period equilibria (convergence of α_N). Back-loading holds

t	$N=4$	$N=8$	$N=16$	$N=32$
0.00	1.73	1.65	1.68	1.61
0.25	2.00	1.96	1.94	1.93
0.50	2.41	2.45	2.33	2.38
0.75	4.00	3.41	3.47	3.30

Table 3: The product $\alpha_N(0,t) \cdot \sqrt{1-t}$ at fixed time points. The self-similar ansatz $\alpha = A_0/\sqrt{1-t}$ is asymptotically valid near $t = T$; at early times ($t = 0, 0.25$) the product is approximately 1.65, while at $t = 0.75$ the equilibrium has not yet entered the terminal layer and the product exceeds A_0. As N increases, the terminal-layer regime extends further from T, and the product at $t = 0.75$ decreases (from 4.00 at $N = 4$ to 3.30 at $N = 32$), trending toward A_0.

for all N tested, extending the $N = 2$ analytical result. The Wright-Fisher structure $\alpha \sim A_0/\sqrt{T-t}$ with $A_0 \approx 1.65$ characterizes the terminal behavior.

4.8 Competition among insiders

The preceding analysis restricts attention to a single informed trader. Prediction markets typically have multiple informed participants. This subsection extends the one-round model to N_{ins} symmetric insiders, each observing $v \in \{0,1\}$.

Setup. Each insider $i \in \{1,\ldots,N_{\text{ins}}\}$ observes v and chooses order $x_i \in \mathbb{R}$. Noise traders submit $u \sim N(0,\sigma_u^2)$. The market maker observes total order flow $Y = \sum_{i=1}^{N_{\text{ins}}} x_i + u$ and sets $p_1 = \mathbb{E}[v \mid Y]$. Each insider maximizes $\mathbb{E}[x_i(v - p_1(Y)) \mid v]$. A symmetric equilibrium has $x_i(1) = a$, $x_i(0) = -a$ for all i.

Proposition 26 (Multi-Insider Equilibrium)**.** *In the symmetric one-round model with N_{ins} insiders and $p_0 = 1/2$:*

(a) The pricing rule is logistic: $p_1(Y) = \sigma(\lambda_z^{(N_{ins})} Y)$ with $\lambda_z^{(N_{ins})} = 2N_{ins}a/\sigma_u^2$. Price impact is $dp_1/dY = \lambda_z^{(N_{ins})} \cdot p_1(1-p_1)$.

(b) Define $k_{N_{ins}} = 2N_{ins}a/\sigma_u$, so $a = k_{N_{ins}}\sigma_u/(2N_{ins})$. The equilibrium constant $k_{N_{ins}}$ solves

$$G_{N_{ins}}(k) := \mathbb{E}[S] - \frac{k^2}{2N_{ins}}\mathbb{E}[S(1-S)] = 0, \quad S = \sigma(-k^2/2 - kZ). \tag{10}$$

(c) $k_{N_{ins}}$ is increasing in N_{ins}, and $k_{N_{ins}} > k^*$ for all $N_{ins} \geq 2$.

(d) Each insider's order $a_{N_{ins}} = k_{N_{ins}}\sigma_u/(2N_{ins})$ is decreasing in N_{ins}.

(e) *Total informed order* $N_{ins} \cdot a_{N_{ins}} = k_{N_{ins}}\sigma_u/2$ *is increasing in* N_{ins}.

(f) *Each insider's profit* $\Pi_{N_{ins}} = a_{N_{ins}}\,\mathbb{E}[S_{N_{ins}}]$ *is decreasing in* N_{ins}.

Proof. Part (a). Under $v = 1$, the total informed order is $N_{\text{ins}}a$, so $Y \sim N(N_{\text{ins}}a, \sigma_u^2)$. The log-likelihood ratio is $4N_{\text{ins}}a \cdot Y/(2\sigma_u^2) = 2N_{\text{ins}}aY/\sigma_u^2$. The posterior is $p_1 = \sigma(\lambda_z^{(N_{\text{ins}})}Y)$ with $\lambda_z^{(N_{\text{ins}})} = 2N_{\text{ins}}a/\sigma_u^2$. The price impact follows from $dp_1/dY = \lambda_z^{(N_{\text{ins}})}p_1(1-p_1)$.

Part (b). Insider 1 deviates to $\tilde{a}$ while the remaining $N_{\text{ins}}-1$ insiders play a. The market maker uses the equilibrium pricing rule. Under $v = 1$, order flow is $Y = \tilde{a} + (N_{\text{ins}}-1)a + u$, and insider 1's profit is

$$\Pi_1(\tilde{a}) = \tilde{a}\,\mathbb{E}\big[\sigma\big(-\lambda_z^{(N_{\text{ins}})}(\tilde{a} + (N_{\text{ins}}-1)a + u)\big)\big].$$

At the symmetric equilibrium $\tilde{a} = a$: the argument of σ is $-\lambda_z^{(N_{\text{ins}})}(N_{\text{ins}}a + u) = -k_{N_{\text{ins}}}^2/2 - k_{N_{\text{ins}}}Z$, and the first-order condition coefficient is $a \cdot \lambda_z^{(N_{\text{ins}})} = 2N_{\text{ins}}a^2/\sigma_u^2 = k_{N_{\text{ins}}}^2/(2N_{\text{ins}})$. The first-order condition is $G_{N_{\text{ins}}}(k_{N_{\text{ins}}}) = 0$.

Part (c). $G_{N_{\text{ins}}}(k) = G_1(k) + (k^2/2)(1 - 1/N_{\text{ins}})\mathbb{E}[S(1-S)]$. At $k = k^*$: $G_1(k^*) = 0$, so $G_{N_{\text{ins}}}(k^*) > 0$ for $N_{\text{ins}} \geq 2$. Since $G_{N_{\text{ins}}}(k) < 0$ for large k (because $\mathbb{E}[S]/\mathbb{E}[S(1-S)] \to 1$ while $k^2/(2N_{\text{ins}}) \to \infty$), the zero of $G_{N_{\text{ins}}}$ exceeds k^*. The same argument applied to $N'_{\text{ins}} > N_{\text{ins}}$ gives $k_{N'_{\text{ins}}} > k_{N_{\text{ins}}}$.

Parts (d)–(f). The equilibrium equation $H(k) = k^2/(2N_{\text{ins}})$, where $H(k) = \mathbb{E}[S]/\mathbb{E}[S(1-S)]$, combined with $H(k) \to 1$ as $k \to \infty$, gives $k_{N_{\text{ins}}} \sim \sqrt{2N_{\text{ins}}}$ for large N_{ins}. Individual order $a_{N_{\text{ins}}} = k_{N_{\text{ins}}}\sigma_u/(2N_{\text{ins}}) \sim \sigma_u/\sqrt{2N_{\text{ins}}}$, which is decreasing. Total informed order $k_{N_{\text{ins}}}\sigma_u/2$ is increasing since $k_{N_{\text{ins}}}$ is increasing. Individual profit $\Pi_{N_{\text{ins}}} = a_{N_{\text{ins}}}\mathbb{E}[S_{N_{\text{ins}}}]$ vanishes because both $a_{N_{\text{ins}}} \to 0$ and $\mathbb{E}[S_{N_{\text{ins}}}] \to 0$ (since $S = \sigma(-k^2/2 - kZ) \to 0$ pointwise as $k \to \infty$). □

The equilibrium constant $G_{N_{\text{ins}}}(k) = 0$ has the same functional form as the single-insider equation but with the coefficient $k^2/(2N_{\text{ins}})$ replacing $k^2/2$. Because $S = \sigma(-k^2/2 - kZ)$ is the same function of k in both cases, all the qualitative properties (existence, uniqueness, boundary behavior) carry over from the single-insider analysis.

Numerical values. Table 4 reports equilibrium quantities for $N_{\text{ins}} \in \{1, 2, 5, 10\}$. For $N_{\text{ins}} = 2$, the identity $\mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$ and the equilibrium equation $G_2(k) = \mathbb{E}[S](1 - k^2/8) = 0$ give $k_2 = 2\sqrt{2}$ exactly.

Competition reduces individual order sizes and profits while increasing total informed trading and information aggregation. Each insider's order scales as $\sigma_u/\sqrt{2N_{\text{ins}}}$, while total informed trading scales as $\sigma_u\sqrt{N_{\text{ins}}/2}$. Individual profit vanishes faster than $1/N_{\text{ins}}$ because $\mathbb{E}[S_{N_{\text{ins}}}]$ decays exponentially in $k_{N_{\text{ins}}}^2$. These patterns mirror the Gaussian multi-insider

	$N_{\text{ins}}=1$	$N_{\text{ins}}=2$	$N_{\text{ins}}=5$	$N_{\text{ins}}=10$
Equilibrium constant $k_{N_{\text{ins}}}$	2	$2\sqrt{2}\approx 2.83$	4.47	6.40
Each insider's order $a_{N_{\text{ins}}}/\sigma_u$	1.000	0.707	0.447	0.320
Total informed order $/\sigma_u$	1.000	1.414	2.236	3.198
Price impact $\lambda_z^{(N_{\text{ins}})}\cdot\sigma_u$	2	2.83	4.47	6.40
Individual profit $/\sigma_u$	0.225	0.082	0.009	0.0003
Log-odds shift $k_{N_{\text{ins}}}^2/2$	2	4	10	20

Table 4: Equilibrium quantities in the one-round model with N_{ins} symmetric insiders and $p_0=1/2$. For $N_{\text{ins}}\geq 5$, the asymptotic approximation $k_{N_{\text{ins}}}\approx\sqrt{2N_{\text{ins}}}$ applies. The log-odds shift $k_{N_{\text{ins}}}^2/2$ measures information aggregation per round.

model of Holden and Subrahmanyam [1992], with one structural difference: in the Gaussian one-round model, individual order also scales as $1/\sqrt{N_{\text{ins}}}$, but individual profit decays as $1/N_{\text{ins}}$ (polynomial) rather than exponentially.

Proposition 27 (Perfect Competition Limit). *As $N_{ins}\to\infty$:*

(a) $k_{N_{ins}}\sim\sqrt{2N_{ins}}$.

(b) Each insider's order $a_{N_{ins}}\sim\sigma_u/\sqrt{2N_{ins}}\to 0$.

(c) Total informed order $\to\infty$.

(d) Individual and total informed profit $\to 0$.

(e) The price converges to the true value: $p_1\to v$ *almost surely.*

Proof. Part (a): the equilibrium equation $H(k)=k^2/(2N_{\text{ins}})$ with $H(k)\to 1$ as $k\to\infty$ gives $k^2/(2N_{\text{ins}})\to 1$. Parts (b)–(d) follow from (a). Part (e): the posterior log-odds shift $k_{N_{\text{ins}}}^2/2+k_{N_{\text{ins}}}Z$ has mean $k_{N_{\text{ins}}}^2/2\to\infty$ under $v=1$, so $p_1\to 1$ almost surely. By symmetry, $p_1\to 0$ under $v=0$. □

Structural preservation. Three features of the single-insider equilibrium survive competition unchanged. The pricing rule remains logistic (Bayes' rule for binary states with Gaussian signals does not depend on the number of insiders). Price impact retains the $p(1-p)$ state-dependent form. The impossibility of state-independent strategies in continuous time (Proposition 20) depends on the binary boundary condition, not the number of insiders. Competition changes the level of price impact ($\lambda_z^{(N_{\text{ins}})}$ increases with N_{ins}) but not the functional form.

Binary vs. Gaussian comparison. Table 5 compares the one-round equilibrium with N_{ins} insiders across the two models.

Feature	Gaussian Kyle	Binary Kyle
Price impact form	Constant λ	State-dependent $\lambda_z \cdot p(1-p)$
Individual order scaling	$O(1/\sqrt{N_{\text{ins}}})$	$O(1/\sqrt{N_{\text{ins}}})$
Individual profit scaling	$O(1/N_{\text{ins}})$	Faster than $O(1/N_{\text{ins}})$
Full revelation at $N_{\text{ins}} \to \infty$	Yes	Yes
Competition externality	Information stealing only	Stealing + price impact reduction

Table 5: One-round equilibrium comparison: binary vs. Gaussian Kyle with N_{ins} insiders. Individual order scaling is the same in both models. The key structural difference is the competition externality in the last row.

The positive externality of waiting. The last row of Table 5 identifies a mechanism unique to the binary model. In the Gaussian Kyle model, price impact λ does not depend on the price level. When insider 1 trades, the only effect on insider 2 is information revelation: prices move closer to v, reducing insider 2's informational advantage. Competition is purely destructive.

In the binary model, price impact is $\lambda_z \cdot p(1-p)$. When insider 1 trades and moves p toward v, two things happen: (i) insider 2's informational advantage shrinks (the same stealing effect), but (ii) $p(1-p)$ decreases, reducing the price impact that all insiders face in subsequent trading. Effect (ii) is a positive externality of trading: each insider's activity improves market conditions for everyone.

A natural conjecture is that the positive externality might allow back-loading to survive competition. In the Gaussian model, competition produces front-loading: insiders race to trade before rivals reveal information [Holden and Subrahmanyam, 1992]. The binary model's state-dependent price impact creates an offsetting force: as any insider trades and moves the price, the declining $p(1-p)$ factor makes future trading cheaper. However, the information-stealing motive dominates the positive externality, producing front-loading under competition (Proposition 28).

Proposition 28 (Front-Loading under Competition)**.** *In the two-round binary Kyle model with $N_{ins} = 2$ symmetric insiders, competition produces front-loading: each insider's expected round-1 profit exceeds the expected round-2 profit.*

Proof. Backward induction: the round-2 game is a one-round two-insider equilibrium at the updated prior z_1. The round-1 problem for each insider maximizes current-round profit plus the continuation value $\mathbb{E}[V_2(z_1)]$. Numerical solution of the two-round, two-insider equilibrium (using Gauss-Hermite quadrature on a 241-point log-odds grid) gives round-1 profit $0.077\sigma_u$ and round-2 profit $0.064\sigma_u$: each insider earns more in round 1 than round 2. The optimal round-1 order is $0.213\sigma_u$, compared to the one-round optimum of $0.707\sigma_u$, confirming that insiders trade less than the static optimum (information conservation holds under competition) but that the split of profits favors early trading (front-loading). □

The positive externality of waiting (declining $p(1-p)$ benefits all insiders) exists but is weaker than the information-stealing incentive (each insider races to trade before the rival reveals information). The transition from back-loading (single insider) to front-loading (two insiders) is sharp: adding one competitor reverses the dynamic pattern entirely.

This result contrasts with the Gaussian benchmark, where competition also produces front-loading [Holden and Subrahmanyam, 1992], but for a simpler reason: in the Gaussian model there is no offsetting positive externality. The binary model has an additional force favoring patience (the $p(1-p)$ reduction), yet the information-stealing motive dominates. The prediction is testable: in markets where a single large insider is dominant, informed trading should concentrate near resolution; in markets with multiple informed participants, informed trading should concentrate early.

5 Discussion

5.1 Testable predictions

The model generates six testable predictions, ordered by discriminating power.

1. State-dependent price impact. Price impact in prediction markets should follow $\lambda(p) = \lambda_z \cdot p(1-p)$, with a peak at $p = 1/2$ and vanishing impact at the boundaries. Equivalently, in log-odds space, price impact should be constant: $\Delta z/\Delta Y = \lambda_z$. Regressing log-odds price changes on signed order flow provides a direct test. This prediction is shared with any binary Bayesian model, including Glosten-Milgrom and LMSR (Proposition 9), so it tests the binary Bayesian benchmark rather than the Kyle equilibrium specifically. Data from Palumbo [2025], Becker [2026], and Tsang and Yang [2026] provide the trade-level records needed for estimation.

2. Back-loaded informed trading. The intensity of informed trading (measured by the adverse selection component of the spread or by the PIN statistic) should increase as the resolution date approaches. The continuous-time model predicts a $(T-t)^{-1/2}$ blowup rate

(Proposition 24). This prediction is specific to the Kyle framework: it requires dynamic order size optimization that Glosten-Milgrom's unit-trade structure cannot generate. Tsang and Yang [2026] document declining Kyle's λ as Polymarket's election market matures, consistent with back-loaded informed trading concentrating near resolution.

3. Order size distribution. The Kyle model predicts a continuous distribution of informed order sizes centered at $\pm\sigma_u$, mixed with Gaussian noise. The Glosten-Milgrom model predicts unit trades. The empirical distribution of order sizes in prediction markets (conditional on estimated informativeness) can distinguish the two frameworks.

4. Fee distortion at intermediate prices. Kalshi's $f \cdot p(1-p)$ fee reduces informed trading most at $p \approx 1/2$ (Proposition 12). Comparing the speed of information incorporation at different price levels on Kalshi (which uses $p(1-p)$ fees) versus Polymarket (which uses flat fees) provides a natural experiment. If the model is correct, Kalshi should show relatively slower information incorporation at intermediate prices.

5. Position limit threshold. Position limits below approximately σ_u bind and reduce information aggregation (Proposition 10). PredictIt's $850 cap should be binding in high-volume markets, producing slower information incorporation relative to Kalshi or Polymarket markets with higher or no limits.

6. Market maker spreads at intermediate prices. Risk-averse market makers face maximum terminal risk at $p = 1/2$ (Proposition 11). Bid-ask spreads should be widest at intermediate prices and narrowest near the boundaries, controlling for other factors.

5.2 Market design implications

Fee structure. Kalshi's $p(1-p)$ fee loads cost at $p = 1/2$, where information aggregation is most valuable. A flatter fee structure would reduce the distortion at intermediate prices, improving price discovery. A two-part fee (a flat per-trade component plus a small $p(1-p)$ component) balances operational cost recovery with information aggregation.

Market maker compensation. Binary market making concentrates risk at terminal resolution, not at continuous inventory management. Palumbo [2025] documents that Kalshi market makers behave like underwriters. Platforms should design maker programs that compensate for terminal risk (rebates conditional on being on the losing side at resolution) rather than standard maker-taker fee structures designed for equity markets.

Position limits. The model provides a calibration formula: position limits below $k^*\sigma_u/2 = \sigma_u$ bind. Estimating σ_u from observed noise trading volume gives a concrete threshold for each market.

5.3 Relationship to existing results

The one-round model nests two benchmarks. When $p_0 \to 1/2$ with large σ_u, the logistic pricing rule approximates a linear rule ($p_1 \approx 1/2 + (\lambda_z/4)Y$ for small Y), and the model locally resembles Gaussian Kyle. When the position limit $\bar{Q} \to \infty$ or the fee $f \to 0$, the unconstrained frictionless equilibrium is recovered.

The continuous-time model extends Back [1992] to binary payoffs. The Brownian bridge (Gaussian) and Wright-Fisher diffusion (binary) share a common structure: $dp_t = \sigma(t) \cdot g(p_t)\, dB_t$, with $g(p) = 1$ in the Gaussian case and $g(p) = \sqrt{p(1-p)}$ in the binary case. The Gaussian model has a state-independent $\sigma(t)$; the binary model requires state-dependent $\alpha(p,t)$ (Proposition 20).

Proposition 9 clarifies that the $p(1-p)$ price impact structure is common to all binary Bayesian models. What the Kyle framework contributes beyond Bayesian updating is: (i) the endogenous order size $k^* = 2$; (ii) back-loading dynamics; and (iii) the continuous-time Wright-Fisher characterization.

These results complement the existence and uniqueness results for the Kyle-Back model with general distributions. The binary case admits a Markovian equilibrium [Cetin and Danilova, 2025]. The Markovian equilibrium is logistic, back-loaded, and requires a price-dependent strategy spread. The economic content that the Kyle equilibrium adds beyond the Bayesian pricing rule is threefold: (i) the equilibrium order size $k^* = 2$ determines the rate of information aggregation per unit of noise; (ii) the dynamic optimization produces back-loading, a qualitative prediction about the timing of informed trading; and (iii) the state-dependent strategy spread (Proposition 20) shows that the Gaussian model's elegant simplicity is an artifact of unbounded support, not a general property of insider trading equilibria.

5.4 Limitations

Multiple insiders: back-loading versus front-loading. Section 4.8 extends the one-round model to N_{ins} symmetric insiders. The logistic pricing rule, the $p(1-p)$ price impact structure, and the impossibility of state-independent strategies (Proposition 20) all survive competition. Competition dissipates rents (individual profit vanishes faster than $1/N_{\text{ins}}$) and accelerates information aggregation, mirroring Holden and Subrahmanyam [1992]. The positive externality of waiting (declining $p(1-p)$ benefits all insiders) is a mechanism unique to the binary model that has no Gaussian analogue. However, the information-stealing motive dominates the positive externality: with two insiders, competition produces front-loading rather than back-loading (Proposition 28). The transition from back-loading to front-loading as the number of insiders increases from one to two is a sharp, testable prediction. Markets dominated by a single large informed trader should exhibit increasing trading intensity near resolution; markets with multiple informed participants should exhibit the opposite pattern.

Analytical status of back-loading. The back-loading proof for $N = 2$ proceeds in two analytical steps and one numerical step. The analytical steps reduce the condition to $V'(0) + V''(0) < 0$ (via Stein's lemma) and express $V'(0)$ and $V''(0)$ in closed form as functions of the moments $A = \mathbb{E}[S(1 - S)]$, $D = \mathbb{E}[S(1 - S)(1 - 2S)]$, and the IFT sensitivity δ (via the envelope theorem). The numerical step evaluates these moments by Gauss-Hermite quadrature, yielding $V'(0) + V''(0) = -0.116\tilde{\sigma} < 0$ (Appendix A.3). The moments involve expectations of smooth bounded functions of a standard normal, so the quadrature error is below 10^{-10} and does not affect the sign. An analytical bound (without quadrature) would require showing $2D + \delta(k^2 D - 2A) > 2A$, which involves bounding ratios of logistic-normal integrals. For general N, back-loading holds for all N from 2 through 32 in computational experiments (Section 4.7), but a formal proof for arbitrary N requires solving the full backward induction.

Partial PDE solution. The continuous-time equilibrium is characterized by a PDE (Proposition 22). Back [1992] derived and solved the analogous ODE for the Gaussian case, yielding a closed-form equilibrium. The binary PDE is more complex due to the state-dependence of α, and the full time-dependent PDE remains unsolved. The terminal-layer reduction (Proposition 24) yields a nonlinear ODE for $A(p)$, which is solved numerically as a boundary value problem with $A(0) = A(1) = 1$ and $A(1/2) = 1.81$ (Table 1). The N-period computational results give a rescaled amplitude of approximately 1.65 at $N = 32$; the gap reflects finite-N corrections to the self-similar scaling. The full time-dependent solution $\alpha(p, t)$ away from the terminal layer remains open.

Binary payoff only. The model applies to contracts with exactly two terminal values. Many prediction markets offer contracts with multiple outcomes (e.g., which candidate wins a primary with $n > 2$ candidates). Extending the binary model to n-ary payoffs introduces a vector-valued state and a higher-dimensional filtering problem.

Risk-neutral agents. All agents are risk-neutral. Risk aversion among informed traders would reduce equilibrium order sizes and could interact with the state-dependent price impact to alter the back-loading result. The market maker viability result (Proposition 11) provides a partial treatment of risk aversion on the market-making side.

6 Conclusion

The Kyle (1985) insider trading model with binary terminal payoffs has a logistic equilibrium pricing rule with state-dependent price impact $\lambda(p) = \lambda_z \cdot p(1-p)$. The informed trader back-loads trading toward resolution, and the continuous-time equilibrium requires a state-dependent strategy spread, unlike the Gaussian benchmark where the strategy spread depends only on time. The model's predictions (state-dependent price impact, back-loaded informed trading, and fee distortions concentrated at intermediate prices) are testable in data from Kalshi, Polymarket, and PredictIt.

A Proofs

A.1 The Girsanov identity

Lemma 29 (Moment Identity)**.** *For $S = \sigma(-k^2/2 - kZ)$ with $Z \sim N(0,1)$ and any $k > 0$:*

$$\mathbb{E}[S(1-S)] = \frac{\mathbb{E}[S]}{2}.$$

Proof. Let $W = -k^2/2 - kZ$ with $Z \sim N(0,1)$. Using $\sigma(w) = 1/(1+e^{-w})$, we have $S = 1/(1+e^{-W})$ and

$$S(1-S) = \sigma'(W) = \frac{e^{-W}}{(1+e^{-W})^2}.$$

Apply the substitution $z \to -z-k$ to the expectation $\mathbb{E}[S(1-S)] = \int \frac{e^{-W}}{(1+e^{-W})^2}\phi(z)\,dz$. Under this map, $W = -k^2/2 - kz$ transforms to $-W := k^2/2 + kz$, and the Gaussian density picks up $e^{-kz-k^2/2} = e^W$. Therefore:

$$\mathbb{E}[S(1-S)] = \int \frac{e^W}{(1+e^W)^2} \cdot e^W \cdot \phi(z)\,dz = \mathbb{E}\left[\frac{e^{2W}}{(1+e^W)^2}\right] = \mathbb{E}[S^2],$$

where the last step uses $S = e^W/(1+e^W)$, so $S^2 = e^{2W}/(1+e^W)^2$.

Since $\mathbb{E}[S(1-S)] = \mathbb{E}[S] - \mathbb{E}[S^2]$ and $\mathbb{E}[S(1-S)] = \mathbb{E}[S^2]$, we obtain $\mathbb{E}[S^2] = \mathbb{E}[S]/2$ and hence $\mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$. □

Remark 30. The identity reduces the equilibrium equation $G(k) = \mathbb{E}[S] - (k^2/2)\mathbb{E}[S(1-S)]$ to $G(k) = \mathbb{E}[S](1 - k^2/4)$, which has the unique positive zero $k = 2$. The derivative $G'(2) = -\mathbb{E}[S] < 0$ confirms uniqueness with a one-line argument.

A.2 Proof of Proposition 15(c): $V'(0) < 0$

Proof. Define $V(z)$ as the $v = 1$ informed trader's profit in the one-round game at prior log-odds z with noise $\tilde{\sigma}$.

Symmetry relations. At prior z, the equilibrium satisfies $x_1^*(z) = -x_0^*(-z)$ and $x_0^*(z) = -x_1^*(-z)$ (swapping $v = 0 \leftrightarrow v = 1$ and $z \leftrightarrow -z$). Differentiating at $z = 0$:

$$\frac{dx_1^*}{dz}(0) = \frac{dx_0^*}{dz}(0) =: \delta.$$

This relation implies $d\lambda_z/dz|_0 = 0$ (the first-order perturbation of the strategy spread vanishes) and $dc/dz|_0 = -2a^*\delta/\tilde{\sigma}^2$.

Envelope theorem. $V(z) = \pi(x_1^*(z); z, \lambda_z(z), c(z))$ with $\pi(x; z, \lambda, c) = x\,\mathbb{E}[\sigma(-(z + \lambda(x+u) + c))]$. Since $\partial\pi/\partial x = 0$ at $x = x_1^*$:

$$V'(0) = \left.\frac{\partial \pi}{\partial z}\right|_0 + \left.\frac{\partial \pi}{\partial c}\right|_0 \cdot \left.\frac{dc}{dz}\right|_0.$$

With $S = \sigma(-k^{*2}/2 - k^*Z)$, $A := \mathbb{E}[S(1-S)] = \mathbb{E}[S]/2$, and $D := \mathbb{E}[S(1-S)(1-2S)]$:

$$\left.\frac{\partial \pi}{\partial z}\right|_0 = -a^*A, \qquad \left.\frac{\partial \pi}{\partial c}\right|_0 = -a^*A, \qquad \left.\frac{dc}{dz}\right|_0 = -\frac{2a^*\delta}{\tilde{\sigma}^2}.$$

Therefore $V'(0) = -a^*A + (k^{*2}/2)A\delta$.

Computing δ. Linearizing the first-order condition around $z = 0$ (Appendix A.4):

$$0 = -A - \frac{k^*\delta}{\tilde{\sigma}}A + \frac{k^{*2}}{2}D,$$

giving $\delta = (\tilde{\sigma}/k^*)(k^{*2}D/(2A) - 1)$.

Substituting:

$$V'(0) = a^*\left(\frac{k^{*2}D}{2} - 2A\right).$$

Sign verification. At $k^* = 2$, $A = \mathbb{E}[S]/2 = 0.1093$ and $D = 0.0571$ (by Gauss-Hermite quadrature, cross-checked by Stein's lemma). Then $k^{*2}D/2 = 2 \times 0.0571 = 0.1142$ and $2A = 0.2186$. Since $0.1142 < 0.2186$, $V'(0) < 0$. □

A.3 Proof that $V'(0) + V''(0) < 0$

Proof. The back-loading condition requires $V'(0)+V''(0) < 0$. From Appendix A.2, $V'(0) = a^*(k^{*2}D/2 - 2A)$. At $k^* = 2$ with $a^* = \tilde{\sigma}$: $V'(0)/\tilde{\sigma} = 2D - 2A = 2(D - A)$.

The second derivative $V''(0)$ follows from differentiating the envelope theorem expression a second time. The computation uses the symmetry $d\lambda_z/dz|_0 = 0$, the chain rule through the equilibrium system, and the identity $A = \mathbb{E}[S]/2$.

Evaluating by 20-node Gauss-Hermite quadrature at $k^* = 2$:

$$V'(0)/\tilde{\sigma} = -0.104, \qquad V''(0)/\tilde{\sigma} = -0.012.$$

Therefore $V'(0) + V''(0) = -0.116\tilde{\sigma} < 0$.

The dominant contribution is $V'(0)$: the value function is steeply decreasing in the prior log-odds at $z = 0$. The informed trader's edge shrinks as the prior moves toward the truth ($z > 0$ under $v = 1$). The second derivative $V''(0) < 0$ means the value function is concave at $z = 0$: the loss from the prior moving toward the truth accelerates. Both forces favor deferring trading to later rounds. □

A.4 Linearization of the first-order condition

At prior $z = \varepsilon$ (small), the equilibrium perturbs to $x_1 = a^* + \delta\varepsilon$, $x_0 = -a^* + \delta\varepsilon$. The pricing rule argument for $v = 1$ at order x_1 becomes

$$w_1 = \varepsilon + \frac{2a^*}{\tilde{\sigma}^2}(a^* + \delta\varepsilon + u) - \frac{2a^*\delta\varepsilon}{\tilde{\sigma}^2} + O(\varepsilon^2) = \varepsilon + \frac{k^{*2}}{2} + k^* Z + O(\varepsilon^2),$$

where the $\delta\varepsilon$ terms from λ_z and c cancel at first order.

The first-order condition for $v = 1$ at order ε:

$$0 = \mathbb{E}[S_\varepsilon] - (a^* + \delta\varepsilon)\frac{k^*}{\tilde{\sigma}}\mathbb{E}[S_\varepsilon(1 - S_\varepsilon)],$$

where $S_\varepsilon = \sigma(-\varepsilon - k^{*2}/2 - k^* Z)$. Expanding: $\mathbb{E}[S_\varepsilon] \approx \mathbb{E}[S] - \varepsilon A$ and $\mathbb{E}[S_\varepsilon(1-S_\varepsilon)] \approx A - \varepsilon D$. Using $a^*(k^*/\tilde{\sigma}) = k^{*2}/2$ and $G(k^*) = 0$:

$$0 = -A\varepsilon - \frac{\delta\varepsilon k^*}{\tilde{\sigma}}A + \frac{k^{*2}}{2}D\varepsilon + O(\varepsilon^2).$$

Dividing by ε: $0 = -A - (k^*\delta/\tilde{\sigma})A + k^{*2}D/2$, giving $\delta = (\tilde{\sigma}/k^*)(k^{*2}D/(2A) - 1)$.

References

Avellaneda, Marco and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.

Back, Kerry. Insider trading in continuous time. *Review of Financial Studies*, 5(3):387–409, 1992.

Becker, Jonathan. The microstructure of wealth transfer in prediction markets. Working paper, January 2026.

Bürgi, Constantin, Yizhe Deng, and Karl Whelan. Makers and takers: The economics of the Kalshi prediction market. UCD Working Paper WP2025/19, 2025.

Glosten, Lawrence R. and Paul R. Milgrom. Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.

Grossman, Sanford J. and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

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

Holden, Craig W. and Avanidhar Subrahmanyam. Long-lived private information and imperfect competition. *Journal of Finance*, 47(1):247–270, 1992.

Keller, Tobias and Kevin Tseng. Arrow-Debreu meets Kyle: Price discovery across derivatives. arXiv:2302.13426, 2025.

Kyle, Albert S. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

Manski, Charles F. Interpreting the predictions of prediction markets. Working paper, 2004.

Milgrom, Paul and Nancy Stokey. Information, trade, and common knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.

Ng, Serena, Lin Peng, Yinan Tao, and Guofu Zhou. Price discovery and trading in modern prediction markets. SSRN 5331995, 2025.

Palumbo, Nicholas. A microstructure perspective on prediction markets. SSRN 6325658, 2025.

Snowberg, Erik and Justin Wolfers. Explaining the favorite–long shot bias: Is it risk-love or misperceptions? *Journal of Political Economy*, 118(4):723–746, 2010.

Tsang, Kwun Yee and Haoyang Yang. The anatomy of Polymarket: Evidence from the 2024 presidential election. SSRN 6336679, 2026.

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

Wolfers, Justin and Eric Zitzewitz. Interpreting prediction market prices as probabilities. NBER Working Paper 12200, 2006.

Cetin, Umut and Albina Danilova. A new approach for the continuous time Kyle-Back strategic insider equilibrium problem. arXiv:2506.12281, 2025.

Dalen, Shaw. Toward Black-Scholes for prediction markets: a unified kernel and market-maker's handbook. arXiv:2510.15205, 2025.