Prediction Markets as Data Marketplaces: The Price of Information under Automated Market Making*

Ping Liu† Zhan Pang‡

October 2026

Abstract

We study prediction markets as data marketplaces in which trading produces valuable forecast data. An informed trader or a noise trader submits an order to an automated market maker using the logarithmic market scoring rule. A Bayesian arbitrageur observes the order but not trader type and corrects the price. Liquidity equals the informed trader's expected trading profit per unit of acquired information and the market maker's expected net payout per unit of information reflected in the final price. Because informed orders scale with liquidity while noise orders do not, liquidity affects how well the arbitrageur can distinguish the two types. Greater liquidity therefore need not deliver more information. We endogenize information acquisition and characterize optimal liquidity and market entry.

*We thank Josh Mollner (discussant) and participants at the 2026 Wabash River Finance Conference for helpful comments.

†Purdue University: liu2554@purdue.edu

‡Purdue University: zpang@purdue.edu

1. Introduction

The expansion of prediction markets has intensified debate over the boundary between federal derivatives regulation and state gaming law, particularly whether contracts on sporting events should be treated differently from those on economic events.[1,2] Beyond their legal classification, an important question is whether these markets produce economically useful information. The Commodity Futures Trading Commission's proposed public-interest assessment recognizes this informational benefit alongside the risks these contracts pose.[3] Our analysis bears on that assessment through one economic function of a prediction market: the production and financing of forecasts that users value. A binary event contract pays one dollar if the event occurs and zero otherwise; its price is commonly interpreted as a market-implied probability that can inform decisions beyond the market itself.

We view prediction markets as data marketplaces: informed traders can profit from their information through trading, while their orders convey information to prices that users can draw upon without taking positions.[4] Because acquiring information is costly, users who

[1]Election contracts have traded on the Iowa Electronic Markets since 1988 (Berg et al., 1997). More recently, combined monthly trading volume on Kalshi and Polymarket increased from less than $5 billion in September 2025 to $53 billion in July 2026, according to Pew Research Center analyses of data compiled by The Block (Radde, 2026a,b).

[2]The Third Circuit held that Kalshi had a reasonable chance of establishing that its event contracts are swaps within the exclusive jurisdiction of the Commodity Exchange Act (KalshiEX, LLC v. Flaherty, 2026). The Ninth and Sixth Circuits held that the contracts before them had not been shown to be swaps and that the Act had not been shown to preempt state gaming law (KalshiEX, LLC v. Assad, 2026; KalshiEX LLC v. Schuler, 2026). A petition for certiorari from the Third Circuit's decision is pending (Flaherty v. KalshiEX, LLC, 2026).

[3]The proposal considers hedging or price-basing utility and economically useful information alongside risks involving manipulation, settlement integrity, and insider information (Commodity Futures Trading Commission, 2026). Its discussion of price discovery and information aggregation appears on pp. 35829–35830.

[4]Federal Reserve staff use contracts on macroeconomic releases to track expectations (Diercks et al., 2026), and Intercontinental Exchange distributes normalized market probabilities to institutional investors (Intercontinental Exchange, 2026). Market questions, outcomes, and crowd forecasts also provide data for training and evaluating forecasting models (Halawi et al., 2024). Appendix E provides further institutional background.

value informative prices may be willing to subsidize informed trading. We examine how an automated market maker provides this subsidy through the trades it executes, and how the sequential interaction of informed trading, noise trading, and arbitrage determines the information ultimately reflected in prices. Arbitrageurs infer information from earlier orders without observing whether those orders are informed. The informational value of the final price therefore depends on how well informed orders can be distinguished from noise, as well as on traders' incentives to acquire information.

Subsidizing information acquisition is particularly important in thin markets, where demand for a forecast need not translate into demand to trade. Those who value the forecast may be unwilling to take positions on the event, and many listed contracts attract little or no trading.[5] In an order book, the return to an informed trade depends on the quotes supplied by other traders. A funded automated market maker can instead commit to a pricing rule before information is acquired and take the other side of each order. Users who value the forecast can therefore subsidize information acquisition by funding the market maker without trading themselves.[6]

In this paper, we study an automated market maker using the logarithmic market scoring rule (LMSR), which is well suited to data marketplaces whose objective is information production. Like other funded automated market makers, it quotes executable prices from the opening. Its bounded loss allows a sponsor to back these quotes with a finite reserve. Strict convexity of the LMSR trading cost in the net order makes truthful revelation optimal: a risk-neutral trader who trades once uniquely maximizes expected profit by moving the price to its posterior. The market maker can thus elicit a forecast without observing the trader's

[5]Trading is thin even at the largest venues, Kalshi and Polymarket: Chetia (2026) reports that roughly 70 percent of all Polymarket's closed markets from 2021 through the end of May 2026 had under $10,000 in reported volume, that over 45,000, nearly 5 percent, had none, and that Kalshi had many shallow markets.

[6]Kalshi runs an order book, while Polymarket began with an automated market maker (Commodity Futures Trading Commission, 2022) and then adopted an order book (Polymarket, 2026a).

underlying data or model. The LMSR's particular advantage for our analysis is that it links trading profits linearly to Shannon information. The informed trader's expected trading profit equals liquidity times the expected Kullback–Leibler divergence of its posterior from the common prior. This expected divergence measures both the mutual information between the signal and settlement and the reduction in expected log-loss from forecasting with the posterior rather than the prior, so the market maker pays for information in proportion to the improvement in forecast accuracy. Liquidity is therefore the price paid per unit of acquired information, allowing us to summarize the trader's acquisition response by its information supply.[7] The LMSR has also been used in sponsored markets, in which the party funding the market maker sets its liquidity and pays for the resulting forecast.[8]

To characterize the price of information, we consider a market maker that opens a binary LMSR market at the common prior and commits to liquidity before trading. The first order comes from a risk-neutral informed trader with a given probability and otherwise from a noise trader whose order contains no information about the event and whose order distribution is independent of liquidity. A risk-neutral arbitrageur then observes the order but not the trader's type. It has no private information about the event and no external benchmark for its probability, so it updates its belief from the order alone, accounting for the possibility that the observed order is noise. The arbitrageur maximizes its expected trading profit by moving the price to its posterior probability of the event, improving the accuracy of the final price as a forecast. The market maker chooses liquidity to maximize the value of this

[7]Under constant-product rules, trading rewards take the form of a different Bregman divergence and lack this linear relation to Shannon information, making the acquisition problem less tractable.

[8]Inkling Markets, Consensus Point, Yahoo!, Microsoft, and Carnegie Mellon's Gates Hillman market have run the LMSR with a liquidity set before trading began (Othman et al., 2013); at the Gates Hillman market, designed to predict the opening day of the Gates and Hillman Centers, the administrator fixed the liquidity and funded the prizes through grants (Othman and Sandholm, 2013). Gnosis implemented the LMSR on chain, where the owner's funding sets the liquidity and can be changed only while trading is paused, and described that funding as a subsidy paid by a party interested in a prediction on the event and expecting only public information in return (Gnosis, 2020; Koeppelmann, 2020). Appendix E collects the record.

forecast accuracy net of the costs of operating the market.

Liquidity equals the informed trader's expected trading profit per unit of acquired information, conditional on trading. The arbitrageur's subsequent trade does not affect this profit. The market maker's expected net payout to the first-round trader and the arbitrageur likewise equals liquidity per unit of delivered information. Uncertainty about trader type, however, limits how much of the acquired information the final price reflects. When a nonzero order could come from either type, the arbitrageur chooses a final price between the first-round price and the prior. If the order is informed, the final price is closer to the prior than the informed trader's belief; if it is noise, the final price still differs from the prior. Both discrepancies contribute to a type-confusion loss. Delivered information therefore equals acquired information multiplied by the probability of informed arrival, less this loss. If trader type is observable, the arbitrageur leaves informed orders unchanged and fully reverses noise orders, eliminating the loss.

Liquidity affects delivered information through both acquisition and the type-confusion loss. Informed orders scale with liquidity, whereas the noise-order distribution remains fixed. Holding acquisition fixed, the loss vanishes as liquidity tends to zero or infinity because nonzero informed orders become distinguishable from noise. At intermediate liquidity levels, the arbitrageur may be unable to distinguish informed orders from noise, resulting in a positive type-confusion loss. Therefore, delivered information can be nonmonotone in liquidity. Greater liquidity incentivizes the informed trader to acquire more information, but it can also increase the type-confusion loss, so the final price need not become more informative.

Whether type confusion arises also depends on the information acquisition technology. We study Bernoulli sampling, in which the informed trader chooses sample size, and a Gaussian measurement, in which it chooses precision. Bernoulli sampling generates finitely many possible informed orders. If noise orders assign zero probability to each individual

order, trader type is identifiable and the type-confusion loss vanishes. At positive Gaussian precision, informed orders have a positive density at every order size and therefore overlap with noise orders drawn from a distribution with a probability density. The arbitrageur cannot perfectly distinguish trader types, resulting in a positive type-confusion loss.

We then characterize the market maker's optimal liquidity and its decision to open the market. Opening is profitable when, at some liquidity, the value of a unit of delivered information exceeds its total cost, including trading payouts, settlement overhead, and the capital charge. A higher valuation of accuracy leads the market maker to choose more delivered information, while a higher capital charge leads it to supply less liquidity. When the capital charge is sufficiently large relative to settlement overhead, a higher valuation of accuracy also increases optimal liquidity. Disclosure of trader type raises the market maker's optimized profit and lowers the valuation of accuracy required for profitable entry. It can therefore make a market profitable that would remain closed when trader type is hidden.

We provide a microfoundation for the market maker's accuracy weight through a contract with a principal who purchases the forecast. The principal pays the market maker in proportion to the improvement in the final price's logarithmic score over the opening price. The accuracy weight thus represents the contractual payment per unit of delivered information. We then study how a planner or a private buyer chooses this payment rate, anticipating the market maker's liquidity choice and the informed trader's acquisition response.

More generally, our analysis informs the policy debate over prediction markets by examining how demand for forecasts can finance costly information acquisition. Users pay for the accuracy of the forecast, and the market maker compensates informed traders through their trading profits. How much information the final price delivers then depends on the incentive to research the event, on how often informed traders arrive, and on how well their orders can be told from noise. Contracts on sporting events and contracts on economic events may

differ along each of these dimensions, but the category of the event does not by itself fix the informational value of a contract. Our results therefore bear on the informational benefit that the proposed public-interest assessment recognizes by identifying when a market produces a useful forecast and at what cost.

Section 2 sets up the model. Section 3 characterizes trading with an informed trader, a noise trader, and an arbitrageur at a given liquidity and relates delivered information to the market maker's net payout. Section 4 endogenizes information acquisition, develops the Bernoulli and Gaussian acquisition technologies, and illustrates trading over the market's life. Section 5 characterizes optimal liquidity and provides a contractual foundation for the market maker's valuation of forecast accuracy. Section 6 relates the results to the regulatory debate over prediction markets. Proofs, the specification of the multi-period market, and institutional background are in the appendix.

1.1 Related Literature

This paper contributes to three strands of literature. The first studies automated market makers built on scoring rules. Hanson (2003) introduces market scoring rules, which reward improvements in forecasts, and Abernethy et al. (2013) and Chen and Vaughan (2010) develop their cost-function representations and connections to online learning. Chen et al. (2010) show that truthful trading under the LMSR depends on the information structure. Ostrovsky (2012) establishes information aggregation for separable securities in both Kyle-type markets and market scoring rules. These studies take traders' information as given and study its revelation at a fixed liquidity. Othman et al. (2013) allow liquidity to increase with shares sold, enabling the market maker to earn a profit. We contribute to this literature by placing the LMSR in a sequential-trade market in which an informed trader, a noise trader, and an arbitrageur interact, and we characterize how acquired information, the information

the final price delivers, and the market maker's payment are related.

We also contribute to the microstructure literature on trading under asymmetric information and its extension to automated market makers. In Glosten and Milgrom (1985), a specialist who cannot tell whether the next customer is informed sets bid and ask prices at conditional expectations, so adverse selection opens a spread and transaction prices convey information. In Kyle (1985), an informed trader submits orders alongside noise traders, and competitive market makers price the order flow, so private information enters the price gradually. In both models, the party that supplies liquidity also infers from the orders. Research on automated market makers in decentralized finance separates the two roles: a pool executes orders by a fixed formula, while arbitrageurs align pool prices with externally observed values, imposing losses on liquidity providers (Capponi and Jia, 2025; Lehar and Parlour, 2025). Our market maker likewise commits to a pricing rule. Prediction markets, however, often lack a readily observable external benchmark. We therefore model an arbitrageur who infers the event probability from an order without observing the trader's type. We examine how liquidity affects both information acquisition and the information conveyed by the final price. Greater liquidity encourages information acquisition but also changes the overlap between informed and noise orders, so it need not improve price discovery. The market maker acts as a monopsonistic buyer of information, choosing liquidity to balance the value of the resulting forecast against its expected payments.

Finally, our paper relates to the literature on payment mechanisms for information acquisition. Chen and Yu (2021) and Li et al. (2022) design scoring rules to encourage costly information acquisition; Bergemann et al. (2018) study the design and sale of information; and Grossman and Stiglitz (1980) show that fully revealing prices cannot sustain costly acquisition. We contribute to this literature by using a sequential trading model à la Glosten–Milgrom to study how an automated market maker finances information acquisition through

liquidity in a prediction market, holding the logarithmic scoring rule fixed.

2. Model setup

A market maker opens a market on a binary event using the logarithmic market scoring rule (LMSR) of Hanson (2003) and commits to its liquidity parameter before trading. The initial order is submitted either by an informed trader with private information about the event or by a noise trader whose order carries no information.[9] An arbitrageur observes the order but not the trader's type. It then trades to align the price with its Bayesian posterior probability of the event.

A binary event has settlement indicator $B \in \{0, 1\}$. Agents share a common prior $\delta_0 = \Pr(B = 1) \in (0, 1)$, which is also the market's initial reference price for the YES contract. A belief $\delta \in [0, 1]$ denotes a probability that $B = 1$ (the prediction-market convention), and δ_0 is its prior value. Throughout, log denotes the natural logarithm and $\operatorname{logit} x := \log \frac{x}{1-x}$ the log-odds of $x \in (0, 1)$.

For beliefs $p, q \in (0, 1)$, $\mathrm{KL}(p \,\|\, q) := p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}$ denotes the Bernoulli relative entropy, the divergence of p from q; it is nonnegative and zero only at $p = q$. The binary entropy is $h(x) := -x \log x - (1 - x) \log(1 - x)$.

Let P denote the informed trader's posterior probability that $B = 1$, conditional on its private information. We take the distribution of P as given.

Assumption 1. *The informed trader's belief P takes values in $(0, 1)$ almost surely, and $\Pr(B = 1 \mid P) = P$ almost surely.*

Assumption 1 requires the trader's belief to be calibrated: conditional on $P = p$, the event occurs with probability p. By the law of iterated expectations, posterior beliefs average to

[9] In Section 4, we endogenize the informed trader's costly information acquisition.

the common prior, $\mathbb{E}[P] = \delta_0$, the Bayes-plausibility property.[10]

For liquidity $b > 0$, the market maker posts the LMSR cost function $C(q_1, q_0) = b \log(e^{q_1/b} + e^{q_0/b})$, where q_1, q_0 are the outstanding YES and NO shares. The corresponding scoring-rule potential is b times the negative Shannon entropy,

$$V(\delta) \;=\; b[\delta \log \delta + (1-\delta)\log(1-\delta)], \qquad V'(\delta) = b \operatorname{logit} \delta, \quad V''(\delta) = \frac{b}{\delta(1-\delta)} > 0. \quad (1)$$

The instantaneous YES price is $\partial_{q_1} C = e^{q_1/b}/(e^{q_1/b} + e^{q_0/b}) = \sigma((q_1 - q_0)/b)$ with $\sigma(z) = 1/(1+e^{-z})$; equivalently the price δ and the inventory satisfy $\operatorname{logit} \delta = (q_1 - q_0)/b$.

The order that moves the price from δ_0 to $\delta \in (0,1)$ is $x = b(\operatorname{logit} \delta - \operatorname{logit} \delta_0)$. Each interior target price therefore corresponds to a unique finite order.

3. Trading and information

Trading has two rounds. In the first round, the informed trader submits an order with probability $\vartheta \in [0,1]$; otherwise, the noise trader submits an order. The trader's type is private information.[11] In the second round, the arbitrageur observes the order and trades against the LMSR schedule, starting from the price resulting from that order. Trading then ends, and the contracts settle when the event outcome becomes known.

We now characterize trading when the arbitrageur cannot observe whether an order comes from the informed trader or the noise trader. We take liquidity $b > 0$ and the distribution of the informed trader's belief P as given and determine the final price, the information it reveals, and the resulting trading profits.

[10] The restriction to interior beliefs ensures that the informed trader's optimal order is finite (Lemma 1).
[11] The single-trader market corresponds to $\vartheta = 1$.

3.1 The informed trader

The informed trader holds posterior belief $P \in (0,1)$ and chooses a net order $x \in \mathbb{R}$ of YES shares, starting from an inventory (q_1, q_0) priced at the reference δ_0: it pays the market maker the cost-function increment and collects one unit per YES share if the event occurs. Its objective is the expected settlement value of the position net of that payment,

$$\max_{x \in \mathbb{R}} \; P\,x - [C(q_1 + x, q_0) - C(q_1, q_0)]. \tag{2}$$

A YES share and a NO share together cost one unit and pay one unit at settlement, so any order in the two contracts has the same net payoff as an order in YES shares alone.[12]

The next lemma gives the trader's profit as a function of the target price and characterizes its optimal choice.

Lemma 1 (Truthful trade and its value)**.** *A trader holding posterior belief $P \in (0,1)$ who moves the market price from the reference δ_0 to an interior target $\delta \in (0,1)$ realizes the outcome-contingent profit*

$$b \log \frac{\delta}{\delta_0} \quad \text{if } B = 1, \qquad b \log \frac{1-\delta}{1-\delta_0} \quad \text{if } B = 0.$$

The expected profit under P, $\pi(\delta; P) = P\, b \log \dfrac{\delta}{\delta_0} + (1-P)\, b \log \dfrac{1-\delta}{1-\delta_0}$, is uniquely maximized at the truthful trade $\delta = P$, with maximal value

$$\pi(P; P) \;=\; b\,\mathrm{KL}(P \,\|\, \delta_0) \;\geq\; 0. \tag{3}$$

Proof. Appendix A.1.

[12] Factoring $e^{w/b}$ out of the cost function gives $C(q_1 + w, q_0 + w) = w + C(q_1, q_0)$ for any $w \in \mathbb{R}$, while w units of each leg jointly pay w regardless of B.

LMSR is strictly proper: the trader maximizes its expected profit by moving the price to its own posterior belief. Its surplus $b\,\mathrm{KL}(P \,\|\, \delta_0)$ is the rent it earns for holding a belief the market does not.

The trader earns a finite rent because it is a schedule-taker rather than a price-taker. A risk-neutral trader facing a fixed price $p \neq P$ would demand an unbounded long or short position. Under LMSR, the marginal cost of acquiring YES shares increases with the net order, and the trader chooses a finite position that equates the marginal price to its posterior belief.[13]

Conditional on informed arrival, the event B and the belief P have the joint distribution specified in Section 2. By Lemma 1, the informed trader optimally moves the price to P. The corresponding net order of YES shares is

$$Q_I := b(\operatorname{logit} P - \operatorname{logit} \delta_0).$$

The informed trader does not trade again, and its settlement profit depends only on its own order and the event outcome. The arbitrageur's subsequent trade therefore does not alter the informed trader's incentive to trade truthfully.

Let $y := \mathbb{E}[\mathrm{KL}(P \,\|\, \delta_0)]$ denote the informed trader's information. By Lemma 1, its expected trading profit conditional on arrival is $b\,y$.

3.2 The noise trader

The noise trader submits an exogenous net order U of YES shares, with cumulative distribution function F_U.[14] Let Q denote the first-round net order: $Q = Q_I$ if the informed trader

[13] The profit formulas in Lemma 1 also apply when the trader finds the market at a price $\delta' \in (0, 1)$, with δ' replacing δ_0, since the proof in Appendix A.1 uses δ_0 only as the price at which the trader finds the market. The optimal target remains P.

[14] We allow F_U to have point masses and impose no moment restrictions on U.

arrives and $Q = U$ if the noise trader arrives.

Assumption 2. *Conditional on noise arrival, U has distribution F_U and $\Pr(B = 1 \mid U) = \delta_0$ almost surely. The distribution F_U does not depend on liquidity b.*

The noise order contains no information about the event: an observer who knew that the order came from the noise trader would retain the prior belief δ_0. Moreover, noise orders are drawn from the same distribution at every liquidity level, whereas the informed trader adjusts its order to the LMSR schedule.

After a first-round order q, the LMSR sets the first-round price

$$\delta_1(q) := \sigma(\operatorname{logit} \delta_0 + q/b),$$

which is strictly increasing in q. An informed order moves this price to P, so $\delta_1(Q_I) = P$. A noise order can move the first-round price even though it contains no information about the event.

Conditional on a noise order $U = u$, the event probability remains δ_0. The profit formula in Lemma 1 therefore gives the noise trader an expected loss of $b \operatorname{KL}(\delta_0 \,\|\, \delta_1(u))$. Averaging over the noise-order distribution F_U, define the expected displacement divergence

$$\nu_b := \mathbb{E}[\operatorname{KL}(\delta_0 \,\|\, \delta_1(U))].$$

Thus, $b\,\nu_b$ is the noise trader's expected loss conditional on arrival.

Liquidity and the noise trader's loss. When $\mathbb{E}|U| < \infty$, higher liquidity reduces the price impact of a given noise order, so the noise trader's expected loss $b\,\nu_b$ falls as b rises and vanishes as $b \to \infty$. As $b \downarrow 0$, the average execution price of a fixed order approaches 1 for

a purchase and 0 for a sale, so the noise trader's loss per YES share approaches $1 - \delta_0$ on a purchase and δ_0 on a sale (Lemma 3 in Appendix A.2).

3.3 The arbitrageur

The arbitrageur is risk neutral and has no private information about B. It observes q but not the trader's type and uses the order to update its belief about the event.[15] Let $\Psi(q) := \Pr(\text{informed arrival} \mid Q = q)$ denote its posterior probability that the order is informed. When Q_I and U have probability density functions $f_{I,b}$ and f_U, respectively, Bayes' rule gives

$$\Psi(q) = \frac{\vartheta f_{I,b}(q)}{\vartheta f_{I,b}(q) + (1-\vartheta) f_U(q)}. \tag{4}$$

This expression applies wherever the denominator is positive.[16]

If the order is informed, it reveals the belief $P = \delta_1(q)$; if it is noise, the event probability remains δ_0. The arbitrageur's posterior probability of the event is therefore

$$\delta_A(q) := \Pr(B = 1 \mid Q = q) = \Psi(q)\delta_1(q) + (1 - \Psi(q))\delta_0. \tag{5}$$

The arbitrageur thus weights the information implied by an informed order by the probability that the order is informed.

By Lemma 1, applied at the starting price $\delta_1(q)$ (footnote 13), the arbitrageur optimally moves the price to $\delta_A(q)$. Its net order of YES shares is

$$x_A(q) = b(\operatorname{logit} \delta_A(q) - \operatorname{logit} \delta_1(q)),$$

[15] The arbitrageur knows δ_0, b, ϑ, the distribution of informed beliefs, the noise-order distribution F_U, and the informed trader's optimal order as a function of P.

[16] Appendix A.3 defines Ψ for general order distributions, including distributions with point masses, and establishes that the prices are interior and the orders are finite almost surely.

and its conditional expected profit is $b\,\mathrm{KL}(\delta_A(q) \,\|\, \delta_1(q)) \geq 0$.[17] The final price is $\delta_A(q)$. The market maker executes both orders according to the LMSR schedule.

Define the *arbitrageur's rent* per unit of liquidity as

$$D_A(b) := \mathbb{E}[\mathrm{KL}(\delta_A(Q) \,\|\, \delta_1(Q))],$$

where the expectation is over the distribution of first-round orders.[18]

The *type-confusion loss* measures the information lost because the arbitrageur cannot observe the first-round trader's type:

$$\mathscr{L}_T(b) \;:=\; \vartheta\,\mathbb{E}[\mathrm{KL}(P \,\|\, \delta_A(Q_I))] + (1-\vartheta)\,\mathbb{E}[\mathrm{KL}(\delta_0 \,\|\, \delta_A(U))]. \tag{6}$$

The two terms compare the final price with the event probability conditional on both the order and the trader's type: P for an informed order and δ_0 for a noise order.

The final price lies between the prior and the first-round price. The arbitrageur therefore trades in the opposite direction to the first order, with an order no larger in absolute value. It leaves an order that it identifies as informed unchanged and fully reverses an order that it identifies as noise. Uncertainty about trader type limits this correction: reversing a possible noise order also risks reversing an informative price change.

The next lemma relates the arbitrageur's rent to the expected displacement divergence and the type-confusion loss.

Lemma 2 (The arbitrageur's rent)**.** *Under Assumptions 1 and 2, at any liquidity $b > 0$,*

[17] The arbitrageur's profit is nonnegative in conditional expectation; it need not be nonnegative in each settlement state.

[18] Both ν_b and $D_A(b)$ may be infinite. The identities in Lemma 2 and Proposition 1(iv) remain valid in that case. Arrival-weighted terms corresponding to a trader type with zero arrival probability are taken to be zero.

$D_A(b) = (1 - \vartheta)\,\nu_b - \mathscr{L}_T(b)$.

Proof. Appendix A.5.

The arbitrageur earns a profit by correcting price changes induced by noise orders, but uncertainty about trader type limits its ability to recover the noise trader's losses.

Relation to the literature. The model adopts a sequential trading structure with unobserved trader types, following Glosten and Milgrom (1985). In their model, the specialist observes trades but not traders' information, and sets quotes accounting for the possibility that an order reflects liquidity demand rather than private information. In the present model, the market maker executes orders according to the committed LMSR schedule, while the arbitrageur infers the event probability from the observed order. Arbitrage following liquidity trades also appears in the automated market maker model of Lehar and Parlour (2025). Their constant-product market maker supplies liquidity according to a predetermined pricing rule, and an arbitrageur corrects price deviations caused by liquidity trades using an observed fundamental value as a benchmark. This assumption reflects cryptocurrency trading across decentralized and centralized exchanges. Liquid centralized exchanges provide reference prices for tokens traded in such pools, allowing arbitrageurs to exploit price differences by trading the same tokens on the centralized exchange and against the pool.

Prediction-market contracts often lack such an external benchmark. An event contract has no traded underlying asset, and the event outcome becomes verifiable only at settlement; before then, its probability must be inferred from the available information.[19] The model therefore assumes that the arbitrageur has neither a private signal nor an external price benchmark. It observes the first-round order but not whether the trader is informed or

[19]When a deep market for the same contract exists elsewhere, its price can play the role that the centralized-exchange price plays in Lehar and Parlour (2025); we focus on the contracts for which no such market exists.

noise. Thus, its belief in (5) is a weighted average of the posterior implied by an informed order and the prior, with weights determined by the posterior probability of informed arrival in (4). As in Lehar and Parlour (2025), the arbitrageur trades against the automated market maker to correct the price, but here the correction rests on an inference about trader type rather than on an independently observed value.

3.4 Delivered information and the market maker's payout

In the trading equilibrium, the informed trader and the arbitrageur maximize their respective conditional expected profits, and the arbitrageur's beliefs follow Bayes' rule under the informed trader's strategy and the exogenous noise-order distribution. In Section 4, we endogenize the informed trader's information acquisition: the arbitrageur uses the distribution of informed orders implied by the informed trader's equilibrium information acquisition choice to update its beliefs after observing the first-round order.

Define *delivered information* as the expected divergence of the final price from the prior,

$$\mathcal{I}(b) \;:=\; \mathbb{E}[\mathrm{KL}(\delta_A(Q) \,\|\, \delta_0)]. \tag{7}$$

The *expected log-loss* of a forecast $t(Q) \in (0,1)$ is $\mathbb{E}[-B \log t(Q) - (1-B)\log(1-t(Q))]$. A lower expected log-loss indicates a more accurate forecast. The following proposition relates delivered information to forecast accuracy and the informed trader's information.

Proposition 1 (Information and delivered accuracy)**.** *Under Assumptions 1 and 2, at any liquidity $b > 0$:*

(i) (Posterior beliefs.) $\Pr(B = 1 \mid \delta_A(Q)) = \delta_A(Q)$ *almost surely, and* $\mathbb{E}[\delta_A(Q)] = \delta_0$.

(ii) (Accuracy.) *The expected log-loss of the final price is* $\mathbb{E}[h(\delta_A(Q))] = h(\delta_0) - \mathcal{I}(b)$.

(iii) (Decomposition.) $\mathcal{I}(b) = \vartheta\, y - \mathcal{L}_T(b)$, *with* $0 \le \mathcal{L}_T(b) \le \vartheta\, y$ *and* $0 \le \mathcal{I}(b) \le h(\delta_0)$.

(iv) (Forecast correction.) *The expected log-loss of the first-round price is* $h(\delta_0) - \mathcal{I}(b) + D_A(b)$.

Proof. Appendix A.4.

A buyer who rewards improvements in forecast accuracy values the reduction in expected log-loss relative to the prior. Part (ii) identifies $\mathcal{I}(b)$ as that reduction, providing the basis for the market maker's revenue in Section 5.

Uncertainty about trader type limits how much of the informed trader's information the final price reflects. For a nonzero order that could come from either type, the arbitrageur chooses a final price strictly between the first-round price and the prior, as specified in (5). If such an order is informed, the final price is closer to the prior than P; if it is noise, the final price still differs from the prior. Both discrepancies contribute to the type-confusion loss.[20]

The LMSR sets the first-round price from the submitted order without distinguishing its source. The arbitrageur accounts for the possibility that the order is noise. Its trade reduces expected log-loss by $D_A(b)$, so its expected rent equals b times the improvement in forecast accuracy. This improvement uses the information in the observed order; the arbitrageur receives no additional signal about the event.

Let π_1 denote the realized settlement profit of the first-round trader, who is either the informed trader or the noise trader, and let π_A denote the arbitrageur's realized settlement profit. The following proposition gives their combined expected profit.

[20] When $\vartheta = 1$, the arbitrageur knows that the order is informed and $\mathcal{I}(b) = y$. When $\vartheta = 0$, it restores the price to the prior and $\mathcal{I}(b) = 0$.

Proposition 2 (Net payout)**.** *Under Assumptions 1 and 2, at any liquidity $b > 0$, the combined profit $\pi_1 + \pi_A$ is integrable, and $\mathbb{E}[\pi_1 + \pi_A] = b\,\mathcal{I}(b)$.*

Proof. Appendix A.6.

The market maker's expected net payout to the first-round trader and the arbitrageur equals b per unit of delivered information. Noise trading affects this payout through the information reflected in the final price: uncertainty about whether an order comes from the informed trader or the noise trader reduces delivered information by $\mathcal{L}_T(b)$ and hence the expected net payout by $b\,\mathcal{L}_T(b)$, relative to a market in which trader type is disclosed.

We examine how the informed arrival probability ϑ and liquidity b affect delivered information, holding the distribution of P and the noise-order distribution F_U fixed. The informed order $Q_I = b(\operatorname{logit} P - \operatorname{logit} \delta_0)$ is the trader's optimal response to the LMSR schedule and therefore adjusts with b.

The informed and noise order distributions *overlap* at b if nonzero informed orders cannot be distinguished from noise orders with probability one.[21] The following proposition describes the effect of increasing the probability of informed arrival.

Proposition 3 (The arrival probability)**.** *Under Assumptions 1 and 2, fix $b > 0$, the distribution of P, and F_U, and regard $\mathcal{I}$ and $\mathcal{L}_T$ as functions of $\vartheta \in [0, 1]$.*

(i) *$\mathcal{I}$ is continuous, increasing, and convex on $[0, 1]$, with $\mathcal{I}(0) = 0$ and $\mathcal{I}(1) = y$. The ratio $\mathcal{I}(\vartheta)/\vartheta$ is increasing on $(0, 1]$. Delivered information is strictly increasing in ϑ if and only if $y > 0$.*

(ii) *$\mathcal{I}$ is strictly convex if and only if the order distributions overlap at b; otherwise $\mathcal{I}(\vartheta) = \vartheta y$ for every ϑ.*

[21] Formally, overlap means that no measurable set A satisfies both $\Pr(Q_I \notin A,\, Q_I \neq 0) = 0$ and $\Pr(U \in A) = 0$. With densities, overlap occurs when $f_{I,b}$ and f_U are both positive on a set of nonzero orders of positive measure.

(iii) $\mathscr{L}_T$ *is concave on* $[0,1]$*, with* $\mathscr{L}_T(0) = \mathscr{L}_T(1) = 0$*. Under overlap, it is strictly concave and positive on* $(0,1)$*; otherwise it is identically zero.*

Proof. Appendix A.7.

A higher informed arrival probability increases both the frequency of informative orders and the probability that the arbitrageur assigns to informed arrival after observing an order consistent with either type. The arbitrageur therefore moves the first-round price less toward the prior. Without overlap, observing a nonzero order identifies the trader's type almost surely, so uncertainty about type causes no information loss.[22]

Liquidity and delivered information. Delivered information need not be monotone in liquidity. For fixed ϑ, Lemma 5 establishes that

$$\lim_{b \downarrow 0} \mathcal{I}(b) = \lim_{b \to \infty} \mathcal{I}(b) = \vartheta y.$$

For fixed $P \neq \delta_0$, the optimal informed order approaches zero as $b \to 0$. Its absolute value tends to infinity as $b \to \infty$. The noise-order distribution remains fixed. This separation makes the type-confusion loss vanish in both limits. For $0 < \vartheta < 1$, $\mathcal{I}(b) < \vartheta y$ if and only if the informed and noise order distributions overlap at b. In this case, $\mathcal{I}$ is neither monotone, concave, nor convex in b.

The informed trader's information level y alone does not determine delivered information: two belief distributions with the same y can generate different overlap with noise orders and hence different $\mathcal{I}(b)$ at the same b, ϑ, and F_U (Lemma 5). Panels (a) and (b) of Figure 1 illustrate these comparative statics for the Gaussian measurement of Section 4.1.

[22] A zero order implies the event probability δ_0 under either trader type.

4. Information acquisition

We now endogenize the informed trader's information by allowing costly acquisition before trading. We specify the acquisition technology and characterize optimal acquisition at a given liquidity b. Section 5 studies the market maker's choice of liquidity in anticipation of this response.

The informed trader acquires a signal about B and uses its realization to form a posterior belief. We index the available signals by their information level y and write P_y for the resulting posterior, where

$$y \;=\; \mathbb{E}[\mathrm{KL}(P_y \,\|\, \delta_0)]. \tag{8}$$

By Lemma 1, a signal with level y yields expected trading profit $b\,y$. Its acquisition cost is $c(y)$, which is increasing and lower semicontinuous, with $c(0) = 0$ and $c(y) > 0$ for $y > 0$.[23]

Assumption 3 (Admissible acquisition levels)**.** *The feasible acquisition levels are* $[0, \bar{y}]$*, where* $0 < \bar{y} < h(\delta_0)$*. Each feasible* y *has a specified posterior belief* P_y *satisfying Assumption 1 and* (8)*.*

The cap $\bar{y}$ excludes full revelation. Specifying a posterior distribution at each level is necessary because signals with the same level can generate different informed-order distributions and hence different delivered information (Section 3.4). The arbitrageur anticipates the informed trader's optimal acquisition choice and uses the corresponding order distribution in its Bayesian update.[24]

For the informed trader, signals with the same level yield the same expected trading profit, so only acquisition cost distinguishes them. Appendix C formalizes the reduction

[23] Lower semicontinuity means $c(y) \le \liminf_{y' \to y} c(y')$. Together with the compact feasible set below, this technical condition ensures that the informed trader's set of optimal acquisition levels is nonempty and compact (Appendix A.10).

[24] If several acquisition levels are optimal, we select the largest, as in Section 5.

from a choice among signals to a choice of y.

We interpret informed arrival as an opportunity to research the event and trade on the resulting information. The probability ϑ governs whether the first-round trader has this opportunity; otherwise, the first-round order comes from the noise trader. The informed trader chooses y upon arrival, before observing the signal and submitting its order, and maximizes $b\,y - c(y)$.[25] The arbitrageur's subsequent trade does not affect this payoff (Section 3.1).

We characterize the informed trader's optimal acquisition and its response to liquidity. Under Assumption 3, Lemma 1 gives its expected payoff net of acquisition cost as

$$g_b(y) := \mathbb{E}[b\,\mathrm{KL}(P_y \,\|\, \delta_0)] - c(y) = b\,y - c(y), \tag{9}$$

and the informed trader chooses $y^*_{\mathrm{LMSR}}(b) \in \arg\max_{y\in[0,\bar{y}]} g_b(y)$.

Proposition 4 (Optimal LMSR acquisition)**.** *For $b > 0$:*

(i) *Suppose c is continuous on $[0,\bar{y}]$ and differentiable on $(0,\bar{y})$. Any interior maximizer satisfies*

$$b = c'(y^*). \tag{10}$$

(ii) *Under the hypotheses of part* (i)*, let $h_b(y) := g_b'(y) = b - c'(y)$. If some $\hat{y} \in (0,\bar{y})$ satisfies $h_b(y) > 0$ for $0 < y < \hat{y}$ and $h_b(y) < 0$ for $\hat{y} < y < \bar{y}$, then $\hat{y} = y^*_{\mathrm{LMSR}}(b)$ is the unique global maximizer.*

(iii) *The set $M(b) := \arg\max_{y\in[0,\bar{y}]} g_b(y)$ is nonempty and compact, and $\underline{y}(b) := \min M(b)$ and $\overline{y}(b) := \max M(b)$ are increasing in b on $(0,\infty)$. If the maximizer is unique for every $b > 0$, $y^*_{\mathrm{LMSR}}(b)$ is increasing in b.*

[25] For the acquisition problem, this timing is without loss of generality after rescaling the reward: if acquisition precedes arrival, expected profit is $\vartheta b\,y - c(y)$, so for $\vartheta > 0$ the characterization in Proposition 4 applies with ϑb replacing b.

Proof. Appendix A.9.

Liquidity b equals the informed trader's expected trading profit per unit of acquired information. Information acquisition and liquidity are complements: higher liquidity raises the return to acquiring additional information while leaving its cost unchanged. The informed trader consequently chooses a weakly higher acquisition level as liquidity increases.[26]

4.1 Two microfoundations of information acquisition

We now provide two microfoundations for the abstract acquisition technology: Bernoulli sampling, in which the informed trader chooses the sample size, and a Gaussian measurement, in which it chooses the precision.[27] Throughout, expectations are taken under the joint distribution of the event outcome and the observations generated by the experiment.

The two technologies differ in how much acquired information the final price reveals. With Bernoulli sampling, the informed trader's posterior belief takes finitely many values at each sample size. If the noise-order distribution assigns zero probability to every individual order, the arbitrageur can identify the trader's type almost surely, and there is no type-confusion loss (Proposition 10). With a Gaussian measurement of positive precision, the informed order has a positive density on $\mathbb{R}$. If noise orders also have a density and $0 < \vartheta < 1$, the arbitrageur cannot identify the trader's type on the overlapping orders, and the type-confusion loss is positive at every $b > 0$ (Proposition 13).

[26] When several acquisition levels are optimal, both the smallest and the largest optimal levels are increasing in liquidity b.

[27] Appendix C formalizes the general experiment model and its reduction to a choice of information level. Formal statements and proofs for the two technologies are in Appendix A.

Acquisition by sample size: a Beta–Bernoulli experiment. An unknown event probability θ has a Beta prior with mean δ_0 and concentration $\kappa > 0$:

$$\theta \sim \mathrm{Beta}(\kappa\delta_0,\ \kappa(1-\delta_0)).$$

A larger κ corresponds to less prior uncertainty about θ. Conditional on θ, the event outcome B and the observations $X_1, X_2, \dots$ are independent Bernoulli(θ) draws. Before observing the data, the informed trader chooses a sample size $N \in \mathbb{N}_0 := \{0, 1, 2, \dots\}$ at cost $c_N(N)$, where $c_N : \mathbb{N}_0 \to [0, \infty)$ is increasing and $c_N(0) = 0$.

Let $X_{1:N} := (X_1, \dots, X_N)$ and $S_N := \sum_{i=1}^{N} X_i$, with $S_0 := 0$. After observing S_N successes in N draws, the informed trader's posterior probability of the event is

$$P_N := \Pr(B = 1 \mid X_{1:N}) = \mathbb{E}[\theta \mid X_{1:N}] = \frac{\kappa\delta_0 + S_N}{\kappa + N},$$

with $P_0 = \delta_0$. The trader thus updates its forecast using the sample frequency, with the prior receiving less weight as the sample size increases.

The acquired information is $y(N) := \mathbb{E}[\mathrm{KL}(P_N \,\|\, \delta_0)]$. As shown in Proposition 9, $y(N)$ is strictly increasing in N and converges to $y_\infty := h(\delta_0) - \mathbb{E}[h(\theta)] < h(\delta_0)$. Additional observations improve the trader's estimate of the event probability, but even knowing θ leaves uncertainty about the event outcome.[28]

At liquidity b, the informed trader chooses N to maximize $b\,y(N) - c_N(N)$. Let $N^*(b)$ denote the largest optimal sample size and $S(b) := y(N^*(b))$ the supplied information, whenever this choice exists.[29] At an optimum with positive sample size, the expected trading profit attributable to the last observation is at least its incremental cost. The expected trading

[28] The achievable levels $\{y(N) : N \in \mathbb{N}_0\}$ form a countable set. Appendix C explains how public lotteries over sample sizes implement intermediate levels.

[29] The condition $c_N(N) \to \infty$ guarantees a finite, nonempty set of optimal sample sizes; see Corollary 2.

profit from one additional observation is no greater than its incremental cost (Corollary 2). These are the discrete counterparts of (10).

When the noise-order distribution assigns zero probability to every individual order, the arbitrageur leaves the informed trader's price unchanged and fully reverses the noise trader's order. Consequently, $\mathcal{I}(b) = \vartheta\, S(b)$ (Proposition 10): the final price reveals all acquired information whenever the informed trader arrives.

Acquisition by precision: a Gaussian experiment. The second technology models information acquisition through the precision of a Gaussian measurement of B. The informed trader chooses $\rho \geq 0$ at cost $c_\rho(\rho)$, where $c_\rho : [0,\infty) \to [0,\infty)$ is continuous and weakly increasing, with $c_\rho(0) = 0$. For $\rho > 0$, it observes $X_\rho = B + \varepsilon_\rho$, where $\varepsilon_\rho \sim N(0, \rho^{-2})$ is independent of B. Thus ρ is the inverse standard deviation of the measurement error. The sufficient statistic $T_\rho := \rho X_\rho = \rho B + Z$, where $Z := \rho\varepsilon_\rho \sim N(0,1)$, contains the same information as X_ρ. At $\rho = 0$, the trader acquires no information, and we set $T_0 := Z$. Bayes' rule gives

$$P_\rho(t) = \Pr(B = 1 \mid T_\rho = t) = \frac{1}{1 + \frac{1-\delta_0}{\delta_0} e^{-\rho t + \rho^2/2}}, \qquad \operatorname{logit} P_\rho(t) = \operatorname{logit} \delta_0 + \rho t - \frac{\rho^2}{2}. \quad (11)$$

Write $\hat{P}_\rho := P_\rho(T_\rho)$ for the informed trader's posterior belief, with $\hat{P}_0 = \delta_0$, and

$$m(\rho) := \mathbb{E}[\hat{P}_\rho(1 - \hat{P}_\rho)] \quad (12)$$

for the mean posterior variance of B. This variance measures the uncertainty that remains after the informed trader observes the signal.

The acquired information is $y(\rho) := \mathbb{E}[\mathrm{KL}(\hat{P}_\rho \,\|\, \delta_0)]$. As Proposition 11 establishes, $y(\rho)$ increases continuously and strictly from zero toward $h(\delta_0)$, with marginal information gain

$y'(\rho) = \rho m(\rho)$. Because the signal measures B directly, arbitrarily precise observations approach full revelation.

At liquidity b, the informed trader chooses precision to maximize its expected trading profit net of acquisition cost:

$$\max_{\rho \geq 0} \quad b\, y(\rho) - c_\rho(\rho). \tag{13}$$

When c_ρ is differentiable, an interior optimum equates the marginal trading profit from greater precision, $b\rho m(\rho)$, to its marginal cost $c'_\rho(\rho)$ (Corollary 3). The corresponding implementation schedule is

$$\mathfrak{b}(\rho) := \frac{c'_\rho(\rho)}{\rho m(\rho)}, \qquad \rho > 0. \tag{14}$$

When $\mathfrak{b}$ is continuous and strictly increasing with range $(0, \infty)$, the informed trader chooses $\rho^*(b) = \mathfrak{b}^{-1}(b)$ (Proposition 12). The market maker can therefore induce a target precision by posting the corresponding liquidity. Power costs $c_\rho(\rho) = \kappa_c \rho^p / p$, with $\kappa_c > 0$ and $p > 2$, satisfy these conditions.[30]

At any positive precision and liquidity, informed orders have a density that is positive at every order size. For $0 < \vartheta < 1$, when noise orders are drawn from a distribution with a probability density, the informed and noise order distributions overlap at every positive liquidity (Proposition 13).[31] The arbitrageur cannot perfectly distinguish the two trader types, so the final price reveals less information than it would if trader type were observable.

Under the conditions of Proposition 12, delivered information is continuous in liquidity. As $b \to 0$, acquisition and delivered information vanish. As $b \to \infty$, the informed trader chooses increasingly precise measurements and the type-confusion loss vanishes, so delivered

[30] A positive lower limit $\mathfrak{b}(0^+)$ creates an acquisition threshold: the informed trader acquires no information at or below that liquidity. Appendix A.16 states the conditions for this threshold response.

[31] This condition allows mixtures of discrete and continuously distributed noise orders. The exact condition is a nonzero absolutely continuous component of F_U, defined in Appendix A.17.

information converges to $\vartheta h(\delta_0)$ (Proposition 13).[32] Panel (c) of Figure 1 plots ϑS and $\mathcal{I}$ against liquidity under a power cost; the gap between them is the type-confusion loss.

4.2 Trading over the market's life

We extend the trading stage of Section 3 to multiple periods to examine how delivered information and the type-confusion loss accumulate over the market's life. The market maker commits to a fixed liquidity before trading begins, and each subsequent period opens at the previous period's final price. In each period, the informed trader arrives with probability ϑ and the noise trader otherwise. The informed trader chooses the precision of its measurement given the opening price, as in Section 4.1. The arbitrageur observes the resulting order and trades to its posterior. Each trader trades only once and holds its position until settlement after the final period. Each period therefore follows the same trading structure, with the opening price serving as the prior.[33]

Figure 3 plots expected cumulative delivered information and the expected cumulative type-confusion loss over time for different liquidity levels and informed-arrival probabilities. As shown in the figure, higher liquidity and more frequent informed arrival increase cumulative delivered information at every horizon. At a given opening price, higher liquidity raises the return to information acquisition, inducing the informed trader to choose greater precision (Proposition 12). A higher arrival probability increases both the frequency of informative orders and the weight the arbitrageur assigns to the information implied by an observed order (Proposition 3). The cumulative type-confusion loss is largest at intermediate liquidity and arrival probabilities. At low liquidity, the informed trader acquires little

[32] These limits describe the experiment without an acquisition cap. Under the cap in Assumption 3, let $y(\bar{\rho}) = \bar{y}$. Under the same conditions on $\mathfrak{b}$, the trader chooses $\min\{\mathfrak{b}^{-1}(b), \bar{\rho}\}$, and the large-liquidity limit of delivered information is $\vartheta\bar{y}$.

[33] Appendix D gives the distributions, number of periods, and parameter values used in the numerical illustration.

information, leaving little information to be lost through type confusion. As liquidity grows, informed orders become larger while the noise-order distribution remains fixed, and the loss vanishes in the limit (Proposition 13). When informed arrival is rare, few orders contain information, so the average loss from type confusion is small. When informed arrival is nearly certain, the arbitrageur faces little uncertainty about trader type, so the loss is again small (Proposition 3). The curves flatten over time as trading reduces uncertainty about the event. Both delivered information and the type-confusion loss in a period are bounded by ϑ times the entropy of its opening price (Proposition 1). When the opening price is close to zero or one, little uncertainty about the event remains, so both the information delivered and the type-confusion loss in that period are small.

5. The market maker's problem and the optimal liquidity

The preceding sections characterize information acquisition and trading at a given liquidity. We now study the market maker's choice of liquidity, anticipating how that choice affects information acquisition, delivered information, and trading profits. At date 0, the market maker commits to $b \geq 0$, with $b = 0$ corresponding to a closed market.

Let $S(b)$ denote the informed trader's optimal acquisition level in (9).[34] Write $\mathcal{I}(b)$ for delivered information in (7) under the resulting belief $P_{S(b)}$. We first set the market maker's objective, then characterize optimal liquidity and the value of disclosing trader type. We then provide a procurement foundation for the market maker's revenue from forecast accuracy.

The market maker earns $\lambda > 0$, its accuracy weight, per unit of reduction in expected

[34]When several levels are optimal, we select the largest. The results hold for any fixed selection whose delivered-information schedule satisfies their hypotheses (Appendix A.19).

log-loss relative to the prior.[35] Since the final price reduces expected log-loss by $\mathcal{J}(b)$, its expected revenue is $\lambda\,\mathcal{J}(b)$.

The market maker incurs a proportional settlement overhead $g \geq 0$ on its expected payments to the informed trader and the arbitrageur. Each dollar of these payments therefore costs the gross-up $\eta := 1 + g$, while receipts from the noise trader incur no overhead. The market maker also holds a reserve of $b\,M(\delta_0)$ to cover the LMSR's maximum possible loss, where $M(\delta_0) := -\log\min\{\delta_0, 1-\delta_0\}$. Carrying the reserve until settlement costs $\gamma > 0$ per unit of capital, giving a total capital charge $c_0 b$, where $c_0 := \gamma M(\delta_0) > 0$.

Suppose $\mathbb{E}|U| < \infty$, so the noise trader's expected loss is finite. The settlement accounts give

$$b\,\mathcal{J}(b) \;=\; \vartheta\, b\, S(b) + b\, D_A(b) - (1-\vartheta)\, b\, \nu_b. \tag{15}$$

The market maker pays the informed trader $\vartheta b S(b)$ and the arbitrageur $bD_A(b)$ in expectation, and receives $(1-\vartheta)b\nu_b$ from the noise trader. Its expected net payout is therefore $b\mathcal{J}(b)$. The overhead applies to the payments before deducting these receipts, so the total settlement overhead is

$$g[\vartheta b S(b) + bD_A(b)] = g\, b\mathcal{J}(b) + g(1-\vartheta)b\nu_b.$$

[35]The rule pays out in expectation whenever the trader brings information, so an LMSR is run at an expected loss by a sponsor whose return is the forecast itself. Polymarket's early stage is the concrete analogue: before moving to an order book, its markets ran on an automated market maker whose liquidity pools were supplied by the platform and its employees among the liquidity providers (Commodity Futures Trading Commission, 2022) and a platform-administered rewards program subsidized liquidity provision on top (Polymarket, 2022). Its prices — routinely quoted as event probabilities (Clinton and Huang, 2025) — are its product. λ prices this accuracy in reduced form: the credibility and adoption an accurate feed earns the platform, or, exactly, the rate at which a client pays for the delivered forecast under a logarithmic scoring contract.

Subtracting the net payout, settlement overhead, and capital charge from revenue gives

$$V_\lambda(b) \;:=\; (\lambda - \eta\, b)\, \mathcal{J}(b) - A(b), \qquad A(b) \;:=\; (\eta - 1)(1 - \vartheta) b \nu_b + c_0 b, \tag{16}$$

for $b > 0$, with $V_\lambda(0) := 0$ for the closed market (Appendix A.18).[36]

The margin $\lambda - \eta b$ is revenue per unit of delivered information less the net payout and its proportional overhead. The term $A(b)$ includes the capital charge and the remaining overhead on payments financed by receipts from the noise trader. No overhead is charged on the receipts themselves.[37]

At a fixed liquidity, the arbitrageur improves forecast accuracy but also receives a payment from the market maker. Its correction reduces expected log-loss by $D_A(b)$, increasing revenue by $\lambda D_A(b)$, and costs $\eta b D_A(b)$ including overhead. Relative to the same market without the arbitrageur, the market maker's expected profit therefore changes by $(\lambda - \eta b) D_A(b)$ (Appendix A.18). The correction is profitable when $b < \lambda/\eta$ and $D_A(b) > 0$.[38]

5.1 Optimal liquidity

We now choose liquidity to maximize the market maker's profit in (16), accounting for the informed trader's acquisition response and the information revealed by the resulting order flow. Set $\mathcal{J}(0) := 0$ for the closed market. Let $B^*(\lambda)$ denote the set of maximizers of V_λ

[36] When $\eta = 1$, settlement overhead is zero and $A(b) = c_0 b$. The objective then follows directly from the expected net payout in Proposition 2.

[37] When $\vartheta = 1$, the arbitrageur earns no profit, $\mathcal{J}(b) = S(b)$, and $A(b) = c_0 b$. Equation (16) then gives the single-trader objective.

[38] This comparison assumes $\mathbb{E}|U| < \infty$ and holds liquidity fixed. Without the arbitrageur, the market maker can profit from noise-trader losses even when the final price is less accurate than the prior. The fixed-liquidity comparison therefore does not imply a ranking of maximized profits or market-entry thresholds; Appendix A.18 provides a counterexample.

over $[0,\infty)$ and $V^*(\lambda) := \sup_{b\geq 0} V_\lambda(b)$. Define the entry threshold

$$\lambda_{\text{entry}} \;:=\; \inf\Big\{\eta\, b + \frac{A(b)}{\mathcal{F}(b)} :\; b>0,\; \mathcal{F}(b)>0\Big\}, \qquad \inf\varnothing := +\infty. \tag{17}$$

For any liquidity that delivers information, $\eta b + A(b)/\mathcal{F}(b)$ is the total cost per unit of delivered information: the grossed-up payment plus the overhead on payments financed by noise losses and the capital charge, both allocated across the information delivered.

We now characterize the conditions for profitable market entry and the market maker's optimal liquidity.

Theorem 1 (Optimal liquidity and market entry)**.** *Under Assumptions 1 and 2, fix $\lambda>0$, $\eta\geq 1$, $c_0>0$, the technology and the selection, and suppose $\mathbb{E}|U|<\infty$ when $\eta>1$.*

(i) (Liquidity cap.) *For every $b\geq\lambda/\eta$, $V_\lambda(b)\leq -A(b)<0=V_\lambda(0)$. Hence every maximizer of V_λ over $[0,\infty)$ lies in $[0,\lambda/\eta)$, and a maximizer $b^*>0$ satisfies $(\lambda-\eta\, b^*)\,\mathcal{F}(b^*)\geq A(b^*)>0$.*

(ii) (Existence.) *If $\mathcal{F}$ is upper semicontinuous on $(0,\lambda/\eta)$ and $\lambda\limsup_{b\downarrow 0}\mathcal{F}(b)\leq\lim_{b\downarrow 0}A(b)$, then V_λ attains its maximum over $[0,\infty)$, and $B^*(\lambda)$ is nonempty and compact.*

(iii) (Entry.) *$V^*(\lambda)>0$ if and only if $\lambda>\lambda_{\text{entry}}$, and $V^*(\lambda)=0$ at $\lambda=\lambda_{\text{entry}}$, whether or not the infimum in (17) is attained. If $\lambda<\lambda_{\text{entry}}$, then $B^*(\lambda)=\{0\}$.*

(iv) (Marginal condition.) *If $b^*\in(0,\lambda/\eta)$ maximizes V_λ and $\mathcal{F}$ is differentiable at b^*, then*

$$(\lambda-\eta\, b^*)\,\mathcal{F}'(b^*) \;=\; \eta\,\mathcal{F}(b^*)+A'(b^*), \qquad A'(b) \;=\; c_0-(\eta-1)\,(1-\vartheta)\,\mathbb{E}[\mathrm{KL}(\delta_1(U)\,\|\,\delta_0)]. \tag{18}$$

Proof. Appendix A.19.

We next examine how the market maker's valuation of accuracy λ affects delivered information at the optimum and its optimized profit.

Proposition 5. *Under the hypotheses of Theorem 1:*

(i) *If $\lambda_1 < \lambda_2$, then $\mathcal{I}(b_2) \geq \mathcal{I}(b_1)$ for every $b_1 \in B^*(\lambda_1)$ and $b_2 \in B^*(\lambda_2)$.*

(ii) *The function $\lambda \mapsto V^*(\lambda)$ is increasing, convex, and Lipschitz with constant $\vartheta\, h(\delta_0)$ on $(0, \infty)$.*

Proof. Appendix A.19.

A higher valuation of accuracy λ also increases optimal liquidity when the capital charge c_0 is sufficiently large.

Proposition 6. *Under the hypotheses of Theorem 1, suppose*

$$c_0 \geq (\eta - 1)(1 - \vartheta)\, M(\delta_0).$$

(i) *For any $\lambda_1 < \lambda_2$ and $b_i \in B^*(\lambda_i)$, optimal liquidity satisfies $b_1 \leq b_2$.*

(ii) *If $b < b'$ are optimal at valuations $\lambda \leq \lambda'$, respectively, then $\mathcal{I}(b) < \mathcal{I}(b')$.*

(iii) *For all but countably many $\lambda > \lambda_{\text{entry}}$, V_λ has at most one maximizer, and exactly one under the hypotheses of Theorem 1*(ii).

Proof. Appendix A.19.

Higher liquidity encourages information acquisition but raises the payment on every unit of delivered information. It also reduces the noise trader's expected loss and the settlement overhead on the payments that loss finances. The market maker therefore chooses liquidity

both to induce information acquisition and to reduce settlement costs. The bound in Proposition 6 ensures that the capital charge offsets these overhead savings, so a higher valuation of accuracy increases optimal liquidity.[39, 40]

Figure 2 plots V^* and the optimal liquidity against λ under the bound of Proposition 6.

Finally, we examine how the capital charge affects optimal liquidity.

Proposition 7. *Under the hypotheses of Theorem 1, fix λ and η. If $c_1 < c_2$ are two capital charges and b_i maximizes V_λ at c_i, then $b_2 \leq b_1$.*

Proof. Appendix A.19.

A higher capital charge c_0 makes larger liquidity commitments relatively more costly, leading the market maker to choose weakly lower liquidity.

We next consider how the probability of informed arrival affects optimized profit. The informed trader chooses acquisition conditional on arrival, so the selected experiment at each liquidity is unchanged as ϑ varies (Section 4).

Corollary 1. *Under Assumptions 1 and 2, fix λ, η, c_0, the noise law, and the experiment selected at each liquidity, suppose $\mathbb{E}|U| < \infty$ when $\eta > 1$, and regard V^* as a function of $\vartheta \in [0,1]$; let $\vartheta_{\text{entry}} := \inf\{\vartheta : V^*(\vartheta) > 0\}$, $\inf \varnothing := +\infty$.*

(i) *V^* is continuous, increasing and convex on $[0,1]$, with $V^*(0) = 0$ and $V^*(\vartheta) \leq \lambda\, \vartheta\, h(\delta_0)$, and $V^*(\vartheta)/\vartheta$ is increasing on $(0,1]$.*

(ii) *The set $\{\vartheta : V^*(\vartheta) > 0\}$ is empty or equals $(\vartheta_{\text{entry}}, 1]$ with $\vartheta_{\text{entry}} < 1$ and $V^*(\vartheta_{\text{entry}}) = 0$, and V^* is strictly increasing on it.*

[39] The technical conditions for existence in Theorem 1(ii) hold in the precision technology under the hypotheses of Proposition 12. They also hold in the sampling technology when F_U has no atoms, $c_N(N) \to \infty$, and $c_N(1) > 0$, with the largest optimal sample size selected. Appendix A.19 verifies these conditions.

[40] The single-trader case, $\vartheta = 1$, has $\mathcal{I} = S$ and $A(b) = c_0 b$; Appendix A.20 gives its characterization and the associated monopsony markdown.

Proof. Appendix A.25.

A higher probability of informed arrival improves final-price accuracy at each liquidity and reduces the overhead on payments financed by noise losses. Both effects raise the market maker's profit at liquidities below λ/η. As informed arrival becomes more likely, the arbitrageur also assigns greater weight to the information implied by an order, so the market delivers more information per informed arrival.

Comparison with the benchmark without hidden trader type. The preceding results characterize liquidity and entry when trader type is not disclosed. We now examine how disclosing type changes the market maker's profit and entry decision. In the disclosed market, trader type becomes public after the first order and before the arbitrageur trades. The informed trader's payoff is determined by its own trade, so disclosure leaves its acquisition incentives and the supply $S(b)$ unchanged. The arbitrageur leaves an informed order unchanged and fully reverses a noise order. Delivered information is therefore $\vartheta S(b)$.

Let V_λ^D and λ^D_{entry} denote the disclosed market's objective and entry threshold, obtained from (16) and (17) by replacing $\mathcal{I}(b)$ with $\vartheta S(b)$. The following proposition compares the two markets at a fixed liquidity and after the market maker chooses liquidity.

Proposition 8 (Disclosure)**.** *Under Assumptions 1 and 2, fix $\lambda > 0$, $\eta \geq 1$, $c_0 > 0$, the technology and the selection, and suppose $\mathbb{E}|U| < \infty$ when $\eta > 1$.*

(i) *For every $b > 0$,*

$$V_\lambda^D(b) - V_\lambda(b) = (\lambda - \eta b)\,\mathcal{L}_T(b).$$

For $b < \lambda/\eta$, the difference is nonnegative and is strictly positive if and only if $\mathcal{L}_T(b) > 0$.

(ii) $\sup_{b\geq 0} V_\lambda^D(b) \geq \sup_{b\geq 0} V_\lambda(b)$ *and* $\lambda_{\text{entry}}^D \leq \lambda_{\text{entry}}$. *The first inequality is strict if* V_λ *attains its supremum at some* $b^* > 0$ *with* $\mathscr{L}_T(b^*) > 0$*; the second is strict if the infimum defining* λ_{entry} *is attained at some* b_0 *with* $\mathscr{L}_T(b_0) > 0$.

Proof. Appendix A.26.

Disclosure allows the arbitrageur to recover more information from the same first-round order. At a given liquidity, the additional information raises accuracy revenue by $\lambda\mathscr{L}_T(b)$ and the expected net payout to traders by $b\mathscr{L}_T(b)$. Including settlement overhead, the additional payout costs the market maker $\eta b\mathscr{L}_T(b)$. The noise trader's losses and the required reserve are unchanged. Disclosure thus raises profit whenever type confusion occurs and the value of a unit of information exceeds its cost, $\lambda > \eta b$.

Disclosure can also make a market profitable that would remain closed with trader type hidden. At a fixed liquidity, disclosure increases delivered information while leaving the capital charge and the settlement overhead on payments financed by the noise trader's losses unchanged. These costs therefore fall per unit of delivered information, allowing the market to become profitable at a lower valuation of accuracy. For valuations strictly between the two entry thresholds, the market maker earns positive profit under disclosure but keeps the market closed when trader type is hidden.

5.2 A procurement foundation for the accuracy weight

The preceding analysis takes the market maker's valuation of accuracy λ as given. We now derive it from a contract with a principal who buys the forecast. The principal pays the market maker λ times the improvement in the final price's logarithmic score over the opening price. The expected payment is $\lambda\,\mathscr{I}(b)$, so λ is the price per unit of delivered information. The market maker chooses liquidity to maximize contract revenue less net trading payouts,

settlement overhead, and the capital charge, as in (16).

Appendix B first studies the single-trader benchmark, $\vartheta = 1$. A planner values accuracy and includes the trader's surplus and the market maker's profit in welfare. With costless public funds, the planner can raise the payment per unit of information above its own valuation to offset the market maker's monopsony markdown. Proposition 16 characterizes the required payment rate and the condition under which it induces the liquidity the planner would choose directly. Costly public funds reduce the optimal payment rate in the closed-form cases of Corollary 5. When public funds are costly, delegating market operation requires financing the market maker's profit in addition to the costs of running the market. A private buyer instead maximizes its own surplus and pays less than its valuation. The market maker applies a further markdown when compensating the trader, compounding the reduction in information acquisition (Corollary 6).

With noise trading, the contract still pays for delivered information, but the planner's choice must account for how liquidity affects the type-confusion loss (Lemma 8). The same result separates the arbitrageur's contribution to accuracy from its compensation. At a fixed liquidity and acquisition choice, the correction improves accuracy but incurs settlement overhead on the payment to the arbitrageur. The payment itself is a transfer among participants and cancels from total welfare. With costless public funds and positive arbitrageur profit, the correction raises welfare exactly when the principal's valuation of accuracy exceeds $(\eta - 1)b$, the settlement overhead per unit of accuracy gained.

6. Policy implications

The policy debate over prediction markets concerns whether event contracts should fall under federal derivatives regulation or state gaming law, and whether the information they

produce serves an economic purpose. The CFTC's June 2026 proposal would subject sports contracts to a special public-interest review while generally excluding contracts on economic and financial indicators from that review. It nevertheless recognizes that sports contracts can produce useful information, to be weighed against risks to market integrity (Commodity Futures Trading Commission, 2026, pp. 35825–35830).

Our model gives this informational benefit a contractual foundation. A user pays the market maker for improvements in the forecast's logarithmic score, evaluated against the realized outcome (Section 5.2). The market maker commits liquidity before research is undertaken, thereby determining the informed trader's expected trading profit per unit of acquired information. The trader earns this compensation by trading on what it learns, with its order conveying information to prices. Users can therefore finance information acquisition through trading profits without observing the underlying research. The market maker's monopsony separates the payment it receives for delivered information from the compensation it provides for information acquisition.

Sports and economic contracts may differ in the value of their forecasts and the prevalence of informed trading. More frequent informed arrival raises delivered information, while overlap with noise orders reduces it through type confusion (Propositions 1 and 3). Expected net trading payouts nevertheless equal liquidity times delivered information (Proposition 2); noise trading therefore does not remove the need to fund information production. The market opens profitably only when the contractual payment for accuracy covers trading payouts and operating costs (Theorem 1). Our analysis informs the public-interest assessment by characterizing when forecast demand can finance costly information acquisition and how much of that information reaches users.

7. Conclusion

We study liquidity provision in prediction markets as data marketplaces, where an automated market maker subsidizes information acquisition through trading. Under the logarithmic market scoring rule (LMSR), liquidity equals both the informed trader's expected trading profit per unit of acquired information and the market maker's expected net payout per unit of delivered information. Delivered information equals acquired information weighted by the probability of informed arrival, less a type-confusion loss arising from the arbitrageur's uncertainty about trader type. Disclosure eliminates this loss and can induce the market maker to open a market that is unprofitable when trader type is hidden.

Higher liquidity encourages information acquisition but need not increase delivered information. Informed orders scale with liquidity, whereas the noise-order distribution remains fixed. Holding acquisition fixed, the type-confusion loss vanishes as liquidity tends to zero or infinity. A positive loss at an intermediate liquidity therefore makes delivered information nonmonotone in liquidity. Bernoulli sampling eliminates type confusion when noise assigns zero probability to each individual order, whereas Gaussian measurement at positive precision leaves a positive loss when the distribution of noise orders has a density.

The market maker chooses liquidity to maximize the value of delivered information net of trading payouts, settlement overhead, and capital costs. Entry is profitable when the value of a unit of delivered information exceeds its total cost at some liquidity. A higher valuation of accuracy increases delivered information at the optimum. It also increases optimal liquidity when the capital charge is sufficiently large relative to settlement overhead.

Our results inform the regulatory debate over prediction markets by linking their informational benefits to the financing of research. Users who value forecasts can compensate information acquisition through the trading profits offered by a market maker. The informa-

tion delivered depends on research incentives, informed participation, and how well informed orders can be distinguished from noise. Sports and economic contracts may differ in these respects, but event category alone does not establish informational value. Our analysis therefore helps assess when forecast demand can sustain information production and how market design affects its cost and the information users receive.

We model a market with one initial order followed by an arbitrageur's correction. The market maker uses LMSR and commits to liquidity before trading. A dynamic extension could allow the market maker to choose liquidity as trading proceeds, accounting for outstanding positions and traders' anticipation of liquidity changes. A second extension would let the market maker choose the scoring rule jointly with liquidity. We leave these questions for future research.

REFERENCES

Abernethy, Jacob, Yiling Chen, and Jennifer Wortman Vaughan, 2013, Efficient market making via convex optimization, and a connection to online learning, *ACM Transactions on Economics and Computation* 1, 12:1–12:39.

Amazon Web Services, 2026, AWS Data Exchange, Product documentation, `https://aws.amazon.com/data-exchange/`, accessed September 17, 2026.

Arrow, Kenneth J., 1962, Economic welfare and the allocation of resources for invention, in Richard R. Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–626 (Princeton University Press, Princeton).

Berg, Joyce E., Robert Forsythe, and Thomas A. Rietz, 1997, Iowa Electronic Market, in Dean Paxson, and Douglas Wood, eds., *The Blackwell Encyclopedic Dictionary of Finance*, 111–113 (Blackwell, Oxford), `https://www.biz.uiowa.edu/faculty/trietz/papers/iemenc.pdf`, accessed October 9, 2026.

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin, 2018, The design and price of information, *American Economic Review* 108, 1–48.

Capponi, Agostino, and Ruizhe Jia, 2025, Liquidity provision on blockchain-based decentralized exchanges, *Review of Financial Studies* 38, 3040–3085.

Chen, Yiling, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen, 2010, Gaming prediction markets: Equilibrium strategies with a market maker, *Algorithmica* 58, 930–969.

Chen, Yiling, and Jennifer Wortman Vaughan, 2010, A new understanding of prediction

markets via no-regret learning, in *Proceedings of the 11th ACM Conference on Electronic Commerce (EC 2010)*, 189–198.

Chen, Yiling, and Fang-Yi Yu, 2021, Optimal scoring rule design under partial knowledge, *arXiv preprint arXiv:2107.07420* .

Chetia, Ananya, 2026, Most prediction market contracts have low volume, leaving users exposed to volatility and bots, CNBC, July 2, 2026, `https://www.cnbc.com/2026/07/02/prediction-markets-mostly-have-thinly-traded-contracts-.html`, accessed October 9, 2026.

Clinton, Joshua D., and TzuFeng Huang, 2025, Prediction markets? the accuracy and efficiency of $2.4 billion in the 2024 presidential election, Working paper, SocArXiv d5yx2.

Commodity Futures Trading Commission, 2020, CFTC designates KalshiEX LLC as a contract market, Press release 8302-20, November 4, 2020, `https://www.cftc.gov/PressRoom/PressReleases/8302-20`, accessed September 14, 2026.

Commodity Futures Trading Commission, 2022, In the matter of Blockratize, Inc., d/b/a Polymarket.com, CFTC Docket No. 22-09, order instituting proceedings, `https://www.cftc.gov/media/6891/enfblockratizeorder010322/download`.

Commodity Futures Trading Commission, 2026, Prediction Markets; Public Interest Determinations, Notice of proposed rulemaking, 17 CFR Part 40, RIN 3038-AF65, 91 Fed. Reg. 35806 (June 12, 2026), Federal Register document 2026-11854, `https://www.federalregister.gov/d/2026-11854`.

Databricks, 2026, What is Databricks Marketplace?, Databricks documentation, `https://docs.databricks.com/aws/en/marketplace`, accessed September 17, 2026.

Diercks, Anthony M., Jared Dean Katz, and Jonathan H. Wright, 2026, Kalshi and the rise of macro markets, Finance and Economics Discussion Series 2026-010, Board of Governors of the Federal Reserve System.

Flaherty v. KalshiEX, LLC, 2026, Docket of the Supreme Court of the United States, No. 26-299, petition for a writ of certiorari filed September 2, 2026, docketed September 8, 2026, `https://www.supremecourt.gov/docket/docketfiles/html/public/26-299.html`, accessed October 9, 2026.

Futuur, 2024, Expanding the forecast horizon, Substack post, March 8, 2024, `https://futuur.substack.com/p/expanding-the-forecast-horizon`, accessed September 14, 2026.

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

Gneiting, Tilmann, and Adrian E. Raftery, 2007, Strictly proper scoring rules, prediction, and estimation, *Journal of the American Statistical Association* 102, 359–378.

Gnosis, 2020, Conditional tokens market makers: LMSRMarketMaker, Smart contract source, `https://github.com/gnosis/conditional-tokens-market-makers`, accessed July 30, 2026.

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

Guo, Dongning, Shlomo Shamai, and Sergio Verdú, 2005, Mutual information and minimum

mean-square error in Gaussian channels, *IEEE Transactions on Information Theory* 51, 1261–1282.

Halawi, Danny, Fred Zhang, Chen Yueh-Han, and Jacob Steinhardt, 2024, Approaching human-level forecasting with language models, in *Advances in Neural Information Processing Systems 37 (NeurIPS 2024)*.

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

Innes, Robert D., 1990, Limited liability and incentive contracting with ex-ante action choices, *Journal of Economic Theory* 52, 45–67.

Intercontinental Exchange, 2026, ICE launches Polymarket signals and sentiment tool turning crowd-sourced dynamic views into market opportunities, Press release, February 11, 2026, `https://ir.theice.com/press/news-details/2026/ICE-Launches-Polymarket-Signals-and-Sentiment-Tool-Turning-Crowd-Sourced-Dynamic-Views-into-Market-Opportunities/default.aspx`, accessed September 14, 2026.

Kalshi, 2026a, Kalshi API, Developer documentation, `https://docs.kalshi.com/welcome`, accessed September 14, 2026.

Kalshi, 2026b, Liquidity incentive program, Kalshi Help Center, `https://help.kalshi.com/en/articles/13823851-liquidity-incentive-program`, accessed September 5, 2026.

Kalshi, 2026c, The orderbook, Kalshi Help Center, `https://help.kalshi.com/en/articles/13823828-the-orderbook`, accessed September 5, 2026.

KalshiEX, LLC v. Assad, 2026, Opinion of the United States Court of Appeals for the Ninth Circuit, No. 25-7516, filed August 28, 2026, https://cdn.ca9.uscourts.gov/datastore/opinions/2026/08/28/25-7516.pdf.

KalshiEX, LLC v. Flaherty, 2026, Opinion of the United States Court of Appeals for the Third Circuit, No. 25-1922, decided April 6, 2026, https://www2.ca3.uscourts.gov/opinarch/251922p.pdf.

KalshiEX LLC v. Schuler, 2026, Opinion of the United States Court of Appeals for the Sixth Circuit, Nos. 26-3196/5235, consolidated with KalshiEX LLC v. Orgel, decided September 25, 2026, https://www.opn.ca6.uscourts.gov/opinions.pdf/26a0272p-06.pdf.

Koeppelmann, Martin, 2020, On-chain automated market makers for prediction markets, Gnosis forum post, June 11, 2020, https://forum.gnosis.io/t/on-chain-automated-market-makers-for-prediction-markets/662, accessed September 14, 2026.

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

Laffont, Jean-Jacques, and Jean Tirole, 1993, *A Theory of Incentives in Procurement and Regulation* (MIT Press, Cambridge, MA).

Lehar, Alfred, and Christine A. Parlour, 2025, Decentralized exchange: The uniswap automated market maker, *Journal of Finance* 80, 321–374.

Li, Yingkai, Jason D. Hartline, Liren Shan, and Yifan Wu, 2022, Optimization of scoring rules, in *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22)*, 988–989.

Milgrom, Paul, and Chris Shannon, 1994, Monotone comparative statics, *Econometrica* 62, 157–180.

Ostrovsky, Michael, 2012, Information aggregation in dynamic markets with strategic traders, *Econometrica* 80, 2595–2647.

Othman, Abraham, David M. Pennock, Daniel M. Reeves, and Tuomas Sandholm, 2013, A practical liquidity-sensitive automated market maker, *ACM Transactions on Economics and Computation* 1, 14:1–14:25.

Othman, Abraham, and Tuomas Sandholm, 2013, The Gates Hillman prediction market, *Review of Economic Design* 17, 95–128.

Polymarket, 2022, Liquidity mining & trading rewards, Platform documentation, `https://legacy-docs.polymarket.com/liquidity-mining-and-trading-rewards`, accessed July 14, 2026.

Polymarket, 2026a, CTF exchange, Source repository, archived May 11, 2026, `https://github.com/Polymarket/ctf-exchange`, accessed September 14, 2026.

Polymarket, 2026b, Liquidity rewards, Platform documentation, `https://docs.polymarket.com/programs/liquidity-rewards`, accessed September 5, 2026.

Polymarket, 2026c, Market data, Platform documentation, `https://docs.polymarket.com/market-data/overview`, accessed September 14, 2026.

Polymarket and Dow Jones, 2026, Polymarket and Dow Jones, publisher of The Wall Street Journal, announce exclusive prediction market partnership, Press release, Business Wire, January 7, 2026, `https://www.businesswire.com/news/home/20260107511213/en/Polymarket-and-Dow-Jones-Publisher-of-The-Wall-Street-Journal-Announce-Exclusive-Prediction-Market-Partnership`.

Radde, Kaitlyn, 2026a, Prediction markets' trading volume doubled between May and July, largely driven by sports, Pew Research Center short read, September 23, 2026, https://www.pewresearch.org/short-reads/2026/09/23/prediction-markets-trading-volume-doubled-between-may-and-july-largely-driven-by-sports/, accessed October 9, 2026.

Radde, Kaitlyn, 2026b, Trading volume on prediction markets has soared in recent months, Pew Research Center short read, May 27, 2026, https://www.pewresearch.org/short-reads/2026/05/27/trading-volume-on-prediction-markets-has-soared-in-recent-months/, accessed October 9, 2026.

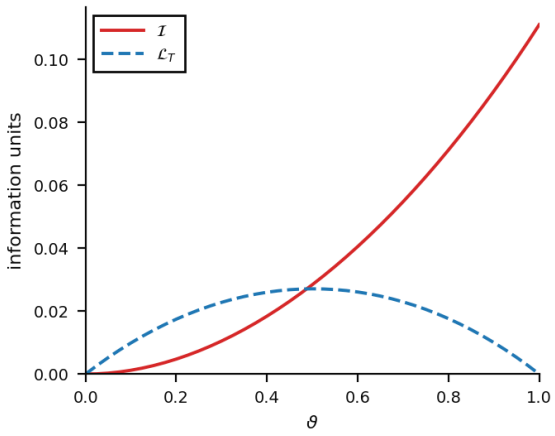
Savage, Leonard J., 1971, Elicitation of personal probabilities and expectations, *Journal of the American Statistical Association* 66, 783–801.

Snowflake, 2026, Snowflake Marketplace, Product documentation, https://www.snowflake.com/en/product/features/marketplace/, accessed September 17, 2026.

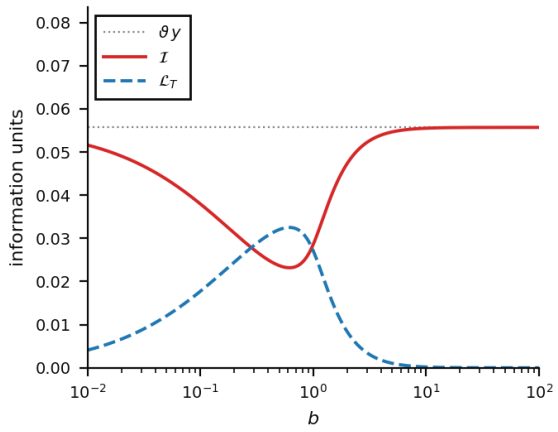
Spengler, Joseph J., 1950, Vertical integration and antitrust policy, *Journal of Political Economy* 58, 347–352.

Topkis, Donald M., 1998, *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).

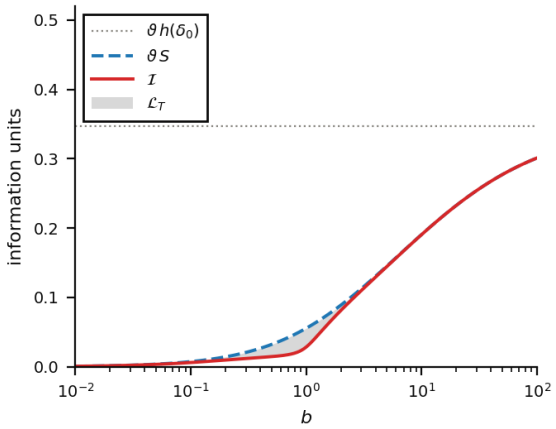
U.S.-China Economic and Security Review Commission, 2026, The people's republic of data: How China is turning data into capital, Issue Brief, August 18, 2026, https://www.uscc.gov/research/peoples-republic-data-how-china-turning-data-capital, accessed September 17, 2026.



(a) Arrival probability

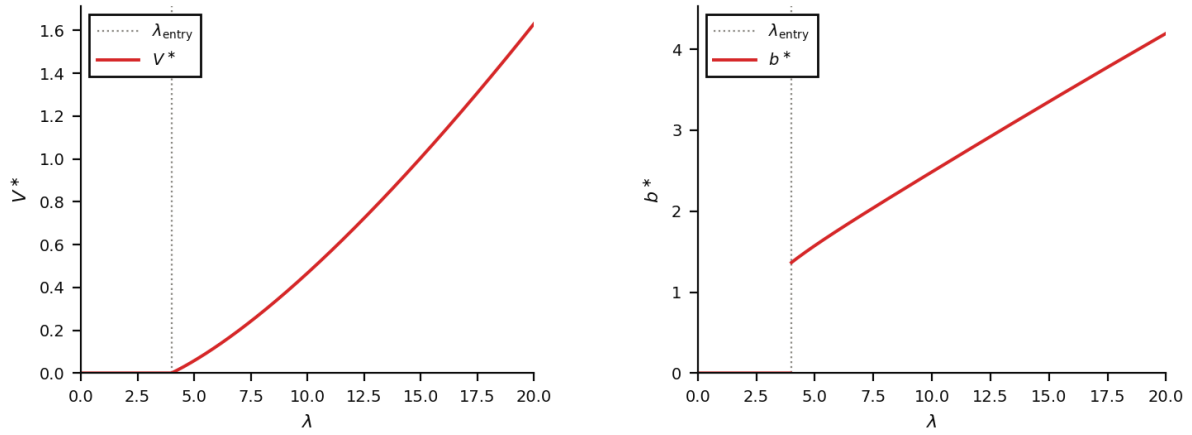


(b) Liquidity, fixed precision



(c) Liquidity, endogenous precision

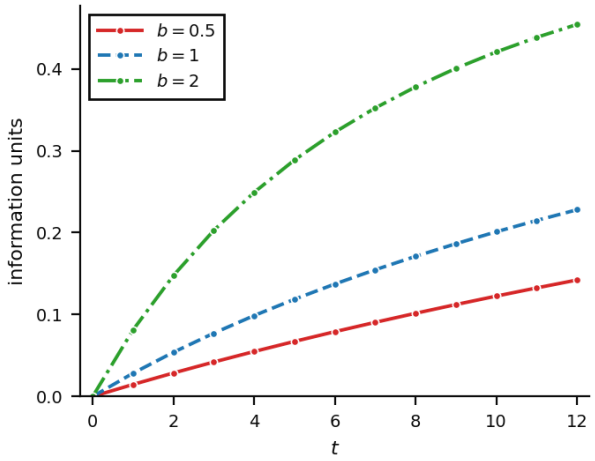
Figure 1: Delivered information $\mathcal{I}$ and the type-confusion loss $\mathcal{L}_T$. Panels (a) and (b) plot both against the arrival probability ϑ and against the liquidity b, respectively, with the remaining parameters and the precision held fixed; the dotted line in panel (b) is ϑy, delivered information without type confusion. Panel (c) plots $\mathcal{I}$ and ϑS, acquired information scaled by the arrival probability, against b with the precision the informed trader chooses at each liquidity; the shaded gap is $\mathcal{L}_T$, and the dotted line is the large-liquidity limit $\vartheta h(\delta_0)$. Parameter values are stated in Appendix D.



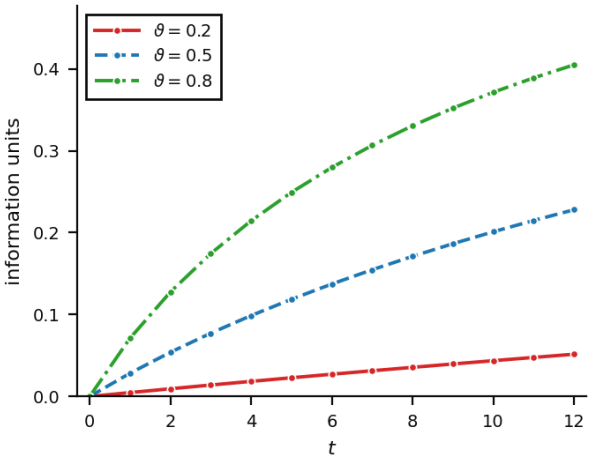
(a) Optimized profit

(b) Optimal liquidity

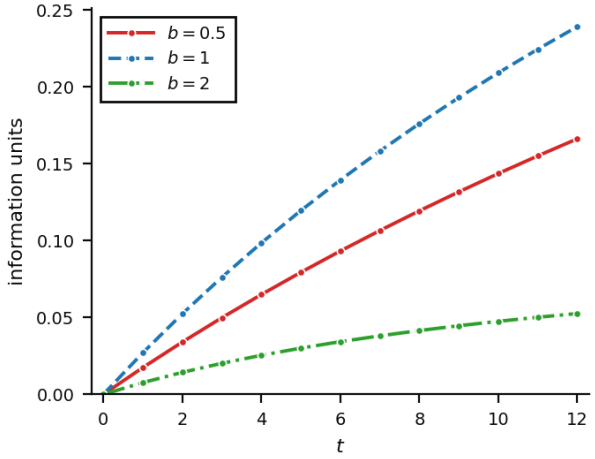
Figure 2: The market maker's optimized profit V^*, in panel (a), and optimal liquidity b^*, in panel (b), against the valuation of accuracy λ. The dotted line marks the entry threshold λ_{entry}. Parameter values are stated in Appendix D.



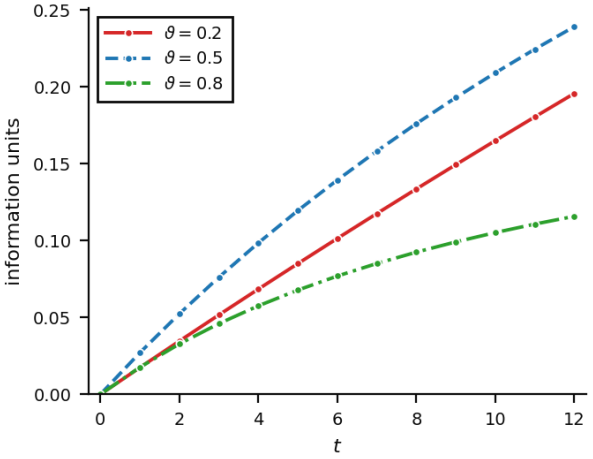
(a) Delivered information by liquidity



(b) Delivered information by arrival probability



(c) Type-confusion loss by liquidity



(d) Type-confusion loss by arrival probability

Figure 3: Expected cumulative delivered information, in panels (a) and (b), and expected cumulative type-confusion loss, in panels (c) and (d), over the periods of the multi-period market. Panels (a) and (c) vary the liquidity at the medium arrival probability; panels (b) and (d) vary the arrival probability at the medium liquidity. The horizontal axis is the period t. Parameter values are stated in Appendix D.

A. Omitted Proofs

This appendix collects the proofs, in the order the body states the results, together with the auxiliary lemmas the proofs use and the single-trader results of Appendix A.20. The proof of Lemma 2 alone departs from the body's order: it follows the proof of Proposition 1, whose bound on the type-confusion loss it uses. Throughout, for beliefs $r, a, d \in (0,1)$ we use the expansion

$$\mathrm{KL}(r \,\|\, d) - \mathrm{KL}(r \,\|\, a) \;=\; r \log \frac{a}{d} + (1-r) \log \frac{1-a}{1-d}, \tag{19}$$

affine in r, which follows from the definition of the divergence by cancelling the $r \log r$ and $(1-r)\log(1-r)$ terms, and its two consequences

$$\mathrm{KL}(r \,\|\, d) \;=\; -h(r) - r \log d - (1-r)\log(1-d) \quad \text{and} \quad -a \log d - (1-a)\log(1-d) \;=\; h(a) + \mathrm{KL}(a \,\|\, d), \tag{20}$$

the first by expanding the divergence, the second by applying the first at $r = a$. Since $-\log d$ and $-\log(1-d)$ are at most $M(d) := -\log \min(d, 1-d)$, the first identity gives $0 \le \mathrm{KL}(r \,\|\, d) \le M(d)$ for every $r \in [0,1]$. In the proofs of Section 3 the sender of the first order is written $J \in \{I, N\}$, the informed trader or the noise trader, so that $\Pr(J = I) = \vartheta$; $T_b(p) := b\,(\mathrm{logit}\, p - \mathrm{logit}\, \delta_0)$ is the map from a belief to the truthful order, a strictly increasing bijection of $(0,1)$ onto $\mathbb{R}$ with inverse δ_1; G is the law of the informed order $Q_I = T_b(P)$; and $M := \vartheta\, G + (1-\vartheta)\, F_U$ is the law of the first order Q.

A.1 Proof of Lemma 1

Let q be an inventory vector whose YES price is δ_0, and let $\tilde{q}$ be an inventory vector whose YES price is δ. If the realized state is $i \in \{0,1\}$, the trader's cost-function profit from moving q to $\tilde{q}$ is

$$(\tilde{q}_i - q_i) - (C(\tilde{q}) - C(q)).$$

For LMSR,

$$\log p_i(q) = \frac{q_i}{b} - \log\left(e^{q_1/b} + e^{q_0/b}\right),$$

where $p_1(q)$ is the YES price and $p_0(q) = 1 - p_1(q)$ is the NO price. Therefore

$$b(\log p_i(\tilde{q}) - \log p_i(q)) = (\tilde{q}_i - q_i) - (C(\tilde{q}) - C(q)).$$

Since $p_1(q) = \delta_0$, $p_1(\tilde{q}) = \delta$, $p_0(q) = 1 - \delta_0$, and $p_0(\tilde{q}) = 1 - \delta$, the state-contingent profits are exactly $b \log(\delta/\delta_0)$ in state $B = 1$ and $b \log((1-\delta)/(1-\delta_0))$ in state $B = 0$.

In particular, for the net YES order $\tilde{q} = (q_1 + x, q_0)$ the profits are $x - [C(\tilde{q}) - C(q)]$ in state $B = 1$ and $-[C(\tilde{q}) - C(q)]$ in state $B = 0$, whose expectation under P is the objective

in (2). Since the post-trade price

$$\delta = p_1(\tilde{q}) = \sigma(\operatorname{logit} \delta_0 + x/b)$$

is strictly increasing in x and maps $\mathbb{R}$ onto $(0,1)$, maximizing (2) over the order x is the same as maximizing the expected profit over the target price $\delta \in (0,1)$.

Taking expectation under posterior P gives $\pi(\delta; P)$. Its first derivative is

$$\partial_\delta \pi(\delta; P) = b\left(\frac{P}{\delta} - \frac{1-P}{1-\delta}\right),$$

and its second derivative is

$$\partial_\delta^2 \pi(\delta; P) = -\frac{bP}{\delta^2} - \frac{b(1-P)}{(1-\delta)^2} < 0$$

for $P, \delta \in (0,1)$. Thus $\pi(\cdot; P)$ is strictly concave. The first-order condition $\partial_\delta \pi(\delta; P) = 0$ is equivalent to $P(1-\delta) = (1-P)\delta$, hence to $\delta = P$. Strict concavity makes this the unique maximizer. Substituting $\delta = P$ gives

$$\pi(P; P) = b\left[P \log \frac{P}{\delta_0} + (1-P) \log \frac{1-P}{1-\delta_0}\right] = b\,\mathrm{KL}(P \,\|\, \delta_0),$$

which is nonnegative by Gibbs' inequality. The argument uses δ_0 only as the YES price of the starting inventory q, so every statement holds with any interior price in its place, as footnote 13 records. Q.E.D.

A.2 Statement and proof of Lemma 3

We next examine how liquidity affects the noise trader's expected loss.

Lemma 3 (The noise trader's loss)**.** *Under Assumption 2, suppose* $\mathbb{E}|U| < \infty$ *and fix* F_U*.*

(i) $\nu_b < \infty$ *for every* $b > 0$*. The noise trader's expected trading profit, weighted by its arrival probability, is* $-(1-\vartheta) b \nu_b$*.*

(ii) $b\nu_b$ *is decreasing and convex in* $b > 0$*, strictly so if* $\Pr(U \neq 0) > 0$*, and identically zero otherwise. Moreover,*

$$\lim_{b \to \infty} b\nu_b = 0, \qquad \lim_{b \downarrow 0} b\nu_b = \mathbb{E}[(1-\delta_0)U^+ + \delta_0 U^-],$$

where $U^+ := \max\{U, 0\}$ *and* $U^- := \max\{-U, 0\}$*.*

(iii) *$b\nu_b$ is differentiable in $b > 0$, with*

$$\frac{d}{db}(b\nu_b) = -\mathbb{E}[\mathrm{KL}(\delta_1(U) \,\|\, \delta_0)] \in [-M(\delta_0), 0], \qquad M(\delta_0) := -\log\min\{\delta_0, 1-\delta_0\}.$$

Higher liquidity reduces the price impact of a given noise order, bringing its average execution price closer to the prior δ_0. The noise trader's expected loss falls, with smaller reductions from successive increases in liquidity. As $b \downarrow 0$, the average execution price approaches 1 for a fixed positive order and 0 for a fixed negative order. The limiting expected loss is therefore $1-\delta_0$ per YES share purchased and δ_0 per YES share sold.

Proof. The displacement divergence. For a log-odds shift $a \in \mathbb{R}$ write $w(a) := \sigma(\operatorname{logit}\delta_0 + a)$ and

$$\varphi(a) \;:=\; \mathrm{KL}(\delta_0 \,\|\, w(a)) \;=\; \log(1-\delta_0+\delta_0 e^a) - \delta_0\, a, \tag{21}$$

the identity by substituting $w(a) = \delta_0 e^a/(1-\delta_0+\delta_0 e^a)$ and $1-w(a) = (1-\delta_0)/(1-\delta_0+\delta_0 e^a)$ into the divergence. Differentiating, $\varphi'(a) = w(a) - \delta_0$ and $\varphi''(a) = w(a)(1-w(a)) > 0$, so φ is strictly convex with $\varphi(0) = \varphi'(0) = 0$, hence $\varphi \geq 0$; and since $w \in (0,1)$, $|\varphi'| < \max\{\delta_0, 1-\delta_0\} \leq 1$, so $\varphi(a) \leq \max\{\delta_0, 1-\delta_0\}\,|a| \leq |a|$. As $a \to \infty$, $\varphi'(a) \to 1-\delta_0$ and $\varphi(a)/a \to 1-\delta_0$; as $a \to -\infty$, $\varphi'(a) \to -\delta_0$ and $\varphi(a)/a \to -\delta_0$, by l'Hôpital's rule. Since $\operatorname{logit}\delta_1(u) = \operatorname{logit}\delta_0 + u/b$, the noise order u leaves the quote $\delta_1(u) = w(u/b)$, and $\nu_b = \mathbb{E}[\varphi(U/b)]$.

(i) $\nu_b \leq \mathbb{E}|U|/b < \infty$ by the bound $\varphi(a) \leq |a|$. Conditional on $J = N$ and $U = u$, the event has probability δ_0 (Assumption 2), so the expected settlement profit of the move $\delta_0 \to \delta_1(u)$ is, by Lemma 1's state-contingent payoffs, $b\,[\delta_0 \log(\delta_1(u)/\delta_0) + (1-\delta_0)\log((1-\delta_1(u))/(1-\delta_0))] = -b\,\mathrm{KL}(\delta_0 \,\|\, \delta_1(u)) = -b\,\varphi(u/b)$. The realized profit is integrable: in state $B = 1$ it is $b\,[\log\sigma(\ell + u/b) - \log\sigma(\ell)]$ with $\ell = \operatorname{logit}\delta_0$, and $\frac{d}{da}\log\sigma(a) = 1-\sigma(a) \in (0,1)$, so its absolute value is at most $|u|$; likewise in state $B = 0$. Averaging over U and weighting by the arrival probability $1-\vartheta$ gives the noise trader's expected trading profit $-(1-\vartheta)\,b\,\nu_b$.

(ii) Fix $u \neq 0$ and set $\psi_u(b) := b\,\varphi(u/b)$, the perspective of φ. With $a = u/b$,

$$\psi_u'(b) \;=\; \varphi(a) - a\,\varphi'(a) \;<\; 0, \qquad \psi_u''(b) \;=\; \frac{a^2\,\varphi''(a)}{b} \;>\; 0,$$

the sign of ψ_u' because the tangent to the strictly convex φ at $a \neq 0$ lies strictly below the graph at 0: $0 = \varphi(0) > \varphi(a) + \varphi'(a)(0-a)$. So ψ_u is strictly decreasing and strictly convex on $(0,\infty)$, while $\psi_0 \equiv 0$. Since $b\,\nu_b = \mathbb{E}[\psi_U(b)]$ is finite, it is nonincreasing and convex, strictly on both counts when $\Pr(U \neq 0) > 0$, since the strict inequalities hold on a set of positive probability and the weak ones everywhere, and identically zero when $U = 0$ almost surely. Limits: as $b \to \infty$, $\psi_u(b) = u\,\varphi(u/b)/(u/b) \to u\,\varphi'(0) = 0$ for each $u \neq 0$, and $|\psi_u(b)| \leq \max\{\delta_0, 1-\delta_0\}|u|$ is integrable, so $b\,\nu_b \to 0$ by dominated convergence. As $b \downarrow 0$, $\psi_u(b) = u\,\varphi(a)/a$ with $a = u/b \to \pm\infty$ according to the sign of u, so $\psi_u(b) \to (1-\delta_0)\,u$

for $u > 0$ and $\psi_u(b) \to (-\delta_0)\, u = \delta_0|u|$ for $u < 0$; the convergence is monotone, ψ_u being nonincreasing in b, so the monotone convergence theorem gives $b\,\nu_b \uparrow \mathbb{E}[(1-\delta_0)U^+ + \delta_0 U^-]$.

(iii) By (21) and the first identity in (20), with $L := \log(1-\delta_0+\delta_0 e^a)$, $\log(w/\delta_0) = a - L$ and $\log((1-w)/(1-\delta_0)) = -L$, so $\mathrm{KL}(w(a) \,\|\, \delta_0) = w(a-L) - (1-w)L = aw - L$, while $\varphi(a) - a\varphi'(a) = L - \delta_0 a - a(w - \delta_0) = L - aw$. Hence

$$\psi_u'(b) \;=\; -\mathrm{KL}(\delta_1(u) \,\|\, \delta_0), \qquad 0 \;\le\; \mathrm{KL}(\delta_1(u) \,\|\, \delta_0) \;<\; M(\delta_0),$$

the strict bound because $\mathrm{KL}(w \,\|\, \delta_0) = -h(w) - w\log\delta_0 - (1-w)\log(1-\delta_0)$ with $h(w) > 0$ at interior w and $-\log\delta_0, -\log(1-\delta_0) \le M(\delta_0)$. So $|\psi_U'(b)| \le M(\delta_0)$ for every b, the difference quotients of ψ_U are bounded by $M(\delta_0)$ by the mean value theorem, and dominated convergence lets us differentiate $b\,\nu_b = \mathbb{E}[\psi_U(b)]$ under the expectation: $\frac{d}{db}(b\,\nu_b) = -\mathbb{E}[\mathrm{KL}(\delta_1(U) \,\|\, \delta_0)] \in [-M(\delta_0), 0]$. No smaller constant serves: for a sure order $U = u > 0$, $\mathrm{KL}(w(u/b) \,\|\, \delta_0) \to -\log\delta_0$ as $b \downarrow 0$, since $w \to 1$ and $h(w) \to 0$; for $u < 0$ the limit is $-\log(1-\delta_0)$; so the supremum of $-\frac{d}{db}(b\nu_b)$ over noise laws and liquidities is $M(\delta_0)$. Q.E.D.

A.3 The trading model: the type posterior and finiteness

The next lemma supplies the objects of Section 3 in general, with no density and no moment assumed: the type posterior as a Radon–Nikodym derivative, the final price as the conditional probability of the event, and the finiteness of every order and quote.

Lemma 4 (The type posterior)**.** *Under Assumptions 1 and 2, fix $b > 0$.*

(i) *Q_I is finite almost surely, $\delta_1(q) \in (0,1)$ for every real q, and $\delta_1(Q_I) = P$.*

(ii) *There is a measurable $\Psi : \mathbb{R} \to [0,1]$, unique up to M-null sets, with $\int_E \Psi\, dM = \vartheta\, G(E)$ for every Borel set E, and $\Psi(Q) = \Pr(J = I \mid Q)$ almost surely; $\mathbb{E}[\Psi(Q)] = \vartheta$. If G and F_U have densities $f_{I,b}$ and f_U with respect to a common σ-finite measure, then Ψ is given by (4) wherever the denominator is positive. If a Borel set A carries every informed order and no noise order, $G(A^c) = 0$ and $F_U(A) = 0$, then $\Psi = \mathbf{1}_A$ M-almost everywhere. If q_0 is an atom of both laws and $0 < \vartheta < 1$, then $\Psi(q_0) = \vartheta\, G(\{q_0\})/M(\{q_0\}) \in (0,1)$.*

(iii) *$\delta_A(q) := \Psi(q)\,\delta_1(q) + (1 - \Psi(q))\,\delta_0$ satisfies $\delta_A(Q) = \Pr(B = 1 \mid Q)$ almost surely, and $\delta_A(q) \in (0,1)$ for every q at which Ψ is defined.*

(iv) *The arbitrage order $x_A(q) = b\,(\mathrm{logit}\,\delta_A(q) - \mathrm{logit}\,\delta_1(q))$ is finite wherever Ψ is defined, hence M-almost everywhere; it moves the price to $\delta_A(q)$, which by Lemma 1, read from the price $\delta_1(q)$ (footnote 13) with $\delta_A(q)$ as its belief, is its unique best response, with conditional expected profit $b\,\mathrm{KL}(\delta_A(q) \,\|\, \delta_1(q))$.*

(v) *No moment of the belief is implied: at* $\delta_0 = \dfrac{1}{2}$ *the belief taking the values* $\sigma(n^2)$ *and* $\sigma(-n^2)$ *with probability* $3/(\pi^2 n^2)$ *each,* $n \geq 1$*, is interior and Bayes plausible with level at most* $\log 2$*, and* $\mathbb{E}|\operatorname{logit} P| = \infty$*.*

Proof. (i) T_b is finite on $(0,1)$, and $P \in (0,1)$ almost surely by Assumption 1, so $Q_I = T_b(P)$ is finite almost surely. For real q, $\delta_1(q) = \sigma(\operatorname{logit}\delta_0 + q/b)$ is the value of the logistic function at a real argument, hence interior. Finally $\delta_1(T_b(p)) = \sigma(\operatorname{logit} p) = p$.

(ii) For every Borel E, $\vartheta\, G(E) \leq \vartheta\, G(E) + (1-\vartheta)\, F_U(E) = M(E)$, so ϑG is absolutely continuous with respect to M and the Radon–Nikodym theorem gives a measurable $\Psi \geq 0$, unique up to M-null sets, with $\int_E \Psi\, dM = \vartheta G(E)$; since $\vartheta G(E) \leq M(E)$ for every E, $\Psi \leq 1$ M-almost everywhere, and we take the version with values in $[0,1]$. The law of Q is M, and $\Pr(J = I,\, Q \in E) = \vartheta\, \Pr(Q_I \in E) = \vartheta G(E) = \int_E \Psi\, dM = \mathbb{E}[\Psi(Q)\mathbf{1}\{Q \in E\}]$, which is the defining property of $\Pr(J = I \mid Q) = \Psi(Q)$; taking $E = \mathbb{R}$ gives $\mathbb{E}[\Psi(Q)] = \vartheta$. With densities against a common measure ν, M has density $\vartheta f_{I,b} + (1-\vartheta) f_U$ and ϑG has density $\vartheta f_{I,b}$, so the ratio (4) satisfies the defining integral identity wherever the denominator is positive, and the denominator vanishes only on an M-null set. If A carries every informed order and no noise order, then for every E, $\int_E \mathbf{1}_A\, dM = M(E \cap A) = \vartheta G(E \cap A) + (1-\vartheta) F_U(E \cap A) = \vartheta G(E)$, using $G(A^c) = 0$ and $F_U(A) = 0$, so $\mathbf{1}_A$ is a version of Ψ. At a shared atom q_0, taking $E = \{q_0\}$ gives $\Psi(q_0) M(\{q_0\}) = \vartheta G(\{q_0\})$, and both $\vartheta G(\{q_0\})$ and $(1-\vartheta) F_U(\{q_0\})$ are positive when $0 < \vartheta < 1$, so the ratio is interior.

(iii) Fix a Borel E. Conditional on $J = I$, $\Pr(B = 1 \mid P) = P$ almost surely by Assumption 1, which the pair (B, P) satisfies conditional on $J = I$ (Section 3.1), and Q_I is a function of P, so $\Pr(B = 1,\, Q \in E \mid J = I) = \mathbb{E}[P\, \mathbf{1}\{Q_I \in E\}] = \mathbb{E}[\delta_1(Q_I)\mathbf{1}\{Q_I \in E\}] = \int_E \delta_1\, dG$ by (i). Conditional on $J = N$, $\Pr(B = 1 \mid U) = \delta_0$ almost surely by Assumption 2, so $\Pr(B = 1,\, Q \in E \mid J = N) = \delta_0 F_U(E)$. Hence

$$\begin{aligned} \Pr(B = 1,\, Q \in E) \;&=\; \vartheta \int_E \delta_1\, dG + (1-\vartheta)\, \delta_0\, F_U(E) \\ &=\; \int_E \delta_1\, \Psi\, dM + \delta_0[M(E) - \vartheta G(E)] \;=\; \int_E [\Psi\, \delta_1 + (1-\Psi)\, \delta_0]\, dM, \end{aligned}$$

the second equality by the defining property of Ψ applied to the bounded measurable $\delta_1 \mathbf{1}_E$, which extends from indicators by the usual approximation, and by $(1-\vartheta) F_U(E) = M(E) - \vartheta G(E)$. This is the defining property of $\Pr(B = 1 \mid Q) = \delta_A(Q)$. Interiority: $\delta_A(q)$ is a convex combination of the interior values $\delta_1(q)$ and δ_0.

(iv) Both logits are finite by (i) and (iii), so $x_A(q)$ is finite. By the order formula of Section 2, from the price $\delta_1(q)$ the net order $b\,(\operatorname{logit}\delta - \operatorname{logit}\delta_1(q))$ moves the price to δ, and by Lemma 1 read from that price (footnote 13) a trader whose probability of $B = 1$ is $\delta_A(q)$ uniquely maximizes its expected settlement profit at $\delta = \delta_A(q)$, with value $b\,\mathrm{KL}(\delta_A(q) \,\|\, \delta_1(q))$; conditional on $Q = q$ the arbitrageur's probability of $B = 1$ is $\delta_A(q)$ by (iii).

(v) The weights sum to one, $\sum_{n\geq 1} 6/(\pi^2 n^2) = 1$. Both values are interior. Bayes plausibility: $\sigma(n^2) + \sigma(-n^2) = 1$, so each pair averages to $\frac{1}{2} = \delta_0$. The level: $\mathrm{KL}(p \,\|\, \frac{1}{2}) = \log 2 - h(p) \leq \log 2$ for every p. Finally $|\mathrm{logit}\, P| = n^2$ on the nth pair, so $\mathbb{E}|\mathrm{logit}\, P| = \sum_n 6/\pi^2 = \infty$. Since $\Pr(B = 1 \mid P) = P$ can be imposed by construction of the joint law of (B, P), the pair satisfies Assumption 1. Q.E.D.

Settlement profits are nonetheless integrable through their conditional expectations. For a trader who moves the price from δ_0 to a random interior price R' and whose conditional probability of $B = 1$ given R' is R', the conditional expected absolute profit is $b\,[R'|\log(R'/\delta_0)| + (1 - R')|\log((1 - R')/(1 - \delta_0))|] \leq b\,[2/e + |\log \delta_0| + |\log(1 - \delta_0)|]$, since $p|\log p| \leq 1/e$ on $(0, 1]$; this bound is used in Appendices A.6 and A.18 with $R' = P$ and $R' = \delta_A(Q)$.

A.4 Proof of Proposition 1

Throughout, write $R := \Pr(B = 1 \mid Q, J)$, the probability of the event given the order and its sender: $R = P$ on $\{J = I\}$, because conditional on $J = I$ the pair (B, P) satisfies Assumption 1 and $P = \delta_1(Q_I)$ is a function of Q (Lemma 4(i)); and $R = \delta_0$ on $\{J = N\}$ by Assumption 2. By the tower property and Lemma 4(iii), $\mathbb{E}[R \mid Q] = \mathbb{E}[B \mid Q] = \delta_A(Q)$.

(i) Since $\delta_A(Q)$ is a function of Q, the tower property gives $\Pr(B = 1 \mid \delta_A(Q)) = \mathbb{E}[\mathbb{E}[B \mid Q] \mid \delta_A(Q)] = \delta_A(Q)$ almost surely. Taking expectations, $\mathbb{E}[\delta_A(Q)] = \mathbb{E}[B] = \vartheta\,\mathbb{E}[P] + (1 - \vartheta)\,\delta_0 = \delta_0$, since on the informed branch $\mathbb{E}[P] = \delta_0$ by the Bayes plausibility of Assumption 1.

(ii) Let $t := \delta_A(Q)$, interior by Lemma 4(iii). Conditioning on Q, $\mathbb{E}[-B \log t - (1 - B)\log(1 - t) \mid Q] = -\delta_A(Q)\log t - (1 - \delta_A(Q))\log(1 - t) = h(\delta_A(Q))$, so the expected log-loss of the final price is $\mathbb{E}[h(\delta_A(Q))]$, finite since h is bounded. By the first identity in (20) at $r = \delta_A(Q)$ and $d = \delta_0$, and by $\mathbb{E}[\delta_A(Q)] = \delta_0$ from part (i),

$$\mathcal{F}(b) \;=\; \mathbb{E}[\mathrm{KL}(\delta_A(Q) \,\|\, \delta_0)] \;=\; -\mathbb{E}[h(\delta_A(Q))] - \delta_0 \log \delta_0 - (1 - \delta_0)\log(1 - \delta_0),$$

and the last two terms are $h(\delta_0)$, which is the claim.

(iii) Apply (19) at $r = R$, $d = \delta_0$, $a = \delta_A(Q)$, and write $g := R \log(\delta_A(Q)/\delta_0) + (1 - R)\log((1 - \delta_A(Q))/(1 - \delta_0))$, so that $\mathrm{KL}(R \,\|\, \delta_0) = \mathrm{KL}(R \,\|\, \delta_A(Q)) + g$ pointwise. The coefficients of R and $1 - R$ in g are functions of Q; splitting each into its positive and negative parts, which are nonnegative functions of Q, and using $\mathbb{E}[R \mid Q] = \delta_A(Q)$ on each part separately, $\mathbb{E}[g \mid Q] = \delta_A(Q)\log(\delta_A(Q)/\delta_0) + (1 - \delta_A(Q))\log((1 - \delta_A(Q))/(1 - \delta_0)) = \mathrm{KL}(\delta_A(Q) \,\|\, \delta_0)$, an M-almost everywhere finite quantity in $[0, M(\delta_0)]$; and $g^+ \leq \mathrm{KL}(R \,\|\, \delta_0) \leq M(\delta_0)$ is bounded, so g is integrable and $\mathbb{E}[g] = \mathcal{F}(b)$. The left side has $\mathbb{E}[\mathrm{KL}(R \,\|\, \delta_0)] = \vartheta\,\mathbb{E}[\mathrm{KL}(P \,\|\, \delta_0)] + (1 - \vartheta)\,\mathrm{KL}(\delta_0 \,\|\, \delta_0) = \vartheta\, y$, finite since $y = h(\delta_0) - \mathbb{E}[h(P)] \leq h(\delta_0)$ by the first identity in (20) and $\mathbb{E}[P] = \delta_0$. Hence $\mathbb{E}[\mathrm{KL}(R \,\|\, \delta_A(Q))] = \vartheta y - \mathcal{F}(b)$ is finite, and split-

ting the expectation over the sender gives the two-branch form (6): $\mathbb{E}[\mathrm{KL}(R \,\|\, \delta_A(Q))] = \vartheta\, \mathbb{E}[\mathrm{KL}(P \,\|\, \delta_A(Q_I))] + (1-\vartheta)\, \mathbb{E}[\mathrm{KL}(\delta_0 \,\|\, \delta_A(U))] = \mathscr{L}_T(b)$. So $\mathscr{I}(b) = \vartheta y - \mathscr{L}_T(b)$, and since both $\mathscr{L}_T(b)$ and $\mathscr{I}(b)$ are expectations of divergences, both are nonnegative and each is at most their sum ϑy. The upper bound $\mathscr{I}(b) \le h(\delta_0)$ is part (ii) with $\mathbb{E}[h(\delta_A(Q))] \ge 0$.

(iv) Conditional on Q, the expected log-loss of the first-round price is $-\delta_A(Q)\log\delta_1(Q) - (1-\delta_A(Q))\log(1-\delta_1(Q)) = h(\delta_A(Q)) + \mathrm{KL}(\delta_A(Q) \,\|\, \delta_1(Q))$ by the second identity in (20). Taking expectations in $[0,\infty]$ and using part (ii), the expected log-loss is $h(\delta_0) - \mathscr{I}(b) + D_A(b)$. Q.E.D.

A.5 Proof of Lemma 2

As in Appendix A.4, $R := \Pr(B = 1 \mid Q, J)$ equals P on the informed branch and δ_0 on the noise branch, and $\mathbb{E}[R \mid Q] = \delta_A(Q)$.

Conditional on $Q = q$ the variable R takes the value $\delta_1(q)$ with probability $\Psi(q)$ and the value δ_0 with probability $1 - \Psi(q)$, so every divergence with R as its first argument is conditionally bounded. Apply (19) at $r = R$, $d = \delta_1(Q)$, $a = \delta_A(Q)$: $\mathrm{KL}(R \,\|\, \delta_1(Q)) = \mathrm{KL}(R \,\|\, \delta_A(Q)) + g'$ with $g' := R\log(\delta_A/\delta_1) + (1-R)\log((1-\delta_A)/(1-\delta_1))$, and conditioning on Q, using $\mathbb{E}[R \mid Q] = \delta_A(Q)$, gives $\mathbb{E}[g' \mid Q] = \mathrm{KL}(\delta_A(Q) \,\|\, \delta_1(Q))$. Hence, conditional on Q,

$$\mathbb{E}[\mathrm{KL}(R \,\|\, \delta_1(Q)) \,|\, Q] \;=\; \mathbb{E}[\mathrm{KL}(R \,\|\, \delta_A(Q)) \,|\, Q] + \mathrm{KL}(\delta_A(Q) \,\|\, \delta_1(Q)),$$

an identity between finite nonnegative quantities. Taking expectations, which is legitimate in $[0,\infty]$ for nonnegative integrands, the right side becomes $\mathscr{L}_T(b) + D_A(b)$ by (6) and the definition of D_A. The left side splits over the sender: on the informed branch $R = P = \delta_1(Q_I)$ and the divergence vanishes; on the noise branch it is $\mathrm{KL}(\delta_0 \,\|\, \delta_1(U))$, with expectation ν_b. So $(1-\vartheta)\,\nu_b = \mathscr{L}_T(b) + D_A(b)$ in $[0,\infty]$, and since $\mathscr{L}_T(b) \le \vartheta y < \infty$ by Proposition 1(iii), this rearranges to $D_A(b) = (1-\vartheta)\,\nu_b - \mathscr{L}_T(b)$, with $D_A(b) = \infty$ exactly when $\nu_b = \infty$ and $\vartheta < 1$. Q.E.D.

A.6 Proof of Proposition 2

By Appendix A.1, applied to any two inventories, a trader who moves the LMSR price from p to p', both interior, realizes the settlement profit $b\log(p'/p)$ if $B = 1$ and $b\log((1-p')/(1-p))$ if $B = 0$. The first order moves the price from δ_0 to $\delta_1(Q)$ and the arbitrageur's from $\delta_1(Q)$ to $\delta_A(Q)$ (Lemma 4(iv)), so the two profits telescope state by state:

$$\pi_1 + \pi_A \;=\; b\log\frac{\delta_A(Q)}{\delta_0} \;\text{ if } B = 1, \qquad \pi_1 + \pi_A \;=\; b\log\frac{1-\delta_A(Q)}{1-\delta_0} \;\text{ if } B = 0. \tag{22}$$

Conditional on Q, the probability of $B = 1$ is $\delta_A(Q)$ (Lemma 4(iii)), so by (22) the conditional expected absolute value of $\pi_1 + \pi_A$ is $b\,[\delta_A|\log(\delta_A/\delta_0)| + (1-\delta_A)|\log((1-\delta_A)/(1-$

$\delta_0))|]$, at most $b\,[2/e+|\log\delta_0|+|\log(1-\delta_0)|]$ by the bound recorded after Lemma 4; hence $\pi_1+\pi_A$ is integrable. Its conditional expectation is $b\,[\delta_A(Q)\log(\delta_A(Q)/\delta_0)+(1-\delta_A(Q))\log((1-\delta_A(Q))/(1-\delta_0))] = b\,\mathrm{KL}(\delta_A(Q)\,\|\,\delta_0)$, and taking expectations gives $\mathbb{E}[\pi_1+\pi_A] = b\,\mathcal{I}(b)$. Q.E.D.

A.7 Proof of Proposition 3

Fix $b>0$, the law of P, and F_U; the dependence on ϑ is made explicit by writing $\mathcal{I}(\vartheta)$, $\mathscr{L}_T(\vartheta)$, Ψ^ϑ, δ_A^ϑ. Conditional on $J=I$ and $B=j$, the belief P has the law μ_j with $\mu_1(dp)=(p/\delta_0)\,\mu(dp)$ and $\mu_0(dp)=((1-p)/(1-\delta_0))\,\mu(dp)$, μ the law of P: this is Bayes' rule for the joint law $\Pr(B=1,P\in dp)=p\,\mu(dp)$ that Assumption 1 fixes, with $\Pr(B=1\mid J=I)=\mathbb{E}[P]=\delta_0$. Write G_j for the law of Q_I given $B=j$, the image of μ_j under T_b. Conditional on $J=N$, U is independent of B, since $\Pr(B=1\mid U)=\delta_0$ almost surely and B is binary. Let $\nu:=G_1+G_0+F_U$, a finite measure dominating all three laws, with densities g_1, g_0, f; write $g:=\delta_0 g_1+(1-\delta_0)g_0$ for the density of G, and note $g_1=(\delta_1/\delta_0)\,g$ and $g_0=((1-\delta_1)/(1-\delta_0))\,g$ ν-almost everywhere, by the change of variables that defines μ_j. The law of Q given $B=j$ then has density $k_j:=\vartheta g_j+(1-\vartheta)f$, and the law of Q has density $k:=\delta_0 k_1+(1-\delta_0)k_0$, both affine in ϑ.

An integral form of delivered information. Bayes' rule gives $\delta_A^\vartheta(q)=\Pr(B=1\mid Q=q)=\delta_0 k_1(q)/k(q)$ for ν-almost every q with $k(q)>0$, a version of the conditional probability of Lemma 4(iii). Substituting into the divergence,

$$\mathcal{I}(\vartheta) \;=\; \int F(k_1(q),k_0(q))\,\nu(dq), \qquad F(x_1,x_0) \;:=\; \sum_{j\in\{0,1\}} \pi_j x_j \log\frac{x_j}{\pi_1 x_1+\pi_0 x_0}, \tag{23}$$

with $\pi_1=\delta_0$, $\pi_0=1-\delta_0$ and the convention $0\log 0=0$, so that $F(k_1,k_0)=\delta_0 k_1\log(k_1/k)+(1-\delta_0)k_0\log(k_0/k)$: multiplying $\mathrm{KL}(\delta_A\,\|\,\delta_0)=\delta_A\log(\delta_A/\delta_0)+(1-\delta_A)\log((1-\delta_A)/(1-\delta_0))$ by k and using $\delta_A/\delta_0=k_1/k$ and $(1-\delta_A)/(1-\delta_0)=k_0/k$ gives the integrand. The integrand is nonnegative: writing $a_j:=\pi_j x_j$ and $c:=a_1+a_0$, $F=\sum_j a_j\log(a_j/c)-\sum_j a_j\log\pi_j$, and the log-sum inequality below gives $\sum_j a_j\log(a_j/(\pi_j c))\geq c\log(c/c)=0$.

The log-sum inequality and the convexity of F. For $a_1,a_0\geq 0$ and $c_1,c_0>0$, with $\bar a:=a_1+a_0$ and $\bar c:=c_1+c_0$,

$$\sum_j a_j\log\frac{a_j}{c_j} \;\geq\; \bar a\log\frac{\bar a}{\bar c}:$$

with $t\log t$ convex on $[0,\infty)$, Jensen's inequality with the weights $c_j/\bar c$ applied to the points a_j/c_j gives $\sum_j c_j\,(a_j/c_j)\log(a_j/c_j)\;\geq\;\bar c\,(\bar a/\bar c)\log(\bar a/\bar c)$, which is the claim. Consequently $(a,c)\mapsto a\log(a/c)$ is jointly convex on $[0,\infty)\times(0,\infty)$: for two points (a_j,c_j) and weights α_j, apply the inequality to $(\alpha_j a_j,\alpha_j c_j)$. Since $F(x_1,x_0)=\sum_j a_j\log(a_j/c)+$ (linear) with a_j and c linear in (x_1,x_0), F is convex on $[0,\infty)^2\setminus\{0\}$, and $F(0,0)=0$.

Convexity and monotonicity of $\mathcal{I}$. For each q, $\vartheta\mapsto(k_1(q),k_0(q))$ is affine, so $\vartheta\mapsto$

$F(k_1(q), k_0(q))$ is convex and nonnegative; the integral of a family of convex nonnegative functions that is finite at every ϑ is convex, and $\mathcal{I}(\vartheta) \le h(\delta_0)$ by Proposition 1(iii). At $\vartheta = 0$, $k_1 = k_0 = f$ and $F(f, f) = 0$, so $\mathcal{I}(0) = 0$; at $\vartheta = 1$, $\mathcal{I}(1) = y$ by Proposition 1(iii), since $\mathcal{L}_T(1) = 0$ there ($\Psi^1 \equiv 1$ by Lemma 4(ii) with $A = \mathbb{R}$, so $\delta_A^1(Q) = \delta_1(Q) = P$ on the informed branch, which has probability one). A convex function on $[0, 1]$ that attains its minimum at 0, as $\mathcal{I} \ge 0 = \mathcal{I}(0)$ does, is nondecreasing: for $0 \le s < t$, $\mathcal{I}(s) \le (1 - s/t)\,\mathcal{I}(0) + (s/t)\,\mathcal{I}(t) \le \mathcal{I}(t)$. The same inequality reads $\mathcal{I}(s)/s \le \mathcal{I}(t)/t$, so $\mathcal{I}(\vartheta)/\vartheta$ is nondecreasing on $(0, 1]$.

Continuity. A convex function is continuous on the interior of its domain. At the endpoints we bound the loss. For a random variable R' with values in $[0, 1]$ and conditional mean $m := \mathbb{E}[R' \mid \mathcal{G}]$, the second identity in (20) gives $\mathbb{E}[\mathrm{KL}(R' \,\|\, m) \mid \mathcal{G}] = h(m) - \mathbb{E}[h(R') \mid \mathcal{G}]$. Conditional on $Q = q$, R takes the value $\delta_1(q)$ with probability $\Psi^\vartheta(q)$ and δ_0 with probability $1 - \Psi^\vartheta(q)$, and its mean is $\delta_A^\vartheta(q)$, so $\mathbb{E}[\mathrm{KL}(R \,\|\, \delta_A^\vartheta(Q)) \mid Q] = h(\delta_A^\vartheta) - \Psi^\vartheta h(\delta_1) - (1 - \Psi^\vartheta) h(\delta_0)$, the Jensen gap of h at a two-point distribution. That gap is at most $h(\Psi^\vartheta)$: with H the Shannon entropy of a discrete random variable and $H(\cdot \mid \cdot)$ its conditional entropy, let J' be a Bernoulli(Ψ) index and B' Bernoulli with parameter $\delta_1(q)$ or δ_0 according to J'; the gap is $H(B') - H(B' \mid J')$, and the two chain-rule decompositions of the joint entropy, $H(B', J') = H(J') + H(B' \mid J') = H(B') + H(J' \mid B')$, give $H(B') - H(B' \mid J') = H(J') - H(J' \mid B') \le H(J') = h(\Psi)$. Taking expectations and applying Jensen's inequality to the concave h with $\mathbb{E}[\Psi^\vartheta(Q)] = \vartheta$ (Lemma 4(ii)),

$$0 \;\le\; \mathcal{L}_T(\vartheta) \;\le\; \min\{\vartheta\, y,\; h(\vartheta)\}, \tag{24}$$

the first bound from Proposition 1(iii). Hence $\mathcal{L}_T(\vartheta) \to 0$ as $\vartheta \to 0$ and as $\vartheta \to 1$, so $\mathcal{I} = \vartheta y - \mathcal{L}_T$ is continuous at both endpoints.

Strict increase. If $y = 0$ then $P = \delta_0$ almost surely, $Q_I = 0$, $\delta_1(Q_I) = \delta_0$, and $\delta_A^\vartheta(Q) = \delta_0$ almost surely by Lemma 4(iii), so $\mathcal{I} \equiv 0$. Let $y > 0$ and suppose $\mathcal{I}(\vartheta_1) = \mathcal{I}(\vartheta_2)$ for some $\vartheta_1 < \vartheta_2$. For $s \in (0, \vartheta_1)$ write $\vartheta_1 = \alpha s + (1 - \alpha)\vartheta_2$ with $\alpha \in (0, 1)$; convexity gives $\mathcal{I}(\vartheta_1) \le \alpha\,\mathcal{I}(s) + (1 - \alpha)\,\mathcal{I}(\vartheta_2) = \alpha\,\mathcal{I}(s) + (1 - \alpha)\,\mathcal{I}(\vartheta_1)$, so $\mathcal{I}(\vartheta_1) \le \mathcal{I}(s)$, while monotonicity gives the reverse; hence $\mathcal{I}$ is constant on $[0, \vartheta_2]$ and $\mathcal{I}(\vartheta_2) = \mathcal{I}(0) = 0$. Then $\mathrm{KL}(\delta_A^{\vartheta_2}(Q) \,\|\, \delta_0) = 0$ almost surely, so $\delta_A^{\vartheta_2}(Q) = \delta_0$, that is, $\Psi^{\vartheta_2}(Q)\,(\delta_1(Q) - \delta_0) = 0$ almost surely. On the informed branch $\Psi^{\vartheta_2} > 0$ G-almost everywhere, since $\vartheta_2 G(\{\Psi^{\vartheta_2} = 0\}) = \int_{\{\Psi^{\vartheta_2} = 0\}} \Psi^{\vartheta_2}\, dM = 0$ with $\vartheta_2 > 0$; so $\delta_1(Q_I) = \delta_0$ almost surely, $P = \delta_0$ almost surely, and $y = 0$, a contradiction.

(ii) Strict convexity and overlap. Fix q with $f(q) > 0$ and $g(q) > 0$. Then $k_1, k_0, k > 0$ on $[0, 1)$, and $\varphi_q(\vartheta) := F(k_1(q), k_0(q))$ is twice continuously differentiable there. Writing $d_j := g_j - f$ for the slope of k_j and $d := \delta_0 d_1 + (1 - \delta_0) d_0$ for that of k, differentiation gives $\varphi_q' = \sum_j \pi_j d_j \log(k_j/k) + \sum_j \pi_j d_j - d = \sum_j \pi_j d_j \log(k_j/k)$, and

$$\varphi_q'' \;=\; \sum_j \pi_j \frac{d_j^2}{k_j} - \frac{d^2}{k} \;=\; \frac{\pi_1 \pi_0\,(d_1 k_0 - d_0 k_1)^2}{k_1 k_0 k} \;=\; \frac{\delta_0(1 - \delta_0)\, f^2\,(g_1 - g_0)^2}{k_1 k_0 k},$$

the middle equality by the Lagrange identity $(u_1^2+u_0^2)(v_1^2+v_0^2)-(u_1v_1+u_0v_0)^2 = (u_1v_0-u_0v_1)^2$ with $u_j = \sqrt{\pi_j}\, d_j/\sqrt{k_j}$ and $v_j = \sqrt{\pi_j k_j}$, and the last by expanding $d_1k_0-d_0k_1 = (g_1-f)(\vartheta g_0+(1-\vartheta)f)-(g_0-f)(\vartheta g_1+(1-\vartheta)f) = f\,(g_1-g_0)$. Since $g_1-g_0 = g\,[\delta_1/\delta_0-(1-\delta_1)/(1-\delta_0)]$ vanishes exactly where $g=0$ or $\delta_1(q)=\delta_0$, that is $q=0$, the integrand φ_q is strictly convex in ϑ on $(0,1)$ exactly on the set $O := \{q \neq 0 :\ f(q)>0,\ g(q)>0\}$. Off O: where $f=0$ and $g>0$, $\varphi_q(\vartheta) = \vartheta\sum_j \pi_j g_j \log(g_j/g)$ is linear; where $g=0$, $\varphi_q \equiv 0$; and at $q=0$ with $f,g>0$, $g_1=g_0=g$ and $\varphi_q\equiv 0$. For $\vartheta_1<\vartheta_2$ and $\bar\vartheta := (\vartheta_1+\vartheta_2)/2$, the midpoint slack $\frac{1}{2}[\mathscr{I}(\vartheta_1)+\mathscr{I}(\vartheta_2)]-\mathscr{I}(\bar\vartheta)$ is the integral over ν of the pointwise slacks, each nonnegative and positive exactly on O; so $\mathscr{I}$ is strictly convex if $\nu(O)>0$, and if $\nu(O)=0$ then φ_q is affine for ν-almost every q with $\varphi_q(0)=0$, so $\mathscr{I}(\vartheta)=\vartheta\,\mathscr{I}(1)=\vartheta y$. It remains to identify $\nu(O)>0$ with the overlap condition of Section 3.4. If $\nu(O)=0$, the set $A := \{q\neq 0:\ g(q)>0\}\setminus O$ carries every nonzero informed order, $G(\{q\neq 0\}\setminus A) = \int_{\{q\neq 0,\, g>0\}\cap O} g\,d\nu = 0$, and no noise order, $F_U(A) = \int_A f\,d\nu = 0$ since $f=0$ on A: the laws do not overlap. Conversely, if a set A carries every nonzero informed order and no noise order, then $f=0$ ν-almost everywhere on A and $g=0$ ν-almost everywhere on $\{q\neq 0\}\setminus A$, so $\nu(O)=0$.

(iii) $\mathscr{L}_T = \vartheta y - \mathscr{I}$ is concave as the difference of an affine and a convex function, with $\mathscr{L}_T(0)=0$ and $\mathscr{L}_T(1) = y-\mathscr{I}(1) = 0$. Under overlap $\mathscr{I}$ is strictly convex, so $\mathscr{L}_T$ is strictly concave, and a strictly concave function that vanishes at both endpoints is positive in between. Without overlap $\mathscr{I}(\vartheta)=\vartheta y$ and $\mathscr{L}_T\equiv 0$. Q.E.D.

A.8 Delivered information and the liquidity

The paragraph closing Section 3.4 rests on the next lemma. Throughout, the law μ of P and the noise law F_U are fixed, and $\mathscr{I}(b)$, $\mathscr{L}_T(b)$ are the objects of Section 3 at the liquidity b; no moment of U is assumed.

Lemma 5 (The liquidity at a fixed experiment)**.** *Under Assumptions 1 and 2:*

(i) $\mathscr{L}_T(b)\to 0$, *and so* $\mathscr{I}(b)\to\vartheta\, y$, *both as* $b\downarrow 0$ *and as* $b\to\infty$.

(ii) *Let* $0<\vartheta<1$. *Then* $\mathscr{I}(b)<\vartheta y$ *if and only if the order laws overlap at* b. *They overlap at every* $b>0$ *if* F_U *has an absolutely continuous component whose density is positive Lebesgue-almost everywhere and* μ *has a nonzero absolutely continuous component. If* P *and* U *each take finitely many values, they overlap exactly at the finitely many liquidities* $b = u/(\operatorname{logit} p - \operatorname{logit}\delta_0)$ *formed from a noise atom* $u\neq 0$ *and a belief atom* $p\neq\delta_0$ *with the same sign of numerator and denominator.*

(iii) *A bounded function* f *on* $(0,\infty)$ *with* $\lim_{b\downarrow 0} f(b) = \lim_{b\to\infty} f(b) = L$ *and* $f(b_0)<L$ *at some* b_0 *is neither monotone, nor concave, nor convex.*

(iv) *Two experiments of the same level can deliver different information at the same liquidity and noise law: at* $\delta_0 = \frac{1}{2}$, $b = 1$, $\vartheta = \frac{1}{2}$ *and* U *uniform on* $\{-\log 3, \log 3\}$, *the belief uniform on* $\{\frac{1}{4}, \frac{3}{4}\}$ *and the belief that puts mass* $\mathrm{KL}(\frac{3}{4} \,\|\, \frac{1}{2})/\mathrm{KL}(\frac{9}{10} \,\|\, \frac{1}{2})$ *split equally on* $\{\frac{1}{10}, \frac{9}{10}\}$ *and the rest at* $\frac{1}{2}$ *have the same level, and the first delivers strictly less than* ϑy *while the second delivers* ϑy.

Proof. (i) With $R := \Pr(B = 1 \mid Q, J)$ as in Appendix A.4, $\mathbb{E}[h(R)] = \vartheta\,\mathbb{E}[h(P)] + (1 - \vartheta)\,h(\delta_0) = h(\delta_0) - \vartheta y$, so Proposition 1(ii)–(iii) give $\mathscr{L}_T(b) = \vartheta y - \mathscr{I}(b) = \mathbb{E}[h(\delta_A(Q))] - \mathbb{E}[h(R)]$. The binary entropy satisfies $|h(x) - h(x')| \le \omega(|x - x'|)$ with $\omega(s) := h(\min\{s, \frac{1}{2}\})$: h is concave with $h(0) = 0$, hence subadditive, so for $x < x'$ with $d := x' - x \le \frac{1}{2}$, $h(x') = h(x + d) \le h(x) + h(d)$ and $h(x) = h(1 - x) = h((1 - x') + d) \le h(x') + h(d)$, while for $d > \frac{1}{2}$ the difference is at most $\log 2 = \omega(d)$. The function ω is nondecreasing, concave and continuous on $[0, \infty)$ with $\omega(0) = 0$. Let t be any measurable forecast $t(Q) \in [0, 1]$. Then

$$\begin{aligned} 0 \;\le\; \mathscr{L}_T(b) \;&\le\; \mathbb{E}[|h(\delta_A(Q)) - h(t(Q))|] + \mathbb{E}[|h(t(Q)) - h(R)|] \\ &\le\; \omega(\mathbb{E}|\delta_A(Q) - t(Q)|) + \omega(\mathbb{E}|t(Q) - R|) \;\le\; 2\,\omega(\mathbb{E}|R - t(Q)|), \end{aligned}$$

by Jensen's inequality for the concave ω and, for the last step, $\mathbb{E}|\delta_A(Q) - t(Q)| = \mathbb{E}|\mathbb{E}[R - t(Q) \mid Q]| \le \mathbb{E}|R - t(Q)|$ with ω nondecreasing. It remains to exhibit forecasts whose error $\mathbb{E}|R - t(Q)|$ vanishes at each end. For small b take $t_b(q) := \delta_1(q)$ on $|q| \le \sqrt{b}$ and $t_b(q) := \delta_0$ elsewhere. On the informed branch $R = P = \delta_1(Q_I)$, so the error is $|P - \delta_0|\mathbf{1}\{|Q_I| > \sqrt{b}\} \le \mathbf{1}\{|\mathrm{logit}\, P - \mathrm{logit}\, \delta_0| > b^{-1/2}\}$, whose expectation tends to zero as $b \downarrow 0$ because $\mathrm{logit}\, P$ is finite almost surely; on the noise branch $R = \delta_0$ and the error is $|\delta_0 - \delta_1(U)|\mathbf{1}\{|U| \le \sqrt{b}\} \le \mathbf{1}\{0 < |U| \le \sqrt{b}\}$, since $\delta_1(0) = \delta_0$, whose expectation tends to zero by continuity of measure. For large b take $t_b(q) := \delta_0$ on $|q| \le \sqrt{b}$ and $t_b(q) := \delta_1(q)$ elsewhere. On the informed branch the error is $|P - \delta_0|\mathbf{1}\{|\mathrm{logit}\, P - \mathrm{logit}\, \delta_0| \le b^{-1/2}\}$, which is zero on $\{P = \delta_0\}$ and tends to zero pointwise on $\{P \ne \delta_0\}$, and is bounded by one, so its expectation tends to zero by dominated convergence; on the noise branch the error is at most $\mathbf{1}\{|U| > \sqrt{b}\}$, whose expectation tends to zero since U is finite. Hence $\mathscr{L}_T(b) \to 0$ at both ends, and $\mathscr{I} = \vartheta y - \mathscr{L}_T \to \vartheta y$.

(ii) The first claim is Proposition 3(iii) at the given ϑ and b: $\mathscr{L}_T(b) > 0$ exactly under overlap at b. For the sufficient condition, let $F_U = F_U^{\mathrm{ac}} + F_U^{\mathrm{s}}$ with a density $w > 0$ Lebesgue-almost everywhere, and $\mu = \mu^{\mathrm{ac}} + \mu^{\mathrm{s}}$ with $\mu^{\mathrm{ac}} \ne 0$. The image G^{ac} of μ^{ac} under T_b is absolutely continuous: for a Lebesgue-null N, $G^{\mathrm{ac}}(N) = \mu^{\mathrm{ac}}(\delta_1(N))$ and δ_1 is locally Lipschitz, so $\delta_1(N)$ is Lebesgue-null and $\mu^{\mathrm{ac}}(\delta_1(N)) = 0$. Let A be any Borel set with $F_U(A) = 0$. Then $\int_A w\,dq = 0$ forces A to be Lebesgue-null, so $G^{\mathrm{ac}}(A) = 0$, and $G(\{q \ne 0\} \setminus A) \ge$

$G^{\mathrm{ac}}(\mathbb{R} \setminus A) - G^{\mathrm{ac}}(\{0\}) = G^{\mathrm{ac}}(\mathbb{R}) > 0$: no set carries every nonzero informed order and no noise order. For finitely supported laws, the nonzero informed orders at b are the finitely many values $b\,(\operatorname{logit} p - \operatorname{logit} \delta_0)$ with $p \neq \delta_0$ a belief atom; the finite set of them carries every nonzero informed order, and it carries no noise order unless some noise atom $u \neq 0$ equals one of them, that is $b = u/(\operatorname{logit} p - \operatorname{logit} \delta_0) > 0$. So overlap holds exactly at those finitely many liquidities, and at every other b Proposition 3(ii) gives $\mathcal{I}(b) = \vartheta y$.

(iii) If f were nondecreasing, $f(b_0) \geq f(b)$ for every $b < b_0$ would give $f(b_0) \geq L$; nonincreasing is symmetric. If f were concave, then for $b_1 < b_0 < b_2$, $f(b_0) \geq \min\{f(b_1), f(b_2)\}$, and letting $b_1 \downarrow 0$, $b_2 \to \infty$ gives $f(b_0) \geq L$. If f were convex and bounded, it would be nonincreasing: if $f(b_1) < f(b_2)$ for some $b_1 < b_2$, convexity gives $f(b) \geq f(b_2) + (b - b_2)\,[f(b_2) - f(b_1)]/(b_2 - b_1) \to \infty$ as $b \to \infty$; and a nonincreasing f with $\lim_{b\to\infty} f = L$ satisfies $f \geq L$ everywhere.

(iv) Write $k(p) := \mathrm{KL}(p \,\|\, \frac{1}{2})$, symmetric about $\frac{1}{2}$, strictly convex, with $k(\frac{1}{2}) = 0$. Both beliefs are interior and Bayes plausible by symmetry, and their levels are $k(\frac{3}{4})$ and $[k(\frac{3}{4})/k(\frac{9}{10})]\,k(\frac{9}{10}) = k(\frac{3}{4})$. At $b = 1$ the first belief's orders are $\operatorname{logit}\frac{3}{4} = \log 3$ and $\operatorname{logit}\frac{1}{4} = -\log 3$, the two noise atoms. At the order $\log 3$ both laws have mass $\frac{1}{2}$, so $\Psi = \vartheta \cdot \frac{1}{2}/(\vartheta \cdot \frac{1}{2} + (1-\vartheta)\cdot\frac{1}{2}) = \frac{1}{2}$ (Lemma 4(ii)) and $\delta_A = \frac{1}{2}\cdot\frac{3}{4} + \frac{1}{2}\cdot\frac{1}{2} = \frac{5}{8}$; at $-\log 3$ likewise $\delta_A = \frac{3}{8}$; so $\mathcal{I}(1) = k(\frac{5}{8})$, and $k(\frac{5}{8}) = k(\frac{1}{2}\cdot\frac{3}{4} + \frac{1}{2}\cdot\frac{1}{2}) < \frac{1}{2}k(\frac{3}{4}) + \frac{1}{2}k(\frac{1}{2}) = \frac{1}{2}k(\frac{3}{4}) = \vartheta y$ by strict convexity. The second belief's orders at $b = 1$ are $\pm\log 9$ and 0, none of which is a noise atom, so the laws do not overlap and $\mathcal{I}(1) = \vartheta y$ by Proposition 3(ii). The comparison is a property of the pair at the liquidity: the first belief's orders $\pm b\log 3$ avoid the noise atoms at every $b \neq 1$, and the second belief's orders $\pm b\log 9$ meet them at $b = \frac{1}{2}$. Q.E.D.

A.9 Proof of Proposition 4

(i) By the hypothesis of part (i), c — and hence g_b — is differentiable on $(0, \bar{y})$. If y^* is an interior maximizer, Fermat's theorem gives $g_b'(y^*) = 0$. Because $g_b'(y) = b - c'(y)$, this is exactly (10).

(ii) Suppose the single-crossing sign conditions hold at $\hat{y}$. By the hypothesis of part (i), g_b is continuous on $[0, \bar{y}]$ and differentiable on $(0, \bar{y})$, so the mean value theorem applies on each closed subinterval of $[0, \bar{y}]$. For $y < \hat{y}$ there is $\xi \in (y, \hat{y})$ with

$$g_b(\hat{y}) - g_b(y) = h_b(\xi)\,(\hat{y} - y) > 0,$$

since $h_b(\xi) > 0$ for $\xi < \hat{y}$; for $y > \hat{y}$ there is $\xi \in (\hat{y}, y)$ with

$$g_b(y) - g_b(\hat{y}) = h_b(\xi)\,(y - \hat{y}) < 0,$$

since $h_b(\xi) < 0$ for $\xi > \hat{y}$. Hence every $y \neq \hat{y}$ yields $g_b(y) < g_b(\hat{y})$, so $\hat{y}$ is the unique global maximizer. Because $\hat{y} \in (0, \bar{y})$ and g_b is differentiable there, Fermat's theorem gives $h_b(\hat{y}) = g_b'(\hat{y}) = 0$.

(iii) *Increasing differences.* For $b_2 > b_1 > 0$ and $y_2 \geq y_1$ in $[0, \bar{y}]$, the cost terms cancel, leaving

$$[g_{b_2}(y_2) - g_{b_2}(y_1)] - [g_{b_1}(y_2) - g_{b_1}(y_1)] = (b_2 - b_1)\,(y_2 - y_1) \geq 0.$$

Extremal selections exist. For each $b > 0$, g_b is upper semicontinuous on the nonempty compact $[0, \bar{y}]$ — $b\,y$ is continuous and the cost is lower semicontinuous by the properties stated in Section 4 — so it attains its maximum by the Weierstrass theorem and $M(b)$ is nonempty; the superlevel set $M(b) = \{y \in [0, \bar{y}] : g_b(y) \geq \max_{[0,\bar{y}]} g_b\}$ is closed, hence compact, so $\underline{y}(b)$ and $\overline{y}(b)$ are well defined.

Monotone selections. Fix $0 < b_1 < b_2$. Write $\bar{y}_i = \overline{y}(b_i)$ and suppose, for contradiction, that $\bar{y}_1 > \bar{y}_2$. Then

$$0 \;\leq\; g_{b_1}(\bar{y}_1) - g_{b_1}(\bar{y}_2) \;\leq\; g_{b_2}(\bar{y}_1) - g_{b_2}(\bar{y}_2) \;\leq\; 0,$$

the first inequality by optimality of $\bar{y}_1$ at b_1, the middle by increasing differences applied to $\bar{y}_1 > \bar{y}_2$ and $b_2 > b_1$, and the last by optimality of $\bar{y}_2$ at b_2. Both differences therefore vanish, so $\bar{y}_1 \in M(b_2)$; but $\bar{y}_1 > \bar{y}_2 = \max M(b_2)$ is impossible. Hence $\overline{y}(b_1) \leq \overline{y}(b_2)$. The same chain with $\underline{y}_i := \underline{y}(b_i)$ shows that $\underline{y}_1 > \underline{y}_2$ would force $\underline{y}_2 \in M(b_1)$, contradicting $\underline{y}_2 < \underline{y}_1 = \min M(b_1)$; so $\underline{y}(b_1) \leq \underline{y}(b_2)$. If $M(b)$ is a singleton for every b, then $\underline{y} = \overline{y} = y^*_{\mathrm{LMSR}}$ is nondecreasing. This is the monotone-selection conclusion for a payoff with increasing differences on a chain (Topkis, 1998; Milgrom and Shannon, 1994). Q.E.D.

A.10 The one-visit supply schedule

The maximal selection is well defined and upper semicontinuous. The map $(b, y) \mapsto g_b(y) = b\,y - c(y)$ is jointly upper semicontinuous on $[0, \infty) \times [0, \bar{y}]$: $b\,y$ is continuous, and $-c$ is upper semicontinuous — the cost is lower semicontinuous by the properties stated in Section 4 — and constant in b. For fixed b, g_b is upper semicontinuous on the compact $[0, \bar{y}]$, so its maximizer set $M(b)$ is nonempty and compact by the semicontinuous Weierstrass argument in (ii) of Appendix A.22; hence $y^*(b) = \max M(b)$ is well defined. The trader's value $v(b) := \max_{[0,\bar{y}]} g_b$ is Lipschitz with constant $\bar{y}$: for $b \geq b'$ every $y \in [0, \bar{y}]$ satisfies $b'\,y - c(y) \leq b\,y - c(y) \leq b'\,y - c(y) + (b - b')\,\bar{y}$, since $0 \leq y \leq \bar{y}$, and maximizing the three expressions over y gives $0 \leq v(b) - v(b') \leq (b - b')\,\bar{y}$; in particular v is continuous on $[0, \infty)$, the endpoint $b = 0$ included. The graph of M is closed: if

$y_n \in M(b_n)$, $b_n \to b$, and $y_n \to y$, then joint upper semicontinuity and continuity of v give $g_b(y) \geq \limsup_n g_{b_n}(y_n) = \lim_n v(b_n) = v(b)$, so $y \in M(b)$. Upper semicontinuity of y^* follows: if $b_n \to b$ with $y^*(b_n)$ converging to $\limsup_n y^*(b_n)$ along a subsequence, the closed graph puts the limit in $M(b)$, so $\limsup_n y^*(b_n) \leq \max M(b) = y^*(b)$. Finally $g_0 = -c$, and the properties $c(0) = 0$ and $c(y) > 0$ for $y > 0$ of Section 4 give $M(0) = \{0\}$, so $y^*(0) = 0$.

The schedule qualifies. $S := y^*$ maps $[0, \infty)$ into $[0, \bar{y}]$, satisfies $S(0) = 0$, and is upper semicontinuous, so Propositions 14 and 15 apply with $\bar{S} = \bar{y}$ at $S = y^*$. They state their extrema and selections over the cap $[0, \lambda/\eta]$; by Proposition 14(i) the optima over $b \geq 0$ coincide with the cap optima, so the conclusions transfer to the market maker's unrestricted problem verbatim. Q.E.D.

A.11 Statement and proof of Proposition 9

In the Beta–Bernoulli experiment of Section 4.1, write

$$\pi_N(s) \;:=\; \Pr(S_N = s) \;=\; \mathbb{E}\Big[\tbinom{N}{s}\,\theta^s(1-\theta)^{N-s}\Big], \qquad s \in \{0, \ldots, N\},$$

for the (beta-binomial) prior-predictive law of S_N.

Proposition 9 (Acquisition by sample size)**.** *In the Beta–Bernoulli experiment the implemented level*

$$y(N) \;=\; \mathbb{E}[\mathrm{KL}(P_N \,\|\, \delta_0)] \;=\; h(\delta_0) - \sum_{s=0}^{N} \pi_N(s)\, h\!\left(\frac{\kappa\delta_0 + s}{\kappa + N}\right)$$

satisfies:

(i) *every $P_N \in (0,1)$ and the posterior is Bayes-plausible, $\mathbb{E}[P_N] = \delta_0$;*

(ii) *$N \mapsto y(N)$ is strictly increasing, with $y(0) = 0$ and $y(N) \uparrow y_\infty := h(\delta_0) - \mathbb{E}[h(\theta)]$, and $0 < y_\infty < h(\delta_0)$;*

(iii) *the trader's choice of the sample size at cost c_N and the direct choice of the level over the achievable set $\{y(N) : N \in \mathbb{N}_0\}$ at the induced cost $\tilde{c}(y(N)) = c_N(N)$ have the same value and the same supplied information.*

Each additional observation improves the informed trader's forecast because the observation and the event outcome depend on the same unknown probability θ. Learning θ cannot eliminate the uncertainty in the subsequent Bernoulli outcome. The residual uncertainty $\mathbb{E}[h(\theta)]$ therefore limits the information obtainable through sampling.[41]

[41] Public lotteries over sample sizes implement levels between those attainable with a fixed sample size; see Appendix C.

Proof. (i) The Beta prior is conjugate to the Bernoulli likelihood, so $\theta \mid X_{1:N} \sim \text{Beta}(\kappa\delta_0 + S_N,\ \kappa(1-\delta_0) + N - S_N)$, with mean $(\kappa\delta_0 + S_N)/(\kappa + N)$. Because $B \mid \theta \sim \text{Bernoulli}(\theta)$ is independent of $X_{1:N}$ given θ, $\Pr(B = 1 \mid X_{1:N}) = \mathbb{E}[\theta \mid X_{1:N}] = P_N$; the posterior therefore depends on the data only through S_N. Since $\kappa\delta_0 > 0$ and $\kappa(1-\delta_0) > 0$, $P_N \in (0,1)$ for every S_N and N. Bayes plausibility is the tower property, $\mathbb{E}[P_N] = \mathbb{E}[\mathbb{E}[B \mid X_{1:N}]] = \mathbb{E}[B] = \delta_0$.

(ii) Because P_N is S_N-measurable and S_N is sufficient for B (part (i)), Lemma 9 applied to the experiment $X_{1:N}$ gives $y(N) = I(B; X_{1:N}) = h(\delta_0) - \mathbb{E}[h(P_N)]$, which is the displayed formula, and $y(0) = 0$ because $P_0 = \delta_0$. For strict monotonicity, the chain rule for mutual information gives

$$y(N+1) - y(N) \;=\; I(B; X_{N+1} \mid X_{1:N}) \;\geq\; 0,$$

and this conditional mutual information is strictly positive because B and X_{N+1} are conditionally correlated: given $X_{1:N}$ they are exchangeable Bernoulli(θ) draws with common conditional mean θ, so

$$\text{Cov}(B, X_{N+1} \mid X_{1:N}) \;=\; \text{Var}(\theta \mid X_{1:N}) \;>\; 0,$$

the Beta posterior being nondegenerate. Hence y is strictly increasing. For the limit, $P_N = \mathbb{E}[B \mid X_{1:N}]$ is a bounded martingale, so it converges almost surely; the strong law of large numbers identifies the limit as θ (since $S_N/N \to \theta$ a.s.). As h is bounded, bounded convergence gives $\mathbb{E}[h(P_N)] \to \mathbb{E}[h(\theta)]$, so $y(N) \to h(\delta_0) - \mathbb{E}[h(\theta)] = y_\infty$. Finally $\theta \in (0,1)$ almost surely gives $\mathbb{E}[h(\theta)] > 0$, so $y_\infty < h(\delta_0)$; and h is strictly concave with θ nondegenerate, so strict Jensen gives $\mathbb{E}[h(\theta)] < h(\mathbb{E}[\theta]) = h(\delta_0)$, i.e. $y_\infty > 0$.

(iii) By (ii) the map $N \mapsto y(N)$ is injective, so each level set $\{M \in \mathbb{N}_0 : y(M) = y(N)\} = \{N\}$ is a singleton and the induced cost is $\tilde{c}(y(N)) = \inf\{c_N(M) : y(M) = y(N)\} = c_N(N)$, with the infimum attained. Lemma 11, applied to the family indexed by $\mathbb{N}_0$ with implemented levels $y(N)$ and $Y = \{y(N) : N \in \mathbb{N}_0\}$, gives the equality of values and the maximizer correspondence, which is the stated equivalence. Q.E.D.

A.12 Statement and proof of Corollary 2

We next characterize the informed trader's ex ante choice of sample size in $\max_{N \in \mathbb{N}_0}\{b\, y(N) - c_N(N)\}$.

Corollary 2 (Optimal sample size)**.** *Fix $b > 0$, suppose $c_N(N) \to \infty$ as $N \to \infty$, and write $\Delta f(N) := f(N+1) - f(N)$. Then $\max_{N \in \mathbb{N}_0}\{b\, y(N) - c_N(N)\}$ is attained, and:*

(i) (Marginal conditions.) *Any maximizer N^* satisfies $b\,\Delta y(N^*) \leq \Delta c_N(N^*)$ and, if $N^* \geq 1$, $b\,\Delta y(N^*-1) \geq \Delta c_N(N^*-1)$.*

(ii) (Stopping rule.) *If $N \mapsto b\,\Delta y(N) - \Delta c_N(N)$ changes sign once from positive to negative on $\mathbb{N}_0$, then $N^* = \min\{N \in \mathbb{N}_0 : b\,\Delta y(N) \leq \Delta c_N(N)\}$ is a maximizer, and it is unique*

if no term $b\,\Delta y(N) - \Delta c_N(N)$ *vanishes.*

Liquidity determines the expected trading profit from the information in each additional observation. Under the condition in part (ii), the first sample size at which the expected profit from another observation no longer exceeds its incremental cost is optimal. The trader selects this sample size before observing the data; the rule does not describe a decision to stop sampling in response to realized observations.

Proof. Existence. Write $\nu(N) := b\,y(N) - c_N(N)$. Since $y(N) < y_\infty$ (Proposition 9(ii)) and $c_N(N) \to \infty$, $\nu(N) \leq b\,y_\infty - c_N(N) \to -\infty$, so there is N_1 with $\nu(N) < \nu(0)$ for all $N > N_1$; the maximum of ν over the finite set $\{0, \ldots, N_1\}$ is attained and is the maximum over $\mathbb{N}_0$.

(i) If N^* is a maximizer, then $\nu(N^*) \geq \nu(N^* + 1)$, which rearranges to $b\,\Delta y(N^*) \leq \Delta c_N(N^*)$; and if $N^* \geq 1$, $\nu(N^*) \geq \nu(N^* - 1)$ rearranges to $b\,\Delta y(N^* - 1) \geq \Delta c_N(N^* - 1)$.

(ii) Write $\zeta(N) := b\,\Delta y(N) - \Delta c_N(N) = \Delta\nu(N)$. The set $\{N : \zeta(N) \leq 0\}$ is nonempty: otherwise ν would be strictly increasing, contradicting $\nu(N) \to -\infty$. Under the single-crossing hypothesis, with $\hat{N} := \min\{N : \zeta(N) \leq 0\}$ we have $\zeta(N) > 0$ for $N < \hat{N}$ and $\zeta(N) \leq 0$ for $N \geq \hat{N}$. Telescoping, for $N < \hat{N}$,

$$\nu(\hat{N}) - \nu(N) = \sum_{M=N}^{\hat{N}-1} \zeta(M) > 0,$$

and for $N > \hat{N}$,

$$\nu(N) - \nu(\hat{N}) = \sum_{M=\hat{N}}^{N-1} \zeta(M) \leq 0.$$

Hence $\hat{N}$ is a maximizer. If no $\zeta(M)$ vanishes, the second sum is strictly negative, so $\nu(N) < \nu(\hat{N})$ for every $N \neq \hat{N}$ and $\hat{N}$ is the unique maximizer. Q.E.D.

A.13 Statement and proof of Proposition 10

We now characterize the final price and trading profits when the noise-order distribution assigns zero probability to every individual order. Let $N^*(b)$ and $S(b)$ be as defined in Section 4.1.

Proposition 10 (Exact separation in the sampling technology)**.** *Under Assumptions 1 and 2, suppose* F_U *has no atoms, fix* $b > 0$ *and* $N \in \mathbb{N}_0$*, and let the informed trader hold the belief* P_N*.*

(i) (Identification.) $\Psi(Q) = 1$ *almost surely when the informed trader sent the order and* $\Psi(Q) = 0$ *almost surely when the noise trader sent it; hence* $\delta_A(Q_I) = P_N$ *and* $\delta_A(U) = \delta_0$ *almost surely.*

(ii) (The correction.) $x_A(Q_I) = 0$ *and* $x_A(U) = -U$ *almost surely; when the noise trader sent the order, its settlement profit and the arbitrageur's cancel in every state; and* $D_A(b) = (1-\vartheta)\,\nu_b$.

(iii) (Delivered information.) $\mathcal{L}_T(b) = 0$ *and* $\mathcal{I}(b) = \vartheta\, y(N)$. *If* $c_N(N) \to \infty$, *then* $N^*(b)$ *exists and, with the sample size chosen at liquidity* b, $\mathcal{I}(b) = \vartheta\, S(b)$.

At a fixed sample size N, the informed trader's belief, and hence its order, takes one of $N+1$ values. The noise trader submits an order in this finite set with probability zero. The arbitrageur can therefore distinguish informed orders from noise orders almost surely.[42]

The arbitrageur preserves the price change due to private information and reverses the price change due to noise. Its settlement profit on a noise order exactly offsets the noise trader's loss, so these two trades impose no net settlement payment on the market maker. The informed trader's acquired information therefore determines delivered information through $\mathcal{I}(b) = \vartheta\, y(N)$.

Proof. Fix $b > 0$ and N, and write $A_N := \{T_b(p_{N,s}) : 0 \le s \le N\}$ with $p_{N,s} := (\kappa\delta_0 + s)/(\kappa + N)$, the at most $N+1$ orders the informed trader can send: $Q_I = T_b(P_N) \in A_N$ surely. An atomless law on $\mathbb{R}$ charges no finite set, by additivity, so $F_U(A_N) = 0$.

(i) A_N carries every informed order and no noise order, so $\Psi = \mathbf{1}_{A_N}$ M-almost everywhere by Lemma 4(ii), at every $\vartheta \in [0,1]$. Hence $\Psi(Q) = 1$ almost surely on $\{J = I\}$, since $Q_I \in A_N$, and $\Psi(Q) = 0$ almost surely on $\{J = N\}$, since $U \notin A_N$ almost surely. By (5), $\delta_A(Q_I) = \delta_1(Q_I) = P_N$ (Lemma 4(i)) and $\delta_A(U) = \delta_0$ almost surely.

(ii) On the informed branch $\delta_A = \delta_1$, so $x_A(Q_I) = 0$. On the noise branch $x_A(U) = b\,(\operatorname{logit}\delta_0 - \operatorname{logit}\delta_1(U)) = -U$, since $\operatorname{logit}\delta_1(U) = \operatorname{logit}\delta_0 + U/b$. The noise trader's order moves the price from δ_0 to $\delta_1(U)$ and the arbitrageur's moves it back from $\delta_1(U)$ to δ_0, so by the state-contingent profit of Appendix A.1 the two settlement profits are negatives of each other in each state. Finally $D_A(b) = \mathbb{E}[\mathrm{KL}(\delta_A(Q) \,\|\, \delta_1(Q))] = \vartheta\cdot 0 + (1-\vartheta)\,\mathbb{E}[\mathrm{KL}(\delta_0 \,\|\, \delta_1(U))] = (1-\vartheta)\,\nu_b$.

(iii) By (i), $P_N = \delta_A(Q_I)$ and $\delta_0 = \delta_A(U)$ almost surely, so both terms of (6) vanish and $\mathcal{L}_T(b) = 0$; Proposition 1(iii) then gives $\mathcal{I}(b) = \vartheta\, y(N)$. If $c_N(N) \to \infty$, the existence step of Appendix A.12 shows that the maximizers of $b\, y(N) - c_N(N)$ over $\mathbb{N}_0$ lie in a finite set, so the maximizer set is finite and nonempty and its largest element $N^*(b)$ exists; applying the first two parts at $N = N^*(b)$ gives $\mathcal{I}(b) = \vartheta\, y(N^*(b)) = \vartheta\, S(b)$. Q.E.D.

The exact condition. The footnote to Proposition 10 states two equivalences. At fixed b, N and $0 < \vartheta < 1$, the informed order law is carried by the finite set A_N, so the order laws

[42] At fixed b and N, with $0 < \vartheta < 1$, zero type-confusion loss requires only that the noise-order distribution assign zero probability to each nonzero informed order. Across all liquidities and sample sizes, the loss is zero if and only if the noise-order distribution has no point mass away from zero. A shared order of zero entails no information loss because the informed belief then equals the prior.

overlap at b exactly when some nonzero element of A_N is a noise atom, and $T_b(p_{N,s}) \neq 0$ exactly when $p_{N,s} \neq \delta_0$; by Proposition 3(iii) the type-confusion loss is positive exactly under overlap, so $\mathscr{L}_T(b) = 0$ if and only if F_U puts no mass on $T_b(p_{N,s})$ for every s with $p_{N,s} \neq \delta_0$. Across all liquidities and sizes with F_U fixed in b, the loss vanishes at every (b, N) if F_U has no atom off zero, since each nonzero informed order is a single point; conversely, if F_U has an atom $u \neq 0$, then at $N = 1$ the beliefs $p_{1,1} = (\kappa\delta_0+1)/(\kappa+1) > \delta_0$ and $p_{1,0} = \kappa\delta_0/(\kappa+1) < \delta_0$ lie on either side of the prior, so with s chosen to make $\operatorname{logit} p_{1,s} - \operatorname{logit} \delta_0$ of the sign of u, the liquidity $b = u/(\operatorname{logit} p_{1,s} - \operatorname{logit} \delta_0) > 0$ puts an informed order on the atom and the loss is positive there.

A.14 Statement and proof of Proposition 11

The next proposition relates measurement precision to acquired information.

Proposition 11 (Acquisition by precision)**.** *In the Gaussian experiment the implemented level*

$$y(\rho) \;=\; \mathbb{E}[\mathrm{KL}(\hat{P}_\rho \,\|\, \delta_0)] \;=\; h(\delta_0) - \mathbb{E}[h(\hat{P}_\rho)]$$

satisfies:

(i) *every $P_\rho(t) \in (0, 1)$ and the posterior is Bayes-plausible, $\mathbb{E}[\hat{P}_\rho] = \delta_0$;*

(ii) *the family is Blackwell-ordered in ρ; the mean posterior variance m is continuous and strictly decreasing on $[0, \infty)$, with $m(0) = \delta_0(1 - \delta_0)$ and $m(\rho) \to 0$ as $\rho \to \infty$; and y is a continuous, strictly increasing bijection from $[0, \infty)$ onto $[0, h(\delta_0))$, continuously differentiable on $(0, \infty)$ with*

$$y'(\rho) \;=\; \rho\, m(\rho) \;>\; 0 \qquad (\rho > 0); \tag{25}$$

(iii) *for each $\bar{y} \in (0, h(\delta_0))$ a unique $\bar{\rho} > 0$ implements it ($y(\bar{\rho}) = \bar{y}$); writing $\rho(\cdot)$ for the inverse of y, the choice of the precision on $[0, \bar{\rho}]$ at cost c_ρ and the direct choice of the level on $[0, \bar{y}]$ at the induced acquisition cost $\tilde{c}(y) := c_\rho(\rho(y))$ have the same value and the same supplied information.*

Greater precision reduces the measurement error in a signal about the event itself. The informed trader can therefore approach full information by increasing precision. The continuous choice of precision implements every information level below full revelation, which gives the abstract acquisition technology a direct foundation in signal quality.

Proof. (i) Given $B = \beta$, $\beta \in \{0, 1\}$, $T_\rho \sim N(\rho\beta, 1)$ has density $\phi(t-\rho\beta)$, with ϕ the standard normal density. For $\rho > 0$ the posterior odds are $\frac{P_\rho(t)}{1-P_\rho(t)} = \frac{\delta_0}{1-\delta_0}\frac{\phi(t-\rho)}{\phi(t)} = \frac{\delta_0}{1-\delta_0}e^{\rho t-\rho^2/2}$, which rearranges to the displayed $P_\rho(t) \in (0, 1)$ for every finite t, ρ; for $\rho = 0$, $T_0 = Z$ is independent

of B, so $\hat{P}_0 = \delta_0 \in (0,1)$ and $y(0) = 0$. Bayes plausibility is $\mathbb{E}[\hat{P}_\rho] = \mathbb{E}[\Pr(B = 1 \mid T_\rho)] = \delta_0$, and Lemma 9 gives $y(\rho) = I(B; T_\rho) = h(\delta_0) - \mathbb{E}[h(\hat{P}_\rho)]$.

(ii) Blackwell order. Fix $0 \le \rho_1 < \rho_2$ and set $r := \rho_1/\rho_2 \in [0,1)$. With $Z' \sim N(0,1)$ independent of (B, T_{ρ_2}), $r\,T_{\rho_2} + \sqrt{1-r^2}\,Z' = \rho_1 B + (rZ + \sqrt{1-r^2}\,Z')$, and the noise $rZ + \sqrt{1-r^2}Z' \sim N(0,1)$ is independent of B, so $(B,\, r\,T_{\rho_2} + \sqrt{1-r^2}\,Z')$ and (B, T_{ρ_1}) have the same law. Hence T_{ρ_1} is a garbling of T_{ρ_2}, $B \to T_{\rho_2} \to T_{\rho_1}$ is a Markov chain, and the data-processing inequality of Lemma 9(ii) gives $y(\rho_1) \le y(\rho_2)$: y is nondecreasing.

Properties of m. Substituting $T_\rho = \rho B + Z$ into (11), $\operatorname{logit} \hat{P}_\rho = \operatorname{logit} \delta_0 + \rho Z + \rho^2(B - \frac{1}{2})$, which is continuous in $\rho \in [0,\infty)$ for each realization of (B, Z); since $\hat{P}_\rho(1 - \hat{P}_\rho) \le \frac{1}{4}$, dominated convergence makes $m(\rho) = \mathbb{E}[\hat{P}_\rho(1 - \hat{P}_\rho)]$ continuous on $[0,\infty)$. At zero, $\hat{P}_0 = \delta_0$ gives $m(0) = \delta_0(1 - \delta_0)$. As $\rho \to \infty$, the displayed log-odds tends to $+\infty$ almost surely on $\{B = 1\}$ and to $-\infty$ almost surely on $\{B = 0\}$, so $\hat{P}_\rho(1 - \hat{P}_\rho) \to 0$ almost surely and, by dominated convergence, $m(\rho) \to 0$.

m is strictly decreasing. Fix $0 \le \rho_1 < \rho_2$, let r, Z' be as above, and set $G := r\,T_{\rho_2} + \sqrt{1-r^2}\,Z'$, so that (B, G) has the law of (B, T_{ρ_1}) and hence $m(\rho_1) = \mathbb{E}[\operatorname{Var}(B \mid G)]$. Because $\sigma(T_{\rho_2}, G) = \sigma(T_{\rho_2}, Z')$ and Z' is independent of (B, T_{ρ_2}), conditioning on (T_{ρ_2}, G) is conditioning on T_{ρ_2} as far as B is concerned: $\mathbb{E}[B \mid T_{\rho_2}, G] = \hat{P}_{\rho_2}$ and $\operatorname{Var}(B \mid T_{\rho_2}, G) = \operatorname{Var}(B \mid T_{\rho_2})$. The law of total variance conditional on G then gives

$$\operatorname{Var}(B \mid G) = \mathbb{E}[\operatorname{Var}(B \mid T_{\rho_2}, G) \mid G] + \operatorname{Var}(\mathbb{E}[B \mid T_{\rho_2}, G] \mid G) = \mathbb{E}[\operatorname{Var}(B \mid T_{\rho_2}) \mid G] + \operatorname{Var}(\hat{P}_{\rho_2} \mid G),$$

and taking expectations, $m(\rho_1) = m(\rho_2) + \mathbb{E}[\operatorname{Var}(\hat{P}_{\rho_2} \mid G)]$. The added term is strictly positive: given (B, G), the pair (T_{ρ_2}, G) is bivariate normal with $\operatorname{Var}(T_{\rho_2} \mid B) = \operatorname{Var}(G \mid B) = 1$ and covariance r, so $\operatorname{Var}(T_{\rho_2} \mid B, G) = 1 - r^2 > 0$; the map $t \mapsto P_{\rho_2}(t)$ is strictly increasing (its log-odds is affine in t with slope $\rho_2 > 0$), so $\operatorname{Var}(\hat{P}_{\rho_2} \mid B, G) > 0$ almost surely, and by the law of total variance $\operatorname{Var}(\hat{P}_{\rho_2} \mid G) \ge \mathbb{E}[\operatorname{Var}(\hat{P}_{\rho_2} \mid B, G) \mid G] > 0$ almost surely. Hence $m(\rho_1) > m(\rho_2)$.

Derivative and strict monotonicity of y. y is continuously differentiable on $(0,\infty)$ with $y'(\rho) = \rho\, m(\rho)$ by the I–MMSE identity of Guo et al. (2005): for a Gaussian channel $Y = \sqrt{\gamma}\,X + Z$ with $Z \sim N(0,1)$ independent of X, $\frac{d}{d\gamma} I(X; Y) = \frac{1}{2}\mathbb{E}[(X - \mathbb{E}[X \mid Y])^2]$. Applying it with $X = B$ and $\gamma = \rho^2$ (so $T_\rho = \sqrt{\gamma}\,B + Z$) and using the chain rule $\frac{dy}{d\rho} = 2\rho \frac{dI}{d\gamma}$,

$$y'(\rho) \;=\; \rho\,\mathbb{E}[(B - \hat{P}_\rho)^2] \;=\; \rho\, m(\rho) \;>\; 0 \qquad (\rho > 0),$$

the middle equality because B is Bernoulli, so $\mathbb{E}[(B - \hat{P}_\rho)^2 \mid T_\rho] = \hat{P}_\rho(1 - \hat{P}_\rho)$; y' is continuous because m is, and it is strictly positive because m is strictly decreasing to its limit 0, hence everywhere positive.

Endpoints and bijection. As $\rho \to 0^+$ the conditional laws $N(\rho, 1)$ and $N(0, 1)$ merge, so $\hat{P}_\rho \to \delta_0$ and, h being bounded, $y(\rho) = h(\delta_0) - \mathbb{E}[h(\hat{P}_\rho)] \to 0 = y(0)$: y is continuous at zero. As $\rho \to \infty$, $\hat{P}_\rho \to \mathbf{1}\{B = 1\}$ almost surely (the log-odds display above) and $\mathbb{E}[h(\hat{P}_\rho)] \to 0$, giving $y(\rho) \to h(\delta_0)$. A continuous, strictly increasing map on $[0, \infty)$ with $y(0) = 0$ and unattained supremum $h(\delta_0)$ is a bijection onto $[0, h(\delta_0))$.

(iii) Since y is a continuous strictly increasing bijection onto $[0, h(\delta_0))$, every $\bar{y} \in (0, h(\delta_0))$ has a unique preimage $\bar{\rho} = \rho(\bar{y}) > 0$, and injectivity makes each level set a singleton, so the induced cost of the family is $\tilde{c}(y) = c_\rho(\rho(y))$, with the infimum attained. Lemma 11, applied to the family indexed by $\rho \in [0, \bar{\rho}]$ with implemented levels $y(\rho)$ and $Y = [0, \bar{y}]$, gives the equality of values and the maximizer correspondence, which is the stated equivalence. Q.E.D.

A.15 Statement and proof of Corollary 3

The next corollary characterizes the maximizer of (13) and ties it to the optimal level.

Corollary 3 (Optimal precision)**.** *Fix $b > 0$ and suppose c_ρ is differentiable on $(0, \infty)$.*

(i) (First-order condition.) *Any interior maximizer $\rho^* > 0$ of* (13) *satisfies*

$$b\,\rho^*\,m(\rho^*) \;=\; c_\rho'(\rho^*), \tag{26}$$

and the supplied level is $y^* = y(\rho^*)$, at which the induced cost of Proposition 11(iii) *is differentiable and satisfies the level-space condition* (10)*: $b = \tilde{c}'(y^*)$.*

(ii) (Existence, uniqueness, and participation.) *Suppose $\rho \mapsto c_\rho'(\rho)/\rho$ is nondecreasing on $(0, \infty)$ and c_ρ is not identically zero, and set $L_0 := \lim_{\rho \downarrow 0} c_\rho'(\rho)/\rho \in [0, \infty)$. Then* (13) *has a unique maximizer ρ^*: if $b\,\delta_0(1 - \delta_0) > L_0$, the maximizer is interior and is the unique solution of* (26)*; if $b\,\delta_0(1 - \delta_0) \le L_0$, it is $\rho^* = 0$.*

(iii) (The precision–level relation.) *For every $\rho > 0$,*

$$y(\rho) \;<\; \frac{1}{2}\,\delta_0(1 - \delta_0)\,\rho^2, \qquad \text{and} \qquad \frac{y(\rho)}{\dfrac{1}{2}\,\delta_0(1 - \delta_0)\,\rho^2} \;\to\; 1 \quad \text{as } \rho \to 0^+. \tag{27}$$

The informed trader values greater precision through the additional expected trading profit it generates. The cost condition in part (ii) makes acquisition cost convex in squared precision, ρ^2. In these units, the marginal information gain decreases as less uncertainty remains, while marginal cost increases. Acquisition begins only when liquidity is high enough to make a small increase in squared precision profitable. Greater prior uncertainty raises this initial return and lowers the liquidity required for acquisition.

Proof. Throughout, write $g(\rho) := b\,y(\rho) - c_\rho(\rho)$, continuous on $[0,\infty)$ (Proposition 11(ii); c_ρ continuous) and differentiable on $(0,\infty)$ with

$$g'(\rho) = b\,y'(\rho) - c'_\rho(\rho) = b\,\rho\,m(\rho) - c'_\rho(\rho),$$

using (25) and the hypothesis on c_ρ.

(i) If $\rho^* > 0$ is a maximizer of (13), Fermat's theorem gives $g'(\rho^*) = 0$, which is (26); the supplied level is $y^* = y(\rho^*)$ by definition. For the level-space form: y is strictly increasing and continuously differentiable at ρ^* with $y'(\rho^*) > 0$, so by the inverse function theorem its inverse $\rho(\cdot)$ is differentiable at y^* with $\rho'(y^*) = 1/y'(\rho^*)$; hence $\tilde{c} = c_\rho \circ \rho$ is differentiable at y^* with

$$\tilde{c}'(y^*) = \frac{c'_\rho(\rho^*)}{y'(\rho^*)} = b,$$

the last equality by (26).

(ii) Write $q(\rho) := c'_\rho(\rho)/\rho$ and $D(\rho) := b\,m(\rho) - q(\rho)$, so that $g'(\rho) = \rho\,D(\rho)$ on $(0,\infty)$ by the display above. The map q is nondecreasing by hypothesis and nonnegative (c_ρ is nondecreasing), so $L_0 = \lim_{\rho\downarrow 0} q(\rho)$ exists in $[0,\infty)$; and D is strictly decreasing on $(0,\infty)$ — the sum of the strictly decreasing $b\,m$ and the nonincreasing $-q$ — with $\lim_{\rho\downarrow 0} D(\rho) = b\,\delta_0(1-\delta_0) - L_0$, since m is continuous at zero with $m(0) = \delta_0(1-\delta_0)$ (Proposition 11(ii) for both properties of m).

Suppose first $b\,\delta_0(1-\delta_0) \le L_0$. For every $\rho > 0$, strict decrease of D gives $D(\rho) < D(\rho/2) \le \lim_{s\downarrow 0} D(s) = b\,\delta_0(1-\delta_0) - L_0 \le 0$, so $g' < 0$ on $(0,\infty)$; by the mean value theorem (and continuity of g at 0), g is strictly decreasing on $[0,\infty)$, so $\rho^* = 0$ is the unique maximizer.

Suppose instead $b\,\delta_0(1-\delta_0) > L_0$. Near zero D is positive: $D(s) \to b\,\delta_0(1-\delta_0) - L_0 > 0$ as $s \downarrow 0$. For all large ρ it is negative: $c_\rho \not\equiv 0$ with $c_\rho(0) = 0$ gives $a > 0$ with $c_\rho(a) > 0$; the mean value theorem on $[0,a]$ (c_ρ continuous there, differentiable inside) gives $\xi \in (0,a)$ with $c'_\rho(\xi) = c_\rho(a)/a > 0$; monotonicity of q gives $q(\rho) \ge q(\xi) > 0$ for $\rho \ge \xi$; and since $m(\rho) \to 0$ (Proposition 11(ii)), $D(\rho) \le b\,m(\rho) - q(\xi) < 0$ for all large ρ. Hence $\hat{\rho} := \sup\{\rho > 0 : D(\rho) > 0\}$ lies in $(0,\infty)$, and strict decrease of D sorts the signs around it: for $\rho < \hat{\rho}$ some $\rho' \in (\rho,\hat{\rho}]$ has $D(\rho') > 0$, so $D(\rho) > D(\rho') > 0$; and if some $\rho > \hat{\rho}$ had $D(\rho) \ge 0$, then $D > 0$ on $(\hat{\rho},\rho)$, contradicting the definition of $\hat{\rho}$. Thus $g' = \rho\,D > 0$ on $(0,\hat{\rho})$ and $g' < 0$ on $(\hat{\rho},\infty)$, and the mean value argument makes $\hat{\rho}$ the unique maximizer: for $0 \le \rho < \hat{\rho}$ the mean value theorem on $[\rho,\hat{\rho}]$ gives $\xi \in (\rho,\hat{\rho})$ with $g(\hat{\rho}) - g(\rho) = g'(\xi)(\hat{\rho} - \rho) > 0$, and for $\rho > \hat{\rho}$ it gives $\xi \in (\hat{\rho},\rho)$ with $g(\rho) - g(\hat{\rho}) = g'(\xi)(\rho - \hat{\rho}) < 0$. The maximizer is interior, so it solves (26) by part (i); and it is the unique solution, because any solution of (26) in $(0,\infty)$ is a zero of the strictly decreasing D, of which there is at most one.

(iii) By Proposition 11(ii), y is continuous on $[0,\infty)$ with $y(0) = 0$ and continuously differentiable on $(0,\infty)$, and $y'(s) = s\,m(s)$ extends continuously to $s = 0$ with value 0; the

fundamental theorem of calculus therefore gives, for every $\rho > 0$,

$$y(\rho) = \int_0^\rho s\, m(s)\, ds.$$

Since m is strictly decreasing with $m(0) = \delta_0(1-\delta_0)$, the integrand satisfies $s\, m(s) < s\, \delta_0(1-\delta_0)$ for every $s > 0$, so $y(\rho) < \frac{1}{2}\delta_0(1-\delta_0)\rho^2$. For the ratio: given $\epsilon \in (0,1)$, continuity of m at zero gives $\delta > 0$ with $m(s) \geq (1-\epsilon)\,\delta_0(1-\delta_0)$ on $[0,\delta]$; then for $\rho \in (0,\delta]$, $y(\rho) \geq (1-\epsilon)\,\frac{1}{2}\delta_0(1-\delta_0)\rho^2$, so the ratio in (27) lies in $[1-\epsilon, 1)$ for all such ρ. As ϵ was arbitrary, the ratio tends to 1 as $\rho \to 0^+$. Q.E.D.

A.16 Statement and proof of Proposition 12

The next proposition gives conditions under which the implementation schedule $\mathfrak{b}$ of (14) is the inverse of the informed trader's response to the liquidity.

Proposition 12 (Implementing a target precision)**.** *Suppose c_ρ is continuously differentiable on $(0,\infty)$ and that $\mathfrak{b}$ is strictly increasing on $(0,\infty)$ with range $(0,\infty)$. For every target $\rho_0 > 0$, at $b = \mathfrak{b}(\rho_0)$ the trader's problem (13) has the unique maximizer ρ_0; hence the best response to any $b > 0$ is $\rho^*(b) = \mathfrak{b}^{-1}(b)$.*

The implementation schedule gives the liquidity needed to induce the informed trader to choose a target precision. A higher target requires a higher reward per unit of information because the induced marginal acquisition cost rises. The market maker can therefore express its liquidity choice as a choice of the precision it induces.[43]

A positive lower limit of the implementation schedule creates a minimum liquidity for information acquisition.[44]

Proof. Two preliminary observations. First, a strictly increasing function from $(0,\infty)$ onto the interval $(0,\infty)$ has no jumps, so $\mathfrak{b}$ is continuous, and its inverse $\mathfrak{b}^{-1}$ is a continuous, strictly increasing bijection of $(0,\infty)$; moreover $\lim_{\rho\to 0^+} \mathfrak{b}(\rho) = 0$ and $\lim_{\rho\to\infty} \mathfrak{b}(\rho) = \infty$ (the limits are the infimum and supremum of the range). Second, by Proposition 11(ii), m is strictly positive on $[0,\infty)$ (it is strictly decreasing with limit 0), so $c'_\rho(\rho) = \mathfrak{b}(\rho)\,\rho\, m(\rho) > 0$ on $(0,\infty)$ and c_ρ is strictly increasing.

[43] The chain rule gives $\tilde{c}'(y) = \mathfrak{b}(\rho(y))$. For $c_\rho(\rho) = \kappa_c \rho^p/p$ with $\kappa_c > 0$ and $p > 2$, $\mathfrak{b}(\rho) = \kappa_c \rho^{p-2}/m(\rho)$ is strictly increasing from zero to infinity. More generally, the same conclusion holds when $c'_\rho(\rho)/\rho$ is positive and increasing with limit zero at the origin: m is positive and strictly decreasing to zero.

[44] Suppose c_ρ is continuously differentiable on $(0,\infty)$ and $\mathfrak{b}$ is strictly increasing with range $(\mathfrak{b}(0^+),\infty)$, where $\mathfrak{b}(0^+) > 0$. Then the response has the *threshold form*: $\rho^*(b) = 0$ for $b \leq \mathfrak{b}(0^+)$ and $\rho^*(b) = \mathfrak{b}^{-1}(b)$ for $b > \mathfrak{b}(0^+)$. For the quadratic cost, $\mathfrak{b}(0^+) = \kappa_c/[\delta_0(1-\delta_0)]$. The proof below establishes this extension.

Fix $\rho_0 > 0$ and set $b = \mathfrak{b}(\rho_0)$. For $\rho > 0$ the trader objective $g(\rho) = b\,y(\rho) - c_\rho(\rho)$ has derivative

$$g'(\rho) = b\,\rho\,m(\rho) - c'_\rho(\rho) = \rho\,m(\rho)\,[b - \mathfrak{b}(\rho)],$$

by (25) and (14). Since $\rho\,m(\rho) > 0$ and $\mathfrak{b}$ is strictly increasing, $g' > 0$ on $(0, \rho_0)$ and $g' < 0$ on (ρ_0, ∞); the mean-value argument of Appendix A.15(ii) then makes ρ_0 the unique maximizer of g on $[0, \infty)$. Conversely, for any $b > 0$ the target $\rho_0 = \mathfrak{b}^{-1}(b)$ exists, and the same argument gives $\rho^*(b) = \mathfrak{b}^{-1}(b)$.

If instead $\mathfrak{b}$ is strictly increasing on $(0, \infty)$ with range $(\mathfrak{b}(0^+), \infty)$ for some $\mathfrak{b}(0^+) > 0$, the threshold form of the footnote to Proposition 12 follows by the same argument: $\mathfrak{b}$ is again continuous, now with $\lim_{\rho \to 0^+} \mathfrak{b}(\rho) = \mathfrak{b}(0^+)$; for $b > \mathfrak{b}(0^+)$ the target $\rho_0 = \mathfrak{b}^{-1}(b)$ exists and is the unique maximizer as above; and for $b \le \mathfrak{b}(0^+)$ the derivative $g'(\rho) = \rho\,m(\rho)\,[b - \mathfrak{b}(\rho)]$ is negative on $(0, \infty)$, since $\mathfrak{b}(\rho) > \mathfrak{b}(0^+) \ge b$, so the same mean-value argument makes g strictly decreasing on $[0, \infty)$ and $\rho^*(b) = 0$. Q.E.D.

A.17 Statement and proof of Proposition 13

We next characterize what the final price reveals when the informed trader observes the Gaussian signal. Write the noise-order distribution as $F_U = F_U^{\mathrm{ac}} + F_U^{\mathrm{s}}$. The component F_U^{ac} has density w, so $F_U^{\mathrm{ac}}(A) = \int_A w(q)\,dq$ for every Borel set A. The remaining component F_U^{s} assigns all its probability mass to a set of zero Lebesgue measure; it includes any point masses. Either component may be zero. Thus $F_U^{\mathrm{ac}} \neq 0$ means that, with positive probability, the noise order is drawn from a distribution with a probability density.

Write $\mathscr{I}(b, \rho)$ and $\mathscr{L}_T(b, \rho)$ for delivered information and the type-confusion loss when the informed trader chooses precision ρ. Under the hypotheses of Proposition 12, or the threshold form stated in its footnote, let $\mathscr{I}(b) := \mathscr{I}(b, \rho^*(b))$.

Proposition 13 (Delivered information in the precision technology)**.** *Under Assumptions 1 and 2, let the informed trader hold the measurement of precision $\rho > 0$ and fix $b > 0$.*

(i) (The order law.) *Given $B = j \in \{0, 1\}$, the informed order Q_I is Gaussian with mean $(2j-1)\,b\rho^2/2$ and variance $b^2\rho^2$; its density*

$$f_{I,b}(q) \;=\; \frac{1}{b\rho}\Big[\delta_0\,\phi\Big(\frac{q/b - \rho^2/2}{\rho}\Big) + (1-\delta_0)\,\phi\Big(\frac{q/b + \rho^2/2}{\rho}\Big)\Big],$$

with ϕ the standard normal density, is positive on the whole line.

(ii) (Delivered information.) *For $\vartheta > 0$, the type posterior is $\Psi(q) = \vartheta f_{I,b}(q)/[\vartheta f_{I,b}(q) + (1-\vartheta)\,w(q)]$ for Lebesgue-almost every q, with $\Psi = 0$ F_U^{s}-almost everywhere, and, with*

$$\delta_A(q) = \Psi(q)\,\delta_1(q) + (1-\Psi(q))\,\delta_0,$$

$$\mathcal{I}(b,\rho) \;=\; \int_{\mathbb{R}} [\vartheta\, f_{I,b}(q) + (1-\vartheta)\, w(q)]\, \mathrm{KL}(\delta_A(q) \,\|\, \delta_0)\, dq. \tag{28}$$

(iii) (Overlap.) *For $0 < \vartheta < 1$, $\mathcal{L}_T(b,\rho) > 0$ if and only if $F_U^{\mathrm{ac}} \neq 0$.*

(iv) (Continuity and the two limits.) *Under the hypotheses of Proposition 12, or under the threshold form of its footnote, $b \mapsto \mathcal{I}(b)$ is continuous on $(0,\infty)$, $\mathcal{I}(b) \to 0$ as $b \downarrow 0$, and $\mathcal{I}(b) \to \vartheta\, h(\delta_0)$ as $b \to \infty$.*

The Gaussian signal induces a positive density of informed orders at every order size. Wherever the noise density is also positive, the arbitrageur assigns positive probability to both trader types. It therefore only partly reverses the first-round price change, which creates a type-confusion loss. Noise orders drawn from the singular component are distinguishable from informed orders almost surely and contribute no such loss. In particular, a discrete noise-order distribution permits perfect identification in the Gaussian technology, whereas a noise-order distribution with a density permits perfect identification in the sampling technology.

When precision is endogenous, low liquidity weakens the incentive to acquire information, so delivered information vanishes as liquidity approaches zero. This differs from the fixed-experiment limit, where the informed trader retains its information as liquidity falls. At large liquidity, the informed trader approaches full information and the type-confusion loss vanishes. The limiting information delivered is therefore full information weighted by the probability of informed arrival.

Throughout, n_j denotes the density of Q_I given $B = j$, $n := \delta_0 n_1 + (1-\delta_0) n_0$ the marginal density $f_{I,b}$, w a finite version of the density of F_U^{ac}, and N a Lebesgue-null Borel set carrying F_U^{s}, so that $F_U^{\mathrm{s}}(N^c) = 0$. Given $B = j$ the noise order has density w as well, since U is independent of B on the noise branch (Appendix A.7), so the law of Q given $B = j$ restricted to N^c has the Lebesgue density $k_j := \vartheta n_j + (1-\vartheta) w$, and the law of Q restricted to N^c has density $k := \delta_0 k_1 + (1-\delta_0) k_0 = \vartheta n + (1-\vartheta) w$.

Proof of parts (i)–(iii). *(i)* By (11), $Q_I = b\,(\operatorname{logit} P_\rho(T_\rho) - \operatorname{logit}\delta_0) = b\,(\rho T_\rho - \rho^2/2)$, and given $B = j$ the statistic $T_\rho = \rho B + Z$ is $N(j\rho, 1)$ (Appendix A.14), so Q_I is Gaussian with mean $b\,(j\rho^2 - \rho^2/2) = (2j-1)\, b\rho^2/2$ and variance $b^2\rho^2$. Its marginal density is the mixture of the two Gaussian densities with the prior weights, which is the displayed formula, positive on the whole line as a positive combination of positive functions.

(ii) Let $\vartheta > 0$. The informed order law G has the density n, so $G(N) = 0$. Define $\tilde{\Psi} := \vartheta n/(\vartheta n + (1-\vartheta) w)$ on N^c and $\tilde{\Psi} := 0$ on N. For every Borel E,

$$\int_E \tilde{\Psi}\, dM = \int_{E \cap N^c} \frac{\vartheta n}{\vartheta n + (1-\vartheta) w}\, [\vartheta n + (1-\vartheta) w]\, dq + 0 = \vartheta\, G(E \cap N^c) = \vartheta\, G(E),$$

since on N^c the law M has the Lebesgue density $\vartheta n + (1-\vartheta)w$. By Lemma 4(ii), $\tilde{\Psi}$ is a version of Ψ: the type posterior is the displayed ratio for Lebesgue-almost every q and zero F_U^{s}-almost everywhere. Then $\delta_A = \Psi\delta_1 + (1-\Psi)\delta_0$ equals δ_0 on N, where the divergence from the prior vanishes, so $\mathscr{I}(b,\rho) = \int \mathrm{KL}(\delta_A \,\|\, \delta_0)\, dM = \int_{N^c} [\vartheta n + (1-\vartheta)w] \mathrm{KL}(\delta_A(q) \,\|\, \delta_0)\, dq$, which is (28) because the integrand vanishes on the null set N.

(iii) Let $0 < \vartheta < 1$. By Proposition 3(iii), $\mathscr{L}_T(b,\rho) > 0$ exactly when the order laws overlap at b. If $F_U^{\mathrm{ac}} \neq 0$, the set $\{w > 0\}$ has positive Lebesgue measure; for any Borel A with $F_U(A) = 0$, $\int_A w\, dq = 0$ forces $\{w > 0\} \cap A$ to be Lebesgue-null, so $G(\{q \neq 0\} \setminus A) \geq \int_{\{w>0\}\setminus(A\cup\{0\})} n\, dq > 0$ since $n > 0$ everywhere: no set carries every nonzero informed order and no noise order, and the laws overlap. If $F_U^{\mathrm{ac}} = 0$, then $F_U = F_U^{\mathrm{s}}$ is carried by N, the set N^c carries every informed order, $G(N) = 0$, and no noise order, $F_U(N^c) = 0$, so the laws do not overlap. Q.E.D.

The remaining part uses the regularity of $\mathscr{I}$ in both arguments, which the next lemma records; its second part supplies the differentiability that Theorem 1(iv) needs in this technology.

Lemma 6 (Regularity of delivered information in the precision technology)**.** *Under Assumptions 1 and 2, for every noise law:*

(i) $(b,\rho) \mapsto \mathscr{I}(b,\rho)$ *is continuously differentiable on* $(0,\infty)^2$*, with, for* $\theta \in \{b,\rho\}$*,*

$$\partial_\theta \mathscr{I}(b,\rho) \;=\; \vartheta \sum_{j\in\{0,1\}} \pi_j \int_{\mathbb{R}} \partial_\theta n_j(q)\, \log \frac{k_j(q)}{k(q)}\, dq, \qquad \pi_1 = \delta_0,\ \pi_0 = 1-\delta_0.$$

(ii) *Under the hypotheses of Proposition 12, if* $\mathfrak{b}$ *is differentiable at* $\rho^*(b)$ *with* $\mathfrak{b}'(\rho^*(b)) > 0$*, then* $b \mapsto \mathscr{I}(b) = \mathscr{I}(b,\rho^*(b))$ *is differentiable at* b*, with* $\mathscr{I}'(b) = \partial_b \mathscr{I}(b,\rho^*(b)) + \partial_\rho \mathscr{I}(b,\rho^*(b))/\mathfrak{b}'(\rho^*(b))$*.*

Proof. (i) If $\vartheta = 0$ then $\mathscr{I} \equiv 0$; let $\vartheta > 0$. As in Appendix A.7, with Lebesgue measure on N^c as the dominating measure, $\delta_A = \delta_0 k_1/k$ and $\mathscr{I}(b,\rho) = \int \iota(b,\rho,q)\, dq$ with $\iota := \sum_j \pi_j k_j \log(k_j/k) \geq 0$, the singular part contributing nothing by part (ii); $k \geq \vartheta n > 0$ everywhere. Two bounds hold for every (b,ρ) and q. First, $k_j/k \leq 1/\pi_j$, since $k \geq \pi_j k_j$. Second, k_j/k is the mediant of n_j/n and 1 when $w > 0$, and equals n_j/n when $w = 0$, so $k_j/k \geq \min\{n_j/n, 1\}$; and $\log n_1 - \log n_0 = q/b$ by the two Gaussian exponents, so $n_1/n = 1/(\delta_0 + (1-\delta_0)e^{-q/b})$ and $n_0/n = 1/((1-\delta_0) + \delta_0 e^{q/b})$, both at least $1/(1+e^{|q|/b}) \geq e^{-|q|/b}/2$. Hence

$$\left|\log \frac{k_j(q)}{k(q)}\right| \;\leq\; \frac{|q|}{b} + \log 2 + \log \frac{1}{\min\{\delta_0, 1-\delta_0\}}.$$

For fixed q, $\theta \mapsto \iota$ is differentiable with $\partial_\theta \iota = \sum_j \pi_j\, \partial_\theta k_j \log(k_j/k)$, the nonlogarithmic terms cancelling because $\sum_j \pi_j \partial_\theta k_j = \partial_\theta k$ (the computation of Appendix A.7), and $\partial_\theta k_j = \vartheta\, \partial_\theta n_j$.

Each n_j is a Gaussian density in q whose mean and standard deviation are smooth positive functions of (b,ρ), so on a compact $K \subset (0,\infty)^2$ the derivatives $\partial_\theta n_j$ are bounded by $C_K(1+q^2)\,e^{-(|q|-m_K)_+^2/(2s_K^2)}$ for constants C_K, m_K, s_K, and the product with the logarithmic bound above is dominated on K by an integrable function of q. By the mean value theorem the difference quotients of ι in θ are dominated by the same function, so dominated convergence gives differentiation under the integral sign, the displayed formula, and, since $\partial_\theta \iota$ is continuous in (b,ρ) for fixed q and $k>0$, the continuity of the partial derivatives on K. As K is arbitrary, $\mathcal{I}$ is continuously differentiable on $(0,\infty)^2$.

(ii) $\mathfrak{b}^{-1}$ is continuous and strictly increasing (Appendix A.16), so if $\mathfrak{b}$ is differentiable at $\rho_0 := \rho^*(b_0)$ with $\mathfrak{b}'(\rho_0) > 0$, then $\mathfrak{b}^{-1}$ is differentiable at b_0 with derivative $1/\mathfrak{b}'(\rho_0)$: for $b \to b_0$ the points $\rho := \mathfrak{b}^{-1}(b)$ tend to ρ_0 with $\rho \neq \rho_0$, and $[\mathfrak{b}^{-1}(b) - \mathfrak{b}^{-1}(b_0)]/(b-b_0) = (\rho-\rho_0)/(\mathfrak{b}(\rho)-\mathfrak{b}(\rho_0)) \to 1/\mathfrak{b}'(\rho_0)$. The map $b \mapsto (b, \mathfrak{b}^{-1}(b))$ is therefore differentiable at b_0, and the chain rule for the continuously differentiable $\mathcal{I}$ of part (i) gives the formula. Q.E.D.

Proof of part (iv). Under the hypotheses of Proposition 12, $\rho^* = \mathfrak{b}^{-1}$ is continuous on $(0,\infty)$ with $\rho^*(b) \to 0$ as $b \downarrow 0$ and $\rho^*(b) \to \infty$ as $b \to \infty$ (Appendix A.16), so $b \mapsto \mathcal{I}(b, \rho^*(b))$ is continuous on $(0,\infty)$ by Lemma 6(i). Small liquidity: $0 \le \mathcal{I}(b,\rho^*(b)) \le \vartheta\, y(\rho^*(b))$ by Proposition 1(iii), and $y(\rho^*(b)) \to y(0) = 0$ by the continuity of y at zero (Proposition 11(ii)). Large liquidity: $\mathcal{I}(b) = \vartheta\, y(\rho^*(b)) - \mathcal{L}_T(b,\rho^*(b))$ with $y(\rho^*(b)) \to h(\delta_0)$ (Proposition 11(ii)), and the type-confusion loss vanishes uniformly in the experiment: the large-liquidity forecast of Lemma 5(i), $t_b(q) = \delta_0$ on $|q| \le \sqrt{b}$ and $\delta_1(q)$ elsewhere, has informed-branch error $\mathbb{E}[|P-\delta_0|\mathbf{1}\{|\operatorname{logit} P - \operatorname{logit}\delta_0| \le b^{-1/2}\}] \le b^{-1/2}/4$, because σ is $\dfrac{1}{4}$-Lipschitz so $|P-\delta_0| \le |\operatorname{logit} P - \operatorname{logit}\delta_0|/4$, whatever the law of P, and noise-branch error at most $\Pr(|U| > \sqrt{b})$; so $\mathcal{L}_T(b,\rho) \le 2\,\omega(b^{-1/2}/4 + \Pr(|U| > \sqrt{b})) \to 0$ uniformly in ρ, with ω the modulus of that lemma. Hence $\mathcal{I}(b) \to \vartheta\, h(\delta_0)$. If the response is instead the threshold form of the footnote to Proposition 12, $\rho^*(b) = 0$ for $b \le \mathfrak{b}(0^+) > 0$ and $\rho^*(b) = \mathfrak{b}^{-1}(b)$ above it, the same argument applies above the threshold, $\mathcal{I}(b) = 0$ at and below it, since a belief of precision zero equals the prior, and $\mathcal{I}$ is continuous at the threshold by the squeeze $0 \le \mathcal{I}(b) \le \vartheta\, y(\rho^*(b)) \to 0$. Q.E.D.

A.18 The ledger, the objective, and the market with no arbitrageur

The closed market. At $b = 0$ no trade executes: the informed trader's payoff is zero at every level and it acquires nothing, since $c(y) > 0$ for $y > 0$, no order moves the price, and the market maker pays and collects nothing, so $V_\lambda(0) = 0$.

The ledger (15). Let $b > 0$ and $\mathbb{E}|U| < \infty$. On the informed branch the first order's profit π_1 is the truthful trade's, integrable by the bound recorded after Lemma 4 with $R' = P$, and $\mathbb{E}[\pi_1 \mid P] = b\,\mathrm{KL}(P \,\|\, \delta_0)$ by Lemma 1; at the selected level its expectation is $b\,S(b)$. On the

noise branch $\mathbb{E}[\pi_1] = -b\,\nu_b$, finite, by Lemma 3(i). Hence $\mathbb{E}[\pi_1] = \vartheta\,b\,S(b) - (1-\vartheta)\,b\,\nu_b$. The arbitrageur's profit $\pi_A = (\pi_1 + \pi_A) - \pi_1$ is integrable as a difference of integrable variables, with $\mathbb{E}[\pi_1 + \pi_A] = b\,\mathcal{I}(b)$ by Proposition 2, so

$$\mathbb{E}[\pi_A] \;=\; b\,\mathcal{I}(b) - \vartheta\,b\,S(b) + (1-\vartheta)\,b\,\nu_b \;=\; b\,[(1-\vartheta)\,\nu_b - \mathcal{L}_T(b)] \;=\; b\,D_A(b),$$

the second equality by $\mathcal{I} = \vartheta S - \mathcal{L}_T$ (Proposition 1(iii)) and the third by the rent identity of Lemma 2, in which every term is finite here. Rearranging, $b\,\mathcal{I}(b) = \vartheta\,b\,S(b) + b\,D_A(b) - (1-\vartheta)\,b\,\nu_b$, which is (15).

The objective (16). The market maker earns $\lambda\,\mathcal{I}(b)$ (Proposition 1(ii)), pays the expected net payout $\mathbb{E}[\pi_1 + \pi_A] = b\,\mathcal{I}(b)$, pays the overhead $\eta - 1$ on its expected payments to the informed trader and to the arbitrageur, $\vartheta\,b\,S(b) + b\,D_A(b) \;=\; b\,\mathcal{I}(b) + (1-\vartheta)\,b\,\nu_b$ by the ledger, and carries the capital charge $c_0\,b$. Summing,

$$V_\lambda(b) \;=\; \lambda\,\mathcal{I}(b) - b\,\mathcal{I}(b) - (\eta-1)\,b\,\mathcal{I}(b) - (\eta-1)(1-\vartheta)\,b\,\nu_b - c_0\,b \;=\; (\lambda - \eta\,b)\,\mathcal{I}(b) - A(b),$$

with $A(b) = (\eta-1)(1-\vartheta)\,b\,\nu_b + c_0\,b$. At $\eta = 1$ the overhead term is absent, the objective is $(\lambda - b)\,\mathcal{I}(b) - c_0\,b$ by Proposition 2 alone, and no moment of U is used. By Lemma 3(ii), $b\,\nu_b$ is convex on $(0,\infty)$, hence continuous, with the stated limit as $b \downarrow 0$, so A is continuous on $(0,\infty)$ with $A(b) \to (\eta-1)(1-\vartheta)\,\mathbb{E}[(1-\delta_0)U^+ + \delta_0 U^-]$.

The market with no arbitrageur. Keep the setup and the first order and remove the arbitrageur, so that the market delivers the first-round price $\delta_1(Q)$; let $\vartheta < 1$ and $\mathbb{E}|U| < \infty$. Its expected log-loss is $h(\delta_0) - \mathcal{I}(b) + D_A(b)$ (Proposition 1(iv)), so the accuracy revenue is $\lambda\,[\mathcal{I}(b) - D_A(b)]$; the net payout is $\mathbb{E}[\pi_1] = \vartheta\,b\,S(b) - (1-\vartheta)\,b\,\nu_b$; the overhead falls on the informed trader's payment alone; and the informed trader's acquisition is unchanged, its payoff being fixed by its own trade (Section 3.1). The market maker's expected profit is therefore

$$V^0_\lambda(b) \;=\; \lambda\,[\mathcal{I}(b) - D_A(b)] - \eta\,\vartheta\,b\,S(b) + (1-\vartheta)\,b\,\nu_b - c_0\,b,$$

and subtracting it from (16), with $b\,\mathcal{I}(b)$ replaced through the ledger,

$$V_\lambda(b) - V^0_\lambda(b) \;=\; \lambda\,D_A(b) - \eta\,b\,\mathcal{I}(b) + \eta\,\vartheta\,b\,S(b) - \eta\,(1-\vartheta)\,b\,\nu_b \;=\; (\lambda - \eta\,b)\,D_A(b),$$

the correction's value at a fixed liquidity, nonnegative inside the cap $b < \lambda/\eta$ and of the margin's sign elsewhere.

The comparison stops at a fixed liquidity. Beyond $b = \lambda$ the market maker with no arbitrageur collects the noise trader's losses at par while the quote it delivers scores worse than the prior, and it can profit from doing so. Take $\delta_0 = \dfrac{1}{2}$, $\vartheta = 0$, a sure noise order $U = 8$, $\lambda = \eta = 1$ and $c_0 = \dfrac{1}{4}$. With no informed arrival $\mathcal{I} \equiv 0$, so $V_\lambda(b) = -c_0 b < 0$ for every $b > 0$ and the market maker with an arbitrageur closes, with value 0; while at $b = 2$,

with $\nu_2 = \varphi(4) = \log((1+e^4)/2) - 2 = \log\cosh 2$ by (21),

$$V_\lambda^0(2) \;=\; -\lambda\,\nu_2 + 2\,\nu_2 - 2\,c_0 \;=\; \log\cosh 2 - \frac{1}{2} \;>\; 0,$$

since $\cosh 2 \geq 1 + 2^2/2 = 3 > e^{1/2}$; the delivered accuracy there is $-\nu_2 < 0$. So neither the optimized profit nor the set of valuations at which the market is worth running is ordered across the two markets without a bound on the capital charge.

A.19 Proofs of Theorem 1, Proposition 5, Proposition 6, and Proposition 7

Throughout, the schedule $\mathcal{I} : (0,\infty) \to [0, \vartheta\, h(\delta_0)]$ is the delivered information generated by the technology and the selection (its range by Proposition 1(iii)), extended by $\mathcal{I}(0) := 0$, and $A(0) := 0$, so that $V_\lambda(b) = \lambda\,\mathcal{I}(b) - [\eta\, b\,\mathcal{I}(b) + A(b)]$ on $[0,\infty)$ is affine in λ with slope $\mathcal{I}(b)$ and affine in c_0 with slope $-b$. By Appendix A.18, A is continuous on $(0,\infty)$ with $A(b) \geq c_0 b > 0$ for $b > 0$ and the finite limit $A(0^+) = (\eta-1)(1-\vartheta)\,\mathbb{E}[(1-\delta_0)U^+ + \delta_0 U^-]$, the moment $\mathbb{E}|U| < \infty$ being used only when $\eta > 1$. The proofs use no property of the schedule beyond those named in each part, so the conclusions hold for any selection whose schedule meets the part's hypotheses, as the footnote to Section 5 states.

Proof of Theorem 1. (i) Liquidity cap. For $b \geq \lambda/\eta$, $\lambda - \eta b \leq 0$ and $\mathcal{I}(b) \geq 0$ give $(\lambda - \eta b)\,\mathcal{I}(b) \leq 0$, so $V_\lambda(b) \leq -A(b) < 0 = V_\lambda(0)$. Hence no $b \geq \lambda/\eta$ maximizes V_λ over $[0,\infty)$, and a maximizer $b^* > 0$ satisfies $V_\lambda(b^*) \geq V_\lambda(0) = 0$, that is, $(\lambda - \eta b^*)\,\mathcal{I}(b^*) \geq A(b^*) > 0$.

(ii) Existence. On $(0, \lambda/\eta)$ the factor $\lambda - \eta b$ is positive and continuous and $\mathcal{I}$ is nonnegative and upper semicontinuous, so their product is upper semicontinuous, as in Appendix A.22(ii). At the cap, for $b_n \uparrow \lambda/\eta$, $0 \leq (\lambda - \eta b_n)\,\mathcal{I}(b_n) \leq \vartheta h(\delta_0)\,(\lambda - \eta b_n) \to 0$, the value of the product at λ/η, so the product is upper semicontinuous on $(0, \lambda/\eta]$ and nothing is asked of $\mathcal{I}$ at the cap. Subtracting the continuous A keeps upper semicontinuity on $(0, \lambda/\eta]$. At the closed market, $\limsup_{b\downarrow 0} V_\lambda(b) = \lambda \limsup_{b\downarrow 0} \mathcal{I}(b) - A(0^+) \leq 0 = V_\lambda(0)$ by hypothesis, which is upper semicontinuity at 0. An upper semicontinuous function on the compact interval $[0, \lambda/\eta]$ attains its supremum, by the nested-superlevel-set argument of Appendix A.22(ii); by part (i) the attained maximum equals $\sup_{b\geq 0} V_\lambda$; and $B^*(\lambda)$, the intersection of $[0, \lambda/\eta]$ with the closed superlevel set of V_λ at its maximum, is nonempty and compact.

(iii) Entry. Where $\mathcal{I}(b) > 0$, $V_\lambda(b) = \mathcal{I}(b)\,[\lambda - \eta b - A(b)/\mathcal{I}(b)]$; where $\mathcal{I}(b) = 0$ and $b > 0$, $V_\lambda(b) = -A(b) < 0$. The argument of Appendix A.22(iii) then applies verbatim with $A(b)$ in place of $c_0 b$: if $\lambda > \lambda_{\text{entry}}$ some b with $\mathcal{I}(b) > 0$ has $\eta b + A(b)/\mathcal{I}(b) < \lambda$, so $V_\lambda(b) > 0$; conversely $V_\lambda(b) > 0$ forces $b > 0$, $\mathcal{I}(b) > 0$ and $\lambda > \eta b + A(b)/\mathcal{I}(b) \geq \lambda_{\text{entry}}$; at $\lambda = \lambda_{\text{entry}}$ no b has $V_\lambda(b) > 0$ while $V_\lambda(0) = 0$, so $V^* = 0$, attained or not; and for $\lambda < \lambda_{\text{entry}}$ every $b > 0$ has $V_\lambda(b) < 0 = V_\lambda(0)$, so $B^*(\lambda) = \{0\}$.

(iv) Marginal condition. By Lemma 3(iii), $b \mapsto b\nu_b$ is differentiable on $(0,\infty)$ with derivative $-\mathbb{E}[\mathrm{KL}(\delta_1(U) \,\|\, \delta_0)]$, so A is differentiable with the derivative displayed in (18); at $\eta = 1$, $A' = c_0$. If $\mathcal{I}$ is differentiable at the interior maximizer b^*, then V_λ is differentiable there by the product rule, and Fermat's rule $V'_\lambda(b^*) = (\lambda - \eta b^*)\,\mathcal{I}'(b^*) - \eta\,\mathcal{I}(b^*) - A'(b^*) = 0$ is (18). Q.E.D.

Proof of Proposition 5. (i) Let $\lambda_1 < \lambda_2$ and $b_i \in B^*(\lambda_i)$. Optimality gives $V_{\lambda_1}(b_1) \geq V_{\lambda_1}(b_2)$ and $V_{\lambda_2}(b_2) \geq V_{\lambda_2}(b_1)$; adding the two and cancelling the terms that do not involve the valuation leaves $(\lambda_2 - \lambda_1)\,[\mathcal{I}(b_2) - \mathcal{I}(b_1)] \geq 0$. *(ii)* For the value, V^* is the pointwise supremum over $b \geq 0$ of the affine functions $\lambda \mapsto \lambda\,\mathcal{I}(b) - [\eta b\,\mathcal{I}(b) + A(b)]$, whose slopes lie in $[0, \vartheta h(\delta_0)]$: it is convex and nondecreasing, finite since $V_\lambda(b) \leq \lambda\,\vartheta h(\delta_0)$, and Lipschitz with constant $\vartheta h(\delta_0)$, because for $\lambda < \lambda'$ each affine function rises by at most $(\lambda' - \lambda)\,\vartheta h(\delta_0)$ and so does their supremum. Q.E.D.

Proof of Proposition 6. (i) Under $c_0 \geq (\eta - 1)(1 - \vartheta)M(\delta_0)$ the loss term is nondecreasing: $A'(b) = c_0 - (\eta - 1)(1 - \vartheta)\,\mathbb{E}[\mathrm{KL}(\delta_1(U) \,\|\, \delta_0)] \geq c_0 - (\eta - 1)(1 - \vartheta)M(\delta_0) \geq 0$ by Lemma 3(iii), and at $\eta = 1$ or $\vartheta = 1$ the bound holds automatically. Suppose $\lambda_1 < \lambda_2$, $b_i \in B^*(\lambda_i)$, and $b_2 < b_1$. Then $b_1 > 0$, so $\mathcal{I}(b_1) > 0$ and $\eta b_1 < \lambda_1$ by Theorem 1(i), and $\mathcal{I}(b_2) \geq \mathcal{I}(b_1)$ by Proposition 5(i). Decomposing,

$$V_{\lambda_1}(b_2) - V_{\lambda_1}(b_1) \;=\; (\lambda_1 - \eta b_2)\,[\mathcal{I}(b_2) - \mathcal{I}(b_1)] + \eta\,(b_1 - b_2)\,\mathcal{I}(b_1) + [A(b_1) - A(b_2)] \;>\; 0, \quad (29)$$

the first term nonnegative since $\lambda_1 - \eta b_2 > \lambda_1 - \eta b_1 > 0$, the second positive, the third nonnegative since A is nondecreasing; this contradicts $b_1 \in B^*(\lambda_1)$. *(ii)* Strictness at one valuation: let $b < b'$ both lie in $B^*(\lambda)$, so $\mathcal{I}(b') > 0$ and $\eta b' < \lambda$ by Theorem 1(i); if $\mathcal{I}(b) \geq \mathcal{I}(b')$, then (29) with (b, b') in the roles of (b_2, b_1) gives $V_\lambda(b) > V_\lambda(b')$, contradicting the optimality of b'; so $\mathcal{I}(b) < \mathcal{I}(b')$. Across $\lambda < \lambda'$ with $b \in B^*(\lambda)$, $b' \in B^*(\lambda')$ and $b < b'$: Proposition 5(i) gives $\mathcal{I}(b') \geq \mathcal{I}(b)$, and if the two were equal to a common s, then $s = \mathcal{I}(b') > 0$ and $V_{\lambda'}(b) - V_{\lambda'}(b') = \eta s\,(b' - b) + [A(b') - A(b)] > 0$, contradicting $b' \in B^*(\lambda')$. *(iii)* Generic uniqueness: V^* is convex and finite, hence differentiable off a countable set, and at a point of differentiability $\lambda > \lambda_{\text{entry}}$ every maximizer b^* satisfies $\mathcal{I}(b^*) = V^{*\prime}(\lambda)$, by the subgradient argument of Appendix A.23 with $\mathcal{I}$ in place of S; two maximizers would then deliver the same information, which the strictness just shown excludes unless they coincide. Under the hypotheses of Theorem 1(ii) a maximizer exists, so there is exactly one. Q.E.D.

Proof of Proposition 7. Write V_i for the objective at the capital charge c_i, so that $V_2(b) = V_1(b) - (c_2 - c_1)\,b$ on $[0,\infty)$. Optimality gives $V_1(b_1) \geq V_1(b_2)$ and $V_2(b_2) \geq V_2(b_1)$; adding the two and cancelling the terms common to both sides leaves $-(c_2 - c_1)\,(b_2 - b_1) \geq 0$, so $b_2 \leq b_1$. Q.E.D.

The technologies. In the precision technology under the hypotheses of Proposition 12, $\mathcal{I}$ is continuous on $(0,\infty)$ with $\mathcal{I}(b) \to 0$ as $b \downarrow 0$ (Proposition 13(iv)), so both hypotheses of Theorem 1(ii) hold for every noise law, and $\mathcal{I}$ is differentiable wherever $\mathfrak{b}$ has a positive derivative (Lemma 6(ii)). In the sampling technology against an atomless noise law with $c_N(N) \to \infty$ and the largest optimal size selected, $\mathcal{I} = \vartheta S$ (Proposition 10(iii)), and S is nondecreasing and upper semicontinuous on $(0,\infty)$. Nondecreasing: for $b < b'$ with selected sizes N, N', adding the two optimality inequalities $b\,y(N) - c_N(N) \geq b\,y(N') - c_N(N')$ and $b'y(N') - c_N(N') \geq b'y(N) - c_N(N)$ gives $(b'-b)\,[y(N') - y(N)] \geq 0$, so $y(N') \geq y(N)$. Upper semicontinuous: at $b_0 > 0$, every size optimal at some $b \in [b_0/2, 2b_0]$ has $c_N(N) \leq b\,y(N) \leq 2b_0 y_\infty$, since the value at $N = 0$ is zero, so finitely many sizes are candidates on that interval by $c_N(N) \to \infty$; along a sequence $b_n \to b_0$ the selected sizes take finitely many values, and any value N' taken infinitely often is optimal at b_0, the objective being continuous in b and the candidates finite; so $y(N') \leq y(N^*(b_0)) = S(b_0)$, the largest maximizer carrying the largest level because y is strictly increasing (Proposition 9(ii)), and $\limsup_n S(b_n) \leq S(b_0)$. At the boundary, $S(b) \to 0$ as $b \downarrow 0$ exactly when $c_N(1) > 0$: if $c_N(1) > 0$, then for $b < c_N(1)/y_\infty$ every $N \geq 1$ has $b\,y(N) - c_N(N) \leq b\,y_\infty - c_N(1) < 0$, so $N^*(b) = 0$; if $c_N(1) = 0$, a first sample beats none at every $b > 0$, so $S(b) \geq y(1) > 0$.

The right side of the marginal condition. The sampling technology under separation supplies an optimum at which delivered information is flat and the right side of (18) is zero. Let $\delta_0 = \dfrac{1}{2}$, $\kappa = \dfrac{1}{4}$, so that one sample yields the belief $\dfrac{9}{10}$ or $\dfrac{1}{10}$ and the level $s := y(1) = \mathrm{KL}(\dfrac{9}{10} \,\|\, \dfrac{1}{2}) = \dfrac{9}{10}\log 1.8 - \dfrac{1}{10}\log 5 \in (0.36, 0.38)$, since $\log 1.8 \in (0.58, 0.59)$ and $\log 5 \in (1.60, 1.61)$; let $c_N(1) = s/10$ and $c_N(N) = 100 + N$ for $N \geq 2$; let the noise order be uniform on $\{u : 1 - \epsilon \leq |u| \leq 1 + \epsilon\}$ with $\epsilon = 10^{-3}$, an atomless law with every moment; and let $\eta = 3$, $c_0 = \dfrac{1}{100}$, $\lambda = 10$, $\vartheta \in [\dfrac{2}{5}, \dfrac{1}{2}]$. One sample is bought exactly for $b \geq \dfrac{1}{10}$, and no second sample at any $b \leq 100/\log 2$, so $S = s$ on $[\dfrac{1}{10}, \lambda/\eta)$ and $S = 0$ below $\dfrac{1}{10}$; separation gives $\mathcal{I} = \vartheta S$. Write $k(b) := b\,\nu_b$ and $r(x) := x\tanh x - \log\cosh x$, so that at $\delta_0 = \dfrac{1}{2}$, $\mathrm{KL}(\delta_1(u) \,\|\, \delta_0) = r(|u|/(2b))$ by the computation of Appendix A.2 and $-k'(b) = \mathbb{E}[r(|U|/(2b))]$ by Lemma 3(iii); r is strictly increasing on $(0,\infty)$, since $r'(x) = x\,\mathrm{sech}^2 x$, with $r(x) \leq x^2/2$ and $r(x) \geq \log 2 - (2x+1)e^{-2x}$, the latter from $\tanh x \geq 1 - 2e^{-2x}$ and $\log\cosh x \leq x - \log 2 + e^{-2x}$. On the step $(\dfrac{1}{10}, \lambda/\eta)$,

$$\partial_b V_\lambda(b;\vartheta) \;=\; -\eta\,\vartheta\,s - (\eta-1)(1-\vartheta)\,k'(b) - c_0, \qquad \partial_b^2 V_\lambda(b;\vartheta) \;=\; -(\eta-1)(1-\vartheta)\,k''(b) \;<\; 0,$$

the second derivative existing because $|U|$ is bounded, so that k' differentiates under the expectation, and negative because k is strictly convex (Lemma 3(ii)). At $b = \frac{1}{10}$, $-k' \geq r(4.995) > 0.69$, so $\partial_b V_\lambda \geq -\frac{3}{2}(0.38) + 0.69 - 0.01 > 0$; at $b = 1$, $-k' \leq r(0.5005) < 0.126$, so $\partial_b V_\lambda \leq -\frac{6}{5}(0.36) + \frac{6}{5}(0.126) - 0.01 < 0$. Hence $V_\lambda(\cdot;\vartheta)$ is strictly concave on the step with a unique stationary point $b^*(\vartheta) \in (\frac{1}{10}, 1)$, its maximizer on the step. It is the global maximizer: below $\frac{1}{10}$ the supply is zero and $V_\lambda = -A < 0$; at and beyond the cap $V_\lambda < 0$ by Theorem 1(i); and $V_\lambda(\frac{1}{10};\vartheta) \geq (10 - 0.3)\,\frac{2}{5}\,(0.36) - 2 \cdot \frac{3}{5}\,(0.5005) - 0.001 > 0$, using $k(\frac{1}{10}) \leq \frac{1}{2}\mathbb{E}|U| \leq 0.5005$ from the bound $b\varphi(u/b) \leq \max\{\delta_0, 1-\delta_0\}|u|$ of Appendix A.2. At $b^*(\vartheta)$ delivered information is flat, $\mathcal{I}' = 0$, and (18) reads $0 = \eta\,\mathcal{I}(b^*) + A'(b^*)$ with $A'(b^*) = -\eta\vartheta s < 0$: the falling overhead on noise-funded payments exactly offsets the grossed-up subsidy on the inframarginal information.

A.20 The single-trader market

At $\vartheta = 1$ there is no noise order and no rent, $\mathcal{I} = S$ and $A(b) = c_0\,b$, and the market maker's objective (16) is $V_\lambda(b) = (\lambda - \eta\,b)\,S(b) - c_0\,b$ on the supply S alone. This appendix collects the single-trader statements that Theorem 1 generalizes, in the form in which the proofs of Appendix B and of the closed forms below use them; the supply is written y^* where a proof reads the trader's selected level directly, $S = y^*$ being the one-visit supply of Appendix A.10, the largest maximizer of (9) at each liquidity. Delivered information is delivered accuracy here as in the mixed market: Proposition 1(ii) at $\vartheta = 1$ is the identity between the level and the log-loss reduction, which Appendix B cites in the following form.

Lemma 7 (Log-loss reduction)**.** *For every feasible y, $\mathbb{E}[h(P_y)] = h(\delta_0) - y$.*

Proof. Appendix A.21.

The single-trader liquidity problem is solved for a generic supply schedule.

Proposition 14 (Optimal liquidity in the single-trader market)**.** *Fix $\lambda > 0$, $\eta \geq 1$, and $c_0 > 0$. Let $S : [0,\infty) \to [0,\bar{S}]$ satisfy $S(0) = 0$, write $V_\lambda(b) := (\lambda - \eta\,b)\,S(b) - c_0\,b$, and let*

$$\lambda_{\text{entry}} \;:=\; \inf\Big\{\eta\,b + \frac{c_0\,b}{S(b)} :\; b > 0,\; S(b) > 0\Big\}, \qquad \inf\varnothing := +\infty. \tag{30}$$

(i) (Cap.) *Every $b > \lambda/\eta$ satisfies $V_\lambda(b) < V_\lambda(0) = 0$; consequently $\sup_{b\geq 0} V_\lambda = \sup_{[0,\lambda/\eta]} V_\lambda$, and every maximizer of V_λ over $[0,\infty)$ lies in $[0,\lambda/\eta]$.*

(ii) (Existence.) *If S is upper semicontinuous on $[0, \lambda/\eta]$, then V_λ attains its maximum on $[0, \lambda/\eta]$, and that maximum equals* $\sup_{b\geq 0} V_\lambda$.

(iii) (Entry.) $\sup_{b\geq 0} V_\lambda > 0$ *if and only if $\lambda > \lambda_{\text{entry}}$, and the supremum is zero at $\lambda = \lambda_{\text{entry}}$, whether or not the infimum in* (30) *is attained. If $\lambda < \lambda_{\text{entry}}$, the unique maximizer of V_λ is $b = 0$.*

(iv) (Marginal condition.) *If $b^* \in (0, \lambda/\eta)$ maximizes V_λ and S is differentiable at b^*, then*

$$(\lambda - \eta\, b^*)\, S'(b^*) \;=\; \eta\, S(b^*) + c_0. \tag{31}$$

Proof. Appendix A.22.

The threshold (30) and the condition (31) are (17) and (18) at $\vartheta = 1$. Rearranged, the marginal condition is a Lerner condition with a capital wedge,

$$\frac{\lambda - \eta\, b^*}{\eta\, b^*} \;=\; \frac{1}{\varepsilon(b^*)}\Big(1 + \frac{c_0}{\eta\, S(b^*)}\Big), \qquad \varepsilon(b) := \frac{b\, S'(b)}{S(b)},$$

since the right side of (31) is positive, so that $S'(b^*) > 0$ at a differentiable interior optimum: the proportional markdown of the grossed-up price of information below the value of accuracy, measured against the price the market maker pays, equals the inverse elasticity of the information supply, scaled up by the capital charge relative to the grossed-up supply. At $\eta = 1$ and $\gamma = 0$ this is the textbook monopsony markdown, with information in place of labor; the settlement overhead raises the effective price $\eta\, b$ of every unit bought and adds no wedge of its own, while the capital charge deepens the measured markdown.

Proposition 15 (Comparative statics in the single-trader market)**.** *Under the hypotheses and notation of Proposition 14, the value $V^*(\lambda) := \sup_{b\geq 0} V_\lambda(b)$ is nondecreasing and convex on $(0, \infty)$. If $\lambda_1 < \lambda_2$ and S is upper semicontinuous on the caps as in Proposition 14*(ii)*, then the maximizer sets $B^*(\lambda_i)$ of V_{λ_i} over $[0, \lambda_i/\eta]$ are nonempty and compact with* $\min B^*(\lambda_1) \leq \min B^*(\lambda_2)$ *and* $\max B^*(\lambda_1) \leq \max B^*(\lambda_2)$*; and for any selections $b_i^* \in B^*(\lambda_i)$, $S(b_1^*) \leq S(b_2^*)$, strictly whenever $b_1^* \neq b_2^*$. For all but countably many $\lambda > \lambda_{\text{entry}}$, V_λ has at most one maximizer over $[0, \infty)$ — exactly one under the hypothesis of Proposition 14*(ii)*.*

Proof. Appendix A.23.

Proposition 6(i) at $\vartheta = 1$ strengthens the extremal ordering to every pair of maximizers, with no monotonicity or concavity of the supply.

The one-visit supply qualifies for Propositions 14 and 15 with $\bar{S} = \bar{y}$ (Appendix A.10), and Corollary 4 solves the iso-elastic case in closed form.

The proofs follow, each in its own subsection.

A.21 Proof of Lemma 7

Fix a feasible y with implementing posterior P_y as in Assumption 3. Expanding the Bernoulli relative entropy in (8),

$$\mathrm{KL}(P_y \,\|\, \delta_0) = -h(P_y) - P_y \log \delta_0 - (1 - P_y) \log(1 - \delta_0),$$

where $h(x) = -x \log x - (1 - x) \log(1 - x)$ is the binary entropy. Taking expectations and using Bayes plausibility $\mathbb{E}[P_y] = \delta_0$ (Assumption 3),

$$y = -\mathbb{E}[h(P_y)] - \delta_0 \log \delta_0 - (1 - \delta_0) \log(1 - \delta_0) = h(\delta_0) - \mathbb{E}[h(P_y)],$$

which is the stated identity $\mathbb{E}[h(P_y)] = h(\delta_0) - y$. Q.E.D.

A.22 Proof of Proposition 14

Throughout, the standing objects are those of the statement: $S : [0, \infty) \to [0, \bar{S}]$ with $S(0) = 0$, and $V_\lambda(b) = (\lambda - \eta\, b)\, S(b) - c_0\, b$ with $\lambda > 0$, $\eta \geq 1$, $c_0 > 0$; we write $V := V_\lambda$ when λ is fixed, and $B^*(\lambda)$ for the set of maximizers of V_λ over $[0, \lambda/\eta]$.

(i) Cap. $V(0) = \lambda\, S(0) = 0$ by $S(0) = 0$. Fix $b > \lambda/\eta$. Then $\lambda - \eta b < 0$ and $S(b) \geq 0$ give $(\lambda - \eta b) S(b) \leq 0$; adding $-c_0 b < 0$ yields $V(b) < 0 = V(0)$. Hence $\sup_{[0,\lambda/\eta]} V \geq V(0) = 0 > \sup_{b > \lambda/\eta} V(b)$, which is the supremum identity; and no $b > \lambda/\eta$ attains $\sup_{b \geq 0} V \geq 0$, so every maximizer over $[0, \infty)$ lies in $[0, \lambda/\eta]$.

(ii) Existence. We first show $b \mapsto (\lambda - \eta b) S(b)$ is upper semicontinuous on $[0, \lambda/\eta]$. Fix b in the interval and a sequence $b_n \to b$ in it. If $\lambda - \eta b > 0$, then

$$\limsup_n\, (\lambda - \eta b_n) S(b_n) = (\lambda - \eta b)\, \limsup_n\, S(b_n) \leq (\lambda - \eta b)\, S(b),$$

using $\lambda - \eta b_n \to \lambda - \eta b > 0$, $S \geq 0$, and upper semicontinuity of S. If $\lambda - \eta b = 0$, then $0 \leq \limsup_n (\lambda - \eta b_n) S(b_n) \leq \bar{S}\, \lim_n (\lambda - \eta b_n) = 0 = (\lambda - \eta b) S(b)$, using $0 \leq S \leq \bar{S}$. Subtracting the continuous function $c_0 b$ preserves upper semicontinuity, so V is upper semicontinuous on the compact interval $[0, \lambda/\eta]$. An upper semicontinuous function attains its supremum on a nonempty compact set: the superlevel sets $\{V \geq \alpha\}$, $\alpha < \sup_{[0,\lambda/\eta]} V$, are nonempty, closed by upper semicontinuity, and nested, so by the finite-intersection property of the compact interval their intersection $\{V \geq \sup_{[0,\lambda/\eta]} V\}$ is nonempty, and any of its points is a maximizer. By (i) the attained maximum, being $\geq V(0) = 0 \geq V(b)$ for all $b > \lambda/\eta$, equals $\sup_{b \geq 0} V$.

(iii) Entry. Suppose $\lambda > \lambda_{\text{entry}}$. By the definition (30) of the infimum there is $b > 0$ with

$S(b) > 0$ and $\eta b + c_0 b/S(b) < \lambda$; then

$$V(b) = S(b)\Big[\lambda - \eta\, b - \frac{c_0\, b}{S(b)}\Big] > 0, \tag{32}$$

so $\sup V > 0$. Conversely, suppose $V(b) > 0$ for some b. Then $b > 0$ because $V(0) = 0$, and $S(b) > 0$, since $S(b) = 0$ would give $V(b) = -c_0 b \le 0$; so (32) applies and gives $\lambda > \eta b + c_0 b/S(b) \ge \lambda_{\text{entry}}$. Hence $\sup V > 0$ implies $\lambda > \lambda_{\text{entry}}$. At $\lambda = \lambda_{\text{entry}}$, no b has $V(b) > 0$ — that would give $\lambda > \lambda_{\text{entry}}$ — while $V(0) = 0$, so $\sup V = 0$, whether or not the infimum in (30) is attained. Finally, let $\lambda < \lambda_{\text{entry}}$, and fix $b > 0$: if $S(b) = 0$ then $V(b) = -c_0 b < 0$; if $S(b) > 0$ then $\eta b + c_0 b/S(b) \ge \lambda_{\text{entry}} > \lambda$ and (32) give $V(b) < 0$. So $V(b) < 0 = V(0)$ for every $b > 0$, and $b = 0$ is the unique maximizer.

(iv) Marginal condition. At b^* the function V is differentiable, with $V'(b^*) = (\lambda - \eta b^*)S'(b^*) - \eta S(b^*) - c_0$ by the product rule. An interior maximizer of V over $[0, \infty)$ is a local maximum, so Fermat's rule gives $V'(b^*) = 0$, which is (31).

A.23 Proof of Proposition 15

The value. For fixed b, the map $\lambda \mapsto V_\lambda(b) = S(b)\,\lambda - [\eta\, b\, S(b) + c_0\, b]$ is affine with slope $S(b) \in [0, \bar{S}]$, so $V^* = \sup_{b \ge 0} V_\cdot(b)$ is a pointwise supremum of nondecreasing affine functions of λ: convex, nondecreasing, and finite ($V_\lambda(b) \le \lambda \bar{S}$). The sign claim is Proposition 14(iii).

Extremal maximizers. Let $\lambda_1 < \lambda_2$ and let S be upper semicontinuous on the caps. By Proposition 14(ii) each $B^*(\lambda_i)$ is nonempty, and it is compact, being the intersection of the compact cap with the closed superlevel set of V_{λ_i} at its attained maximum; so its minimum and maximum exist. For the maxima, suppose toward a contradiction $b_2 := \max B^*(\lambda_2) < b_1 := \max B^*(\lambda_1)$. Since $b_1 \le \lambda_1/\eta < \lambda_2/\eta$, b_1 is feasible at λ_2, and $b_2 < b_1 \le \lambda_1/\eta$ is feasible at λ_1, so optimality gives $V_{\lambda_2}(b_2) \ge V_{\lambda_2}(b_1)$ and $V_{\lambda_1}(b_1) \ge V_{\lambda_1}(b_2)$. Adding the two and cancelling the common terms leaves $(\lambda_2 - \lambda_1)[S(b_2) - S(b_1)] \ge 0$, so $S(b_2) \ge S(b_1)$. If $S(b_2) > S(b_1)$: since $\lambda_1 - \eta b_1 \ge 0$ and $b_2 < b_1$,

$$V_{\lambda_1}(b_1) \;\le\; (\lambda_1 - \eta b_1)\, S(b_2) - c_0\, b_1 \;\le\; (\lambda_1 - \eta b_2)\, S(b_2) - c_0\, b_2 \;=\; V_{\lambda_1}(b_2),$$

the first inequality by $S(b_1) \le S(b_2)$ and $\lambda_1 - \eta b_1 \ge 0$, the second by $\eta\,(b_1 - b_2)\, S(b_2) \ge 0$ and $c_0(b_1 - b_2) \ge 0$. Optimality of b_1 at λ_1 forces equality throughout; equality in the second inequality requires $\eta\,(b_1 - b_2)\, S(b_2) = 0$, hence $S(b_2) = 0$, contradicting $S(b_2) > S(b_1) \ge 0$. If $S(b_2) = S(b_1)$: then

$$V_{\lambda_2}(b_1) - V_{\lambda_2}(b_2) = [V_{\lambda_1}(b_1) - V_{\lambda_1}(b_2)] + (\lambda_2 - \lambda_1)[S(b_1) - S(b_2)] = V_{\lambda_1}(b_1) - V_{\lambda_1}(b_2) \ge 0,$$

so $b_1 \in B^*(\lambda_2)$, contradicting $b_1 > b_2 = \max B^*(\lambda_2)$. Both cases being impossible, $\max B^*(\lambda_1) \le \max B^*(\lambda_2)$. For the minima, suppose $b_2 := \min B^*(\lambda_2) < b_1 := \min B^*(\lambda_1)$. The same two

optimality inequalities give $S(b_2) \geq S(b_1)$; the case $S(b_2) > S(b_1)$ is excluded by the displayed chain verbatim, which uses only $b_2 < b_1$, feasibility, and optimality of b_1 at λ_1; and if $S(b_2) = S(b_1)$, then $V_{\lambda_1}(b_2) - V_{\lambda_1}(b_1) = [V_{\lambda_2}(b_2) - V_{\lambda_2}(b_1)] - (\lambda_2 - \lambda_1)[S(b_2) - S(b_1)] \geq 0$, so $b_2 \in B^*(\lambda_1)$, contradicting $b_2 < b_1 = \min B^*(\lambda_1)$. This is the monotone-selection conclusion in the language of Topkis (1998); Milgrom and Shannon (1994), obtained with no monotonicity of S.

Delivered supply. Let $b_i^* \in B^*(\lambda_i)$ be arbitrary. If $b_2^* \leq \lambda_1/\eta$, both points are feasible at both weights, and adding $V_{\lambda_2}(b_2^*) \geq V_{\lambda_2}(b_1^*)$ to $V_{\lambda_1}(b_1^*) \geq V_{\lambda_1}(b_2^*)$ gives $(\lambda_2 - \lambda_1)[S(b_2^*) - S(b_1^*)] \geq 0$, hence $S(b_2^*) \geq S(b_1^*)$. If $b_2^* > \lambda_1/\eta \geq b_1^*$, suppose toward a contradiction $S(b_2^*) < S(b_1^*)$. Since $\lambda_2 - \eta b_2^* \geq 0$ and $b_2^* > b_1^*$,

$$V_{\lambda_2}(b_2^*) \;\leq\; (\lambda_2 - \eta b_2^*)\, S(b_1^*) - c_0\, b_2^* \;\leq\; (\lambda_2 - \eta b_1^*)\, S(b_1^*) - c_0\, b_1^* \;=\; V_{\lambda_2}(b_1^*),$$

and optimality of b_2^* at λ_2 forces equality; equality in the second inequality requires $\eta\,(b_2^* - b_1^*)\, S(b_1^*) = 0$, hence $S(b_1^*) = 0$, contradicting $S(b_1^*) > S(b_2^*) \geq 0$.

Strictness. Suppose $b_1^* \neq b_2^*$ and $S(b_1^*) = S(b_2^*) =: s$. If $b_2^* > b_1^*$, then b_1^* is feasible at λ_2, and $V_{\lambda_2}(b_2^*) \geq V_{\lambda_2}(b_1^*)$ reads $-(\eta s + c_0)\, b_2^* \geq -(\eta s + c_0)\, b_1^*$, so $(\eta s + c_0)(b_2^* - b_1^*) \leq 0$ — impossible, since $\eta s + c_0 \geq c_0 > 0$. If $b_2^* < b_1^*$, then $b_2^* < b_1^* \leq \lambda_1/\eta$ is feasible at λ_1, and $V_{\lambda_1}(b_1^*) \geq V_{\lambda_1}(b_2^*)$ gives $(\eta s + c_0)(b_1^* - b_2^*) \leq 0$, impossible the same way. Hence $S(b_1^*) < S(b_2^*)$ whenever $b_1^* \neq b_2^*$.

Generic uniqueness. V^* is convex and finite on $(0, \infty)$ by the value step above, so at every λ its one-sided derivatives exist, are finite, and satisfy $D^- V^*(\lambda) \leq D^+ V^*(\lambda)$, with $D^+ V^*(\lambda_1) \leq D^- V^*(\lambda_2)$ for $\lambda_1 < \lambda_2$ (monotonicity of difference quotients of a convex function). The open intervals $(D^- V^*(\lambda), D^+ V^*(\lambda))$ at distinct nondifferentiability points are therefore nonempty and pairwise disjoint; each contains a distinct rational, so there are at most countably many such points. Fix $\lambda > \lambda_{\text{entry}}$ at which V^* is differentiable, and let b^* be any maximizer of V_λ over $[0, \infty)$, if one exists. For every $\lambda' > 0$,

$$V^*(\lambda') \;\geq\; V_{\lambda'}(b^*) \;=\; V_\lambda(b^*) + (\lambda' - \lambda)\, S(b^*) \;=\; V^*(\lambda) + (\lambda' - \lambda)\, S(b^*),$$

so $S(b^*)$ is a subgradient of V^* at λ, and differentiability gives $S(b^*) = V^{*\prime}(\lambda)$ — one value shared by every maximizer. Any maximizer has $V_\lambda(b^*) = V^*(\lambda) > 0$ by Proposition 14(iii), hence $S(b^*) > 0$ as in Proposition 14(iii); so if b', b'' are two maximizers, $s := S(b') = S(b'') > 0$, and $V_\lambda(b') = V_\lambda(b'')$ reads $-(\eta s + c_0)\, b' = -(\eta s + c_0)\, b''$ with $\eta s + c_0 > 0$, giving $b' = b''$. So V_λ has at most one maximizer; under the hypothesis of Proposition 14(ii) one exists, hence exactly one.

A.24 Statement and proof of Corollary 4

An iso-elastic acquisition cost closes the characterization of Proposition 14 in explicit form.

Corollary 4 (Closed forms under an iso-elastic acquisition cost)**.** *Let $c(y) = k\,y^{1+\chi}/(1+\chi)$ on $[0,\bar{y}]$ with $k>0$ and $\chi>0$, and let $k\,\bar{y}^{\chi} \geq \lambda/\eta$. Then for every $b \in [0,\lambda/\eta]$ the acquired level $y^*(b) = (b/k)^{1/\chi}$ is the unique maximizer of g_b, and the trader's ex-ante surplus is $g_b(y^*(b)) = \dfrac{\chi}{1+\chi}\,b\,y^*(b)$.*

(i) (Quadratic case.) *For $\chi = 1$, $\lambda_{\text{entry}} = c_0\,k$ and the optimal liquidity is unique at every $\lambda > 0$:*

$$b^* = \max\Big\{0,\ \frac{\lambda - c_0\,k}{2\,\eta}\Big\}. \tag{33}$$

(ii) (General exponent.) *For any $\chi > 0$, every optimum with strictly positive value lies in $(0,\lambda/\eta)$ and solves*

$$(b/k)^{1/\chi}\,[\lambda - \eta\,(1+\chi)\,b] \;=\; \chi\,c_0\,b. \tag{34}$$

Proof. Throughout, $V(b) := (\lambda - \eta\,b)\,y^*(b) - c_0\,b$ is the market maker's objective (16) at $\vartheta = 1$ and $V^* := \sup_{b\geq 0} V(b)$.

The trader's response and surplus. On $[0,\bar{y}]$, $g_b(y) = b\,y - k\,y^{1+\chi}/(1+\chi)$ has $g_b'(y) = b - k\,y^{\chi}$, strictly decreasing in y since $\chi > 0$, so g_b is strictly concave and its maximizer on $[0,\bar{y}]$ is unique, equal to $\min\{(b/k)^{1/\chi},\bar{y}\}$; it is $y^*(b) = (b/k)^{1/\chi}$ whenever $(b/k)^{1/\chi} \leq \bar{y}$, which holds for every $b \in [0,\lambda/\eta]$ by the cap condition $k\,\bar{y}^{\chi} \geq \lambda/\eta$. Uniqueness makes every selection — in particular the maximal one — coincide with this formula, so the one-visit schedule is $S(b) = \min\{(b/k)^{1/\chi},\bar{y}\}$ — continuous on $[0,\infty)$, with $S(0) = 0$ — and Proposition 14 applies with $\bar{S} = \bar{y}$. For the surplus, at any $b \in (0,\lambda/\eta]$ the trader's first-order condition $k\,y^*(b)^{\chi} = b$ gives

$$g_b(y^*(b)) = b\,y^*(b) - \frac{k\,y^*(b)^{1+\chi}}{1+\chi} = b\,y^*(b) - \frac{b\,y^*(b)}{1+\chi} = \frac{\chi}{1+\chi}\,b\,y^*(b),$$

and at $b = 0$ both sides vanish because $y^*(0) = 0$.

(i) Quadratic case. At $\chi = 1$ the schedule is $S(b) = \min\{b/k,\bar{y}\}$, strictly positive exactly for $b > 0$, so the objective of (30) equals $\eta\,b + c_0\,k$ for $0 < b \leq k\,\bar{y}$ and $\eta\,b + c_0\,b/\bar{y} > \eta\,k\,\bar{y} + c_0\,k$ for $b > k\,\bar{y}$; its infimum over $b > 0$ is therefore $c_0\,k$, approached as $b \downarrow 0$ and not attained, so $\lambda_{\text{entry}} = c_0\,k$. For the optimum, at $\chi = 1$ and $b \in [0,\lambda/\eta]$,

$$V(b) = (\lambda - \eta b)\,\frac{b}{k} - c_0\,b = \frac{(\lambda - c_0 k)\,b - \eta\,b^2}{k}.$$

If $\lambda \leq c_0 k$, both terms are nonpositive and the second is strictly negative for $b > 0$, so $V(b) < 0 = V(0)$ on $(0,\lambda/\eta]$; and for $b > \lambda/\eta$, $V(b) = (\lambda - \eta b)\,y^*(b) - c_0 b < 0$ because $\lambda - \eta b < 0$ and $y^*(b) > 0$ for $b > 0$ [$g_b(y) > 0$ for small $y > 0$ when $b > 0$, so the maximizer is not 0]. Hence $b^* = 0$ is the unique optimum, which is (33) with the maximum attained

at 0, and the value is 0. If $\lambda > c_0 k$, then on $[0, \lambda/\eta]$, V is a strictly concave quadratic with vertex $b^\dagger = (\lambda - c_0 k)/(2\eta)$, which lies in $(0, \lambda/\eta)$ because $0 < \lambda - c_0 k < 2\lambda$; so $b^\dagger$ is the unique maximizer on the cap, with value $V(b^\dagger) = (\lambda - c_0 k)^2/(4\eta k) > 0$; every b beyond the cap has $V(b) < 0$ as above, so $b^* = b^\dagger$ is the unique optimum over $b \geq 0$, which is (33).

(ii) General exponent. Let b^* be an optimum with $V(b^*) > 0$. By Proposition 14(i), $b^* \leq \lambda/\eta$; and $b^* \notin \{0, \lambda/\eta\}$ because $V(0) = 0$ and $V(\lambda/\eta) = -c_0 \lambda/\eta \leq 0$. On $(0, \lambda/\eta)$,

$$V'(b) = -\eta\,(b/k)^{1/\chi} + (\lambda - \eta b)\,\frac{(b/k)^{1/\chi}}{\chi\, b} - c_0 = \frac{(b/k)^{1/\chi}}{\chi\, b}\,[\lambda - \eta\,(1+\chi)\,b] - c_0,$$

the first equality by the product rule with $\frac{d}{db}(b/k)^{1/\chi} = (b/k)^{1/\chi}/(\chi b)$, so Fermat's rule $V'(b^*) = 0$, multiplied through by $\chi\, b^* > 0$, is (34). Q.E.D.

A.25 Proof of Corollary 1

Write $V(b; \vartheta)$ for the objective at the arrival probability ϑ, with the experiment selected at b and the noise law fixed as ϑ varies, and $k(b) := b\,\nu_b$, finite when $\eta > 1$ by the moment hypothesis and absent from the objective at $\eta = 1$. For $b \in (0, \lambda/\eta)$,

$$V(b; \vartheta) \;=\; (\lambda - \eta b)\,\mathcal{J}(b; \vartheta) + \vartheta\,(\eta - 1)\,k(b) - (\eta - 1)\,k(b) - c_0\,b.$$

By Proposition 3(i), applied at the experiment and noise law of b, $\vartheta \mapsto \mathcal{J}(b; \vartheta)$ is nondecreasing, convex and continuous on $[0, 1]$, and $\lambda - \eta b > 0$; the remaining terms are affine and nondecreasing in ϑ. So $V(b; \cdot)$ is nondecreasing and convex on $[0, 1]$ for each b in the open cap. For $b \geq \lambda/\eta$, $V(b; \vartheta) \leq -A(b) \leq -c_0 b < 0$ for every ϑ (Theorem 1(i)), while closure gives 0; hence $V^*(\vartheta) = \sup_{b \in [0, \lambda/\eta)} V(b; \vartheta)$, the supremum over a fixed index set of nondecreasing convex functions, which is nondecreasing and convex.

(i) $V^*(0) = 0$, since $V(b; 0) = -A(b) < 0$ for $b > 0$. The bound $V(b; \vartheta) \leq (\lambda - \eta b)\,\mathcal{J}(b; \vartheta) \leq \lambda\,\vartheta\,h(\delta_0)$ (Proposition 1(iii)) gives $0 \leq V^*(\vartheta) \leq \lambda \vartheta h(\delta_0)$ and continuity at 0. Convexity gives continuity on $(0, 1)$. At 1: given $\varepsilon > 0$ pick b_ε with $V(b_\varepsilon; 1) > V^*(1) - \varepsilon$; then $V(b_\varepsilon; \vartheta) \to V(b_\varepsilon; 1)$ as $\vartheta \uparrow 1$ by the continuity of $V(b_\varepsilon; \cdot)$, so $\liminf_{\vartheta \uparrow 1} V^*(\vartheta) \geq V^*(1) - \varepsilon$, and monotonicity gives $V^*(\vartheta) \leq V^*(1)$. Finally, a convex function with $V^*(0) = 0$ has $V^*(\vartheta)/\vartheta$ nondecreasing on $(0, 1]$: for $0 < s < t$, $V^*(s) \leq (s/t)\,V^*(t)$.

(ii) The set $\{\vartheta : V^*(\vartheta) > 0\}$ is an up-set by monotonicity and relatively open in $[0, 1]$ by continuity, so it is empty or of the form $(\vartheta_{\text{entry}}, 1]$; when nonempty it contains 1, so $\vartheta_{\text{entry}} < 1$, and $V^*(\vartheta_{\text{entry}}) = 0$ by continuity, since $V^* \leq 0$, hence $= 0$, on $[0, \vartheta_{\text{entry}}]$. For $\vartheta_{\text{entry}} < \vartheta_1 < \vartheta_2$, convexity between 0 and ϑ_2 gives $V^*(\vartheta_2) \geq (\vartheta_2/\vartheta_1)\,V^*(\vartheta_1) > V^*(\vartheta_1)$, since $V^*(\vartheta_1) > 0$. Q.E.D.

A.26 Proof of Proposition 8

The disclosed market. With the sender disclosed after the first order, the arbitrageur conditions on (Q, J): its probability of the event is $\Pr(B = 1 \mid Q, J) = R$, which is $P = \delta_1(Q_I)$ on the informed branch, by the calibration of Assumption 1, and δ_0 on the noise branch, by Assumption 2. By Lemma 1 read from the first-round price (footnote 13) it moves the price to R: it leaves an informed order alone and moves the price back from $\delta_1(U)$ to δ_0 on a noise order, with the order $-U$, so the noise trader's and its own settlement profits cancel state by state as in Appendix A.13. The final price is R, whose divergence from the prior has expectation $\vartheta\,\mathbb{E}[\mathrm{KL}(P \,\|\, \delta_0)] + (1-\vartheta)\cdot 0 = \vartheta S(b)$ at the selected level, and whose type-confusion loss, the divergence of R from itself, is zero. The informed trader's acquisition is unchanged, its payoff being fixed by its own trade (Section 3.1), so both markets share S. The payments: the informed trader's expected profit is $\vartheta b S(b)$, the noise trader's $-(1-\vartheta)b\nu_b$ (Lemma 3(i)), the arbitrageur's $(1-\vartheta)b\nu_b$ by the cancellation, and the net payout is $\vartheta b S(b)$; the overhead falls on $\vartheta b S(b) + (1-\vartheta)b\nu_b$. As in Appendix A.18,

$$V_\lambda^D(b) \;=\; \lambda\,\vartheta S(b) - \eta\,\vartheta b S(b) - (\eta-1)(1-\vartheta)\,b\nu_b - c_0 b \;=\; (\lambda - \eta b)\,\vartheta S(b) - A(b),$$

with $V_\lambda^D(0) = 0$, the objective (16) with ϑS in the place of $\mathscr{I}$; at $\eta = 1$ no moment is used.

(i) Subtracting, $V_\lambda^D(b) - V_\lambda(b) = (\lambda - \eta b)\,[\vartheta S(b) - \mathscr{I}(b)] = (\lambda - \eta b)\,\mathscr{L}_T(b)$ by Proposition 1(iii). For $b < \lambda/\eta$ the factor is positive and $\mathscr{L}_T \geq 0$, so the difference is nonnegative, and positive exactly when $\mathscr{L}_T(b) > 0$.

(ii) For $b \geq \lambda/\eta$ both objectives are at most $-A(b) < 0 = V_\lambda^D(0) = V_\lambda(0)$ (Theorem 1(i) applies to both), so each supremum is the supremum over $[0, \lambda/\eta)$, where $V_\lambda^D \geq V_\lambda$ by part (i); hence $\sup V_\lambda^D \geq \sup V_\lambda$. If V_λ attains its supremum at $b^* > 0$ with $\mathscr{L}_T(b^*) > 0$, then $V_\lambda(b^*) \geq 0$ forces $b^* < \lambda/\eta$, and $\sup V_\lambda^D \geq V_\lambda^D(b^*) > V_\lambda(b^*) = \sup V_\lambda$ by part (i). For the thresholds, write $r^D(b) := \eta b + A(b)/(\vartheta S(b))$ and $r(b) := \eta b + A(b)/\mathscr{I}(b)$ on their index sets. Since $\vartheta S = \mathscr{I} + \mathscr{L}_T \geq \mathscr{I}$, every b in the index set of r lies in that of r^D with $r^D(b) \leq r(b)$; taking infima, $\lambda_{\text{entry}}^D \leq \lambda_{\text{entry}}$. If the infimum defining λ_{entry} is attained at b_0 with $\mathscr{L}_T(b_0) > 0$, then

$$r(b_0) - r^D(b_0) \;=\; A(b_0)\Big[\frac{1}{\mathscr{I}(b_0)} - \frac{1}{\vartheta S(b_0)}\Big] \;=\; \frac{A(b_0)\,\mathscr{L}_T(b_0)}{\mathscr{I}(b_0)\,\vartheta S(b_0)} \;>\; 0,$$

so $\lambda_{\text{entry}}^D \leq r^D(b_0) < r(b_0) = \lambda_{\text{entry}}$.

Q.E.D.

B. A procurement foundation for the accuracy weight

We derive the accuracy weight λ from a contract under which a *principal* purchases the forecast from the market maker. We first consider the single-trader benchmark, $\vartheta = 1$, where delivered information equals $y^*(b)$. Lemma 8 extends the contract and welfare analysis to

noise trading.

The market maker earns revenue from the contract and has no separate preference for accuracy. The acquisition technology, LMSR rule, and trader response remain as specified in the body. Paying the trader costs $\eta \geq 1$ per dollar, and the capital charge is $c_0 = \gamma M(\delta_0) > 0$. The principal's payment to the market maker incurs no settlement overhead. We treat the settlement overhead and capital charge as resource costs in welfare.

The principal values each unit of log-loss reduction at $\omega > 0$. By Lemma 7, information level y therefore has value ωy. A public planner finances payments with an excess burden $\mu \geq 0$: each dollar paid costs society $1 + \mu$ dollars (Laffont and Tirole, 1993). The principal observes the opening price δ_0, the final price δ, and the outcome B, but neither the market maker's rule nor the trader's information acquisition.

Before the market maker chooses liquidity, the principal commits to the contract

$$F \;+\; \lambda\Big[B \log \frac{\delta}{\delta_0} + (1 - B) \log \frac{1 - \delta}{1 - \delta_0}\Big], \qquad F \geq 0, \quad \lambda \geq 0. \tag{35}$$

The payment rewards the improvement in the final price's logarithmic score over the opening price (Savage, 1971; Gneiting and Raftery, 2007). It has the same form as the trader's payoff in Lemma 1, with λ replacing b. The restriction $F \geq 0$ prevents the principal from extracting the market maker's profit through an upfront fee (Innes, 1990).[45] All parties are risk neutral, and the market maker's outside option is zero. Assumption 3 ensures that the realized contract payment is finite.

The principal posts the contract, the market maker chooses liquidity, the trader acquires information and trades, and the event settles. The final price is P_y, so at $F = 0$ the expected payment is

$$\mathbb{E}\Big[\lambda\Big(B \log \frac{P_y}{\delta_0} + (1 - B) \log \frac{1 - P_y}{1 - \delta_0}\Big)\Big] \;=\; \lambda\, \mathbb{E}[\mathrm{KL}(P_y \,\|\, \delta_0)] \;=\; \lambda\, y : \tag{36}$$

the payment rate times delivered information. The market maker therefore earns $(\lambda - \eta b) y^*(b) - c_0 b$, the objective (16) at $\vartheta = 1$. We call a maximizing liquidity a *maker best response* to λ. The market maker receives λ per unit of information and pays ηb, including settlement overhead, while bearing the capital charge. Without payment for accuracy, it optimally closes the market.

The public planner includes its own payoff, the market maker's profit, and the trader's surplus in welfare. The contract payment and trading subsidy cancel as transfers. Acquisition, settlement overhead, capital, and the excess burden of public funds remain as costs. We set $F = 0$: a positive fixed payment does not change the liquidity choice and adds the

[45] This restriction does not impose limited liability: the scoring payment can be negative when the final price scores worse than the prior on the realized outcome.

fiscal cost μF.[46] The planner chooses λ to maximize

$$\max_{\lambda \geq 0} \mathscr{W}(\lambda) \;:=\; \omega\, y - c(y) - \mu\, \lambda\, y - (\eta - 1)\, b(\lambda)\, y - c_0\, b(\lambda), \qquad y = y^*(b(\lambda)), \tag{37}$$

where $b(\lambda)$ is the market maker's response. The term $\omega y - c(y)$ is the value of information net of acquisition cost; the remaining terms are the excess burden of financing the contract, settlement overhead, and the capital charge.[47] A private buyer counts only its own surplus and instead maximizes $(\omega - \lambda)y$, as in Corollary 6.

For comparison, an *integrated planner* operates the LMSR market and chooses liquidity directly, anticipating the same trader response and bearing the same settlement and capital costs. The next proposition characterizes the contract that can induce its preferred liquidity when public funds are costless.

Proposition 16 (The planner's stake with costless funds)**.** *Let $\mu = 0$, and let y^* be a selection of trader maximizers as in Section 5. Define the integrated welfare*

$$\mathscr{W}_I(b) \;:=\; \omega\, y^*(b) - c(y^*(b)) - (\eta - 1)\, b\, y^*(b) - c_0\, b, \qquad b \geq 0, \tag{38}$$

and let b_I maximize $\mathscr{W}_I$ over $b \geq 0$, with $y^*(b_I) > 0$, y^* differentiable at b_I, and $y^{*\prime}(b_I) \in (0, \infty)$. Define

$$\lambda^\dagger \;:=\; \eta\, b_I + \frac{\eta\, y^*(b_I) + c_0}{y^{*\prime}(b_I)}. \tag{39}$$

(i) *If b_I is the unique maker best response to $\lambda^\dagger$, then posting $\lambda^\dagger$ attains the integrated optimum: planner welfare equals $\mathscr{W}_I(b_I)$, and no stake yields strictly higher planner welfare under any maker best response.*

(ii) *Conversely, if b_I is a maker best response to some stake λ with $b_I \in (0, \lambda/\eta)$, then $\lambda = \lambda^\dagger$.*

(iii) *If, in addition, $b_I > 0$, c is differentiable at $y^*(b_I)$, and $y^*(b_I) < \bar{y}$, then*

$$\lambda^\dagger \;=\; \omega + \frac{y^*(b_I)}{y^{*\prime}(b_I)}. \tag{40}$$

Proof. Throughout, $\mu = 0$ and $F = 0$. At a stake λ met by a maker best response b, the planner's welfare (37) is

$$\mathscr{W} = \omega\, y^*(b) - c(y^*(b)) - (\eta - 1)\, b\, y^*(b) - c_0\, b = \mathscr{W}_I(b),$$

[46] At $\mu = 0$, the fixed payment is a transfer with no effect on welfare.

[47] The principal receives $\omega y - (1+\mu)\lambda y$, the market maker receives $\lambda y - \eta b y - c_0 b$, and the trader receives $by - c(y)$. Their sum gives (37).

the integrated welfare (38) evaluated at the market maker's response.

(i) Upper bound and attainment. Every stake is met by some maker best response $b \geq 0$, and the display gives $\mathscr{W} = \mathscr{W}_I(b) \leq \mathscr{W}_I(b_I)$, since b_I maximizes $\mathscr{W}_I$ over $b \geq 0$. By hypothesis b_I is the unique maker best response to $\lambda^\dagger$, so posting $\lambda^\dagger$ delivers $b = b_I$ and $\mathscr{W} = \mathscr{W}_I(b_I)$ — the upper bound. Hence planner welfare equals $\mathscr{W}_I(b_I)$, and no stake yields strictly higher planner welfare under any maker best response, which is part (i).

(ii) Converse. Let b_I be a maker best response to a stake λ with $b_I \in (0, \lambda/\eta)$, so that b_I maximizes $\varphi_\lambda(b) := (\lambda - \eta\, b)\, y^*(b) - c_0\, b$ over $b \geq 0$. By hypothesis y^* is differentiable at b_I, so φ_λ is differentiable at the interior maximizer b_I — the sum of the product of the affine factor $\lambda - \eta\, b$ with y^* and the linear term $-c_0\, b$ — and Fermat's rule $\varphi'_\lambda(b_I) = 0$ reads

$$(\lambda - \eta\, b_I)\, y^{*\prime}(b_I) \;=\; \eta\, y^*(b_I) + c_0.$$

Since $y^{*\prime}(b_I) \in (0, \infty)$, solving for λ gives $\lambda = \eta\, b_I + [\eta\, y^*(b_I) + c_0]/y^{*\prime}(b_I) = \lambda^\dagger$, the definition (39). The same display shows $\lambda^\dagger > \eta\, b_I$: the stake sits strictly above the grossed-up liquidity it induces, consistent with interiority.

(iii) The valuation form. Suppose in addition $b_I > 0$, c is differentiable at $y^*(b_I)$, and $y^*(b_I) < \bar{y}$. Two first-order conditions hold. First, the trader's: $y^*(b_I)$ maximizes $b_I\, y - c(y)$ over $[0, \bar{y}]$ by the definition of the selection, lies in the open interval $(0, \bar{y})$ by hypothesis, and c is differentiable there, so Fermat's rule gives $c'(y^*(b_I)) = b_I$. Second, the integrated planner's: $b_I > 0$ maximizes $\mathscr{W}_I$ over $b \geq 0$, and $\mathscr{W}_I$ is differentiable at b_I — y^* is differentiable at b_I and c at $y^*(b_I)$, so the chain rule applies — hence Fermat's rule gives

$$0 = \mathscr{W}'_I(b_I) = [\omega - c'(y^*(b_I))]\, y^{*\prime}(b_I) - (\eta - 1)\, [y^*(b_I) + b_I\, y^{*\prime}(b_I)] - c_0.$$

Substituting $c'(y^*(b_I)) = b_I$ and collecting the $y^{*\prime}(b_I)$ terms,

$$(\omega - \eta\, b_I)\, y^{*\prime}(b_I) \;=\; (\eta - 1)\, y^*(b_I) + c_0.$$

Dividing by $y^{*\prime}(b_I)$ and substituting into (39),

$$\lambda^\dagger = \eta\, b_I + \frac{(\eta - 1)\, y^*(b_I) + c_0}{y^{*\prime}(b_I)} + \frac{y^*(b_I)}{y^{*\prime}(b_I)} = \eta\, b_I + (\omega - \eta\, b_I) + \frac{y^*(b_I)}{y^{*\prime}(b_I)} = \omega + \frac{y^*(b_I)}{y^{*\prime}(b_I)},$$

which is (40). Q.E.D.

The market maker limits liquidity because a higher payment to the trader applies to all acquired information. The planner counts the trader's surplus and seeks to offset this monopsony markdown. The payment rate in (40) exceeds ω by the compensation needed to align the market maker's marginal incentive with the planner's. Because public funds are costless, the resulting profit for the market maker is a transfer. Delegation therefore

entails no welfare loss when the contract induces b_I; part (i) states the condition for this implementation.

Iso-elastic acquisition costs yield explicit comparisons between delegation and direct operation when public funds are costly.

Corollary 5 (Closed forms and the cost of delegation)**.** *Let the acquisition cost be iso-elastic as in Corollary 4, $c(y) = k\,y^{1+\chi}/(1+\chi)$ on $[0,\bar{y}]$ with $k\,\bar{y}^{\chi} \geq (1+\chi)\,\omega$; write $n := 1+\chi$, $D_0 := 1+n\,(\eta-1)$, and $d := c_0\,k$ in this corollary and its proof. Let $\mathscr{W}_D$ denote the maximum of (37) over $\lambda \geq 0$, and let*

$$\mathscr{W}_I(b;\mu) \;:=\; \mathscr{W}_I(b) - \mu\,[\eta\,b\,y^*(b) + c_0\,b], \qquad \mathscr{W}_I^* \;:=\; \max_{b\geq 0}\; \mathscr{W}_I(b;\mu), \tag{41}$$

the problem of an integrated planner running the rule itself — choosing b directly and financing its own grossed-up outlay.

(i) (Settlement overhead only.) *Set $\gamma = 0$. The planner's problem (37) has the unique solution*

$$\lambda^* = \frac{\eta\,n\,\omega}{D_0 + \mu\,\eta\,n^2}, \qquad \textit{inducing} \qquad b^* = \frac{\lambda^*}{\eta\,n}, \quad y^*(b^*) = \Big(\frac{\omega}{k\,(D_0+\mu\,\eta\,n^2)}\Big)^{1/\chi}, \tag{42}$$

and λ^* does not depend on the cost scale k; $\lambda^* > \omega$ if and only if $\mu < \chi/(\eta\,n^2)$; the integrated planner uniquely posts $b_I = \omega/(D_0+\mu\,\eta\,n)$; and delegation delivers the welfare share

$$\frac{\mathscr{W}_D}{\mathscr{W}_I^*} \;=\; \Big(\frac{D_0+\mu\,\eta\,n}{D_0+\mu\,\eta\,n^2}\Big)^{1/\chi} \;\leq\; 1, \tag{43}$$

with equality if and only if $\mu = 0$.

(ii) (Quadratic cost.) *Let $\chi = 1$. If $\omega > (1+\mu)\,d$, the planner's problem has the unique solution*

$$b^* = \frac{\omega - (1+\mu)\,d}{2\eta - 1 + 4\mu\eta}, \qquad \lambda^* = 2\eta\,b^* + d, \qquad y^*(b^*) = b^*/k, \tag{44}$$

with $\lambda^* > \omega$ if and only if $\mu < (\omega-d)/[2\eta\,(2\omega-d)]$; the integrated planner uniquely posts $b_I = [\omega-(1+\mu)\,d]/(2\eta-1+2\mu\eta)$; and

$$\frac{\mathscr{W}_D}{\mathscr{W}_I^*} \;=\; \frac{2\eta-1+2\mu\eta}{2\eta-1+4\mu\eta} \;\leq\; 1, \tag{45}$$

with equality if and only if $\mu = 0$. If $\omega \leq (1+\mu)\,d$, both problems have value zero, attained by inducing no acquisition.

Proof. Throughout, $n = 1+\chi$, $D_0 = 1+n(\eta-1)$, $d = c_0 k$, the cost is $c(y) = k\,y^n/n$ on $[0,\bar{y}]$, and the cap satisfies $k\,\bar{y}^{\chi} \geq n\,\omega$. Write $D_D := D_0 + \mu\,\eta\,n^2$ and $D_I := D_0 + \mu\,\eta\,n$.

Step 1: the trader's response at every liquidity. On $[0,\bar{y}]$ the derivative $g_b'(y) = b - k\,y^\chi$ is strictly decreasing in y, so g_b is strictly concave and has a unique maximizer. For $b \in [0, k\bar{y}^\chi]$ the derivative vanishes at $(b/k)^{1/\chi} \le \bar{y}$, so $y^*(b) = (b/k)^{1/\chi}$. For $b > k\bar{y}^\chi$, $g_b'(y) \ge b - k\,\bar{y}^\chi > 0$ on $[0,\bar{y}]$, so g_b is strictly increasing and $y^*(b) = \bar{y}$.

Step 2: the market maker's response. Write $\varphi_\lambda(b) := (\lambda - \eta\,b)\,y^*(b) - c_0\,b$. On $b > k\bar{y}^\chi$ the level is frozen at $\bar{y}$, so φ_λ has slope $-\eta\,\bar{y} - c_0 < 0$ there and every best response lies in $[0, k\bar{y}^\chi]$. At an interior best response $b \in (0, k\bar{y}^\chi)$ the first-order condition (34) holds with the stake in place of λ; solving it for the stake, using $(b/k)^{-1/\chi}\,b = k^{1/\chi} b^{1-1/\chi}$,

$$\lambda \;=\; \eta\,n\,b + \chi\,c_0\,k^{1/\chi}\,b^{1-1/\chi} \;=:\; \Lambda(b). \tag{46}$$

At $\gamma = 0$ the response is explicit for every $\chi > 0$: on $(0, \min\{\lambda/\eta, k\bar{y}^\chi\})$,

$$\log\varphi_\lambda(b) = \log(\lambda - \eta\,b) + \frac{1}{\chi}\log b - \frac{1}{\chi}\log k$$

is strictly concave — the first summand is strictly concave in b, the second concave — with stationary point $b = \lambda/(\eta\,n)$ from $-\dfrac{\eta}{\lambda - \eta b} + \dfrac{1}{\chi b} = 0$. Hence, if $\lambda \le \eta\,n\,k\bar{y}^\chi$, the unique maker best response is $b(\lambda) = \lambda/(\eta\,n)$, which satisfies $\lambda/(\eta n) \le k\bar{y}^\chi$; if $\lambda > \eta\,n\,k\bar{y}^\chi$, then $\log\varphi_\lambda$ is strictly increasing on $(0, k\bar{y}^\chi)$ — its stationary point lies to the right — so the unique response is the corner $b(\lambda) = k\bar{y}^\chi$ with delivered level $\bar{y}$.

Step 3: reduction to a tilted maker problem. Fix $b \in (0, k\bar{y}^\chi]$ and suppose the market maker meets the stake $\Lambda(b)$ with the response b. Two identities follow from Step 1: the trader's first-order condition $k\,y^\chi = b$ at $y = y^*(b)$ gives $c(y) = k\,y^n/n = b\,y/n$ and $k^{1/\chi}b^{1-1/\chi}\,y = b$. The planner's welfare (37) at this stake and response is therefore

$$\begin{aligned}\mathscr{W} &= \omega\,y - \frac{b\,y}{n} - \mu\,[\eta\,n\,b + \chi\,c_0\,k^{1/\chi}b^{1-1/\chi}]\,y - (\eta-1)\,b\,y - c_0\,b \\ &= \omega\,y - \Big[\frac{1}{n} + (\eta-1) + \mu\,\eta\,n\Big]\,b\,y - (1+\mu\chi)\,c_0\,b \;=\; \Big(\omega - \frac{D_D}{n}\,b\Big)\,y^*(b) - (1+\mu\chi)\,c_0\,b \;=:\; w_D(b),\end{aligned}$$

since $\dfrac{1}{n} + (\eta-1) + \mu\,\eta\,n = D_D/n$: on the pass-through branch the planner's problem is the market maker's problem (16) at accuracy value ω, settlement gross-up D_D/n, and capital charge $(1+\mu\chi)\,c_0$. The same two identities reduce the integrated objective (41) on $(0, k\bar{y}^\chi]$ to

$$\mathscr{W}_I(b;\mu) = \Big(\omega - \frac{D_I}{n}\,b\Big)\,y^*(b) - (1+\mu)\,c_0\,b \;=:\; w_I(b),$$

using $\dfrac{1}{n} + (\eta-1) + \mu\,\eta = D_I/n$.

Step 4: settlement overhead only (part i). Set $\gamma = 0$, so $c_0 = 0$ and, by Step 2, every

stake up to $\eta\, n\, k\bar{y}^{\chi}$ passes through as $b(\lambda) = \lambda/(\eta n) = \Lambda^{-1}(\lambda)$. On $0 < b < n\,\omega/D_D$,

$$\log w_D(b) = \frac{1}{\chi}\log b + \log\Big(\omega - \frac{D_D}{n}\, b\Big) - \frac{1}{\chi}\log k$$

is strictly concave — the first summand is concave, the second strictly concave — with stationary point solving $\dfrac{1}{\chi b} = \dfrac{D_D/n}{\omega - (D_D/n)b}$, that is, $\omega = n\,\dfrac{D_D}{n}\, b = D_D\, b$ after collecting the χ terms: the unique on-branch maximizer is $b^* = \omega/D_D$, with stake $\lambda^* = \Lambda(b^*) = \eta\, n\,\omega/D_D$, level $y^*(b^*) = (\omega/(kD_D))^{1/\chi}$, margin $\omega - \dfrac{D_D}{n} b^* = \omega\chi/n > 0$ — so $w_D(b^*) > 0$ — and no k in λ^*; for larger b on the branch, $w_D \le 0$. The candidate is in range: $b^* \le \omega \le k\bar{y}^{\chi}/n < k\bar{y}^{\chi}$ and $\lambda^* \le \eta\, n\,\omega \le \eta\, n\, k\bar{y}^{\chi}$. Stakes with a capped response deliver the frozen level $\bar{y}$, so $\mathscr{W}(\lambda) = \omega\bar{y} - c(\bar{y}) - \mu\lambda\bar{y} - (\eta - 1)\, k\bar{y}^{\chi}\,\bar{y}$ is nonincreasing in λ there and at the branch boundary equals $w_D(k\bar{y}^{\chi}) < w_D(b^*)$ — strict by uniqueness of the on-branch maximizer and $b^* \neq k\bar{y}^{\chi}$ — so capped stakes are dominated and (42) is the unique solution. The markup comparison: $\lambda^* > \omega$ iff $\eta\, n > D_D$ iff $\mu\,\eta\, n^2 < \eta\, n - D_0 = n - 1 = \chi$ — using $\eta n - D_0 = \eta n - 1 - n\eta + n = n - 1$ — that is, $\mu < \chi/(\eta n^2)$. The same argument applied to w_I — log-concave with stationarity $\omega = D_I\, b$ — gives the unique integrated optimum $b_I = \omega/D_I$ with margin $\omega\chi/n$, capped liquidities dominated as above. At the optima, $\mathscr{W}_D = y^*(b^*)\,\omega\chi/n$ and $\mathscr{W}_I^* = y^*(b_I)\,\omega\chi/n$, so

$$\frac{\mathscr{W}_D}{\mathscr{W}_I^*} = \frac{y^*(b^*)}{y^*(b_I)} = \Big(\frac{b^*}{b_I}\Big)^{1/\chi} = \Big(\frac{D_I}{D_D}\Big)^{1/\chi} \;\le\; 1,$$

with equality iff $\mu = 0$, since $D_D - D_I = \mu\,\eta\, n\,(n-1) \ge 0$ vanishes exactly at $\mu = 0$.

Step 5: quadratic cost (part ii). Let $\chi = 1$ and $n = 2$, so $D_D = 2\eta - 1 + 4\mu\eta$, $D_I = 2\eta - 1 + 2\mu\eta$, and write $A := \omega - (1+\mu)\, d$. On the uncapped range $\varphi_\lambda(b) = (\lambda - \eta b)\, b/k - c_0\, b$ is a strictly concave quadratic with vertex $(\lambda - d)/(2\eta)$, so the unique response is $b(\lambda) = \max\{0,\ (\lambda - d)/(2\eta)\}$ for $\lambda \le 2\eta\, k\bar{y} + d$ and the corner $k\bar{y}$ beyond. By Step 3, with $y = b/k$,

$$k\, w_D(b) = A\, b - \frac{D_D}{2}\, b^2, \qquad k\, w_I(b) = A\, b - \frac{D_I}{2}\, b^2.$$

If $A > 0$, the strictly concave quadratic $k\, w_D$ has the unique maximizer $b^* = A/D_D$, with value $A^2/(2D_D) > 0$ and stake $\lambda^* = \Lambda(b^*) = 2\eta\, b^* + d$; it is in range since $b^* \le \omega \le k\bar{y}/2 < k\bar{y}$, and capped stakes are dominated as in Step 4. The markup comparison: $\lambda^* > \omega$ iff $2\eta A > (\omega - d)\, D_D$, and

$$2\eta A - (\omega - d)\, D_D = 2\eta\omega - 2\eta(1+\mu)d - (2\eta - 1)(\omega - d) - 4\mu\eta(\omega - d) = (\omega - d) - 2\mu\eta\,(2\omega - d),$$

so $\lambda^* > \omega$ iff $\mu < (\omega - d)/[2\eta(2\omega - d)]$. The integrated problem gives $b_I = A/D_I$ with

value $A^2/(2kD_I)$, whence $\mathscr{W}_D/\mathscr{W}_I^* = D_I/D_D \leq 1$ with equality iff $\mu = 0$. If $A \leq 0$, then $k\,w_D(b) \leq A\,b < 0$ for every $b > 0$ and likewise $k\,w_I$, while capped responses are again dominated by the branch boundary; stakes $\lambda \leq d$ induce $b = 0$ and welfare zero, so both problems have value zero, attained by inducing no acquisition. Q.E.D.

Costly public funds reduce the planner's willingness to compensate the market maker for accuracy. In the closed-form cases, the optimal payment rate falls below the planner's valuation once the excess burden exceeds the stated threshold. Delegation also requires the planner to finance the market maker's profit in addition to the trading subsidy and operating costs. Financing that profit explains the welfare loss relative to direct operation in (43) and (45).

We next consider a private buyer that values the forecast but excludes the trader's surplus and the market maker's profit from its objective.

Corollary 6 (Double marginalization)**.** *Under the hypotheses of Corollary 5, a private buyer chooses the stake to maximize $(\omega - \lambda)\,y^*(b(\lambda))$, with $b(\lambda)$ the maker best response.*

(i) (Settlement overhead only.) *Set $\gamma = 0$. The buyer uniquely posts $\lambda_B = \omega/n$, independent of η; the market maker responds $b_B = \omega/(\eta\,n^2)$, and the induced acquisition is $y^*(b_B) = (\omega/(\eta\,k\,n^2))^{1/\chi}$.*

(ii) (Quadratic cost.) *Let $\chi = 1$. If $\omega > d$, the buyer uniquely posts $\lambda_B = (\omega + d)/2$; the market maker responds $b_B = (\omega - d)/(4\eta)$, and the induced acquisition is $y^*(b_B) = b_B/k$. If $\omega \leq d$, the buyer's value is zero and it optimally induces no acquisition.*

Proof. Write $\mathscr{W}_B(\lambda) := (\omega - \lambda)\,y^*(b(\lambda))$, with $b(\lambda)$ the maker best response of Step 2 of the proof of Corollary 5.

(i) At $\gamma = 0$ the response is $b(\lambda) = \lambda/(\eta\,n)$ for $\lambda \leq \eta\,n\,k\bar{y}^\chi$, with delivered level $y(\lambda) = (\lambda/(\eta nk))^{1/\chi}$. $\mathscr{W}_B(0) = 0$ and $\mathscr{W}_B(\lambda) > 0$ on $(0, \omega)$; for $\lambda > \omega$ the delivered level is strictly positive while $\omega - \lambda < 0$, so $\mathscr{W}_B(\lambda) < 0$ whether or not the stake is capped. Every maximizer therefore lies in $(0, \omega)$, inside the pass-through range because $\omega \leq k\bar{y}^\chi \leq \eta\,n\,k\bar{y}^\chi$. On $(0, \omega)$,

$$\log \mathscr{W}_B(\lambda) = \log(\omega - \lambda) + \frac{1}{\chi}\log\lambda - \frac{1}{\chi}\log(\eta\,n\,k)$$

is strictly concave — the first summand is strictly concave, the second concave — with stationary point solving $\frac{1}{\chi\lambda} = \frac{1}{\omega - \lambda}$, that is, $\omega - \lambda = \chi\lambda$: the unique maximizer is $\lambda_B = \omega/n$, and η enters $\mathscr{W}_B$ only through the positive constant $(\eta nk)^{-1/\chi}$, so the stake does not depend on it. The response and the induced level follow.

(ii) At $\chi = 1$ the response is $b(\lambda) = \max\{0,\, (\lambda - d)/(2\eta)\}$ for $\lambda \leq 2\eta\,k\bar{y} + d$ and the corner $k\bar{y}$ beyond, so $\mathscr{W}_B(\lambda) = (\omega - \lambda)(\lambda - d)/(2\eta k)$ on $(d,\, 2\eta k\bar{y} + d]$ and $\mathscr{W}_B(\lambda) = 0$ for $\lambda \leq d$. Suppose $\omega > d$. Then $\mathscr{W}_B > 0$ exactly on (d, ω), and every maximizer lies there:

stakes beyond ω deliver a positive level at a negative margin, capped ones included, since $\omega \le k\bar{y}/2 < 2\eta\, k\bar{y} + d$. On (d, ω) the quadratic $2\eta k\, \mathscr{W}_B(\lambda) = (\omega - \lambda)(\lambda - d)$ is strictly concave with vertex $\lambda_B = (\omega + d)/2$: the unique maximizer, with $b_B = (\lambda_B - d)/(2\eta) = (\omega - d)/(4\eta)$ and $y^*(b_B) = b_B/k$. If $\omega \le d$, stakes $\lambda \le d$ give $\mathscr{W}_B = 0$ while stakes $\lambda > d \ge \omega$ give a positive level at a negative margin, so $\mathscr{W}_B(\lambda) < 0$: the maximum is zero, attained by inducing no acquisition. Q.E.D.

The private buyer pays less than its valuation to retain a surplus from the forecast. The market maker then pays the trader less than it receives, after allowing for settlement overhead and capital costs. Each decision reduces compensation for information acquisition. These successive markdowns are the procurement counterpart of double marginalization (Spengler, 1950). The planner with costless public funds instead counts both the trader's surplus and the market maker's profit and raises the payment rate to offset the market maker's markdown.

The mixed market. We now allow noise trading and use the information supply $S(b)$ and delivered information $\mathscr{I}(b)$ of Section 5. The market maker's total expected payment to the informed trader and the arbitrageur is $T(b) := b[\mathscr{I}(b) + (1 - \vartheta)\nu_b] = \vartheta b S(b) + b D_A(b)$ (Appendix A.18). Welfare weights all participants' expected monetary payoffs equally. Contract and trading payments are transfers, while acquisition, settlement overhead, and capital are resource costs. Acquisition costs arise only upon informed arrival (Section 4).

The next lemma characterizes contract revenue, the planner's liquidity incentives, and the welfare contribution of the arbitrageur.

Lemma 8 (The procurement layer in the mixed market)**.** *Under Assumptions 1 and 2 with $\mathbb{E}|U| < \infty$, let the principal value delivered accuracy at $\omega > 0$, finance its payments with costless funds, $\mu = 0$, and set $F = 0$. Define the integrated welfare*

$$\mathscr{W}_I(b) \;:=\; \omega\, \mathscr{I}(b) - \vartheta\, c(S(b)) - (\eta - 1)\, T(b) - c_0\, b, \qquad b > 0, \qquad \mathscr{W}_I(0) := 0. \tag{47}$$

(i) (The contract.) *The contract* (35) *pays on the final price $\delta_A(Q)$, and its expected payment at stake λ is $\lambda\, \mathscr{I}(b)$; the market maker's profit at stake λ is $V_\lambda(b)$ of* (16)*; and the sum of the expected payoffs of the principal, the market maker and the three traders is $\mathscr{W}_I(b)$.*

(ii) (The wedge.) $\mathscr{W}_I(b) - V_\omega(b) = b\, \mathscr{I}(b) - \vartheta\, c(S(b)) = \vartheta\, [b\, S(b) - c(S(b))] - b\, \mathscr{L}_T(b)$.

(iii) (The stake.) *Let $b_I > 0$ maximize $\mathscr{W}_I$ over $[0, \infty)$, let $\mathscr{I}$, S and A be differentiable at b_I, and let b_I be a maker best response to some stake λ with $b_I \in (0, \lambda/\eta)$. Then*

$$(\lambda - \omega)\, \mathscr{I}'(b_I) \;=\; \mathscr{I}(b_I) - b_I\, \mathscr{L}_T'(b_I). \tag{48}$$

If $\mathcal{I}'(b_I) \neq 0$, the stake is $\lambda^\dagger := \omega + [\mathcal{I}(b_I) - b_I\,\mathcal{L}_T'(b_I)]/\mathcal{I}'(b_I)$, which at $\vartheta = 1$ is (40)*; and if b_I is the unique maker best response to $\lambda^\dagger$, posting $\lambda^\dagger$ attains the integrated optimum.*[48]

(iv) (Delegation.) *At any stake, the planner's welfare under any maker best response is at most $\sup_{b\geq 0} \mathcal{W}_I(b)$, with equality if and only if the induced liquidity maximizes $\mathcal{W}_I$.*

(v) (The correction.) *Let $\mathcal{W}_I^0$ be the integrated welfare of the market with no arbitrageur of Appendix A.18, at the same liquidity and acquisition. Then $\mathcal{W}_I(b) - \mathcal{W}_I^0(b) = [\omega - (\eta - 1)\,b]\,D_A(b)$, of the sign of the bracket whenever $D_A(b) > 0$, and $\sup_{b\geq 0} \mathcal{W}_I \geq \sup_{b\geq 0} \mathcal{W}_I^0$.*

Proof. (i) No trade follows the arbitrageur's, so the price standing at settlement is $\delta_A(Q)$, and the contract pays $\lambda\,[B\log(\delta_A(Q)/\delta_0) + (1-B)\log((1-\delta_A(Q))/(1-\delta_0))]$. Conditional on Q the probability of $B = 1$ is $\delta_A(Q)$ (Lemma 4(iii)), so the conditional expected payment is $\lambda\,\mathrm{KL}(\delta_A(Q)\,\|\,\delta_0)$ and the expected payment is $\lambda\,\mathcal{I}(b)$, finite. The market maker receives it and bears the payments of Appendix A.18, so its profit is $V_\lambda(b) = \lambda\,\mathcal{I}(b) - b\,\mathcal{I}(b) - (\eta - 1)\,T(b) - c_0\,b$. The five expected payoffs are $\omega\,\mathcal{I}(b) - \lambda\,\mathcal{I}(b)$ for the principal, $V_\lambda(b)$ for the market maker, $\vartheta\,[b\,S(b) - c(S(b))]$ for the informed trader, whose expected trading profit is $\vartheta b S(b)$ by the ledger and whose acquisition cost is charged on the informed branch, $b\,D_A(b)$ for the arbitrageur and $-(1-\vartheta)\,b\,\nu_b$ for the noise trader, the last three by (15). Summing, the stake cancels against the market maker's receipt, and the three traders' expected profits sum to $b\,\mathcal{I}(b)$ by the ledger, cancelling against the net payout; what remains is $\omega\,\mathcal{I}(b) - \vartheta\,c(S(b)) - (\eta-1)\,T(b) - c_0\,b = \mathcal{W}_I(b)$.

(ii) Subtracting $V_\omega(b) = \omega\,\mathcal{I}(b) - b\,\mathcal{I}(b) - (\eta-1)\,T(b) - c_0\,b$ from (47) leaves $b\,\mathcal{I}(b) - \vartheta\,c(S(b))$, and $b\,\mathcal{I}(b) = \vartheta\,b\,S(b) - b\,\mathcal{L}_T(b)$ by Proposition 1(iii).

(iii) Let $g(b) := b\,S(b) - c(S(b)) = \max_{y\in[0,\bar{y}]}\{b\,y - c(y)\}$, the informed trader's value. For $b' \neq b$, the optimality of $S(b)$ at b and of $S(b')$ at b' give $g(b') \geq b'S(b) - c(S(b)) = g(b) + (b'-b)\,S(b)$ and $g(b) \geq g(b') + (b-b')\,S(b')$, so the difference quotient $[g(b')-g(b)]/(b'-b)$ lies between $S(b)$ and $S(b')$; as S is differentiable, hence continuous, at b_I, g is differentiable at b_I with $g'(b_I) = S(b_I)$. By part (ii), $\mathcal{W}_I = V_\omega + \vartheta g - b\,\mathcal{L}_T$ with $\mathcal{L}_T = \vartheta S - \mathcal{I}$ differentiable at b_I, and V_ω is differentiable at b_I by the product rule, since $\mathcal{I}$ and A are. Fermat's rule at the interior maximizer b_I of $\mathcal{W}_I$ gives

$$0 \;=\; \mathcal{W}_I'(b_I) \;=\; (\omega - \eta b_I)\,\mathcal{I}'(b_I) - \eta\,\mathcal{I}(b_I) - A'(b_I) + \vartheta S(b_I) - \mathcal{L}_T(b_I) - b_I\,\mathcal{L}_T'(b_I),$$

and $\vartheta S(b_I) - \mathcal{L}_T(b_I) = \mathcal{I}(b_I)$. The market maker's marginal condition at b_I, Theorem 1(iv), reads $(\lambda - \eta b_I)\,\mathcal{I}'(b_I) = \eta\,\mathcal{I}(b_I) + A'(b_I)$; substituting it for $\eta\mathcal{I}(b_I) + A'(b_I)$ gives $(\omega -$

[48] If $\mathcal{I}'(b_I) = 0$, which Appendix A.19 shows possible at a maker optimum below the capital bound, (48) forces $\mathcal{I}(b_I) = b_I\mathcal{L}_T'(b_I)$, and the market maker's marginal condition (18) at b_I reads $\eta\,\mathcal{I}(b_I) + A'(b_I) = 0$ at every stake, so b_I is stationary for the market maker at every stake whose cap exceeds it and no stake is singled out.

$\lambda)\,\mathcal{I}'(b_I) + \mathcal{I}(b_I) - b_I\mathcal{L}_T'(b_I) = 0$, which is (48). At $\vartheta = 1$, $\mathcal{L}_T \equiv 0$ and $\mathcal{I} = S = y^*$, so (48) reads $(\lambda - \omega)\,y^{*\prime}(b_I) = y^*(b_I)$, and $\lambda^\dagger = \omega + y^*(b_I)/y^{*\prime}(b_I)$ is (40). The implementation claim is part (iv): if b_I is the unique best response to $\lambda^\dagger$, posting $\lambda^\dagger$ induces b_I and the planner's welfare equals $\mathcal{W}_I(b_I) = \sup \mathcal{W}_I$.

(iv) At a stake met by the best response b, the planner's welfare is the sum of the five expected payoffs, which is $\mathcal{W}_I(b)$ by part (i), and $\mathcal{W}_I(b) \leq \sup_{b' \geq 0} \mathcal{W}_I(b')$ with equality if and only if b maximizes $\mathcal{W}_I$.

(v) In the market with no arbitrageur the delivered accuracy is $\mathcal{I}(b) - D_A(b)$ (Proposition 1(iv)), the overhead falls on the informed trader's payment $\vartheta b S(b)$ alone, the noise trader's loss is collected at par and is a transfer, and the acquisition is the same, so $\mathcal{W}_I^0(b) = \omega\,[\mathcal{I}(b) - D_A(b)] - \vartheta\,c(S(b)) - (\eta - 1)\,\vartheta b S(b) - c_0 b$. Subtracting from (47) and using $T(b) - \vartheta b S(b) = b\,D_A(b)$ gives $\omega\,D_A(b) - (\eta - 1)\,b\,D_A(b)$. For the suprema: at $\eta = 1$ the difference is $\omega D_A(b) \geq 0$ at every b. At $\eta > 1$, by Lemma 2 and Proposition 1(iii), $\mathcal{I} - D_A = \vartheta S - (1 - \vartheta)\nu_b$, so $\mathcal{W}_I^0(b) = [\omega - (\eta - 1)b]\,\vartheta S(b) - \vartheta c(S(b)) - \omega(1 - \vartheta)\nu_b - c_0 b$, which is at most $-c_0 b < 0 = \mathcal{W}_I^0(0)$ for $b \geq \omega/(\eta - 1)$; so $\sup \mathcal{W}_I^0$ is its supremum over $[0, \omega/(\eta - 1))$, where $\mathcal{W}_I \geq \mathcal{W}_I^0$ pointwise since the bracket is positive and $D_A \geq 0$. In either case $\sup \mathcal{W}_I \geq \sup \mathcal{W}_I^0$. Q.E.D.

Noise trading separates the trader's compensation for acquisition from the information the principal receives. The stake condition (48) therefore depends on how liquidity changes the type-confusion loss. The single-trader markup alone no longer captures the planner's incentive to change liquidity.

The arbitrageur improves the final price without changing the informed trader's acquisition choice. At a fixed liquidity, its contribution to welfare is the value of that accuracy improvement less the additional settlement overhead. Its trading profit cancels against the other participants' payments when their payoffs are added. Thus the relevant cost per unit of accuracy gained is $(\eta - 1)b$, rather than the arbitrageur's full compensation. Allowing the planner to choose liquidity, the highest attainable welfare is at least as large with the arbitrageur as without it.

C. The general experiment model behind the acquisition level

The acquisition technology in Section 4 assigns a cost to each feasible information level. We derive this representation from a family of experiments about B. For each level y, the induced cost $\tilde{c}(y)$ is the infimum of acquisition costs among experiments satisfying $\mathbb{E}[\mathrm{KL}(P \,\|\, \delta_0)] = y$. Because the informed trader's expected trading profit is $b\,y$, it has no incentive to choose a more costly experiment that supplies the same level.

Lemma 10 constructs a family with interior posteriors that implements every level in $[0, \bar{y}]$, and Lemma 11 establishes the equivalence between choosing an experiment and choosing a level at its induced cost.[49] This reduction concerns the informed trader's acquisition problem. The distribution of posterior beliefs remains relevant to the arbitrageur's inference and to delivered information when trader type is unobserved.

Experiments. An *experiment* about B is a signal S_α with realization s, whose law depends on B; it yields the posterior $P_\alpha(s) = \Pr(B = 1 \mid S_\alpha = s) \in [0, 1]$, a random variable through s. Experiments are ranked by informativeness in the Blackwell order[50] and carry an acquisition cost $c_{\mathcal{A}}(\alpha) \geq 0$ over a family indexed by $\alpha \in \mathcal{A}$. An experiment *implements* the level $I(\alpha) := \mathbb{E}_s[\mathrm{KL}(P_\alpha(s) \,\|\, \delta_0)]$, the expectation taken over the prior-predictive (marginal) law of s.

The next lemma records the posterior identities behind the acquisition level.

Lemma 9 (Posterior identities)**.** *For any experiment with posterior $P_\alpha(s)$ and implemented level $I(\alpha)$:*

(i) $\mathbb{E}_s[P_\alpha(s)] = \delta_0$*;*

(ii) $I(\alpha) = I(B; S_\alpha)$*, the Shannon mutual information; consequently $I(\alpha) \geq 0$, with equality iff S_α is uninformative ($P_\alpha(s) = \delta_0$ almost surely), and the implemented level is nondecreasing in the Blackwell order;*

(iii) $\mathbb{E}_s[h(P_\alpha(s))] = h(\delta_0) - I(\alpha)$*.*

Proof. (i) By the law of total probability, $\mathbb{E}_s[P_\alpha(s)] = \mathbb{E}[\Pr(B = 1 \mid S_\alpha)] = \Pr(B = 1) = \delta_0$.

(ii) Let μ_α be the prior-predictive law of S_α. Mutual information can be written as

$$I(B; S_\alpha) = \int \mathrm{KL}(P_{B|S_\alpha=s} \,\|\, P_B)\, d\mu_\alpha(s).$$

For each realization s, the posterior law $P_{B|S_\alpha=s}$ is Bernoulli with parameter $P_\alpha(s) = \Pr(B = 1 \mid S_\alpha = s)$, while the prior law P_B is Bernoulli with parameter $\delta_0 = \Pr(B = 1)$. Hence

$$\mathrm{KL}(P_{B|S_\alpha=s} \,\|\, P_B) = P_\alpha(s) \log \frac{P_\alpha(s)}{\delta_0} + (1 - P_\alpha(s)) \log \frac{1 - P_\alpha(s)}{1 - \delta_0} = \mathrm{KL}(P_\alpha(s) \,\|\, \delta_0),$$

[49] The equivalence applies to the specified set of feasible levels. The cap $\bar{y} < h(\delta_0)$ and any additional restrictions on acquisition costs remain separate assumptions.

[50] The Blackwell order ranks experiments by informativeness: S_{α_2} is Blackwell more informative than S_{α_1} if S_{α_1} is a *garbling* of S_{α_2}—its realization passed through an independent noisy channel, so that $B \to S_{\alpha_2} \to S_{\alpha_1}$ is a Markov chain: S_{α_1} is conditionally independent of B given S_{α_2}, i.e. $\Pr(S_{\alpha_1} \in A \mid S_{\alpha_2}, B) = \Pr(S_{\alpha_1} \in A \mid S_{\alpha_2})$ for every event A.

with the usual convention $0 \log 0 = 0$. Integrating over μ_α gives $I(B; S_\alpha) = I(\alpha)$. Gibbs' inequality gives $\mathrm{KL}(P_{B|S_\alpha=s} \,\|\, P_B) \geq 0$ for μ_α-almost every s, with equality if and only if $P_{B|S_\alpha=s} = P_B$. Therefore $I(\alpha) = 0$ if and only if $P_\alpha(s) = \delta_0$ almost surely, i.e. the signal is uninformative about B. Finally, if S_{α_1} is Blackwell less informative than S_{α_2}, then S_{α_1} is a garbling of S_{α_2}, so $B \to S_{\alpha_2} \to S_{\alpha_1}$ is a Markov chain. The data-processing inequality then gives $I(B; S_{\alpha_1}) \leq I(B; S_{\alpha_2})$.

(iii) Expanding the Bernoulli relative entropy,

$$\mathrm{KL}(P_\alpha(s) \,\|\, \delta_0) = -h(P_\alpha(s)) - P_\alpha(s) \log \delta_0 - (1 - P_\alpha(s)) \log(1 - \delta_0).$$

Taking the prior-predictive expectation and using *(i)*,

$$I(\alpha) = -\mathbb{E}_s[h(P_\alpha(s))] - \delta_0 \log \delta_0 - (1 - \delta_0) \log(1 - \delta_0) = h(\delta_0) - \mathbb{E}_s[h(P_\alpha(s))]. \qquad \text{Q.E.D.}$$

The information level measures the improvement in the informed trader's forecast relative to the prior, evaluated by expected log-loss. The final market price need not achieve the same improvement because the arbitrageur does not observe trader type. The next lemma constructs experiments that implement every level in the feasible interval with interior posterior beliefs.

Lemma 10 (Feasibility of every level)**.** *Fix $\bar{y} \in (0, h(\delta_0))$. There is a Blackwell-ordered family of experiments $\{S_\alpha\}_{\alpha \in [0,1)}$ with almost-surely interior posteriors whose implemented level $I(\alpha)$ is continuous and nondecreasing, with $I(0) = 0$ and $I(\alpha) \to h(\delta_0)$ as $\alpha \to 1$. Consequently every $y \in [0, \bar{y}]$ is implemented by an experiment with almost-surely interior, Bayes-plausible posteriors, and $\alpha(y) := \min\{\alpha \in [0, 1) : I(\alpha) = y\}$ indexes a Blackwell-ordered family by the level itself.*

Proof. Take the symmetric binary channel: $S_\alpha \in \{+, -\}$ with

$$\Pr(S_\alpha = + \mid B = 1) = \Pr(S_\alpha = - \mid B = 0) = a(\alpha) := \frac{1}{2} + \frac{\alpha}{2}, \qquad \alpha \in [0, 1).$$

Bayes' rule gives the posteriors and signal probabilities

$$P_\alpha(+) = \frac{\delta_0 \, a(\alpha)}{w(\alpha)}, \quad P_\alpha(-) = \frac{\delta_0 \, (1 - a(\alpha))}{1 - w(\alpha)}, \quad w(\alpha) := \Pr(S_\alpha = +) = \delta_0 \, a(\alpha) + (1 - \delta_0)(1 - a(\alpha)).$$

Since $a(\alpha) \in [\frac{1}{2}, 1)$ and $\delta_0 \in (0, 1)$, both posteriors lie in $(0, 1)$: the posteriors are interior for every $\alpha \in [0, 1)$, and they are Bayes-plausible by Lemma 9(i).

Blackwell order. Fix $0 \leq \alpha_1 < \alpha_2 < 1$ and write $a_i = a(\alpha_i)$. Passing S_{α_2} through an

independent symmetric binary channel with accuracy

$$a_g = \frac{a_1 + a_2 - 1}{2a_2 - 1}$$

reproduces the conditional law of S_{α_1} given each state: flipping a symmetric signal of accuracy a_2 through an independent symmetric channel of accuracy a_g yields a symmetric signal of accuracy $a_2 a_g + (1 - a_2)(1 - a_g)$ (correct exactly when both stages are correct or both are wrong), and substituting a_g gives $a_2 a_g + (1-a_2)(1-a_g) = a_1$. Moreover $a_g \in [\frac{1}{2}, 1]$: it equals 1 at $a_1 = a_2$, equals $\frac{1}{2}$ at $a_1 = \frac{1}{2}$, and is increasing in a_1. Hence S_{α_1} is a garbling of S_{α_2}, so the family is Blackwell-ordered in α, and I is nondecreasing by Lemma 9(ii).

Level map. $I(\alpha) = w(\alpha)\,\mathrm{KL}(P_\alpha(+) \,\|\, \delta_0) + (1 - w(\alpha))\,\mathrm{KL}(P_\alpha(-) \,\|\, \delta_0)$ is a finite composition of continuous maps of α on $[0, 1)$, hence continuous. At $\alpha = 0$ both posteriors equal δ_0, so $I(0) = 0$. As $\alpha \to 1$, $P_\alpha(+) \to 1$, $P_\alpha(-) \to 0$, and $w(\alpha) \to \delta_0$, so

$$I(\alpha) \;\to\; \delta_0 \log \frac{1}{\delta_0} + (1 - \delta_0) \log \frac{1}{1 - \delta_0} \;=\; h(\delta_0).$$

Conclusion. Fix $y \in [0, \bar{y}]$. Since $\bar{y} < h(\delta_0)$, continuity of I, $I(0) = 0$, and $I(\alpha) \to h(\delta_0)$ give some α with $I(\alpha) = y$ by the intermediate value theorem. The set $\{\alpha \in [0, 1) : I(\alpha) = y\}$ is closed in $[0, 1)$ by continuity and bounded away from 1 (levels near 1 exceed $\bar{y}$), so its infimum is attained and $\alpha(y)$ is well defined. Because I is nondecreasing, $y \mapsto \alpha(y)$ is nondecreasing, so the reindexed family $\{S_{\alpha(y)}\}_{y \in [0, \bar{y}]}$ inherits the Blackwell order. Q.E.D.

Lemma 11 (Exactness of the induced-cost reduction)**.** *Let $\{S_\alpha\}_{\alpha \in \mathscr{A}}$ be an experiment family with cost $c_{\mathscr{A}}(\alpha) \geq 0$ and implemented levels $I(\alpha)$, let $Y \subseteq I(\mathscr{A})$ be nonempty, and let $\mathscr{A}_Y := \{\alpha \in \mathscr{A} : I(\alpha) \in Y\}$. Define the induced cost $\tilde{c}(y) := \inf\{c_{\mathscr{A}}(\alpha) : \alpha \in \mathscr{A},\ I(\alpha) = y\}$. Then, for every b,*

$$\sup_{\alpha \in \mathscr{A}_Y} \{b\, I(\alpha) - c_{\mathscr{A}}(\alpha)\} \;=\; \sup_{y \in Y} \{b\, y - \tilde{c}(y)\}.$$

If in addition the infimum defining $\tilde{c}(y)$ is attained for every $y \in Y$, the two problems have the same maximizers up to implementation: $\hat{\alpha}$ maximizes the left side iff $\hat{y} := I(\hat{\alpha})$ maximizes the right side and $c_{\mathscr{A}}(\hat{\alpha}) = \tilde{c}(\hat{y})$. In particular the supplied information coincides.

The trader values differences among experiments with the same information level only through their acquisition costs. When the infimum cost at each level is attained, choosing a level at its induced cost preserves the trader's optimal payoff and information acquisition.

Proof. For $y \in Y$ the level set $\{\alpha \in \mathscr{A} : I(\alpha) = y\}$ is nonempty (because $Y \subseteq I(\mathscr{A})$), and on

it $b\,I(\alpha) - c_{\mathscr{A}}(\alpha) = b\,y - c_{\mathscr{A}}(\alpha)$, so

$$\sup_{\alpha:\, I(\alpha)=y} \{b\,I(\alpha) - c_{\mathscr{A}}(\alpha)\} = b\,y - \tilde{c}(y),$$

by the definition of $\tilde{c}(y)$ as the infimum of $c_{\mathscr{A}}$ on the level set. Since $\mathscr{A}_Y = \bigcup_{y\in Y}\{\alpha : I(\alpha) = y\}$, taking the supremum over $y \in Y$ of both sides gives the displayed equality (a supremum over a union is the supremum of the per-set suprema).

For the maximizer correspondence, suppose the infimum is attained for every $y \in Y$. If $\hat{\alpha}$ maximizes the left side, set $\hat{y} = I(\hat{\alpha})$; then

$$b\,\hat{y} - c_{\mathscr{A}}(\hat{\alpha}) \;\le\; b\,\hat{y} - \tilde{c}(\hat{y}) \;\le\; \sup_{y\in Y}\{b\,y - \tilde{c}(y)\} \;=\; b\,\hat{y} - c_{\mathscr{A}}(\hat{\alpha}),$$

the first inequality by definition of the induced cost, the second because $\hat{y} \in Y$, and the equality by the displayed identity and the optimality of $\hat{\alpha}$. Hence both inequalities are equalities: $c_{\mathscr{A}}(\hat{\alpha}) = \tilde{c}(\hat{y})$ and $\hat{y}$ maximizes the right side. Conversely, if $\hat{y}$ maximizes the right side and $\hat{\alpha}$ attains the induced cost at $\hat{y}$ (so $I(\hat{\alpha}) = \hat{y}$ and $c_{\mathscr{A}}(\hat{\alpha}) = \tilde{c}(\hat{y})$), then $b\,I(\hat{\alpha}) - c_{\mathscr{A}}(\hat{\alpha}) = b\,\hat{y} - \tilde{c}(\hat{y})$ equals the common supremum, so $\hat{\alpha}$ maximizes the left side. Q.E.D.

Countable menus and public lotteries. The sampling technology of Section 4.1 implements a countable set of levels, whereas Assumption 3 takes the feasible set to be the interval $[0, \bar{y}]$. Public lotteries close the gap. Let the trader draw $J \in \{N, N+1\}$ independently of B, reveal J, and take J samples: the level map extends to a continuous, strictly increasing function onto $[0, y_\infty)$, which covers $[0, \bar{y}]$ once κ is small enough that $y_\infty > \bar{y}$. Allowing the lottery to mix any two sample sizes, the induced cost over lotteries is the lower convex envelope of the points $\{(y(N), c_N(N)) : N \in \mathbb{N}_0\}$.

D. Numerical illustrations: setup and parameter values

This section states the design of the multi-period market of Section 4.2 and the parameter values of the numerical illustrations. In the multi-period market, the market maker commits b before trading, n periods follow, and the contracts settle once, after period n, when B is revealed. Period t opens at the public price δ_{t-1} and is the trading of Section 3 with δ_{t-1} as the prior: an informed trader, who chooses its precision by solving (13) at that prior, or the noise trader submits the first order, an arbitrageur trades once to the posterior (5), and the final price becomes δ_t. The standardized innovations Z of the measurements, noise orders, and arrivals are independent across periods and of the event.

The informed trader and the arbitrageur in each period are new participants who trade

once and hold their positions until settlement. Their payoffs therefore do not depend on subsequent prices, and Lemma 1 applies conditional on the public history. The history affects trading in period t only through the opening price δ_{t-1}, so the expressions in Section 3 apply with δ_{t-1} as the prior. We denote the resulting delivered information and type-confusion loss by $\mathcal{I}(b;\delta_{t-1})$ and $\mathcal{L}_T(b;\delta_{t-1})$, respectively.

Cumulative delivered information satisfies $\mathbb{E}[\mathrm{KL}(\delta_n \,\|\, \delta_0)] = \sum_{t=1}^{n} \mathbb{E}[\mathcal{I}(b;\delta_{t-1})]$, with expectations over the public price: Proposition 1(ii) at each prior gives $\mathbb{E}[h(\delta_t) \mid \delta_{t-1}] = h(\delta_{t-1}) - \mathcal{I}(b;\delta_{t-1})$, so $\sum_{t=1}^{n} \mathbb{E}[\mathcal{I}(b;\delta_{t-1})] = h(\delta_0) - \mathbb{E}[h(\delta_n)]$, and (20) with $\mathbb{E}[\delta_n] = \delta_0$, from part (i) at each prior, gives the same value for $\mathbb{E}[\mathrm{KL}(\delta_n \,\|\, \delta_0)]$. The cumulative type-confusion loss is the corresponding sum $\sum_{t=1}^{n} \mathbb{E}[\mathcal{L}_T(b;\delta_{t-1})]$.

Figures 1 to 3 use the following baseline parameter values, with one parameter varied in each panel: $\delta_0 = 1/2$, $\vartheta = 1/2$, $b = 1$, $U \sim N(0,1)$, and the power cost $c_\rho(\rho) = \kappa_c \rho^p / p$ with $\kappa_c = 0.2$ and $p = 4$; panel (b) of Figure 1 fixes $\rho = 1$; Figure 2 takes $\eta = 1.2$ and $c_0 = 0.12 \log 2$; and the multi-period market has $n = 12$ periods, with b varied over $\{0.5, 1, 2\}$ and ϑ over $\{0.2, 0.5, 0.8\}$.

E. Prediction Markets in Practice

A data marketplace distributes products its providers produce. Databricks Marketplace, Snowflake Marketplace, and AWS Data Exchange catalog datasets, models, applications, and updating feeds that providers publish and consumers discover, license, or subscribe to (Databricks, 2026; Snowflake, 2026; Amazon Web Services, 2026), and China's data exchanges, from Guiyang in 2015 to Shanghai, Beijing, and Shenzhen, register and standardize transactions in data products that firms list (U.S.-China Economic and Security Review Commission, 2026). In each, the provider produces the content outside the marketplace, which lowers the cost of finding, contracting for, and delivering it. A prediction market's product is produced by the trading itself. Traders make it through the positions they take, are paid by what those positions earn rather than by a license fee, and keep the raw information behind their orders, so the marketplace buys the implication of a private signal without the signal changing hands, the disclosure problem of Arrow (1962) sidestepped in part by contingent compensation. The settlement then scores the product against the outcome, a label no listed dataset carries.

The product's consumers are visible. Intercontinental Exchange distributes normalized Polymarket probabilities to institutional traders in near real time and as historical series (Intercontinental Exchange, 2026), venues provide programmatic access to prices, order books,

and trade histories,[51] and financial news publishers present prices as forecasts to readers.[52] Who finances the depth a trader faces differs by venue. Under an order book the traders' own bids and asks supply it: Kalshi, designated a contract market in 2020 (Commodity Futures Trading Commission, 2020), runs an order book, and Polymarket began with an automated market maker whose liquidity pools participants and the platform funded (Commodity Futures Trading Commission, 2022), then adopted an order book matching orders off chain and settling them on chain (Polymarket, 2026a); both venues reward traders for resting orders.[53] An automated market maker instead fills every order at prices its rule computes, so a sponsor that values the forecast can fund the depth itself. Inkling Markets, Consensus Point, Yahoo!, and Microsoft have used the LMSR (Othman et al., 2013), and Futuur reported running a liquidity-sensitive variant in 2024.[54] Gnosis implemented the LMSR on chain and subsequently a fixed-product market maker that outside liquidity providers can fund; it identifies potentially costly on-chain logarithms and difficulty crowdfunding liquidity dynamically as drawbacks of the LMSR (Gnosis, 2020; Koeppelmann, 2020).

In every LMSR deployment, someone must choose a number for the liquidity. At Carnegie Mellon's Gates Hillman market, the administrator fixed the liquidity before trading began, but it proved too low late in the market's life (Othman and Sandholm, 2013; Othman et al., 2013). The Gnosis owner sets the liquidity through committed funding and may add or withdraw funding, pause the rule, or close the market before resolution (Gnosis, 2020; Koeppelmann, 2020). Implementers describe setting the liquidity as "more art than science" (Othman et al., 2013).

That choice is the model's object. The market maker posts the rule, commits the liquidity before the trader acquires information, and holds it for the market's life. Because deployed code of this kind leaves the owner free to change the funding, the commitment rests on the venue rather than on the technology, and the characterization gives the number the venue commits to.

[51] See Polymarket (2026c) and Kalshi (2026a).

[52] Polymarket and Dow Jones announced on January 7, 2026 that Polymarket's prices would appear across *The Wall Street Journal*, *Barron's*, *MarketWatch*, and *Investor's Business Daily* (Polymarket and Dow Jones, 2026).

[53] See Kalshi (2026b,c) and Polymarket (2026b).

[54] See Futuur (2024), which names the liquidity-sensitive market maker of Othman et al. (2013).