

The Cost of Being First

Aryan Ayyar

Manipal Academy of Higher Education

In a price-time-priority market, an order ahead in the queue executes more often but can earn less. Priority adds the fills in which demand stops at that order, and such small purchases can come from informed traders. In our limit order book, dealers cancel and resubmit at random times, and customers can arrive between their moves. We construct an equilibrium in which an order fills more often than the one behind it yet earns less, while every dealer profits from supplying liquidity. Dealers also pay to give up a place that is worth more than the back of the queue at today's book. They do so to avoid being moved up into a worse position when a rival ahead cancels. In our example this requires intermediate revision speed. A two-round extension adds learning, repricing and repeated informed trading. The model yields testable restrictions on fills and cancellations by queue position.

October 2026

Keywords: queue priority; adverse selection; limit order book; cancellation; liquidity supply.

JEL classification: D82, G14.

1. Introduction

Most equity and futures markets run a limit order book with price-time priority. At a given price, the order that arrived first executes first. Traders and exchanges treat a place at the front of the queue as valuable: it raises the chance that a resting order is filled. Yet a fill is worth only what the trade behind it reveals. This paper asks when being ahead in the queue is costly, and how dealers who supply liquidity respond when it is.

The answer starts from a simple comparison. Take two equal sell orders at the same price, one directly behind the other. Both are filled when a buyer's demand reaches past them, and neither is filled when demand stops short of both. They differ in one case only: when demand stops exactly at the first order. Moving one step forward therefore adds exactly the fills in which the buyer stops at you. If those small purchases tend to come from an informed trader who knows the asset is worth more than the price, the extra fills lose money. The order farther back is protected from them. Informed traders often limit the size of their orders because larger trades are costly to them; in our model this comes from rising trading and inventory costs. Their demand then tends to stop early in the queue, while large liquidity needs sweep it.

This paper shows three things. (i) *Harmful priority with profitable supply*. In equilibrium, an earlier order can fill more often than the order behind it and still earn less, while every dealer chooses to enter and earns a positive profit. (ii) *Strategic priority surrender*. Dealers who can revise their orders at random times pay a fee to cancel an order that is worth more than the back of the queue at today's book, and rejoin the back. They do this to avoid being moved up into a worse position when a rival ahead of them cancels. We call this paid move to the back *priority surrender*. (iii) *Learning and repeated informed trading*. In a short second round, the same stopping events that make priority costly also move beliefs about the asset and about the informed trader's inventory; dealers reprice, and the same informed trader can buy again.

We study a one-sided market in which a finite number of dealers each offer one unit for sale at a price chosen from a grid. A single customer arrives at a random time, buys, and then the asset pays off. The customer is either an informed trader, who knows the asset's value and faces rising trading and inventory costs. Otherwise it is a liquidity trader, who wants up to a fixed number of units and values them at the public mean plus an urgency premium. Both buy with a market order and take the cheapest and earliest units first. Before the customer arrives, each dealer receives revision opportunities at random times. At each opportunity it can keep its order, cancel it, move it to another price, or cancel and resubmit at the same price for a small fee, which places it at the back of the queue. We call the last action a *reset*. Dealers cannot react instantly: a customer can arrive between any two moves, and a rival can move again after you.

The first result concerns the value of queue position. Large liquidity purchases sweep the whole queue and pay every dealer a profit. A purchase by the informed trader loses money at every order it reaches. What matters for priority is where it stops: the order just behind the stop escapes the loss. In our main example, four dealers sell at a price of 30, and the informed trader, when the asset is worth 100, buys exactly two units. The order at rank two is therefore filled by the informed trader, while the order at rank three is not. Rank two is worth 3.2 and rank three 4.1, although rank two is filled more often. All four ranks are profitable, because liquidity traders often buy the whole queue. We prove that this pattern is part of an equilibrium in which dealers choose whether to enter, which price to post, and when to revise. It survives small changes in the probabilities, values, costs, rates and fees, holding fixed the price grid, the set of customer types and the symmetry of dealers.

Supporting this equilibrium requires ruling out every alternative a dealer has: posting another price and repairing it later, waiting outside and entering later, withdrawing, and selling immediately. One idea handles all of these. A dealer cannot fix a bad quote before its next revision opportunity, and that opportunity arrives before the customer only with a probability below one. Simple upper bounds on what any alternative can earn then show that posting 30 and taking a place in the queue is best. Undercutting to 25 would buy the front of the queue,

but it halves the profit on every liquidity fill. Several prices on the grid can make money, yet in our example no dealer gains by leaving 30.

The second result explains why a dealer would give up a valuable position. Consider the dealer at rank three in the example. Moving to the back today looks like a mistake: rank four is worth 4, rank three 4.1, and the reset costs a fee of 0.01. The immediate gain is -0.11 . But the dealer at rank two resets whenever it can. If our dealer stays, a single such reset moves it up into rank two, where the informed trader's demand stops. If it moves to the back, it must be pushed up twice before it reaches rank two; the customer may well arrive first, and our dealer can reset again on the way. The fee buys time away from rank two. Once rivals' future moves and the chance that the customer arrives first are taken into account, resetting gains 0.02.

In equilibrium, the dealers at ranks two and three reset at every opportunity, and the dealers at ranks one and four keep their orders. The front order never moves; the other three dealers keep cycling through ranks two to four. The reset at rank three is strategic priority surrender: the dealer pays to leave a place that is worth more than the back at today's book. The reset at rank two is not, because rank two is already the worst place at today's book.

This behavior depends on how fast dealers can revise relative to the arrival of demand. Speed is relative here: what matters is a dealer's revision rate against the rate of demand opportunities, including those that end without a trade. A closed form for the example shows that, under this policy, the dealer at rank three resets only when its revision rate exceeds one seventh of that rate and the fee is small enough. If revision is too slow, the protection rarely matters before trading occurs. If revision is very fast and costly, this policy is no longer an equilibrium, because the fees paid in repeated cancellations outweigh the gain. Resetting does not change the total of the dealers' terminal surplus claims; it changes who bears which fills. Fees and waiting costs reduce the dealers' total surplus.

The third result adds a short second round of trading. After the first customer, the public sees how many units were bought but not who bought them. Two fills can come from a liquidity trader who wanted two units or from an informed trader, so beliefs must track both the asset's value and the informed trader's inventory. After no fill, the price stays at 30. After two fills, the expected value rises enough that every dealer reprices to 80. In some histories the same informed trader buys two units in the first round and one more in the second. The equilibrium holds if the chance of continuation is small enough that today's incentives dominate the future.

The model gives a precise signature for the mechanism. It predicts a same-price cancellation with no new public information, by a dealer whose immediate move to the back loses money. At the same time, a rival ahead of it resets at a positive rate, so staying risks being moved up into a worse position. Each order remains profitable overall, yet the margin between two adjacent positions is negative where informed demand stops. Testing this requires order-level data

that link cancellations and resubmissions to the same owner and record queue positions, not only depth at each price. The model also has an implication for markouts. A markout measures the value of a step forward without bias only if its benchmark is the expected value given information that includes the event that demand stopped there. Otherwise the error can have either sign.

The results rest on a stylized market. Prices lie on a coarse grid, and our bounds keep all dealers at one quote, so the model does not study fine price competition; with a finer tick, undercutting can pay in a simpler simultaneous game without revision (Appendix B). There are finitely many dealers, each displaying one unit, and the book is one-sided. Customers use only market orders. The learning extension has two rounds and a small chance of continuation rather than an ongoing stream of trades. Fine-tick competition, endogenous depth and a stationary market with repeated informed trading remain open.

Related literature. Adverse selection is the classic cost of supplying liquidity. Glosten and Milgrom (1985) show that a dealer facing informed traders must set an ask above the expected value to break even, and Kyle (1985) studies how an informed trader times and sizes orders when prices respond to order flow. Glosten (1994) derives the price schedule that an open limit order book supports under asymmetric information. We take adverse selection as given and ask how it is allocated across positions in the queue at a single price.

A second strand models order placement in limit order books. Parlour (1998) shows that an order changes the decisions of later traders and so its own execution probability, in a market where orders cannot be revised. Foucault et al. (2005) study how patient and impatient traders shape the book when waiting is costly, and Goettler et al. (2005) solve a dynamic limit order market numerically. Roşu (2009) builds a continuous-time model in which traders arrive at random and can cancel and revise instantly and at no cost, under symmetric information. Roşu (2020) adds informed traders who choose between market and limit orders. Our game uses Roşu’s random arrival of events and Parlour’s sequential, forward-looking behavior. It is not a special case of either. Dealers revise at a finite rate, and the informed trader chooses how much to buy. A literature on speed studies how faster trading and quote updating change liquidity suppliers’ exposure to being picked off by fast traders (Menkveld and Zoican, 2017; Baldauf and Mollner, 2020). Speed matters here for a different reason: it sets how often rivals reshuffle the queue before the next customer arrives.

Closest to our question are models of queue position itself. Moallemi and Yuan (2017) separate the immediate value of an execution from the option value of a place in the queue, holding quotes fixed. Garriott et al. (2025) study how queue position and adverse selection interact with the quantities market makers quote and with their inventories. Yueshen (2026) studies uncertainty about where an order lands in the queue, optional revision once positions are known, and positions that end up losing money. In our model positions are known and every order is profitable as a whole. The cost lies in the step between two

adjacent positions, and dealers manage it by resetting repeatedly in anticipation of their rivals' cancellations.

Section 2 presents the model. Section 3 derives the value of queue position and the main equilibrium. Section 4 studies strategic priority surrender. Section 5 adds learning and repeated informed trading. Section 6 develops empirical implications, and Section 7 concludes. Proofs are in Appendix A; Appendix B discusses the boundaries of the results.

2. The model

The model has three parts. Dealers form a one-sided limit order book and then revise their quotes at random times while they wait for a customer. A customer, who may be informed, arrives at a random time and buys from the front of the queue. This section sets out each part, the timing, and the equilibrium concept.

2.1 Asset and traders

There is one risky asset and a cash numeraire. The asset has terminal value v , with finite support in $[0, H]$, and all traders share the common prior π for v , whose mean is m . There are $N < \infty$ dealers. Each dealer is risk neutral and uninformed, and holds one unit of the asset. A dealer cannot sell short, display more than one unit, or switch to buying. An outside bid $B \leq m$ lets any dealer sell its unit immediately at B and leave the market for good.

The book therefore has a single side: dealers sell, and customers buy. We take this restriction from the tractable side of the analysis in Roşu (2009). A finite, nonnegative price grid $\mathcal{P}$ is common knowledge; its spacing is the tick size. All random objects are independent: the value v , each dealer's revision clock, the customer's arrival time and the customer's type. No public signal about v appears before the customer arrives, and two events never occur at the same moment.

2.2 Order book, priority and revision

Execution follows price-time priority: lower asks execute first, and at a common price the earlier order executes first. The state of the book, b , records each dealer's live quote, its relative time priority at that price, and whether the dealer is active, empty or permanently gone. A dealer is *active* if it has a live quote and *empty* if it does not. Relative order at each price is all that matters, so exact timestamps are not needed. We call a dealer's place in the queue its *rank*, and write "the dealer at rank k " for the owner of the k th unit in line.

Dealers first form the book. They move one at a time, in a publicly known order, and each either abstains, posts one quote $p \in \mathcal{P}$ for a display fee $f \geq 0$, or sells at B . After formation, dealer i can act only at the jumps of its own Poisson clock, which has rate $\nu_i > 0$. At each such revision opportunity a dealer who has not exited may

- (a) retain its quote;

- (b) withdraw it at no fee;
- (c) post or replace a quote at any price in $\mathcal{P}$ for a fee $c_i \geq 0$;
- (d) sell at B and exit.

A replacement cancels the old order and takes the last time priority at the new price.

One replacement matters most. A *reset* is a cancel and resubmit at the same price, so the dealer joins the back of the queue at that price. Withdrawn and initially absent dealers may enter at any later revision opportunity of their own, and there is no limit on the number of revisions. A dealer pays a waiting cost $h_i \geq 0$ per unit of time while its quote is live.

The finite revision rate is the key friction. Between two of its own opportunities a dealer cannot touch its order, so a customer can arrive after a rival has moved but before the dealer can respond. The game also differs from a fixed revision schedule, because no one is guaranteed to move last: a rival may always move again after you.

2.3 Customers

At time τ , which is exponential with rate $\lambda > 0$, a single customer opportunity arrives, and the customer type is drawn at that moment. With probability a the customer is an informed trader who knows v ; we also call it the *insider*. With probability $b_{d,u}$ the customer is a liquidity trader with capacity d and urgency u . With the remaining probability b_0 there is no customer. The subscripted $b_{d,u}$ and b_0 are these type probabilities, not books. These probabilities sum to one and are independent of v and of all clocks. Thus λ is the rate at which a demand opportunity ends the trading stage, and that opportunity may end with no trade. When λ is matched to data, it is therefore not the rate of observed trades.

Customers use market orders only. A customer buys a prefix of the book in price-time order, paying each order's own ask, or buys nothing. It cannot post a limit order, and dealers cannot cancel after seeing demand or value. We call the book at the moment the customer arrives the *final book*. The asset then settles. A dealer whose unit is sold earns $p_i - v$ relative to keeping the unit, and zero if the unit is unsold; a market sale earns $B - v$. This trading surplus, before display, replacement and waiting costs, is the dealer's *gross surplus*. Net surplus subtracts the display fee, any replacement fees, and h_i times the time the quote was live. All comparisons below concern expected net surplus.

Let the book hold n asks ordered as $p_1 \leq \dots \leq p_n$. An informed trader with inventory z chooses $q \in \{0, \dots, n\}$ to maximize

$$qv - \sum_{j=1}^q p_j - \{C(z+q) - C(z)\} - K(q). \quad (1)$$

Here C is an inventory cost and K a cost of trading quantity q at once. Write $\kappa_j(z)$ for the informed trader's marginal cost of the j th unit, which we assume

is nonnegative and nondecreasing in j . A liquidity trader of type (d, u) chooses $q \leq \min(d, n)$ to maximize

$$q(m + u) - \sum_{j=1}^q p_j. \quad (2)$$

At a tie, each customer takes the smaller quantity. The next lemma is standard; it fixes how customer demand reaches each rank.

Lemma 1 (Optimal demand). *Let marginal costs $\kappa_j(z)$ be nonnegative and nondecreasing. Then:*

- (i) *the informed trader buys exactly $\#\{j : v > p_j + \kappa_j(z)\}$ units;*
- (ii) *a liquidity trader buys exactly $\#\{j \leq \min(d, n) : m + u > p_j\}$ units;*
- (iii) *every unit bought by the informed trader has $p_j < v$;*
- (iv) *whether demand reaches rank k does not change when more expensive or later units are appended behind it.*

Proof. See Appendix A.1.

Both customer types buy units in order until the next unit is no longer worth its price plus marginal cost. Because the costs are nonnegative, an informed purchase is always a loss to the dealer who sells it. Part (iv) says that the value of a rank does not depend on the depth behind it. Our running example uses the costs


$$\begin{aligned} C(z) &= z^2, & K(q) &= 14q^2 + q/2, \\ C(z + q) - C(z) + K(q) &= 2zq + 15q^2 + q/2. \end{aligned} \quad (3)$$

In the example these rising costs make the insider stop after two units at a price of 30, while large liquidity needs sweep the queue. The small linear term $q/2$ breaks quantity ties on the grid of five-unit ticks; it does not generate the reversal of rank values, which holds without it, and it is convenient for the two-round extension in Section 5.

2.4 Timing

Figure 1 summarizes the timing. Dealers first form the book in sequence, and then the clocks start. Revision opportunities, each dealer's at rate v_i , and the customer opportunity, at rate λ , arrive in random order. At τ the customer buys a prefix, the value is revealed and payoffs are paid. Revisions are the only events that change ranks before τ . A dealer moves back when it resets, and moves forward when a rival ahead of it leaves its place, for example by resetting or withdrawing.

A. Initial formation



In a public order, each dealer abstains, quotes one unit at a grid price p (fee f), or sells at the outside bid B . Clocks start once formation ends.

B. Continuous-time revision and the customer




Figure 1: Timing. Panel A: dealers form the book one at a time. Panel B: after formation, each dealer i 's revision opportunities arrive at rate ν_i and the single customer opportunity arrives at rate λ , all independently, so they occur in random order. No dealer moves between its own opportunities, so the customer can arrive between any two revisions. The marks show one possible sequence of events.

2.5 Equilibrium

A dealer's strategy gives a (possibly random) action at its formation turn and at every revision opportunity. The action may depend on the whole public history of quotes and actions, not only on the current rank. The customer chooses a quantity at every feasible book and private type. Hidden orders, cancellation after the customer arrives, larger displayed quantities and dealer purchases are not part of the game.

No dealer has private information, and dealer actions reveal nothing about v , so dealers keep the prior until τ . We can therefore replace the customer by the expected terminal surplus that optimal demand gives dealer i at book b ; call this terminal reward $R_i(b)$. The result is a public stochastic game in which the customer's arrival ends play. Let $J_i(b)$ be dealer i 's continuation value at book b under a strategy profile σ , and let $T(b, \sigma_j)$ be the book after dealer j takes its action. Then

$$0 = \lambda [R_i(b) - J_i(b)] - h_i \mathbf{1}_{\{i \text{ active}\}} + \sum_j \nu_j \left[J_i(T(b, \sigma_j)) - J_i(b) - c_j \mathbf{1}_{\{j=i, \sigma_j \text{ submits}\}} \right], \quad (4)$$

where a market sale replaces the seller's continuation by $B - m$.

At rate λ the customer arrives and dealer i receives its terminal reward instead of its continuation. While its quote is live it pays waiting cost h_i . At rate ν_j dealer j moves and the book changes, and dealer i pays a fee only if it is the

one who submits. At its own opportunity, each dealer picks the action that maximizes *its own* continuation less its fee; history-dependent strategies satisfy the same equation, indexed by history.

Definition 1 (Equilibrium). An *equilibrium* is a subgame-perfect equilibrium of the public stochastic game with terminal rewards $R_i(b)$, combined with optimal customer demand and the prior before τ . Equivalently, it is a public-strategy Bayesian equilibrium that is sequentially optimal at every information set.

The equilibria we construct are stationary on an invariant class of books, where actions depend only on the current book; off that class, strategies may depend on history. The next lemma shows that such continuations always exist.

Lemma 2 (Equilibrium continuations exist). *Suppose there are finitely many dealers, each with one unit, a finite price grid, bounded waiting and action costs, and $\lambda > 0$. Then:*

- (i) *the public game has a pure, history-dependent subgame-perfect equilibrium;*
- (ii) *the same holds outside any given invariant class of books, for the game that ends when play first reaches the class and pays bounded continuation values there.*

Proof. See Appendix A.2.

Lemma 2 lets us specify behavior on the path of interest and complete it with optimal play everywhere else. It works because at every event the customer arrives, and play ends, with positive probability, so distant future payoffs matter little.

2.6 Relation to Parlour and Roşu

The game borrows from two benchmarks. From Roşu (2009) it takes continuous time, independent Poisson events and the one-sided book. From Parlour (1998) it takes sequential, anticipating behavior: each dealer chooses its order knowing how later traders will respond. It departs from both in revision. Parlour's orders cannot be revised, and Roşu's revision is immediate and free. Here each dealer revises at random times, as often as it likes, for a fee. Customers also differ. Value is common and privately known, and the informed trader picks an optimal quantity, whereas in Parlour (1998) traders differ in private tastes and trade one unit, and in Roşu (2009) information is symmetric.

Neither benchmark is obtained as a special case. Setting the informed probability to zero does not produce Parlour's arriving buyers and sellers or its terminal date. Letting ν grow without bound does not by itself yield Roşu's immediate-response game; that would require a convergence result we do not have. Our customers use market orders only, unlike in Roşu (2020), where informed traders choose between market and limit orders. We therefore prove our results directly in this game rather than importing them from either model.

3. Harmful priority with profitable supply

This section derives the paper's first result. We begin with a fixed book and ask what one step forward in the queue is worth; the answer depends only on what is known about the asset when demand stops at a given rank. We then let dealers revise repeatedly and characterize equilibrium resetting at a common price, and give simple bounds that rule out every price and participation deviation. Finally, a four-dealer economy brings the pieces together. In it, rank two fills more often than rank three but earns less, every dealer enters voluntarily, and the dealer at rank three pays to give up a position that is worth more than the back of the queue at today's book.

3.1 The value of one step forward

Fix a full book of N units, all quoted at a common price p , and let D be the number of units the customer buys, as given by Lemma 1. Demand D includes zero, both for customers who buy nothing and for the no-customer type. The *immediate value* of rank k is the dealer's expected surplus if the book stays as it is until the customer arrives:

$$V_k = \mathbb{E}[(p - v)\mathbf{1}_{\{D \geq k\}}], \quad F_k = \mathbb{P}(D \geq k).$$

Here F_k is the probability that demand reaches rank k , the execution probability at that rank. Because the units are identical, relabeling who owns which rank changes neither prices nor demand. The *priority margin* $\Delta_k = V_k - V_{k+1}$ is the value of moving one step forward, from rank $k + 1$ to rank k .

Two equal orders at the same price are filled together or not at all, except in one event: the order at rank k fills and the order at rank $k + 1$ does not exactly when demand stops at k . Priority therefore adds only the fills where demand stops at you. The next proposition turns this observation into an identity and a sign test.

Proposition 1 (Priority margin and its sign). *For adjacent equal units at a common price, under a fixed law of the final book,*

$$\Delta_k = V_k - V_{k+1} = \mathbb{E}[(p - v)\mathbf{1}_{\{D=k\}}]. \quad (5)$$

If $\mathbb{P}(D = k) > 0$, the margin Δ_k is negative exactly when $\mathbb{E}[v \mid D = k] > p$, while the gain in execution probability is $F_k - F_{k+1} = \mathbb{P}(D = k) > 0$.

Proof. See Appendix A.3.

The proposition separates two effects that are often bundled together. A step forward weakly raises the execution probability, by the probability $\mathbb{P}(D = k)$ that demand stops at k . It raises expected profit only if those extra fills are, on average, sales at a price above the asset's expected value given that demand stops there. If demand that stops at rank k tends to come from a buyer who knows the asset is worth more than p , the extra fill is a loss, and being ahead

then has a cost. In the four-dealer economy studied below, high-value informed demand always stops at rank two at the all-30 book, so $\Delta_2 = -0.9$ although $F_2 - F_3 > 0$. Our contribution is the equilibrium analysis built on this identity.

The sign of each margin depends on the law of stopping. The next result asks when information imposes a pattern across ranks. Let $\varphi_k(v) = \mathbb{P}(D = k \mid v)$ be the probability that demand stops at k when the value is v . Split the margin into the part earned when the asset is worth less than the price and the part lost when it is worth more:

$$G_k = \mathbb{E}[(p - v)^+ \varphi_k(v)], \quad L_k = \mathbb{E}[(v - p)^+ \varphi_k(v)], \quad \Delta_k = G_k - L_k.$$

Proposition 2 (Information ordering and shape). *Suppose that for every pair of ranks $i < j$ and almost every $x < p < y$,*

$$\varphi_j(y)\varphi_i(x) \geq \varphi_i(y)\varphi_j(x).$$

Then $L_j G_i \geq L_i G_j$, and a negative nonzero margin cannot precede a positive one. If the inequality holds for almost every $x < y$ (full likelihood-ratio ordering), the conditional means $\mathbb{E}[v \mid D = k]$ are nondecreasing in k over the ranks at which demand stops with positive probability.

Proof. See Appendix A.3.

The condition says that, compared with a stop at an earlier rank, a stop at a later rank is relatively more likely when the asset is worth more, so later stops carry worse news for the dealer. Once a margin turns negative, no later margin is positive: harmful steps forward come after helpful ones, never before. The condition must hold for every pair of ranks, not only adjacent ones, because a rank at which no informative demand stops cannot pass the ordering along. These restrictions apply to the immediate values V at a fixed book. They do not carry over to the equilibrium values w defined below, which also contain each dealer's option to change rank, the moves its rivals make and the fees everyone pays.

The four-dealer example below violates the ordering condition. In it, stops at rank three come only from liquidity traders, while stops at rank two are mostly informed, so $\Delta_3 > 0$ follows $\Delta_2 < 0$. Proposition 2 is the benchmark that disciplines the shape of the margins when the ordering holds: negative margins can then occur only behind positive ones, never in front of them.

3.2 Equilibrium at a common price

A dealer does not hold its rank until the customer arrives. At each of its revision opportunities it can keep its order or reset it to the back, and rivals can do the same. We therefore need the value of a rank when everyone follows a rule for resetting. Suppose all dealers share the revision rate ν , the replacement fee c and the waiting cost h . Consider the policy “reset exactly when your current

rank lies in S ", for a set $S \subseteq \{1, \dots, N-1\}$. Let $s_k = |S \cap \{1, \dots, k-1\}|$ be the number of resetting ranks ahead of rank k .

Under this policy, rank changes are simple. When a rival ahead of you resets, it moves to the back and you move up one rank, from k to $k-1$. A reset by a rival behind you leaves your rank unchanged, and your own reset sends you to rank N . Let w_k be the *value of rank k under the policy*: the expected net surplus of holding rank k when every dealer follows the policy thereafter, including all future fees, waiting costs and rank changes. Once the policy condition (7) below holds, w_k is the equilibrium value of rank k . These values solve

$$0 = \lambda(V_k - w_k) - h + \nu \mathbf{1}_{\{k \in S\}}(w_N - c - w_k) + \nu s_k(w_{k-1} - w_k), \quad (6)$$

where the last term is zero for $k = 1$.

Equation (6) balances the flow of value at rank k . The first term says that the customer arrives at rate λ ; the dealer then receives the immediate value V_k in place of its continuation value w_k . The second term is the waiting cost h that a live quote pays per unit of time. The third term applies only at resetting ranks: at rate ν the dealer's own clock rings, and it moves to the back, worth w_N , after paying the fee c . The fourth term collects the rivals' resets. Each of the s_k resetting ranks ahead of the dealer resets at rate ν , and each such reset moves the dealer from k to $k-1$.

The equations value a policy; they do not yet make it optimal. A dealer resets when the back of the queue, net of the fee, is worth more than staying:

$$k \in S \implies w_N - c > w_k, \quad k \notin S \implies w_N - c \leq w_k. \quad (7)$$

Let Q_S be the matrix of rates at which a dealer's rank changes under the policy (its generator), and let $r_S(k) = \mathbf{1}_{\{k \in S\}}$. Then the rank equations have the closed-form solution

$$w = (\lambda I - Q_S)^{-1}(\lambda V - h\mathbf{1} - c\nu r_S). \quad (8)$$

Equation (8) is the matrix form of the N linear equations (6).

Theorem 1 (Equilibrium resetting at a common price). *Consider a full N -unit book at a common price, with common revision rate ν , replacement fee c and waiting cost h . For any fixed reset set $S \subseteq \{1, \dots, N-1\}$, equation (6) has a unique finite solution, given by (8). If this solution satisfies (7), the stationary policy "reset exactly at ranks in S " is a subgame-perfect equilibrium of the restricted game in which dealers can only retain their orders or reset them at the same price. No history-dependent retain-or-reset strategy does better for the dealer. If the inequalities in (7) are strict, each prescribed action is strictly optimal; retaining a quote and passing count as the same action. The policy is part of an equilibrium of the full game only if entry, price, withdrawal and market-sale deviations are also unprofitable.*

Proof. See Appendix A.4.

The theorem reduces a dealer's strategic problem to N linear equations and N inequalities. The values account for the customer arriving between any two revisions and for every rival's later resets; no dealer moves last, and none has a single chance to act. The waiting cost enters very simply: by (8), a cost h lowers every rank value by h/λ . All dealers wait for the same customer, so waiting does not change any comparison in (7) and does not change reset choices at the common-price book. The last sentence of the theorem matters, because a dealer could also change its price, withdraw, sell at the outside bid or stay out. The next subsection deals with these deviations.

3.3 Prices and entry

To rule out deviations from the common price, we use two uniform bounds on what a dealer can earn at the customer's arrival. Suppose that at every final book, any dealer's expected gross surplus is at most $\bar{M} \geq 0$. Suppose also that any final quote at a price $p \neq p^*$ earns expected gross surplus at most $\bar{K}$, with $0 \leq \bar{K} \leq \bar{M}$. The bound $\bar{M}$ also covers the market sale, $B - m \leq \bar{M}$, and the zero payoff of staying out. Both bounds apply to the final book only, and a dealer that starts at the wrong price may repair its quote later, so the bounds alone do not settle the question.

The repair question is a race. Before its next revision opportunity, a dealer cannot fix a bad quote, and its next revision comes before the customer with probability $\alpha = \nu/(\lambda + \nu)$, whatever its rivals do in between. If the customer arrives first, an off-price quote earns at most $\bar{K}$, and an absent dealer earns nothing. If the dealer's revision comes first, no continuation earns more than $\bar{M}$. This gives the bounds

$$U_{\text{off}} = (1 - \alpha)\bar{K} + \alpha\bar{M}, \quad U_{\text{empty}} = \alpha\bar{M}. \quad (9)$$

These bounds place no restriction on how rivals respond: $\bar{M}$ covers every later repair, reset or withdrawal by the dealer and every action of its rivals. Fees and waiting costs are nonnegative and can only lower the payoff further.

Theorem 2 (Prices and entry). *Suppose a rank policy satisfies the conditions of Theorem 1 at a full book in which every dealer quotes p^* , the bounds $\bar{K}$ and $\bar{M}$ hold, and*

$$\underline{w} := \min_k w_k > U_* := \max\{U_{\text{off}}, U_{\text{empty}}, B - m, 0\}, \quad 0 \leq f < \underline{w} - U_*. \quad (10)$$

Then the full game has an equilibrium in which every dealer quotes p^* at the start, all dealers participate, and each then follows the rank policy until the customer arrives. The deviations ruled out include every initial quote, quotes repaired later, staying out and entering later, withdrawing and re-entering, replacement at every price, market sale, and repeated history-dependent deviations. Off the equilibrium path, the strategies can be chosen pure and history-dependent; they need not be stationary Markov strategies.

Proof. See Appendix A.5.

The logic is a comparison of two numbers. Each dealer at the common price earns at least $\underline{w}$, its worst equilibrium rank value. Any deviation earns at most U_* , however it continues, because of the race. A deviation leaves the common-price book, and Lemma 2 supplies equilibrium play from there until the book returns to it. Condition (10) then makes every deviation unprofitable and leaves room for a display fee f . The fee condition is sufficient, not necessary: it compares entry net of the fee with bounds that do not charge the deviating dealer any initial fee.

3.4 The four-dealer economy

We now apply these results to a specific economy with $N = 4$ dealers. The asset is worth $H = 100$ with probability $\mu = 1/5$ and 0 otherwise, so its mean is $m = 20$. Prices lie on the grid $\mathcal{P} = \{0, 5, \dots, 105\}$, the outside bid is $B = 0$, and the target quote is $p^* = 30$. Informed costs follow (3). The customer types and their probabilities are:

		Liquidity, (d, u)				
Customer type	Informed	(1, 13)	(2, 13)	(3, 13)	(4, 13)	(4, 18)
Probability	1/14	1/10	1/100	1/100	3/10	1/10

With the remaining probability $b_0 = 143/350$, no customer arrives. Liquidity types have total probability $13/25$, and their reservation values $m + u$ are 33 and 38.

At 30, a high-value insider buys two units and a low-value insider buys none. For the high-value insider, the utilities of buying $q = 0, \dots, 4$ units are 0, $109/2$, 79, $147/2$ and 38, so two units is a strict optimum. High-value informed demand therefore always stops at rank two at the all-30 book. Each of the first two ranks loses $a\mu(H - p^*) = (1/14)(1/5)(70) = 1$ to the insider. Liquidity demand reaches the four ranks with probabilities (0.52, 0.42, 0.41, 0.40), and each liquidity fill earns $30 - 20 = 10$ in expectation. Hence

$$V = (4.2, 3.2, 4.1, 4), \quad \Delta = (1, -0.9, 0.1).$$

The second margin is negative because demand that stops at rank two is mostly informed. The type table is built for this: the high-value insider stops at two with probability $a\mu = 1/70$, while liquidity demand for exactly two units has probability only $1/100$. No informed demand stops at rank three, which is why $\Delta_3 > 0$. These values come directly from the primitives, not from assumed rank payoffs.

The price bounds are simple. At any rival book, a final quote at 20 or below earns nonpositive expected surplus. A quote at 25 earns at most $5(13/25) = 13/5$ before informed losses. A quote at 35 reaches only urgency-18 liquidity buyers and earns at most $15/10 = 3/2$. Quotes at 40 or above receive no liquidity demand and only weakly loss-making informed fills. Hence $\bar{K} = 13/5$. At

30, the gross bound is $10(13/25) = 26/5$, which also bounds every other action, so $\bar{M} = 26/5$. The next theorem puts these pieces together. It shows that the four-dealer economy has an equilibrium with harmful priority, profitable entry and strategic priority surrender: the strategic resetting in its title is the reset at rank three, out of a place worth more than the back at today's book.

Theorem 3 (Harmful priority, profitable supply and strategic resetting). *In this economy, with $\lambda = 1$, $\nu = 1/4$, $c = 1/100$, $h = 1/100$ and $f = 1/10$, the full game has an equilibrium in which all four dealers quote 30 at the start. At every later opportunity of its own, a dealer at rank two or three resets to the back at the same price, and a dealer at rank one or four retains. The immediate values, margins and equilibrium values before the display fee are*

$$V = (4.2, 3.2, 4.1, 4), \quad \Delta = (1, -0.9, 0.1), \quad w = (4.19, 3.345, 3.945, 3.975). \quad (11)$$

Every participating dealer earns strictly positive net surplus after the display fee. Rank two has a higher execution probability than rank three, both immediately and in equilibrium, but a lower immediate value and a lower equilibrium value. The dealer at rank three strictly prefers to reset although its immediate gain from moving to the back is negative. The conclusions hold for all primitives close to these values, holding fixed the finite set of customer types and identical dealers.

Proof. See Appendix A.6.

The reset set is $S = \{2, 3\}$. The values in (11) use $h = 0.01$. With $h = 0$ the equilibrium values are $w = (21/5, 671/200, 791/200, 797/200) = (4.2, 3.355, 3.955, 3.985)$, each 0.01 higher. From here on, equilibrium values in the text are quoted at $h = 0$; a waiting cost h lowers each of them by h/λ . Figure 2 plots the immediate values, the equilibrium values at $h = 0$, and the margins. With either waiting cost, the gains from resetting are

$$w_4 - c - w_k = (-9/40, 31/50, 1/50, -1/100), \quad k = 1, \dots, 4. \quad (12)$$

They are positive exactly at ranks two and three, as (7) requires. For prices and entry, $\alpha = 1/5$, so $U_{\text{off}} = 78/25 = 3.12$ and $U_{\text{empty}} = 26/25 = 1.04$, while the market sale yields $B - m = -20$. Hence $U_* = 3.12$, and with $h = 0.01$, $\underline{w} - f = 3.345 - 0.1 = 3.245 > 3.12$, so condition (10) holds.

Execution probabilities show that harmful priority is not a statement about fills. Holding the book fixed, the fill probabilities by rank are

$$(187/350, 76/175, 41/100, 2/5) \approx (0.534, 0.434, 0.410, 0.400).$$

Following the equilibrium strategy from each initial rank, they are

$$\begin{aligned} & (187/350, 1259/2940, 2019/4900, 5939/14700) \\ & \approx (0.534, 0.428, 0.412, 0.404). \end{aligned} \quad (13)$$




Figure 2: Values by queue rank in the four-dealer economy, with waiting cost $h = 0$. Panel A plots the immediate values $V = (4.2, 3.2, 4.1, 4)$, which hold the full book fixed until the customer arrives, and the equilibrium values $w = (4.2, 3.355, 3.955, 3.985)$. Each w_k is the value of holding rank k and following the equilibrium thereafter; it includes all future replacement fees and all moves caused by rivals' resets. A positive waiting cost h lowers every w_k by h/λ . Panel B plots the priority margins $\Delta = (1, -0.9, 0.1)$. Rank two fills more often than rank three (fill probabilities are in (13)) but is worth less under both measures.

The first vector fixes the final book; the second follows the whole equilibrium strategy, including rank changes from resets. In both, rank two fills more often than rank three, about 0.43 against 0.41. Yet rank two is worth less: 3.2 against 4.1 immediately, and 3.355 against 3.955 in equilibrium. The extra fills that rank two receives are mostly the insider's purchases. Large liquidity purchases fund every rank, so supply remains profitable.

The grid offers more than one profitable price: a lone dealer earns gross trading profits of $107/70$ at 25, $21/5$ at 30 and $4/7$ at 35. Yet no dealer gains by undercutting. A quote at 25 buys the front of the queue, but it halves the profit on every liquidity fill, from 10 to 5, so a final quote at 25 earns at most $13/5 = 2.6$. Even allowing for a later repair, an off-price start earns at most $U_{\text{off}} = 3.12$, below the at least 3.245 that every rank earns at 30, net of the display fee. A finer tick can change this: Appendix B shows that undercutting to 29 pays in a simultaneous game without revision.

The dealer at rank three resets although, at today's book, the move to the back loses $4 - 4.1 - 0.01 = -0.11$; in equilibrium the same move gains 0.02. Section 4 explains why, and when this is worth its fee.

4. Strategic priority surrender

Theorem 3 shows that dealers pay to give up a place in the queue that is worth more than the back at today's book. This is strategic priority surrender, the second result of the paper. This section explains why it happens and when. The answer turns on how fast dealers can revise relative to the arrival of demand

opportunities, and on the fee for each revision. We first solve the four-dealer rank game in closed form. We then derive the region of speeds and fees in which the dealer at rank three resets, account for what this activity costs dealers as a group, and close with what repeated access implies.

4.1 Why a dealer gives up a valuable position

Start with the mechanism in words. Under the equilibrium policy, the front order never moves: it never resets, and no one ahead of it can. The other three dealers keep cycling through ranks two to four, each resetting whenever it is at rank two or three. The dealer at rank two resets at every opportunity, because high-value informed demand stops exactly there and makes rank two the worst place in the queue. Now consider the dealer at rank three. At today's book, moving to the back looks like a mistake. Rank four is worth 4 and rank three 4.1, and the reset costs a fee of 0.01, so the immediate gain is $4 - 4.1 - 0.01 = -0.11$.

But staying has a cost that today's book does not show. If the dealer stays at rank three, a single reset by the dealer ahead pushes it into rank two. From the back, it must be pushed up twice before it reaches rank two. At rank four it moves up at rate 2ν , since either dealer ahead may reset; once at rank three, it resets again at rate ν . The customer may well arrive first, so the fee buys time away from rank two. Once every rival move and the chance that the customer arrives at any moment are counted, the reset gains 0.02. In equilibrium the back is the best place behind the front: $w_4 > w_3$ (3.985 against 3.955 at $h = 0$).

The rank equation (6) turns this reasoning into numbers. Write $x = \nu/\lambda$ for the revision rate relative to the customer rate. Speed is relative here: at $x = 1/4$ a dealer expects one revision opportunity per four demand opportunities, counting those that end without a trade. Set the waiting cost $h = 0$ for now, and take the candidate policy $S = \{2, 3\}$: ranks two and three reset, ranks one and four retain. Dividing (6) by λ gives four equations,

$$\begin{aligned} w_1 &= V_1, \\ (1+x)w_2 &= V_2 + x(w_4 - c), \\ (1+2x)w_3 &= V_3 + x(w_4 - c) + xw_2, \\ (1+2x)w_4 &= V_4 + 2xw_3. \end{aligned} \tag{14}$$

Each line reads directly. The dealer at rank one never moves, so it keeps its immediate value. The dealer at rank two leaves for the back at its own rate. The dealer at rank three also leaves for the back, and is moved up to rank two when the rival ahead resets. The dealer at rank four is moved up to rank three when either rival ahead resets.

The system has a convenient explicit solution:

$$\begin{aligned} w_3 &= \frac{xV_2 + (1+x)V_3 + xV_4 - cx(1+2x)}{1+3x}, \\ w_4 &= \frac{V_4 + 2xw_3}{1+2x}, \quad w_2 = \frac{V_2 + x(w_4 - c)}{1+x}. \end{aligned} \quad (15)$$

These values describe the candidate policy, not yet an equilibrium. The policy is an equilibrium only if each rank's reset gain $g_k = w_4 - c - w_k$ has the right sign: positive at the ranks that reset and nonpositive at the ranks that retain. The next proposition gives these signs in closed form.

Proposition 3 (Region of strategic priority surrender). *Let $N = 4$, $h = 0$, $x = \nu/\lambda$ and $\Gamma_x = (1+2x)(1+3x)$. Under the candidate policy $S = \{2, 3\}$, the reset gains at ranks two and three are*

$$g_2 = \frac{V_4 - V_2 + x(2V_3 + 2V_4 - 4V_2) - c(1+2x)^2}{\Gamma_x}, \quad (16)$$

$$g_3 = \frac{V_4 - V_3 + x(2V_4 - V_2 - V_3) - c(1+2x)^2}{\Gamma_x}. \quad (17)$$

If $g_2 > 0$, $g_3 > 0$ and $w_4 - c < V_1$, the policy is a strict equilibrium of the same-price retain-or-reset game of Theorem 1. Theorem 2 then adds optimality in the full game. In particular, if $V_3 > V_4$ and $g_3 > 0$, the dealer at rank three resets although the immediate move to the back lowers its value. At the example values $V = (4.2, 3.2, 4.1, 4)$,

$$g_3 > 0 \iff c < \frac{0.7x - 0.1}{(1+2x)^2}, \quad (18)$$

which requires $x > 1/7$. At equality the dealer at rank three retains.

Proof. See Appendix A.7.

Each gain in (16)–(17) has two parts. The terms without x compare the back of the queue with the current rank at the fixed book. The terms in x capture the rivals' resets: how often the dealer would be pushed up if it stayed, and how often it is pushed up from the back if it leaves. The fee enters with weight $(1+2x)^2$. For rank three the immediate part $V_4 - V_3$ is negative in the example, so only the second part can make resetting worthwhile. The dealer at rank four never resets when $c > 0$, and the condition $w_4 - c < V_1$ keeps the dealer at rank one in place.

At the example, $x = 1/4$ and $c = 0.01$. The threshold in (18) equals $0.075/2.25 \approx 0.033$, well above the fee. The waiting cost does not change these comparisons, since h lowers every rank value by the same h/λ ; the gains (12) are the same at $h = 0$ and $h = 0.01$. Figure 3 plots the threshold against x .




Figure 3: Strategic priority surrender at rank three. The solid curve is the right side of (18), the largest fee compatible with the policy $S = \{2, 3\}$ at rank three: below it (shaded), the reset gain g_3 at the example values is positive. The curve is an inverted U, positive only for $x > 1/7$ (dotted line) and falling toward zero as x grows. The dashed line is the example fee $c = 0.01$, and the marked point is the equilibrium of Theorem 3.

Rivals' revisions are essential. If rivals never revise, a dealer at rank three with $V_3 > V_4$ and a positive fee retains, even if it can revise as often as it likes. The reset pays here only because a rival ahead resets at a positive rate.

4.2 Speed, fees and aggregate surplus

Resetting moves exposure from one dealer to another, but it cannot change the total of the terminal surplus claims $\sum_k V_k$. At any instant the four owners hold the four ranks, so summing their values counts each final rank once. At the example, $\sum_k V_k = 31/2$, and

$$\sum_k w_k = \sum_k V_k - \frac{2vc + 4h}{\lambda}. \quad (19)$$

Two ranks reset at rate v each, so fees accrue at total rate $2vc$; all four dealers wait at total rate $4h$; and the customer arrives after an expected time $1/\lambda$. Resetting changes who bears the insider's fill, not the four surplus claims at the end. Fees and waiting therefore lower total dealer surplus.

Speed matters in both directions. Hold the fee $c > 0$ fixed. The right side of (18) tends to zero as $x \rightarrow \infty$, so once revision is fast enough the policy $S = \{2, 3\}$ is no longer an equilibrium. With $c = 0$ the three lower-rank values converge to $(V_2 + V_3 + V_4)/3$ as $x \rightarrow \infty$, and the strict reset gains vanish: the ranks behind the front become alike.

Figure 3 shows the resulting inverted U in the largest fee compatible with the policy. When revision is slow relative to demand, the customer usually arrives before any rival resets, so protection from being moved up rarely matters and

the fee is not worth paying; for $x \leq 1/7$ no fee is low enough. When revision is very fast and costly, dealers churn: each position is reshuffled many times before the customer comes, and the fees eat the gain. Under this policy and at the example values, the dealer at rank three gives up its position only between these extremes. Outside this region $S = \{2, 3\}$ is no longer an equilibrium, and we do not characterize the reset pattern that replaces it.

4.3 Myopic resetting and cycles

One might think that repeated access makes the queue cycle forever. Suppose each dealer evaluates a zero-fee reset at the current book and ignores rivals' responses. Then an ordering in which no dealer gains from resetting requires $V_N \leq \min_k V_k$, so the back must be the worst place; if this fails, every ordering of owners has some dealer who gains by moving to the back.

With $V = (1, 2, 2)$, for example, the front dealer gains one by moving to the back, and successive front owners rotate the queue from (A, B, C) to (B, C, A) to (C, A, B) and back. This is a cycle of provisional payoffs, computed without rivals' responses. The four-dealer example fails the no-cycle condition, since $V_4 = 4$ exceeds $V_2 = 3.2$, so under myopic zero-fee resetting some dealer would gain from moving to the back at every ordering of owners. Yet in the equilibrium of Theorem 3, where dealers anticipate their rivals' delayed responses, the front order never moves and the reset policy is stationary. Repeated access therefore need not destroy strategic priority surrender.

5. Learning and repeated informed trading

The stop at rank two that makes rank two costly is also the event the market can learn from: in the example below, after two fills the ask moves from 30 to 80. So far the game ends with a single customer, so nobody learns from a fill and the insider trades once. Yet a fill is public news, and if buyers who stop after two units are often informed, a later round should open at a higher price. This section adds a second round to the game of Section 2. In one equilibrium, executions move beliefs about value and informed inventory, dealers reprice on that news, and the insider who bought earlier buys again.

5.1 A two-round game

The first round is the complete game of Section 2, with dealers revising at their Poisson clocks until the customer arrives. Every dealer now has two endowed units, and it can display and sell at most one unit in each round. The insider's private state (v, z) persists across rounds, and its inventory z starts at zero. Let Y be the number of units executed in the first round, from zero (no trade) to four (the book clears). The public observes Y but not who the customer was. After the first customer opportunity, the market continues to a second round with probability ρ , drawn independently of everything else; otherwise the asset settles.

If the market continues, the quotes that were not filled keep their prices and timestamps. We call a dealer whose first-round quote is still live an *incumbent*. A dealer whose first unit was filled is empty and can use its second unit. So can a dealer that sold its first unit at the outside bid: it can still use its remaining endowment in the second round, although that first-round sale still ends its first-round moves. A uniformly random order of the four dealers is then drawn and made public. Each dealer gets exactly one move, in that order, and no customer arrives during this sequence. An empty dealer may abstain or post at any grid price for free. An incumbent may retain its quote, withdraw it for free, or replace it at any grid price for the fee c . Any dealer may instead sell its unit to the market at the outside bid B .

A second customer type is then drawn, independently of the first and with the same probabilities. If it is informed, it is the same insider, and it maximizes (1) at its actual inventory z . A liquidity trader values a unit at the updated public mean of v plus the same urgency u as before. The final round has one move per dealer and no customer in between, a simplification that keeps the game finite.

5.2 Beliefs and the forward-looking insider

Dealers learn from the execution count. Fix the first book b . Let $q^0(b, v)$ be the insider's quantity when it maximizes current utility only, and let t denote the first customer type. Then $y(b, v, t)$ is the number of units executed, and $z(b, v, t)$ is the insider's inventory afterwards: $q^0(b, v)$ if the customer was informed, zero otherwise. With prior $\pi_0 = \pi$, the public belief about value and inventory after observing y is

$$\pi_b(v, z | y) = \frac{\pi_0(v) \sum_t \mathbb{P}(t) \mathbf{1}_{\{y(b, v, t)=y, z(b, v, t)=z\}}}{\sum_{v', t} \pi_0(v') \mathbb{P}(t) \mathbf{1}_{\{y(b, v', t)=y\}}}. \quad (20)$$

The numerator counts the ways of producing count y with value v and inventory z , and the denominator counts all ways of producing y . We use (20) whenever the denominator is positive. At a count that is physically possible but has zero probability, we specify some feasible posterior. The insider still trades optimally at its actual inventory, even if that inventory lies outside the public posterior's support. Because dealers do not observe a deviation by the insider, they keep updating as if it had bought its equilibrium quantity.

The insider is forward-looking: buying more today raises its inventory and its marginal cost tomorrow. Two bounds keep this effect small. First, current-optimal quantities are strict. In the first round, where $z = 0$, at every grid book the high-value insider's marginal gain from unit j is $100 - p_j - 15(2j - 1) - 1/2$, and at low value the gain is negative, so every other quantity earns at least $g_I = 1/2$ less than the current optimum. Second, at any inventory $z \geq 0$ the extra utility the insider can obtain in the final round lies in $[0, U_I]$, where

$$U_I = \max_{0 \leq q \leq 4} \{100q - 15q^2 - q/2\} = 327/2.$$

Dealers' continuation payoffs are bounded in the same way. Every dealer has a sure move before the second customer, and it can always withdraw or abstain, so it can guarantee zero. Informed fills never pay a dealer, and each filled liquidity unit pays at most the largest urgency in conditional expectation. Each dealer's optimized second-round payoff therefore lies in $[0, \bar{D}]$, where the $b_{d,u}$ below are the liquidity-type probabilities of Section 2, not books:

$$\bar{D} = u_{\max} \sum_{d,u} b_{d,u} = 18(13/25) = 234/25. \quad (21)$$

Both bounds hold at every book and every feasible posterior, not only at the histories studied below.

5.3 Equilibrium with learning

The next theorem states that the first-round equilibrium of Theorem 3 survives when dealers and the insider look ahead to a short second round, and that the second round shows learning, repricing and repeated informed trading. We use the perfect Bayesian equilibrium (PBE): every player acts optimally at every information set, and beliefs follow (20) whenever it applies.

Theorem 4 (Learning and repeated informed trading). *Take the example economy of Theorem 3, with $h = 0.01$, $f = 0.1$, and continuation probability $\rho = 1/10000$. There is a PBE of the two-round game with the following properties.*

- (i) *All dealers initially quote 30. In the first round, at each of their revision opportunities, the dealers at ranks two and three reset, as in Theorem 3.*
- (ii) *At every possible first book, the insider buys its current-optimal quantity.*
- (iii) *The public updates its joint belief about value and informed inventory by Bayes' rule, and second-round prices, participation and revisions are sequentially optimal at every history.*
- (iv) *With positive probability, the same high-value insider buys two units in the first round and one unit in the second.*
- (v) *After $Y = 2$, every dealer's final quote is 80 in every realized order of second-round moves.*

The first-round policies are optimal against forward-looking quantity deviations by the insider and against all dealer deviations, with the second round's payoffs included.

Proof. See Appendix A.8.

The example is one case of a more general sufficient condition. Let $\delta > 0$ be the smallest strict gap in the first-round retain-or-reset comparisons of

Theorem 3, and let U_* be the price and entry bound in (10). The condition is

$$\rho U_I < g_I, \quad 2\rho\bar{D} < \delta, \quad \min_k w_k - f > U_* + \rho\bar{D}. \quad (22)$$

In the example, $\delta = 1/100$, the gap at rank four in (12). The first two inequalities hold because $\rho U_I = 327/20000 < 1/2$ and $2\rho\bar{D} = 468/250000 < 1/100$. The third holds because $\min_k w_k - f = 3.245$ exceeds $U_* = 3.12$ by more than $\rho\bar{D} = 234/250000$. All three inequalities are strict.

Each part of (22) says that current incentives outweigh the most the future could change them. First, buying a different quantity today can raise the insider's later utility by at most ρU_I , less than the g_I it loses today, so the insider keeps its current-optimal quantity, which is the rule dealers use to update beliefs. Second, adding the second round raises each first-round rank value by between 0 and $\rho\bar{D}$. A retain-or-reset comparison therefore moves by at most $2\rho\bar{D}$, which is too little to reverse any first-round choice. Third, every deviation's bound rises by at most $\rho\bar{D}$, so quoting 30 and entering remain strictly optimal. A small ρ makes all three comparisons go through: the second round is real, but too unlikely to undo first-round incentives.

5.4 Learning in the example

Two first-round histories show how beliefs and prices respond. Recall that the prior probability of high value is $\mu = 1/5$, so the prior mean is 20. At 30 the high-value insider buys two units and the low-value insider buys none. Counts of one, three or four units therefore come only from liquidity traders, so after $Y = 1, 3$ or 4 the posterior equals the prior and no informed inventory has accumulated. Below we drop the subscript b , since the first book is the all-30 book.

No fill. A first round without trade occurs when no customer arrives or when the informed customer has low value. Then

$$\begin{aligned} \mathbb{P}(Y = 0) &= 163/350, & m_0 &= 2860/163, \\ \pi(100, 0 \mid 0) &= 143/815, & \pi(0, 0 \mid 0) &= 672/815. \end{aligned}$$

No informed inventory has accumulated. The posterior mean $m_0 \approx 17.5$ is slightly below the prior mean of 20, because a high-value insider would have bought, but the shift is too small to move prices, and all dealers finish the second round at 30. Rank values at that book still differ. In the final all-30 book, the terminal value of final rank two is about 4.35, less than final ranks three and four (about 5.11 and 4.98). Table 1 reports different numbers for $Y = 0$, because it averages over move orders and indexes dealers by their first-round rank. In some move orders a dealer therefore resets at 30 to move back in the queue, so relative timestamps can change.

Two fills. A first round with two executions comes either from a liquidity trader who wanted two units or from the high-value insider. The posterior is

$$\begin{aligned} \mathbb{P}(Y = 2) &= 17/700, \\ \pi(0, 0 \mid 2) &= 28/85, \quad \pi(100, 0 \mid 2) = 7/85, \quad \pi(100, 2 \mid 2) = 10/17. \end{aligned} \tag{23}$$

The buyer was informed with probability $10/17$, and in that case the insider now holds two units. The probability of high value rises to $57/85$, and the posterior mean to $m_2 = 1140/17 \approx 67.1$. Liquidity reservation values become $1361/17 \approx 80.1$ and $1446/17 \approx 85.1$. A quote at 80 therefore serves both liquidity types, while 85 serves only the more urgent type.

At 80 the high-value insider buys exactly one unit, whether its inventory is zero or two. With $z = 2$, its first marginal gain is $1/2 > 0$ and its second is $-59/2 < 0$; with $z = 0$ it also buys one, and at 85 it buys none in either case. The low-value insider buys nothing. At an all-80 book the gross rank values are $(202/35, 462/85, 451/85, 88/17) \approx (5.77, 5.44, 5.31, 5.18)$. Any final quote other than 80 earns at most $351/85 \approx 4.13$; for example, 70 or 75 earns at most $(75 - m_2)(13/25)$. Since $88/17 - c > 351/85$, every dealer posts 80 when its turn comes, whatever the order. Incumbents hold stale quotes at 30, so they pay c to replace them, while filled dealers post for free. Two fills move the price from 30 to 80, and the insider who bought two units buys one more.

Table 1 reports each dealer's expected second-round payoff $\Pi_k(Y)$. It averages over all 24 public move orders and is indexed by the dealer's rank immediately before the first execution. Because the move order is uniformly random and tie rules treat dealers alike, a dealer's expected payoff depends only on its rank before the first execution. After $Y = 2$, ranks one and two were filled and post at 80 for free; ranks three and four are incumbents and pay c to replace. Each group's payoff is the average all-80 rank value, less c for the incumbents.

Table 1: Expected second-round payoffs by first-round rank immediately before the first execution, for each first-round execution count Y

Y	$\Pi_1(Y)$	$\Pi_2(Y)$	$\Pi_3(Y)$	$\Pi_4(Y)$
0	5.599	4.786	4.831	4.814
1	3.750	4.200	3.760	3.780
2	5.422	5.422	5.412	5.412
3	3.767	3.767	3.767	4.200
4	3.875	3.875	3.875	3.875

Note. Entries are values of the finite second-round game, rounded to three decimals. Exact values are in Table 2 of Appendix A.

The table enters the first round directly. Let $\bar{\Pi}_k = \sum_y \mathbb{P}(Y = y) \Pi_k(y)$ be the expected second-round payoff of the dealer at rank k . The terminal reward of

the first-round dealer at rank k becomes

$$\tilde{V}_k = V_k + \rho \bar{\Pi}_k,$$

and solving (6) with $\tilde{V}$ in place of V gives the dealers' total first-round values. These differ from w by at most $\rho \bar{D}$, so resetting at ranks two and three and quoting 30 remain optimal.

The repeat purchase in part (iv) requires high value, an informed first customer, continuation and an informed second customer. Its probability is $\mu a^2 \rho = 1/9800000$. The probability is tiny because ρ is chosen small; what matters is that repeated informed trading is part of equilibrium play.

6. Empirical implications

We first derive how fills and markouts by queue rank respond to the customer mix, and how often dealers cancel and resubmit as information, fees and speed change. We then describe what separates the mechanism from ordinary cancellation after news, and when a markout measures a priority margin.

6.1 The value of a queue position at a given book

Let the asset value be either 0 or H , with $\mu = \mathbb{P}(v = H)$. By Lemma 1, only the high-value insider buys. Write T_k^L for the probability that the customer is a liquidity trader whose demand reaches rank k ; this is unconditional, so it includes the probability that a liquidity trader arrives. Write $T_k^I \in \{0, 1\}$ for whether the high-value insider's demand reaches rank k . At a common-price book with price p ,

$$V_k = (p - \mu H) T_k^L - a \mu (H - p) T_k^I. \quad (24)$$

The first term is the profit from liquidity trades, which bring no news about value; the second is the loss to the insider, who buys only when the asset is worth H . In the example, $p - \mu H = 10$, $a \mu (H - p) = 1$, $T^L = (0.52, 0.42, 0.41, 0.40)$ and $T^I = (1, 1, 0, 0)$, which gives $V = (4.2, 3.2, 4.1, 4)$. In data, rank-level fill rates and markouts (see Proposition 4 below) are the observable counterparts of these reach probabilities and values. Holding the book, the prior and the other customer types fixed, equation (24) yields three predictions about rank-level profits.

Prediction 1. *More informed trading lowers profits at the ranks the insider reaches.* Raise the informed probability a by taking mass from the no-customer outcome, keeping liquidity flow and the insider's chosen quantity unchanged. Then each V_k falls at rate $\mu(H - p)T_k^I$. In the example, only ranks one and two lose; ranks three and four are unaffected.

Prediction 2. *Larger informed orders move the loss, not just enlarge it.* A larger informed order weakly lowers every V_k , but it has no uni-

form effect on the margins Δ_k . If the insider buys three units instead of two, the adverse stop moves from rank two to rank three: Δ_2 rises and Δ_3 falls. A larger insider inventory raises the insider's marginal costs, so it can move the adverse stop toward the front of the queue.

Prediction 3. *Sweeps fund every rank; exact-size demand shapes the margins.* An increase in liquidity demand that adds profitable executions deep in the queue raises the value of every rank it reaches. The mass of liquidity orders of exactly size k affects the individual margin Δ_k .

While no customer changes quantity, $\partial V_k / \partial p = F_k$ and $\partial \Delta_k / \partial p = \mathbb{P}(D = k)$. Crossing a quantity or reservation threshold changes discretely who stops at each rank, so these derivatives do not predict equilibrium spreads.

6.2 Equilibrium responses

When dealers revise, the relevant values are the equilibrium values w . For a fixed reset set S , they solve the linear system (8), $w = (\lambda I - Q_S)^{-1}(\lambda V - h\mathbf{1} - c\nu r_S)$. Differentiating it gives local predictions, valid while the reset set and the price stay the same, stated at the example of Theorem 3.

Prediction 4. *More informed trading makes rank three more willing to give up its place.* A higher a lowers V_2 , the value of the position into which rank three would be moved up. Since V_2 depends on a through (24), a higher a raises the numerator of the rank-three reset gain g_3 in equation (17). At $x = 1/4$ and $c = 0.01$, that numerator is $3.5a - 0.1975$, so rank three resets when $a > 79/1400$, provided the other equilibrium inequalities still hold.

Prediction 5. *A higher replacement fee discourages resets.* For fixed x and V , a common increase in c lowers g_3 at rate $(1+2x)/(1+3x)$, which is $6/7$ at the example. A change in one dealer's fee alone breaks the symmetry across ranks and requires the system with labeled dealers.

Prediction 6. *Anticipatory resets need intermediate revision speed:* in the example, the rank-three reset under $S = \{2, 3\}$ requires $x = \nu/\lambda > 1/7$ (equation (18)), and with a positive fee that policy fails once revision is very fast (Section 4).

Prediction 7. *Cancellation activity is proportional to revision speed:* in the equilibrium of Theorem 3, the two middle ranks reset at total rate 2ν , the expected number of resets before the customer arrives is $2\nu/\lambda$, and prices and depth stay constant near the example.

Scope of the predictions. All of these are local. When a primitive crosses a policy boundary, the response can be a discrete jump. A waiting cost lowers every w_k by h/λ but leaves reset choices unchanged while the same policy remains in force (Theorem 1); higher entry or waiting costs eventually violate the participation condition of Theorem 2. Adding dealers changes the rank transitions, the possible stopping ranks and the entry incentives. The rank equation (6) applies for any finite number of dealers, and four is the depth of the example.

6.3 Distinguishing the mechanism

Dealers cancel for many reasons, such as a public revision in value. The mechanism of this paper predicts five facts together:

- (i) dealers earn positive total value from supplying liquidity;
- (ii) the fills in which demand stops at some interior rank lose money on average (a negative priority margin);
- (iii) a dealer cancels and resubmits at the same price with no new public signal about value;
- (iv) at the book it faces, the dealer's immediate gain from moving to the back is negative;
- (v) a rival ahead of it resets at a positive rate that can be identified, for example from its past cancellations and resubmissions at this price, and such a reset would move the dealer up into a worse position.

Ordinary cancellation after a public value revision does not need this combination. No single cancellation establishes the mechanism; the test concerns the pattern.

Testing it requires order-level data that link orders to owners and record cancellation and replacement timestamps, fills and markouts. Rank-level fill rates and markouts must share a common denominator of demand opportunities, and periods with no trade belong in it, because the model's λ counts them. Since such opportunities are not directly observed, the ratio x is hard to measure. The test also needs an estimate of the rate at which rivals can respond. A dealer's actual queue rank cannot be recovered from aggregate depth at a price alone. A trade that clears the whole price level shows only that demand reached its last unit. Such trades identify tail probabilities, not the probabilities $\mathbb{P}(D = k)$ that demand stops at an interior rank. Linking a given dealer's fills and markouts to its rank requires order-level data with queue positions.

6.4 Measuring margins with markouts

The margin $\Delta_k = \mathbb{E}[(p-v)\mathbf{1}_{\{D=k\}}]$ uses the terminal value v , which a researcher does not observe. Empirical work uses a markout instead, comparing the price with a benchmark m_s observed at a horizon s after the trade. The next proposition states when this measures the margin.

Proposition 4 (Markout benchmark). (i) Let $\mathcal{G}$ be an information set and Z a bounded allocation, such as $\mathbf{1}_{\{D=k\}}$, that is known given $\mathcal{G}$. If $p = \mathbb{E}[v \mid \mathcal{G}]$, then $\mathbb{E}[(p - v)Z \mid \mathcal{H}] = 0$ for every coarser information set $\mathcal{H} \subseteq \mathcal{G}$.

(ii) Let m_s be an integrable benchmark observed at horizon s . The measured margin satisfies

$$\mathbb{E}[(p - m_s)\mathbf{1}_{\{D=k\}}] = \Delta_k + \mathbb{E}[(v - m_s)\mathbf{1}_{\{D=k\}}]. \quad (25)$$

The second term is zero if $m_s = \mathbb{E}[v \mid \mathcal{G}_s]$ for an information set $\mathcal{G}_s$ that contains the event $\{D = k\}$. Otherwise its sign is unrestricted.

Proof. See Appendix A.3.

Part (i) explains why a negative margin is possible at all. A price equal to the efficient price given information that includes the fill would earn nothing on that fill on average. But a resting quote is set before the customer arrives and cannot depend on whether demand stops at the dealer's rank, so the fill itself carries information. Both parts are identities about conditional means, not results about equilibrium pricing. By part (ii), a markout measures the theoretical margin when its benchmark is the expected value given information that includes the stopping event. A benchmark taken before prices have absorbed the information in the order flow adds a wedge of unknown sign. In practice this argues for benchmarks taken late enough that the benchmark price is an efficient price given information that includes the trade; very short markouts can misstate the sign of the priority margin.

In the example, demand stops at rank two exactly when two units trade, with probability $17/700$. Take as benchmark m_s the posterior mean after two fills from Section 5, $1140/17$. It conditions on the stopping event, so the wedge in equation (25) is zero, and the measured margin $(30 - 1140/17) \cdot 17/700 = -0.9$ equals Δ_2 .

7. Conclusion

Price-time priority allocates fills. The extra fills that a step forward in the queue delivers are those in which demand stops at the dealer's rank. If such stops tend to come from informed traders, being ahead in the queue costs money, even when every order in the queue is profitable. When dealers can revise at repeated, random times, this has a strategic consequence. A dealer pays to give up a place worth more than the back at today's book, to avoid being moved up into a worse position after a rival ahead resets. In our example this requires intermediate revision speed: slow revision makes the protection rarely useful, and with very fast revision and a positive fee, the fees outweigh the gain. In a two-round extension, fills teach dealers about value and informed inventory, dealers reprice, and the same informed trader buys again.

These results come from a deliberately simple market, with a coarse grid on which a single price is selected, a finite number of unit dealers on one side

of the book, market orders only, and a second round that occurs with small probability; extending them to fine ticks, endogenous depth and an ongoing stream of informed trades is the natural next step.

The broader point concerns what queue data should measure. The value of priority is not the value of the average fill at a price. It is the value of the fills that priority adds, net of what those fills reveal, and it depends on what rivals will do next. Rank-level markouts, cancellations at an unchanged price without news, and revision speed relative to the arrival of demand opportunities are the observables that bear on it.

References

- Baldauf, M., and J. Mollner (2020). High-frequency trading and market performance. *Journal of Finance*, 75(3), 1495–1526. doi:10.1111/jofi.12882.
- Foucault, T., O. Kadan, and E. Kandel (2005). Limit order book as a market for liquidity. *Review of Financial Studies*, 18(4), 1171–1217. doi:10.1093/rfs/hhi029.
- Garriott, C., V. van Kervel, and M. A. Zoican (2025). Queuing and inventories in limit order markets. *Journal of Financial Markets*, 75, 100982. doi:10.1016/j.finmar.2025.100982.
- Glosten, L. R. (1994). Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4), 1127–1161. doi:10.1111/j.1540-6261.1994.tb02450.x.
- Glosten, L. R., and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100. doi:10.1016/0304-405X(85)90044-3.
- Goettler, R. L., C. A. Parlour, and U. Rajan (2005). Equilibrium in a dynamic limit order market. *Journal of Finance*, 60(5), 2149–2192. doi:10.1111/j.1540-6261.2005.00795.x.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335. doi:10.2307/1913210.
- Menkveld, A. J., and M. A. Zoican (2017). Need for speed? Exchange latency and liquidity. *Review of Financial Studies*, 30(4), 1188–1228. doi:10.1093/rfs/hhx006.
- Moallemi, C. C., and K. Yuan (2017). A model for queue position valuation in a limit order book. Working paper, Columbia Business School, revised June 2017. doi:10.2139/ssrn.2996221.

- Parlour, C. A. (1998). Price dynamics in limit order markets. *Review of Financial Studies*, 11(4), 789–816. doi:10.1093/rfs/11.4.789.
- Roşu, I. (2009). A dynamic model of the limit order book. *Review of Financial Studies*, 22(11), 4601–4641. doi:10.1093/rfs/hhp011.
- Roşu, I. (2020). Liquidity and information in limit order markets. *Journal of Financial and Quantitative Analysis*, 55(6), 1792–1839. doi:10.1017/S0022109019000759.
- Yueshen, B. Z. (2026). Queuing uncertainty of limit orders. *Management Science*, 72(6), 4760–4779. doi:10.1287/mnsc.2023.03371.

A. Proofs

The proofs use the game, notation and tie conventions of the main text. Payoff bounds on final quotes apply to a quote that stays in the book until the customer arrives; quotes that are repaired later and private quantity deviations are handled by separate arguments below.

A.1 Proof of Lemma 1

Proof. The argument compares successive utility increments. Write the insider's increment from unit j as $v - p_j - \kappa_j(z)$; since prices and marginal costs are nondecreasing in j , these increments are nonincreasing. Utility strictly increases while the increments are positive, and it cannot increase after the first nonpositive increment. Selecting the smallest maximizer therefore chooses exactly the positive increments, which gives part (i).

The liquidity calculation uses the increments $m + u - p_j$ up to the capacity bound $\min(d, n)$ and the same argument, which gives part (ii). Since $\kappa_j \geq 0$, every unit the insider selects has $v > p_j$, which is part (iii). The condition for demand to reach rank k involves only that unit's own threshold and the earlier, weakly lower, thresholds, so units appended later cannot alter it, which is part (iv). $\square$

A.2 Proof of Lemma 2

Proof. We build the equilibrium from finite truncations of an embedded discrete game and pass to a limit along a diagonal subsequence. The argument does not need the Bellman iteration of this general-sum game to be a contraction.

Embedded game. Superpose all revision clocks, keeping null events for dealers who have exited. At any event, the customer arrives and play ends with probability $\eta = \lambda / (\lambda + \sum_i v_i) > 0$; otherwise exactly one dealer receives a revision opportunity. The time to the event is an independent exponential. Waiting costs are linear in time and past waiting is sunk. We may therefore replace each waiting cost between events by its conditional expectation $h_i \mathbf{1}_{\{i \text{ active}\}} / (\lambda + \sum_i v_i)$. The realized duration conveys no private information.

Truncation. The embedded game has a finite set of actions and event identities, finitely many public physical states, and countably many finite event histories. Allow strategies that depend arbitrarily on the event history. Truncate after ℓ events at which play does not end. At the cutoff, assign zero continuation outside the prescribed boundary, and keep the bounded boundary payoffs when play reaches it. Finite backward induction in this perfect-information game, with a fixed deterministic tie rule, gives a pure subgame-perfect profile. At the initial formation stage, which is finite, ordinary backward induction is also available once continuations are assigned.

Uniform control of the tail. Let $\bar{R}$ bound the absolute terminal and boundary payoffs, and let $\bar{c}$ bound the absolute expected cost per event. Conditional

on surviving ℓ events, the expected magnitude of the remaining payoff is at most $\bar{R} + \bar{c}/\eta$. The unconditional discrepancy in the tail of any strategy profile, including any profile with a deviation, is therefore at most

$$(1 - \eta)^\ell (\bar{R} + \bar{c}/\eta).$$

The same bound applies from every finite history, after reindexing its remaining horizon. Cutoff effects therefore vanish uniformly, not only along an equilibrium path.

Limit. Enumerate all finite event histories. Take a diagonal subsequence of the truncated pure profiles that stabilizes at the first history, then at the second, and so on; finiteness of each action set ensures that such a subsequence exists. Its limiting pure profile is optimal against every unilateral deviation of finite horizon at every history. Otherwise a strictly profitable finite deviation, with an error tolerance below its gain, would contradict the corresponding sufficiently long truncated equilibrium. Any deviation can be truncated, and its tail bounded by the same geometric bound, so no deviation is profitable either. This proves subgame perfection.

Continuous time. The constructed strategies ignore exact elapsed durations. They remain optimal against strategies that use those durations. By memorylessness, at any realized continuous-time history the remaining continuation problem is the same embedded problem. Randomizing actions on durations cannot exceed a conditional continuation payoff that is already maximized, and there is no additional privately informative clock. This justifies working with countably many event histories. The argument for part (ii), with bounded prescribed boundary payoffs, is identical. $\square$

A.3 Proof of Propositions 1, 2 and 4

Proof. All three results follow from indicator identities and conditional averaging; none uses equilibrium play.

Proposition 1. Pathwise, $\mathbf{1}_{\{D \geq k\}} - \mathbf{1}_{\{D \geq k+1\}} = \mathbf{1}_{\{D=k\}}$. Multiply by $p - v$ and take expectations to obtain (5). If $\mathbb{P}(D = k) = 0$, the margin is zero and no conditional mean needs to be defined. If $\mathbb{P}(D = k) > 0$, write $\Delta_k = \mathbb{P}(D = k) (p - \mathbb{E}[v \mid D = k])$; this gives the sign criterion. Taking expectations of the same indicator identity gives $F_k - F_{k+1} = \mathbb{P}(D = k)$. Summing the identity over adjacent ranks gives the corresponding identity for a block of ranks, $V_k - V_{l+1} = \mathbb{E}[(p - v)\mathbf{1}_{\{k \leq D \leq l\}}]$ for $k \leq l$.

Proposition 2, weighted ordering. Integrate the assumed kernel inequality over $x < p < y$ against $(p - x)(y - p) d\pi(x) d\pi(y)$. The left and right integrals factor as $L_j G_i$ and $L_i G_j$, which gives $L_j G_i \geq L_i G_j$. Now suppose the margin at rank i is negative and nonzero, so $L_i > G_i \geq 0$, and let $j > i$. If $G_i > 0$ and $G_j > 0$, division gives $L_j/G_j \geq L_i/G_i > 1$, so $\Delta_j < 0$. If $G_i = 0 < L_i$, the undivided inequality forces $G_j = 0$. If $G_j = 0$ directly, the margin at j is nonpositive. These cases exclude a later positive margin. If some rank

between i and j has zero stopping probability, an adjacent-pair condition would say nothing across it; hence the condition is imposed on all pairs.

Proposition 2, full likelihood-ratio ordering. Let $i < j$ be ranks with $\mathbb{P}(D = i) > 0$ and $\mathbb{P}(D = j) > 0$. Symmetrizing the double integral gives

$$\begin{aligned} & \mathbb{P}(D = i)\mathbb{P}(D = j)(\mathbb{E}[v \mid D = j] - \mathbb{E}[v \mid D = i]) \\ &= \int_{x < y} (y - x) [\varphi_j(y)\varphi_i(x) - \varphi_i(y)\varphi_j(x)] d\pi(x) d\pi(y) \geq 0. \end{aligned}$$

The value is bounded, so the integrals exist.

Proposition 4. For part (i), let $p = \mathbb{E}[v \mid \mathcal{G}]$ and let the bounded allocation Z be $\mathcal{G}$ -measurable. For $\mathcal{H} \subseteq \mathcal{G}$, condition first on $\mathcal{G}$:

$$\mathbb{E}[(p - v)Z \mid \mathcal{H}] = \mathbb{E}[Z \mathbb{E}[p - v \mid \mathcal{G}] \mid \mathcal{H}] = 0.$$

For part (ii), add and subtract v : $(p - m_s)\mathbf{1}_{\{D=k\}} = (p - v)\mathbf{1}_{\{D=k\}} + (v - m_s)\mathbf{1}_{\{D=k\}}$. Taking expectations and using (5) gives (25). If $m_s = \mathbb{E}[v \mid \mathcal{G}_s]$ and $\{D = k\} \in \mathcal{G}_s$, part (i) with p replaced by m_s , $Z = \mathbf{1}_{\{D=k\}}$ and the trivial information set $\mathcal{H}$ shows that the wedge $\mathbb{E}[(v - m_s)\mathbf{1}_{\{D=k\}}]$ is zero. Without these measurability conditions the wedge can take either sign. $\square$

A.4 Proof of Theorem 1

Proof. The proof writes the rank values as the resolvent of a killed Markov chain and then applies a dynamic-programming argument to deviations.

Existence and uniqueness. Under a fixed reset policy, the rank of a given dealer follows a finite-state continuous-time chain with generator Q_S . Killing the chain at the independent rate $\lambda > 0$ gives the resolvent

$$(\lambda I - Q_S)^{-1} = \int_0^\infty e^{-\lambda t} e^{tQ_S} dt.$$

The matrix integral exists, is nonnegative, and has row sums $1/\lambda$. Applying it to the vector of flow and terminal rewards gives the unique finite value (8), which satisfies (6). Uniqueness also follows from the maximum principle. A nonzero solution of the homogeneous equation would attain a positive maximum or a negative minimum, which is incompatible with the killing term.

Optimality. At a revision opportunity, retaining leaves the dealer at value w_k , and resetting places it at the back at value $w_N - c$. All other dealers use the same current-rank policy after either action. Condition (7) therefore makes the prescribed action optimal at each rank. For deviations of any form, apply the dynamic-programming inequality on successive intervals between events under an arbitrary deviating strategy, and iterate up to a finite cutoff in the number of events. The candidate value at the start bounds the expected accumulated rewards of the deviation plus the candidate value at the cutoff. The geometric tail bound in the proof of Lemma 2 removes the cutoff. Equality holds under

the prescribed policy, so no history-dependent or mixed deviation improves on it. All permutations of owners share the same transition law for a given dealer's rank, which proves subgame optimality throughout the class.

Waiting cost. Subtracting h/λ from every value leaves all reset differences and transition differences unchanged, and subtracts h in the killing term of the reward equation. This establishes the shift by h/λ . $\square$

A.5 Proof of Theorem 2

Proof. The proof has three parts: the race bounds (9), deviations during the trading stage, and the formation stage.

Race bounds. Let τ_i be dealer i 's next revision time, which is independent of τ . On the event $\{\tau < \tau_i\}$, the dealer cannot repair a quote that is not at p^* . Rival actions, including responsive price changes or withdrawals, reveal no information about value, and they leave the terminal bound $\bar{K}$ valid at every resulting book. On the event $\{\tau_i < \tau\}$, every future strategy of the dealer has gross terminal expectation at most $\bar{M}$. This holds even if its eventual quote, market sale or abstention depends on history. Conditioning on the public history before the customer arrives preserves the prior and the uniform bounds. Nonnegative future costs can be dropped for an upper bound. Independence and the competition between exponential clocks then give (9). A dealer with no quote before τ_i earns zero, instead of at most $\bar{K}$, if the customer arrives first. An immediate market sale gives $B - m$.

Deviations during trading. When play leaves the class of full common-price books, construct an equilibrium of the auxiliary outside game from Lemma 2(ii). This game stops on the first return to that class, with payoffs given by the rank vector w . These boundary payoffs are feasible continuation payoffs of the rank policy, which satisfies Theorem 1. The gross upper bound $\bar{M}$ continues to apply to an outside strategy followed by that policy, since this is a feasible strategy with only nonnegative costs. Every action that leaves the class from a common-price book is therefore worth at most U_* . Retaining already gives $w_k > U_*$, and Theorem 1 determines whether resetting does better still. Hence no deviation in the full game during trading is profitable.

Patch together the outside equilibrium and the policy on the invariant class. A deviation over several stages either stays within the rank game or has a first action that leaves the class. Equivalently, apply the same one-shot deviation argument and the uniform tail bound. This proves sequential optimality at every history.

Formation. Consider a prefix of the formation order in which every dealer so far has quoted p^* . The last dealer earns its rank value less f by quoting p^* , which exceeds every alternative by (10). An initial quote away from p^* remains subject to the repair race bound, abstention to the bound for an empty dealer, and market sale to $B - m$. Induct backward over these all- p^* prefixes. Quoting p^* therefore induces every later dealer to quote p^* and yields the correctly anticipated rank value.

All other initial histories are completed by finite backward induction, using the stochastic-game continuations already constructed. The bounds do not presume that responses at those histories are passive. The belief before trading remains the prior, and customers act optimally at every resulting book by Lemma 1. $\square$

A.6 Proof of Theorem 3

Proof. The proof computes the immediate values and price bounds from the primitives, solves the rank equations, and then shows that every inequality is strict and survives small perturbations.

Values and price bounds. The six positive customer-type probabilities sum to $207/350$, leaving $b_0 = 143/350$. The customer utilities and liquidity tails computed in Section 3 give V and Δ . The bound over all books follows by separating liquidity fills, which are independent of value, from informed fills. If $p \leq 20$, both contributions are nonpositive. At 25, liquidity profit is at most $5(13/25)$; at 35, at most $15/10$; and above 38 there is no liquidity fill. Informed fills have nonpositive surplus by Lemma 1. The bound $26/5$ at 30 is uniform, and so are the resulting $\bar{K}$ and $\bar{M}$. Responsive rival actions cannot overturn these bounds, because they carry no information about value.

Rank values and resetting. Substituting V , $x = 1/4$ and $c = 0.01$ into (14) gives exactly $w = (21/5, 671/200, 791/200, 797/200)$ at $h = 0$. Direct subtraction gives (12), hence strict optimality of the reset policy under (7). A positive waiting cost enters through the common shift by h/λ (Theorem 1), which gives (11). The inequalities $\underline{w} - f = 3.245 > 3.12$ and $3.12 > 1.04 > 0 > -20$ establish condition (10) of Theorem 2 with $f = 0.1$.

Execution probabilities. For the immediate fill probability, add the liquidity reach to the high-value informed reach $a\mu\mathbf{1}_{\{k \leq 2\}}$. For the eventual fill probability, replace the reward vector V in the same killed chain by this reach vector, remove fees and waiting costs, and solve. This yields (13). In particular, $1259/2940 > 2019/4900$, whereas $671/200 < 791/200$. Common waiting costs and initial fees leave this ordering of surplus unchanged. At rank three, the immediate gain from moving to the back is $-11/100$ and the equilibrium gain is $1/50 > 0$. All participation is voluntary and strictly profitable.

Open neighborhood. Keep the finite grid and the finite set of customer types, but vary the positive probabilities, values, costs, rates, waiting cost and entry fee in a neighborhood. At the example, every customer choice over the finite set of possible books has a strict gap. Informed gains are half-integer multiples relative to the grid of five-unit ticks, and the liquidity reservations 33 and 38 are not grid prices. Hence the demand maps are locally constant.

Terminal expected payoffs are continuous in these primitives. Their maxima over finitely many books give continuous bounds, and the coarse bounds displayed above have slack relative to the exact maxima. The inverse of the killed chain is continuous for $\lambda > 0$. Every action, price, fee and participation inequality on which the result rests is strict, with $h = 0.01$, $c = 0.01$ and $f = 0.1$

in the interior of their nonnegative domains. All of them persist in some open neighborhood, relative to the probability simplex. No continuity of an arbitrarily selected off-path equilibrium is needed, because the uniform deviation bounds allow it to be selected anew. $\square$

A.7 Proof of Proposition 3 and equation (19)

Proof. The reset gains follow by solving the four rank equations; the aggregate identity follows by summing over owners.

Reset gains. Divide (6) by λ and use $S = \{2, 3\}$ to obtain (14). Substitute the last equation into the middle two and solve the resulting nonsingular system to obtain (15). Subtract $w_2 + c$ and $w_3 + c$ from w_4 and collect the strictly positive denominator Γ_x to obtain (16)–(17). Rank four always strictly retains when $c > 0$, since resetting returns it to rank four at cost c . Rank one retains under the stated inequality $w_4 - c < V_1 = w_1$. These four comparisons are exactly the action choices of the restricted game. With $V_2 = 3.2$, $V_3 = 4.1$ and $V_4 = 4$, the numerator of g_3 is $0.7x - 0.1 - c(1 + 2x)^2$, which proves (18); it can be positive only if $x > 1/7$. If $V_3 > V_4$, the immediate move to the back loses even before costs. A positive equilibrium gain therefore requires the responsive continuation.

Aggregate surplus. At any instant, summing the values of all owners over the ranks they hold assigns each terminal rank exactly once. At all times before the customer arrives, two ranks incur reset fees at total expected rate $2vc$, and all four incur waiting costs at rate $4h$. The expected duration is $1/\lambda$, which gives (19). Alternatively, summing the four linear equations cancels all transfers between ranks and gives the same expression.

Limits. For $c = 0$, the leading terms of (15) give the common limit $(V_2 + V_3 + V_4)/3$ of the three lower-rank values as $x \rightarrow \infty$. For $c > 0$, the numerator in (17) is eventually negative. So the candidate policy eventually fails. $\square$

A.8 Proof of Theorem 4

Proof. The proof proceeds in five steps. Step 1 specifies every second-round subgame and bounds its payoffs. Steps 2 and 3 show that these bounded continuation payoffs do not change first-round behavior. Step 4 assembles a PBE, and Step 5 derives the repricing and the repeat purchase.

Step 1: All terminal subgames. For each possible first book and observed execution count, use (20) under the proposed current-optimal first-round demand rule. There are finitely many physical first books, counts and feasible private inventory states. At every possible posterior and every publicly drawn order of terminal moves, solve the finite terminal dealer game by backward induction. At every leaf, solve customer quantities by direct maximization at every actual private state. This specifies play after every intermediate move, on and off the equilibrium path.

Withdrawal is free, and every dealer has a sure move before the second customer, so zero is always available. Terminal informed fills give the dealer non-positive surplus, and a liquidity unit yields at most $u_{\max}$ in conditional expectation. Unit capacity and the bound on liquidity mass give $0 \leq \Pi_i \leq \bar{D}$, with $\bar{D}$ as in (21). The argument holds even when earlier dealers deviate or stale quotes remain.

The initial order in the second round is defined relative to the observed first book and the list of filled owners. At a full common-price first book, choose the terminal tie conventions so that they are equivariant under relabeling of owners, and draw every permutation of dealers with equal probability. Expected second-round reward then depends on an owner's rank immediately before the first execution, not on its arbitrary name. The convention used breaks ties in favor of retaining or passing first, then withdrawal, then quotes in increasing order of price. Every dealer, including one that sold its first-round unit to the market, may use its remaining unit in the second round. A first-round market sale still ends that dealer's first-round moves permanently.

Step 2: Forward-looking first-round quantities. At every sorted grid book, the high-value insider's marginal gains are half-integers with absolute value at least $1/2$, nonincreasing in rank. The low-value insider's marginal losses are at least $15 + 1/2$. Hence every alternative quantity loses at least $g_I = 1/2$ relative to the current optimum.

When evaluating a private deviation, hold fixed the public posterior and continuation maps from Step 1. The deviation's new inventory may lie outside the support of the public posterior, but terminal optimization is defined there as well. Any such deviation can raise the insider's subsequent expected utility by at most ρU_I . Since $\rho U_I = 327/20000 < 1/2$, current-optimal demand remains strictly optimal at every first book, so it is optimal for the insider to follow the demand rule that dealers use in Bayesian updating. Uncertainty about the customer's identity, no trade, and the censoring at a full clear are all included in the sum over customer types.

Step 3: Recurrent first-round revisions. Set $\bar{\Pi}_k = \mathbb{E}[\Pi_k(Y)] \in [0, \bar{D}]$ at the all-target book. Solve the same candidate rank system with reward $\tilde{V} = V + \rho \bar{\Pi}$, and write its solution as $\tilde{w}$. The resolvent of the fixed reset generator has row sums $1/\lambda$, so

$$0 \leq \tilde{w}_k - w_k \leq \rho \bar{D}.$$

A retain-or-reset payoff difference therefore changes in absolute value by at most $2\rho \bar{D}$, bounding the two terms separately. The common waiting shift cancels from every difference. The strict gap $\delta = 0.01$ therefore preserves every prescribed action under (22).

For deviations in price and other actions, replace the terminal bounds by $\bar{M} + \rho \bar{D}$ and $\bar{K} + \rho \bar{D}$. The upper bound for a repaired price rises by at most $\rho \bar{D}$.

The bound for a dealer that is initially empty also rises by at most $\rho\bar{D}$, because it can still participate in the second round. A first-round market sale has upper bound $B - m + \rho\bar{D}$. By (22), all of these lie strictly below the equilibrium values at the all-target book, net of the initial fee. Lemma 2 and the formation induction in the proof of Theorem 2 now apply with these augmented rewards. This gives complete first-round strategies.

Step 4: Bayesian consistency and a PBE. Before the first execution, the uninformed dealers' actions and the independent survival of the clocks reveal nothing about v , so the prior applies at every first book. At first observations of positive probability, Step 2 makes (20) consistent with Bayes' rule. At unsupported observations, select the specified feasible posterior and the sequentially optimal terminal continuation. Dealers use that posterior; the insider uses its actual (v, z) . Steps 1–3 establish optimality at every information set. The profile is therefore a PBE.

Step 5: Repricing and the repeat purchase. The posterior in (23) follows by adding the liquidity and informed contributions to each count. At $Y = 2$ the probability of high value is $57/85$, and the high-value insider's inventory is either zero or two. At 80 both inventories lead to a purchase of exactly one unit; at 85 neither buys. The gross rank values at an all-80 book are

$$(202/35, 462/85, 451/85, 88/17).$$

For any final quote other than 80, a uniform upper bound is $351/85$. Quotes at 65 or below give nonpositive surplus. Quotes at 70 and 75 give at most $(75 - m_2)(13/25) = 351/85$; 85 gives at most $61/34$; and quotes at 90 or above have no liquidity demand. Informed contributions are nonpositive.

At a final book with mixed prices, an order at 80 at rank r has the same reach condition as rank r in an all-80 book. Its threshold is $80 + \kappa_r(z)$, and every preceding threshold is weakly lower. It therefore earns the corresponding displayed rank value, whatever the prices elsewhere; at depth below four, the same rank values apply by truncation. Since $88/17 - c > 351/85$, every dealer optimally finishes at 80 along every prefix of the move order on which earlier movers have done so, whatever the order of later movers.

Initially no surviving quote is at 80: all incumbents hold 30, and filled dealers are empty. Backward induction therefore makes every participant submit or replace at 80, with final time priority equal to the realized move order. All other intermediate histories are completed by the finite solution specified in Step 1. Averaging the all-80 rank values and deducting c only for the two incumbents gives the row $Y = 2$ of Table 1.

At $Y = 0$, the terminal values at an all-30 book are

$$(4563/815, 3548/815, 8323/1630, 812/163).$$

The uniform bound for quotes other than 30 is $3159/815$, which is below $3548/815 - c$. Thus all dealers finish at 30 in every move order; their relative time priority is chosen by finite backward induction and may change. For $Y = 1, 3$ and 4 , the posterior equals the prior, informed inventory is zero, and the bound $\bar{K} = 13/5 < \min_k V_k - c$ proves the same conclusion at the target price. The remaining entries of Table 1, given exactly in Table 2, are computed from the finite game.

Table 2: Exact expected second-round payoffs by first-round rank (decimal values in Table 1)

Y	$\Pi_1(Y)$	$\Pi_2(Y)$	$\Pi_3(Y)$	$\Pi_4(Y)$
0	4563/815	234049/48900	118121/24450	23539/4890
1	15/4	21/5	94/25	189/50
2	2581/476	2581/476	32203/5950	32203/5950
3	113/30	113/30	113/30	21/5
4	31/8	31/8	31/8	31/8

Finally, at $Y = 2$ and actual inventory $z = 2$, the high-value insider's strict terminal optimum at 80 is one unit. The event in which this happens has probability $\mu a^2 \rho > 0$, which proves repeated optimal informed purchasing. $\square$

B. Boundaries of the results and exact computations

This appendix collects short results that mark where the conclusions of the main text stop, and then describes the exact computations behind the numbers in the example.

A finer tick can break the common-price book. Consider the customer economy of Section 3 in a different game. There are no revision opportunities, dealers submit unit quotes simultaneously, equal quotes receive an exogenous time priority, and the grid also permits the price 29. Suppose all rivals quote 30. A dealer assigned rank two at 30 earns $16/5$. Undercutting at 29 places it first and gives

$$(29 - 20)(13/25) - (1/14)(1/5)(100 - 29) = 1283/350 > 16/5.$$

The all-30 profile is therefore not a Nash equilibrium of that terminal quote game. Privately optimal informed quantities do not by themselves protect a common price from undercutting on a finer grid. The example concerns only this simultaneous game without revision: it leaves open how later dealers would respond in a sequential quote game, and whether the model with recurrent revision and a one-unit tick has an equilibrium, and if so, one with harmful priority.

Why depth alone is not a sufficient state. Take two states with the same four owners and four units, in which a given owner is at rank two in one and at rank three in the other. The owner's equilibrium values w and its conditional exposure to the insider differ between them. A state that records only the number of units in the book cannot represent the owner's incentives. With persistent information, the same physical depth after zero or two past fills can also come with different joint posteriors and different optimal quotes. A recursion on depth alone is therefore not closed once revision and learning are added. A richer state may still be tractable.

Free entry and indivisible capacity. The equilibrium has positive entry rents because the pool of dealers and the displayed capacity are bounded. Suppose an identical new dealer could add a last unit at no cost and earn strictly positive continuation surplus. Excluding it cannot then be a free-entry equilibrium, because entering is a profitable feasible deviation. Suppose instead a dealer could split or duplicate its unit. The present action space, and the price and priority bounds built on it, would then no longer suffice. For these reasons the model does not make depth fully endogenous and does not derive dealers' optimal order sizes.

Without adverse stopping information. Suppose $p \geq \mathbb{E}[v \mid D = k]$ at every interior rank where demand stops with positive probability. Then all immediate margins are nonnegative, and the mechanism behind harmful priority is absent; informed trading alone does not produce it. Suppose instead that all V_k are equal and $c > 0$. Then retaining at every rank is the equilibrium of the rank game. A reset cannot change the terminal reward or the waiting time, and it strictly costs more. These cases show that the model does not mechanically force dealers to give up priority.

Local derivatives. Within a fixed set of customer choices in which every choice is strict, the demand indicators are locally constant. Differentiating their finite expectations gives the derivatives $\partial V_k / \partial p = F_k$ and $\partial \Delta_k / \partial p = \mathbb{P}(D = k)$ in Section 6. For a change in informed order size, subtract (24) before and after the change. Its effect on Δ_k is proportional to the difference between the reach changes at adjacent ranks, which has no uniform sign. For a fixed policy S , differentiating $(\lambda I - Q_S)w = \lambda V - h\mathbf{1} - c\nu r_S$ gives the local derivative of the equilibrium values. Crossing a policy boundary or a customer-choice boundary invalidates this fixed-matrix derivative, and no global monotonicity follows from it.

Exact computations. All numbers in the example are computed in exact rational arithmetic from the primitives, using three separate representations:

- (i) an enumeration of the 14,949 terminal books of depth one to four on the 22-price grid, with 29,898 informed quantity decisions;
- (ii) the full generator over the 24 ownership permutations, which gives all 96 retain-or-reset comparisons and the eventual fill probabilities; and
- (iii) a symbolic solution of the rank system (14), which gives the reset gains, the failures of the policy outside the region of Proposition 3 and the fine-tick example above.

The second round is solved as a finite game with 77,871 nodes, over the five first-round fill counts with positive probability and the 24 move orders; prices outside $(m, m + u_{\max})$ are excluded by a proved uniform bound, and retaining a stale quote is kept as an action. These computations supplement the proofs in Appendix A: the second-round pricing conclusions are proved analytically ($Y = 0$ and $Y = 2$ directly, $Y = 1, 3$ and 4 through $\bar{K}$), and the exact payoff entries of Table 1 are computed.