Price-Time Priority and Fine-Tick Non-Convergence

Abstract

In a price-time priority limit order book, a liquidity supplier who improves price by one tick becomes the front quote at the improved price and acquires its first execution states. When those states are adversely selected, improvement transfers adverse-selection exposure, not only price. I derive this margin and the implied price-time frontier from queue-priority primitives. Fine ticks remove the mechanical price concession but not the cost of owning toxic front priority, so the no-undercutting region need not vanish as the tick shrinks to zero. Credible, separating order-path disclosure shifts the frontier outward: a limit-order-book analogue of Sunshine Trading.

JEL classification: G10, G14, D82.

Keywords: limit order book; price-time priority; adverse selection; tick size; price improvement; sunshine trading.

1 Introduction

Consolidated limit order books are the dominant trading protocol in modern exchange-based markets. Their defining allocation rule, price-time priority, determines not only who trades first but which execution states each liquidity supplier owns. This paper studies a local but consequential part of that design: the decision of a passive liquidity supplier to improve price by one tick rather than remain behind existing depth.

The textbook intuition is that price improvement trades a small price concession for faster execution. This intuition is incomplete in a price-time priority book, and the way it is incomplete is the subject of this paper. A trader who improves one tick does not only move ahead of the old queue. She becomes the front quote at the improved price and therefore acquires the first execution states at that price. If those front states are adversely selected, price improvement transfers adverse-selection exposure to the trader who improves. The cost of improvement is not only the tick.

Models of tick size and queue-jumping already study a tradeoff between price improvement and time priority, so the contribution must be located precisely. In existing models, undercutting is disciplined by the tick cost, by fill probability, or by the option value of queue position. In this paper, undercutting is disciplined by a different object: the adverse-selection value of the newly acquired front-priority block,

$$\Gamma^{\text{front}} = \sum_{m \in B^{\text{front}}} \Delta_m, \qquad T^{\text{front}} = \left[-\Gamma^{\text{front}}\right]_+, \tag{1}$$

where Δ_m is the queue-priority value of marginal front execution state m (price and state arguments are suppressed). This object is derived from a queue-priority value function, not assumed. It is also why the paper's central comparative static is possible: the tick cost can vanish while T^{front} remains. No relabeling of queue-jumping delivers that separation, because in the standard accounts the cost of undercutting is the tick itself.

The economics is a relocation result. In dealer markets, adverse-selection exposure is priced explicitly through a quote or spread. In a consolidated limit order book there is no single dealer who absorbs that exposure for the market. Instead, exposure is allocated across queue positions by the priority rule. A passive trader chooses not only a price but a location on a price-time surface, and that location determines which execution states the trader owns. No participant quotes the dealer's spread in a book; its adverse-selection component is borne, position by position, along the queue.

The paper formalizes this frontier. Consider an ask-side trader deciding whether to hold at ask a or improve one tick to $a - \delta$. Let G^P denote the non-front execution-access gain from improving, δF^I the tick concession on expected improved fills, C^I the cost of improving or reposting, and Γ^{front} the signed value of the newly acquired front-priority block. The improvement margin is

$$\Phi = G^P - \delta F^I + \Gamma^{\text{front}} - C^I. \tag{2}$$

The price-time frontier is the condition $\Phi = 0$. When $\Phi > 0$, improvement is attractive; when $\Phi < 0$, holding is preferred conditional on active participation. The decomposition is incremental and carefully avoids double counting: the value of holding already contains the exposure of the current quote, and Γ^{front} records only the signed exposure newly acquired by moving to the improved price. Lemma 2 makes this convention precise, and a companion queue-priority primitive, $\Delta_m = \pi b_m[\tau(\pi) - A_m]$, signs each marginal front state: a front state is toxic exactly when its rank-specific adverse-selection intensity A_m exceeds the noise-to-informed odds $\tau(\pi)$.

The paper delivers four sets of results on this frontier.

First, a static characterization. Proposition 4 shows that holding weakly dominates one-tick improvement if and only if $\Phi \leq 0$, so no-undercutting is the local response to a negative incremental priority transfer rather than a failure of competition. Proposition 6 adds an outside option and characterizes improve, hold, and exit regions; cancellation emerges as an equilibrium participation decision, not as an anomaly. Proposition 7 shows that heterogeneous liquidity suppliers sort across these regions, so displayed depth is the cross-section of types for whom the price-time exposure at that location is acceptable.

Second, the paper's central result. Theorem 9 is a fine-tick non-convergence theorem. Consider a sequence of markets with tick size $\delta_n \downarrow 0$. If there is a positive-measure region of states in which the transferred front block is sniping dominated, then queue-priority primitives imply a negative signed front-block value that is bounded away from zero independently of the tick. When the non-front access gain is bounded by this primitive toxicity gap, the no-undercutting region does not vanish as the tick converges to zero. The limiting frontier is $\Phi_0 = G^P + \Gamma^{\text{front}} - C^I$: fine ticks remove the mechanical price concession but not the adverse-selection cost of owning the front execution states. The theorem is deliberately a non-convergence-of-region result, not a spread-floor theorem; it does not require, and does not claim, a full dynamic equilibrium of the book.

Third, a robustness result. Proposition 11 embeds the improve-hold-exit margin in a dynamic quote-placement problem and shows that, under single-crossing restrictions on continuation-adjusted action values, the static frontier survives as an ordered threshold policy. The dynamic section is a robustness check on the static margin, not a separate equilibrium theory of the book.

Fourth, a mechanism transfer. Theorem 13 gives a limit-order-book analogue of Sunshine Trading. In the dealer-market original, credible disclosure that an order is liquidity motivated lowers the dealer's adverse-selection cost. In a consolidated book, the analogous public object is order-path credibility: displayed depth, slicing regularity, timing, routing, and willingness to bear front priority. The theorem supplies the missing incentive condition. A high-credibility path is informative only if liquidity-motivated traders are willing to choose it and toxic traders do not find it profitable to mimic. Under this separation condition, order-path credibility raises the signed posterior value Γ^{front}, shifts the price-time frontier outward, and expands the region in which liquidity suppliers improve or remain active. The value of the result is the translation, from a dealer disclosure mechanism to a price-time priority object, rather than the comparative static itself.

The paper is pure theory, and its claims are calibrated accordingly. It does not solve for a full dynamic equilibrium of the limit order book, rank priority rules by welfare, or provide an empirical test for manipulation. Its contribution is to isolate a price-time mechanism that the standard language of tick size and queue-jumping conceals: price improvement is the purchase of front priority, front priority can be toxic, and therefore fine ticks do not eliminate the price-time frontier.

1.1 Related literature

The paper sits at the intersection of four strands of market microstructure, and it differs from each in a specific way.

Adverse selection and order-flow inference. Classic models price adverse selection into a dealer quote or a competitive spread (Glosten and Milgrom, 1985; Kyle, 1985; Glosten, 1994), and empirical work documents adverse-selection costs borne by limit orders (Sandås, 2001). This paper keeps the adverse-selection problem but relocates the exposure from a single quote to positions on a price-time surface. The relevant object is not only the price at which liquidity is supplied but the execution states assigned by priority. Relative to this literature, the paper adds queue-location exposure: two quotes at the same price with different ranks bear different adverse-selection costs.

Limit order book theory. Equilibrium models of order placement study the choice between market and limit orders, execution risk, and the dynamics of the book (Foucault, 1999; Parlour, 1998; Foucault et al., 2005; Goettler et al., 2005; Rosu, 2009; Hollifield et al., 2004). These models price execution access; the front of the queue is valuable because it trades first. This paper separates execution access from front exposure: the first execution states are not merely earlier draws from the same distribution, they are the states most likely to be on the wrong side of information. The signed front-block value Γ^{front} can therefore be negative even when access is valuable, which is the wedge none of the placement models carry.

Tick size and queue-jumping. Work on tick size and time priority studies how a coarse price grid rations queue position and rewards speed (Yao and Ye, 2018), and speed-race models show that adverse-selection races survive as mechanical latencies shrink (Budish et al., 2015; Menkveld, 2013; Hoffmann, 2014). The fine-tick theorem is complementary but distinct: it is not about the race to the front, it is about whether the front is worth owning. Even when fine ticks eliminate the mechanical cost of undercutting, the signed value of the transferred front block need not vanish, so the no-undercutting region can survive the fine-tick limit. In tick-size models the binding object disappears with the tick; here it does not.

Sunshine Trading and disclosure. Admati and Pfleiderer (1991) show that credible disclosure of liquidity motives can lower adverse-selection costs and coordinate liquidity; related work studies how predictable uninformed flow concentrates trading (Admati and Pfleiderer, 1988). This paper transfers the sunshine mechanism from a dealer-market announcement to order-path credibility in a book, and adds the incentive condition the transfer requires: the high-credibility path must be separating, otherwise toxic mimicry undoes the inference. The claim is a mechanism-transfer claim, not an equivalence claim. A consolidated book is not a dealer market, but the dealer's spread has a book analogue in the price-time frontier, the adverse-selection component of the spread has an analogue in toxic-front exposure, and the sunshine channel has an analogue in credible order-path geometry.

The fee implication relates to work on maker-taker pricing and limit-order execution quality (Colliard and Foucault, 2012; Malinova and Park, 2015; Battalio et al., 2016). Here the rebate is not a welfare instrument by itself; it shifts a particular frontier, the margin for bearing signed front-priority exposure. The queue-priority machinery builds on the companion priority valuation model (Anonymous, 2026).

The remainder of the paper proceeds as follows. Section 2 sets up the local quote-placement model and derives the improvement margin from queue-priority primitives. Section 3 characterizes the static price-time frontier. Section 4 adds participation and type sorting. Section 5 states and proves the fine-tick non-convergence theorem. Section 6 establishes dynamic robustness. Section 7 develops the order-path credibility result. Section 8 collects implications, including maker-taker fees and empirical content, and Section 9 concludes. All proofs are in Appendix A.

2 Model

The model is intentionally local. It studies a same-side quote-placement problem at a given public order-book state, which keeps the price-time mechanism separate from a full equilibrium theory of the entire book. Section 6 shows that the same margin survives when payoffs are replaced by continuation-adjusted action values.

2.1 Environment

Let $s \in S$ be the public state, where (S, Σ, μ) is a measure space. The state includes the current best prices, queue depths, volatility, recent order flow, and any public signals relevant for adverse selection. A passive ask quote is described by a price p and rank k at that price, where $k = 1$ is the front of the queue and higher time priority corresponds to lower k. Let

$$V(p, k; s) \tag{3}$$

denote the continuation value of posting at (p, k) in state s, over a fixed decision horizon that ends when the trader next revises the quote.

The decision problem is the following. Quoted prices lie on a grid of width $\delta > 0$, the tick. An ask-side trader with displayed size q currently holds, or can join, the queue at the prevailing ask a. The trader chooses between holding at a (action H) and improving one tick to $a - \delta$ (action I); Section 4 adds an outside option (action O). Improving makes the trader the front quote at the improved price. Over the horizon, random marketable buy volume $M(p; s)$ reaches price p; this is the reduced-form analogue of the execution-arrival variable in the companion queue-priority model (Anonymous, 2026), indexed by price to distinguish the current quote from the improved quote.

2.2 Queue-priority increments and the front block

Definition 1 (Queue-priority increment)**.** The same-price priority increment at price p and rank k is

$$\Delta_k(p; s) = V(p, k; s) - V(p, k+1; s). \tag{4}$$

The increment Δ_k is the value of owning the marginal execution states assigned to rank k rather than rank $k+1$. It can take either sign. A positive increment means that moving one rank forward at the same price is valuable. A negative increment means that the marginal step forward exposes the trader to adverse-selection states whose conditional value is worse than the execution-access gain.

Summing increments recovers the value of front priority from the value function itself, which is the sense in which everything below is derived rather than assumed.

Lemma 1 (Front-block telescoping)**.** *For any price p, size $q \geq 1$, and state s,*

$$V(p, 1; s) - V(p, q+1; s) = \sum_{m=1}^{q} \Delta_m(p; s). \tag{5}$$

For a quote at price p, the first q units of queue priority are the states in which the quote executes because marketable volume reaches the front of the queue. Write

$$B^{\text{front}}(p, q; s) = \{m : 1 \leq m \leq q,\ m \leq M(p; s)\}, \tag{6}$$

where the index m labels marginal front units or fill states. The notation is reduced form by design: B^{front} is the priority block transferred to the trader when the trader becomes the front quote at price p.

If the trader improves from a to $a-\delta$, the trader receives the block $B^{\text{front}}(a-\delta, q; s)$. Define the signed front-block value and its negative part, the toxic-front exposure:

$$\Gamma^{\text{front}}(a-\delta, q; s) = \sum_{m \in B^{\text{front}}(a-\delta, q; s)} \Delta_m(a-\delta; s), \tag{7}$$

$$T^{\text{front}}(a-\delta, q; s) = \left[-\Gamma^{\text{front}}(a-\delta, q; s)\right]_+. \tag{8}$$

These objects are functionals of the queue-priority value function through Definition 1 and Lemma 1; they are not additional primitives. The signed object Γ^{front} records whether the transferred front block helps or hurts the improving trader. When the block has positive value, it adds to the improvement margin. When it has negative value, T^{front} is the cost of being forced to own those front execution states.

The execution-access gain is defined separately. Let $B^{\text{access}}(a-\delta, q; s)$ be the set of non-front execution states newly reachable because the trader improves price, excluding the transferred front block. For $m \in B^{\text{access}}$, let $g_m(a-\delta; s) \geq 0$ denote the incremental surplus from accessing that non-front fill state relative to holding at a. Define

$$G^P(s) = \sum_{m \in B^{\text{access}}(a-\delta, q; s)} g_m(a-\delta; s). \tag{9}$$

Thus G^P is not the residual that makes the decomposition balance. It is the value of non-front access created by price improvement, and the signed front block Γ^{front} accounts for the first execution states separately.

2.3 The improvement margin

Let $V_H(s)$ be the value of holding or joining at the current ask, $V_I(s)$ the value of improving one tick, $F^I(s) \geq 0$ the expected fill quantity at the improved quote, and $C^I(s) \geq 0$ the cost of improving or reposting. The improvement margin is

$$\Phi(s) = V_I(s) - V_H(s). \tag{10}$$

Lemma 2 (Improvement decomposition; no double counting)**.** *Suppose the gross value of the improved quote differs from the holding value through three channels only: the transferred front block, the newly accessible non-front fill states, and the price concession δ per expected fill unit, and that improvement incurs cost $C^I(s)$. Then*

$$\Phi(s) = G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(a - \delta, q; s) - C^I(s). \tag{11}$$

Moreover, the decomposition is incremental: any adverse-selection exposure already borne by the holding quote remains in $V_H(s)$, and Γ^{front} is only the signed value of the newly acquired front block at the improved price.

Lemma 2 is the accounting backbone of the paper, and it is what blocks the concern that front toxicity is smuggled in as an extra friction. The model does not add front toxicity on top of the full value of the improved quote; it decomposes the difference $V_I - V_H$. In the toxic-front region, where $\Gamma^{\text{front}} = -T^{\text{front}}$, the familiar expression $G^P - \delta F^I - T^{\text{front}} - C^I$ is recovered.

Remark 1 (What the three channels exclude). The exhaustiveness hypothesis of Lemma 2 is an accounting convention, not an economic restriction. Channels that operate through rivals' responses to the improvement, such as re-undercutting or queue replenishment behind the new quote, are not assumed away; they are embedded in the continuation values that define Δ_m, g_m, and F^I. What the hypothesis rules out is a fourth channel outside those objects, for example a direct payoff to the act of improving that operates through neither execution states nor the price of fills.

When the improved price and size are clear from context, $\Gamma^{\text{front}}(s)$ and $T^{\text{front}}(s)$ abbreviate $\Gamma^{\text{front}}(a-\delta, q; s)$ and $T^{\text{front}}(a-\delta, q; s)$. The explicit arguments are restored whenever the transferred block matters.

Definition 2 (Price-time frontier)**.** The price-time frontier is the state set

$$\mathcal{F} = \{s \in S : \Phi(s) = 0\}. \tag{12}$$

States with $\Phi(s) > 0$ favor price improvement. States with $\Phi(s) < 0$ favor holding, conditional on active participation.

2.4 Microfoundation of the increment sign

The companion queue-priority primitive gives a microfoundation for the sign of Δ_m, which the fine-tick theorem uses.

Assumption 1 (Informed/noise mixture)**.** Each marginal front state m at price p is reached with probability $p_m(s) > 0$. Conditional on being reached, the executing flow is informed with probability $\pi(s) \in (0, 1)$ and liquidity motivated (noise) with probability $1 - \pi(s)$. Execution against noise flow yields a gain $\beta_m(s) > 0$; execution against informed flow yields a loss $A_m(s)\, \beta_m(s)$, where $A_m(s) \geq 0$ is the rank-specific adverse-selection intensity. Write $b_m(s) = p_m(s)\, \beta_m(s) > 0$ for the reach-weighted noise gain.

Lemma 3 (Increment sign)**.** *Under Assumption 1, with prior noise-to-informed odds $\tau(\pi) = (1-\pi)/\pi$,*

$$\Delta_m(p; s) = \pi(s)\, b_m(s) \left[\tau(\pi(s)) - A_m(s)\right]. \tag{13}$$

In particular, the marginal front state m is toxic, $\Delta_m(p; s) < 0$, if and only if $A_m(s) > \tau(\pi(s))$.

A sniping-dominated state is one in which the first execution states have high adverse-selection intensity relative to the noise-flow odds. Rank specificity matters: A_m is typically largest for the earliest front states, because flow that executes instantly against a freshly improved quote is disproportionately likely to be reacting to information the quote has not yet impounded.

Assumption 2 (Local boundedness)**.** For each state studied below, V_H, G^P, F^I, Γ^{front}, and C^I are finite, with $C^I \geq 0$ and $F^I \geq 0$.

The local model can be read either as a reduced-form one-period game or as a Bellman action-value comparison; Section 6 makes the second reading explicit.

3 The Static Price-Time Frontier

This section gives the first result in its most economical form. The object is a local best response, not a full entry game. A liquidity supplier at a given book state compares holding at the current ask with improving one tick and becoming the front quote at the improved price.

Fix state s, current ask a, tick size δ, and displayed size q. By Lemma 2, the deviation from holding to improving has four components: non-front execution access, the tick concession on fills, the signed value of the acquired front-priority block, and the cost of improving,

$$\Phi(s) = G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(a - \delta, q; s) - C^I(s). \tag{14}$$

Proposition 4 (Price-time best response)**.** *Holding weakly dominates one-tick improvement if and only if* $\Phi(s) \leq 0$*, equivalently*

$$G^P(s) + \Gamma^{\text{front}}(a - \delta, q; s) \leq \delta F^I(s) + C^I(s). \tag{15}$$

Improvement weakly dominates holding if and only if the reverse inequality holds.

Proposition 4 is a best-response statement rather than an equilibrium claim. Its role is to isolate the price-time frontier that enters the fine-tick and disclosure results, not to extract strategic content from a two-by-two matrix. A trader improves only if the non-front execution-access gain plus the signed value of the transferred front block is large enough to pay for the lower execution price and the improvement friction. In the toxic-front region, where $\Gamma^{\text{front}}(a-\delta,q;s) = -T^{\text{front}}(a-\delta,q;s) < 0$, the hold condition becomes

$$G^P(s) \leq \delta F^I(s) + T^{\text{front}}(a-\delta,q;s) + C^I(s). \tag{16}$$

No-undercutting is therefore not a failure of competition. It is the local response to a negative incremental priority transfer.

Corollary 5 (Toxic-front brake)**.** *Fix s, a, and q. If*

$$G^P(s) + \Gamma^{\text{front}}(a-\delta,q;s) < C^I(s), \tag{17}$$

then holding strictly dominates one-tick improvement for every $\delta F^I(s) \geq 0$. No-undercutting then survives arbitrarily small tick concessions, including a tick concession of zero.

Corollary 5 previews the central comparative static. The brake on undercutting in inequality (17) contains no tick term: it compares access value plus signed front-block value with improvement cost. Section 5 turns this observation into a limit theorem by deriving a uniform negative bound on Γ^{front} from queue-priority primitives.

This formulation leaves the strategic order-placement game for future work. A full game would have to specify how ranks are assigned when multiple traders improve, how congestion costs are generated, and how price priority interacts with entry. Those objects are not needed for the paper's central claim, which is that the local improvement margin contains a signed front-priority transfer and that this transfer can survive the fine-tick limit.

4 Participation, Cancellation, and Type Sorting

The two-action comparison explains movement along the price-time surface. A liquidity supplier also decides whether to participate on that surface at all. Let $V_O(s)$ be the outside option: cancel, stay out, or repost away from the toxic priority block. The individual problem is

$$\max\{V_H(s),\, V_H(s) + \Phi(s),\, V_O(s)\}. \tag{18}$$

Proposition 6 (Participation regions)**.** *The optimal action is characterized by*

$$H \iff V_H(s) \geq V_O(s) \text{ and } \Phi(s) \leq 0, \tag{19}$$

$$I \iff V_H(s) + \Phi(s) \geq V_O(s) \text{ and } \Phi(s) \geq 0, \tag{20}$$

$$O \iff V_H(s) < V_O(s) \text{ and } V_H(s) + \Phi(s) < V_O(s). \tag{21}$$

Proposition 6 separates two margins. The improvement margin Φ determines movement along the price-time surface; the participation margin determines whether the trader remains on the surface at all. The separation gives a clean interpretation of cancellation: cancellation is not mechanically suspicious or off-equilibrium. It is the equilibrium response when both active quote values are dominated by the outside option. In states where holding and improving are both dominated, a bona fide liquidity supplier exits or reposts away from the toxic front block.

Now let liquidity suppliers have types

$$\theta = (C^I, \lambda, \rho, \sigma, K^O), \tag{22}$$

where C^I is improvement cost, λ parameterizes toxicity sensitivity, ρ inventory motive, σ

speed or monitoring ability, and K^O reposting friction. Write

$$\Phi(s,\theta) = G^P(s,\theta) - \delta F^I(s,\theta) + \Gamma^{\text{front}}(s,\theta) - C^I(\theta), \tag{23}$$

and let $V_O(\theta) = -K^O(\theta)$ for notational simplicity.

Proposition 7 (Type sorting)**.** *For fixed s, the type space is partitioned into improve, hold, and exit regions:*

$$\Theta_I(s) = \{\theta : V_H(s,\theta) + \Phi(s,\theta) \geq V_O(\theta),\ \Phi(s,\theta) \geq 0\}, \tag{24}$$

$$\Theta_H(s) = \{\theta : V_H(s,\theta) \geq V_O(\theta),\ \Phi(s,\theta) \leq 0\}, \tag{25}$$

$$\Theta_O(s) = \{\theta : V_H(s,\theta) < V_O(\theta),\ V_H(s,\theta) + \Phi(s,\theta) < V_O(\theta)\}. \tag{26}$$

Holding the state and the remaining components of θ fixed, higher improvement costs and lower signed front-block values weakly shrink the improve region; in toxic-front states, higher toxic-front exposure weakly shrinks the improve region. Higher reposting frictions weakly shrink the exit region.

The book is therefore a sorted ecology of liquidity suppliers. Observed depth at a price is not just aggregate willingness to trade. It is the cross-section of types for whom the price-time exposure at that location is acceptable. This sorting result is what later gives the sunshine theorem its participation content: anything that raises the signed posterior value of front states moves types from exit and hold toward improvement.

5 Fine-Tick Non-Convergence

The previous sections identify the price-time frontier. This section gives the paper's central result and derives the persistent toxic-front region from queue-priority primitives. The point is not merely that a bad region persists if one assumes it persists. The point is that in a sniping-dominated order-flow regime, the front-priority block has negative signed value for reasons unrelated to the tick size.

The usual tick-size intuition is mechanical. If price improvement costs one tick, then smaller ticks make improvement cheaper, and in the limit one might expect price competition to erase no-undercutting regions. That conclusion follows only if the tick concession is the relevant cost of price improvement. Under price-time priority, improving also transfers the first execution states at the improved price, and the value of those states is governed by the rank-specific adverse-selection intensity of front executions, not by the tick.

Consider a sequence of markets indexed by n, with tick size $\delta_n \downarrow 0$. At state s, define

$$\Phi_n(s) = G_n^P(s) - \delta_n F_n^I(s) + \Gamma_n^{\text{front}}(s) - C_n^I(s), \qquad \Gamma_n^{\text{front}}(s) = \sum_{m \in B_n^{\text{front}}(s)} \Delta_{m,n}(s), \tag{27}$$

with non-front access gain

$$G_n^P(s) = \sum_{m \in B_n^{\text{access}}(s)} g_{m,n}(s), \qquad g_{m,n}(s) \geq 0. \tag{28}$$

The no-undercutting region is

$$N_n = \{s \in S : \Phi_n(s) \leq 0\}. \tag{29}$$

Assumption 2 is maintained at each n, so $C_n^I(s) \geq 0$ and $F_n^I(s) \geq 0$ throughout.

Assumption 3 (Primitive sniping region)**.** There exist a threshold $N_0 < \infty$, constants $\bar{G} < \infty$ and $\underline{T} > 0$ with $\bar{G} < \underline{T}$, and a measurable region $A \subseteq S$ with $\mu(A) > 0$, such that for all $n \geq N_0$ and all $s \in A$:

(i) *Bounded non-front access:* $\sum_{m \in B_n^{\text{access}}(s)} g_{m,n}(s) \leq \bar{G}$.

(ii) *Sniping-dominated front block:* each transferred front state admits the informed/noise representation of Lemma 3,

$$\Delta_{m,n}(s) = \pi_n(s)\, b_{m,n}(s) \left[\tau(\pi_n(s)) - A_{m,n}(s)\right], \qquad m \in B_n^{\text{front}}(s), \tag{30}$$

and the aggregate toxicity of the block satisfies

$$\sum_{m \in B_n^{\text{front}}(s)} \pi_n(s)\, b_{m,n}(s) \left[A_{m,n}(s) - \tau(\pi_n(s))\right] \geq \underline{T}. \tag{31}$$

Assumption 3 is the primitive content of the theorem, and its structure matters. The region A is not defined by $\Phi_n \leq 0$; that would be circular. It is defined by bounded non-front access and a sniping-dominated front block in the queue-priority primitives. A non-empty improved-level queue, bounded displayed size, or bounded non-front marketable volume delivers the access bound $\bar{G}$. The toxicity bound in inequality (31) says that, on A, the rank-specific adverse-selection intensities of the transferred front states exceed the noise-flow odds by a fixed margin in the aggregate. The economically relevant gap is $\underline{T} - \bar{G} > 0$: front toxicity dominates what non-front access can pay for. Two primitives are conspicuously absent. The assumption places no condition on the tick sequence, and no bound on expected fills: fills enter the theorem only through $\delta_n F_n^I(s) \geq 0$, and enter the limiting frontier only through the local boundedness imposed in Corollary 10.

Lemma 8 (Primitive front-toxicity lower bound)**.** *Under Assumption 3, for all $n \geq N_0$ and all $s \in A$,*

$$\Gamma_n^{\text{front}}(s) \leq -\underline{T}, \qquad \textit{equivalently} \quad T_n^{\text{front}}(s) \geq \underline{T}. \tag{32}$$

The lemma converts the primitive sniping condition into a uniform negative bound on the signed front-block value. The bound is independent of δ_n by construction: nothing in inequality (32) involves the tick.

Theorem 9 (Fine-tick non-convergence)**.** *Under Assumption 3, for all $n \geq N_0$ and all $s \in A$,*

$$\Phi_n(s) \leq -(\underline{T} - \bar{G}) < 0. \tag{33}$$

In particular, $A \subseteq N_n$ and

$$\mu(N_n) \geq \mu(A) > 0 \tag{34}$$

for all $n \geq N_0$. The no-undercutting region does not vanish as $\delta_n \downarrow 0$, and on A the improvement margin is bounded away from the frontier uniformly in the tick.

Corollary 10 (Limiting frontier)**.** *Suppose in addition that $G_n^P(s) \to G^P(s)$, $\Gamma_n^{\text{front}}(s) \to \Gamma^{\text{front}}(s)$, and $C_n^I(s) \to C^I(s)$ pointwise, with $F_n^I(s)$ locally bounded. Then*

$$\Phi_n(s) \to \Phi_0(s) = G^P(s) + \Gamma^{\text{front}}(s) - C^I(s), \tag{35}$$

and the limiting no-undercutting region is

$$N_0 = \{s : G^P(s) + \Gamma^{\text{front}}(s) \leq C^I(s)\}, \tag{36}$$

which on the toxic-front region takes the form $N_0 = \{s : G^P(s) \leq T^{\text{front}}(s) + C^I(s)\}$.

Three remarks calibrate the result.

First, the theorem is a non-convergence-of-region result, not a spread-floor theorem. It does not say that equilibrium spreads remain bounded away from zero as ticks shrink; that claim would require a full equilibrium model of the book, including entry and the assignment of ranks when multiple traders improve. The result is narrower and cleaner: in a primitive sniping-dominated region, rational liquidity suppliers decline to improve price even as the tick goes to zero.

Second, the limiting object identifies what fine ticks do and do not remove. The limiting margin is $G^P + \Gamma^{\text{front}} - C^I$. Fine ticks eliminate the deterministic price concession δF^I. They

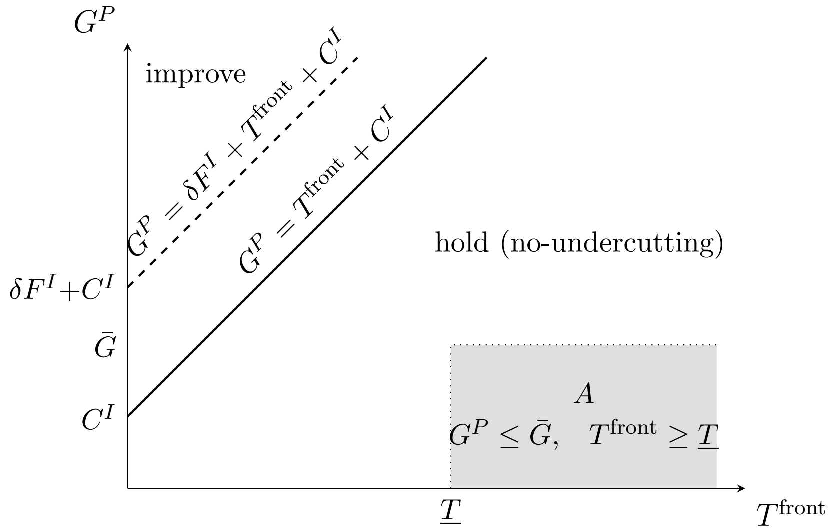

Figure 1: The price-time frontier in toxic-front states, where $\Gamma^{\text{front}} = -T^{\text{front}}$. For a positive tick, improvement requires $G^P \geq \delta F^I + T^{\text{front}} + C^I$ (dashed line). As $\delta_n \downarrow 0$ the frontier shifts to $G^P = T^{\text{front}} + C^I$ (solid line): the tick concession vanishes, but the toxic-front term does not. The shaded rectangle is the primitive sniping region A of Assumption 3: bounded non-front access ($G^P \leq \bar{G}$) and front toxicity at least $\underline{T} > \bar{G}$ place A strictly inside the no-undercutting region for every tick.

Alt text: Schematic with toxic-front exposure on the horizontal axis and the access gain on the vertical axis. Two parallel upward-sloping lines mark the frontier: a dashed line for a positive tick above a solid line for the zero-tick limit. The improve region lies above the lines and the hold region below. A shaded rectangle labeled A, with low access gain and high toxicity, sits in the lower right, strictly below both lines.

do not eliminate the signed value of the front execution states, because that value is priced by A_m relative to $\tau(\pi)$, objects that have nothing to do with the price grid.

Third, the theorem explains why tick-size reforms can disappoint. If the policy diagnosis is that coarse ticks block price competition, then fine ticks should produce pervasive undercutting. The theorem identifies the state region where they will not: wherever the front of the book is toxic enough that the access gain cannot pay for it, the price-time frontier survives the reform. Figure 1 summarizes the geometry.

6 Dynamic Robustness

This section is a robustness check rather than a separate equilibrium theory of the full order book. The static model embeds in a dynamic quote-placement problem without changing the economic margin: the improve-hold-exit comparison survives when payoffs are replaced by continuation-adjusted action values.

Let a trader of type θ choose among H, I, and O each period. Fix a stationary rival policy α_j. The trader's Bellman value is

$$W(s,\theta) = \max_{a\in\{H,I,O\}} Q(s,a;\theta), \qquad Q(s,a;\theta) = u(s,a;\theta,\alpha_j(s)) + \beta\,\mathbb{E}\left[W(s',\theta) \mid s,a,\alpha_j(s)\right], \tag{37}$$

with discount factor $\beta \in (0,1)$ and bounded per-period payoffs u. Define the dynamic improvement margin and the dynamic participation margin:

$$\Phi^D(s,\theta) = Q(s,I;\theta) - Q(s,H;\theta), \qquad \Psi^D(s,\theta) = \max\{Q(s,H;\theta), Q(s,I;\theta)\} - Q(s,O;\theta). \tag{38}$$

Proposition 11 (Dynamic action regions)**.** *For fixed stationary rival policy, bounded payoffs, and $\beta \in (0,1)$, the Bellman operator in equation (37) is a contraction with a unique fixed point, and the best response is characterized by action-value comparisons:*

$$\begin{aligned} I &\iff Q(I) \geq \max\{Q(H), Q(O)\}, \\ H &\iff Q(H) \geq \max\{Q(I), Q(O)\}, \\ O &\iff Q(O) \geq \max\{Q(H), Q(I)\}. \end{aligned} \tag{39}$$

Proposition 11 states that the static margins of Sections 3 and 4 are the one-period shadows of dynamic action-value margins. To recover the ordered improve-hold-exit picture, impose a scalar ordered state x representing front toxicity or cost, alongside residual state y.

Suppose

$$\begin{aligned} Q(H;x,y,\theta) &= A_H(y,\theta) - h(\theta)x, \\ Q(I;x,y,\theta) &= A_I(y,\theta) - i(\theta)x, \\ Q(O;x,y,\theta) &= -K^O(\theta), \end{aligned} \tag{40}$$

with $i(\theta) > h(\theta) \geq 0$. The inequality $i > h$ is the dynamic analogue of the front-block transfer: improving is more exposed to front toxicity than holding.

Corollary 12 (Ordered threshold policy)**.** *Under the single-crossing structure* (40)*, Φ^D is strictly decreasing in x, with improvement cutoff*

$$\bar{x}_I(y,\theta) = \frac{A_I(y,\theta) - A_H(y,\theta)}{i(\theta) - h(\theta)}, \qquad \bar{x}_O(y,\theta) = \frac{A_H(y,\theta) + K^O(\theta)}{h(\theta)} \tag{41}$$

(the latter defined for $h(\theta) > 0$). If $\bar{x}_I \leq \bar{x}_O$, the optimal policy has the ordered threshold form

$$x < \bar{x}_I : \ I, \qquad \bar{x}_I \leq x \leq \bar{x}_O : \ H, \qquad x > \bar{x}_O : \ O. \tag{42}$$

The dynamic results are best-response statements, not equilibrium claims. They do not establish uniqueness of Markov-perfect equilibrium, and they impose rather than derive the single-crossing structure. Their role is to certify that nothing in the static analysis depends on myopia: when the equilibrium action values satisfy the same single-crossing restrictions, the static price-time frontier survives dynamically as an ordered threshold in front toxicity.

7 Order-Path Credibility and the Sunshine Frontier

The second paper-level result transfers the logic of Sunshine Trading (Admati and Pfleiderer, 1991) to a consolidated limit order book. In the dealer-market original, a liquidity trader can credibly reveal that a large order is not information based. The credibility of the announcement is not cosmetic; it is the mechanism. If informed traders could cheaply mimic the announcement, the announcement would not lower adverse-selection costs.

In a consolidated book there is no single dealer who receives a clean preannouncement. The public object is instead the geometry of the order path. Displayed depth, slicing regularity, child-order timing, routing, and willingness to bear priority all affect beliefs about whether front executions are informed or liquidity motivated. The relevant question is therefore not whether an order path is transparent, but whether it is separating.

7.1 Separating order paths

Let $z \in \{L, H\}$ be a public order-path signal, with H the high-credibility path. There are two order-flow types, B and T. Type B is liquidity motivated; type T is toxic or information motivated. A type $j \in \{B, T\}$ who chooses path z pays path cost $k_j(z)$, with $k_j(L) = 0$. The high-credibility path can be costly because it requires patient slicing, displayed exposure, or other behavior that is natural for liquidity demand but costly for toxic demand.

The implied game has the following timing. Nature draws the order-flow type $j \in \{B, T\}$. The type chooses a path $z \in \{L, H\}$ and pays $k_j(z)$. Passive liquidity suppliers observe z but not j, update beliefs about front-execution states, and choose among improving, holding, and exiting as in Sections 3 and 4. Execution then occurs and payoffs realize. The analysis below does not solve this signaling game in full generality. It imposes, in Assumption 4, that the game admits a separating equilibrium, and it characterizes how the frontier moves across the on-path posteriors.

Let $U_j(z)$ be the type-j continuation value before liquidity suppliers update beliefs about front executions, and let $X_T(z)$ denote the additional payoff a toxic trader obtains from inducing more aggressive liquidity supply through path z.

Definition 3 (Separating high-credibility path)**.** The path H is separating if

$$U_B(H) - k_B(H) \geq U_B(L) \tag{43}$$

and

$$U_T(H) + X_T(H) - k_T(H) \leq U_T(L) + X_T(L), \tag{44}$$

with at least one inequality strict.

Inequality (43) says the liquidity type is willing to use the high-credibility path. Inequality (44) says the toxic type does not want to mimic it, even after accounting for the benefit $X_T(H)$ of relaxing the price-time frontier. The second condition is what makes the signal informative rather than decorative.

Assumption 4 (Order-path separation)**.** The signaling game admits a separating equilibrium: the high-credibility path satisfies Definition 3. On the equilibrium path, when $z = H$, passive liquidity suppliers assign the posterior associated with type B; when $z = L$, they assign the lower-credibility posterior.

7.2 From posteriors to the frontier

To link the posterior to the price-time frontier, let $\nu(\omega \mid s, z)$ denote the posterior distribution over front-execution states ω; the notation is distinct from the informed-flow probability $\pi(s)$ of Assumption 1. The queue-priority increment is evaluated under that posterior,

$$\Delta_m(p; s, z) = \mathbb{E}_{\nu(\cdot|s,z)} \left[\text{value of owning marginal front state } m \text{ at price } p\right], \tag{45}$$

so that the signed front block, toxic exposure, and improvement margin become signal dependent:

$$\Gamma^{\text{front}}(a - \delta, q; s, z) = \sum_{m \in B^{\text{front}}(a-\delta,q;s)} \Delta_m(a - \delta; s, z), \tag{46}$$

$$\Phi(s, z) = G^P(s, z) - \delta F^I(s, z) + \Gamma^{\text{front}}(a - \delta, q; s, z) - C^I(s, z), \tag{47}$$

with $T^{\text{front}}(a - \delta, q; s, z) = [-\Gamma^{\text{front}}(a - \delta, q; s, z)]_+$.

Assumption 5 (Posterior monotonicity)**.** Under the separating high-credibility path,

$$\Gamma^{\text{front}}(a-\delta, q; s, H) \geq \Gamma^{\text{front}}(a-\delta, q; s, L), \tag{48}$$

and

$$G^P(s,H) - \delta F^I(s,H) - C^I(s,H) \geq G^P(s,L) - \delta F^I(s,L) - C^I(s,L), \tag{49}$$

with at least one inequality strict on the state region of interest.

Assumption 5 is where the inference does its work: a credible liquidity-motive signal raises the posterior value of owning front execution states. Under the microfoundation of Lemma 3, inequality (48) holds whenever the H-posterior lowers the inferred adverse-selection intensities A_m or raises the noise-flow odds $\tau(\pi)$ on the transferred block. The second inequality (49) is permitted to hold with equality everywhere; nothing in the theorem requires the access channel to move. The economic content flows through the front channel, inequality (48), which is the limit-order-book translation of the dealer's lower adverse-selection cost.

Theorem 13 (Sunshine frontier shift)**.** *Under Assumptions 4 and 5,*

$$\Phi(s,H) \geq \Phi(s,L) \tag{50}$$

for every state s, with strict inequality wherever posterior monotonicity is strict. Hence the improvement region weakly expands,

$$\{s : \Phi(s,L) \geq 0\} \subseteq \{s : \Phi(s,H) \geq 0\}. \tag{51}$$

If in addition $V_H(s,H) \geq V_H(s,L)$, the active participation region also weakly expands under the high-credibility path.

The content of Theorem 13 is the mechanism transfer, not the calculus. Read as algebra, a monotone shift in Γ^{front} shifting Φ is immediate. Read as economics, the theorem maps each piece of the dealer-market sunshine mechanism into a price-time priority object: the dealer's adverse-selection cost becomes the toxic-front exposure T^{front}, the credible preannouncement becomes a separating order path, and the dealer's improved quote becomes an outward shift of the price-time frontier. The separation requirement is what blocks the manipulation objection. A toxic trader could try to manufacture the high-credibility path, but if the path is sufficiently costly or incompatible with toxic monetization, inequality (44) fails for the mimic and the posterior survives. Without the incentive condition, the result would collapse to a monotone comparative static and should not be called Sunshine.

Transparency is therefore not enough. Public order-path information shifts the frontier outward only when it credibly raises the signed posterior value of front priority; in toxic-front regions, only when it lowers inferred T^{front}. If the same visibility makes toxic mimicry easier, reveals exploitable execution demand, or lowers the value of bearing front priority, the separating condition fails and the frontier does not move, or moves inward.

8 Implications

This section records implications of the price-time frontier that either follow immediately from the signed priority value Δ_m or carry over from the companion queue-priority model. They are not separate headline results. Their role is to show what the new margin changes and to identify which parts of the analysis are inherited from the underlying priority valuation model.

8.1 Results inherited from queue-priority valuation

The companion queue-priority model shows that time priority can have negative value when early execution states are adversely selected. In the notation of this paper, Lemma 1 gives

$$V(p,1;s) - V(p,K+1;s) = \sum_{m=1}^{K} \Delta_m(p;s). \tag{52}$$

If the first K marginal priority states have negative total value, a trader prefers rank $K+1$ to the front rank at the same price. This is the depth-as-insurance result, and it explains why price improvement can be costly in the present paper: improving removes the insurance provided by depth ahead and transfers the signed block Γ^{front} to the improver.

The same logic applies to priority rules. A matching rule r assigns ownership weights $w^r_m(q;s)$ over marginal execution states, with rule-specific signed front-block value

$$\Gamma^{\text{front},r}(p,q;s) = \sum_m w^r_m(q;s)\,\Delta_m(p;s), \qquad T^{\text{front},r} = \left[-\Gamma^{\text{front},r}\right]_+ . \tag{53}$$

FIFO concentrates ownership of the earliest execution states; sharing rules distribute them. This does not imply a welfare ranking of FIFO and pro-rata. It isolates one channel: priority rules allocate adverse-selection exposure, not only fill probability.

The participation result inherits the companion paper's reposting logic. A trader exits when $V_O(s) > \max\{V_H(s), V_H(s)+\Phi(s)\}$. Because $\Phi = G^P - \delta F^I + \Gamma^{\text{front}} - C^I$, a lower signed front-block value weakly lowers the value of improving and can move states from improvement into holding or exit. Cancellation is therefore an endogenous participation response to adverse-selection exposure, not a separating statistic by itself.

8.2 Maker-taker fees

Maker-taker fees are the one implication that directly modifies the price-time frontier. Let r be a per-share rebate paid to executed passive liquidity at the improved quote. The improvement margin becomes

$$\Phi_r(s) = G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(s) - C^I(s) + rF^I(s) = \Phi(s) + rF^I(s). \tag{54}$$

Proposition 14 (Rebate frontier shift)**.** *If $F^I(s) > 0$, the trader improves under rebate r if and only if*

$$r \geq r^*(s) = \frac{\delta F^I(s) + C^I(s) - G^P(s) - \Gamma^{\text{front}}(s)}{F^I(s)}. \tag{55}$$

Moreover, the improve region $\{s : \Phi_r(s) \geq 0\}$ is weakly increasing in r.

The rebate is therefore a subsidy to bearing front exposure. In toxic-front states, where $\Gamma^{\text{front}} = -T^{\text{front}}$, the cutoff becomes

$$r^*(s) = \frac{\delta F^I(s) + T^{\text{front}}(s) + C^I(s) - G^P(s)}{F^I(s)}. \tag{56}$$

This does not imply that rebates are good or bad in welfare terms. It identifies the margin they move: the willingness of liquidity suppliers to own toxic front execution states.

8.3 Fine ticks and residual price-time economics

The fine-tick theorem implies that price-time economics remains active even when the mechanical tick concession vanishes. Holding is not automatic in the surviving region; it also requires $V_H(s) \geq V_O(s)$. Exit is not automatic either; it requires the outside option to dominate both active quote values.

Corollary 15 (Residual hold and exit regions)**.** *Under Assumption 3, the no-undercutting region has positive measure for all sufficiently small ticks. If $V_H(s) \geq V_O(s)$ on a positive-measure subset $A_H \subseteq A$, then the hold region has positive measure. If*

$$V_O(s) > \max\{V_H(s), V_H(s) + \Phi_n(s)\} \tag{57}$$

on a positive-measure subset $A_O \subseteq A$, then the exit region also has positive measure for all sufficiently small ticks.

The limiting margin is $\Phi_0 = G^P + \Gamma^{\text{front}} - C^I$: fine ticks eliminate the deterministic price concession but not signed front-priority exposure.

8.4 Empirical content

The model is pure theory, but its central objects have observable counterparts, and the two main results discipline what one should expect to see in data.

First, the fine-tick theorem predicts that undercutting intensity after a tick-size reduction should be state dependent in a specific way: undercutting should increase least, or not at all, in instruments and periods where front executions are most adversely selected, for example where the short-horizon markout of trades that execute immediately against a freshly improved quote is most negative. A purely mechanical tick story predicts undercutting wherever the tick was binding; the price-time frontier predicts a residual no-undercutting region indexed by front toxicity, not by the old tick.

Second, the model reinterprets cancellation. Because exit is the optimal participation response when both active quote values are dominated, cancellation rates should covary with measured front toxicity, holding fill rates fixed. Cancellation alone is therefore uninformative about intent, a point with regulatory relevance that this paper deliberately does not develop further.

Third, the sunshine theorem predicts that credible order-path geometry, such as regular slicing by identifiable liquidity traders, should be followed by tighter improvement behavior by liquidity suppliers, but only where mimicry is costly. Where the high-credibility path is cheap to imitate, the same geometry should produce no frontier shift.

Fourth, the maker-taker proposition predicts that rebate changes should move improvement behavior most in states with high measured front toxicity, because the rebate is effectively a subsidy to bearing T^{front}.

9 Conclusion

This paper studies how price improvement works when price priority and time priority jointly determine not only execution access but adverse-selection exposure. In a consolidated limit order book, a trader who improves one tick does not simply pay a smaller price concession to obtain faster execution. The trader acquires the front execution states at the improved price. If those states are toxic, improvement transfers adverse-selection exposure to the trader who becomes first in the queue.

The central object is the price-time frontier, $\Phi = G^P - \delta F^I + \Gamma^{\text{front}} - C^I$. Improvement is attractive when the non-front execution-access gain plus the signed value of the newly acquired priority block exceeds the tick concession and improvement friction. In toxic-front states, $\Gamma^{\text{front}} = -T^{\text{front}}$, so the same margin says that improvement must overcome the toxic-front cost of the acquired block. This one margin gives a common logic for no-undercutting, participation, cancellation, type sorting, and dynamic thresholds. Those are not separate frictions; they are different regions of the same price-time surface.

The fine-tick theorem is the paper's main discipline. Smaller ticks reduce the mechanical cost of price improvement, but they do not remove the value of the front execution states. Using the queue-priority primitive $\Delta_m = \pi b_m[\tau(\pi) - A_m]$, the paper shows that a sniping-dominated front block has negative signed value for reasons unrelated to tick size. If that primitive toxicity gap dominates the non-front access gain on a positive-measure region of states, the no-undercutting region does not vanish as the tick converges to zero. The result is local by design: it concerns the persistence of a state region in which rational liquidity suppliers decline to improve price, not an equilibrium spread floor.

The order-path credibility theorem translates a classic dealer-market disclosure mechanism into the book. In a dealer market, credible liquidity-motive disclosure lowers the dealer's adverse-selection cost. In a consolidated book, the analogous object is public credibility over the order path, and the translation requires an incentive condition: toxic traders must not find it profitable to mimic the high-credibility path. When separation holds, credibility raises the posterior signed value of front execution states, shifts the price-time frontier outward, and expands the region in which liquidity suppliers improve or remain active. Transparency is valuable only through this channel; it must change beliefs about the value of front priority in an incentive-compatible way.

The model is intentionally narrower than a full theory of the limit order book. It does not establish uniqueness of a dynamic equilibrium, rank priority rules by welfare, or identify manipulation. Its contribution is to isolate a price-time mechanism that the standard language of tick size and queue-jumping conceals. In a dealer market, adverse-selection exposure is priced through a spread. In a consolidated limit order book, that exposure is unbundled across queue positions. The dealer has not disappeared; the dealer's adverse-selection spread has become a price-time frontier.

References

Anat R. Admati and Paul Pfleiderer. A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies*, 1(1):3–40, 1988.

Anat R. Admati and Paul Pfleiderer. Sunshine trading and financial market equilibrium. *The Review of Financial Studies*, 4(3):443–481, 1991.

Anonymous. Queue priority as adverse-selection exposure. Working paper. Reference anonymized for review, 2026.

Robert Battalio, Shane A. Corwin, and Robert Jennings. Can brokers have it all? on the relation between make-take fees and limit order execution quality. *The Journal of Finance*, 71(5):2193–2238, 2016. doi: 10.1111/jofi.12422.

Eric Budish, Peter Cramton, and John Shim. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4): 1547–1621, 2015.

Jean-Edouard Colliard and Thierry Foucault. Trading fees and efficiency in limit order markets. *The Review of Financial Studies*, 25(11):3389–3421, 2012. doi: 10.1093/rfs/hhs089.

Thierry Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets*, 2(2):99–134, 1999.

Thierry Foucault, Ohad Kadan, and Eugene Kandel. Limit order book as a market for liquidity. *The Review of Financial Studies*, 18(4):1171–1217, 2005.

Lawrence R. Glosten. Is the electronic open limit order book inevitable? *The Journal of Finance*, 49(4):1127–1161, 1994.

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

Ronald L. Goettler, Christine A. Parlour, and Uday Rajan. Equilibrium in a dynamic limit order market. *The Journal of Finance*, 60(5):2149–2192, 2005.

Peter Hoffmann. A dynamic limit order market with fast and slow traders. *Journal of Financial Economics*, 113(1):156–169, 2014.

Burton Hollifield, Robert A. Miller, and Patrik Sandås. Empirical analysis of limit order markets. *The Review of Economic Studies*, 71(4):1027–1063, 2004.

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

Katya Malinova and Andreas Park. Subsidizing liquidity: The impact of make/take fees on market quality. *The Journal of Finance*, 70(2):509–536, 2015. doi: 10.1111/jofi.12230.

Albert J. Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.

Christine A. Parlour. Price dynamics in limit order markets. *The Review of Financial Studies*, 11(4):789–816, 1998.

Ioanid Rosu. A dynamic model of the limit order book. *The Review of Financial Studies*, 22 (11):4601–4641, 2009.

Patrik Sandås. Adverse selection and competitive market making: Empirical evidence from a pure limit order market. *The Review of Financial Studies*, 14(3):705–734, 2001.

Chen Yao and Mao Ye. Why trading speed matters: A tale of queue rationing under price controls. *The Review of Financial Studies*, 31(6):2157–2183, 2018. doi: 10.1093/rfs/hhy002.

A Proofs

Proof of Lemma 1. By Definition 1, $\Delta_m(p;s) = V(p,m;s) - V(p,m+1;s)$. Summing over $m = 1,\dots,q$,

$$\sum_{m=1}^{q} \Delta_m(p;s) = \sum_{m=1}^{q} [V(p,m;s) - V(p,m+1;s)] = V(p,1;s) - V(p,q+1;s),$$

since the sum telescopes. □

Proof of Lemma 2. Write $V_I(s) = \widetilde{V}_I(s) - C^I(s)$, where $\widetilde{V}_I(s)$ is the gross-of-cost value of the improved quote. By hypothesis, $\widetilde{V}_I(s) - V_H(s)$ operates through three channels only. First, the trader acquires the transferred front block $B^{\text{front}}(a-\delta,q;s)$; by Definition 1 applied at the improved price, the signed value of the marginal front states in that block is $\sum_{m\in B^{\text{front}}(a-\delta,q;s)} \Delta_m(a-\delta;s) = \Gamma^{\text{front}}(a-\delta,q;s)$, as in equation (7). Second, the trader gains access to the non-front fill states $B^{\text{access}}(a-\delta,q;s)$, each valued at its incremental surplus

$g_m(a-\delta; s) \geq 0$ relative to holding at a, summing to $G^P(s)$ as in equation (9). Third, every expected fill unit at the improved quote executes at $a-\delta$ rather than a, a price concession of δ per unit on $F^I(s)$ expected units, contributing $-\delta F^I(s)$. Summing the three channels and subtracting the cost,

$$\Phi(s) = V_I(s) - V_H(s) = G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(a-\delta, q; s) - C^I(s).$$

For the no-double-counting claim, observe that the decomposition is applied to the difference $V_I - V_H$, so any exposure common to both actions, in particular the adverse-selection exposure of the current quote at a, cancels and remains embedded in the level $V_H(s)$. The front block entering Γ^{front} consists exclusively of execution states at the improved price that the trader does not own under H, so no exposure is counted twice. □

Proof of Lemma 3. Under Assumption 1, marginal front state m is reached with probability $p_m(s) > 0$. Conditional on being reached, the state yields the gain $\beta_m(s)$ when the executing flow is noise, which occurs with probability $1-\pi(s)$, and the loss $A_m(s)\beta_m(s)$ when the executing flow is informed, which occurs with probability $\pi(s)$. Hence, with $b_m(s) = p_m(s)\beta_m(s)$,

$$\begin{aligned}
\Delta_m(p; s) &= p_m(s)\left[(1-\pi(s))\,\beta_m(s) - \pi(s)\,A_m(s)\,\beta_m(s)\right] \\
&= (1-\pi(s))\,b_m(s) - \pi(s)\,A_m(s)\,b_m(s) \\
&= \pi(s)\,b_m(s)\left[\frac{1-\pi(s)}{\pi(s)} - A_m(s)\right] = \pi(s)\,b_m(s)\left[\tau(\pi(s)) - A_m(s)\right].
\end{aligned}$$

Since $\pi(s) \in (0,1)$ and $b_m(s) > 0$, the sign of Δ_m is the sign of $\tau(\pi(s)) - A_m(s)$, so $\Delta_m(p; s) < 0$ if and only if $A_m(s) > \tau(\pi(s))$. □

Proof of Proposition 4. The trader holds when $V_H(s) \geq V_I(s)$. Since $\Phi(s) = V_I(s) - V_H(s)$ by equation (10), this is equivalent to $\Phi(s) \leq 0$. Substituting the decomposition of Lemma 2 gives

$$G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(a-\delta, q; s) - C^I(s) \le 0,$$

which rearranges to inequality (15). The statement for improvement is the reverse inequality. □

Proof of Corollary 5. Suppose inequality (17) holds. For any $\delta F^I(s) \ge 0$,

$$\Phi(s) = G^P(s) + \Gamma^{\text{front}}(a-\delta, q; s) - C^I(s) - \delta F^I(s) < 0 - \delta F^I(s) \le 0,$$

so holding strictly dominates improvement for every nonnegative tick concession, including a tick concession of zero. □

Proof of Proposition 6. Holding is optimal when V_H is at least the outside option and at least the improvement value $V_H + \Phi$, which is equivalent to $V_H \ge V_O$ and $\Phi \le 0$. Improving is optimal when $V_H + \Phi$ is at least the outside option and at least V_H, which is equivalent to $V_H + \Phi \ge V_O$ and $\Phi \ge 0$. Exit is optimal when V_O exceeds both active quote values, which is conditions (21). □

Proof of Proposition 7. The partition follows from Proposition 6 applied type by type. Since C^I enters Φ in equation (23) negatively and Γ^{front} enters positively, increasing improvement cost or decreasing the signed front-block value weakly lowers Φ and therefore weakly shrinks the set where $\Phi \ge 0$. In a toxic-front region, $\Gamma^{\text{front}} = -T^{\text{front}}$, so increasing T^{front} has the same effect. Increasing reposting friction K^O lowers $V_O(\theta) = -K^O(\theta)$ and therefore weakly shrinks the set where exit is optimal. □

Proof of Lemma 8. By the representation (30) in condition (ii) of Assumption 3, for all

$n \geq N_0$ and $s \in A$,

$$\Gamma_n^{\text{front}}(s) = \sum_{m \in B_n^{\text{front}}(s)} \pi_n(s)\, b_{m,n}(s) \left[\tau(\pi_n(s)) - A_{m,n}(s)\right]$$
$$= - \sum_{m \in B_n^{\text{front}}(s)} \pi_n(s)\, b_{m,n}(s) \left[A_{m,n}(s) - \tau(\pi_n(s))\right].$$

Inequality (31) states that the final sum is at least $\underline{T}$. Therefore $\Gamma_n^{\text{front}}(s) \leq -\underline{T}$, and $T_n^{\text{front}}(s) = [-\Gamma_n^{\text{front}}(s)]_+ \geq \underline{T}$. □

Proof of Theorem 9. Fix $n \geq N_0$ and $s \in A$. By Lemma 8, $\Gamma_n^{\text{front}}(s) \leq -\underline{T}$. By condition (i) of Assumption 3, $G_n^P(s) \leq \bar{G}$. By Assumption 2, $C_n^I(s) \geq 0$ and $F_n^I(s) \geq 0$, so $\delta_n F_n^I(s) \geq 0$. Hence

$$\Phi_n(s) = G_n^P(s) - \delta_n F_n^I(s) + \Gamma_n^{\text{front}}(s) - C_n^I(s) \leq \bar{G} - 0 - \underline{T} - 0 = -(\underline{T} - \bar{G}) < 0,$$

using $\bar{G} < \underline{T}$ from Assumption 3. This is the uniform bound (33); the bound does not depend on n or s. In particular $\Phi_n(s) \leq 0$, so $s \in N_n$. Since $s \in A$ was arbitrary, $A \subseteq N_n$, and therefore $\mu(N_n) \geq \mu(A) > 0$ for all $n \geq N_0$. □

Proof of Corollary 10. Since $\delta_n \to 0$ and $F_n^I(s)$ is locally bounded, $\delta_n F_n^I(s) \to 0$. Taking limits in equation (27) along the assumed pointwise limits gives $\Phi_n(s) \to G^P(s) + \Gamma^{\text{front}}(s) - C^I(s) = \Phi_0(s)$. The limiting no-undercutting region is the set where $\Phi_0(s) \leq 0$, which is the first set in equation (36); the second form follows from $\Gamma^{\text{front}} = -T^{\text{front}}$ on the toxic-front region. □

Proof of Proposition 11. For fixed stationary rival policy α_j and bounded u, the Bellman operator

$$(\mathcal{T}W)(s, \theta) = \max_{a \in \{H,I,O\}} \left\{ u(s, a; \theta, \alpha_j(s)) + \beta\, \mathbb{E}[W(s', \theta) \mid s, a, \alpha_j(s)] \right\}$$

maps bounded functions into bounded functions and satisfies

$$\|\mathcal{T}W - \mathcal{T}W'\|_\infty \leq \beta \|W - W'\|_\infty,$$

since the maximum and the conditional expectation are both nonexpansive. With $\beta \in (0, 1)$, the Banach fixed-point theorem gives a unique fixed point W. At the fixed point, optimality is exactly maximization over the three action values, which is the characterization (39). □

Proof of Corollary 12. Under the structure (40),

$$\Phi^D(x, y, \theta) = [A_I(y, \theta) - A_H(y, \theta)] - [i(\theta) - h(\theta)]\, x,$$

which is strictly decreasing in x because $i(\theta) > h(\theta)$. Hence $\Phi^D \geq 0$ exactly on $x \leq \bar{x}_I(y, \theta)$ with $\bar{x}_I$ as in equation (41). In the hold region, holding beats exit exactly when $A_H(y, \theta) - h(\theta)x \geq -K^O(\theta)$, that is, $x \leq \bar{x}_O(y, \theta)$ for $h(\theta) > 0$. If $\bar{x}_I \leq \bar{x}_O$, then for $x < \bar{x}_I$ improvement dominates holding and, a fortiori, exit; for $\bar{x}_I \leq x \leq \bar{x}_O$ holding dominates improvement and exit; and for $x > \bar{x}_O$ exit dominates both. This is the ordered policy (42). □

Proof of Theorem 13. Assumption 4 implies that H is a credible liquidity-motive signal: by inequality (43) the liquidity type is willing to choose it, and by inequality (44) the toxic type is not willing to mimic it. Liquidity suppliers therefore condition front-execution values on the H-posterior after observing H. Subtracting the improvement margins in equation (47),

$$\begin{aligned}\Phi(s, H) - \Phi(s, L) = &\left[G^P(s, H) - \delta F^I(s, H) - C^I(s, H)\right] \\ &- \left[G^P(s, L) - \delta F^I(s, L) - C^I(s, L)\right] \\ &+ \left[\Gamma^{\text{front}}(a - \delta, q; s, H) - \Gamma^{\text{front}}(a - \delta, q; s, L)\right].\end{aligned}$$

Both bracketed differences are nonnegative by Assumption 5, so $\Phi(s, H) \geq \Phi(s, L)$, with strict inequality when either channel is strict. The set inclusion (51) follows immediately. Active participation is determined by $\max\{V_H(s,z), V_H(s,z)+\Phi(s,z), V_O(s)\}$; if $V_H(s,H) \geq V_H(s,L)$, both active values weakly increase under H while the outside option is unchanged, so the active region cannot shrink. □

Proof of Proposition 14. Improvement is optimal against holding when $\Phi_r(s) \geq 0$. Since $F^I(s) > 0$, equation (54) gives

$$G^P(s) - \delta F^I(s) + \Gamma^{\text{front}}(s) - C^I(s) + rF^I(s) \geq 0 \iff r \geq r^*(s)$$

with $r^*(s)$ as in equation (55). For monotonicity, if $r' \geq r$ then $\Phi_{r'}(s) - \Phi_r(s) = (r'-r)F^I(s) \geq 0$, so any state in the improve region under r remains in the improve region under r'. □

Proof of Corollary 15. The first statement is Theorem 9. If $s \in A_H$, then for all sufficiently large n, $\Phi_n(s) \leq 0$ by the theorem and $V_H(s) \geq V_O(s)$ by assumption, so Proposition 6 implies holding on A_H. If inequality (57) holds on A_O, Proposition 6 implies exit on A_O. Since $\mu(A_H) > 0$ and $\mu(A_O) > 0$, the corresponding regions have positive measure. □