

Order Flow, Liquidation Risk, and Margin

Aryan Ayyar

Manipal Academy of Higher Education, Bengaluru

Why should margin change after order flow has already moved the price? A prime broker cares about something the price need not reveal: how much informed trading remains if the position must be closed. In a Kyle-style model with staggered execution, markets with the same price-update variance can have different margin sensitivity to signed order-flow surprises. That sensitivity depends on the broker's inference weight on order flow and the persistence of informed execution. In the benchmark, a one-standard-deviation adverse surprise raises the per-unit requirement by about 23%. Binding requirements induce forced position reductions.

JEL classification: G14, G23, G24, G33.

Keywords: adverse selection, margin, funding liquidity, prime brokerage, price discovery, forced liquidation.

26 September 2026

1. Introduction

A price can be informative about value and still say little about the cost of unwinding a particular position. That distinction matters for a prime broker. Once a leveraged position has settled, the broker no longer asks what the position is worth at the current price; it asks what the account would lose if the position had to be closed. For a long position, closing means selling into the orders that remain. For a short, it means buying back. We study how information in order flow enters that liquidation problem and, through it, margin.

The model separates the information in the date-0 price from the trading that remains after date 0. There are N informed accounts and transitory hedgers. Dealers observe aggregate order flow and set the date-0 price. Informed accounts, however, trade through finite execution programs, so part of their desired quantity is still waiting to be executed after the price has been set. A prime broker observes its own client's order as well as aggregate flow. It therefore has a different inference problem from the dealer: conditional on what it knows about its client, what does the remaining order flow say about the trades that other informed accounts will still submit during a closeout?

This distinction produces the paper’s main result. Hold the variance of the date-0 price update fixed and change the composition of order flow by making private signals noisier while thinning transitory hedging. The price reveals the same amount of information in variance terms, but the broker puts more weight on aggregate flow when forecasting other informed orders. At the same time, informed execution becomes more persistent. If K is the broker’s Bayesian weight on the order-flow innovation and ψ is execution persistence, the local response of margin to a signed innovation is simply ψK . Both terms rise along the comparison.

The result is easiest to read as a statement about what price discovery leaves out. Price-update variance records how much information entered the price. It does not record who generated that information or how much of their trading is still to come. Those features matter to a broker that may have to unwind a position. Two markets can therefore display the same price-update variance and yet produce different funding consequences after an equally adverse order-flow surprise. In the benchmark used in the paper, $\psi K \approx .29$ and the innovation standard deviation is about .563. A one-standard-deviation adverse innovation raises the per-unit requirement by about .163 from a zero-innovation level of .705, or roughly 23%.

Margin levels need not move monotonically with the informativeness of prices. In a separate experiment, reducing transitory hedging can lower residual uncertainty enough to reduce the required buffer even as expected liquidation costs rise. That possibility matters because it separates the paper’s central composition result from the broader claim that more informative trading must always raise margin. The robust object here is the loading of margin on signed order-flow news.

When equity falls short of the requirement, the broker cancels the client’s remaining execution program and reduces the position. The smallest feasible reduction trades off two forces: selling releases margin on the position that remains, but the sale itself consumes equity through execution costs. On the interior branch, a more adverse order-flow innovation therefore increases the forced quantity. The same quantity can matter beyond the account. If liquidity-provider execution costs rise with inventory, a larger forced reduction leaves more inventory with those providers and raises the cost of a fixed subsequent trade before the inventory is reallocated. We treat this last step as a conditional implication of the execution technology, not as a separate equilibrium theory of market depth.

The paper sits between market microstructure and funding-liquidity theory. In Kyle (1985) and Glosten and Milgrom (1985), order flow reveals private information to the market and moves prices. We ask what the same flow reveals to an intermediary that sees an additional piece of information—its own client’s order—and cares about liquidation rather than only current value. This timing also distinguishes the model from work in which financing constraints shape information acquisition before trading, such as Glebkin, Gondhi, and Kuong

(2021), and from settings in which a financing price is itself a public signal, as in Nezafat and Schroder (2022).

The connection to funding liquidity is complementary. Brunnermeier and Pedersen (2009) and Gârleanu and Pedersen (2011) link margins and constrained balance sheets to positions and market liquidity; Brunnermeier and Pedersen (2005) studies strategic trading around distressed investors. Here the new object is the composition of order flow behind a settled price. The broker uses that composition to assess position-specific liquidation risk and then applies a margin rule. Work on risk-sensitive margins and procyclicality provides the institutional benchmark (Murphy, Vasios, and Vause 2014; Biais, Heider, and Hoerova 2016; Glasserman and Wu 2018).

The model is intentionally partial equilibrium after date 0. Continuation trades execute under a prearranged temporary-impact technology and do not generate another round of public price discovery; the margin rule covers benchmark-relative execution losses rather than the full mark-to-market risk of the position. These assumptions keep the analysis on one question: whether information that has already moved the market price can still matter for financing because liquidation depends on trading that remains.




Figure 1
Order-flow information, margin, and forced liquidation.

Section 2 sets out the model. Section 3 derives the information and margin results. Sections 4 and 5 study forced reductions and their execution consequences. Section 6 discusses empirical implementation.

2. The Economy and Timing

The model has one round of price discovery and a later liquidation window. Dealers set the date-0 price from aggregate order flow. After that trade settles,

one prime broker observes its client's order and uses it to reinterpret the same aggregate flow. The broker now cares about the cost of closing the client's position if funding is withdrawn, not about re-estimating the asset's current value.

A risky claim pays v at date 2. Date 0 is the pricing date, date 1 is the margin reset, and $J \geq 1$ execution slices lie between the reset and maturity. The riskless numeraire has gross return one. No new public fundamental signal arrives between dates 0 and 1.

2.1 Participants and Commitments

There are $N \geq 1$ informed accounts. Account i observes s_i and begins a finite execution program. Signals and transitory hedging are Gaussian:

$$v \sim N(0, \sigma_v^2), \quad s_i = v + \eta_i, \quad \eta_i \stackrel{\text{iid}}{\sim} N(0, \sigma_\eta^2), \quad u_0 \sim N(0, \sigma_u^2), \quad (2.1)$$

The shocks in (2.1) are mutually independent and have strictly positive variances. The hedging order u_0 is exhausted at date 0; informed accounts may continue trading afterward.

Broker i observes its client's entering position q_0^i and equity e_0^i . These balance-sheet states carry no additional information about the Gaussian primitives. Initial finance supports the linear order $x_{i0} = \beta s_i$, with common $\beta > 0$; the execution problem below determines both β and the remaining schedule.

Execution access and financing are arranged before signals arrive. Other informed accounts keep their schedules through the window. The focal broker alone has a reset right: if its client fails the margin test, the remaining client schedule is canceled and replaced by the broker-directed reduction described below.

Competitive risk-neutral dealers observe aggregate initial flow and set

$$y_0 = \sum_{i=1}^N x_{i0} + u_0, \quad p_0 = \mathbb{E}[v | y_0] = \kappa y_0. \quad (2.2)$$

The dealers clear only the date-0 market. Later client trades use a prearranged temporary-impact technology, while designated liquidity providers execute any broker-directed reduction. Thus the date-0 pricing problem and the later closeout service are distinct.

Before signals arrive, the broker fixes uniform closeout weights $r_j = 1/J$. The reset determines only how much to liquidate; the weights determine how that quantity is spread across the J slices. Table 1 gives the timing.

2.2 Finite Execution and Initial Pricing

Each informed account chooses a target and splits it across $J+1$ trades. Temporary impact favors gradual execution; delay favors trading sooner, as in finite-horizon execution models (Bertsimas and Lo 1998; Almgren and Chriss

Table 1
Timing, information, and commitments

Stage	Information and commitments	Decision or outcome
Before signals	Common execution parameters, program finance, closeout access, and broker weights $r_j = 1/J$ are fixed.	Peers commit to continuation; the focal program remains subject to its broker's reset override.
Date 0	Each account observes s_i and chooses its target and complete planned schedule. Dealers observe only y_0 .	Initial orders execute at p_0 . Settlement determines the pre-reset position and equity. Hedgers exit.
Date 1	The focal broker observes its client's order and account state, together with (y_0, p_0) .	It calibrates and locks the pre-reset-position-size margin rate, then authorizes ordinary continuation or cancels the mandate and chooses a reduction.
Slices 1, ..., J	Peer orders and execution disturbances reach the venue. Chosen quantities and schedules are committed.	Peers execute; the focal account either continues its remaining client execution program or completes its selected reduction. Actual costs change cash, debt, and collateral.
End of window; date 2	The last slice completes all authorized trades. The remaining claim subsequently pays v .	Any collateral breach enters workout; remaining positions settle at maturity, with creditor recoveries determined from actual resources.

2001). Let x_{it} denote trade t and $\mathcal{R}_{it} = T_i - \sum_{h=0}^t x_{ih}$ the remaining target. The reduced-form objective is

$$\begin{aligned}
 \max_{T_i, x_{i0}, \dots, x_{iJ}} \quad & \alpha_i T_i - \frac{\gamma}{2} T_i^2 - \sum_{t=0}^J \left(\kappa x_{it}^2 + \frac{\gamma}{2J} \mathcal{R}_{it}^2 \right), \\
 \text{subject to} \quad & \sum_{t=0}^J x_{it} = T_i, \quad \gamma \geq 0, \quad \kappa > 0.
 \end{aligned} \tag{2.3}$$

The entire schedule is chosen before y_0 is observed and must complete the target by slice J . The parameter γ governs both target curvature and the cost of waiting. This is an execution benchmark, not a repeated Kyle game: once date 0 has cleared, p_0 is taken as given and continuation trades face temporary impact without generating another public-price update. That restriction lets the model isolate how much informed trading remains after price formation.

The linear benefit from the target is

$$\alpha_i = \rho [1 - \kappa \beta (N - 1)] s_i, \quad \rho = \frac{1}{1 + r}, \quad r = \frac{\sigma_\eta^2}{\sigma_v^2}. \tag{2.4}$$

Each account takes (β, κ) as given when solving its execution problem. Equilibrium requires the implied initial order to agree with dealer pricing. We

use the same κ for date-0 price impact and temporary execution cost; this normalization ties the pace of execution to the price-impact parameter and is part of the benchmark.

Proposition 1 (Finite execution and pricing). For $\gamma > 0$, define $\theta > 0$ by $\cosh \theta = 1 + \gamma / (4J\kappa)$. The unique solution of (2.3) has

$$g_j = \frac{\cosh((J + \frac{1}{2} - j)\theta)}{\cosh((J + \frac{1}{2})\theta)}, \quad j = 1, \dots, J, \quad (2.5)$$

$$\ell_0 = \left(1 + \sum_{j=1}^J g_j\right)^{-1}, \quad T_i = \frac{\alpha_i}{\gamma + 2\kappa\ell_0}, \quad x_{i0} = \ell_0 T_i, \quad x_{ij} = g_j x_{i0}.$$

Writing $\ell_j = \ell_0 g_j$ for $j \geq 1$, all weights are positive and the $J+1$ trades sum to T_i . For $\gamma = 0$, these expressions have the finite limit $\ell_t = 1/(J+1)$, $g_j = 1$, and $T_i = (J+1)\alpha_i/(2\kappa)$. For every $\gamma \geq 0$, there is a unique positive pair (β, κ) satisfying

$$\beta[\kappa\{2 + \rho(N-1)\} + \Gamma_J] = \rho, \quad \Gamma_J = \frac{\gamma}{\ell_0}, \quad (2.6)$$

$$\kappa = \frac{N\beta\sigma_v^2}{\beta^2\sigma_v^2 N(N+r) + \sigma_u^2}.$$

Its initial orders are $x_{i0} = \beta s_i$.

Appendix B.1 gives the recursion and uniqueness proof. The schedule has the familiar trade-off: impact spreads the order, while delay pulls it forward. The statistic that matters for liquidation is the average continuation ratio

$$\psi = \frac{1}{J} \sum_{j=1}^J g_j = \frac{\ell_0^{-1} - 1}{J}. \quad (2.7)$$

With positive delay costs, $0 < \psi < 1$; with zero delay costs, $\psi = 1$. An initial order x_{i0} therefore predicts $J\psi x_{i0}$ of additional execution. We use ψ as the measure of order-flow persistence.

2.3 Settlement and Broker Information

Let B_1^i denote net cash after the initial order. Settlement gives

$$B_1^i = e_0^i - p_0 x_{i0}, \quad q_1^i = q_0^i + x_{i0}, \quad Q^i = |q_1^i|, \quad (2.8)$$

$$e^i = B_1^i + p_0 q_1^i = e_0^i + q_0^i p_0.$$

Because the initial trade is marked at its execution price, its cash payment and marked value offset. Settlement leaves e^i as the equity available at the reset.

The broker sees more than the dealer. Since $x_{i0} = \beta s_i$ and β is known, its client's order reveals s_i . Let $\mathcal{G}_i$ collect the broker's client information before it observes aggregate flow. The relevant information sets are

$$\mathcal{D}_0 = \sigma(y_0), \quad \mathcal{G}_i = \sigma(q_0^i, e_0^i, x_{i0}), \quad \mathcal{H}_i = \mathcal{G}_i \vee \sigma(y_0, p_0). \quad (2.9)$$

The broker does not observe other accounts' orders. Its own client's signal nevertheless predicts them through the common value. Aggregate flow supplies the second signal. Because $p_0 = \kappa y_0$, price and flow contain the same public information; what is distinctive is the broker's private conditioning information.

Suppress the focal index on its position and equity, and write $Q = |q_1|$ and $\varsigma = \text{sgn}(q_1)$ for $Q > 0$. Define the position-signed peer-order exposure

$$P_i = \sum_{\ell \neq i} x_{\ell 0}, \quad \Theta_i = -\varsigma \kappa P_i, \quad A_i = \mathbb{E}[\Theta_i | \mathcal{H}_i]. \quad (2.10)$$

The sign convention makes Θ_i adverse for either position direction: peer selling hurts a long liquidation, while peer buying hurts a short cover. Section 3 derives the posterior for this object.

2.4 Liquidation Execution

A forced reduction executes against a precommitted service schedule. Let $\lambda_0 \geq 0$ denote the initial cost state and $d \geq 0$ its loading. If the broker chooses a reduction $x \in [0, Q]$, the per-unit concession in slice j is

$$U_j(x) = d\lambda_0 + \eta_c R_{j-1}^r x + \nu_c r_j x + g_j \Theta_i + \epsilon_j, \quad (2.11)$$

$$p_j(x) = p_0 - \varsigma U_j(x), \quad R_{j-1}^r = \sum_{h < j} r_h, \quad r_j = \frac{1}{J}.$$

The coefficients η_c and ν_c capture inventory already absorbed from the reduction and the size of the current slice. The term $d\lambda_0$ shifts the cost level. These coefficients are exogenous to the date-0 pricing block.

The complete execution-disturbance vector satisfies

$$\epsilon \sim N(\mathbf{0}, \Omega), \quad \Omega \geq 0, \quad \epsilon \perp (v, \eta, u_0, \text{entering account states}). \quad (2.12)$$

Execution shocks arrive after the margin decision. Peer flow is common across slices through the continuation schedule, so uncertainty about other informed accounts creates correlated closeout costs. A sufficiently favorable state can produce a negative concession relative to p_0 .

Once the broker selects D , the reduction is executed across all J slices. Let $U(D) = J^{-1} \sum_j U_j(D)$ denote the average concession and $L = DU(D)$ the total execution bill. A provider that absorbs signed quantity $\varsigma r_j D$ receives

$$\varsigma r_j D \{v - p_j(D)\} = \varsigma r_j D (v - p_0) + r_j D U_j(D). \quad (2.13)$$

The providers keep the acquired asset exposure. Margin is therefore computed against execution losses relative to p_0 ; any subsequent move in the fundamental value is accounted for later in terminal resources.

2.5 The Margin Reset and Account Books

The broker sets margin against the expected uncovered cost of liquidating the full pre-reset position. Given an execution-loss limit $\bar{\ell} > 0$, it chooses the smallest nonnegative rate m such that

$$\mathbb{E}[(QU(Q) - mQ)^+ | \mathcal{H}_i] \leq \bar{\ell}, \quad m \geq 0. \quad (2.14)$$

The rate is calibrated at the entering position size Q and cost state λ_0 . Once set, it is held fixed through the liquidation window; after a reduction D , required collateral is $m(Q - D)$. Section 3 solves for m .

If $e \geq mQ$, the account continues. If $0 < e < mQ$, the client schedule is canceled and the broker chooses the smallest reduction that restores the condition in expectation,

$$e - D\mathbb{E}[U(D) | \mathcal{H}_i] \geq m(Q - D).$$

If no regular interior solution exists, the position is fully wound down. Section 4 derives this rule and then distinguishes the ex ante margin calculation from realized collateral shortfalls.

Let C indicate continuation. For a positive position, $C = 1$ when $e \geq mQ$ and $C = 0$ otherwise; at $Q = 0$, set $m = D = 0$. Ordinary trades are $x_j^P = Cg_jx_{i0}$. A positive reduction replaces the remaining client program with forced orders $f_j = -\varsigma r_j D$.

For bookkeeping, separate the settled position from any continuation program. The latter begins with zero holdings and zero drawn finance. If continuation is authorized, trade x_j^P executes at $p_0 + \kappa x_j^P$, so its cumulative financing and benchmark-marked equity are

$$\begin{aligned} q_j^P &= \sum_{h \leq j} x_h^P, & H_j^P &= \sum_{h \leq j} (p_0 + \kappa x_h^P) x_h^P, \\ E_j^P &= -H_j^P + p_0 q_j^P = -\kappa \sum_{h \leq j} (x_h^P)^2. \end{aligned} \quad (2.15)$$

Positive H_j^P is program debt and negative H_j^P is program cash. The continuation book finances its own principal and execution costs and remains subordinate to the settled position during the window.

Let $L_j = \sum_{h \leq j} r_h D U_h(D)$ be the cumulative forced invoice, defined as zero directly whenever $D = 0$. Pre-reset marked equity is $E_j^I = e - L_j$. The two books then satisfy

$$E_j^I + E_j^P = e - L_j - \kappa \sum_{h \leq j} (x_h^P)^2. \quad (2.16)$$

Equation (2.16) makes the accounting transparent: liquidation invoices reduce the settled book, while ordinary execution costs sit in the continuation book. At maturity, write $q^P = q_J^P$, $C_{\text{exec}} = \kappa \sum_j (x_j^P)^2$, and $L = L_J$. Total contractual resources are

$$\begin{aligned} q_2 &= q_1 - \varsigma D + q^P, & Q_2 &= |q_2|, \\ N &= e - L - C_{\text{exec}} + q_2(v - p_0), \end{aligned} \quad (2.17)$$

where ςD is zero directly at $Q=0$. The full program draw is included, and ordinary trading may cross the pre-reset sign. Actual recoveries and the client's limited-liability payout follow book priorities and settlement resources.

Collateral remains the client's asset even when it is restricted. Let M_{j-1} denote the amount tied up before slice j and Δ_j the corresponding time weight, with $\sum_j \Delta_j = 1$. Its opportunity cost is $k \sum_j \Delta_j M_{j-1}$. Sections 4 and 5 track realized collateral, payments, and recoveries.

The rest of the paper revolves around five objects: K , the broker's weight on the order-flow innovation; ψ , execution persistence; z , the innovation signed against the position; m , the margin rate; and D , the forced reduction. Appendix A collects the remaining notation. We next derive how order flow maps into m .

3. Order Flow, Liquidation Risk, and Margin

The broker's problem has two steps. Its client's order predicts what other informed accounts are likely to have traded, and aggregate flow tells the broker how far actual peer trading departed from that prediction. The second step would be redundant if the broker knew no more than the dealer. It matters here because the broker sees its own client's order.

That inference is then multiplied by persistence. Order flow that reveals a large peer position matters more for liquidation when much of that position is still scheduled to trade. The product ψK therefore governs the response of margin to signed order-flow news.

3.1 Client Information and the Order-Flow Innovation

Let $n = N - 1$. From its client's order, the broker recovers $s_i = x_{i0}/\beta$. Before looking at aggregate flow, its conditional distribution for the other n informed orders is

$$\mu_i = \beta n \rho s_i, \quad V_i = \beta^2 \left(n^2 \frac{\sigma_v^2 \sigma_\eta^2}{\sigma_v^2 + \sigma_\eta^2} + n \sigma_\eta^2 \right). \quad (3.1)$$

The first term in V_i comes from the common value and the second from idiosyncratic signal noise. The broker then removes its client's order from total flow and observes

$$Y_i = y_0 - x_{i0} = P_i + u_0, \quad \iota_i = Y_i - \mu_i, \quad \iota_i | \mathcal{G}_i \sim N(0, V_i + \sigma_u^2). \quad (3.2)$$

Thus ι_i is the part of client-netted order flow that the broker did not predict from its own customer.

Proposition 2 (The broker's posterior). The peer-order distribution conditional on the full reset information $\mathcal{H}_i$ is

$$P_i | \mathcal{H}_i \sim N(\mu_i + K_i \iota_i, R_i), \quad (3.3)$$

$$K_i = \frac{V_i}{V_i + \sigma_u^2}, \quad R_i = \frac{V_i \sigma_u^2}{V_i + \sigma_u^2}.$$

For a pre-reset position of size $Q > 0$ and direction ς , define

$$A_i^0 = -\varsigma \kappa \mu_i, \quad z_i = -\varsigma \kappa \iota_i, \quad A_i = A_i^0 + K_i z_i. \quad (3.4)$$

Then $\Theta_i | \mathcal{H}_i \sim N(A_i, \kappa^2 R_i)$. Within this economy, a unit increase in the observed z_i raises expected adverse pressure by K_i and leaves its conditional variance unchanged.

The updating rule is standard Gaussian projection, but the conditioning set is not the dealer's. With two accounts, $\beta = 1$, and unit primitive variances, the broker forecasts peer flow as $0.2s_i + 0.6Y_i$, with residual variance 0.6. If $s_i = 2$, then $Y_i = 0$ and $Y_i = 1$ imply forecasts of 0.4 and 1.0; because $\mu_i = 1$, the corresponding innovations are -1 and 0 . A negative surprise is adverse for a long because it predicts more peer selling; the sign reverses for a short.

The same logic gives a simple information-coverage result. Suppose the broker observes n_o accounts, including its own, while $n_u = N - n_o$ remain hidden. Holding trading policies fixed, the variance of unobserved peer orders and the weight placed on aggregate flow are

$$V(n_o) = \beta^2 \sigma_v^2 \left(\frac{n_u^2 r}{r + n_o} + n_u r \right), \quad K(n_o) = \frac{V(n_o)}{V(n_o) + \sigma_u^2}. \quad (3.5)$$

Coverage substitutes for public inference. As the broker observes more accounts, both $K(n_o)$ and residual peer-order uncertainty fall. With full coverage, aggregate flow adds nothing about peers, although known peer orders can still make liquidation expensive. Figure 2 plots this comparative static.


3.2 Execution Loss and the Margin Requirement

Because the same hidden peer order affects every execution slice, its uncertainty survives aggregation. Averaging the service prices in (2.11) gives

$$U(x) = a + b_J x + Z, \quad Z | \mathcal{H}_i \sim N(0, \sigma_U^2),$$

$$a = d\lambda_0 + \psi A_i, \quad b_J = \frac{\eta_c(J-1)}{2J} + \frac{v_c}{J}, \quad (3.6)$$

$$\sigma_U^2 = \psi^2 \kappa^2 R_i + v_\epsilon, \quad v_\epsilon = J^{-2} \mathbf{1}' \Omega \mathbf{1}.$$


Figure 2

Broker coverage and unobserved competing order flow. The number of observed accounts, including the focal account, varies from one to eight while the trading population and its pricing and execution policies remain fixed. The left panel plots the dimensionless loading $K(n_o)$ in (3.5); the right panel plots residual competing-order-flow variance in squared currency-per-unit terms, $\kappa^2 V(n_o) \sigma_u^2 / [V(n_o) + \sigma_u^2]$. Parameters are $N=8$, $r=1$, $\sigma_v^2=1$, $\gamma=0.1$, $J=4$, and $\sigma_u=1$. Full coverage eliminates uncertainty about peer orders and their public update; it leaves any known pressure and independent execution disturbances in the cost calculation. Parameters are illustrative.

The residual term is $Z = \psi(\Theta_i - A_i) + J^{-1} \sum_j \epsilon_j$. A hidden peer order therefore affects all slices in the same direction, so its contribution to average execution risk is scaled by ψ^2 . Persistence raises the exposure to both the posterior mean and the posterior uncertainty.

For a full closeout, expected per-unit cost is $\mu_U = a + b_J Q$. Using the normal stop-loss transform

$$\mathcal{L}(t) = \varphi(t) - t\bar{\Phi}(t), \quad \mathcal{L}'(t) = -\bar{\Phi}(t) < 0.$$

turns the expected-loss constraint into a mean-plus-buffer margin rule.

Lemma 1 (Nonnegative margin requirement). For $Q > 0$, $\sigma_U > 0$, and $0 < \bar{\ell} < \infty$, the least rate satisfying (2.14) is

$$t_Q = \mathcal{L}^{-1}\left(\frac{\bar{\ell}}{Q\sigma_U}\right), \quad B = \sigma_U t_Q, \quad m = \max\{0, \mu_U + B\}. \quad (3.7)$$

On the positive-rate branch, at fixed Q and $\bar{\ell}$,

$$\frac{\partial m}{\partial \mu_U} = 1, \quad M_\sigma := \frac{\partial m}{\partial \sigma_U} = \frac{\varphi(t_Q)}{\bar{\Phi}(t_Q)} > 0, \quad \frac{\partial m}{\partial z_i} = \psi K_i. \quad (3.8)$$

Inside the strict zero-rate region these local derivatives are zero.

On the interior branch, margin is expected closeout cost plus the risk buffer B . The key derivative is therefore immediate: an adverse innovation moves the posterior mean of peer flow by K_i , and persistence carries that forecast into closeout with weight ψ , giving $m_{z_i} = \psi K_i$. The zero floor and deterministic boundary cases are collected in Appendix B.3.

Comparative statics that change the trading environment generally move both the mean and variance of closeout costs. For a differentiable path τ , with Q and $\bar{\ell}$ fixed and $m > 0$,

$$\begin{aligned} m_\tau &= a_\tau + Q(b_J)_\tau + M_\sigma(\sigma_U)_\tau, \\ a_\tau &= d_\tau \lambda_0 + d(\lambda_0)_\tau + \psi_\tau A_i + \psi(A_i)_\tau, \\ 2\sigma_U(\sigma_U)_\tau &= 2\psi\psi_\tau \kappa^2 R_i + \psi^2(\kappa^2 R_i)_\tau + (v_\epsilon)_\tau. \end{aligned} \quad (3.9)$$

A realized order-flow innovation is simpler: within a fixed economy it shifts the mean only. Changes in signal quality or hedging affect both terms.

3.3 Order-Flow Composition at Fixed Price-Update Variance

We now compare economies with the same price-update variance. Fix the account, contract, execution technology, and $N \geq 2$. Increase signal noise r while reducing transitory hedging σ_u so that $\mathcal{V}_p := \text{Var}(p_0)$ does not change. The experiment holds constant how much information prices reveal in variance terms while changing the composition of the order flow that generated that information.

Write the informed share of initial flow variance and its zero-delay-cost ceiling as

$$\chi = \frac{\beta^2 \sigma_v^2 N(N+r)}{\beta^2 \sigma_v^2 N(N+r) + \sigma_u^2}, \quad \bar{\chi}_N = \frac{N+r}{N+1+2r}. \quad (3.10)$$

Suppress the focal index on K, z, S, A and A^0 . Define $h = (N-1)r/[N(1+r)]$ and $t_\chi = 1 - (1-h)\chi$. The relevant variance identities are

$$\begin{aligned} \mathcal{V}_p &= \frac{\sigma_v^2 N}{N+r} \chi, \quad K = \frac{h\chi}{t_\chi}, \\ S^2 &:= \text{Var}(z | \mathcal{G}_i) = \kappa^2 (V_i + \sigma_u^2) = \mathcal{V}_p t_\chi, \\ C_P &:= \kappa^2 R_i = \mathcal{V}_p K(1-\chi), \quad \sigma_U^2 = \psi^2 C_P + v_\epsilon. \end{aligned} \quad (3.11)$$

Because $\mathcal{V}_p = \sigma_v^2 - \text{Var}(v | y_0)$, the comparison keeps the dealer's aggregate learning fixed. It does not keep the broker's innovation variance S^2 or residual peer uncertainty C_P fixed.

Let $\bar{\mathcal{V}}_p > 0$ be the target public variance. The admissible composition path satisfies

$$\chi(r) = \frac{\bar{\mathcal{V}}_p(N+r)}{N\sigma_v^2} < \bar{\chi}_N(r), \quad r > 0. \quad (3.12)$$

Along an admissible interval, the pricing block delivers a unique $\sigma_u(r) > 0$. Higher r means noisier private signals, less transitory hedging, and a larger informed share of initial flow.

We compare states with the same own signal, position, direction, and settled equity.¹ The signed innovation z is held fixed in price units; its standardized value $w = z/S$ may differ across economies.

Proposition 3 (Order-flow composition and margin sensitivity). Along (3.12), both K and ψ increase with r . At any positive requirement, the response to the signed order-flow innovation is $m_z = \psi K$, and

$$\frac{d(\psi K)}{dr} > 0.$$

For any $z_* > 0$ and fixed own signal, if both endpoint requirements are positive, their contrast satisfies

$$m(r, z_*) - m(r, -z_*) = 2z_*\psi(r)K(r), \quad \frac{d}{dr}\{m(r, z_*) - m(r, -z_*)\} > 0. \quad (3.13)$$

The proposition has a simple economic decomposition. Noisier private signals make the client's own order less informative about peers, so aggregate flow receives more weight: K rises. Thinner hedging raises price impact, which slows the remaining informed schedule in this execution technology: ψ also rises. The margin loading is their product. Appendix B.4 proves that both monotonicities hold along the fixed- $\mathcal{V}_p$ path.

In the benchmark used later, $\psi K \simeq .29$ and $S \simeq .563$, so a one-standard-deviation adverse innovation raises m by about .163 from .70482, or roughly 23%. A rule that reacts only to price-update variance would not move along this path. A rule based only on conditional execution-cost variance would react symmetrically to z and $-z$. The signed difference in (3.13) comes from the conditional mean of closeout costs.

Corollary 1 (A margin-level region). Along (3.12), suppose $s_i = 0$, $z \geq 0$ is fixed, and $\chi \leq 1/2$. Then σ_U increases strictly with r , the positive margin rate increases strictly, and the nonnegative requirement is nondecreasing over any interval satisfying these restrictions.

This corollary concerns levels rather than sensitivity. With $s_i = 0$, $z \geq 0$, and $\chi \leq 1/2$, both the expected-cost term and the risk buffer rise. Because $K < \chi$,

$$\frac{(C_P)_r}{C_P} = (1-K)\frac{h_r}{h} + \frac{\chi r(1-K-\chi)}{\chi(1-\chi)} > 0,$$

so residual execution risk also rises in this region. At higher informed shares the variance effect can reverse, and margin levels need not rise. The signed loading result in Proposition 3 does not rely on this level restriction.

¹ At the designated signal and order-flow state, choose $q_0 = \varsigma Q - \beta s_i$ and $e_0 = e - q_0 p_0$ as that economy's entering endowments, then condition on that state. This matches Q and e without a signal-dependent initial-cap selection. The two endpoints z and $-z$ can require different entering equity to preserve the same settled e . Appendix B.4 supplies the regular-conditional-state construction.

3.4 Transitory Hedging and the Margin Response

A different experiment varies only the scale of transitory hedging, $u := \sigma_u$, holding signal quality fixed. We match the standardized state $w = z/S$ rather than the price-unit innovation. Unlike the previous exercise, price-update variance is free to move.

Less hedging makes the initial price more informative and raises execution persistence in the benchmark. Those changes push expected closeout cost upward, but they can also reduce residual uncertainty about peer flow. Margin levels therefore depend on which force dominates.

Use χ to parameterize this path, and let primes denote total derivatives with respect to χ along its pricing solution. Put

$$\begin{aligned} P_0 &= \frac{\sigma_v^2 N}{N+r}, & \vartheta_i &= \frac{-\varsigma s_i (N-1)}{(N+r)(1+r)}, \\ \mathcal{Z}(\chi) &= KS = \frac{\sqrt{P_0} h \chi^{3/2}}{\sqrt{t_\chi}}, & A &= \vartheta_i \chi + w \mathcal{Z}(\chi), \\ C_P &= P_0 \frac{h \chi^2 (1-\chi)}{t_\chi}, & \sigma'_U &= \frac{2\psi \psi' C_P + \psi^2 C'_P}{2\sigma_U}. \end{aligned} \quad (3.14)$$

The derivative makes that trade-off explicit:

$$m' = \vartheta_i (\psi + \chi \psi') + w (\psi \mathcal{Z})' + M_\sigma \sigma'_U. \quad (3.15)$$


Proposition 4 (Transitory hedging and the state threshold). In this positive-delay-cost hedging experiment, $d\chi/du < 0$, $d\kappa/du < 0$, and $0 < d\log\beta/d\log u < 1$. At a positive requirement with $Q > 0$ and $0 < \bar{\ell} < \infty$, define

$$w_* = -\frac{\vartheta_i (\psi + \chi \psi') + M_\sigma \sigma'_U}{(\psi \mathcal{Z})'}, \quad \text{sign}\left(-\frac{dm}{du}\right) = \text{sign}(w - w_*). \quad (3.16)$$

The denominator is strictly positive. If $\chi \leq 1/2$, $-\varsigma s_i \geq 0$ and $w \geq 0$, thinner hedging strictly raises the positive requirement.

When $\chi \leq 1/2$ and the signed state is adverse, both terms move in the same direction and thinner hedging raises margin. At high informed shares, residual risk can fall fast enough to overturn the mean effect. The threshold w_* marks the point at which the expected-cost effect catches up with that risk reduction. Appendix B.5 gives the derivatives.

Panel B of Figure 3 shows this reversal at a high informed share, $\chi = .91$. At $w = .01$, thinner hedging lowers margin: the decline in residual execution risk outweighs the increase in expected closeout cost. The effect is small in the calibration, which is why we treat it as a boundary result rather than the paper's main quantitative finding.


Figure 3

Composition and hedging scale are different comparisons. Panel A holds $\mathcal{V}_p = 0.20$, $N = 5$, $\sigma_v^2 = 1$, $\gamma = 0.12$, $J = 4$, $s_i = 0$, and the price-unit state $z = 0.5$ fixed. As r varies from 0.15 to 2, the finite pricing solution adjusts σ_u , β , κ and execution persistence. The solid curve is margin m in currency per asset unit; the dashed curve is the dimensionless response coefficient ψK . Their levels have different units. Panel B varies the hedging standard deviation with $N = 20$, $r = 1$, $\sigma_v^2 = 1$, $\gamma = 0.1$, $J = 4$, $s_i = 0$, and $w = 0.01$ fixed; the increasing curve illustrates a region where thinner hedging lowers margin. Both panels use $Q = 1.5$, $e = 0.75$, $\bar{\ell} = 0.02$, $d\lambda_0 = 0.2$, $\eta_c = 0.12$, and $\nu_c = 0.14$. Independent slice disturbances have covariance $0.16I_4$ in A and $0.04I_4$ in B, where I_4 is the identity matrix. Conditional peer uncertainty and σ_U are recomputed throughout. Parameters are illustrative, not estimated.

3.5 Zero-Delay-Cost Benchmark

The zero-delay-cost case is a useful benchmark. When $\gamma = 0$, the target is split equally across the $J + 1$ trading dates. Informed intensity scales with hedging and price impact moves inversely.

Proposition 5 (Matched-state invariance). Set $\gamma = 0$ and vary $u > 0$ at fixed N , r , σ_v^2 , and J . Then $\chi = \bar{\chi}_N$, $g_j = \psi = 1$, and $\beta \propto u$, $\kappa \propto u^{-1}$. For $Q > 0$, if own signal, position direction, z , pre-reset (Q, e) , the expected-loss limit and the service inputs are matched, then A , σ_U , the nonnegative margin m and the selected reduction D are invariant to u . Matching $w = z/S$ is equivalent because S is constant on this path. At zero pre-reset exposure, $m = D = 0$ directly in every economy.

Under the matched-state comparison in Proposition 5, $\mathcal{V}_p$, K , S , $\kappa\beta$, and C_p are invariant, so both margin and the chosen reduction are unchanged. Appendix B.6 gives the proof. The next section turns the margin requirement into an actual liquidation quantity.

4. Margin Calls and Forced Reductions

A margin call does not translate one-for-one into a sale. Selling reduces the collateral required on the position that remains, but the sale itself costs money. The broker therefore chooses the smallest reduction that restores the margin

condition in expectation. We solve that quantity and then ask how it changes with order-flow news.

4.1 The Minimum Margin-Restoring Reduction

Fix the pre-reset position Q , settled equity e , and the margin rate m determined in Section 3. For a candidate reduction x , expected per-unit execution cost is $c(x) = a + b_J x$, so forecast equity after the sale is $E^f(x) = e - xc(x)$. Relative to the locked requirement on the remaining position, slack is

$$F(x) = E^f(x) - m(Q - x) = e - mQ + (m - a)x - b_J x^2. \quad (4.1)$$

If $e < mQ$, the broker starts with negative slack. It chooses the first root of F at which another unit sold would still improve support, $F'(x) > 0$.

Proposition 6 (The margin-call reduction rule). The operational contract assigns $D \in [0, Q]$ as follows.

- (i) At $Q=0$, set $m=D=0$. Nonnegative equity, including $e=0$, permits ordinary continuation; negative equity cancels it and leaves a cash deficit.
- (ii) At $Q > 0$, nonpositive equity or an infeasible finite-rate contract triggers full wind-down, $D=Q$, and cancels ordinary continuation.
- (iii) At $Q > 0$, $e > 0$, and finite m , the condition $e \geq mQ$ gives $D=0$ and permits ordinary continuation, including equality.
- (iv) At $0 < e < mQ$, put $\delta = mQ - e > 0$. If $b_J > 0$, define $\Delta_F = (m - a)^2 - 4b_J\delta$. When $m > a$, $\Delta_F > 0$, and

$$D_- = \frac{2\delta}{m - a + \sqrt{\Delta_F}} < Q, \quad (4.2)$$

choose $D = D_-$. It is the smallest forecast-feasible quantity, with

$$H := F'(D) = m - a - 2b_J D = \sqrt{\Delta_F} > 0. \quad (4.3)$$

If $b_J = 0$, choose $D = \delta / (m - a)$ when $m > a$ and this quantity is below Q . In every remaining deficient state, choose $D = Q$. Every positive reduction cancels the entire remaining client execution program.

Because F is concave, the relevant interior solution is the smaller quadratic root. Equation (4.2) writes that root in numerically stable form. Tangencies and states with no regular interior cure are assigned full wind-down; Appendix B.7 collects the boundary cases.

The minimum matters. With $Q = 1$, $e = .95$, $a = .10$, $b_J = 1$, and $m = 1$, a sale of only $D = .059488$ restores expected support even though a full unwind would leave $F(Q) = -.15$. Looking only at the economics of a complete liquidation would therefore misclassify this state.

The remaining branches are straightforward. A zero margin rate with positive equity permits continuation; a zero expected-loss limit under

nondegenerate Gaussian execution risk admits no finite rate and therefore triggers full wind-down. At $Q=0$, both the requirement and forced quantity are zero. Moving across a branch boundary can produce a discrete jump in D .

4.2 Information, Equity Erosion, and Quantity Responses

On the interior branch, the liquidation quantity is differentiable. Let $H = F'(D) > 0$. Along any parameter path τ , implicit differentiation gives

$$D_\tau = \frac{(Q-D)m_\tau + Da_\tau + D^2(b_J)_\tau - e_\tau + mQ_\tau}{H}. \quad (4.4)$$

The numerator has an economic interpretation: parameter changes alter the requirement on retained units, the cost of units sold, equity, and the starting position.

For the signed order-flow experiment, only the conditional mean moves. Holding the account and trading environment fixed gives $m_z = a_z = \psi K$, so

$$D_z = \frac{\psi K Q}{H}. \quad (4.5)$$

The result is especially simple: adverse news raises the required collateral on the entire pre-reset position, so the minimum restorative sale rises at rate $\psi K Q / H$. The denominator H measures how much net collateral relief the next unit of selling provides.

It is useful to write that denominator in relative terms:

$$\Xi = \frac{a + 2b_J D}{m}, \quad H = m(1 - \Xi), \quad Q_m = \frac{E^f(D)}{m^2} = \frac{Q - D}{m}. \quad (4.6)$$

Here Ξ is marginal execution cost relative to the margin released by one more unit sold, while Q_m is the direct quantity response to the margin rate. Then

$$D_z = \frac{\psi K (Q_m + D/m)}{1 - \Xi}.$$

When $0 < \Xi < 1$, execution costs amplify the quantity response because some of the margin released by selling is immediately lost to execution. Appendix B.8 gives the stability and boundary details.

If a parameter moves both expected cost and risk, write $m = a + b_J Q + B$. At fixed Q, e, b_J ,

$$D_\tau = \frac{Qa_\tau + (Q - D)B_\tau}{H}. \quad (4.7)$$

Expected cost and the risk buffer enter the liquidation decision with different weights. A rise in the margin rate alone is therefore not enough to sign the change in D .

Corollary 2 (Composition and regular reductions). Along the fixed-public-price-variance path in (3.12), suppose $s_i=0$, $z \geq 0$, and $\chi \leq 1/2$, with the other inputs held as in Corollary 1. At an interior margin-restoring reduction,

$$D_r = \frac{Qa_r + (Q-D)M_\sigma(\sigma_U)_r}{H} > 0. \quad (4.8)$$

In the low-informed-share region of the composition experiment, both the mean and risk terms raise the forced quantity. Outside that region, the two can oppose each other, just as they do for the margin level. On the continuation and full-wind-down branches, D is locally fixed at 0 and Q , respectively.

4.3 Margin-Call Regions

Table 2 and figure 4 show the three liquidation regions at fixed trading parameters. Moving right makes the order-flow state more adverse and raises the required reduction; moving up adds equity and reduces it. Conditional execution risk is held fixed in this illustration.

Table 2
Continuation, partial reduction, and full wind-down

z	e	m	D	$Q-D$	$E^f(D)$	Decision
-1.0	1.20	.41480	.00000	1.50000	1.20000	Continue
0.0	.75	.70482	.68239	.81761	.57627	Partial reduction
.5	.75	.84983	1.31245	.18755	.15939	Partial reduction
1.0	.75	.99485	1.50000	.00000	-.16504	Full wind-down

All rows use the parameters of Figure 4. They imply $b_J = .08$, $\sigma_U = .2951763$ and $B = .3848220$. Position quantities Q, D are in asset units; z and m are currency per unit, and e and forecast equity are currency. The partial reductions satisfy $E^f(D) = m(Q-D)$; the full wind-down leaves negative forecast equity. These are illustrative conditional states, with entering endowments matched as in Section 3.3.

4.4 Posted Collateral and End-of-Window Shortfalls

The reduction is chosen before execution shocks are realized. After slice j , cumulative execution cost is $L_j = \sum_{h \leq j} r_h D U_h(D)$. For $Q > 0$, the forced trade is $f_j = -\varsigma r_j D$ and changes cash by


$$-p_j(D)f_j = \varsigma p_0 r_j D - r_j D U_j(D).$$

Benchmark-marked equity in the settled book is therefore $E_j^I = e - L_j$; the same notation covers long liquidation and short covering.

The intended residual collateral amount and actual restricted principal are

$$M^* = \begin{cases} m(Q-D), & Q-D > 0, \\ 0, & Q-D = 0, \end{cases} \quad (4.9)$$

$$M_j = \min\{M^*, (E_j^I)^+\}, \quad j=0, \dots, J, \quad E_0^I = e.$$


Figure 4

Chosen reduction across order-flow states and equity. Color gives the fraction D/Q selected by Proposition 6, from zero under continuation to one under full wind-down. The horizontal axis is the signed price-unit innovation z ; the vertical axis is post-settlement pre-reset account equity e . Fixed parameters are $N=5$, $r=1$, $\sigma_v^2=1$, $\gamma=.12$, $J=4$, $\sigma_u=1$, $s_i=0$, $Q=1.5$, $\bar{\ell}=.02$, $d\lambda_0=.2$, $\eta_c=.12$, $v_c=.14$, and disturbance covariance $\Omega=.16I_4$. The finite pricing solution, peer persistence and full conditional variance determine each rate; only the realized state and available equity vary. Full wind-down can leave debt. The color scale measures a quantity fraction; default depends on subsequent execution and value realizations. No parameter is estimated.

After the sale, intended collateral is $M^*=m(Q-D)$ and actual restricted collateral cannot exceed remaining equity. Full wind-down leaves no residual requirement. The chosen reduction is completed before the broker resolves any realized shortfall; any unsold position is then held to maturity.

If the account continues, there is no forced invoice and the settled position keeps collateral mQ . Any costs generated by the continuation program remain in its separate book, as in (2.16).

For intervals Δ_j before each slice, with $\sum_j \Delta_j = 1$, the pre-reset collateral's opportunity cost is

$$F_{\text{coll}} = k \sum_{j=1}^J \Delta_j M_{j-1}. \quad (4.10)$$

Posting collateral does not itself destroy resources. The only resource cost attached to restricted collateral is the opportunity cost in (4.10); execution losses are already recorded in L_j .

Let $E_D = e - L_j$ be end-window pre-reset account equity. Define its collateral deficiency and benchmark-marked deficit by

$$C_D = [M^* - E_D]^+, \quad S_D = [-E_D]^+. \quad (4.11)$$

At a regular root, $E_D = M^* - DZ$. If $D\sigma_U > 0$, the centered Gaussian shock in (3.6) gives

$$\begin{aligned} \Pr(C_D > 0 | \mathcal{H}_i) &= \frac{1}{2}, & \mathbb{E}[C_D | \mathcal{H}_i] &= \frac{D\sigma_U}{\sqrt{2\pi}}, \\ \Pr(S_D > 0 | \mathcal{H}_i) &= \overline{\Phi}\left(\frac{m(Q-D)}{D\sigma_U}\right). \end{aligned} \quad (4.12)$$

At an interior cure, expected equity is set exactly equal to the residual requirement. A centered execution shock therefore puts realized collateral above or below that requirement with equal probability, even though the probability of negative marked equity can be tiny. For example, $Q = 1.5$, $e = .75$, $a = .20$, $b_J = .08$, $\sigma_U = .25$, and $B = .35$ imply $m = .67$ and $D = .604818$: the collateral-deficiency probability is one half, while the probability of negative benchmark-marked equity is 3.64×10^{-5} .

On interruption, contractual legal resources at maturity are

$$N_I = E_D + (q_1 + f)(v - p_0), \quad f = \sum_j f_j. \quad (4.13)$$

A partial reduction leaves fundamental exposure in the unsold position. Full wind-down removes that exposure, leaving client cash $[e - QU(Q)]^+$ and broker shortfall

$$S_{\text{full}} = [QU(Q) - e]^+, \quad \mathbb{E}[S_{\text{full}} | \mathcal{H}_i] = Q\sigma_U \mathcal{L}\left(\frac{e/Q - a - b_J Q}{\sigma_U}\right) \quad (4.14)$$

for $Q\sigma_U > 0$. The ex ante margin rule and the ex post shortfall are therefore different objects: the first is calibrated to a hypothetical closeout of Q , while the second depends on realized execution costs and available equity. Appendix B.9 gives the remaining boundary cases.

4.5 Liquidation Inventory and Subsequent Trading Costs

Finally, suppose liquidity-provider execution costs depend on the inventory already sitting in the relevant pool. Let X_0 be that inventory before the forced sale. With no reallocation during the closeout window, the sale moves the state according to

$$X_1 = X_0 + D, \quad \lambda_0 = \lambda_{\min} + \omega_X X_0, \quad \lambda_1 = \lambda_0 + \omega_X D, \quad \omega_X \geq 0. \quad (4.15)$$

The slope ω_X translates liquidation inventory into the cost level faced by later trades. It does not change the already locked margin rate.

At a regular matched order-flow state,

$$(\lambda_1)_z = \omega_X \frac{\psi K Q}{H}.$$

Before inventory is reallocated, a subsequent trader with exposure d_L pays an additional $d_L \omega_X D$ per unit, holding its other execution conditions fixed. The next section writes that implication for a fixed trade size.

If providers can shed inventory before the next trader arrives, only $\underline{F}$ part of the burden remains. With retention $\varphi_X \in [0, 1]$ and placement capacity $\bar{F} \geq 0$,

$$X^{\text{next}} = X_1 - \min\{(1 - \varphi_X)X_1, \bar{F}\}.$$

Nonbinding placement capacity leaves $\varphi_X X_1$ in the pool; otherwise capacity determines the residual inventory.

5. Forced Liquidation and Subsequent Trading Costs

In the model, spillovers arise through liquidity-provider inventory. If execution costs rise with that inventory, a larger forced sale leaves a more expensive trading environment for the next user of the same pool. The result below traces that link. We then close the accounting by showing where execution payments, asset risk, and any creditor loss end up.

5.1 Trading Costs for a Subsequent Trader

Consider a trader that arrives after the forced sale but before inventory is reallocated. It trades $y_L \geq 0$ in the same pool and has inventory-cost exposure $d_L \geq 0$. The forced sale raises the pool's cost state by $\omega_X D$, so the trader's extra bill is its quantity times that cost increase.

Proposition 7 (Liquidation inventory and subsequent execution costs).

Use the matched order-flow states of Section 3.3, holding the own signal, position direction, $(Q, e, \bar{\ell})$, entering λ_0 , trading and service parameters, and disturbance covariance fixed. Suppose the focal reset has a positive-margin interior reduction, $0 < D < Q$ and $H > 0$. Hold the subsequent trader's benchmark price, signed direction, quantity and schedule, other service states, and execution shocks fixed. Before further placement, its invoice increment attributable to the focal inventory is

$$I_L^X = y_L d_L (\lambda_1 - \lambda_0) = y_L d_L \omega_X D, \quad \frac{\partial I_L^X}{\partial z} = y_L d_L \omega_X \frac{\psi K Q}{H}. \quad (5.1)$$

The derivative is strictly positive when $y_L, d_L, \omega_X, K > 0$.

Combining (4.5) with (4.15) gives the second equality. When H is small, each unit sold releases little net collateral, so the broker must sell more in response to the same adverse information. The same amplification then appears in the subsequent trader's bill. Appendix B.10 gives the boundary cases.

Between two matched reset states, the implied difference in the subsequent invoice is $y_L d_L \omega_X (D_2 - D_1)$. The subsequent quantity is held fixed; the model does not determine that trader's demand.

If the focal account continues, there is no liquidation inventory. A full wind-down leaves inventory Q , although the local derivative with respect to z is then zero because the quantity is already at its upper bound. Setting $\omega_X = 0$ shuts off the subsequent-cost channel without changing the forced quantity.

More generally, any parameter that changes the forced quantity changes the subsequent invoice by

$$(\mathcal{I}_L^X)_\tau = y_L d_L \omega_X D_\tau.$$

Thus the low-share composition region in Corollary 2 raises the invoice whenever $y_L d_L \omega_X > 0$.

5.2 Execution Payments and Transferred Risk

The remaining accounting is easiest to write relative to the settled price. Let $\delta_v = v - p_0$. Ordinary continuation produces total quantity $q^P = \sum_j x_j^P$ and execution concession $C_{\text{exec}} = \kappa \sum_j (x_j^P)^2$. A forced reduction has total order $f = -\varsigma D$ and invoice L .

The account's contractual legal net resources in (2.17) and the counterparties' incremental terminal trading profits are

$$\begin{aligned} \mathcal{N} &= e - L - C_{\text{exec}} + (q_1 + f + q^P) \delta_v, \\ \Pi_F &= L - f \delta_v, \\ \Pi_P &= C_{\text{exec}} - q^P \delta_v. \end{aligned} \tag{5.2}$$

Here Π_F is the trading profit of the providers that absorb the forced order and Π_P the trading profit of ordinary counterparties, before their own financing costs.

The cash transfers and asset positions cancel across these three accounts:

$$\mathcal{N} + \Pi_F + \Pi_P = e + q_1 \delta_v. \tag{5.3}$$

Equation (5.3) is an accounting identity: execution concessions are transfers to counterparties, while the underlying asset risk moves with the position. Appendix B.11 gives the detailed derivation.

The same accounting applies to the subsequent trader. Holding its quantity and asset payoff fixed, an inventory-induced increase $d_L \omega_X D$ in the concession transfers $\mathcal{I}_L^X$ from the trader to its counterparty.

5.3 Creditor Recoveries and Collateral Resources

If contractual resources are negative, final incidence depends on creditor recoveries. Let $\mathcal{A}$ be distributable assets, H_h the remaining claim of creditor h , R_h its recovery, and $c \geq 0$ the client's residual payout. Recoveries respect the priority structure in Section 2.5.

With committed execution completed, no new outside injection and no modeled bankruptcy-resource destruction, complete distribution satisfies

$$\begin{aligned} c + \sum_h R_h &= \mathcal{A}, & \mathcal{N} &= \mathcal{A} - \sum_h H_h, & S &= \sum_h (H_h - R_h), \\ c - S &= \mathcal{N}, & c - S + \Pi_F + \Pi_P &= e + q_1 \delta_v. \end{aligned} \tag{5.4}$$

The identity $c - \mathcal{S} = \mathcal{N}$ replaces contractual net resources with the client's payout minus creditor losses without counting unpaid debt twice.

After interruption, the client receives $(\mathcal{N}_I)^+$ and creditors lose $(-\mathcal{N}_I)^+$. Under full wind-down this becomes $[L - e]^+$. A partial reduction is different because the unsold position retains fundamental risk, so an end-window collateral deficiency need not equal the eventual creditor loss.

For example, let $D = .5$, $Q = 2$, $e = 1.175$, $v = p_0$, and $U(D) = 2.45$. The liquidation bill is 1.225, leaving contractual resources of $-.05$. The client receives zero, creditors lose .05, and providers receive trading profit of 1.225. The three amounts sum to entering equity, as (5.4) requires.

On continuation, the contractual books are

$$\mathcal{N}_I = e + q_1 \delta_v, \quad \mathcal{N}_P = q^P \delta_v - C_{\text{exec}}, \quad \mathcal{N} = \mathcal{N}_I + \mathcal{N}_P.$$

Continuation-program creditors recover first from that program's positions and proceeds; they do not absorb pre-reset collateral during the window. Equation (5.4) imposes aggregate consistency without requiring a particular allocation across individual creditors.

Restricted collateral remains a client asset. Execution invoices are already deducted through L_j , while F_{coll} measures only the opportunity cost of tying up collateral. Adding that resource cost to (5.4) gives

$$c - \mathcal{S} - F_{\text{coll}} + \Pi_F + \Pi_P = e + q_1 \delta_v - F_{\text{coll}}. \quad (5.5)$$

Appendix B.11 verifies the accounting. Proposition 7 is an incidence result: it traces how liquidation inventory changes a subsequent trader's execution bill under the specified cost technology. A welfare analysis would additionally require a model of providers' resource costs, risk bearing, and the value of canceled trading opportunities.

6. Empirical Predictions and Implementation

The model's empirical prediction is an interaction. After controlling for volatility and the account's balance sheet, margin should respond more strongly to a signed order-flow surprise when that flow is both informative about hidden peer trading and persistent over the closeout horizon. A structural test would require the broker's information set. A reduced-form test can instead proxy the two model objects, K and ψ , using information available before the margin reset.

A natural specification is

$$\Delta m_{it} = \alpha + \beta \widehat{z}_{it} + \theta \widehat{z}_{it} \text{Persist}_{it} \text{Info}_{it} + \Gamma X_{it} + \text{FE} + \varepsilon_{it},$$

where X_{it} contains volatility, position size, equity, and contract controls measured before the reset. With positively oriented proxies for persistence and informed-flow content, the model predicts $\theta > 0$. Bilateral prime-broker data would map most closely to the theory; exchange-set margin changes offer a more accessible, though less direct, setting (Abruzzo and Park 2014).

6.1 Margin Sensitivity and Order-Flow Composition

The cleanest test mirrors Proposition 3. Compare equally sized adverse and favorable order-flow innovations after conditioning on the broker's prior information. With positive margin at both endpoints,

$$m(r, z_*) - m(r, -z_*) = 2z_*\psi(r)K(r).$$

The contrast is larger when the broker learns more from aggregate flow and when inferred peer trading is more persistent. Holding price-update variance fixed is one way to isolate that composition effect from a volatility-only margin rule.

The empirical unit should therefore include the dated margin request, the signed position, the client's own order, and the information available to the broker before the reset. Position size, equity, horizon, and execution conditions must be controlled for because each changes the closeout problem directly. If the observed margin measure combines several risk components, the liquidity or closeout component should be separated where possible.

The signed-pair comparison is useful because the broker's prior and the common risk buffer drop out. What remains is the response of expected liquidation cost to the innovation. Price-update variance should be estimated at a common horizon rather than proxied by the realized return on the observation date.

The interaction is not unique to the exact stop-loss contract in (2.14); other margin rules that respond to expected liquidation costs can generate the same signed pattern. The test should therefore be read as evidence on the model's information channel, not as a test of one contractual formula. Public news and changes in broker balance-sheet capacity are natural competing explanations.

Predictions for margin levels are less robust. In the low-informed-share region of Corollary 1, margin rises with the composition shift; at high informed shares, lower residual risk can overturn that sign. For this reason, the signed loading is the cleaner empirical object.

6.2 Constructing the Broker's Order-Flow Forecast

The order-flow innovation is not raw flow. It is the part of aggregate flow that remains after conditioning on the client's own information. In the centered model,

$$\iota_i = y_0 - x_{i0} - \mu_i, \quad z_i = -\zeta K \iota_i, \quad A_i = A_i^0 + K_i z_i.$$

For a long position, client-netted flow that is more negative than expected is adverse; for a short, unexpectedly positive flow is adverse. The prior μ_i must be formed from information that genuinely predates the aggregate-flow observation.

This timing matters. If the client's information is omitted from the prior, correlated private signals can be mistaken for learning from public flow. The identity $s_i = x_{i0}/\beta$ is model-specific, so empirical work would typically replace

it with a broker-side forecast or a proxy for the client's information. Price and flow cannot serve as independent public signals in the model because $p_0 = \kappa \gamma_0$.

Three variances should not be conflated: price-update variance $\mathcal{V}_p$, innovation variance S_i^2 , and conditional execution-loss variance σ_U^2 . They describe different parts of the model. A volatility control addresses the first, but not the information content of order flow or the residual uncertainty of liquidation costs.

Persistence should likewise be measured before the liquidation response. Scheduled orders or mandate records are ideal; pre-reset order-flow autocorrelation is a noisier proxy because it also captures common news and correlated strategies. Post-reset flow is a bad proxy for ψ when it already contains the forced sales being explained.

6.3 Transitory Hedging and Matched-State Benchmarks

The hedging comparative static yields a different prediction. Holding signal quality fixed, changes in transitory hedging affect both the information content of prices and residual liquidation risk. At a positive margin rate, Proposition 4 gives

$$\text{sign}\left(-\frac{dm}{du}\right) = \text{sign}(w - w_*),$$

where w_* depends on both the broker's prior and the change in residual execution risk.

When the informed share is modest, thinner hedging raises margin in adverse states. At very high informed shares, the fall in residual risk can dominate and the sign can reverse, as in Figure 3. That reversal is a useful boundary condition rather than a baseline prediction.

Empirically, this exercise requires a credible measure of transitory flow. Turnover alone is not enough: news, signal quality, and broker funding can move volume and margin together.

The zero-delay-cost benchmark sharpens the same point. When $\gamma=0$, changing hedging scale leaves margin and the chosen reduction unchanged under the matched-state comparison in Proposition 5. That prediction is about the execution schedule, not about raw turnover.

6.4 From the Request to Sales and Shortfalls

The next outcome is the forced sale itself. A clean design needs the margin request and the broker-directed reduction, not merely a change in the client's position. Voluntary trades, preemptive deleveraging, and outside capital injections otherwise contaminate the mapping from m to D .

On the interior branch, the model predicts $D_z = \psi KQ/H$. The denominator H is the net collateral relief from one more unit sold, so a small H amplifies the response. This object should be measured from information available at the reset, not from realized execution costs that arrive afterward.

Branch selection is economically important. Slack accounts have no forced sale, while full wind-downs are already at $D=Q$ and therefore have no local quantity response to additional bad news. Empirical work should separate these regions rather than estimate a single linear response across them.

Requested margin, posted collateral, and creditor loss are also distinct outcomes. The model's one-half deficiency probability in (4.12) applies to the mechanical expected-equity cure under centered Gaussian shocks; it is not a generic default prediction. Terminal credit losses additionally depend on the value of any unsold position and on recoveries.

6.5 Liquidation Inventory and Subsequent Execution Costs

The final prediction follows the forced sale into the next trade executed by the same liquidity-provider pool. Before inventory is reallocated, Proposition 7 attributes $y_L d_L \omega_X D$ of the subsequent invoice to the focal liquidation. Comparing two matched states gives the difference $y_L d_L \omega_X (D_2 - D_1)$.

This is a prediction for execution-cost levels under the specified inventory technology. It is not, by itself, a prediction for market depth or welfare.

If inventory is partly placed before the next trade, the relevant state is

$$\mathcal{T}(X) = X - \min\{(1 - \varphi_X)X, \bar{F}\}.$$

The retained share is proportional only when placement capacity is slack. Common liquidity shocks and sales by other accounts are the main empirical confounds.

The cleanest falsification is $\omega_X = 0$: holding the entering cost state fixed, the focal sale can still occur but it no longer changes the subsequent trader's cost through inventory.

6.6 Information Coverage

Information coverage gives a final cross-sectional prediction. As a broker observes more of the relevant orders directly, aggregate flow should carry less incremental information and the signed-flow loading $\psi K(n_o)$ should fall. Full coverage removes the public update but not known liquidation pressure or other execution costs.

This comparison requires actual information access, not simply the number of client relationships. Consolidation or reporting changes may alter both coverage and the underlying trading environment, so relationship counts are at best a proxy.

Taken together, the empirical restrictions form a sequence: information available before the reset predicts the signed margin response; the margin shortfall predicts broker-directed sales; and those sales predict subsequent execution costs only when inventory remains with the same liquidity providers.

7. Conclusion

Price discovery and liquidation risk are not the same object. A market price reflects information about value, but a broker financing a particular position also cares about the trades that remain if that position must be closed. Because the broker observes its own client's order, aggregate flow can reveal something about those remaining trades even after the same flow has already moved the price.

In the model, the signed margin response is governed by two objects: how much the broker learns from aggregate flow and how persistent informed execution is. Order-flow composition can change both while leaving price-update variance unchanged. In the benchmark, a one-standard-deviation adverse surprise raises the per-unit requirement by about 23%. When equity is insufficient, the higher requirement translates into a larger forced reduction; under the inventory-cost specification, that sale can also raise the execution cost of the next trade served by the same liquidity-provider pool.

The broader point is narrow but useful. Price informativeness does not by itself determine the funding consequences of order flow. For a financed position, who is still likely to trade—and for how long—can matter as much as how much information has already entered the price.

Appendix A. Notation and Measurement

Table 3 follows trading through settlement, suppressing an unambiguous focal index. Currency C and quantity U give price units C/U and slope units C/U^2 . Statuses distinguish *primitive* exogenous inputs, *maintained* contracts or services, *derived* solutions or accounting definitions, and transaction or account *Record* counterparts when available. Records, forecasts and proxies remain distinct.

Table 3
Model objects, units, and interpretation

Symbol	Definition and economic role	Units; status
<i>Fundamentals, initial trading, and finite mandates</i>		
N, J	Informed accounts and post-reset trading slices; $N \geq 1, J \geq 1$ are integers. An ordinary mandate has $J+1$ trades including date 0; a forced reduction has J slices.	Counts; primitive
v, s_i, η_i	Terminal value, private signal and its error: $s_i = v + \eta_i$. The centered Gaussian primitives and their independence are specified in (2.1).	C/U ; primitive
$\sigma_v^2, \sigma_\eta^2, r, \rho$	Value and signal-error variances; $r = \sigma_\eta^2 / \sigma_v^2 > 0$, $\rho = (1+r)^{-1}$. The variances are strictly positive.	Variances: C^2/U^2 ; ratios: unitless; primitive/derived

Continued on the next page

Table 3 (continued)

Symbol	Definition and economic role	Units; status
u_0, σ_u^2	Transitory hedging flow and its strictly positive variance. The transitory-hedging comparison uses the standard deviation $u = \sigma_u$.	U, U^2 ; primitive
x_{i0}, y_0, p_0	Initial client order, aggregate order and settled price: $x_{i0} = \beta s_i$, $y_0 = \sum_i x_{i0} + u_0$, $p_0 = \kappa y_0$.	Orders: U; price: C/U; derived/record
β, κ	Positive own-signal order loading and initial price impact; jointly solved by (2.6).	U^2/C , C/U^2 ; derived
$\alpha_i, \gamma, \Gamma_J$	Mandate benefit coefficient from (2.4); predetermined target/delay-cost curvature $\gamma \geq 0$; initial-order wedge $\Gamma_J = \gamma/\ell_0$.	α_i : C/U; curvatures: C/U^2 ; maintained/derived
$T_i, x_{it}, \mathcal{R}_{it}$	Signed target, planned trades for $t=0, \dots, J$, and unexecuted target $\mathcal{R}_{it} = T_i - \sum_{h=0}^t x_{ih}$. Completion gives $\mathcal{R}_{iJ} = 0$.	U; derived; planned orders have record counterparts
$\theta, \ell_t, g_j, \psi$	Finite schedule in (2.5): $x_{it} = \ell_t T_i$, $\sum_{t=0}^J \ell_t = 1$, $g_j = \ell_j/\ell_0$, $\psi = J^{-1} \sum_{j=1}^J g_j$. For $\gamma > 0$, $\cosh \theta = 1 + \gamma/(4J\kappa)$. At $\gamma = 0$, $\ell_t = 1/(J+1)$ and $g_j = \psi = 1$.	Unitless; derived
<i>Pre-reset positions and information</i>		
q_0, q_1, Q, ς	Entering and pre-reset signed holdings, pre-reset position size $Q = q_1 $, and direction $\varsigma = \text{sgn}(q_1)$ for $Q > 0$. The private signal is s_i ; the position sign is ς .	Holdings: U; sign: unitless; record/derived
e_0, B_1, e	Entering equity, post-trade net cash and settled pledgeable pre-reset account equity: $B_1 = e_0 - p_0 x_{i0}$ and $e = e_0 + q_0 p_0$, as in (2.8).	C; record/accounting
$\mathcal{D}_0, \mathcal{G}_i, \mathcal{H}_i$	Dealer information $\sigma(y_0)$; own-account information $\sigma(q_0, e_0, x_{i0})$; full reset information $\mathcal{G}_i \vee \sigma(y_0, p_0)$. Common parameters are known.	Information sets; maintained
P_i, μ_i, V_i	Peer initial flow $P_i = \sum_{\ell \neq i} x_{\ell 0}$; prior mean $\mu_i = \mathbb{E}[P_i \mathcal{G}_i]$ and variance $V_i = \text{Var}(P_i \mathcal{G}_i)$ in (3.1).	P_i, μ_i : U; V_i : U^2 ; derived
Y_i, ι_i	Client-netted aggregate order flow $Y_i = y_0 - x_{i0}$ and its client-information innovation $\iota_i = Y_i - \mu_i$. Conditional on $\mathcal{G}_i$, its variance is $V_i + \sigma_u^2$.	U; record-based/model-derived
K_i, R_i	Bayesian weight on order flow $K_i = V_i/(V_i + \sigma_u^2)$ and posterior peer variance $R_i = V_i \sigma_u^2/(V_i + \sigma_u^2)$. Thus $\mathbb{E}[P_i \mathcal{H}_i] = \mu_i + K_i \iota_i$.	K_i : unitless; R_i : U^2 ; derived
$\Theta_i, A_i^0, z_i, A_i$	Realized signed pressure $\Theta_i = -\varsigma \kappa P_i$; own prior $A_i^0 = -\varsigma \kappa \mu_i$; order-flow innovation $z_i = -\varsigma \kappa \iota_i$; forecast $A_i = A_i^0 + K_i z_i = \mathbb{E}[\Theta_i \mathcal{H}_i]$.	C/U; realized state/derived forecasts
$\mathcal{V}_p, S_i^2, w_i$	Price-update variance $\mathcal{V}_p = \text{Var}(p_0)$; client-information innovation variance $S_i^2 = \text{Var}(z_i \mathcal{G}_i) = \kappa^2 (V_i + \sigma_u^2)$; standardized state $w_i = z_i/S_i$.	Variances: C^2/U^2 ; w_i : unitless; derived
$\chi, \bar{\chi}_N$	Informed share of initial aggregate flow variance and its zero-delay-cost ceiling, defined in (3.10). χ describes trading composition; K_i describes the broker's conditional public update.	Unitless; derived

Continued on the next page

Table 3 (continued)

Symbol	Definition and economic role	Units; status
n_o, n_u	Numbers of observed accounts, including the focal account, and unobserved accounts: $n_u = N - n_o$. Coverage variance $V(n_o)$ and loading $K(n_o)$ appear in (3.5).	Counts; maintained information coverage
<i>Execution service and conditional losses</i>		
r_j, R_{j-1}^r	Committed forced-execution weight $r_j = 1/J$ and earlier cumulative weight $R_{j-1}^r = \sum_{h < j} r_h$. These weights apply to D , while ℓ_t, g_j describe ordinary execution.	Unitless; maintained
λ_0, d	Entering execution-cost level and its nonnegative per-unit loading. Their product $d\lambda_0$ is the service's baseline concession.	λ_0 : C/U; d : unitless; maintained
η_c, v_c, b_J	Nonnegative service slopes for earlier absorbed units and the current slice. Average size coefficient $b_J = \eta_c(J-1)/(2J) + v_c/J$.	C/U ² ; maintained slopes/derived average
$\epsilon_j, \Omega, v_\epsilon$	Centered Gaussian execution disturbances, independent of initial primitives; $\Omega \geq 0$ is their joint covariance. Averaged disturbance variance is $v_\epsilon = J^{-1} \mathbf{1}' \Omega \mathbf{1}$.	ϵ_j : C/U; variances: C ² /U ² ; maintained/derived
$U_j(x), p_j(x), U(x)$	Actual per-unit concession and execution price for a candidate reduction x : $p_j(x) = p_0 - \varsigma U_j(x)$, $U(x) = J^{-1} \sum_j U_j(x)$. The service in (2.11) uses realized peer flow.	C/U; maintained service/realized execution
$a, c(x), \mu_U$	Conditional cost intercept $a = d\lambda_0 + \psi A_i$, forecast per-unit cost $c(x) = a + b_J x$, and pre-reset-position-size full-closeout mean $\mu_U = c(Q)$.	C/U; derived forecasts
Z, σ_U^2, C_P	Averaged execution residual: $U(x) = a + b_J x + Z$, $Z \mathcal{H}_t \sim N(0, \sigma_U^2)$. Here $C_P = \kappa^2 R_t$ and $\sigma_U^2 = \psi^2 C_P + v_\epsilon$.	Z : C/U; variances: C ² /U ² ; derived
$\bar{\ell}, \mathcal{L}(t), t_Q, B$	Allocated expected execution-loss limit; normal stop-loss transform $\mathcal{L}(t) = \varphi(t) - t\bar{\Phi}(t)$. For $Q, \sigma_U, \bar{\ell} > 0$, cutoff $t_Q = \mathcal{L}^{-1}(\bar{\ell}/(Q\sigma_U))$ and buffer $B = \sigma_U t_Q$. B may be nonpositive.	Budget: C; transform/cutoff: unitless; B : C/U; maintained/derived
m, M_σ	Nonnegative pre-reset-position-size rate $m = \max\{0, a + b_J Q + B\}$, locked for the window. At a positive rate, $M_\sigma = \partial m / \partial \sigma_U = \varphi(t_Q) / \bar{\Phi}(t_Q)$ with mean, size and budget fixed.	m : C/U; M_σ : unitless; derived/request record
<i>Reduction, actual collateral, and the two books</i>		
$D, Q-D, C$	Selected forced quantity, residual pre-reset position size and ordinary-program authorization indicator. $0 \leq D \leq Q$; $C \in \{0, 1\}$ follows Proposition 6 and Section 2.5.	Quantities: U; C : unitless; derived/decision records
$F(x), E^f(x), \xi$	Forecast support slack $F(x) = e - mQ + (m-a)x - b_J x^2$; forecast equity $E^f(x) = e - xc(x)$; initial slack $\xi = e - mQ$.	C; derived forecasts
δ, Δ_F, D_-	Deficiency $\delta = mQ - e$; quadratic discriminant $\Delta_F = (m-a)^2 - 4b_J \delta$; smaller root candidate D_- in (4.2). The linear $b_J = 0$ rule is stated separately there.	C, C ² /U ² , U; derived

Continued on the next page

Table 3 (continued)

Symbol	Definition and economic role	Units; status
H, Ξ, Q_m	Regular relief slope $H = m - a - 2b_J D > 0$; marginal cost ratio $\Xi = (a + 2b_J D)/m$; partial rate sensitivity $Q_m = E^f(D)/m^2 = (Q - D)/m$ at a regular root, holding forecast equity fixed.	H : C/U; Ξ : unitless; Q_m : U ² /C; derived
f_j, f, x_j^P, q_j^P	Signed forced slice $f_j = -\varsigma r_j D$ and total $f = \sum_j f_j$; actual ordinary trade $x_j^P = C g_j x_{j0}$ and cumulative program holding $q_j^P = \sum_{h \leq j} x_h^P$.	U; derived/trade records
L_j, L, C_{exec}	Realized forced invoice $L_j = \sum_{h \leq j} r_h D U_h(D)$, $L = L_J$; ordinary concession $C_{\text{exec}} = \kappa \sum_j (x_j^P)^2$. Set $L_j = 0$ directly when $D = 0$.	C; realized accounting
H_j^P, E_j^P, E_j^I	Program net funding draw $H_j^P = \sum_{h \leq j} (p_0 + \kappa x_h^P) x_h^P$; program equity $E_j^P = -H_j^P + p_0 q_j^P$; pre-reset account equity $E_j^I = e - L_j$. A negative draw is program cash.	C; accounting/record
M^*, M_j	Intended residual requirement $M^* = m(Q - D)$ when $Q - D > 0$, and zero directly at full wind-down; actual restricted principal $M_j = \min\{M^*, (E_j^I)^+\}$.	C; requirement/actual collateral record
$k, \Delta_j, F_{\text{coll}}$	Positive opportunity-cost rate; interval weights $\sum_j \Delta_j = 1$ before execution; single collateral charge $F_{\text{coll}} = k \sum_j \Delta_j M_{j-1}$. Principal remains a client asset.	Rate/weights: unitless; cost: C; primitive/derived
E_D, C_D, S_D	End-window pre-reset account equity $E_D = e - L$; collateral deficiency $C_D = [M^* - E_D]^+$; benchmark-marked deficit $S_D = [-E_D]^+$.	C; realized accounting
q^P, q_2, Q_2, δ_v	$q^P = q_j^P$; actual terminal signed holding $q_2 = q_1 + f + q^P$ and size $Q_2 = q_2 $; value change $\delta_v = v - p_0$. Ordinary continuation can change the holding's sign.	Holdings: U; δ_v : C/U; realized accounting
$\mathcal{N}, \mathcal{N}_I, \mathcal{N}_P$	Combined contractual terminal resources $\mathcal{N} = e - L - C_{\text{exec}} + q_2 \delta_v$. On interruption, $\mathcal{N}_I = E_D + (q_1 + f) \delta_v$. On continuation, $\mathcal{N}_I = e + q_1 \delta_v$, $\mathcal{N}_P = q^P \delta_v - C_{\text{exec}}$.	C; contractual accounting
Π_F, Π_P	Forced providers' and ordinary counterparties' incremental terminal trading profits: $\Pi_F = L - f \delta_v$, $\Pi_P = C_{\text{exec}} - q^P \delta_v$, before their own resource/financing costs.	C; trading accounts
$\mathcal{A}, H_h, R_h, c, \mathcal{S}$	Gross distributable assets; creditor h 's remaining claim and actual recovery, $0 \leq R_h \leq H_h$; nonnegative client payout; aggregate unpaid claims $\mathcal{S} = \sum_h (H_h - R_h)$. Section 5.3 states the recovery priorities.	C; settlement accounting/records
S_{full}	Credit shortfall after full wind-down: $[QU(Q) - e]^+$ for $Q > 0$. At zero exposure, negative pre-reset cash leaves $(-e)^+$ directly.	C; realized credit outcome
<i>Provider inventory and the subsequent trader</i>		
X_0, X_1, λ_1	Entering and post-window gross inventory load in the relevant pool: $X_1 = X_0 + D$. Before later placement, $\lambda_1 = \lambda_0 + \omega_X D$, with $\lambda_0 = \lambda_{\min} + \omega_X X_0$.	Inventory: U; cost levels: C/U; accounting/maintained map

Continued on the next page

Table 3 (continued)

Symbol	Definition and economic role	Units; status
$\omega_X, \lambda_{\min}$	Nonnegative inventory-to-cost slope and baseline cost level. ω_X is independent of η_c, ν_c and γ ; it changes subsequent traders' cost levels.	C/U ² , C/U; maintained
$\varphi_X, \bar{F}, \mathcal{T}(X)$	Placement retention parameter $\varphi_X \in [0, 1]$, nonnegative capacity and remaining burden $\mathcal{T}(X) = X - \min\{(1 - \varphi_X)X, \bar{F}\}$. Thus $X^{\text{next}} = \mathcal{T}(X_1)$.	Retention: unitless; capacity/burden: U; maintained/derived
y_L, s_L, p_L, U_L, d_L	Specified later quantity and direction, benchmark, average concession and nonnegative inventory-cost exposure. Quantity and service conditions are matched in Proposition 7.	Quantity: U; prices: C/U; sign/loading: unitless; conditioned service/records
I_L^X	Invoice increment attributable to focal inventory: $I_L^X = y_L d_L \omega_X D$ before inventory reallocation. After placement, the corresponding burden increment is $\mathcal{T}(X_0 + D) - \mathcal{T}(X_0)$.	C; conditional attribution

The loss inversion uses normal density φ , distribution Φ and survival $\bar{\Phi} = 1 - \Phi$. Its t_Q formula requires $Q > 0$, $\sigma_U > 0$ and $0 < \bar{\ell} < \infty$; deterministic-cost and zero-budget branches are in Section 3.2. At $Q = 0$, set $m = D = f_j = L_j = 0$ without a position sign; nonnegative equity, including zero, permits continuation, whereas negative cash cancels it. Any positive reduction cancels the whole remaining client execution program.

On the positive-margin interior branch, Ξ and Q_m apply: regularity requires $\Xi < 1$, local convergence of the literal quantity iteration requires $|\Xi| < 1$; see (4.6). Auxiliaries h, t_χ in (3.11) and $P_0, \vartheta_i, \mathcal{Z}(\chi)$ in (3.14) are primitive-based. Hedging-comparison primes are total χ derivatives along the pricing solution; w_* in (3.16) is a standardized signed threshold on the stated positive-rate domain.

A.1 Forecasts, Units, and Account Balances

Before realization, $\mathcal{V}_p$ measures public price dispersion; S_i^2 conditions signed-innovation dispersion on own information. After public updating, $\sigma_U^2 = \text{Var}(U(x) | \mathcal{H}_i)$ measures residual per-unit risk, constant across innovations and candidate sizes within one Gaussian economy. A committed positive reduction's invoice variance is $D^2 \sigma_U^2$. Shared peer uncertainty gives

$$\text{Cov}(U_j(x), U_h(x) | \mathcal{H}_i) = \kappa^2 R_i g_j g_h + \Omega_{jh}.$$

Peer uncertainty contributes even with independent service disturbances.

Slice count changes both the finite policy and size coefficient:

$$b_J = \frac{\eta_c}{2} + \frac{\nu_c - \eta_c/2}{J}, \quad b_{J+1} - b_J = -\frac{\nu_c - \eta_c/2}{J(J+1)}.$$

Horizon changes require recomputing pricing, the full g_j schedule and window disturbance covariance. Section 3 comparisons fix J .

Calibrated at pre-reset Q and entering λ_0 , the locked m gives residual requirement $m(Q - D)$; requirement relief need not release posted principal. Invoices change pre-reset account equity E_j^I and restricted principal M_j . Program finance covers principal and concessions; creditors first claim program positions and proceeds, with further claims subordinate until pre-reset obligations are discharged (Section 2.5). The books satisfy

$$E_j^I + E_j^P = e - L_j - \kappa \sum_{h \leq j} (x_h^P)^2.$$

Program costs reduce combined resources while pre-reset collateral stays protected. Actual restricted principal determines the single opportunity charge F_{coll} , outside legal claims and default balances.

Payment and priority-respecting recovery accounting gives

$$\mathcal{N} + \Pi_F + \Pi_P = e + q_1 \delta_v, \quad c - \mathcal{S} = \mathcal{N}.$$

New trading payments and opposite risky holdings cancel in the first identity. The second uses actual priority-respecting recoveries, retaining negative program resources and unpaid claims. End-window covenant deficiency C_D , marked deficit S_D and terminal unpaid claims $\mathcal{S}$ remain distinct.

A.2 Records, Forecasts, and Empirical Proxies

Table 4 maps records to Section 6 comparisons. Contract and service coefficients require measurement from terms or explicit-model estimation; no dataset or estimator is supplied.

Table 4
Measurement objects and information timing

Object	Records and timing	Model-based construction or proxy restriction
Own information and order-flow innovation	Own order and account records; public initial flow and settled price, all available before the reset.	$s_i = x_{i0}/\beta$ follows the centered common-coefficient law. Outside it, order-based signal extraction is a proxy. Construct μ_i from own information before the public outcome, then $\iota_i = Y_i - \mu_i$. Price and flow are collinear public observations here.
Variance scales	A forecast of public price dispersion before the initial outcome; client-information innovation forecasts; conditional closeout-risk forecasts before the reset.	Preserve $\mathcal{V}_p$, S_i^2 and σ_U^2 separately. Return volatility and flow variance target different scales; either requires a model to recover conditional execution risk. A realized price move cannot supply its own ex-ante variance.

Continued on the next page

Table 4 (continued)

Object	Records and timing	Model-based construction or proxy restriction
Ordinary persistence	Committed schedules or separately identified ordinary-flow forecasts available before the reset, over the same closeout horizon.	g_j and ψ describe ordinary continuation. Pre-reset autocorrelation is a proxy affected by common news and correlated programs. Subsequent forced flow is an outcome, unavailable for the reset's contemporaneous forecast.
Margin request and collateral	Locked per-unit request after settlement; pre-reset position size and pledgeable equity before reset; actual restricted principal after each invoice.	Separate the execution-loss bucket from other credit requirements and current settlement transfers. m , M^* and M_j are respectively a rate, demanded residual amount and actual restricted asset.
Selected reduction	Broker-directed orders after the reset, alongside the canceled or authorized ordinary mandate, voluntary orders and any top-up.	Classify slack, regular interior and full wind-down before applying a local derivative. Match settled Q, e ; anticipatory trading and collateral additions can change selection into those branches.
Realized shortfalls	Invoices and pre-reset account equity at window end; asset value and actual creditor recoveries at maturity.	Use C_D for residual collateral deficiency, S_D for a marked deficit, and S for unpaid contractual claims at settlement. A full sale can leave debt; the Gaussian half-probability result concerns end-window deficiency at a regular cure.
Information coverage	Relevant orders/exposures actually visible to the broker before the reset.	Relationship counts proxy coverage and can also reflect competition, selection and concentration. Interpret $K(n_o)$ at fixed trading population and policies. Full coverage leaves known pressure in the loss forecast.
Subsequent-trader execution costs	The same liquidity-provider pool's inventory after the window, any inventory reallocation, and the subsequent trader's dated quantity, direction, benchmark, schedule and invoice.	Match the service and other execution conditions. Before placement use D ; afterward use $\mathcal{T}(X_0 + D) - \mathcal{T}(X_0)$. Common funding/liquidity shocks, other sales and provider choice can also affect the bill.

Composition fixes public variance $\mathcal{V}_p$, own signal and signed price-unit innovation z along (3.12). The adverse–favorable margin difference requires positive endpoint rates; Corollary 1 gives the separate low-share level domain. Hedging instead fixes signal quality and standardized innovation $w = z/S$, allowing public variance to move. Zero-delay-cost neutrality matches own signal and normalized state, not raw flow or own order.

With later quantities and service conditions fixed (Proposition 7), placement leaves focal burden increment $\varphi_X D$ only if both states' capacity is nonbinding; binding capacity or crossings use the table's exact burden difference. The $\omega_X =$

0 shutdown preserves the sale at fixed entering λ_0 within the nonnegative cost domain; fixing baseline cost and entering burden instead can change the current contract. The inventory-to-cost coefficient is independent of current service slopes and does not measure slope-based depth.

Appendix B. Proofs

The proofs begin with the finite execution mandate and initial pricing, then derive the broker's conditional loss law and margin comparisons. The remaining subsections establish the reduction rule, realized account balances, and subsequent-trader incidence. Each derivative is taken on its stated branch. Zero exposure, the collateral floor and termination boundaries are treated directly.

B.1 Finite Execution and Initial Pricing

Proof of Proposition 1. Fix a signed target T and suppress the account index. Let $V_t(R)$ be the minimum execution and delay cost from a residual target R immediately before trade t . Terminal completion requires $x_J = R$, hence $V_J(R) = \kappa R^2$. Write $V_t(R) = h_t R^2/2$ and set $c_J = \gamma/J$. Backward induction gives, for $t = J-1, \dots, 0$,

$$V_t(R) = \min_x \left\{ \kappa x^2 + \frac{1}{2}(c_J + h_{t+1})(R-x)^2 \right\}, \quad (B.1)$$

$$f_t = \frac{2\kappa}{2\kappa + c_J + h_{t+1}}, \quad x_t = (1 - f_t)R, \quad h_t = \frac{2\kappa(c_J + h_{t+1})}{2\kappa + c_J + h_{t+1}}.$$

The minimized quadratic is strictly convex. Set $f_J = 0$ and $h_J = 2\kappa$. The resulting weights are

$$\ell_t = (1 - f_t) \prod_{h=0}^{t-1} f_h > 0, \quad \sum_{t=0}^J \ell_t = 1 - \prod_{h=0}^J f_h = 1.$$

The empty product is one. Thus $x_t = \ell_t T$, and the residual after the final trade is zero, including for a negative target. The inner cost is $h_0 T^2/2$, with $h_0 = 2\kappa \ell_0 > 0$. The outer objective is $\alpha_i T - (\gamma + h_0) T^2/2$; its unique maximizer is $T_i = \alpha_i / (\gamma + 2\kappa \ell_0)$.

For the closed form, use residual coordinates $\mathcal{R}_{-1} = T$, $\mathcal{R}_J = 0$ and $x_t = \mathcal{R}_{t-1} - \mathcal{R}_t$. The first-order conditions of the strictly convex inner problem are

$$\mathcal{R}_{t-1} - \left(2 + \frac{\gamma}{2J\kappa} \right) \mathcal{R}_t + \mathcal{R}_{t+1} = 0, \quad t = 0, \dots, J-1.$$

When $\gamma > 0$, their unique solution is

$$\mathcal{R}_t = T \frac{\sinh((J-t)\theta)}{\sinh((J+1)\theta)}, \quad \cosh \theta = 1 + \frac{\gamma}{4J\kappa}. \quad (B.2)$$

Taking consecutive differences and dividing the coefficient of T in x_j by that in x_0 gives

$$g_j = \frac{\sinh((J+1-j)\theta) - \sinh((J-j)\theta)}{\sinh((J+1)\theta) - \sinh(J\theta)} = \frac{\cosh((J+\frac{1}{2}-j)\theta)}{\cosh((J+\frac{1}{2})\theta)}.$$

These are coefficient ratios, so they remain defined when $\alpha_i=0$ and all realized trades vanish. The weights sum to one, giving $\ell_0 = (1 + \sum_j g_j)^{-1}$ and the schedule in (2.5).

For $j \geq 1$, put $a_j = J + \frac{1}{2} - j$ and $a_0 = J + \frac{1}{2}$. Then

$$\frac{d \log g_j}{d\theta} = a_j \tanh(a_j \theta) - a_0 \tanh(a_0 \theta) < 0,$$

because $a \mapsto a \tanh(a\theta)$ is strictly increasing for positive a, θ . At fixed positive γ , increasing κ reduces θ and raises every g_j . Also $0 < g_J < \dots < g_1 < 1$, so $0 < \psi < 1$ and (2.7) follows. At $\gamma=0$, the recursion instead gives $h_t = 2\kappa/(J-t+1)$, all $J+1$ weights equal $1/(J+1)$, and $g_j = \psi = 1$. The target is $(J+1)\alpha_i/(2\kappa)$, finite for every $\kappa > 0$.

It remains to establish the positive pricing fixed point. Eliminating T_i gives $(2\kappa + \Gamma_J)x_{i0} = \alpha_i$, where $\Gamma_J = \gamma/\ell_0$. Substituting (2.4) yields the first equation of (2.6). The independent Gaussian primitives imply

$$\text{Cov}(v, y_0) = N\beta\sigma_v^2, \quad \text{Var}(y_0) = \beta^2\sigma_v^2 N(N+r) + u^2, \quad u := \sigma_u > 0.$$

Their Gaussian projection gives the second pricing equation.

To solve the two equations jointly, define the positive scalar ζ , share χ and price representation

$$\begin{aligned} \zeta &= \frac{\beta\sigma_v\sqrt{N(N+r)}}{u}, & \chi &= \frac{\zeta^2}{1+\zeta^2}, & \bar{\chi}_N &= \frac{N+r}{N+1+2r}, \\ \beta &= \frac{u\zeta}{\sigma_v\sqrt{N(N+r)}}, & \kappa &= \frac{\sigma_v\sqrt{N/(N+r)}}{u} \frac{\zeta}{1+\zeta^2}, & & \\ \frac{\chi}{\bar{\chi}_N} + G &= 1, & G &= (1+r)\beta\Gamma_J. \end{aligned} \tag{B.3}$$

The last equality follows from $\kappa\beta = \chi/(N+r)$. For $\gamma > 0$, write $v = \gamma/(2J\kappa)$ and $F_E(v) = \ell_0$. The subscript distinguishes this execution-weight function from the liquidation slack $F(x)$. Normalizing (B.1) by 2κ gives

$$\widehat{h}_J = 1, \quad \widehat{h}_t = \frac{v + \widehat{h}_{t+1}}{1 + v + \widehat{h}_{t+1}}, \quad F_E = \widehat{h}_0. \tag{B.4}$$

Backward composition with the increasing concave map $x/(1+x)$ makes F_E smooth, increasing and concave on $[0, \infty)$, strictly increasing for $J \geq 1$.

Moreover $F_E(0) = 1/(J+1)$ and $F_E < 1$. Concavity implies $vF'_E(v) \leq F_E(v) - F_E(0) < F_E(v)$. Consequently

$$\varepsilon_E := \frac{vF'_E}{F_E} \in (0, 1) \quad (v > 0).$$

At fixed primitive parameters, v is proportional to $\zeta + \zeta^{-1}$ and $G = (1+r)\beta\gamma/F_E$. Hence

$$\frac{\partial \log G}{\partial \log \zeta} = 1 - \varepsilon_E \frac{\zeta^2 - 1}{\zeta^2 + 1} > 0.$$

The left side of the scalar equation in (B.3) increases strictly from zero to infinity: at large ζ , $G \geq (1+r)\beta\gamma \rightarrow \infty$. It therefore crosses one exactly once. At $\gamma=0$, $G=0$ and the increasing share crosses $\bar{\chi}_N \in (0, 1)$ exactly once. In both cases the representations recover one positive (β, κ) . All initial orders then satisfy the mandate and Bayesian pricing equations. The predetermined γ remains an input to this fixed point. ■

B.2 Own Information, Public Updating and Coverage

Proof of Proposition 2 and (3.5). Entering account states are uninformative about the Gaussian primitives, and the known positive coefficient β makes the own order reveal s_i . Gaussian conditioning gives

$$v | \mathcal{G}_i \sim N\left(\rho s_i, \frac{\sigma_v^2 r}{1+r}\right).$$

For $n = N - 1$, $P_i = \beta n v + \beta \sum_{\ell \neq i} \eta_\ell$. The peer errors remain independent of (v, s_i) . Its mean and variance conditional on $\mathcal{G}_i$ are therefore exactly μ_i, V_i in (3.1). Given that information, the centered pair $(P_i - \mu_i, \iota_i)$ has covariance matrix

$$\begin{pmatrix} V_i & V_i \\ V_i & V_i + u^2 \end{pmatrix}.$$

The residual $P_i - \mu_i - K_i \iota_i$ has zero covariance with ι_i and variance $V_i - V_i^2/(V_i + u^2) = R_i$. Joint Gaussianity gives its independence from the observation and establishes (3.3). Price is measurable from flow, so it adds no second public observation.

At $Q > 0$ the own account determines ς . Multiplying the posterior by $-\varsigma\kappa$ gives mean $A_i^0 + K_i z_i$ and variance $\kappa^2 R_i$. The coefficients and variance are constant across public realizations in a given economy, so $A_{i,z} = K_i$. For the two-account example with $\beta = \sigma_v^2 = \sigma_\eta^2 = u^2 = 1$, $\mu_i = s_i/2$, $V_i = 3/2$, $K_i = 3/5$ and $R_i = 3/5$. Thus the mean is $s_i/5 + 3Y_i/5$, as printed.

For coverage, let $\mathcal{O}$ contain the n_o observed accounts, including i , and let $n_u = N - n_o$. Conditioning on their recovered signals gives

$$\mathbb{E}[v | (s_\ell)_{\ell \in \mathcal{O}}] = \frac{\sum_{\ell \in \mathcal{O}} s_\ell}{r + n_o}, \quad \text{Var}(v | (s_\ell)_{\ell \in \mathcal{O}}) = \frac{\sigma_v^2 r}{r + n_o}.$$

The unknown aggregate order consequently has variance $V(n_o) = \beta^2 \sigma_v^2 [n_u^2 r / (r + n_o) + n_u r]$ and mean $\beta n_u \sum_{\ell \in O} s_\ell / (r + n_o)$. Subtract all observed orders from y_0 and project on the remaining order-flow innovation. This gives loading $K(n_o)$ in (3.5) and posterior unknown-order variance $V(n_o)u^2 / [V(n_o) + u^2]$. Add the observed peer orders back to obtain the total peer forecast.

As n_o rises, n_u falls and $r + n_o$ rises, so $V(n_o)$ decreases strictly until it reaches zero. Both $V/(V + u^2)$ and $Vu^2/(V + u^2)$ increase strictly in V . Thus their coverage responses have the stated signs. At full coverage, unknown variance and incremental Bayesian weight on order flow are zero, while the sum of observed peer orders can be nonzero. At $N=1$, $P_i=0$ directly. Neither boundary removes independent execution disturbances or the service's deterministic costs. Peers' continuation commitments are fixed before signals; no survival event is added to these conditioning sets. ■

B.3 Execution Loss and the Nonnegative Rate

Proof of (3.6) and Lemma 1. For uniform forced weights, $R_{j-1}^r = (j-1)/J$. Thus

$$\frac{1}{J} \sum_{j=1}^J R_{j-1}^r = \frac{J-1}{2J}, \quad \frac{1}{J} \sum_{j=1}^J r_j = \frac{1}{J}, \quad \frac{1}{J} \sum_{j=1}^J g_j = \psi.$$

Averaging (2.11) gives $U(x) = a + b_J x + Z$, with a, b_J, Z as stated in (3.6). Conditional on $\mathcal{H}_i$, the shared peer residual and the independent Gaussian disturbance vector yield

$$\text{Cov}(U_j(x), U_h(x) | \mathcal{H}_i) = \kappa^2 R_i g_j g_h + \Omega_{jh}.$$

Summing all entries and dividing by J^2 gives $\sigma_U^2 = \psi^2 \kappa^2 R_i + J^{-2} \mathbf{1}' \Omega \mathbf{1}$. This remains the conditional law at a selected D , since D is chosen from reset information before these residuals arrive. The finite horizon also gives $b_{J+1} - b_J = -(v_c - \eta_c/2) / [J(J+1)]$ by subtraction. A change in J additionally changes the remaining client execution schedule, pricing and the specified disturbance vector.

For $Q, \sigma_U > 0$, set $\mu_U = a + b_J Q$ and let W be standard normal. Direct integration gives

$$\mathbb{E}[(W-t)^+] = \int_t^\infty (w-t) \varphi(w) dw = \varphi(t) - t \bar{\Phi}(t) = \mathcal{L}(t). \quad (\text{B.5})$$

Its derivative is $-\bar{\Phi}(t) < 0$. As $t \rightarrow -\infty$, $\mathcal{L}(t) \geq -t \rightarrow \infty$; as $t \rightarrow \infty$, dominated convergence gives zero. It is therefore a strictly decreasing bijection from the real line onto $(0, \infty)$. The conditional expected uncovered loss is

$$Q \sigma_U \mathcal{L}\left(\frac{m - \mu_U}{\sigma_U}\right).$$

For $0 < \bar{\ell} < \infty$, its feasible rates form the half-line $m \geq \mu_U + \sigma_U t_Q$. Intersecting with $m \geq 0$ gives (3.7). At a positive optimum the expected-loss inequality binds; in the strict zero-rate region it is slack. Since $\mathcal{L}(0) = 1/\sqrt{2\pi}$ and its inverse is decreasing, $B > 0$ precisely on the budget domain printed in Section 3.2; equality gives zero buffer.

Holding $Q, \bar{\ell}$ fixed, differentiate $\mathcal{L}(t_Q) = \bar{\ell}/(Q\sigma_U)$ to obtain

$$\frac{\partial B}{\partial \sigma_U} = t_Q + \frac{\mathcal{L}(t_Q)}{\Phi(t_Q)} = \frac{\varphi(t_Q)}{\Phi(t_Q)} = M_\sigma > 0.$$

The positive-rate mean derivative is one. Within an economy, only the mean changes with z_i , at rate ψK_i , proving (3.8). Inside the strict floor all these local rate derivatives vanish. At equality with the floor, the rate is the positive part of a smooth unrestricted rate and has its corresponding one-sided derivatives.

For later use, the full positive-rate derivative along a differentiable path τ is

$$\begin{aligned} m_\tau = & a_\tau + Q(b_J)_\tau + \left(b_J + \frac{\bar{\ell}}{Q^2 \Phi(t_Q)} \right) Q_\tau \\ & + M_\sigma(\sigma_U)_\tau - \frac{\bar{\ell}_\tau}{Q \Phi(t_Q)}. \end{aligned} \quad (\text{B.6})$$

It follows by differentiating the same inverse identity, without fixing its size or budget arguments. In particular,

$$\begin{aligned} a_\tau = & d_\tau \lambda_0 + d(\lambda_0)_\tau + \psi_\tau A_i + \psi(A_i)_\tau, \\ 2\sigma_U(\sigma_U)_\tau = & 2\psi\psi_\tau \kappa^2 R_i + \psi^2(\kappa^2 R_i)_\tau + (v_\epsilon)_\tau. \end{aligned}$$

At fixed $Q, \bar{\ell}$, these formulas reduce to (3.9).

If $\sigma_U = 0$, the constraint is $Q(\mu_U - m)^+ \leq \bar{\ell}$, giving $m = \max\{0, \mu_U - \bar{\ell}/Q\}$ for $Q > 0$ and finite nonnegative budget. If $Q > 0$, $\sigma_U > 0$ and $\bar{\ell} = 0$, every finite threshold leaves a positive Gaussian expected excess, so no finite rate is feasible. If the budget is infinite, the restriction is absent and $m = 0$. At $Q = 0$, collateral and forced invoices are zero directly; no per-unit inversion or position sign is evaluated. These boundaries use their direct constraints rather than the differentiable formula above. ■

B.4 Attainable Composition Comparisons

Proof of Proposition 3 and corollary 1. First derive the scales in (3.11). Write $W_0 = \beta^2 \sigma_v^2 N(N+r)$ for informed aggregate-flow variance. The dealer's Gaussian projection gives

$$\mathcal{V}_p = \kappa^2(W_0 + u^2) = \frac{\sigma_v^2 N}{N+r} \chi = \sigma_v^2 - \text{Var}(v | y_0).$$

The own-conditioned variance in (3.1) satisfies $V_i = hW_0$, where $h = (N-1)r/[N(1+r)]$. Since $\chi = W_0/(W_0 + u^2)$, substitution gives

$$K = \frac{h\chi}{t_\chi}, \quad S^2 = \mathcal{V}_p t_\chi, \quad C_P = \kappa^2 R_i = \mathcal{V}_p K(1-\chi), \quad t_\chi = 1 - (1-h)\chi.$$

For $N \geq 2$ and $r > 0$, $0 < h < 1$ and $0 < K < \chi < 1$.

Attainability. Put $c_p = \overline{\mathcal{V}}_p / \sigma_v^2 > 0$. On the proposed path $\chi = c_p(N+r)/N < \bar{\chi}_N$. Rearranging the first pricing equation using $\kappa\beta = \chi/(N+r)$ gives

$$\frac{2Jv}{F_E(v)} = C(r), \quad C(r) = \frac{(N+r)/\chi - (N+1+2r)}{1+r} = \frac{N/c_p - N+1}{1+r} - 2 > 0. \quad (\text{B.7})$$

The map $v/F_E(v)$ increases strictly, with elasticity $1 - \varepsilon_E > 0$. It starts at zero and diverges, since $F_E < 1$. Thus exactly one positive v solves this equation. Define

$$\kappa = \frac{\gamma}{2Jv}, \quad \beta = \frac{\chi}{(N+r)\kappa}, \quad u = \frac{\sqrt{\mathcal{V}_p(1-\chi)}}{\kappa}. \quad (\text{B.8})$$

Together with $\sigma_\eta^2 = r\sigma_v^2$, these primitives satisfy both original pricing equations and the target public variance. Here u is strictly positive. Smoothness and strict monotonicity give a smooth path on every admissible open interval.

Since $C > 0$, the constant $N/c_p - N+1$ is positive. Consequently $C_r < 0$. Implicit differentiation gives

$$\begin{aligned} \chi_r &= \frac{c_p}{N} > 0, & h_r &= \frac{N-1}{N(1+r)^2} > 0, \\ \frac{\kappa_r}{\kappa} &= -\frac{C_r}{C(1-\varepsilon_E)} > 0, & \frac{\beta_r}{\beta} &= -\frac{\kappa_r}{\kappa}, \\ \frac{u_r}{u} &= -\frac{\chi_r}{2(1-\chi)} - \frac{\kappa_r}{\kappa} < 0, & \frac{\psi_r}{\psi} &= \frac{\varepsilon_E}{1-F_E} \frac{\kappa_r}{\kappa} > 0. \end{aligned} \quad (\text{B.9})$$

The last identity uses $\psi = (F_E^{-1} - 1)/J$. It is consistent with the strict increase of every g_j as κ rises. The posterior loading satisfies

$$\frac{K_r}{K} = (1-K) \left(\frac{h_r}{h} + \frac{\chi_r}{\chi(1-\chi)} \right) > 0.$$

Thus $(\psi K)_r > 0$ along the attainable primitive path.

Matched account states. Fix the designated signal $\bar{s}_i$, direction ς , pre-reset $Q > 0$, equity e and signed price-unit innovation z . In each economy the designated order-flow state is

$$\bar{y}_0 = \beta \bar{s}_i + \mu_i - \frac{\varsigma z}{\kappa}, \quad \bar{p}_0 = \kappa \bar{y}_0.$$

Choose that economy's deterministic entering endowments $q_0 = \varsigma Q - \beta \bar{s}_i$ and $e_0 = e - q_0 \bar{p}_0$. They are constants chosen before the realization; they are not

random functions imposed on realized signals. Conditional on the designated $(s_i, y_0) = (\bar{s}_i, \bar{y}_0)$, settlement then gives the matched pre-reset account state. Positive primitive variances imply a nonsingular, everywhere positive density for (s_i, y_0) ; the displayed Gaussian conditional laws provide its continuous regular conditional version, including at a designated zero signal. No event selecting accounts by a financing cap or peer survival is introduced. The two endpoints z_* and $-z_*$ can require different deterministic e_0 to preserve e . This construction also applies to the hedging and neutrality comparisons after substituting their chosen z .

Signed contrast. At a fixed point on the path, write the unrestricted rate as $M_0 + \psi K z$, where $M_0 = d\lambda_0 + \psi A_i^0 + b_J Q + B$. The two innovation endpoints have the same prior based on client information, conditional variance and buffer. When both rates are positive, their difference is $2z_*\psi K$; its derivative is positive for $z_* > 0$. This proves Proposition 3. At a floor endpoint the exact contrast is $[M_0 + \psi K z_*]^+ - [M_0 - \psi K z_*]^+$, so the linear contrast is used only on the stated positive-endpoint domain.

Level region. At $s_i = 0$ and fixed $z \geq 0$, $A = Kz$ and $a_r = (\psi K)_r z \geq 0$. Holding $\mathcal{V}_p$ fixed and differentiating $C_P = \mathcal{V}_p K(1 - \chi)$ gives

$$\frac{(C_P)_r}{C_P} = (1 - K) \frac{h_r}{h} + \frac{\chi_r(1 - K - \chi)}{\chi(1 - \chi)}.$$

If $\chi \leq 1/2$, then $K < \chi$ makes both terms positive. Therefore $(C_P)_r > 0$, and $\psi_r > 0$ implies $(\sigma_U)_r > 0$. At a positive rate, $m_r = a_r + M_\sigma(\sigma_U)_r > 0$, including at $z = 0$. Taking the positive part makes the nonnegative rate nondecreasing on any interval with these restrictions. All terms are smooth at an interior positive-rate state, so its strict margin derivative persists for own signals in a neighborhood of zero. This continuity statement concerns the margin derivative; the sign of the mean component need not remain nonnegative for every nearby own signal. At higher shares the second term above can change sign, and the proof supplies no unrestricted margin-level ordering. ■

B.5 Hedging Scale and the State Threshold

Proof of Proposition 4. Hold $N \geq 2, r > 0, \sigma_v^2, \gamma > 0, J$ fixed and vary $u = \sigma_u$. Use ζ, G, ε_E from (B.3)–(B.4). Write dots for total derivatives with respect to $\log u$. At fixed ζ , v is proportional to u , so $\partial \log G / \partial \log u = 1 - \varepsilon_E > 0$. At fixed u , its $\log \zeta$ derivative is $1 + \varepsilon_E(1 - 2\chi)$. Differentiating the scalar fixed point gives

$$\begin{aligned} \mathcal{D} &= \frac{2\chi(1 - \chi)}{\bar{\chi}_N} + G[1 + \varepsilon_E(1 - 2\chi)] > 0, \\ \log \dot{\zeta} &= -\frac{G(1 - \varepsilon_E)}{\mathcal{D}} < 0, \quad \dot{\chi} = 2\chi(1 - \chi) \log \dot{\zeta} < 0, \\ \log \dot{\kappa} &= -\frac{2(1 - \chi)}{\mathcal{D}} < 0, \quad \log \dot{\beta} = \frac{2(1 - \chi)(\chi / \bar{\chi}_N + \varepsilon_E G)}{\mathcal{D}} \in (0, 1). \end{aligned} \tag{B.10}$$

For the impact formula, start from $\log \kappa = -1 + (1 - 2\chi) \log \zeta$. Its negative numerator simplifies to $2(1 - \chi)(\chi/\bar{\chi}_N + G) = 2(1 - \chi)$. For the order loading, use $\log \beta = 1 + \log \zeta$; its displayed positive numerator and the strict negativity of $\log \zeta$ give the two bounds. These total elasticities include the changing finite-execution wedge.

Since χ decreases strictly with u , it can parameterize the path. Primes below denote total χ derivatives. Here h is fixed and $\mathcal{V}_p = P_0 \chi$, with $P_0 = \sigma_v^2 N / (N + r)$. Holding the own signal and standardized state $w = z/S$ fixed gives

$$A = \vartheta_i \chi + w \mathcal{Z}, \quad \mathcal{Z} = \frac{\sqrt{P_0} h \chi^{3/2}}{\sqrt{t_\chi}}, \quad C_P = P_0 \frac{h \chi^2 (1 - \chi)}{t_\chi},$$

where ϑ_i is defined in (3.14). In particular,

$$\begin{aligned} \frac{\mathcal{Z}'}{\mathcal{Z}} &= \frac{3}{2\chi} + \frac{1-h}{2t_\chi} > 0, & \frac{C'_P}{C_P} &= \frac{2}{\chi} - \frac{1}{1-\chi} + \frac{1-h}{t_\chi}, \\ C(\chi) &= \frac{(N+r)/\chi - (N+1+2r)}{1+r} > 0, & (B.11) \\ \frac{\kappa'}{\kappa} &= \frac{N+r}{(1+r)\chi^2 C(1-\varepsilon_E)} > 0, & \frac{\psi'}{\psi} &= \frac{\varepsilon_E}{1-F_E} \frac{\kappa'}{\kappa} > 0. \end{aligned}$$

The final two identities follow by differentiating $2Jv/F_E = C(\chi)$. Thus $(\psi \mathcal{Z})' > 0$. The prior based on client information is $\vartheta_i \chi$, so it also moves along this path. Differentiating the loss variance and the positive rate gives

$$\sigma'_U = \frac{2\psi\psi' C_P + \psi^2 C'_P}{2\sigma_U}, \quad m' = \vartheta_i(\psi + \chi\psi') + w(\psi \mathcal{Z})' + M_\sigma \sigma'_U.$$

Here $C_P > 0$ ensures $\sigma_U > 0$, even if $v_\varepsilon = 0$. The coefficients of this derivative, including M_σ , are independent of the realized mean state w . Factoring the expression gives $m' = (\psi \mathcal{Z})'(w - w_*)$ with w_* in (3.16). Since $d\chi/du < 0$, the sign of $-dm/du$ is precisely the sign of $w - w_*$.

If $\chi \leq 1/2$, then $2/\chi - 1/(1-\chi) > 0$, so $C'_P > 0$ and $\sigma'_U > 0$. The restrictions $-\varsigma s_i \geq 0$ and $w \geq 0$ make the remaining mean terms nonnegative. Thus $m' > 0$ on the positive-rate branch. The strict floor has zero local response; a boundary uses the positive-part rule. At high shares the negative term in C'_P/C_P can dominate. Substitution of the Section 3.4 illustration at $\chi = .91$ gives

$$\kappa = 12.9251592, \quad \psi = .9928038, \quad \sigma_U = .2713528, \quad w_* = .2554702.$$

With $w = .01$, the same formulas give $m = .6684894$ and $dm/d\log u = .0018286 > 0$, verifying the printed reversal. The exact threshold, rather than this illustrative magnitude, determines the comparison at other states. ■

B.6 Finite Zero-Delay-Cost Neutrality

Proof of Proposition 5. At $\gamma=0$, the finite solution in Appendix B.1 has $J+1$ equal trades and $g_j=\psi=1$. The scalar pricing equation gives $\chi=\bar{\chi}_N$, independent of u . Hence ζ is constant, $\beta \propto u$ and $\kappa \propto u^{-1}$. With fixed N, r, σ_v^2 , the variance identities give constant $\mathcal{V}_p, h, K, S, C_P$ and $\kappa\beta=\bar{\chi}_N/(N+r)$.

At a matched own signal and direction, $A_i^0 = -\zeta\kappa\beta(N-1)\rho s_i$ is constant. Matching z therefore fixes $A=A_i^0+Kz$; matching w is equivalent because S is constant. Fixed service inputs then give constant a, b_J, σ_U . The same $Q, \bar{\ell}$ yield the same nonnegative margin, including deterministic and floor branches. With Q, e also fixed, every input to Proposition 6 is unchanged. Its selected quantity and continuation indicator are consequently unchanged in the regular, tangent, capped, nonpositive-equity and infeasible-contract branches. At $Q=0$, $m=D=0$ directly. This proves the complete matched-state assertion, including $N=1$, where $K=C_P=0$.

By contrast, fixing raw client-netted flow Y_i gives $z = -\zeta\kappa Y_i + \zeta\kappa\beta(N-1)\rho s_i$; its first term generally changes with u . Fixing the own order gives $s_i = x_{i0}/\beta$, which generally changes as well. Neither is the matching convention of the proposition. Finite target completion is preserved throughout the zero-delay-cost comparison. ■

B.7 The Complete Locked-Rate Quantity Rule

Proof of Proposition 6. The zero-exposure, nonpositive-equity and infeasible-contract cases are direct provisions of the operational contract. At $Q=0$, define forced orders and invoices as zero; the sign is unused, and $C=1$ exactly when $e \geq 0$. At $Q > 0$, $e \leq 0$ or an infeasible finite demanded rate gives $D=Q, C=0$. For positive equity and finite rate, $e \geq mQ$ is supported at $D=0$ and has $C=1$, including equality. In particular $m=0, e > 0$ always continues.

It remains to classify $0 < e < mQ$. Write $\delta = mQ - e > 0$ and $t = m - a$, so $F(x) = -\delta + tx - b_J x^2$ starts strictly below zero. If $b_J > 0$ and $t \leq 0$, it decreases on $[0, Q]$ and has no feasible point. If $t > 0$, its unrestricted maximum is at $x_v = t/(2b_J)$ and has value $\Delta_F/(4b_J)$, where $\Delta_F = t^2 - 4b_J\delta$. A negative discriminant gives no feasible point. A zero discriminant gives only the tangency x_v ; when that point is in $(0, Q)$ it is feasible but nonregular, since its relief slope is zero. The prescribed choice is full wind-down.

For $t > 0$ and $\Delta_F > 0$, both roots

$$D_- = \frac{t - \sqrt{\Delta_F}}{2b_J}, \quad D_+ = \frac{t + \sqrt{\Delta_F}}{2b_J}$$

are positive, and the full feasible set on the real line is $[D_-, D_+]$. Intersect it with $[0, Q]$. If $D_- < Q$, its minimum is the regular upward crossing D_- , with $F'(D_-) = \sqrt{\Delta_F} > 0$. Rationalizing the numerator gives (4.2). If $D_- = Q$, the first feasible quantity is the full-exhaustion endpoint. If $D_- > Q$, the feasible interval misses $[0, Q]$. These cases, as well as tangencies outside the allowed interval, have no interior margin-restoring reduction and receive $D=Q$.

If $b_J = 0$, $F(x) = -\delta + tx$ is linear. For $t > 0$ its first feasible point is δ/t , regular whenever it lies in $(0, Q)$. Equality with Q gives exhaustion, and a larger value is unattainable. For $t \leq 0$, the forecast test is never met. This exhausts all deficient states. Every positive selected reduction has $C = 0$, so no remaining ordinary trade is executed.

The proof uses the entire feasible interval. For the printed negative-buffer example, $F(x) = -.05 + .9x - x^2$ has first root .0594875 in $(0, 1)$ although $F(1) = -.15$. A negative full-unwind forecast balance therefore does not rule out a regular smaller cure. Deterministic execution, nonpositive buffers and the rate floor require no change in this interval argument. The Gaussian zero-budget case has no finite margin by Appendix B.3, and is already assigned to wind-down without evaluating an infinite-rate quadratic. The rule always selects a quantity in $[0, Q]$; full wind-down terminates the position even when its realized costs leave debt. ■

B.8 Regular Quantity Responses

Proof of (4.4)–(4.7) and Corollary 2. At a selected root $0 < D < Q$ with $H = m - a - 2b_J D > 0$, the implicit-function theorem applies to $F(D) = 0$. Its total differential is

$$0 = e_\tau - (Q - D)m_\tau - mQ_\tau - Da_\tau - D^2(b_J)_\tau + HD_\tau.$$

Solving gives (4.4). The pre-reset-position-size margin derivative is (B.6); it includes moving size, budget, mean and risk. The rate has no derivative with respect to the subsequently chosen quantity: the current contract locks the calibration at pre-reset Q and entering λ_0 .

For the matched within-economy public-state path, conditional variance is unchanged and $a_z = m_z = \psi K$. Thus the numerator becomes $(Q - D)\psi K + D\psi K = Q\psi K$, proving (4.5). Equivalently, $m - a = b_J Q + B$ is constant on this path, while the deficiency has derivative $Q\psi K$. Both forms keep the cost of sold units as well as the requirement on retained units.

Let $T(D; m) = Q - E^f(D)/m$ be the literal quantity map at fixed a, b_J, e, Q . At its root,

$$T_D = \frac{a + 2b_J D}{m} = \Xi, \quad T_m = \frac{E^f(D)}{m^2} = \frac{Q - D}{m} = Q_m, \quad H = m(1 - \Xi).$$

Regular relief requires $\Xi < 1$. A strict local contraction of the iteration requires $|\Xi| < 1$; $|\Xi| > 1$ is locally repelling, and equality is outside the strict contraction test. For example, $Q = m = 1, a = -2, b_J = 0, e = .4$ gives the regular root $D = .2$, $H = 3$ and $\Xi = -2$. Thus regularity alone does not validate the iteration. Rewriting the public response gives

$$D_z = \frac{\psi K(Q_m + D/m)}{1 - \Xi}.$$

For $K > 0$ and $0 < \Xi < 1$, the denominator is below one and $D/m > 0$, so this response strictly exceeds the partial rate-only term $\psi K Q_m$. No such amplification claim is needed for favorable-cost negative Ξ .

At fixed Q, e, b_J , substitute $m = a + b_J Q + B$ to obtain $HD_\tau = Qa_\tau + (Q - D)B_\tau$. On the composition region proved in Appendix B.4, $a_r \geq 0$ and $B_r = M_\sigma(\sigma_U)_r > 0$. Since $Q - D > 0$, this proves Corollary 2. On the separate low-share hedging region, Appendix B.5 gives $a' \geq 0$ and $B' > 0$ as χ rises, so thinner hedging also increases a regular reduction. When the two components have opposite signs, their weights Q and $Q - D$ differ; the sign of $m_\tau = a_\tau + B_\tau$ alone cannot determine D_τ .

On a strict continuation branch the selected quantity is identically zero. On a strict full-wind-down branch it equals Q , so its local derivative is zero when Q is fixed, and equals Q_τ if pre-reset position size varies. A tangency has $H = 0$ and no regular implicit derivative. At the loss of a regular root the contract can switch to full wind-down discontinuously. All of the displayed interior derivatives retain their strict branch restrictions. ■

B.9 Actual Collateral and Realized Shortfalls

Proof of (4.9)–(4.14). The pre-reset position starts with net cash $B_1 = e - p_0 q_1$. For $Q > 0$ a forced slice has signed client order $f_j = -\varsigma r_j D$, price $p_j = p_0 - \varsigma U_j(D)$, and cash increment

$$-p_j f_j = \varsigma p_0 r_j D - r_j D U_j(D).$$

The simultaneous holding change is f_j . Marking at p_0 cancels its benchmark cash leg and leaves equity change $-r_j D U_j(D)$. After slice j ,

$$B_j^I = e - p_0 q_1 - \sum_{h \leq j} p_h f_h, \quad q_j^I = q_1 + \sum_{h \leq j} f_h, \quad B_j^I + p_0 q_j^I = e - L_j.$$

This calculation applies to a sale from a long and a purchase covering a short; favorable negative concessions increase equity. At $D = 0$, invoices and forced cash flows are zero directly, and at $Q = 0$ no direction is evaluated.

The collateral-restriction rule sets aside available pre-reset account resources up to the intended requirement. Therefore $M_j = \min\{M^*, (e - L_j)^+\}$. Losses first consume any unrestricted reserve and then restricted collateral; later gains restore restricted collateral up to M^* . The execution commitment completes the selected quantity even if a temporary deficit requires a broker advance. At full wind-down $M^* = 0$ directly, including when the demanded rate is infeasible. At a regular root, $M^* = e - Dc(D)$; if $Dc(D) \geq 0$, it does not exceed entering equity. If instead forecast execution gains are needed to fund that amount, initial restricted collateral is only $\min\{M^*, e^+\}$. These facts establish (4.9) under the stated execution and collateral-restriction conventions.

On continuation, $L_j = 0$ and the supported pre-reset principal stays at mQ . The ordinary facility records $H_j^P = \sum_{h \leq j} (p_0 + \kappa x_h^P) x_h^P$, so its benchmark-marked equity is $-H_j^P + p_0 q_j^P = -\kappa \sum_{h \leq j} (x_h^P)^2$. Adding the two books gives

(2.16). In particular, pre-reset collateral need not finance the ordinary invoices: the corresponding program debt remains in the combined accounts.

Posting or releasing principal transfers an asset between the client's restricted and unrestricted subaccounts. Its use against an execution invoice is already included in L_j . The opportunity cost over the interval before slice j is $k\Delta_j M_{j-1}$; summing proves (4.10), with no second charge for the principal. This opportunity value is outside the legal claims.

At a regular root, the full invoice is $D(a+b_J D+Z)$, and forecast equality gives $E_D = M^* - DZ$. Consequently $C_D = (DZ)^+$ and $S_D = [DZ - M^*]^+$. For $D\sigma_U > 0$, the centered normal residual has equal probabilities of either sign, and

$$\Pr(C_D > 0 | \mathcal{H}_i) = \frac{1}{2}, \quad \mathbb{E}[C_D | \mathcal{H}_i] = \frac{D\sigma_U}{\sqrt{2\pi}}, \quad \Pr(S_D > 0 | \mathcal{H}_i) = \overline{\Phi}\left(\frac{M^*}{D\sigma_U}\right).$$

These are end-window events. At $D=0$ or $\sigma_U=0$, the direct balances apply; zero deviation at a deterministic regular cure is not a positive breach.

After interruption, no remaining client execution program executes. The unsold holding is $q_1 + f$, so maturity changes legal resources by $(q_1 + f)(v - p_0)$, proving (4.13). The held remainder is settled without another sale. At full wind-down it is zero, leaving signed resources $e - QU(Q)$. Client limited liability and the completed pre-reset-book settlement give payout $[e - QU(Q)]^+$ and credit shortfall $[QU(Q) - e]^+$. Applying (B.5) gives

$$\mathbb{E}[S_{\text{full}} | \mathcal{H}_i] = Q\sigma_U \mathcal{L}\left(\frac{e/Q - a - b_J Q}{\sigma_U}\right).$$

If costs are deterministic, use $[Q(a+b_J Q) - e]^+$; at zero exposure use $(-e)^+$ for cash default. A partial reduction retains value risk at maturity. Accordingly its terminal credit shortfall, its end-window marked deficit and its collateral deficiency are different objects. The demanded-rate calibration uses mQ , whereas actual shortage-state credit losses use the available e . ■

B.10 Subsequent-Trader Execution Costs

Proof of Proposition 7. With no inventory reallocation during the focal window, each forced unit adds one unit to the relevant liquidity-provider pool. Thus $X_1 - X_0 = D$, and the maintained map gives $\lambda_1 - \lambda_0 = \omega_X D$. At the specified later quantity and fixed other service conditions, the per-unit invoice changes by $d_L \omega_X D$. Multiplying by y_L proves the level identity in (5.1). Within the positive-margin interior branch, Appendix B.8 gives $D_z = \psi K Q / H$; applying the chain rule gives the derivative and its strictness conditions.

The matched entering-endowment construction in Appendix B.4 preserves the focal pre-reset Q, e and own signal across its order-flow states. Each such state has its own settled $p_0 = \kappa y_0$. The subsequent trader's benchmark, direction, quantity, schedule, other states and shocks are held fixed separately. Between

two admissible reset states, its invoice difference is therefore $y_L d_L \omega_X (D_2 - D_1)$. This comparison does not require a quantity response by the subsequent trader.

Strict continuation gives zero added burden and invoice. Strict full wind-down gives invoice level $y_L d_L \omega_X Q$, with zero local public-state derivative at fixed Q . At $Q=0$, the contribution is zero directly. At tangencies or branch changes the regular derivative is unavailable. If $K=0$, incremental public learning contributes zero local derivative, although known pressure can support a nonzero sale and invoice level. The level vanishes if y_L , d_L or ω_X is zero. For the $\omega_X=0$ comparison that preserves current D , entering λ_0 must stay fixed: only then are all inputs to the current rate and quantity rule unchanged. Since $\lambda_0 = \lambda_{\min} + \omega_X X_0$, fixing instead $\lambda_{\min}$ and X_0 can change the current contract. Any comparison of the primitive representation must also preserve its nonnegative cost-state domain.

For placement before the later trade, let

$$\mathcal{T}(X) = X - \min\{(1 - \varphi_X)X, \bar{F}\} = \max\{\varphi_X X, X - \bar{F}\}.$$

The attributable burden is $\mathcal{T}(X_0 + D) - \mathcal{T}(X_0)$. If both states are nonbinding, it is $\varphi_X D$; if both are capacity-bound, it is D . At a crossing use the displayed difference. For $\varphi_X < 1$ and $X_0 \leq \bar{F}/(1 - \varphi_X) \leq X_0 + D$, it equals $D + (1 - \varphi_X)X_0 - \bar{F}$, interpolating between these two cases. At $\varphi_X = 1$ there is no placement and the increment is D . Multiplying the burden increment by $y_L d_L \omega_X$ supplies the after-placement invoice. This also establishes the placement qualification in Sections 6 and Appendix A.

Finally, any regular parameter path holding subsequent-trader inputs and ω_X fixed has $(\mathcal{I}_L^X)_\tau = y_L d_L \omega_X D_\tau$. The complete derivative proved above carries all moving mean, risk, size and budget terms. In the restricted composition region $D_r > 0$, so the invoice rises strictly when $y_L d_L \omega_X > 0$. ■

B.11 Payments, Recovery Priorities and Collateral Resources

Proof of (5.2)–(5.5). Write $\delta_v = v - p_0$. For a forced signed client order f_j at p_j , the client's incremental terminal payoff is $f_j(v - p_j)$ and the provider's is its negative. Using $f_j = -\zeta r_j D$ when $Q > 0$,

$$f_j(v - p_j) = f_j \delta_v - r_j D U_j(D), \quad -f_j(v - p_j) = r_j D U_j(D) - f_j \delta_v.$$

Summing gives the provider profit $\Pi_F = L - f \delta_v$. For an ordinary signed trade x_j^P at $p_0 + \kappa x_j^P$, the two trading payoffs are

$$x_j^P \delta_v - \kappa (x_j^P)^2, \quad \kappa (x_j^P)^2 - x_j^P \delta_v.$$

Thus $\Pi_P = C_{\text{exec}} - q^P \delta_v$. The client starts with $e + q_1 \delta_v$ at terminal valuation, before these new trading legs. Adding its executed legs gives

$$\mathcal{N} = e - L - C_{\text{exec}} + (q_1 + f + q^P) \delta_v, \quad \mathcal{N} + \Pi_F + \Pi_P = e + q_1 \delta_v.$$

This proves (5.2) and (5.3). In every cancellation branch $x_j^P = 0$; at zero forced quantity $f = L = 0$ directly. Signed orders ensure the same identities for short covering, favorable concessions and an remaining client execution program crossing the pre-reset sign.

The financing ledger retains the full ordinary draw:

$$H_J^P = p_0 q^P + C_{\text{exec}}, \quad -H_J^P + v q^P = q^P \delta_v - C_{\text{exec}} = N_P.$$

A positive H_J^P is debt, and a negative value is program cash. The acquired position and its principal funding offset at the benchmark except for the actual concession. Omitting execution-cost debt or treating sale proceeds as a free transfer to the pre-reset position would break this identity. On continuation $N_I = e + q_1 \delta_v$ and $N = N_I + N_P$; on interruption the latter program is absent. Providers and ordinary counterparties retain the opposite risky holdings, so their fundamental exposures cancel only after all three trading accounts are added.

To pass from contractual balances to actual payouts, gross the existing terminal assets and remaining liabilities after earlier repayments. Positive cash and positive settlement values enter $\mathcal{A}$; debt and settlement obligations of signed risky liabilities enter the nonnegative claims H_h . Negative centered prices or signed holdings are handled by the corresponding positive and negative settlement values, rather than being dropped. Their difference is the already computed N . Let R_h be the actual recoveries, with $0 \leq R_h \leq H_h$, under the book priorities. Committed execution has already occurred. Without an outside injection or bankruptcy-resource destruction, complete distribution gives

$$c + \sum_h R_h = \mathcal{A}, \quad S = \sum_h (H_h - R_h).$$

Subtracting claims from the distribution identity yields $c - S = \mathcal{A} - \sum_h H_h = N$. Substitution into the trading identity proves (5.4) for the actual priority-respecting vector. The identity does not require an invented allocation rule for those individual recoveries.

On interruption there is only the completed pre-reset-book settlement. Claims are met from its resources before any residual client payout. Accordingly $c = (N_I)^+$ and $S = (-N_I)^+$, using $x^+ - (-x)^+ = x$. Full wind-down has loss $[L - e]^+$; at zero exposure and negative cash it is $(-e)^+$. For the Section 5.3 illustration, $L = .5(2.45) = 1.225$ and $N_I = 1.175 - 1.225 = -.05$. Therefore $c = 0$, $S = .05$ and $\Pi_F = 1.225$ at $v = p_0$; the three signed incidences sum to 1.175.

On continuation, program creditors first claim their positions and proceeds, with further claims subordinate until pre-reset obligations have been discharged. Program claims cannot consume pre-reset collateral during the window. These priorities govern the recoveries in the same distribution identity. Their aggregate conservation equation alone does not impose $c = (N)^+$ across the separately financed books. In particular, negative program resources and unpaid program claims remain represented. The end-window covenant

deficiency C_D is neither an additional asset distribution nor a second unpaid claim to subtract.

Collateral principal is already among client assets. Posting and release are internal transfers, and consumption against invoices has already entered L . The sole modeled opportunity charge is $F_{\text{coll}} = k \sum_j \Delta_j M_{j-1}$, outside legal claims. Deducting it once from the payout-and-loss identity gives

$$c - S - F_{\text{coll}} + \Pi_F + \Pi_P = e + q_1 \delta_v - F_{\text{coll}},$$

which is (5.5). No reduced-form mandate objective term or expected-loss limit is a second cash charge here.

For completeness, a specified subsequent trader's signed order $\varsigma_L y_L$ at $p_L + \varsigma_L U_L$ contributes $\varsigma_L y_L (v - p_L) - y_L U_L$ to its payoff and the opposite amount to its counterparty. Increasing U_L by $d_L \omega_X D$ changes these payoffs by $-I_L^X$ and $+I_L^X$. This establishes both recipients of the invoice incidence. The accounting preserves transferred risk and actual credit losses; it does not identify a provider's resource costs or assign a social value to canceled mandates or collateral. ■

References

- Abruzzo, N., and Y.-H. Park. 2014. An Empirical Analysis of Futures Margin Changes: Determinants and Policy Implications. Finance and Economics Discussion Series 2014-86, Board of Governors of the Federal Reserve System.
- Almgren, R., and N. Chriss. 2001. Optimal Execution of Portfolio Transactions. *Journal of Risk* 3:5–39.
- Bertsimas, D., and A. W. Lo. 1998. Optimal Control of Execution Costs. *Journal of Financial Markets* 1:1–50.
- Biais, B., F. Heider, and M. Hoerova. 2016. Risk-Sharing or Risk-Taking? Counterparty Risk, Incentives, and Margins. *Journal of Finance* 71:1669–98.
- Brunnermeier, M. K., and L. H. Pedersen. 2005. Predatory Trading. *Journal of Finance* 60:1825–63.
- . 2009. Market Liquidity and Funding Liquidity. *Review of Financial Studies* 22:2201–38.
- Gârleanu, N., and L. H. Pedersen. 2011. Margin-Based Asset Pricing and Deviations from the Law of One Price. *Review of Financial Studies* 24:1980–2022.
- Glasserman, P., and Q. Wu. 2018. Persistence and Procyclicality in Margin Requirements. *Management Science* 64:5705–24.
- Glebkin, S., N. Gondhi, and J. C.-F. Kuong. 2021. Funding Constraints and Informational Efficiency. *Review of Financial Studies* 34:4269–322.
- 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* 14:71–100.
- Kyle, A. S. 1985. Continuous Auctions and Insider Trading. *Econometrica* 53:1315–35.
- Murphy, D., M. Vasios, and N. Vause. 2014. An Investigation into the Procyclicality of Risk-Based Initial Margin Models. Bank of England Financial Stability Paper No. 29.

Nezafat, M., and M. Schroder. 2022. Private Information, Securities Lending, and Asset Prices. *Review of Financial Studies* 35:1009–63.