Multi-Asset Market Making

Ioanid Roşu, Andriy Shkilko, Elvira Sojli, and Wing Wah Tham*

July 8, 2026

Abstract

We model a risk-averse dealer who manages inventory by setting quotes in multiple correlated assets. In closed form, each midquote equals the fundamental value minus a linear combination of the dealer's inventories in all assets, with coefficients that form a pricing matrix; its off-diagonal entries capture cross-hedging: a long dealer quotes to sell correlated assets. Cross-hedging increases with fundamental correlation, while own-quote shading weakens, a substitution effect. It declines with risk aversion, as she mean-reverts each position too quickly for portfolio rebalancing to matter, and is diluted across many assets. The resulting transitory price pressure lowers return correlations at high frequencies, a pattern known as the Epps effect. Ignoring cross-hedging would bias risk-aversion estimates downward: with N perfectly correlated assets the estimate recovers only a fraction $1/N$ of the true coefficient. In message-level Toronto Stock Exchange data, market makers' quotes display cross-hedging that is directionally sharp and increases with correlation, yet quantitatively modest, consistent with relatively high dealer risk aversion in the model.

KEYWORDS: Market making, inventory management, pricing matrix, cross-hedging, Epps effect.

*Roşu (corresponding author), rosu@hec.fr, HEC Paris; Shkilko, ashkilko@uga.edu, University of Georgia; Sojli, e.sojli@unsw.edu.au and Tham, w.tham@unsw.edu.au, both at UNSW Sydney. We thank Bruno Biais, Dion Bongaerts, Davide Cianciaruso, Hans Degryse, Thierry Foucault, Denis Gromb, Terrence Hendershott, Stefano Lovo, Sophie Moinas, and Daniel Schmidt for their suggestions. We are also grateful to finance seminar participants at Norwegian School of Economics, HEC Paris, Peking University, and Bucharest University of Economic Studies, as well as conference participants at the 10th Erasmus Liquidity Conference, the 10th Women in Market Microstructure Conference, and the 10th Annual Conference of Romanian Academic Economists from Abroad (ERMAS), for valuable comments.

> *"A market maker's success depends on it posting the best available prices and accurately responding to relevant market data in similar and correlated instruments."*
>
> — Virtu Financial, Inc., S-1 Registration Statement, U.S. Securities and Exchange Commission (2014)

Market makers are the liquidity supply side of financial markets, posting bid and ask quotes at which other participants trade. On the liquidity demand side, investors buy and sell assets to form and rebalance portfolios, motivated by differences in views, information, or endowment. While that demand side has been extensively analyzed through the lens of Modern Portfolio Theory, the supply side has received comparatively little attention, especially when the dealer is active in multiple assets simultaneously. Two questions arise naturally. How does a market maker manage inventory risk when positions in different assets are correlated? And what effect does this multi-asset risk management have on prices?

In Modern Portfolio Theory, investors allocate capital across assets to trade off expected return against risk. Much less is known about the corresponding portfolio problem faced by the intermediaries who make those assets tradeable. A market maker typically provides liquidity in many securities at once, and shocks to her inventory in one security change the risk of her entire portfolio. Virtu Financial, one of the largest electronic market makers in the world, puts it plainly in its public filings: a market maker's success depends on accurately responding to market data in *similar and correlated instruments* (Virtu Financial, Inc., 2014). The firm simultaneously quotes prices for over 25,000 securities across 235 venues in 37 countries, managing inventory across all of them as a single correlated portfolio. Yet the canonical inventory-based market-making literature largely analyzes quote setting one traded asset at a time, with inventory control operating through quote shading in that single asset. Even when a dealer's background portfolio risk enters the objective, the theory rarely models dynamically chosen quotes across multiple traded assets that allow active cross-hedging.[1] The

[1]The distinction is between a single traded asset with exogenous background portfolio risk, as in Ho and Stoll (1981) and related work, and a setting with multiple traded assets in which the dealer jointly sets quotes across all of them and can exploit correlations to hedge inventory actively. The

single-asset framework developed by Ho and Stoll (1981) and extended by Hendershott and Menkveld (2014) has been very influential, but it is silent on how a dealer allocates inventory adjustment across correlated assets inside her own portfolio. This cross-asset dimension is precisely the margin that a portfolio perspective makes natural.

We propose a dynamic model of market making in N risky assets. Each trading round, a monopolist dealer sets ask and bid quotes in all assets at which investors trade. The market consists of one risk-free asset and several risky assets, whose fundamental value evolves according to a random walk with multivariate normal increments. Trading takes place at discrete dates. After the dealer sets ask and bid quotes, the investors submit an aggregate market order that is the sum of an inelastic component, an elastic component linear in the fundamental value minus the quote, and a multivariate normal noise. The dealer maximizes her expected profit subject to a penalty that is quadratic in her inventory in the risky assets.

We solve the model in closed form, with a particularly transparent characterization as the discount coefficient approaches one. Each midquote equals the fundamental value minus a linear combination of the dealer's inventories across all assets, with the coefficients forming an N-dimensional pricing matrix. Its diagonal entries capture the familiar own-stock price pressure that pushes a long dealer to lower her own quote, while its off-diagonal entries capture cross-stock price pressure, the channel through which inventory in one asset shifts the dealer's quotes in another. The off-diagonals are positive whenever fundamental returns are positively correlated: a dealer who is long in one stock lowers her midquote in correlated stocks, encouraging investors to buy those stocks from her and thereby building offsetting positions. This cross-stock price pressure, i.e., inventory cross-hedging, is the central object of the paper.

The model delivers a sharp set of predictions about how cross-hedging varies with the underlying correlation structure. The off-diagonal entries are increasing in the fundamental correlation between two stocks: more correlated assets are better substitutes

important exception is Ho and Stoll (1983), in which competing dealers hold multi-stock portfolios and reservation quotes respond to the covariance of each stock with the dealer's inventory portfolio. Their pricing is myopic mean-variance, however: quote adjustments are proportional to risk aversion times covariances, so the ratio of cross-stock to own-stock adjustment is pinned down by the correlation structure and is invariant to risk aversion. In our dynamic model this ratio is endogenous, lies below the static benchmark, and declines with risk aversion.

for inventory adjustment, so the dealer relies on them more. By the same substitution effect, the diagonal entries are decreasing in correlation, since the dealer relies less on the focal asset itself when correlated assets absorb part of the adjustment. In the limit of perfect correlation, the two assets become perfect substitutes and only the sum of inventories matters to the dealer.

A further prediction concerns the role of risk aversion. As the dealer becomes more risk-averse, the pricing matrix approaches a diagonal form. The intuition is that a very risk-averse dealer mean-reverts each inventory to zero so quickly that cross-hedging has no time to operate: in the limit, the liquidity shock in each asset is fully reversed within a single period, pricing discounts are set independently across assets, and the off-diagonal channels close down. Another comparative static concerns the number of assets: with many correlated assets, the hedge is spread across all of them, so each pairwise quote response is inversely proportional to the number of assets.

The final prediction concerns measured return correlations. When the dealer is risk-averse and the pricing matrix is close to diagonal, quotes in different assets respond independently to their respective inventories. At high frequencies, this independent inventory-driven price pressure generates a mean-reverting component that suppresses comovement; at lower frequencies, the random-walk component of fundamental value dominates and return correlations converge to their fundamental level. Our model thus provides an inventory-based theoretical account of the Epps (1979) effect.

We take the model's central predictions to the complete message-level record of the Toronto Stock Exchange STAMP feed, covering 45 trading days from January 3 to March 6, 2006. Every order, cancellation, and trade carries a user identifier, so we reconstruct each market maker's inventory trade by trade, in every stock, and observe the resting quotes she posts alongside it. This combination, the full order lifecycle with market maker identifiers, is rare in available data and is what makes the off-diagonal of the pricing matrix observable. We identify 549 market makers and sort them into three groups: designated market makers with specialist obligations, endogenous market makers whose two-sided liquidity provision is voluntary, and market makers that do both.

Our empirical design follows the theory's main object of interest, quoted prices. For

each market maker, stock, and minute we measure the quoted liquidity imbalance of her resting quotes, ask depth minus bid depth in dollars. The imbalance is positive when the market maker supplies relatively more liquidity on the ask side than on the bid side, shifting her quoted supply schedule toward selling. We then ask whether her inventory in one stock drives the positive imbalance in the rest of her portfolio, with inventory preceding the quote by five minutes. When it does, for a sustained run, we call it a cross-hedging episode, and within each episode an exact covariance decomposition names the stocks that carry the sell-side liquidity imbalance.

We identify 879,942 episodes, run by 537 of the 549 market makers, and the stocks chosen as hedges are more correlated with the focal stock than random stocks from the same portfolio, so the mechanism is widespread and organized by correlation; yet episodes cover on the order of a tenth of the trading day for the average market maker-stock pair, and within the median episode the rest-of-portfolio imbalance offsets 13 percent of the focal inventory covariance, so the mechanism is modest, a magnitude the model attributes to risk aversion and to the hedge being spread across many stocks. A market maker accumulates a long position in one stock and turns her resting quotes to a sell-side liquidity imbalance in a few correlated stocks, e.g., banks against banks, energy against energy. The quoted liquidity imbalance is concentrated, with two or three stocks carrying most of it: long the focal stock while quoting to sell its correlated relatives is the opposite of what a common factor produces.

The positive quoted liquidity imbalance increases with cross-pair correlation. The share of pairs quoted to sell rises with the return correlation of the two stocks, from 0.58 to 0.73 for the most correlated pairs. This finding is in line with the cross-hedging prediction that the off-diagonal grows with correlation. The model's remaining comparative statics, the fall of the diagonal as correlation rises and the attenuation of cross-hedging with risk aversion, are causal claims: testing them requires exogenous variation in correlation and in risk aversion, which a single cross-section does not provide, and we leave them to future research.

The portfolio dimension of inventory management matters for measurement and interpretation. In the single-asset literature, including Ho and Stoll (1981) and Hendershott and Menkveld (2014), observing an inventory half-life in one asset and mapping it

to risk aversion is equivalent to assuming that the dealer does not hedge through correlated positions. When the dealer is managing a correlated portfolio, any one position is unwound more slowly because correlated assets absorb part of the adjustment burden, which weakens the incentive to eliminate the focal position directly. With N perfectly correlated assets, studying one asset in isolation recovers only a fraction $1/N$ of the true risk-aversion coefficient, an underestimation that equals 50 percent for two assets and approaches 100 percent as the number of assets grows. For imperfect correlation, the underestimation increases with both the correlation among the assets and their number. Market makers are therefore more risk-averse than single-asset estimates suggest. More generally, single-asset measures of price impact, market quality, or inventory adjustment speed can be systematically biased when dealers manage correlated portfolios, and the direction and magnitude of the bias vary with the correlation structure in the way our model specifies.

Our paper contributes to several strands of the literature. The first is the theory of inventory-based market making. Stoll (1978) and Amihud and Mendelson (1980) establish the foundational link between inventory costs and the bid-ask spread in single-asset settings. Ho and Stoll (1981) derive the optimal bid-ask spread for a dealer facing return uncertainty in a single asset. Ho and Stoll (1983) extend the framework to competing dealers who make markets in multiple stocks: reservation quotes respond to the covariance of each stock with the dealer's inventory portfolio, and order flow is won by the dealers whose inventories make them most competitive. Hansch, Naik, and Viswanathan (1998) test this framework on London Stock Exchange dealer data and find that relative inventories determine which dealers obtain trades, while Naik and Yadav (2003) find that inventory management is predominantly stock-by-stock, with limited support for the portfolio dimension. Hendershott and Menkveld (2014) bring the single-asset framework to the data, identifying price pressure as the mechanism by which dealers restore inventory balance. Our model extends this line of work by solving the dealer's dynamic portfolio problem in closed form. The single-asset pricing scalar generalizes to an N-dimensional pricing matrix whose off-diagonal elements are theoretically disciplined and empirically detectable, and its dynamics reconcile the theory of Ho and Stoll (1983) with the evidence of Naik and Yadav (2003): optimal cross-hedging is attenuated

relative to the static benchmark and declines with risk aversion, so positions can look managed stock-by-stock even though the portfolio dimension survives in the quotes and remains first-order for the measurement of risk aversion.

In a related but distinct setting, Daures-Lescourret and Moinas (2023) model imperfect competition among two market makers posting prices for the same asset on two exchanges, focusing on how fragmentation affects liquidity; our focus is instead on how a single dealer optimally allocates inventory adjustment across correlated assets, and what that implies for price formation, market quality, and the measurement of risk aversion.

The second strand is the broader literature on market maker behavior. Menkveld (2013) documents the inventory dynamics of a large high-frequency market maker in a single asset, and Brogaard, Hendershott, and Riordan (2014) study high-frequency traders and price discovery across stocks while treating each stock independently. Copeland and Galai (1983) examine quote-setting under adverse selection, also in a single-asset framework. More broadly, the literature on dealer networks, including Hugonnier, Lester, and Weill (2020) and Di Maggio, Kermani, and Song (2017), examines how intermediaries interact with each other across markets, but does not address how an individual dealer manages inventory across correlated positions within her own portfolio. Our empirical contribution is to bring the cross-asset dimension into this literature directly. Rather than studying each stock in isolation, we reconstruct the full multi-asset inventory position of each identified market maker and measure how shocks in one stock propagate into quoting behavior across the dealer's entire portfolio. To our knowledge, we are the first to jointly derive a multi-asset closed-form pricing matrix for inventory management and test its cross-asset predictions using message-level order data with market maker identifiers across a broad equity universe with structurally distinct market maker populations.

Finally, our model provides a complementary theoretical mechanism for the Epps effect. Epps (1979) documents that stock return correlations decline systematically as the measurement frequency increases, a finding replicated across markets and time periods. The leading explanations attribute the attenuation to nonsynchronous trading (Scholes and Williams, 1977), bid-ask bounce (Roll, 1984), and differential information flow (de Jong and Nijman, 1997). Our model adds an inventory-based channel: a risk-

averse dealer who mean-reverts inventory aggressively injects a transitory component into prices that attenuates measured co-movement at high frequencies, but that vanishes at lower frequencies where the random-walk component of fundamental value dominates. This mechanism is complementary to the canonical explanations, and its distinctive prediction is that the degree of attenuation should vary with the aggregate risk aversion of the market makers active in a stock pair.

1 Model

1.1 Environment

Our model is close in spirit to the dealer models of Ho and Stoll (1981) and Hendershott and Menkveld (2014), except that we consider trading in multiple assets, as in Ho and Stoll (1983), but with quotes chosen dynamically rather than from a myopic mean-variance criterion. Specifically, the market consists of one risk-free asset and N risky assets. Trading in the risky assets takes place in a market exchange, at discrete dates $t = 0, 1, 2, \ldots$, according to the mechanism described below. There are two types of market participants: (a) one monopolistic market maker called the dealer ("she") who sets ask and bid quotes at which others trade, and (b) traders, who submit market orders.

Assets. The risk-free asset is used as a numeraire and has a return of zero. Each risky asset $n \in \{1, 2, \ldots, N\}$ has a fundamental value $v_{n,t}$ that follows a continuous multi-variate random walk process for which the increments have instant covariance matrix per unit of time equal to Σ_v. We define the fundamental value v_t of the N assets as the column vector with n'th entry equal to $v_{n,t}$. In Appendix B, we show that $v_{n,t}$ can be interpreted as the cash value that shareholders receive at liquidation, an event which may occur in each period with a constant probability.[2]

Dealer's Quotes and Objective. Suppose the dealer observes the fundamental

[2]Suppose there exists $\pi \in (0, 1)$ such that asset n liquidates in each period with probability π, in which case the shareholders receive $v_{n,t}$ per share. Then we show in Appendix B that the expected profit of a trader with quantities bought and sold at t equal to $-Q_t^b$ and $-Q_t^s$, respectively, has the form described in equation (2) with $\beta = 1 - \pi$.

value v_t before trading at t and sets the ask quote vector a_t and the bid quote vector b_t. Let x_t be the dealer's inventory in the risky assets before trading at t, which is also a column vector with N entries. If Q_t^b is the aggregate buy market order vector at t and Q_t^s is the aggregate sell market order vector at t, the dealer's inventory evolves according to

$$x_{t+1} \;=\; x_t - Q_t^b + Q_t^s. \tag{1}$$

For a given quoting strategy, the dealer's expected utility at $\tau = 1, 2, \ldots$ is:

$$\mathsf{E}_\tau \sum_{t=\tau}^{\infty} \beta^{t-\tau} \Big((v_t - b_t)' Q_t^s \;+\; (a_t - v_t)' Q_t^b \;-\; \gamma\, x_t' \Sigma_v x_t \Big), \tag{2}$$

where $\beta \in (0,1)$ and $\gamma > 0$.[3] The dealer maximizes expected profit, but at each t faces a utility loss that is quadratic in the inventory. In Appendix B, we show that the expected utility of a risk-averse dealer is approximately of the form (2) when her risk aversion coefficient γ is sufficiently small.[4]

Trading. Upon observing the quotes, traders submit at t the following aggregate market order vectors:

$$\begin{aligned} Q_t^b &= \frac{k}{2}(v_t - a_t) + \ell + \varepsilon_t^b, \quad \text{with} \quad \varepsilon_t^b \overset{IID}{\sim} \mathcal{N}\left(0, \Sigma_L/2\right), \\ Q_t^s &= \frac{k}{2}(b_t - v_t) + \ell + \varepsilon_t^s, \quad \text{with} \quad \varepsilon_t^s \overset{IID}{\sim} \mathcal{N}\left(0, \Sigma_L/2\right), \end{aligned} \tag{3}$$

where k and Σ_L are symmetric positive definite matrices, and ℓ is a column vector.[5] The parameter k is the investor elasticity matrix and ℓ is the inelasticity parameter vector.

Dealer's Information. The main results do not depend on any particular information structure, but to obtain explicit formulas we define the dealer's information set $\mathcal{I}_t$ before trading at t as the history of net order flow: $\mathcal{I}_t = \left\{ Q_\tau^b - Q_\tau^s \mid \tau < t \right\}$.[6]

[3] The equilibrium results are the same in the limit when $\beta = 1$ if the dealer's inventory penalty $\gamma x_t' \Sigma_v x_t$ in equation (2) is replaced by $\gamma x_{t+1}' \Sigma_v x_{t+1}$; see Proposition A.1 in the Appendix.

[4] This is done for one asset, but in principle the proof holds also for multiple assets.

[5] By extending the microfoundations from Roşu, Sojli, and Tham (2021), k can be shown to be proportional to $g^{-1}\Sigma_u^{-1}$, where g is the informed trader's risk aversion, $v + u$ is the asset's liquidation value, and v is the informed trader's signal. A similar result obtains from CARA investors with expected values equal to v, in which case k^{-1} is the conditional variance of $v + u$.

[6] Alternatively, the information set could contain both the buy and sell order flow separately, i.e., $\mathcal{I}_t = \left\{ Q_\tau^b, Q_\tau^s | \tau \in \{0, 1, \ldots, t-1\} \right\}$, but the formulas would be more complicated.

2 Equilibrium

2.1 Optimal Quotes

In this section, we solve for the dealer's optimal quotes and show that the results are obtained in closed form. We fix the dealer's risk aversion coefficient γ, the fundamental covariance Σ_v, the elasticity matrix k, the inelasticity parameter ℓ, and the noise covariance Σ_L. The matrices k and Σ_L are assumed symmetric and positive definite. Let I_N be the identity matrix in N dimensions.

Proposition 1 describes the optimal quotes set by the dealer. This result is obtained by applying standard methods in linear-quadratic dynamic programming, see Appendix A.[7] We fix γ, Σ_v, k, ℓ, and Σ_L. The matrices k and Σ_L are symmetric and positive definite. Let I_N denote the identity matrix.

Proposition 1. *If v_t is the fundamental value at t and x_t is the dealer's inventory before trading at t, then the dealer's optimal quotes at $t = 0, 1, \ldots$ are:*

$$a_t \;=\; v_t - \lambda x_t + h, \qquad b_t \;=\; v_t - \lambda x_t - h, \tag{4}$$

where h is an $N \times 1$ vector and λ is an $N \times N$ matrix satisfying:

$$h \;=\; k^{-1}\ell, \qquad \lambda \;=\; k^{-1}\Lambda, \tag{5}$$

where Λ is the unique symmetric positive definite matrix solving the quadratic equation:

$$\Lambda^2 + \Lambda\Big(S + \frac{1-\beta}{\beta} I_N\Big) \;=\; S, \quad \textit{with} \quad S \;=\; \gamma k \Sigma_v. \tag{6}$$

When β converges to 1, the parameter λ converges to:

$$\lambda \;=\; \frac{k^{-1}}{2}\Big((S^2 + 4S)^{1/2} - S\Big). \tag{7}$$

We consider only the limit $\beta = 1$ henceforth, in order to simplify formulas.

To get intuition for Proposition 1, suppose there is only one asset ($N = 1$) and

[7] For a general treatment of such problems, see Ljungqvist and Sargent (2000).

the dealer is risk-neutral ($\gamma = 0$). Equation (7) shows that $\lambda = 0$, hence the dealer's inventory does not affect her strategy. Equation (4) implies that the dealer sets her quotes at equal distance around the value v_t: the ask quote is $a_t = v_t + h$ and the bid quote is $b_t = v_t - h$, where h is the constant half-spread. The equilibrium value $h = \ell/k$ in equation (5) reflects two opposite concerns for the dealer: too large a spread would decrease the profit that comes from the price-elastic investors (with elasticity k), while too small a spread would lower the profit that comes from the inelastic traders (who submit a fixed quantity ℓ at both quotes).

If the dealer is risk-averse, equation (7) shows that the dealer's inventory affects the quotes across assets in a linear way via the coefficient λ, i.e., the midquote price $p_t = (a_t + b_t)/2$ satisfies:

$$p_t \;=\; v_t - \lambda x_t. \tag{8}$$

To understand why, suppose before trading at t, the dealer has accommodated a positive net order flow and thus has a negative inventory ($x_t < 0$). To bring back her inventory to zero, she must encourage new sellers and discourage new buyers, which can be achieved by raising both quotes at t. Thus, the dealer raises the midquote above the value v_t, to a level equal to $p_t = v_t + \lambda(-x_t)$. For one asset, this is the price pressure mechanism identified in Hendershott and Menkveld (2014).

Equation (8) characterizes the direction in which the dealer adjusts her quoted supply schedule in response to inventory changes. Here this adjustment is implemented entirely through prices. In electronic limit-order books dealers jointly choose prices and displayed quantities, so the same inventory-management decision appears empirically as an imbalance in quoted liquidity. Our empirical analysis therefore studies quoted liquidity imbalance, the relative liquidity supplied on the ask and bid sides, as the observable counterpart of the theoretical quote adjustment.

To get intuition for the case with multiple assets, we analyze in more detail the formula (7) for λ. If C is an $N \times N$ matrix, we denote by $|C|$ the infinity norm of C, i.e., the maximum absolute row sum of C. We use the big-O notation extended to matrices in the standard way, with asymptotic behavior considered near the zero matrix.[8]

[8]If $X = X(c)$ is a scalar function with matrix values and $f(c)$ is a positive function (at least for small c), the notation $X = O(f(c))$ means that there exist $M, \eta > 0$ such that for all $c \in (0, \eta)$ the

Corollary 1 provides approximations to the equation (7) in two extreme cases, when the risk aversion coefficient γ is either very high or very low.

Corollary 1. *In the context of Proposition 1, if* $\left|(\gamma k\Sigma_v)^{-1}\right| < 1/4$ *then:*

$$\lambda \;=\; k^{-1} + k^{-1}O\big((\gamma k\Sigma_v)^{-1}\big). \tag{9}$$

If $\left|\gamma k\Sigma_v\right| < 4$ *then:*[9]

$$\lambda \;=\; \gamma^{\frac{1}{2}}k^{-\frac{1}{2}}\Sigma_v^{\frac{1}{2}} + k^{-1}O\big(\gamma k\Sigma_v\big). \tag{10}$$

Equation (9) shows that when γ is large, $\lambda \approx k^{-1}$: the pricing matrix is approximately diagonal. A very risk-averse dealer is more concerned with minimizing the inventory penalty than with exploiting portfolio correlations. In the limit, she sets prices so that the inventory in each asset is expected to drop to zero in the next period. Since the liquidity shocks ε_t are independent across assets, the optimal pricing discounts are set independently, and λ becomes diagonal. A very risk-averse dealer mean-reverts each inventory so quickly that cross-hedging has no value.

2.2 Two Risky Assets

In this section, we focus on $N = 2$ with symmetric fundamental volatility and scalar elasticity to obtain simpler closed-form expressions.

Corollary 2. *Let* $N = 2$ *and suppose that:*

$$\Sigma_v \;=\; \sigma_v^2\begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}, \quad k = \begin{bmatrix} k_0 & 0 \\ 0 & k_0 \end{bmatrix}. \tag{11}$$

ratios $|X(c)|/f(c)|$ are bounded by M. If $X = X(C)$ is a matrix function with matrix values and $f(C)$ is positive definite for C positive definite and sufficiently small (in norm), the notation $X = O(f(C))$ means that the ratios $\big(f(C)\big)^{-1}X(C)$ are uniformly bounded in norm for C sufficiently small.

[9] In Appendix B, we analyze the microfoundations of the dealer's expected utility in the one-dimensional case ($N = 1$), and we show that the inequality $\gamma k\sigma_v^2 \ll 1$ is one of the conditions that are sufficient to obtain a quadratic penalty as in equation (2).

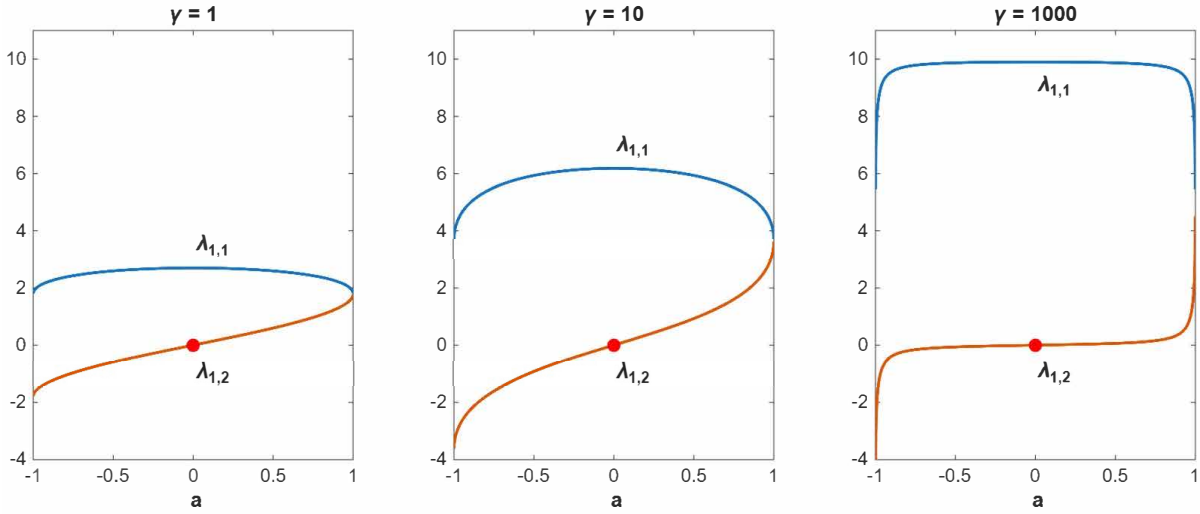
Then the dealer's pricing coefficient is:

$$\lambda = \frac{1}{k_0(m_1+1)(m_2+1)} \begin{bmatrix} m_1+m_2+2 & m_1-m_2 \\ m_1-m_2 & m_1+m_2+2 \end{bmatrix}, \quad \textit{with} \tag{12}$$
$$m_1 = \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1-a)^{-1}+1\right)^{\frac{1}{2}}, \quad m_2 = \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1+a)^{-1}+1\right)^{\frac{1}{2}}.$$

Moreover, the cross-hedging coefficient $\lambda_{1,2}$ is strictly increasing in a, and equals zero at $a = 0$. If $a \in (0,1)$, then viewed as a function of $\gamma \in (0,\infty)$, $\lambda_{1,2}$ has an inverted U-shape: it vanishes at the two extremes and attains a unique interior maximum.

Figure 1 shows comparative statics of the pricing matrix λ with respect to the dealer's risk aversion coefficient (γ) and the fundamental correlation coefficient (a).

Figure 1: Dealer's Pricing Function and Risk Aversion. In the context of Corollary 2, this figure shows how the coefficients $\lambda_{1,1}$ and $\lambda_{1,2}$ of the dealer's pricing function λ vary with the increment correlation between the two components of the fundamental value: $a = \mathsf{corr}(\Delta v_1, \Delta v_2)$. The dealer's risk aversion coefficient is $\gamma = 1$ in the left graph, $\gamma = 10$ in the center graph, and $\gamma = 1000$ in the right graph. In all graphs, the fundamental volatility is $\sigma_v = 1$ and the elasticity matrix is $k = k_0 I_2$, with $k_0 = 0.1$.



As Figure 1 shows, there is significant cross-hedging across the two assets. For instance, suppose the dealer's risk aversion coefficient is $\gamma = 1$, the fundamental volatility is $\sigma_v = 1$, and the elasticity matrix is $k = k_0 I_2$, with $k_0 = 0.1$.[10] Suppose the fundamental correlation between the two assets is $a = 0.6$. Corollary 2 implies that the pricing

[10] We choose a relatively small value of k_0 to indicate that the order flow is not very elastic in the dealer's quotes.

matrix is:

$$\lambda_{(a=0.6)} = \begin{bmatrix} 2.54 & 0.73 \\ 0.73 & 2.54 \end{bmatrix}. \tag{13}$$

Suppose the value of both assets is 10 and the dealer's starting inventory before trading at t is $x_t = [-1, +1]'$, i.e., the dealer is short one unit of asset 1 and long one unit of asset 2. The pricing equation (8) implies that the dealer must set the midquotes at $p_t = v_t - \lambda x_t$:

$$p_t = \begin{bmatrix} 10 \\ 10 \end{bmatrix} - \begin{bmatrix} 2.54 & 0.73 \\ 0.73 & 2.54 \end{bmatrix} \begin{bmatrix} -1 \\ 1 \end{bmatrix} = \begin{bmatrix} 11.81 \\ 8.19 \end{bmatrix}. \tag{14}$$

If instead the dealer had zero units in asset 2, the midquotes would be set at:

$$p_t = \begin{bmatrix} 10 \\ 10 \end{bmatrix} - \begin{bmatrix} 2.54 & 0.73 \\ 0.73 & 2.54 \end{bmatrix} \begin{bmatrix} -1 \\ 0 \end{bmatrix} = \begin{bmatrix} 12.54 \\ 10.73 \end{bmatrix}. \tag{15}$$

Thus, being long one unit of asset 2 helps reduce the pressure to increase inventory in asset 1, and as a result the dealer sets the midquote of asset 1 closer to the fundamental value of 10, at a rate equal to $0.73/2.54 \approx 29\%$ of the extra inventory in asset 2. The rate of cross-hedging increases with the correlation between assets 1 and 2, reaching 100% in the limit when the assets are perfectly correlated.

Another implication of the pricing equation (15) is that, even if the dealer is only short in asset 1, the dealer also increases her midquote in asset 2, by 0.73. The intuition for this fact is similar to that of standard portfolio theory, where in order to reduce the overall risk of a portfolio, it is more efficient to hedge using other assets as well.

If the fundamental correlation were to increase from $a = 0.6$ to $a = 0.8$, the dealer's pricing coefficient would change to:

$$\lambda_{(a=0.8)} = \begin{bmatrix} 2.38 & 1.06 \\ 1.06 & 2.38 \end{bmatrix}. \tag{16}$$

Comparing λ with the pricing matrix from equation (13), the cross-hedging coefficient $\lambda_{1,2}$ increases from 0.73 to 1.06. The intuition is straightforward: When the correlation

between the two assets increases, asset 2 is used at a higher rate for hedging inventory in asset 1. At the same time, the own-hedging coefficient $\lambda_{1,1}$ decreases from 2.54 to 2.38. To understand this result, consider the more extreme case when the fundamental correlation becomes very close to $a = 1$, i.e., when the two assets are perfectly correlated. In that case, the pricing matrix has all entries equal to each other (with a value of 1.79). Thus, only the sum of inventories matters to the dealer and the assets are almost perfect substitutes. This implies that, compared to the case when $a = 0.6$, asset 1 is used less intensely for hedging, since asset 2 has become a better substitute.

If the dealer's risk aversion γ is very large, Corollary 1 implies that λ is approximately diagonal.[11] (This can also be seen in Figure 1 when $\gamma = 1000$.) Intuitively, a very risk-averse dealer is more concerned about minimizing the inventory penalty than maximizing the expected profit component. But in the limit, minimizing the inventory penalty means setting the prices such that the inventory in each asset is expected to drop to approximately zero in each period. This means that the dealer sets prices such that the liquidity shock in each asset is expected to be almost fully reversed in each period. As the liquidity shocks are assumed uncorrelated among the different assets, the pricing discounts are set almost independently of each other, and therefore the pricing coefficient λ is approximately diagonal. Differently said, a very risk-averse dealer mean-reverts her inventory to zero very frequently, and hence hedging across assets has little value.

Note that cross-hedging requires *intermediate* risk aversion. As Corollary 2 shows, for a in $(0, 1)$ the cross-hedging coefficient $\lambda_{1,2}$ has an inverted U-shape in γ: it vanishes both as γ approaches 0 and ∞, and reaches a maximum at an intermediate value of γ. The intuition combines the two limits already discussed: a near-risk-neutral dealer barely hedges at all, while a very risk-averse dealer mean-reverts each inventory to zero so quickly that hedging across assets has no time to add value. Cross-hedging is most valuable in between, when the dealer holds inventory long enough for portfolio considerations to matter, but is averse enough to want to manage it. This is visible in Figure 1: at $a = 0.6$, the amount of cross-hedging $\lambda_{1,2}$ is small at $\gamma = 1$, peaks near $\gamma = 10$, and is small again at $\gamma = 1000$.

[11] In the case of $N = 2$ assets, this can be proved directly using Corollary 2: When γ approaches infinity, m_1 and m_2 are approximately equal to 1, and λ is approximately proportional to the identity matrix, i.e., $\lambda \approx k_0^{-1} I_2$.

2.3 Correlation and the Dealer's Asset Selection

In this section, we analyze how the dealer chooses in which assets to provide quotes. For simplicity, we assume that the dealer already provides quotes in one asset and must choose one more asset, based on the fundamental correlation between them.

To solve this problem, we must calculate the value function of the dealer. The proof of Proposition 1 shows that in the general case with N risky assets, the dealer's value function is a quadratic function of the inventory x:

$$V(x) \;=\; \frac{1}{1-\beta}\Big(\ell' k^{-1}\ell - \beta\alpha \circ \Sigma_L\Big) \;-\; \frac{1}{\beta}x'\alpha x, \tag{17}$$

where $A \circ B = \sum_{i,j=1}^{N} A_{ij}B_{ij}$ denotes the Frobenius inner product (or double dot product) of two $N \times N$ matrices A and B, $\beta < 1$ is the discount rate, α is the $N \times N$ matrix that satisfies equation (A18), and Σ_L is the covariance matrix of the liquidity demand $Q_t^b - Q_t^s$.

To simplify formulas, we define the dealer's normalized value function $W(x)$ as the limit of $(1-\beta)V(x)$ when β converges to 1. According to equation (17), the normalized value function is:

$$W(x) \;=\; \ell' k^{-1}\ell - \alpha \circ \Sigma_L. \tag{18}$$

Corollary 3 provides formulas for the simpler case when $N = 2$, as well as asymptotic limits in the extreme cases when the dealer is either almost risk-neutral or very risk-averse.

Corollary 3. *Let $N = 2$ and suppose that:*

$$\Sigma_v \;=\; \sigma_v^2 \begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}, \quad k \;=\; \begin{bmatrix} k_0 & 0 \\ 0 & k_0 \end{bmatrix}, \quad \ell \;=\; \begin{bmatrix} \ell_1 \\ \ell_2 \end{bmatrix}, \quad \Sigma_L \;=\; L_0 \begin{bmatrix} 1 & b \\ b & 1 \end{bmatrix}. \tag{19}$$

Then the dealer's normalized value function $W(x)$ satisfies:

$$W(x) \;=\; k_0(\ell_1^2 + \ell_2^2) - \frac{\gamma\sigma_v^2 L_0}{2}\big(n_1(1-b) + n_2(1+b) + 2(1+ab)\big), \tag{20}$$

where $n_1 \;=\; m_1(1-a)$ and $n_2 \;=\; m_2(1+a)$, with m_1 and m_2 as in Corollary 2.

When γ approaches zero or infinity, respectively, the dealer's normalized value function satisfies the following asymptotic equations:

$$\begin{aligned} W(x) &= k_0(\ell_1^2+\ell_2^2) - \gamma^{\frac{1}{2}} L_0 \sigma_v k_0^{-\frac{1}{2}} \Big((1-a)^{\frac{1}{2}}(1-b) + (1+a)^{\frac{1}{2}}(1+b) \Big) + O(\gamma^{\frac{3}{2}}), \\ W(x) &= k_0(\ell_1^2+\ell_2^2) - 2\gamma L_0 \sigma_v^2 (1+ab) - \frac{2L_0}{k_0} + O\big(\gamma^{-1}\big), \end{aligned} \tag{21}$$

where the first and second equations correspond to γ small and γ large, respectively.

Moreover, suppose $b > 0$. When γ is small, W is U-shaped in a on $[-1,1]$, with local maxima at $a = \pm 1$ and $W(-1) > W(+1)$. Restricting to $a \in (0,1]$, W has an interior minimum at $a^* = 2b/(1+b^2)$, so its maximum is attained at $a = 1$ if $b < \sqrt{2}-1$ and at $a = 0$ if $b > \sqrt{2}-1$. When γ is large, W is strictly decreasing in a on $[-1,1]$.

Corollary 3 shows that asset selection depends on two correlations: the fundamental correlation a, which enters W through α as a *cross-hedging channel*, and the liquidity-shock correlation b, which enters through Σ_L as a *liquidity-shocks channel*. Throughout this discussion, we assume $b > 0$, the empirically relevant case in which liquidity shocks are positively correlated across the two assets.

The two channels pull the dealer's choice in different directions. Cross-hedging is most effective when the dealer can use inventory in one asset to offset inventory in the other, which requires $|a|$ to be large. But when liquidity shocks are positively correlated, high positive a causes inventories in the two assets to accumulate together, amplifying the inventory penalty. This tension resolves differently depending on the dealer's risk aversion: cross-hedging dominates when γ is small, while the inventory-penalty cost dominates when γ is large. Equation (21) makes this explicit, with the two limits in γ giving opposite comparative statics of W in a.

When the dealer is almost risk-neutral, her dominant concern is the ability to cross-hedge. Cross-hedging is most powerful at extreme correlations, $|a|$ close to 1: at $a = 1$ the two assets are perfect substitutes, and at $a = -1$ they are mirror images, so in either case inventory in one asset can fully offset inventory in the other. The dealer therefore benefits from extreme $|a|$ regardless of sign, giving the U-shape stated in the corollary with maxima at $a = \pm 1$.

The asymmetry $W(-1) > W(+1)$ admits an intuitive interpretation in the basis of

inventory sum and difference coordinates, $s = (x_1 + x_2)/\sqrt{2}$ and $d = (x_1 - x_2)/\sqrt{2}$. Both Σ_v and Σ_L are diagonal in this basis, splitting variance between the two directions according to a and b respectively. At $a = +1$ only the sum direction has nonzero fundamental variance, but that direction also has high liquidity-shock variance when $b > 0$. At $a = -1$ only the difference direction matters, and it has low liquidity-shock variance. The dealer therefore prefers $a = -1$, where her cross-hedging operates along a direction with low inventory accumulation.

In equity markets, fundamental correlations are usually positive, often restricting the dealer's choice to $a \in (0, 1]$. Within this range, the corollary identifies $a = 0$ and $a = 1$ as the candidate maxima, separated by a threshold $b^* = \sqrt{2} - 1 \approx 0.41$. The intuition is a trade-off between two inventory-management strategies: at $a = 0$ the inventory penalty is spread evenly across the two assets (no cross-hedging); at $a = 1$ the assets are perfect substitutes and the penalty concentrates on a single direction, which has high liquidity-shock variance when $b > 0$. Cross-hedging wins when the concentration cost remains below the spread cost, i.e., when $b < b^*$. Empirically, equity liquidity-shock correlations typically lie below b^* under normal trading but may exceed it during periods of market-wide stress, predicting a regime switch in the dealer's optimal asset-selection strategy.

When the dealer is very risk-averse, the picture changes. The dealer mean-reverts inventory so aggressively that at each period her inventory consists almost entirely of the previous period's net demand shock, $x_t \approx \varepsilon^s_{t-1} - \varepsilon^b_{t-1}$. Cross-hedging plays no useful role: the dealer is killing inventory too quickly for portfolio considerations to matter. Instead, the inventory penalty becomes $\gamma \, \Sigma_v \circ \Sigma_L$, which captures the alignment between fundamental volatilities and liquidity-shock variances, and is increasing in a when $b > 0$. The dealer therefore prefers low (or negative) correlation, behaving like a Markowitz investor.

2.4 Inventory Half-Life and Risk Aversion

In this section, we analyze how the dealer's inventory evolves over time and how this evolution is related to the dealer's risk aversion. To get simpler solutions, we consider the

particular case with $N = 2$ assets. As in Corollary 2, we assume that the fundamental volatility is the same for both assets and that the elasticity matrix is proportional to the identity matrix.

We define the inventory half-life as the expected time it takes for the inventory in one asset to be reduced by half. Corollary 4 provides an approximate formula for the inventory half-life when the dealer's risk aversion coefficient is small.

Corollary 4. *Let $N = 2$ and suppose that:*

$$\Sigma_v \;=\; \sigma_v^2 \begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}, \quad k = \begin{bmatrix} k_0 & 0 \\ 0 & k_0 \end{bmatrix}. \tag{22}$$

Let $m_1 = \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1-a)^{-1}+1\right)^{1/2}$ and $m_2 = \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1+a)^{-1}+1\right)^{1/2}$. Then, for any $t \in \mathbb{N}$:

$$\left(I_2 - k_0\lambda\right)^t \;=\; \begin{bmatrix} \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t + \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t & \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t - \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t \\ \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t - \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t & \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t + \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t \end{bmatrix}, \tag{23}$$

and the inventory half-life T starting from an inventory $x_0 = \left[1, 0\right]'$ solves the equation:

$$\left(\frac{m_1-1}{m_1+1}\right)^T + \left(\frac{m_2-1}{m_2+1}\right)^T \;=\; 1. \tag{24}$$

When γ is small, the inventory half-life in asset 1 is approximately:

$$T \;\approx\; \frac{m}{(1+a)^{\frac{1}{2}} + (1-a)^{\frac{1}{2}}}, \quad \textit{with} \quad m \;=\; \gamma^{-\frac{1}{2}}k_0^{-\frac{1}{2}}\sigma_v^{-1}. \tag{25}$$

When assets 1 and 2 are uncorrelated ($a = 0$), equation (25) implies that the inventory half-life $T_{a=0}$ is approximately $m/2$. When assets 1 and 2 are perfectly correlated ($a = 1$), the inventory half-life is $T_{a=1} = m/\sqrt{2} = T_{a=0} \times \sqrt{2}$, which is 41% larger than the inventory half-life in the uncorrelated case. Intuitively, with correlated assets it takes longer to reduce inventory to zero, as the other asset can be used for hedging and the need to reduce the inventory is effectively weaker.

Corollary 4 implies that market makers may appear significantly less risk-averse

than they actually are. Indeed, suppose we observe a dealer's inventory half-life of T. Then, using equation (25) we can estimate the dealer's risk aversion coefficient under the assumption that $a = 0$ or $a = 1$:

$$\gamma_{a=0} \;=\; \frac{k_0^{-1}\sigma_v^{-2}}{4T^2}, \qquad \gamma_{a=1} \;=\; \frac{k_0^{-1}\sigma_v^{-2}}{2T^2}. \tag{26}$$

Thus, if we wrongly assumed that the assets were uncorrelated when in fact they were perfectly correlated, we would be wrong by a factor of $\gamma_{a=0}/\gamma_{a=1} = 1/2$, implying that we would underestimate the dealer's risk aversion by 50%.

The two-asset case understates the magnitude of this bias. A market maker typically trades many correlated assets, and Corollary 5 extends the half-life analysis to N assets with equicorrelated fundamentals.

Corollary 5. *Let $N \geq 2$ and suppose that:*

$$\Sigma_v \;=\; \sigma_v^2\big[(1-a)I_N + a\,\mathbf{1}_N\mathbf{1}_N'\big], \qquad k \;=\; k_0 I_N, \tag{27}$$

where $\mathbf{1}_N$ is the N-dimensional column vector of ones and $a \in \big(-\frac{1}{N-1}, 1\big)$. Let

$$m_1 \;=\; \big(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1-a)^{-1}+1\big)^{1/2}, \qquad m_2 \;=\; \Big(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}\big(1+(N-1)a\big)^{-1}+1\Big)^{1/2}, \tag{28}$$

and denote $\rho_i = \frac{m_i-1}{m_i+1}$ for $i = 1, 2$. Then, for any $t \in \mathbb{N}$:

$$\big(I_N - k_0\lambda\big)^t \;=\; \rho_1^t\, I_N + \frac{\rho_2^t - \rho_1^t}{N}\,\mathbf{1}_N\mathbf{1}_N', \tag{29}$$

and the inventory half-life T starting from an inventory $x_0 = \big[1, 0, \ldots, 0\big]'$ solves the equation:

$$\frac{1}{N}\,\rho_2^T + \frac{N-1}{N}\,\rho_1^T \;=\; \frac{1}{2}. \tag{30}$$

When γ is small, the inventory half-life in asset 1 is approximately:

$$T \;\approx\; \frac{N\,m}{2\Big(\big(1+(N-1)a\big)^{\frac{1}{2}} + (N-1)(1-a)^{\frac{1}{2}}\Big)}, \quad \textit{with} \quad m \;=\; \gamma^{-\frac{1}{2}}k_0^{-\frac{1}{2}}\sigma_v^{-1}. \tag{31}$$

When $N = 2$, Corollary 5 reduces to Corollary 4. Setting $t = 1$ in equation (29) also yields the pricing matrix itself, $\lambda = k_0^{-1}\big[(1-\rho_1)I_N + \frac{\rho_1-\rho_2}{N}\,\mathbf{1}_N\mathbf{1}_N'\big]$: each off-diagonal entry equals $(\rho_1-\rho_2)/(Nk_0)$, positive for $a > 0$ and vanishing in inverse proportion to the number of assets, so with many correlated assets the hedge is spread across all of them and each pairwise quote response is small. To describe the bias for general N, define:

$$\bar{r} \;=\; \frac{\big(1+(N-1)a\big)^{\frac{1}{2}} + (N-1)(1-a)^{\frac{1}{2}}}{N}, \tag{32}$$

so that equation (31) reads $T \approx m/(2\bar{r})$, with $\bar{r} = 1$ when $a = 0$. Repeating the estimation exercise above, an econometrician who observes the inventory half-life T and inverts the single-asset (uncorrelated) formula obtains $\gamma_{a=0} = k_0^{-1}\sigma_v^{-2}/(4T^2)$, while the true risk aversion coefficient is $\gamma = k_0^{-1}\sigma_v^{-2}/(4T^2\bar{r}^2)$. The ratio of the estimated to the true coefficient is therefore:

$$\frac{\gamma_{a=0}}{\gamma} \;=\; \bar{r}^{\,2}, \tag{33}$$

which is decreasing in both the correlation a and the number of assets N. Two limits are instructive. At perfect correlation, $\bar{r}^2 = 1/N$, so the underestimation equals $1-1/N$ and approaches 100% as the number of assets grows: the 50% figure of the two-asset case is the most conservative one. For fixed $a \in (0,1)$, $\bar{r}^2$ converges to $1-a$ as the number of assets grows: the underestimation converges to the correlation itself. At a correlation of $a = 0.44$, the average correlation among the energy stocks that recur as hedge pairs in our sample (Table C.4), the bias is only 5% for $N = 2$, but it reaches 20% at $N = 10$ and 27% at $N = 25$, values that bracket the median and mean number of stocks traded by the market makers in our sample (Table 1).[12] These figures assume a uniform correlation across assets. With heterogeneous correlations the bias depends on the whole distribution of pairwise correlations, and it can be larger: hedged combinations of highly correlated stocks carry little risk and are unwound slowly, lengthening measured half-lives. The recurring hedge pairs in our sample are such stocks (Table C.4).

The mechanism behind the longer half-life also organizes the inventory facts that

[12] At exactly $a = 1$, the half-life defined by equation (30) is infinite for every $N \geq 2$: a fraction $(N-1)/N$ of the initial position lies in directions that the dealer is content to carry, and is expected never to unwind. The approximation (31) measures the initial speed of reversion, which remains finite, so the biases computed from it are conservative.

we document in Section 3. A unit inventory shock to one asset decomposes into a component $1/N$ along the average of all inventories, which mean-reverts at the fast rate associated with m_2, and a component $(N-1)/N$ orthogonal to it, which mean-reverts at the slow rate associated with m_1. As the number of assets grows, the net position across all assets reverts faster, at a rate proportional to $(1+(N-1)a)^{1/2}$, while the position in any single asset reverts more slowly. Slow inventory reversion in individual assets is therefore not evidence of low risk aversion: it is what a risk-averse dealer who trades many correlated assets optimally looks like. Moreover, the ratio of the two reversion rates depends only on N and a, and not on the risk aversion γ, the demand elasticity k_0, or the volatility σ_v.

2.5 Stationary Equilibria

We define a stationary equilibrium as an equilibrium for which the price covariance matrix is constant over time. One can show that there is a unique stationary equilibrium and that any equilibrium converges to the unique stationary equilibrium.

Proposition 2 describes the evolution of the dealer's inventory and price over time and provides formulas for various covariances in a stationary equilibrium.

Proposition 2. *In equilibrium, the dealer's inventory and price satisfy the following formulas:*

$$\begin{aligned} x_{t+1} &= (I_N - k\lambda)x_t - \varepsilon_t, \\ \Delta p_{t+1} &= \lambda\varepsilon_t + \Delta v_{t+1} + \lambda k\lambda x_t, \end{aligned} \tag{34}$$

where $\varepsilon_t = \varepsilon_t^b - \varepsilon_t^s$. Suppose, moreover, that k, Σ_v, and Σ_L commute with each other. Then, in a stationary equilibrium, the inventory variance $\Sigma_x = \mathsf{Var}(x_t)$, price variance $\Sigma_p = \mathsf{Var}(\Delta p_t)$, and price autocovariance $A_p = \mathsf{Cov}(\Delta p_{t+1}, \Delta p_{t+2})$ are constant and satisfy:

$$\begin{aligned} \Sigma_x &= k^{-1}\lambda^{-1}(2I_N - k\lambda)^{-1}\Sigma_L, \\ \Sigma_p &= 2\lambda^2(2I_N - k\lambda)^{-1}\Sigma_L + \Sigma_v, \\ A_p &= -k\lambda^3(2I_N - k\lambda)^{-1}\Sigma_L. \end{aligned} \tag{35}$$

If $\left|(\gamma k \Sigma_v)^{-1}\right| < 1/4$*, we have:*

$$\begin{aligned} \lambda &\approx k^{-1}, \\ \Sigma_x &\approx \Sigma_L, \\ \Sigma_p &\approx 2k^{-2}\Sigma_L + \Sigma_v, \\ A_p &\approx -k^{-2}\Sigma_L. \end{aligned} \tag{36}$$

If $\left|\gamma k \Sigma_v\right| < 4$*, we have:*

$$\begin{aligned} \lambda &\approx \gamma^{\frac{1}{2}} k^{-\frac{1}{2}} \Sigma_v^{\frac{1}{2}}, \\ \Sigma_x &\approx \frac{1}{2}\left(\gamma k \Sigma_v\right)^{-\frac{1}{2}} \Sigma_L, \\ \Sigma_p &\approx \left(\gamma k^{-1}\Sigma_L + I_N\right)\Sigma_v, \\ A_p &\approx -\frac{1}{2}\,\gamma^{\frac{3}{2}} k^{-\frac{1}{2}} \Sigma_L \Sigma_v^{\frac{3}{2}}. \end{aligned} \tag{37}$$

Corollary 6 focuses on the simpler case in which there are only $N = 2$ assets and provides closed-form results. As in Corollary 2, we assume that the fundamental volatility is the same for both assets and that the elasticity matrix is proportional to the identity matrix.

Corollary 6. *Let* $N = 2$ *and suppose that:*

$$\Sigma_v \;=\; \sigma_v^2 \begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}, \quad k = \begin{bmatrix} k_0 & 0 \\ 0 & k_0 \end{bmatrix}, \quad \Sigma_L = \begin{bmatrix} L_0 & 0 \\ 0 & L_0 \end{bmatrix}. \tag{38}$$

Then, in a stationary equilibrium, the inventory variance $\Sigma_x = \mathsf{Var}(x_t)$*, price variance* $\Sigma_p = \mathsf{Var}(\Delta p_t)$*, price autocovariance* $A_p = \mathsf{Cov}(\Delta p_{t+1}, \Delta p_{t+2})$*, and price correlation*

$\rho_p = \mathsf{Corr}(\Delta p_{1,t}, \Delta p_{2,t})$ *are constant and satisfy:*

$$
\begin{aligned}
\Sigma_x &= \frac{L_0}{8m_1m_2}\begin{bmatrix} m_1^2m_2+m_1m_2^2+4m_1m_2+m_1+m_2 & (m_1-m_2)(1-m_1m_2) \\ (m_1-m_2)(1-m_1m_2) & m_1^2m_2+m_1m_2^2+4m_1m_2+m_1+m_2 \end{bmatrix}, \\
\Sigma_p &= \frac{2L_0\begin{bmatrix} m_1^2+m_2^2+m_1+m_2 & (m_1-m_2)(m_1+m_2+1) \\ (m_1-m_2)(m_1+m_2+1) & m_1^2+m_2^2+m_1+m_2 \end{bmatrix}}{k_0^2m_1m_2(m_1+1)(m_2+1)} + \sigma_v^2\begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}, \\
A_p &= \frac{2L_0\begin{bmatrix} m_1^3+m_2^3+2m_1^2+2m_2^2+m_1+m_2 & (m_1-m_2)(m_1^2+m_2^2+2m_1+2m_2+m_1m_2+1) \\ (m_1-m_2)(m_1^2+m_2^2+2m_1+2m_2+m_1m_2+1) & m_1^3+m_2^3+2m_1^2+2m_2^2+m_1+m_2 \end{bmatrix}}{k_0^2m_1m_2(m_1+1)^2(m_2+1)^2}, \\
\rho_p &= \frac{\frac{(m_1-m_2)(m_1+m_2+1)}{m_1m_2(m_1+1)(m_2+1)}+a\nu_0}{\frac{m_1^2+m_2^2+m_1+m_2}{m_1m_2(m_1+1)(m_2+1)}+\nu_0}, \quad \textit{with} \quad \nu_0 = \frac{k_0^2\sigma_v^2}{2L_0}, \\
m_1 &= \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1-a)^{-1}+1\right)^{\frac{1}{2}}, \quad m_2 = \left(4\gamma^{-1}k_0^{-1}\sigma_v^{-2}(1+a)^{-1}+1\right)^{\frac{1}{2}}.
\end{aligned}
\tag{39}
$$

If $L_0 \gg k_0\sigma_v^2 \gg 1/\gamma$*, then* m_1 *is close to* m_2*,* ν_0 *is close to zero, and the price correlation* ρ_p *between assets 1 and 2 is close to zero.*

Thus, when the demand noise coefficient L_0 and the dealer's risk aversion γ are sufficiently high, the asset return correlation is low at higher frequencies. Indeed, a more risk-averse dealer does less inventory hedging across assets, and thus quotes are less correlated across assets. If by assumption the noise demand is also uncorrelated across assets, then the price correlation between assets is also low. In our model therefore we obtain a theoretical version of the Epps effect. Epps (1979) finds that stock returns are relatively uncorrelated at high frequencies (less than 10 minutes) but approach a constant correlation at low frequencies (more than 1 day). In our model, the inventory effect in prices induces a mean-reverting component at higher frequencies that disappears at lower frequencies, for which the random walk component of prices dominates. At low frequencies, the return correlation approaches the fundamental correlation, which is assumed constant over time.

3 Empirical Evidence

The model makes one central prediction. A risk-averse market maker who is long an asset adjusts her quotes to sell correlated assets, so the off-diagonal elements of the pricing matrix are positive. The theory also says at what scale to look: cross-hedging declines with the dealer's risk aversion, and with many correlated stocks the hedge is spread across all of them (Corollary 5), so each pairwise quote response should be small and the empirical object must aggregate across the rest of the portfolio.

We take this prediction, direction and magnitude, to message-level exchange data. In the stylized model environment, the dealer's quote is one-dimensional, so inventory management operates entirely through the midquote. Modern electronic limit order books, however, allow market makers to manage inventory along two dimensions. They jointly choose the prices at which they stand ready to trade and the quantities displayed at those prices. Both margins determine the probability that incoming order flow executes against the market maker and therefore both are instruments for inventory management. A market maker seeking to reduce a long inventory may lower ask prices, increase displayed ask depth, reduce bid depth, or combine these actions. Each shifts expected executions toward selling and serves the same inventory-management objective as quote shading in the model.

Our empirical analysis, therefore, focuses on the market maker's quoted liquidity imbalance, measured as resting ask depth minus resting bid depth in dollars. This measure captures how the market maker shifts her quoted supply schedule toward selling or buying. It is the empirical analogue of the inventory-driven quote adjustment in the model because it reflects the direction of the market maker's inventory-management decision whether it is implemented through prices, displayed quantities, or a combination of the two. The theory predicts a directional change in the dealer's quoted supply schedule, and the data allow us to observe that direction through the imbalance of posted liquidity. We show that such cross-hedging is widespread across market makers, concentrated in a small number of correlated stocks, and modest in intensity, the profile the theory assigns to risk-averse dealers trading many correlated assets. Position-based tests of the portfolio dimension, such as Naik and Yadav (2003), find inventory management that

is predominantly stock-by-stock; the theory explains why one should look in the quotes instead, since a quoted imbalance becomes a position only as order flow arrives to fill it.

The model features a single market maker (the "dealer"), while the data feature multiple market makers. The theory should be read as each market maker's best response to the order flow she faces. The elastic component k is then the slope of her residual demand, the additional flow she executes per unit of quote improvement given the quotes of the other market makers, and the inelastic component is the flow routed to her by obligations and client relationships. Three assumptions make this reading exact: the residual order flow is locally linear in her quotes, which is a good approximation because optimal inventory-induced quote adjustments turn out to be small; inventories are private, so the other market makers' quotes do not respond to her inventory and only add noise to her order flow; and each market maker is small enough to take the flow schedule as given. Under these assumptions, each market maker's optimal policy retains the form of the model, a midquote equal to $v_t - \lambda x_t$, with the pricing matrix λ of equation (7) evaluated at her own risk aversion and residual elasticity. Competition, as in Ho and Stoll (1983), operates through the allocation of order flow across market makers rather than through the mapping from each market maker's inventory to her own quotes. It changes the levels of spreads and price impact, but not the signs and comparative statics of λ that we test. Accordingly, all our tests relate each market maker's quotes to her own inventory.

3.1 Data and Market Maker Identification

We obtain detailed audit trail data from the STAMP 4.0.3 internal feed of the Toronto Stock Exchange over 45 trading days from January 3 to March 6, 2006, a subset of the sample used by Anand and Venkataraman (2016), AV16 hereafter. The TSX is an electronic limit order book market. The data is a message-level record of every order, cancellation, and execution for all TSX-listed securities time-stamped to millisecond resolution. Each message carries information on the member firm (e.g., TD Bank) and the user identification (ID) within a member firm participating on the active and passive side of each trade. Traders at the member firm are uniquely assigned a user ID (`UserId`),

and this is the port through which orders are submitted to TSX. Traders place orders on behalf of their own proprietary (principal) account or serve as brokers and enter orders on behalf of their clients. The TSX data identify whether an order is submitted on behalf of the proprietary account or the client. However, individual clients within a user ID are not identified in the data. We keep all messages related to a principal's activity and reconstruct, from the signed flow, the full intraday inventory path of every user in every stock, which is a continuous signed position carried across days. We restrict our analysis to the exchange's Regular Trading Hours (09:30–16:00 Eastern).

The exchange assigns one member firm as the Designated Market Maker for each stock. A designated market maker that auto-participates is allocated a fixed share of small incoming orders, and the auto-participation flag is recorded in the feed (but not shown to the market). Following AV16, we read the designation from the specialist account type (`AccountType=ST`) and the auto-participation (`RTAutofill`) flag. We sort principal users into three groups. A Designated Market Maker (DMM) is a user that is designated and trades only its assigned stocks. An Endogenous Market Maker (EMM) is a user that carries no DMM designation and provides liquidity through sustained two-sided market participation, with a minimum of ten trades and both buys/sells in at least three separate five-minute windows. A user that is designated in some stocks and quotes voluntarily in others is Both, the counterpart of the specialist trading in non-assigned stocks in AV16. Our EMM group is broader than the AV16 algorithm-free construct, because we identify it from market activity rather than from designation alone. The universe is 549 users, 69 designated, 252 endogenous, and 228 Both. Throughout, we refer to each of these users as a market maker.

We provide a description of our universe of stocks and market maker activity in Table 1. In Panel A of Table 1, we describe the universe of stocks included in our analysis. A median of 1,183 stocks trade each day, with 219 trades and CAD 4.0 million of volume per stock-day, at a price of CAD 15.2 and a relative spread of 2.15%. The spread and the price are in line with AV16. Our per-stock trading activity is lower, because we keep a broader and thinner tail of the market, 1,183 traded stocks a day instead of 900. The daily average return is also higher, 0.25% as opposed to -0.02%, because the first quarter of 2006 was a stronger market than the full year AV16 covers.

Table 1: Sample Characteristics

The table presents characteristics for our Toronto Stock Exchange sample. Panel A presents the stock characteristics across stock-days for stocks matched to Canadian Financial Markets Research Center (CFMRC) data. Panel B presents the market maker-trading day characteristics across the 549 identified market makers. Market makers are split into three groups: *DMM* defined as market making in designated stocks only that carry formal specialist obligations; *EMM* defined as voluntary two-sided quoters endogenously market making in stocks with no obligations; and *Both* defined as market makers that endogenously quote two sides in some stocks while carrying market making obligations in others. Panel A reports means, medians, and quartiles across days; Panel B reports means with medians in parentheses. The sample period is January 3 to March 6, 2006. All variables are defined in Table C.1.

Panel A. Sample Stocks

	Mean	Median	P25	P75
Number of stocks	1,182	1,183	1,176	1,190
Trades per stock	222.6	218.7	211.5	238.7
Share volume per stock	233,139	232,540	219,881	251,367
Dollar volume per stock (CAD)	3,960,111	3,963,908	3,648,858	4,336,689
Closing price (CAD)	15.2	15.2	15.0	15.4
Quoted spread (CAD)	0.191	0.191	0.177	0.208
Relative spread (%)	2.191	2.152	2.101	2.228
Daily return (%)	0.254	0.381	-0.080	0.681

Panel B. Market Makers

	All	DMM	EMM	Both
Trading days	41.0	44.2	37.3	44.1
	(45.0)	(45.0)	(42.0)	(45.0)
Distinct securities (sample)	24.4	17.1	18.4	33.1
	(13)	(10)	(7)	(17)
Stocks traded per day	15.3	12.5	11.0	20.9
	(8)	(9)	(3)	(12)
Trades per day	383.0	431.2	343.5	412.1
	(117.9)	(268.2)	(41.5)	(241.7)
Dollar volume per day (CAD)	4,163,182	2,740,359	3,600,163	5,216,058
	(1,098,522)	(1,330,943)	(425,832)	(1,845,646)
Net portfolio zero reversions	1.858	3.625	0.970	2.305
	(0.622)	(1.419)	(0.250)	(1.100)
Days ending flat (%)	2.805	0.000	5.371	0.817
	(0.000)	(0.000)	(0.000)	(0.000)
Passive-fill share (%)	47.3	30.9	58.8	39.6
	(47.8)	(30.8)	(58.6)	(38.2)

Panel B of Table 1 provides market maker characteristics at the market maker level. The median market maker trades 118 times a day and turns over CAD 1.1 million, near

the specialist averages in AV16, 143 trades and CAD 1.5 million, though our pooled mean of 383 trades reflects a thicker tail of highly active market makers their specialist sample lacks. Participant groups exhibit distinct behaviors. The median DMM trades 268 times a day and fills passively 31% of the time, while crossing zero-inventory multiple times a day. The median endogenous market maker trades 42 times a day and rarely crosses zero inventory. They mostly fill their orders passively, 59% per day. Designated market makers turn inventory, while endogenous market makers post and wait.

3.2 Inventory Management

Market makers manage the inventory they carry, and that inventory moves their quotes, as shown for NYSE specialists by Hendershott and Menkveld (2014) and Comerton-Forde, Hendershott, Jones, Moulton, and Seasholes (2010). We document inventory management on the TSX at two levels, the position in a single stock and the net position across the whole portfolio, in Table 2. Every statistic is built from minute-level positions.

Index a market maker by m, a stock by s, a trading day by d, and a one-minute bin of the trading day by $\tau \in \{0, \ldots, 389\}$. Let $q_{msd\tau}$ be the market maker's signed share position in stock s, cumulated from her fills and carried across days, and $p_{sd\tau}$ the prevailing mark price; the stock-level, mark-to-market inventory in dollars is $I_{msd\tau} = q_{msd\tau}\, p_{sd\tau}$. Summing across the stocks held on a day gives the two portfolio-level series, the net portfolio and the gross portfolio,

$$N_{md\tau} = \sum_s I_{msd\tau}, \qquad G_{md\tau} = \sum_s |I_{msd\tau}|.$$

Every stock-level statistic is formed per stock-day and averaged equal-weight across the market maker's stock-days $\mathcal{D}_m$; every portfolio-level statistic is formed per market-maker-day and averaged across her days. We write $\mathcal{M}_{msd}$ for the minutes of day d in which stock s carries a recorded position and $\bar{\tau}_{msd} = \max \mathcal{M}_{msd}$ for its last such minute, with $\mathcal{M}_{md}$ and $\bar{\tau}_{md}$ the analogues for the portfolio. Each cell of Table 2 is the cross-maker mean of the per-maker statistic, with the cross-maker standard deviation in parentheses, reported for all market makers and separately for the DMM and EMM roles. We also present the median statistics in Appendix Table C.2. The role groups

DMM and EMM are fixed per market maker over the sample; the 228 makers who do both are included in All but not reported separately.

We start with intraday dispersion of positions held. Within a stock-day the range and volatility are the maximum minus the minimum and the population standard deviation of the signed inventory, each averaged over the market maker's stock-days. The net portfolio's range (CAD 794,000) and volatility (CAD 207,000) are a small fraction of the gross portfolio they sit on. This implies that the position the market maker actually bears during the day is far smaller than the gross inventory she intermediates.

Moving on to the end-of-day inventory positions, the stock is the snapshot at the day's last minute, $I^{\text{eod}}_{msd} = I_{msd\bar{\tau}_{msd}}$, and the mean signed and mean absolute inventories are the daily averages across market makers

$$\frac{1}{|\mathcal{D}_m|}\sum_{(s,d)\in\mathcal{D}_m} I^{\text{eod}}_{msd}, \qquad \frac{1}{|\mathcal{D}_m|}\sum_{(s,d)\in\mathcal{D}_m} \left|I^{\text{eod}}_{msd}\right|.$$

The signed mean is near zero when the market maker carries no standing directional position. At the portfolio level the signed mean uses the end-of-day net $N^{\text{eod}}_{md} = N_{md\bar{\tau}_{md}}$ and the absolute mean uses the end-of-day gross G^{eod}_{md}. Therefore, the absolute mean is the position in one stock at the stock level and the gross inventory at the portfolio level.

The single-stock position is carried, not cleared. Endogenous market makers hold the heaviest positions in individual stocks, a mean end-of-day absolute position of CAD 544,000 against CAD 158,000 for designated market makers, with means far above the medians (CAD 179,000 and 35,000) because a few market makers carry large portfolios. At the portfolio level the gross position is an order of magnitude larger, a mean of CAD 12.7 million, while the net is near zero at the center of the distribution: the mean signed net is CAD -3.7 million, a handful of directional makers, but essentially flat at the median (CAD 6,000). The market maker intermediates a large gross position while bearing only a small net inventory. This observation indicates that market makers are holding offsetting positions within their portfolio.

Zooming into the end of day held inventory, we analyze the long and short held positions. On day d let H_{md} be the number of stocks the market maker holds and L_{md} (S_{md}) the number whose end-of-day position is long (short), $\operatorname{sign} q^{\text{eod}}_{msd} > 0$ (< 0); the

Table 2: Inventory Dynamics

The table describes inventory dynamics at two levels: the single stock (the market maker's position in one name) and the portfolio (her net position summed across all names). The stock level is the market maker's signed mark-to-market dollar position in a single stock, $I_{msd\tau} = q_{msd\tau}p_{sd\tau}$; the portfolio level is the net inventory $N_{md\tau} = \sum_s I_{msd\tau}$, and the gross inventory is $G_{md\tau} = \sum_s |I_{msd\tau}|$. Panel A: *mean signed* and *mean absolute* are the end-of-day position (last-minute snapshot) and its absolute value averaged over days, the absolute row using the gross portfolio at the portfolio level. *Intraday range* and *volatility* are the within-day maximum-minus-minimum and population standard deviation of the signed series (stock inventory at the stock level, net portfolio at the portfolio level), averaged over days. *Zero-reversions per day* counts within-day zero inventory touching or crossing; *minutes between reversions* is the mean gap between consecutive reversions. *Share long/short* is the percent of held stocks ending the day long/short. Stock-level statistics are per stock-day averaged equal-weight across the maker's stock-days; portfolio-level per market-maker-day averaged across days; dollar rows in thousands of CAD. Panel B pairs each market maker's portfolio and stock zero-reversion rates and tests $\delta_m = B_m - S_m$ with a paired t-test, $t = \bar{\delta}/(s_\delta/\sqrt{n})$ on $n-1$ degrees of freedom (two-sided), with the Wilcoxon signed-rank test as the skew-robust companion. Parentheses are standard deviations in Panel A and the standard error of $\bar{\delta}$ in Panel B. Columns are All, DMM, and EMM. All variables are defined in Table C.1.

	Stock			Portfolio		
	All	DMM	EMM	All	DMM	EMM
Panel A. Inventory dynamics						
Intraday dispersion (000s CAD)						
Intraday range	69.5	47.6	79.6	794.1	277.9	568.1
	(122.4)	(58.2)	(153.9)	(2,349)	(344.1)	(1,193)
Intraday volatility	17.6	11.6	20.2	207.3	65.6	143.0
	(33.6)	(14.5)	(41.7)	(646.9)	(92.0)	(303.3)
End-of-day position (000s CAD)						
Mean signed	-49.8	-19.1	33.1	-3,729	-218.1	-109.6
	(1,113)	(146.5)	(1,034)	(41,013)	(1,327)	(5,211)
Mean absolute	459.1	157.7	544.4	12,729	1,556	6,923
	(1,334)	(621.0)	(1,141)	(65,411)	(6,302)	(18,873)
End-of-day sign mix (portfolio, %)						
Share long				46.2	49.2	43.1
				(22.0)	(14.8)	(27.6)
Share short				40.6	48.0	39.3
				(21.9)	(14.0)	(27.6)
Turnover						
Zero reversions per day	0.61	1.13	0.38	1.86	3.62	0.97
	(1.09)	(1.64)	(1.14)	(3.31)	(5.60)	(2.47)
Minutes between reversions	62.4	71.2	59.7	43.8	41.4	51.1
	(29.6)	(28.7)	(36.0)	(42.1)	(29.9)	(55.0)

Panel B. Portfolio vs single stock: paired test of zero reversions per day

	All	DMM	EMM
Mean difference (portfolio − stock)	1.25	2.49	0.59
	(0.11)	(0.53)	(0.10)
Median difference	0.29	0.86	0.06
t (df)	11.6 (548)	4.7 (68)	6.1 (251)
p-value	0.00	0.00	0.00
Wilcoxon p	0.00	0.00	0.00

end-of-day share long and short are $100\,L_{md}/H_{md}$ and $100\,S_{md}/H_{md}$, averaged over the market maker's days and defined at the portfolio level only. Long and short need not sum to one hundred because some held stocks end flat. At the close the portfolio is roughly balanced, 46% of held stocks long and 41% short (49% and 48% for DMMs). The inventory is two-sided, long in many stocks and short in others, with a net that is small and worked around zero rather than parked at it.

We finally turn to the within day behavior. We measure the number of time the stock and portfolio inventory reaches or crosses zero within the day in *Zero-reversions per day*. Zero-reversions per day averages over the market maker's stock-days; *Minutes between reversions* averages the gaps $\tau_k - \tau_{k-1}$ over the reversion minutes of a day, over the days on which the maker that reverses back to zero at least once. The portfolio-level counterparts apply the same two definitions to the net portfolio $N_{md\tau}$.

The net portfolio reverts to zero far more often than any single stock, and this is an important contrast. A position in one stock reaches or crosses zero 0.61 times a day (1.13 for DMMs, 0.38 for EMMs), while the net portfolio does so 1.86 times (3.62 for DMMs, 0.97 for EMMs), and the gaps between reversions are shorter for the portfolio, 44 against 62 minutes.

We test the stock and portfolio zero reversion difference within market maker. Writing S_m and B_m for market maker m's stock- and portfolio-level zero-reversion rates as just defined, both observed for every maker, we compare them on the difference $\delta_m = B_m - S_m$ and test $H_0 : \mathbb{E}[\delta_m] = 0$ against a two-sided alternative with a paired (one-sample) t-test. With $\bar{\delta} = n^{-1}\sum_m \delta_m$ and $s_\delta = \big[(n-1)^{-1}\sum_m(\delta_m - \bar{\delta})^2\big]^{1/2}$, the statistic is the standardized mean difference

$$t = \frac{\bar{\delta}}{s_\delta/\sqrt{n}} \underset{H_0}{\sim} t_{n-1}, \qquad p = 2\,\Pr\big(T_{n-1} > |t|\big),$$

with $s_\delta/\sqrt{n}$ the standard error of $\bar{\delta}$ and $n-1$ degrees of freedom.

The net portfolio reverts to zero 1.25 more times per day than a single stock ($\bar{\delta} = 1.25$, $t = 11.6$, $p = 0.00$), and the gap holds within DMMs ($+2.49$, $p = 0.00$) and EMMs ($+0.59$, $p = 0.00$). Because the paired differences are right-skewed (median 0.29) we confirm the result with the Wilcoxon signed-rank test, which ranks the absolute

differences and refers the signed-rank sum to its null distribution. The rank test also rejects the null of no difference in every group. A sum can return to zero more often than its parts only if the parts offset: a short built in one stock pulls the net through zero while each stock continues to be carried in the portfolio. The market maker manages inventory across stocks, not stock by stock.

A position that is carried, only partly turned over, and offset by opposite positions elsewhere is inventory managed across the portfolio. The near-zero net portfolio inventory is not an empty portfolio; it is opposite positions held at once, long in some stocks and short in others. The offsets do not arrive on their own: customer flow in correlated stocks moves together, if anything, and flow that moves together leaves behind positions that move together, a one-sided portfolio. A portfolio that is long some stocks and short others, with the net worked through zero while every stock is carried, has to be built. This is the pattern the theory assigns to a risk-averse dealer who cross-hedges: the net portfolio inventory reverts fast because it is expensive to hold, and the stocks revert slowly because, hedged against one another, they are cheap to hold (Corollary 5). Which stocks carry the offset, and whether it tracks return comovement, is the subject of the cross-hedging analysis that follows.

3.3 Cross-Hedging

So far, we have provided evidence on inventory management within a stock and across the portfolio. Our theoretical object is the dealer's quoted supply schedule. In the model this schedule is represented by the midquote because price is the only strategic variable. In the data, market makers jointly choose prices and displayed quantities, so the same inventory motive is reflected in the balance of liquidity supplied on the two sides of the market. We therefore study the directional imbalance of the market maker's quoted liquidity rather than price alone.

Specifically, for each market maker, stock, and minute we measure the quoted liquidity imbalance as $\text{imb}_i = \text{rest_ask} - \text{rest_bid}$ in dollars, the resting ask depth minus resting bid depth. Positive values indicate that the market maker is supplying relatively more liquidity to buyers than to sellers, making future executions more likely to reduce

a long inventory position. A market maker who is long in stock i and tries to hedge by selling a correlated stock j should therefore shift the quoted liquidity in j toward the ask side, producing a positive imbalance. We define the rest-of-portfolio imbalance as the sum of these imbalances across all other stocks traded by the market maker: $\text{imb}_{\text{not-}i} = \sum_{j \neq i} \text{imb}_j$.

To investigate whether a market maker's position in i is related to the ask-side imbalance in the rest of her portfolio, we compute thirty-minute rolling correlations (rc) between $\text{imb}_{\text{not-}i}(t+5)$ and the market maker's inventory in i at time t, so inventory precedes the quote response by five minutes. With a single regressor the sign of rc is the sign of the response. A *cross-hedging episode*, or simply an episode, is a run of at least five consecutive minutes with $rc > 0$ for a given market maker, day, and focal stock. We retain episodes whose covariance between the lagged focal inventory and the rest-of-portfolio imbalance is positive, in agreement with the selection; about 11 percent fail this check and are excluded. Within an episode we decompose the covariance of the same two series, which share that sign, into named hedge stocks,

$$\text{cov}(\text{inv}_i, \text{imb}_{\text{not-}i}) = \sum_{j \neq i} \text{cov}(\text{inv}_i, \text{imb}_j).$$

This setting allows us to analyze not only the imbalance of the whole portfolio, but also which stocks carry it. Throughout, we call stock i the *focal* stock and a stock j whose imbalance covaries positively with the focal inventory a *hedger*; a focal–hedger pair is one such combination.

We identify 879,942 cross-hedging episodes, run by 537 of the 549 market makers over 12,645 market maker-stock pairs (Table C.3). Table C.4 presents the list of the most common stock cross-pairs. The intensity of cross-hedging is limited, however. The median episode lasts 26 minutes, so episodes cover on the order of a tenth of the trading day for the average market maker-stock pair, and within the median episode the rest-of-portfolio imbalance offsets 13 percent of the focal inventory covariance, with 69 percent of episodes below 30 percent. The imbalance is also concentrated. In the median episode the top three hedge stocks carry 86 percent of the positive covariance. That pooled figure hides the structure. For the DMM and Both groups the imbalance

is tight, the top three stocks carrying 100 and 99 percent over a median of four and six contributing stocks. For endogenous market makers the top three carry 62 percent over a median of 24 stocks.

The key empirical distinction between the model's cross-hedging mechanism and the market maker's own-inventory response is directional. The own-inventory response operates through inventory accumulated in the hedger itself, whereas cross-hedging operates through inventory accumulated in the focal stock. Table 3 exploits this distinction. When customers sell in a sector, the market maker becomes long both the focal stock and the hedger together, and her quoted liquidity in the hedger tilts toward the ask side through the own-inventory response, the diagonal of λ. Because correlated stocks tend to receive correlated order flow, this mechanism also predicts that the fraction of positive focal-inventory/quoted-liquidity responses should increase with the return correlation between the two stocks, even in the absence of cross-hedging. The two mechanisms differ in one observable implication: whether the market maker's inventories in the focal stock and the hedger move together. The own-inventory response requires the two inventories to be positively related, whereas cross-hedging does not.

Table 3 separates these two mechanisms using two different correlations. The rows sort focal–hedger pairs by their intraday return correlation, computed on 30-minute returns for the full sample period, the empirical counterpart of the fundamental correlation that determines the strength of cross-hedging (Corollary 2). Measured at this horizon the correlation is attenuated by the Epps effect of Section 2.5, so sorting on it if anything flattens the gradients the table looks for. The columns instead classify pair-episodes, one focal–hedger pair on one market maker-day, by the correlation between the market maker's inventories in the focal stock and the hedger over the episode.

Inventory correlation is the Pearson correlation between the market maker's minute-level signed dollar inventory in the focal stock, lagged five minutes, and her contemporaneous inventory in the hedger, over the minutes of the episode window in which both positions are observed (at least twenty overlapping minutes; pair-episodes with fewer overlapping minutes, or with no inventory variation, are undetermined and enter only the All column). A pair-episode is Negative if the inventory correlation is at or below -0.1, Near-zero if strictly between -0.1 and 0.1, and Positive if at or above 0.1. Within

each inventory-correlation group, the table reports the share of pair-episodes in which the covariance between the focal inventory and the hedger's quoted liquidity imbalance is positive. The final row ("Gradient") summarizes the increase in this share from the lowest to the highest return-correlation row within each inventory-correlation group, as the difference between the most negative and most positive return correlation within the portfolio correlation. It therefore measures how much more likely the market maker is to increase quoted liquidity on the ask side of the hedger as the two stocks become more highly correlated, holding fixed the relation between her inventories in the two stocks.

If the quoted imbalance is driven primarily by the own-inventory response, it should be strongest when the market maker's inventories move together (the Positive column) and should weaken or disappear when they are unrelated or move in opposite directions. In the Negative column the own-inventory account in fact predicts a share below one half, and it predicts a flat row: once the two inventories move against each other, the return correlation between the stocks has no remaining channel through which to act. In contrast, the model predicts that the increase in positive focal-inventory/quoted-liquidity responses as return correlation rises should persist even when the market maker's inventories are uncorrelated or negatively correlated.

The increase in ask quoted liquidity is present in every inventory-correlation group. Indeed, it is strongest when the market maker's inventories move in opposite directions, rising by 0.14 across the return-correlation rows, compared with 0.11 when the inventories move together, the group in which the own-inventory response has the greatest scope to operate. The level of the share sits above one half in every cell, including where the theory predicts a buy-side imbalance. Two forces, common to all cells, set the level: the episode screen, which selects on a positive relation between the focal inventory and the rest of the portfolio's imbalance and so nudges every hedger's share upward; and a sample-wide ask-side stance, consistent with market makers leaning against the quarter's rising market (Table 1) and with the negative mean signed positions in Table 2.

Neither force can produce a gradient: the theory's content is the gradient, not the level. A quoted liquidity imbalance that becomes increasingly tilted toward the ask side as return correlation rises, even when the market maker's inventory in the hedger is unrelated to or moving opposite the focal inventory, is inconsistent with the own-

inventory response and is the signature of the off-diagonal response the model predicts: the market maker's quoted liquidity in the hedger responds to inventory accumulated in the focal stock rather than to inventory accumulated in the hedger itself.[13] This distinction between the off-diagonal and diagonal responses is the empirical content of the model's central prediction. The two exhibits divide the labor: Table 2 establishes that inventory is managed across the portfolio, and Table 3 establishes that the quoted response is driven by the focal inventory, which separates the off-diagonal of λ from a long-short portfolio selected on relative prices.

To test whether the gradient is different from 0, we re-estimate the gradient market maker by market maker. The resulting gradient is the increase in her share quoted to sell per unit of focal–hedger return correlation, and is defined whenever she has episodes in at least three of the five correlation bins, so fewer market makers enter than the full sample. We then treat each market maker's gradient as one observation and report, across market makers, the mean and median gradient, a two-sided t-test against zero, and the Wilcoxon signed-rank test. Every market maker enters the cross-maker test equally.

The gradient is positive and highly significant in every inventory-correlation group: Negative +0.136 ($t = 7.0$), Near-zero +0.133 ($t = 7.2$), and Positive +0.075 ($t = 4.9$), all $p = 0.00$ across roughly 390 market makers. The share quoted to sell therefore rises with the focal–hedger return correlation not only where the two inventories move together, but at least as steeply where they move independently or oppositely, the opposite of what a pure own-inventory response would produce.

[13] A quote rule that tracks returns rather than inventories would also inherit the pair's return correlation. Its natural sign runs against the pattern: a market maker who quotes to sell into strength while accumulating a long position on weakness generates a negative covariance between the focal inventory and the hedger's ask-side imbalance.

Table 3: Cross-Hedging and the Hedger's Own Inventory

The table separates cross-hedging from the market maker's own-inventory response. Each pair-episode is classified by the correlation between the market maker's inventory in the focal stock and her inventory in the hedger over the episode window. In Panel A, we report the fraction of pair-episodes with a positive covariance between the focal stock inventory and the hedger's quoted liquidity imbalance, by inventory and return correlation. *Inventory correlation* (columns) is the Pearson correlation between the market maker's minute-level signed dollar inventory in the focal stock, lagged five minutes, and her contemporaneous inventory in the hedger, over the minutes of the episode window in which both positions are observed (at least twenty overlapping minutes; pair-episodes with fewer overlapping minutes, or with no inventory variation, are undetermined and enter only the All column). A pair-episode is Negative if this correlation is at or below -0.1, Near-zero if strictly between -0.1 and 0.1, and Positive if at or above 0.1. The own-inventory response can mechanically raise the share quoted to sell only where the two inventories move together, the Positive column. *Return correlation* (rows) is the 30-minute intraday stock-return correlation between i and j, a single pairwise value estimated once over the full sample, so it does not vary across a pair's episodes. Bins are left-closed: < -0.1, $[-0.1, 0)$, $[0, 0.1)$, $[0.1, 0.3)$, and > 0.3. The gradient is the difference between the top and bottom bins. In Panel B, we report a market maker-level version of the Panel A gradient that uses the full return-correlation profile rather than only its endpoints. For each market maker and inventory-correlation group (Negative, Near-zero, Positive), we take her share quoted to sell in each of the five 30-minute return-correlation bins and fit a weighted least-squares line of that share on the bin midpoint (-0.2, -0.05, 0.05, 0.20, 0.40). The slope is the gradient in the market maker's share quoted to sell per unit of return correlation between the focal stock and the hedger; it is defined for a market maker whenever she has pair-episodes in at least three of the five bins. We report, across the market makers for whom it is defined, the mean and median per-market-maker slope, an exact two-sided t-test against zero, and the Wilcoxon signed-rank test, as in the paired test of Table 2. The sample period is 3 January to 6 March 2006.

Panel A. Cross-hedging and inventory

	Inventory correlation			
	Negative	Near-zero	Positive	All
Return correlation				
Below -0.1	0.603	0.566	0.624	0.583
-0.1 to 0	0.644	0.602	0.642	0.623
0 to 0.1	0.667	0.624	0.659	0.647
0.1 to 0.3	0.714	0.673	0.698	0.694
Above 0.3	0.743	0.699	0.730	0.728
Gradient	+0.140	+0.133	+0.106	+0.145
Pair-episodes	10,277,451	3,636,413	10,707,063	27,628,075

Panel B. Gradient statistics

Mean	0.136	0.133	0.075
Median	0.123	0.109	0.085
t-stat.	7.0	7.2	4.9
p-val.	(0.00)	(0.00)	(0.00)
N	389	377	392

4 Conclusion

We develop a model of market making in many correlated assets. A risk-averse dealer sets quotes in every asset through a pricing matrix λ, which maps her inventory in one

asset into the quotes she posts in all of them. The diagonal of λ is the familiar own-inventory price pressure. The off-diagonal is cross-hedging: long one asset, the dealer adjusts her quotes to sell the assets correlated with it. The off-diagonal is positive and rises with fundamental correlation, the diagonal falls as closer substitutes become available, and a more risk-averse dealer holds a more diagonal matrix and hedges less across assets. Reading risk aversion from a single-asset inventory half-life, as if the portfolio were one asset at a time, ignores this structure and understates risk aversion: with N perfectly correlated assets the estimate recovers only a fraction $1/N$ of the true coefficient, and with imperfect correlation the bias converges to the fundamental correlation itself as the number of assets grows.

We take the central prediction to message-level data from the Toronto Stock Exchange. Market makers manage inventory within a stock but rarely clear it, carrying the position across days. Their portfolios are large and two-sided, and the net position is worked through zero while each stock is held, so the stocks offset one another across the portfolio. In identified episodes the mechanism is direct: a market maker accumulates a long position in a stock and shifts quoted liquidity toward the ask side in a few correlated stocks, selling them against the long. The pattern is widespread, present for 537 of 549 market makers, yet modest in intensity, covering about a tenth of the trading day and a minority of the inventory covariance even within episodes, and concentrated, a portfolio-wide imbalance carried by two or three correlated stocks: the profile the theory assigns to risk-averse dealers trading many correlated assets. The direction is the identification. Long the focal stock while quoting to sell its correlated relatives is the opposite of what a common factor produces, and it is what the theory places in the off-diagonal of λ. The sell-side liquidity imbalance varies with stock-pair correlation. The share of pairs quoted to sell rises with the return correlation of the two stocks, from 0.58 to 0.73 for the most correlated pairs, the sign the theory gives the off-diagonal.

The theory carries the economics, and the data provide evidence that the mechanism it describes is present and operates through the quote. The comparative statics of λ, the fall of its diagonal as correlation rises and the attenuation of cross-hedging with risk aversion, are sharper predictions that a single cross-section cannot identify, since correlation and risk aversion are chosen by the market maker or unobserved. Testing

each at the standard the theory sets requires its own source of exogenous variation, which we leave for future research. In the meantime, the portfolio dimension already matters for measurement: single-asset estimates understate dealer risk aversion, the more so the larger and the more correlated the portfolio, and market makers are therefore more risk-averse than existing estimates suggest.

REFERENCES

Amihud, Y. and H. Mendelson (1980). "Dealership market: Market-making with inventory." *Journal of Financial Economics*, 8, 31–53.

Anand, A. and K. Venkataraman (2016). "Market conditions, fragility, and the economics of market making." *Journal of Financial Economics*, 121(2), 327–349.

Brogaard, J., T. Hendershott, and R. Riordan (2014). "High-frequency trading and price discovery." *Review of Financial Studies*, 27, 2267–2306.

Comerton-Forde, C., T. Hendershott, C. M. Jones, P. C. Moulton, and M. S. Seasholes (2010). "Time variation in liquidity: The role of market-maker inventories and revenues." *The Journal of Finance*, 65(1), 295–331.

Copeland, T. E. and D. Galai (1983). "Information effects on the bid-ask spread." *Journal of Finance*, 38, 1457–1469.

Daures-Lescourret, L. and S. Moinas (2023). "Fragmentation and strategic market-making." *Journal of Financial and Quantitative Analysis*, 58(4), 1675–1700.

de Jong, F. and T. Nijman (1997). "High frequency analysis of lead-lag relationships between financial markets." *Journal of Empirical Finance*, 4, 259–277.

Di Maggio, M., A. Kermani, and Z. Song (2017). "The value of trading relationships in turbulent times." *Journal of Financial Economics*, 124, 468–493.

Epps, T. W. (1979). "Comovements in stock prices in the very short run." *Journal of the American Statistical Association*, 74(366a), 291–298.

Hansch, O., N. Y. Naik, and S. Viswanathan (1998). "Do inventories matter in dealership markets? Evidence from the London Stock Exchange." *Journal of Finance*, 53, 1623–1656.

Hendershott, T. and A. J. Menkveld (2014). "Price pressures." *Journal of Financial Economics*, 114(3), 405–423.

Ho, T. and H. R. Stoll (1981). "Optimal dealer pricing under transactions and return uncertainty." *Journal of Financial Economics*, 9, 47–73.

Ho, T. S. Y. and H. R. Stoll (1983). "The dynamics of dealer markets under competition." *Journal of Finance*, 38, 1053–1074.

Hugonnier, J., B. Lester, and P.-O. Weill (2020). "Frictional intermediation in over-the-counter markets." *Review of Economic Studies*, 87, 1750–1796.

Ljungqvist, L. and T. J. Sargent (2000). *Recursive Macroeconomic Theory*. MIT Press, Cambridge, MA.

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

Naik, N. Y. and P. K. Yadav (2003). "Do dealer firms manage inventory on a stock-by-stock or a portfolio basis?" *Journal of Financial Economics*, 69, 325–353.

Roll, R. (1984). "A simple implicit measure of the effective bid-ask spread in an efficient market." *Journal of Finance*, 39, 1127–1139.

Roşu, I., E. Sojli, and W. W. Tham (2021). "Quoting activity and the cost of capital." *Journal of Financial and Quantitative Analysis*, 56(8), 2764–2799.

Scholes, M. and J. Williams (1977). "Estimating betas from nonsynchronous data." *Journal of Financial Economics*, 5, 309–327.

Stoll, H. R. (1978). "The supply of dealer services in securities markets." *Journal of Finance*, 33, 1133–1151.

Virtu Financial, Inc. (2014). "Form S-1 registration statement." U.S. Securities and Exchange Commission. Available at `https://ir.virtu.com/node/9491/html`.

Appendix A. Proofs

Proof of Proposition 1. Denote by $\mathcal{I}_t$ the dealer's information just before trading at t and by E_t the expectation operator conditional on $\mathcal{I}_t$. The dealer's strategy is a set of quotes (a_t, b_t) adapted to the filtration $\mathcal{I}_t$ for $t \in \mathbb{N}$, where a_t and b_t are $N \times 1$ vectors. Equivalently, the dealer chooses (h_t, δ_t), where h_t is the (quoted) half-spread and δ_t is the price discount:

$$h_t \;=\; \frac{a_t - b_t}{2}, \quad \text{and} \quad \delta_t \;=\; v_t - \frac{a_t + b_t}{2}. \tag{A1}$$

We show that $h_t = h$ and $\delta_t = \lambda x_t$, where h and λ are defined in equation (5).

Equation (3) implies that traders' buy and sell demands at t are given, respectively, by $Q_t^b = (k/2)(v_t - a_t) + \ell + \varepsilon_t^b$ and $Q_t^s = (k/2)(b_t - v_t) + \ell + \varepsilon_t^s$, with $\varepsilon_t^b, \varepsilon_t^s \sim \mathcal{N}(0, \Sigma_L/2)$. Define η_t to be the net demand noise at t:

$$\eta_t \;=\; \varepsilon_t^s - \varepsilon_t^b, \tag{A2}$$

which is uncorrelated with the past information and has a multivariate normal distribution $\mathcal{N}(0, \Sigma_L)$. If x_t is the dealer's inventory before trading at t, equation (1) implies that x_t evolves according to $x_{t+1} = x_t - Q_t^b + Q_t^s$, which translates into:

$$x_{t+1} \;=\; x_t - k\delta_t + \eta_t \quad \text{with} \quad \eta_t \overset{IID}{\sim} \mathcal{N}\big(0, \Sigma_L\big). \tag{A3}$$

Substituting Q_t^b and Q_t^s in the dealer's objective function from (2), we get:

$$\mathsf{E} \sum_{t=0}^{\infty} \beta^t \, \mathsf{E}_t \Big(-(a_t - v_t)' \frac{k}{2} (a_t - v_t) - (v_t - b_t)' \frac{k}{2} (v_t - b_t) + 2\ell' h_t - \gamma x_t' \Sigma_v x_t \Big). \tag{A4}$$

With the notation in equation (A1), we compute $\mathsf{E}_t (v_t - a_t)' k (v_t - a_t) = (\delta_t - h_t)' k (\delta_t - h_t)$ and $\mathsf{E}_t (v_t - b_t)' k (v_t - b_t) = (\delta_t + h_t)' k (\delta_t + h_t)$. Thus, the dealer solves:

$$\max_{(h_t, \delta_t)_{t \geq 0}} \; \mathsf{E} \sum_{t=0}^{\infty} \beta^t \Big(-\delta_t' k \delta_t - h_t' k h_t + 2\ell' h_t - \gamma x_t' \Sigma_v x_t \Big), \tag{A5}$$

where the inventory x_t evolves according to equation (A3). Using the Bellman principle of optimization, we reduce the dynamic optimization in (A5) to finding a value function

V that solves the static optimization problem:

$$V(x_t) \;=\; \max_{h_t,\delta_t} \Big(-\delta_t' k \delta_t - h_t' k h_t + 2\ell' h_t - \gamma x_t' \Sigma_v x_t + \beta\, \mathsf{E}_t\, V(x_{t+1})\Big), \tag{A6}$$

We guess that $V : \mathbb{R}^N \to \mathbb{R}^N$ is a quadratic function of the form:

$$V(x) \;=\; W_0 - x' W_2 x, \tag{A7}$$

where W_0 is a scalar and W_2 is a symmetric $N \times N$ matrix. If A and B are two $N \times N$ matrices, denote their Frobenius inner product (or double dot product) by $A \circ B = \sum_{i,j=1}^N A_{ij} B_{ij}$. A direct calculation shows that $\mathsf{E}_t\big(\eta_t' W_2 \eta_t\big) = W_2 \circ \Sigma_L$. Substituting x_{t+1} from equation (A3), the problem becomes:

$$\begin{aligned} V(x_t) \;=\; \max_{h_t,\delta_t} \Big(& -\delta_t' k \delta_t - h_t' k h_t + 2\ell' h_t - \gamma x_t' \Sigma_v x_t + \beta W_0 \\ & - \beta x_t' W_2 x_t - \beta \delta_t' k W_2 k \delta_t + 2\beta \delta_t' k W_2 x_t - \beta W_2 \circ \Sigma_L \Big). \end{aligned} \tag{A8}$$

The first order condition with respect to h_t implies that the equilibrium half-spread satisfies $h_t = k^{-1}\ell$, which is the same as the constant h in the first part of equation (5).

The first order condition with respect to δ_t is $-2k\delta_t - 2\beta k W_2 k \delta_t + 2\beta k W_2 x_t \;=\; 0$, which implies that the equilibrium price discount is:

$$\delta_t \;=\; \lambda x_t, \quad \text{with} \quad \lambda \;=\; (k + \beta k W_2 k)^{-1} \beta k W_2. \tag{A9}$$

The second order condition with respect to δ_t is that the symmetric matrix $I_N + \beta W_2 k$ is positive definite, which we denote with an inequality sign:

$$I_N + \beta W_2 k \;>\; 0. \tag{A10}$$

Combining equations (A7) and (A8) and identifying the quadratic coefficients (of the form $x_t'(...)x_t$), it follows that W_2 satisfies:

$$W_2 \;=\; \gamma \Sigma_v + \beta W_2 - \beta^2 W_2 k (k + \beta k W_2 k)^{-1} k W_2, \tag{A11}$$

or, equivalently, after multiplying by k on both sides:

$$\gamma k\Sigma_v - \frac{1-\beta}{\beta} w_2 \;=\; w_2(I_N + w_2)^{-1} w_2, \quad \text{with} \quad w_2 \;=\; \beta k W_2. \tag{A12}$$

Denote by:

$$Z \;=\; (I_N + w_2)^{-1} w_2. \tag{A13}$$

Simple matrix manipulation shows that $w_2 = (I_N - Z)^{-1} Z$. Multiplying equation (A12) by $(I_N - Z)$ on both sides, we obtain:

$$(I_N - Z)S - \frac{1-\beta}{\beta} Z \;=\; Z^2, \quad \text{with} \quad S \;=\; \gamma k \Sigma_v. \tag{A14}$$

This is equivalent to $Z^2 + Z\big(S + \frac{1-\beta}{\beta} I_N\big) = S$, which is the same equation as (6). Following the existence and uniqueness of the square root of a symmetric positive definite matrix, it is straightforward to show that this equation has a unique solution in symmetric positive definite matrices. Thus, $Z = \Lambda$. Moreover, equations (A9) and (A13) imply that:

$$\lambda \;=\; (k + \beta k W_2 k)^{-1} \beta k W_2 \;=\; k^{-1}\Lambda. \tag{A15}$$

This proves that $\lambda = k^{-1}\Lambda$, which is the second part of equation (5).

For future use, we identify the constant terms (independent of x_t) in equation (A8), and we obtain:

$$W_0 \;=\; \frac{1}{1-\beta}\Big(\ell' k^{-1} \ell - \beta W_2 \circ \Sigma_L\Big). \tag{A16}$$

When β converges to 1, equation (6) becomes $\Lambda^2 + \Lambda S = S$. This equation has a unique symmetric positive definite solution:

$$\Lambda \;=\; \frac{1}{2}\Big((S^2 + 4S)^{1/2} - S\Big). \tag{A17}$$

This proves equation (7).

In the same limit $\beta \to 1$, setting $\alpha \;=\; W_2$ in equation (A9) implies $\gamma\Sigma_v \;=\; \alpha k (I + \alpha k)^{-1}\alpha$, which using the identity $k(I + \alpha k)^{-1} \;=\; (k^{-1} + \alpha)^{-1}$ is equivalent to

$$(k^{-1} + \alpha)^{-1} \;=\; \gamma \alpha^{-1} \Sigma_v \alpha^{-1}. \tag{A18}$$

□

Proposition A.1. *Suppose the dealer's objective in equation (2) is modified by replacing the inventory penalty $\gamma x_t' \Sigma_v x_t$ with $\gamma x_{t+1}' \Sigma_v x_{t+1}$. Then, for any fixed $\beta \in (0,1)$, the dealer's dynamic optimization problem is equivalent to the original problem in Proposition 1, with the same primitives except that the risk-aversion coefficient γ is replaced by γ/β. Consequently, the equilibrium quotes remain of the form*

$$a_t = v_t - \lambda x_t + h, \quad b_t = v_t - \lambda x_t - h, \quad \textit{with} \quad h = k^{-1}\ell, \quad \lambda = k^{-1}\Lambda, \tag{A19}$$

where Λ is the unique symmetric positive definite solution of

$$\Lambda^2 + \Lambda\Big(\frac{\gamma}{\beta} k\Sigma_v + \frac{1-\beta}{\beta} I_N\Big) \;=\; \frac{\gamma}{\beta} k\Sigma_v. \tag{A20}$$

In particular, when $\beta \to 1$, the limiting equilibrium is identical to the one in Proposition 1, namely

$$\lambda \;=\; \frac{k^{-1}}{2}\Big(\big(S^2 + 4S\big)^{1/2} - S\Big), \quad \textit{with} \quad S \;=\; \gamma k \Sigma_v. \tag{A21}$$

Proof. Let $\Pi_t = (v_t - b_t)' Q_t^s + (a_t - v_t)' Q_t^b$ denote the dealer's period-t expected profit term. Consider the modified objective

$$\mathsf{E}_\tau \sum_{t=\tau}^{\infty} \beta^{\,t-\tau} \Big(\Pi_t - \gamma x_{t+1}' \Sigma_v x_{t+1}\Big). \tag{A22}$$

Fix the state x_τ. Re-indexing the penalty term yields

$$\sum_{t=\tau}^{\infty} \beta^{\,t-\tau} x_{t+1}' \Sigma_v x_{t+1} = \sum_{s=\tau+1}^{\infty} \beta^{\,s-1-\tau} x_s' \Sigma_v x_s = \frac{1}{\beta} \sum_{s=\tau+1}^{\infty} \beta^{\,s-\tau} x_s' \Sigma_v x_s. \tag{A23}$$

Hence

$$\sum_{t=\tau}^{\infty} \beta^{\,t-\tau} x_{t+1}' \Sigma_v x_{t+1} = \frac{1}{\beta} \sum_{s=\tau}^{\infty} \beta^{\,s-\tau} x_s' \Sigma_v x_s - \frac{1}{\beta} x_\tau' \Sigma_v x_\tau. \tag{A24}$$

Substituting into (A22), we obtain

$$\mathsf{E}_\tau \sum_{t=\tau}^{\infty} \beta^{t-\tau}\Big(\Pi_t - \gamma x'_{t+1}\Sigma_v x_{t+1}\Big) = \mathsf{E}_\tau \sum_{t=\tau}^{\infty} \beta^{t-\tau}\Big(\Pi_t - \frac{\gamma}{\beta} x'_t\Sigma_v x_t\Big) + \frac{\gamma}{\beta} x'_\tau \Sigma_v x_\tau. \tag{A25}$$

The last term depends only on the current state x_τ and not on the dealer's choice variables at dates $t \geq \tau$. Therefore, for every fixed $\beta \in (0,1)$, maximizing (A22) is equivalent to maximizing the original objective (2) with γ replaced by γ/β.

Proposition 1 then applies directly, with $S = \frac{\gamma}{\beta} k\Sigma_v$. Thus, the optimal quotes are again linear and we get the same formulas as in equations (4) and (5): $a_t = v_t - \lambda x_t + h$, $b_t = v_t - \lambda x_t - h$, $h = k^{-1}\ell$, and $\lambda = k^{-1}\Lambda$, where Λ solves equation (A20).

Finally, when $\beta \to 1$, we have $(\gamma/\beta)k\Sigma_v \to \gamma k\Sigma_v$, so the limiting matrix equation becomes

$$\Lambda^2 + \Lambda S = S, \qquad S = \gamma k \Sigma_v, \tag{A26}$$

which is exactly the limit equation in Proposition 1. Therefore the limiting equilibrium when $\beta \to 1$ is the same under the two specifications of the inventory penalty. □

Proof of Corollary 1. First, assume that $\big|S^{-1}\big| < 1/4$. The Taylor series for $(1+x)^{1/2}$ is $1 + x/2 - x^2/8 + \cdots$ and converges when $|x| < 1$. A similar analysis for matrices shows that the Taylor series for $(1 + 4S^{-1})^{1/2}$ is $1 + 2S^{-1} - 2S^{-2} + \cdots$ and converges when $|S^{-1}| < 1/4$. Equation (7) implies that the matrix $\Lambda = k\lambda$ satisfies:

$$\Lambda \;=\; \frac{S}{2}\Big(\big(I_N + 4S^{-1}\big)^{1/2} - I_N\Big) \;=\; \frac{S}{2}\Big(2S^{-1} - 2S^{-2} + S^{-1}O\big(S^{-2}\big)\Big). \tag{A27}$$

Thus, $\Lambda = I_N - S^{-1} + O\big(S^{-2}\big) = I_N + O\big(S^{-1}\big)$, which proves equation (9).

Now, assume that $|S| < 4$. Rewriting equation (7), we obtain that:

$$\begin{aligned} 2\Lambda \;&=\; \Big(4S\Big(1 + \frac{S}{4}\Big)\Big)^{1/2} - S \;=\; 2S^{1/2}\Big(1 + \frac{S}{8} + O\big(S^2\big)\Big) - S \\ &=\; 2S^{1/2}\Big(1 - S^{1/2} + \frac{S}{8} + O\big(S^2\big)\Big). \end{aligned} \tag{A28}$$

Thus, $\Lambda = S^{1/2} + O(S)$, which proves equation (10). □

Proof of Corollary 2. If we denote by

$$P \;=\; \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \tag{A29}$$

then P diagonalizes the matrices Σ_v and $k = k_0 I_2$:

$$P'\Sigma_v P \;=\; \sigma_v^2 \,\mathrm{diag}(1+a, 1-a), \qquad P'kP \;=\; k_0 I_2. \tag{A30}$$

If we denote the eigenvalues of $S = \gamma k \Sigma_v$ by

$$s_+ \;=\; \gamma k_0 \sigma_v^2 (1+a), \qquad s_- \;=\; \gamma k_0 \sigma_v^2 (1-a), \tag{A31}$$

equation (A30) implies that

$$P'SP \;=\; \gamma k_0 \sigma_v^2 \,\mathrm{diag}(1+a, 1-a) \;=\; \mathrm{diag}(s_+, s_-). \tag{A32}$$

The matrix $\Lambda = k\lambda = \big((S^2 + 4S)^{1/2} - S\big)/2$ has eigenvalues

$$\mu_+ \;=\; \frac{1}{2}\big(\sqrt{s_+^2 + 4s_+} - s_+\big) \quad \text{and} \quad \mu_- \;=\; \frac{1}{2}\big(\sqrt{s_-^2 + 4s_-} - s_-\big), \tag{A33}$$

and is also diagonalized by P. Using the identity

$$\frac{1}{2}\big(\sqrt{s^2 + 4s} - s\big) \;=\; \frac{2}{1 + \sqrt{1 + 4/s}}, \tag{A34}$$

we obtain

$$\mu_+ \;=\; \tfrac{2}{m_2+1}, \quad \mu_- \;=\; \tfrac{2}{m_1+1}, \quad \text{with} \quad m_1 \;=\; \Big(1 + \tfrac{4}{s_-}\Big)^{1/2}, \quad m_2 \;=\; \Big(1 + \tfrac{4}{s_+}\Big)^{1/2}, \tag{A35}$$

which are the same formulas for m_1 and m_2 as in equation (12). Therefore,

$$P'\Lambda P \;=\; \mathrm{diag}\Big(\tfrac{2}{m_2+1}, \tfrac{2}{m_1+1}\Big), \tag{A36}$$

and rotating back gives

$$\Lambda \;=\; P\,\mathrm{diag}\Big(\tfrac{2}{m_2+1},\tfrac{2}{m_1+1}\Big)\,P' \;=\; \tfrac{1}{(m_1+1)(m_2+1)}\begin{bmatrix} m_1+m_2+2 & m_1-m_2 \\ m_1-m_2 & m_1+m_2+2 \end{bmatrix}. \tag{A37}$$

Since $\lambda = k^{-1}\Lambda = k_0^{-1}\Lambda$, this finishes the proof of equation (12).

It remains to establish the comparative statics of $\lambda_{1,2}$. Recall that

$$k_0\,\lambda_{1,2} \;=\; \frac{m_1-m_2}{(m_1+1)(m_2+1)} \;=\; \frac{1}{m_2+1}-\frac{1}{m_1+1}, \tag{A38}$$

where m_1 and m_2 satisfy

$$m_1^2 \;=\; 1+\frac{\xi}{1-a} \quad\text{and}\quad m_2^2 \;=\; 1+\frac{\xi}{1+a}, \quad\text{with}\quad \xi \;=\; 4\gamma^{-1}k_0^{-1}\sigma_v^{-2}. \tag{A39}$$

Let $a \in (-1,1)$. We prove that $\lambda_{1,2}$ is strictly increasing in a. Differentiating m_1^2 and m_2^2 with respect to a for fixed ξ, we get

$$\frac{\partial m_1}{\partial a} \;=\; \frac{\xi}{2(1-a)^2 m_1} > 0, \qquad \frac{\partial m_2}{\partial a} \;=\; -\frac{\xi}{2(1+a)^2 m_2} < 0. \tag{A40}$$

Differentiating equation (A38) with respect to a, we calculate

$$k_0\,\tfrac{\partial \lambda_{1,2}}{\partial a} \;=\; \tfrac{\partial m_1/\partial a}{(m_1+1)^2} - \tfrac{\partial m_2/\partial a}{(m_2+1)^2} \;=\; \tfrac{\xi}{2(1-a)^2 m_1(m_1+1)^2} + \tfrac{\xi}{2(1+a)^2 m_2(m_2+1)^2} \;>\; 0, \tag{A41}$$

which proves that $\lambda_{1,2}$ is strictly increasing in a. Moreover, when $a = 0$, we have $m_1 = m_2$ and $\lambda_{1,2} = 0$.

Let $a \in (0,1)$. We prove that $\lambda_{1,2}$ has an inverted U-shape in $\xi = 4\gamma^{-1}k_0^{-1}\sigma_v^{-2}$, with zero limits at the extremes. Since $a > 0$, we have $m_1 > m_2 \geq 1$, and $\lambda_{1,2} > 0$. First, we show that $\lambda_{1,2}$ has zero limits when ξ converges to 0 or ∞. Indeed, when $\xi \to 0$ (i.e., $\gamma \to \infty$), $m_1, m_2 \to 1$ and $\lambda_{1,2} \to 0$. When $\xi \to \infty$ (i.e., $\gamma \to 0^+$), for $i = 1, 2$, one sees that $m_i \to \infty$, so each $1/(m_i+1) \to 0$, and thus $\lambda_{1,2} \to 0$. Next, we prove that $\lambda_{1,2}$ has an inverted U-shape in ξ. Differentiating $m_i^2 = 1+\xi/(1\mp a)$ in ξ, for $i = 1, 2$, we obtain

$$m_1' \;=\; \frac{\partial m_1}{\partial \xi} \;=\; \frac{1}{2(1-a)m_1}, \qquad m_2' \;=\; \frac{\partial m_2}{\partial \xi} \;=\; \frac{1}{2(1+a)m_2}, \tag{A42}$$

Differentiating equation (A38) with respect to ξ, we calculate

$$k_0 \frac{\partial \lambda_{1,2}}{\partial \xi} \;=\; \frac{m_1'}{(m_1+1)^2} - \frac{m_2'}{(m_2+1)^2} \;=\; \frac{G(\xi)}{2(1-a)(1+a)\, m_1(m_1+1)^2\, m_2(m_2+1)^2}, \tag{A43}$$

where

$$G(\xi) \;=\; (1+a)\, m_2(m_2+1)^2 - (1-a)\, m_1(m_1+1)^2. \tag{A44}$$

When $\xi = 0$, we have $m_1 = m_2 = 1$ and $G(0) = 8a > 0$. A direct computation gives

$$\frac{\partial G}{\partial \xi} \;=\; \frac{h(m_2) - h(m_1)}{2}, \quad \text{with} \quad h(m) \;=\; 3m + 4 + \frac{1}{m}, \tag{A45}$$

which is strictly negative because h is strictly increasing on $[1, \infty)$ and $m_1 > m_2 \geq 1$. Hence G is strictly decreasing on $(0, \infty)$, so it has at most one zero. If G had no zero, then $G > 0$ throughout and $\lambda_{1,2}$ would be strictly increasing on $(0, \infty)$, contradicting $\lambda_{1,2} \to 0$ as $\xi \to \infty$. Therefore G has a unique zero ξ^*, with $G > 0$ on $(0, \xi^*)$ and $G < 0$ on (ξ^*, ∞). Hence $\lambda_{1,2}$ has a unique interior maximum, attained at $\gamma^* \;=\; 4(\xi^* k_0 \sigma_v^2)^{-1}$. $\square$

Proof of Corollary 3. Let P be defined as in equation (A29). By the proof of Corollary 2, $P' \Sigma_v P \;=\; \sigma_v^2 \operatorname{diag}(1+a, 1-a)$ and $P' k P \;=\; k_0 I_2$. Similarly, we also have

$$P' \Sigma_L P \;=\; L_0 \operatorname{diag}(1+b, 1-b). \tag{A46}$$

Let $A \;=\; \gamma^{-1} \Sigma_v^{-1} k^{-1}$. If we denote by

$$a_+ \;=\; \frac{1}{\gamma k_0 \sigma_v^2 (1+a)}, \qquad a_- \;=\; \frac{1}{\gamma k_0 \sigma_v^2 (1-a)}, \tag{A47}$$

the matrix $A \;=\; \gamma^{-1} \Sigma_v^{-1} k^{-1}$ is also diagonalized by P:

$$P' A P \;=\; \operatorname{diag}(a_+, a_-)\,. \tag{A48}$$

Equation (18) implies that for every $x \in \mathbb{R}^2$,

$$\lim_{\beta \to 1} (1-\beta) V(x) \;=\; \ell' k^{-1} \ell - \alpha \circ \Sigma_L. \tag{A49}$$

From the proof of Proposition 1, $\alpha = \lim_{\beta\to 1} W_2$ satisfies equation (A18): $(k^{-1}+\alpha)^{-1} = \gamma\alpha^{-1}\Sigma_v\alpha^{-1}$. Taking the inverse of both sides, we get $k^{-1}+\alpha = \gamma^{-1}\alpha\Sigma_v^{-1}\alpha$. Multiplying on the left by $\gamma^{-1}\Sigma_v^{-1}$ and setting $X = \gamma^{-1}\Sigma_v^{-1}\alpha$, we get the quadratic equation $A + X = X^2$. Since α is positive definite, so is X, hence we get

$$X = \tfrac{1}{2}\big(I_2 + (I_2 + 4A)^{1/2}\big). \tag{A50}$$

The matrix X is also diagonalized by P, with eigenvalues

$$x_+ = \tfrac{1}{2}\big(1 + (1+4a_+)^{1/2}\big) = \tfrac{1+m_2}{2}, \quad x_- = \tfrac{1}{2}\big(1 + (1+4a_-)^{1/2}\big) = \tfrac{1+m_1}{2}, \tag{A51}$$

where the second equalities use the definitions of m_1 and m_2 in Corollary 2. Since $\alpha = \gamma\Sigma_v X$, the matrix α is also diagonalized by P, with

$$\begin{aligned} P'\alpha P &= \tfrac{\gamma\sigma_v^2}{2}\operatorname{diag}\big((1+a)(1+m_2),\,(1-a)(1+m_1)\big) \\ &= \tfrac{\gamma\sigma_v^2}{2}\operatorname{diag}\big((1+a)+n_2,\,(1-a)+n_1\big), \end{aligned} \tag{A52}$$

where the second equality uses $n_1 = m_1(1-a)$ and $n_2 = m_2(1+a)$. Since the Frobenius product is invariant under orthogonal changes of basis,

$$\begin{aligned} \alpha \circ \Sigma_L &= (P'\alpha P) \circ (P'\Sigma_L P) \\ &= \tfrac{\gamma\sigma_v^2 L_0}{2}\Big(\big((1+a)+n_2\big)(1+b) + \big((1-a)+n_1\big)(1-b)\Big) \\ &= \tfrac{\gamma\sigma_v^2 L_0}{2}\big(n_1(1-b) + n_2(1+b) + 2(1+ab)\big). \end{aligned} \tag{A53}$$

Substituting this identity into equation (A49) proves equation (20).

For asymptotics, from Corollary 2 recall that, for $i = 1, 2$,

$$m_i = \sqrt{1 + \tfrac{4}{\gamma k_0 \sigma_v^2(1\mp a)}} = \tfrac{2}{\gamma^{1/2}k_0^{1/2}\sigma_v(1\mp a)^{1/2}}\sqrt{1 + \tfrac{\gamma k_0\sigma_v^2(1\mp a)}{4}}. \tag{A54}$$

When $\gamma \to 0$, $m_i^2 \to \infty$, and factoring out the dominant term gives

$$m_i = 2\gamma^{-\frac{1}{2}}k_0^{-\frac{1}{2}}\sigma_v^{-1}(1\mp a)^{-\frac{1}{2}} + O(\gamma^{\frac{1}{2}}). \tag{A55}$$

Multiplying by $(1 \mp a)$ yields

$$n_1 \;=\; 2\gamma^{-\frac{1}{2}} k_0^{-\frac{1}{2}} \sigma_v^{-1} (1-a)^{\frac{1}{2}} + O(\gamma^{\frac{1}{2}}), \quad n_2 \;=\; 2\gamma^{-\frac{1}{2}} k_0^{-\frac{1}{2}} \sigma_v^{-1} (1+a)^{\frac{1}{2}} + O(\gamma^{\frac{1}{2}}), \quad \text{(A56)}$$

which substituted into equation (20) yields the first line of equation (21).

When $\gamma \to \infty$, $m_i^2 \to 1$, and using the Taylor expansion for the square root gives

$$m_i \;=\; \sqrt{1 + \tfrac{4}{\gamma k_0 \sigma_v^2 (1 \mp a)}} \;=\; 1 + \frac{2}{\gamma k_0 \sigma_v^2 (1 \mp a)} + O(\gamma^{-2}). \quad \text{(A57)}$$

Multiplying by $(1 \mp a)$ yields

$$n_1 \;=\; (1-a) + \frac{2}{\gamma k_0 \sigma_v^2} + O(\gamma^{-2}), \quad n_2 \;=\; (1+a) + \frac{2}{\gamma k_0 \sigma_v^2} + O(\gamma^{-2}), \quad \text{(A58)}$$

which substituted into equation (20) yields the second line of equation (21).

Finally, we derive the comparative statics for $b > 0$. In the small-γ regime, equation (21) gives $W(a) \;=\; C_0 - \gamma^{1/2} D_0 \, f(a,b) + O(\gamma^{3/2})$, where C_0 and $D_0 > 0$ are independent of a, and $f(a,b) \;=\; (1-a)^{1/2}(1-b) + (1+a)^{1/2}(1+b)$. Direct computation gives

$$f''(a) \;=\; -\tfrac{1}{4}(1-a)^{-3/2}(1-b) - \tfrac{1}{4}(1+a)^{-3/2}(1+b) < 0, \quad \text{(A59)}$$

so f is strictly concave on $(-1,1)$ and W is convex (U-shaped) with local maxima at $a = \pm 1$. Since $f(-1,b) - f(+1,b) \;=\; -2\sqrt{2}\, b < 0$ when $b > 0$, $W(-1) > W(+1)$. On $(0,1]$, setting $f'(a) \;=\; 0$ gives the unique interior critical point $a^* \;=\; 2b/(1+b^2)$, which is a local maximum of f and hence a local minimum of W; the maximum of W on $(0,1]$ is therefore at an endpoint. Comparing $f(0,b) \;=\; 2$ and $f(1,b) \;=\; \sqrt{2}(1+b)$ gives the threshold $b^* \;=\; \sqrt{2} - 1$ at which the two endpoint values coincide.

In the large-γ regime, equation (21) gives $\partial W / \partial a \;=\; -2\gamma L_0 \sigma_v^2 \, b + O(\gamma^{-1}) < 0$ for γ sufficiently large, so W is strictly decreasing in a. □

Proof of Corollary 4. Let P be defined as in equation (A29). By the proof of Corollary 2, the matrix $k_0 \lambda$ is diagonalized by P and satisfies

$$P'(k_0 \lambda) P = \operatorname{diag}\left(\frac{2}{m_2 + 1}, \frac{2}{m_1 + 1} \right). \quad \text{(A60)}$$

Therefore,

$$P'(I_2 - k_0\lambda)P = \operatorname{diag}\left(\frac{m_2-1}{m_2+1}, \frac{m_1-1}{m_1+1}\right). \tag{A61}$$

Taking the t-th power and rotating back, we obtain

$$\begin{aligned}(I_2 - k_0\lambda)^t &= P \operatorname{diag}\left(\left(\frac{m_2-1}{m_2+1}\right)^t, \left(\frac{m_1-1}{m_1+1}\right)^t\right) P' \\ &= \begin{bmatrix} \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t + \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t & \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t - \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t \\ \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t - \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t & \frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^t + \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^t \end{bmatrix},\end{aligned}$$

which proves the first displayed formula.

If $x_0 = [1, 0]'$, then the expected inventory after t periods is $(I_2 - k_0\lambda)^t x_0$, so the expected inventory in asset 1 is reduced by half exactly when

$$\frac{1}{2}\left(\frac{m_2-1}{m_2+1}\right)^T + \frac{1}{2}\left(\frac{m_1-1}{m_1+1}\right)^T = \frac{1}{2}, \tag{A62}$$

which is equivalent to

$$\left(\frac{m_1-1}{m_1+1}\right)^T + \left(\frac{m_2-1}{m_2+1}\right)^T = 1. \tag{A63}$$

This proves the equation that characterizes the inventory half-life.

Finally, suppose that γ is small and let $m = \gamma^{-1/2} k_0^{-1/2} \sigma_v^{-1}$. Then

$$m_1 = 2m(1-a)^{-1/2} + O(m^{-1}), \qquad m_2 = 2m(1+a)^{-1/2} + O(m^{-1}). \tag{A64}$$

Moreover,

$$\log\left(\frac{m_i-1}{m_i+1}\right) = -\frac{2}{m_i} + O(m_i^{-3}), \qquad i = 1, 2. \tag{A65}$$

Hence the half-life equation becomes

$$\exp\left(-\frac{2T}{m_1} + O\left(\frac{T}{m_1^3}\right)\right) + \exp\left(-\frac{2T}{m_2} + O\left(\frac{T}{m_2^3}\right)\right) = 1. \tag{A66}$$

Keeping the first-order terms in $1/m_1$ and $1/m_2$, we obtain

$$1 - \frac{2T}{m_1} + 1 - \frac{2T}{m_2} \approx 1, \tag{A67}$$

so that

$$T \approx \frac{1}{2\left(m_1^{-1}+m_2^{-1}\right)} = \frac{m}{(1+a)^{1/2}+(1-a)^{1/2}}. \tag{A68}$$

This proves equation (25). □

Proof of Corollary 5. The vector $\mathbf{1}_N$ is an eigenvector of Σ_v with eigenvalue $\sigma_v^2\big(1+(N-1)a\big)$, and the orthogonal complement of $\mathbf{1}_N$ is an eigenspace of Σ_v with eigenvalue $\sigma_v^2(1-a)$; the condition $a > -\frac{1}{N-1}$ ensures that Σ_v is positive definite. Let P be an orthogonal N-dimensional matrix whose first column is $\mathbf{1}_N/\sqrt{N}$. Then $P'\Sigma_v P = \sigma_v^2\,\mathrm{diag}\big(1+(N-1)a,\,(1-a)I_{N-1}\big)$ and $P'kP = k_0 I_N$. Since equation (7) expresses λ as a matrix function of $S = \gamma k_0 \Sigma_v$, the matrix $k_0\lambda$ is diagonalized by the same P, and the eigenvalue of $k_0\lambda$ that corresponds to an eigenvalue s of S equals

$$\frac{1}{2}\Big(\big(s^2+4s\big)^{1/2}-s\Big) \;=\; \frac{2}{m+1}, \quad \text{with} \quad m \;=\; \big(4s^{-1}+1\big)^{1/2}, \tag{A69}$$

where the equality follows from $s(m^2-1)=4$, as in the proof of Corollary 2. Applying this to the eigenvalue $s = \gamma k_0 \sigma_v^2\big(1+(N-1)a\big)$ on the span of $\mathbf{1}_N$ and to the eigenvalue $s = \gamma k_0 \sigma_v^2(1-a)$ on its orthogonal complement, we obtain:

$$P'\big(I_N - k_0\lambda\big)P \;=\; \mathrm{diag}\big(\rho_2,\,\rho_1 I_{N-1}\big), \tag{A70}$$

with ρ_1 and ρ_2 as in the statement of the corollary. Taking the t-th power and rotating back,

$$\big(I_N - k_0\lambda\big)^t \;=\; P\,\mathrm{diag}\big(\rho_2^t,\,\rho_1^t I_{N-1}\big)P' \;=\; \rho_1^t\, I_N + \big(\rho_2^t - \rho_1^t\big)\,P\,\mathrm{diag}(1,0,\ldots,0)P' \;=\; \rho_1^t\, I_N + \frac{\rho_2^t-\rho_1^t}{N}\,\mathbf{1}_N\mathbf{1}_N', \tag{A71}$$

where the last equality uses the fact that $P\,\mathrm{diag}(1,0,\ldots,0)P'$ is the outer product of the first column of P with itself, which equals $\frac{1}{N}\mathbf{1}_N\mathbf{1}_N'$. This proves equation (29).

If $x_0 = [1,0,\ldots,0]'$, then, as in the proof of Corollary 4, the expected inventory after t periods is $(I_N - k_0\lambda)^t x_0$, so the expected inventory in asset 1 equals the $(1,1)$ entry of the matrix in equation (29), which is $\frac{1}{N}\rho_2^t + \frac{N-1}{N}\rho_1^t$. Setting this expression equal to $\frac{1}{2}$ proves equation (30).

Finally, suppose that γ is small and let $m = \gamma^{-1/2} k_0^{-1/2} \sigma_v^{-1}$. Then

$$m_1 \;=\; 2m(1-a)^{-1/2} + O(m^{-1}), \qquad m_2 \;=\; 2m\big(1+(N-1)a\big)^{-1/2} + O(m^{-1}), \quad \text{(A72)}$$

and, as in the proof of Corollary 4, $\log \rho_i = -\frac{2}{m_i} + O(m_i^{-3})$ for $i = 1, 2$. Keeping the first-order terms in $1/m_1$ and $1/m_2$, the half-life equation (30) becomes

$$\frac{1}{N}\Big(1 - \frac{2T}{m_2}\Big) + \frac{N-1}{N}\Big(1 - \frac{2T}{m_1}\Big) \;\approx\; \frac{1}{2}, \quad \text{(A73)}$$

so that

$$T \;\approx\; \frac{N}{4\big(m_2^{-1} + (N-1)\, m_1^{-1}\big)} \;=\; \frac{N\, m}{2\Big(\big(1+(N-1)a\big)^{1/2} + (N-1)(1-a)^{1/2}\Big)}. \quad \text{(A74)}$$

This proves equation (31). □

Proof of Proposition 2. Proposition 1 implies that $p_t = v_t - \lambda x_t$, hence $v_t - p_t = \lambda x_t$. Also, $Q_t^b - Q_t^s = k(v_t - p_t) + \varepsilon_t = k\lambda x_t + \varepsilon_t$. As the dealer's inventory changes by the amount of the net order flow, $x_{t+1} = x_t - (Q_t^b - Q_t^s) = x_t - k\lambda x_t - \varepsilon_t$, which proves the first part of equation (34).

Denote by $Y = k\lambda$. At $t+1$, the equation $p_t = v_t - \lambda x_t$ becomes $p_{t+1} = v_{t+1} - \lambda x_{t+1}$, hence $\Delta p_{t+1} = \Delta v_{t+1} - \lambda(x_{t+1} - x_t)$. As $x_{t+1} - x_t = -Y x_t - \varepsilon_t$, we obtain $\Delta p_{t+1} = \Delta v_{t+1} + \lambda\varepsilon_t + \lambda Y x_t$, which proves the second part of equation (34).

In a stationary equilibrium, the inventory variance remains constant, i.e., $\mathsf{Var}(x_{t+1}) = \mathsf{Var}(x_t) = \Sigma_x$. By assumption, k and Σ_v commute with each other, hence equation (7) implies that k and λ also commute with each other (they are also both symmetric). The equation $x_{t+1} = (I_N - Y)x_t - \varepsilon_t$ implies that $\Sigma_x = (I_N - Y)\Sigma_x(I_N - Y) + \Sigma_L$. Hence, Σ_x satisfies the linear matrix equation $Y\Sigma_x + \Sigma_x Y - Y\Sigma_x Y = \Sigma_L$. If k and Σ_v also commute with Σ_L, Σ_x satisfies the equation $Y(2I_N - Y)\Sigma_x = \Sigma_L$, hence $\Sigma_x = Y^{-1}(2I_N - Y)^{-1}\Sigma_L$, which proves the first part of equation (35).

The price increment at $t+1$ is $\Delta p_{t+1} = \Delta v_{t+1} + \lambda\varepsilon_t + \lambda Y x_t$, hence $\Sigma_p = \Sigma_v + \lambda\Sigma_L\lambda +$

$\lambda Y \Sigma_x Y \lambda$. Using the formula $Y\Sigma_x Y = Y\Sigma_x + \Sigma_x Y - \Sigma_L$, we compute:

$$\Sigma_p \;=\; \Sigma_v + \lambda(Y\Sigma_x + \Sigma_x Y)\lambda, \quad \text{with} \quad Y \;=\; k\lambda. \tag{A75}$$

If Σ_L commutes with k and Σ_v, equation (A75) becomes $\Sigma_p = \Sigma_v + 2k\lambda^3\Sigma_x = \Sigma_v + 2\lambda^2(2I_N - k\lambda)^{-1}\Sigma_L$, which proves the second part of equation (35).

The price increment at $t+2$ is $\Delta p_{t+2} = \Delta v_{t+2} + \lambda\varepsilon_{t+1} + \lambda Y x_{t+1} = \Delta v_{t+2} + \lambda\varepsilon_{t+1} + \lambda Y(I_N - Y)x_t - \lambda Y\varepsilon_t$. Thus, the autocovariance $A_p = \mathsf{E}\big((\Delta p_{t+2})\,(\Delta p_{t+1})'\big) = \lambda Y(I_N - Y)\Sigma_x Y\lambda - \lambda Y\Sigma_L\lambda = \lambda Y\big((I_N - Y)\Sigma_x Y - \Sigma_L\big)\lambda$. Using the equation $\Sigma_x Y - Y\Sigma_x Y = \Sigma_L - Y\Sigma_x$, we compute $A_p = -\lambda Y^2\Sigma_x\lambda$. If Σ_L commutes with k and Σ_v, we compute $A_p = -k^2\lambda^4\Sigma_x = -k\lambda^3(2I_N - k\lambda)^{-1}\Sigma_L$, which proves the third part of equation (35).

□

Proof of Corollary 6. Let P be defined as in equation (A29). By the proof of Corollary 2, the matrix λ is diagonalized by P and satisfies

$$P'\lambda P = \operatorname{diag}\left(\frac{2}{k_0(m_2+1)}, \frac{2}{k_0(m_1+1)}\right). \tag{A76}$$

Since $\Sigma_L = L_0 I_2$ and $\Sigma_v = \sigma_v^2 \begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}$, we also have

$$P'\Sigma_L P = L_0 I_2, \qquad P'\Sigma_v P = \operatorname{diag}\big(\sigma_v^2(1+a), \sigma_v^2(1-a)\big). \tag{A77}$$

We now apply Proposition 2 in this diagonal basis.

First,

$$P'\Sigma_x P = k_0^{-1}(P'\lambda P)^{-1}(2I_2 - k_0 P'\lambda P)^{-1}L_0 I_2 = \operatorname{diag}\left(\frac{L_0(m_2+1)^2}{4m_2}, \frac{L_0(m_1+1)^2}{4m_1}\right). \tag{A78}$$

Rotating back gives

$$\Sigma_x = P\,\mathrm{diag}\left(\frac{L_0(m_2+1)^2}{4m_2}, \frac{L_0(m_1+1)^2}{4m_1}\right)P'$$
$$= \frac{L_0}{8m_1m_2}\begin{bmatrix} m_1^2m_2+m_1m_2^2+4m_1m_2+m_1+m_2 & (m_1-m_2)(1-m_1m_2) \\ (m_1-m_2)(1-m_1m_2) & m_1^2m_2+m_1m_2^2+4m_1m_2+m_1+m_2 \end{bmatrix}.$$

This proves the formula for Σ_x.

Second, Proposition 2 implies

$$P'(\Sigma_p-\Sigma_v)P = 2(P'\lambda P)^2(2I_2-k_0P'\lambda P)^{-1}L_0I_2 = \mathrm{diag}\left(\frac{4L_0}{k_0^2m_2(m_2+1)}, \frac{4L_0}{k_0^2m_1(m_1+1)}\right). \tag{A79}$$

Hence

$$\Sigma_p = P\,\mathrm{diag}\left(\frac{4L_0}{k_0^2m_2(m_2+1)}, \frac{4L_0}{k_0^2m_1(m_1+1)}\right)P' + \sigma_v^2\begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}$$
$$= \frac{2L_0}{k_0^2m_1m_2(m_1+1)(m_2+1)}\begin{bmatrix} m_1^2+m_2^2+m_1+m_2 & (m_1-m_2)(m_1+m_2+1) \\ (m_1-m_2)(m_1+m_2+1) & m_1^2+m_2^2+m_1+m_2 \end{bmatrix}$$
$$+ \sigma_v^2\begin{bmatrix} 1 & a \\ a & 1 \end{bmatrix}.$$

This proves the formula for Σ_p.

Third,

$$P'A_pP = -k_0(P'\lambda P)^3(2I_2-k_0P'\lambda P)^{-1}L_0I_2 = -\,\mathrm{diag}\left(\frac{4L_0}{k_0^2m_2(m_2+1)^2}, \frac{4L_0}{k_0^2m_1(m_1+1)^2}\right). \tag{A80}$$

Therefore,

$$A_p = -P\,\mathrm{diag}\left(\frac{4L_0}{k_0^2m_2(m_2+1)^2}, \frac{4L_0}{k_0^2m_1(m_1+1)^2}\right)P' = -\frac{2L_0}{k_0^2m_1m_2(m_1+1)^2(m_2+1)^2}\times$$
$$\times\begin{bmatrix} m_1^3+m_2^3+2m_1^2+2m_2^2+m_1+m_2 & (m_1-m_2)(m_1^2+m_2^2+2m_1+2m_2+m_1m_2+1) \\ (m_1-m_2)(m_1^2+m_2^2+2m_1+2m_2+m_1m_2+1) & m_1^3+m_2^3+2m_1^2+2m_2^2+m_1+m_2 \end{bmatrix}. \tag{A81}$$

This proves the formula for A_p.

Finally, the matrix Σ_p has the form $\begin{bmatrix} u & v \\ v & u \end{bmatrix}$, hence $\rho_p = v/u$. Using the formula for Σ_p above, we get

$$u = \frac{2L_0}{k_0^2} \frac{m_1^2 + m_2^2 + m_1 + m_2}{m_1 m_2 (m_1 + 1)(m_2 + 1)} + \sigma_v^2, \tag{A82}$$

$$v = \frac{2L_0}{k_0^2} \frac{(m_1 - m_2)(m_1 + m_2 + 1)}{m_1 m_2 (m_1 + 1)(m_2 + 1)} + a\sigma_v^2. \tag{A83}$$

Dividing numerator and denominator by $2L_0/k_0^2$ and using $\nu_0 = k_0^2 \sigma_v^2/(2L_0)$ proves the stated formula for ρ_p. □

Appendix B. Expected Profit under Liquidation

We consider the case in which there is a single risky asset, i.e., $N = 1$. For simplicity, we assume that the bid-ask spread is zero and trading in each period takes place at the same price. At each time $t \in \mathbb{N}$, the asset liquidates with a probability $\pi \in (0, 1)$. If the asset does not liquidate at t, the investor trades $\Delta x_t = x_{t+1} - x_t$ at a price p_t. If the asset liquidates at t, the investor receives a cash amount of v_t for each of the x_t shares and thus realizes the following profit:[14]

$$\Pi_t \;=\; x_t v_t - \sum_{s=0}^{t-1} p_s \Delta x_s \;=\; x_0 v_t + \sum_{s=0}^{t-1} (v_s - p_s) \Delta x_s. \tag{B1}$$

Define $\beta \in (0, 1)$ by:

$$\beta \;=\; 1 - \pi. \tag{B2}$$

[14] We use the convention that a sum from a larger to a smaller index is zero, e.g., $\sum_{t=0}^{-1} p_t \Delta x_t = 0$.

The investor's expected profit at $t = 0$ is:[15]

$$\begin{aligned}\Pi &= \mathsf{E}_0\Big(\pi\Pi_0 + (1-\pi)\pi\Pi_1 + \cdots + (1-\pi)^t\pi\Pi_t + \cdots\Big) \\ &= \mathsf{E}_0\Big(\sum_{t=0}^{\infty}(1-\beta)\beta^t\Pi_t\Big) \;=\; \mathsf{E}_0\Big(\sum_{t=0}^{\infty}\beta^t\Pi_t - \sum_{t=0}^{\infty}\beta^{t+1}\Pi_t\Big) \\ &= \mathsf{E}_0\Big(\Pi_0 + \sum_{t=0}^{\infty}\beta^{t+1}(\Pi_{t+1} - \Pi_t)\Big),\end{aligned} \tag{B3}$$

Equation (B3) shows that the expected profit can be considered as the expectation of a sum of discounted instantaneous profits: $\Delta\Pi_{t+1} = \Pi_{t+1} - \Pi_t$. From (B1), the instananeous profit at t satisfies:

$$\Delta\Pi_{t+1} \;=\; x_{t+1}v_{t+1} - x_tv_t - p_t\Delta x_t \;=\; x_{t+1}\Delta v_{t+1} + (v_t - p_t)\Delta x_t. \tag{B4}$$

Thus, the update in realized profit at $t+1$ has two sources: an inventory component, $x_t\Delta v_{t+1}$, and a trading component, $(v_{t+1} - p_t)\Delta x_t$. The inventory component measures the reevaluation of the existing inventory, and can be thought as a stochastic dividend paid at $t+1$. The trading component reflects the cash payment made for trading at t, but with a gain estimated at the value v_t.

We model risk aversion by having the dealer incur each period a penalty proportional to the variance of the instantanteous profit. Thus, the dealer's expected utility is of the form:

$$\sum_{t=0}^{\infty}\beta^t\Big(\mathsf{E}_t\big(\Delta\Pi_{t+1}\big) - \gamma\,\mathsf{Var}_t\big(\Delta\Pi_{t+1}\big)\Big), \tag{B5}$$

where γ is the dealer's risk aversion coefficient. As in equations (A1) and (A2), we define:

$$h_t \;=\; \frac{a_t - b_t}{2}, \quad p_t \;=\; \frac{a_t + b_t}{2}, \quad \delta_t \;=\; v_t - p_t, \quad \eta_t \;=\; \varepsilon_t^s - \varepsilon_t^b. \tag{B6}$$

[15]Note that the discount factor at t is β^{t+1}. Indeed, by convention, trading starts at $t = 0$, and even though the first instantaneous profit occurs at zero, it must be still discounted by β. This can be reconciled easily if the time before trading is denoted by $t = -1$.

Equations (1) and (3) imply that:

$$\begin{aligned}
\Delta x_t &= k\Big(\frac{a_t+b_t}{2}-v_t\Big)-\varepsilon_t^b+\varepsilon_t^s \\
&= -k\delta_t+\eta_t, \\
\Delta\Pi_{t+1} &= x_{t+1}\Delta v_{t+1}-(v_t-a_t)Q_t^b+(v_t-b_t)Q_t^s \\
&= x_{t+1}\Delta v_{t+1}-k\delta_t^2-kh_t^2+2h_t\ell-(\delta_t+h_t)\varepsilon_t^s+(\delta_t-h_t)\varepsilon_t^b.
\end{aligned} \tag{B7}$$

At the optimum (see the proof of Proposition 1), we have:

$$\begin{aligned}
\delta_t &= \lambda x_t, \quad \text{with} \quad \lambda = \frac{1}{2}\Big(-\gamma+\sqrt{\gamma^2+\frac{4\gamma}{k}}\Big), \\
h_t &= h = \frac{\ell}{k},
\end{aligned} \tag{B8}$$

where λ is computed in the limit case when $\beta = 1$. Equation (B7) becomes:

$$\begin{aligned}
x_{t+1} &= (1-k\lambda)x_t+\eta_t, \\
\Delta\Pi_{t+1} &= x_{t+1}\Delta v_{t+1}-k\lambda^2x_t^2+\frac{\ell^2}{k}-\lambda x_t\eta_t-\frac{\ell}{k}(\varepsilon_t^s+\varepsilon_t^b).
\end{aligned} \tag{B9}$$

As x_t and v_t are in the dealer's information set at t, so are $a_t = v_t + h$, $b_t = v_t - h$, and $\delta_t = v_t - (a_t + b_t)/2$. Therefore, the variance of $\Delta\Pi_{t+1}$ conditional on the information at t is:

$$\mathsf{Var}_t(\Delta\Pi_{t+1}) = \mathsf{Var}_t\Big(\big((1-k\lambda)x_t+\eta_t\big)\Delta v_{t+1}-\lambda x_t\eta_t-\frac{\ell}{k}(\varepsilon_t^s+\varepsilon_t^b)\Big). \tag{B10}$$

If we define:

$$\sigma_\eta^2 = \mathsf{Var}(\eta_t) = \mathsf{Var}(\varepsilon_t^s-\varepsilon_t^b) \quad \text{and} \quad \sigma_v^2 = \mathsf{Var}(\Delta v_{t+1}), \tag{B11}$$

we have:

$$\begin{aligned}
\mathsf{Var}\Big(\big((1-k\lambda)x_t+\eta_t\big)\Delta v_{t+1}\Big) &= \big((1-k\lambda)^2x_t^2+\sigma_\eta^2\big)\sigma_v^2, \\
\mathsf{Var}\big(\lambda x_t\eta_t\big) &= \lambda^2x_t^2\sigma_\eta^2, \\
\mathsf{Var}\big(h(\varepsilon_t^s+\varepsilon_t^b)\big) &= \frac{\ell^2}{k^2}\sigma_\eta^2.
\end{aligned} \tag{B12}$$

Equations (B10) and (B12) imply:

$$\mathsf{Var}_t(\Delta\Pi_{t+1}) \;=\; (1-k\lambda)^2 x_t^2\sigma_v^2 + \sigma_\eta^2\Big(\sigma_v^2 + \lambda^2 x_t^2 + \frac{\ell^2}{k^2}\Big) + \text{ correlation terms.} \tag{B13}$$

We consider the case in which $k\gamma$ is very small:

$$1 \gg k\gamma. \tag{B14}$$

We compute:

$$\begin{aligned} k\lambda \;&=\; \frac{-k\gamma + \sqrt{(k\gamma)^2 + 4k\gamma}}{2} \;=\; \sqrt{k\gamma} + o(k\gamma) \;\approx\; \sqrt{k\gamma}, \\ \lambda^2 \;&\approx\; \frac{\gamma}{k}. \end{aligned} \tag{B15}$$

Equations (B13) and (B15) imply:

$$\mathsf{Var}_t(\Delta\Pi_{t+1}) \;=\; \big(1-\sqrt{k\gamma}\big) x_t^2\sigma_v^2 + \sigma_\eta^2\Big(\sigma_v^2 + \frac{\gamma}{k} x_t^2 + \frac{\ell^2}{k^2}\Big) + \text{ correlation terms.} \tag{B16}$$

Therefore, if γ and σ_η are sufficiently small, we obtain:

$$\mathsf{Var}_t(\Delta\Pi_{t+1}) \;\approx\; x_t^2\sigma_v^2. \tag{B17}$$

This shows that when the dealer's risk aversion is sufficiently low, the expected utility in (B5) can be approximated by:

$$\mathsf{E}_0 \sum_{t=0}^{\infty} \beta^t \Big(\mathsf{E}_t\, \Delta\Pi_{t+1} - \gamma x_t^2\Big), \tag{B18}$$

where we have replaced $\gamma\sigma_v^2$ by γ. Thus, we can replace the quadratic utility term, $\gamma\,\mathsf{Var}(\Delta\Pi_{t+1})$, by the quadratic inventory penalty, γx_t^2.

Appendix C. Additional Figures and Tables

Table C.1: Variable definitions

Table	Variable	Definition
Primitives	Inventory (shares)	Signed share position of a market maker in a stock, cumulated from her fills and carried across days.
Primitives	Inventory (CAD)	Inventory in shares marked to market at the last trade price.
Primitives	Quoted liquidity imbalance	Resting ask depth minus resting bid depth, measured in dollars. Positive values indicate that the market maker supplies relatively more liquidity on the ask side than on the bid side.
Primitives	Market maker group	DMM, EMM, or Both, fixed per market maker over the sample. The universe is the 549 market makers active in at least two securities.
1, Panel A	Number of stocks	Distinct stocks traded that day, matched to CFMRC. All Panel A variables are CFMRC stock-day values, averaged across stocks each day, then across days.
1, Panel A	Trades per stock	Number of trades in a stock-day.
1, Panel A	Share (dollar) volume per stock	Share volume, and closing price times share volume, per stock-day.
1, Panel A	Closing price	Closing price in CAD.
1, Panel A	Quoted (relative) spread	Closing ask minus closing bid in CAD, and the same divided by the quote midpoint, in percent.
1, Panel A	Daily return	CFMRC daily return, in percent.
1, Panel B	Trading days	Days on which the market maker trades.
1, Panel B	Distinct securities	Securities the market maker trades over the sample.
1, Panel B	Stocks traded per day	Distinct stocks the market maker trades that day, averaged over her days.
1, Panel B	Trades (dollar volume) per day	Trades, and gross dollar volume, per trading day.
1, Panel B	Net portfolio zero reversions	Times per day the market maker's net portfolio, the sum of signed positions across stocks, reaches or crosses zero: a nonzero net moving to zero, or to the opposite sign, at the next minute.
1, Panel B	Days ending flat	Percent of days with zero end-of-day inventory.
1, Panel B	Passive-fill share	Percent of fills where the market maker's resting order is executed.
2	Inventory, stock (CAD)	Signed share position in one stock marked to market at the prevailing price, per market maker, stock, day, and minute.
2	Net portfolio (CAD)	Signed sum across stocks of the stock inventory (CAD), per market maker, day, and minute.

Table	Variable	Definition
2	Gross portfolio (CAD)	Sum across stocks of the absolute stock inventory (CAD).
2	Zero reversions per day	Number of times within a day the signed position reaches or crosses zero. Stock level on the stock inventory, portfolio level on the net portfolio. Computed per day, then averaged (stock: equal-weight across the market maker's stock-days; portfolio: across the market maker's days).
2	Minutes between reversions	Mean number of minutes between consecutive reversions within a day, averaged over days with at least one reversion; days and market makers that never revert are excluded.
2	Mean signed (end of day)	Signed position at the last minute of the day, averaged over days; net portfolio at the portfolio level. Near zero means no standing directional position.
2	Mean absolute (end of day)	Absolute end-of-day position at the stock level; end-of-day gross portfolio at the portfolio level; averaged over days.
2	Intraday range (CAD)	Within-day maximum minus minimum of the signed position (stock inventory at the stock level, net portfolio at the portfolio level), averaged over days.
2	Intraday volatility (CAD)	Within-day standard deviation of that same signed series, averaged over days.
2	End-of-day share long (short)	Percent of a market maker's held stocks that end the day long (short); the remainder end flat. Portfolio level only. Computed per day and averaged over days.
2, Panel B	Paired portfolio−stock difference	Per market maker, its portfolio zero reversions per day minus its stock zero reversions per day.
C.3, C.4	Rest-of-portfolio quoted liquidity imbalance	Sum of the quoted liquidity imbalance over all stocks except the focal one, per market maker-minute.
C.3, C.4	Rolling correlation rc	Thirty-minute rolling correlation between the rest-of-portfolio quoted liquidity imbalance and the market maker's inventory in the focal stock five minutes earlier.
C.3, C.4	Cross-hedging episode	Run of at least five consecutive minutes with $rc > 0$ for a market maker, day, and focal stock. Episode length is the number of such minutes.
C.3, C.4	Retained episode	Episode whose covariance between the lagged focal inventory and the rest-of-portfolio imbalance is positive, in agreement with the selection; about 11 percent fail and are excluded.

Table	Variable	Definition
C.3, C.4	Covariance contribution	Covariance of lagged focal inventory with a single hedge stock's quoted liquidity imbalance. Summed over stocks it reproduces the rest-of-portfolio covariance exactly.
C.3, C.4	Positive covariance share	A positive contributor's covariance as a share of the episode's total positive covariance.
C.3, C.4	Top-1/3/5 share	Summed positive covariance share of the largest one, three, and five contributors.
C.3, C.4	Positive contributors	Number of hedge stocks with positive covariance contribution.
C.3	Covariance offset ratio	Share of the focal inventory covariance offset by the rest-of-portfolio imbalance within the episode.
C.4	Pair episodes (market makers)	Number of episodes, and distinct market makers, in which a focal-hedger pair appears.
C.4	Mean ρ_j	Mean return correlation between the focal stock and the hedge stock.
3	Pair-episode	One focal-hedger pair on one market maker-day; the covariance between the five-minute-lagged focal inventory and the hedger's quoted liquidity imbalance is computed over the day's episode minutes.
3	Inventory correlation	Correlation between the market maker's five-minute-lagged inventory in the focal stock and her contemporaneous inventory in the hedger over the episode window; pair-episodes with fewer than 20 overlapping minutes, or with no inventory variation, are unclassified. Negative if at or below -0.1, Near-zero if strictly between -0.1 and 0.1, Positive if at or above 0.1.
3	Return correlation	Correlation between the 30-minute returns of the focal and hedge stock, a single pairwise value over the sample; the same correlation used in Table C.4. Rows bin it left-closed at $-0.1, 0, 0.1, 0.3$.
3	Gradient	Increase in the share quoted to sell from the bottom to the top return-correlation bin within an inventory-correlation group. Panel B re-estimates it per market maker across bins, over market makers with episodes in at least three bins.

Table C.2: Inventory Dynamics (medians)

Median version of Table 2. We report the median across market makers, with the standard deviation in parentheses. Definitions, construction, and columns are as in Table 2. The paired test (Panel B there) is a single test per group and is not repeated here.

	Stock			Portfolio		
	All	DMM	EMM	All	DMM	EMM
Turnover						
Zero reversions per day	0.27	0.56	0.14	0.62	1.42	0.25
	(1.09)	(1.64)	(1.14)	(3.31)	(5.60)	(2.47)
Minutes between reversions	59.7	64.2	54.4	32.1	31.7	34.0
	(29.6)	(28.7)	(36.0)	(42.1)	(29.9)	(55.0)
End-of-day position (000s CAD)						
Mean signed	1.08	-0.32	0.08	6.38	-2.26	0.66
	(1,113)	(146.5)	(1,034)	(41,013)	(1,327)	(5,211)
Mean absolute	63.5	34.7	178.6	575.6	349.2	820.0
	(1,334)	(621.0)	(1,141)	(65,411)	(6,302)	(18,873)
Intraday dispersion (000s CAD)						
Intraday range	31.7	22.8	37.3	189.2	132.0	131.4
	(122.4)	(58.2)	(153.9)	(2,349)	(344.1)	(1,193)
Intraday volatility	7.39	5.35	8.51	41.4	34.2	29.5
	(33.6)	(14.5)	(41.7)	(646.9)	(92.0)	(303.3)
End-of-day sign mix (portfolio, %)						
Share long				47.9	49.6	44.3
				(22.0)	(14.8)	(27.6)
Share short				41.0	48.0	39.5
				(21.9)	(14.0)	(27.6)

Table C.3: Episode Population Characteristics

Cross-hedging episodes over the 45 trading days. An episode is a run of at least five consecutive minutes in which the thirty-minute rolling correlation between the rest-of-portfolio quoted liquidity imbalance and the market maker's five-minute-lagged inventory in the focal stock is positive. Panel A counts episodes, market makers, and market maker-focal cells. Panel B reports the concentration of the episode's positive covariance across hedge stocks. Panel C reports the strength of the association within episodes and the covariance offset ratio, the share of the focal inventory covariance offset by the rest-of-portfolio imbalance. Variable definitions are presented in Table C.1. The sample period is 3 January to 6 March 2006.

	All	DMM	EMM	Both
A. Coverage				
Episodes	879,942	58,634	449,107	372,201
Market makers	537	68	242	227
Market maker-focal cells	12,645	1,074	4,484	7,087
Episodes per market maker (median)	742	764.5	281	1,140
B. Concentration				
Top-1 cov share (median)	0.489	0.806	0.305	0.669
Top-3 cov share (median)	0.857	1.000	0.615	0.989
Top-5 cov share (median)	0.962	1	0.773	1.000
Share top-3 >= 0.80	0.548	0.966	0.314	0.764
Positive contributors (median)	13	4	24	6
C. Strength				
Episode length in minutes (median)	26	24	26	25
Covariance offset ratio (median)	0.130	0.096	0.131	0.133
Share cov-offset < 0.30	0.690	0.701	0.694	0.683

Table C.4: Recurring focal-hedger pairs

The thirty most frequent focal-hedger pairs across episodes. Episodes and Market makers count the episodes and the distinct market makers in which the pair appears. Med. share is the hedge stock's median positive covariance share within those episodes. Mean ρ_j is the mean return correlation between the focal and the hedge stock, computed on 30-minute returns over the sample period. Tickers are TSX symbols. The recurring pairs are energy and resources stocks, with a mean pairwise return correlation near 0.44.

Focal	Hedge	Episodes	Market makers	Med. share	Mean ρ_j
ECA	CNQ	4,878	70	0.069	0.446
CNQ	ECA	4,792	68	0.088	0.446
PCA	ECA	4,242	53	0.074	0.445
PCA	CNQ	4,206	59	0.061	0.423
ECA	PCA	4,143	55	0.043	0.436
TLM	CNQ	4,023	51	0.060	0.448
AL	ECA	4,017	48	0.066	0.438
TLM	ECA	4,005	51	0.072	0.452
CNQ	PCA	4,003	59	0.049	0.422
CNQ	SU	3,987	57	0.083	0.431
SU	CNQ	3,935	54	0.063	0.446
ECA	SU	3,917	52	0.077	0.429
AL	CNQ	3,894	44	0.054	0.437
ECA	TLM	3,887	53	0.042	0.423
TLM	SU	3,831	47	0.076	0.443
CNQ	TLM	3,785	51	0.044	0.418
SU	ECA	3,781	49	0.071	0.451
PCA	TLM	3,655	52	0.041	0.421
TLM	PCA	3,639	49	0.040	0.436
SU	TLM	3,628	48	0.043	0.430
PCA	SU	3,533	50	0.069	0.429
CCO	ECA	3,391	33	0.066	0.447
CCO	CNQ	3,378	37	0.055	0.451
AL	PCA	3,351	38	0.036	0.424
ECA	AL	3,328	48	0.029	0.415
G	ECA	3,243	34	0.067	0.460
SU	PCA	3,242	47	0.040	0.436
PCA	NXY	3,202	41	0.026	0.415
NXY	TLM	3,176	37	0.036	0.438
ABX	ECA	3,140	38	0.063	0.455