

(very preliminary and incomplete)

Price Pressures¹

Terrence Hendershott and Albert J. Menkveld

May 4, 2009

¹Terrence Hendershott, Haas School of Business, University of California at Berkeley, 545 Student Services Bldg #1900, Berkeley, CA 94720, tel +1 510 643 0619, fax +1 510 643 1412, hender@haas.berkeley.edu. Albert J. Menkveld, VU University Amsterdam, FEWEB, De Boelelaan 1105, 1081 HV, Amsterdam, Netherlands, tel +31 20 598 6130, fax +31 20 598 6020, albertjmenkveld@gmail.com. We benefit from in-depth discussions with Bruno Biais, Sunny (Xiaoyue) Li, and Jan van Schuppen. The project was inspired by the Ljungqvist and Sargent NYU lecture series based on their book *Recursive Macroeconomic Theory*. Menkveld gratefully acknowledges Robert Engle for sponsoring his NYU visit in 08/09 and VU University Amsterdam for a VU talent grant.

Price Pressures

Abstract

We study price pressures—price deviations from fundamental values due to a risk-averse intermediary supplying liquidity to asynchronously arriving investors with idiosyncratic hedging values. In our model, the intermediary uses price pressure to mean-revert costly inventory by trading off the size of the price pressure against the cost of remaining in a risky inventory state. Price pressure is associated with the social cost of lower realization of investor hedging value because the intermediary’s efforts to mean revert inventory substitutes low-hedge-value investors on the side of the market that reduces the risk of the intermediary’s position for high-hedge-value investors on the risk-increasing side. Empirically, twelve years of daily New York Stock Exchange (NYSE) intermediary data reveal economically large price pressures and associated social costs. A \$100,000 inventory shock causes price pressure of 1.01% for the small-capitalization stocks and 0.02% for the large-cap stocks. However, the price pressure conditional on inventory reduces inventory positions leading to the smaller differences in transitory volatility in daily stock returns (average price pressure): 1.20% vs. 0.17% for small and large stocks, respectively. The intermediary’s larger volatility of inventory positions in large stocks results in a greater model-based estimate of social cost due lower unrealized hedge gains for large stocks of \$9.68 million per year vs. \$0.67 million for small stocks. The aggregate lost hedging gains are estimated to be greater than \$50 billion for all NYSE common stocks for our sample period.

Keywords: liquidity, inventory risk, intermediary JEL: G12, G14, D61

Studies of price pressure often measure the magnitude of transitory price deviations from fundamental value, commonly referred to as price pressure, due to asynchronously arriving investors with idiosyncratic hedging values demanding substantial immediacy/liquidity from intermediaries supplying liquidity.¹ These provide evidence that non-informational trading frictions and associated price pressures can be large.² Price pressure is associated with two sources of welfare loss due to inefficient allocation of risk: the intermediaries' cost of bearing risks and the intermediaries' risk reducing behavior lowering the realized hedge values by investors. By examining the liquidity supplying behavior of intermediaries we can calculate the normal costs for representative investors of trading large quantities for non-informational reasons, the average magnitude of the deviation of prices from fundamental values, and the social costs of price pressure. To characterize the magnitude of price pressure in terms of frequency and size and its social costs we use 12 years of a daily measure of liquidity supplied on the New York Stock Exchange (NYSE).

To study the dynamics of price pressure we construct a single-asset theoretical model of liquidity supply with the reduced-form assumption that investors with less than perfectly elastic demands to trade arrive stochastically. As in Grossman and Miller (1988) natural motivations for asynchronous arrivals are idiosyncratic shocks to agents' consumption or investment opportunity sets which lead to liquidity or hedging needs. If agents are impatient and if it is costly to continuously monitor the market for trading opportunities, the market is not always complete as allocative efficiency via perfect risk sharing is not possible. In this case an intermediary will naturally enter to bear the costs of monitoring the market and stand ready to both buy and sell (Townsend (1978)). To focus our results on price pressure the model has no information asymmetry.

In the infinite-horizon recursive model the intermediary dynamically chooses the prices at which she is willing to buy and sell, the bid and ask prices, respectively. When the intermediary is at her desired position in a security the bid and ask prices symmetrically straddle the security's fundamental value. If a seller then arrives the intermediary buys and her position, also referred to as her inventory level, is higher than desired, exposing her to idiosyncratic price risk. To mitigate this risk the intermediary then stochastically mean reverts her inventory back to zero by attempting

¹Kraus and Stoll (1972) provide some of the first evidence on liquidity demands from block trades causing price pressure. Harris and Gurel (1986) and subsequent papers on additions to the S&P 500 index find evidence for price pressure. Greenwood (2005) extends this examining transitory price effects upon a weighting change to the Nikkei 225. Coval and Stafford (2007) examine price pressure due to mutual fund redemptions. These are all demand side effects.

²Theoretical models examine circumstances that lead to substantial price pressure due to shocks to liquidity supply. Brunnermeier and Pedersen (2008) find that when funding liquidity is low liquidity shocks causes prices to deviate more from fundamentals. Acharya, Shin, and Yorulmazer (2008) examine how crises and financing constraint induces price pressure (fire sale prices) and foreign direct investment even when the foreign firms are inefficient owners.

to sell. She adjusts both the bid and ask prices downward to induce buyers to arrive faster than sellers. In doing this the intermediary skews her quotes downward by setting the average of the bid-ask quotes, the midquote price, to be below the fundamental value. The size of the deviation of the midquote price from the fundamental value is our theoretical and empirical measure of price pressure. As emphasized in Stoll (1978) and Grossman and Miller (1988) this is distinct from the measures of the cost of immediacy based on the width of the bid-ask spread.

The main innovation of our theoretical approach is to facilitate calculation of price pressure's social costs. We refer to investors idiosyncratic desire to trade as their hedging values, although the motivations for trade could be more general. The intermediary's use of price pressure to mean revert her inventory results in lower realization of investor hedging value because the price pressure substitutes for low-hedge-value investors on the side of the market that reduces the risk of the intermediary's position for high-hedge-value investors on the risk-increasing side. Assuming a specific distribution for investors' hedge value enables the model to quantify the magnitude of the unrealized investor hedging gains due to price pressure as a function of observables in the data. We also examine the measures of price pressure and the associated social costs by numerically solving the Ho and Stoll (1981) model.

Before empirically studying price pressure we use the Kalman filter to estimate a state-space model that decomposes stock prices into their fundamental value and noise: the random walk and stationary components. The stationary component represents pricing error around the fundamental value. One important observation from the theoretical model is that deviations from fundamentals due to intermediation are not immediately reversed, but rather last as long as it takes the intermediary to mean revert her inventory to its desired level. To account for this we allow for the pricing errors to follow an autoregressive process. Because hedging of risk associated with the common market-wide factor is relatively inexpensive we focus on idiosyncratic effects. The estimation also allows for delayed adjustment to the common factor. Throughout the paper we conduct our analysis yearly at the stock level and report averages across stocks by market capitalization quintiles. There is persistence in the pricing errors with first-order autocorrelation coefficients typically between 0.3 and 0.5. The daily volatility of the stationary component of prices is roughly 0.75% as compared to the daily volatility of the random walk component of 1.73%.

To determine how much of this pricing error is attributable to intermediary-related price pressure we employ intermediary data from the New York Stock Exchange from 1994-2005. Market makers who act as intermediaries to supply liquidity on the NYSE, called specialists, are required

to report their positions in every security every day.³ We extend the above estimation of pricing errors to incorporate the NYSE intermediaries' inventory positions by allowing the idiosyncratic inventory level to enter directly into the price equation. The coefficient on inventory represents the conditional price pressure. The standard deviation of the idiosyncratic inventory characterizes the frequency of price pressure. Combining the frequency and conditional price pressure yields the average price pressure which we refer to as the price pressure transitory volatility.

The conditional price pressure for the largest-quintile stocks is 0.02 basis points per one thousand dollars of intermediary inventory. For the smallest-quintile stocks one thousand dollars of inventory results in 1.01 basis points of price pressure. The small conditional price pressure in large stocks increases the frequency of price pressure: the standard deviation of inventory is \$1.1 million for the largest stocks versus \$165 thousand for the smallest stocks. Combining the frequency and conditional size of price pressure produces estimates of the daily price pressure transitory volatility that range from 0.17% for the largest stocks to 1.20% for smallest stocks.

The social costs associated with price pressure have two components of the inefficient allocation of risk: the intermediary's cost of bearing risks and the intermediary's risk reducing behavior lowering the realized hedge values by investors. The latter cost combines the number of investors and the average reduction in realized investor hedging value due to price pressure. While the conditional and unconditional price pressure is larger for small stocks, the greater number of investors in large stocks results in higher social costs from price pressure in the large stocks: unrealized hedge gains for large stocks is \$9.86 million per year vs. \$0.67 million for small stocks. The average stock's social cost from price pressure is \$3.60 million per year. Given that there are roughly 1,500 common stocks traded on the NYSE each year, the total costs for our 12-year sample period is more than \$50 billion.

Prior empirical work on intermediary inventories find support for risk management via inventory control, but weak support for inventories causing price pressure.⁴ Hasbrouck (1988) uses trade-designed liquidity demand to identify inventory effects. He finds evidence of negative auto-correlation of trade sign consistent with risk management, but finds no evidence of price pressure. Madhavan and Smidt (1991), Madhavan and Smidt (1993), and Hasbrouck and Sofianos (1993) employ NYSE

³Recent technology changes which sped up trading on the NYSE have reduced the role of the specialist in intermediation (Hendershott and Moulton (2007)). Now a more diverse group of traders fulfill the role of temporarily carrying inventory to facilitate trading between asynchronously arriving buyers and sellers. The lack of position data for all these participants makes a broad empirical characterization of price pressure more challenging in the current environment.

⁴Indirect evidence consistent with price pressure exists in the literature on return reversals (Lehmann (1990), Campbell, Grossman, and Wang (1993), and others).

intermediary inventory data and find evidence supporting risk management, but not price pressure.⁵ Hendershott and Seasholes (2007) use a long time series of NYSE data and find evidence of both inventory control and price pressure. Their price pressure findings are cross-sectional: a portfolio of stocks where the intermediary is long outperforms a portfolio of stocks where the intermediary is short by 45.4 basis points over two weeks. Our findings extend this portfolio approach to determine the price pressure per dollar of inventory, the average impact of price pressure on stocks' volatility, and the social costs of price pressure.

1 A simple model to characterize price pressure dynamics and its relation to social cost

In this section, we characterize the conditional and average price pressure for a security as a result of the friction of nonsynchronous arrivals of investors who demand immediacy. We model the supply of immediacy with a representative risk-averse intermediary who produces price quotes. If the intermediary's position differs from her desired position her optimal policy sets the midquote (average of the bid and ask quote) on that side of the fundamental value that creates a nonzero average net (signed) liquidity demand to mean-reverts her inventory.⁶ In this way she skews the midquote downward if she is long relative to her optimal position and upward if she is short. These transitory price effects manifest themselves as price pressures to liquidity demanders who, when trying to trade into a large position, push quotes further away from fundamental values.

The intermediary facilitating trade among asynchronous investor arrivals generates inefficient allocation of risk in two ways. First, when the intermediary takes on position she bears risk until the final investor arrives. Second, the intermediary's effort to reduce her risk decreases the amount of hedging value the investors can capture. If the intermediary is risk neutral with unlimited capital then there is no price pressure and no social costs due to inefficient allocation of risk.

Social cost of the intermediary bearing risk. The first component of the social cost of inefficient allocation is the intermediary bearing idiosyncratic price risk on her inventory position

⁵Ho and Macris (1984) estimate a model similar to Madhavan and Smidt (1991) using American Stock Exchange intermediary inventory data for stock options. Using data on inventories from London Stock Exchange dealers Hansch, Naik, and Viswanathan (1998), Reiss and Werner (1998), and Naik and Yadav (2003) examine intermediary risk management.

⁶One way to rationalize investor heterogeneity (impatient liquidity demander vs. intermediary) as a general equilibrium outcome is that it is costly to monitor the market (e.g., update quotes to prevent picking-off risk) and it is therefore natural to see some agents specialize who pay the market monitoring cost and stand ready to buy and sell (Townsend (1978)).

between counterbalancing liquidity demand arrivals. In a static model, Grossman and Miller (1988) argue that transitory price effects serve to compensate intermediaries who take on idiosyncratic price risk in matching asynchronous liquidity demand. Intermediaries buy at prices below fundamental value if sellers arrive before buyers or sell at prices above fundamental value if sellers arrive before buyers. We refer to these transitory price effects as price pressures.

Grossman and Miller argue that these transitory price effects are the essence of liquidity. They analyze the size of these effects in an industrial organization context thereby fitting early microstructure models of inventory-management (e.g., Stoll (1978)) “into a larger framework that also encompasses the ultimate demanders and suppliers (p.617).” They compare and contrast the price pressure measure to alternative measures and discuss the limitations of two alternative measures: the bid-ask spread and the liquidity ratio:

1. “... (as Stoll (1978) has emphasized) the bid-ask spread only exactly measures the intermediary’s return for providing immediacy if she simultaneously crosses (i.e., executes on both sides of) the trade, one at the bid and the other at the ask (p.628).”
2. “Another widely used empirical measure in inter-market comparison is the liquidity ratio, defined as the ratio of average dollar volume of trading to the average price change during some interval. . . These measures, of course, tell us at best only about past average associations between price changes and volume. They do not answer the critical question of how the sudden arrival of a larger-than-average order would affect price. Nor do they distinguish adequately among the sources of price volatility. A particular market may display high price variability not because it is illiquid but because new fundamental information arrives frequently (p.630).”

Social cost of unrealized hedge value. The second component of the inefficient allocation is that when the intermediary has an inventory position differing from her desired position she skews the quotes to encourage trading by low-hedge-value investors on the side of the market that mean-reverts her inventory, but at the same time the quote skewing discourages trading by higher-hedge-value investors who arrive at the inventory-increasing side of the market. That is, even if transaction rates are unchanged on nonzero inventory positions the *expected* hedge value of the investors on the side of the market that she discourages from trading is higher than the *expected* hedge value of the investors on the other side of the market that she now encourages to trade. We will further develop this argument and illustrate it in the context of a dynamic model inspired by Ho and Stoll (1981).

1.1 Recursive dynamic inventory model setup

It is our aim to characterize price pressure dynamics and analyze how they relate to the two components of social cost relative to a first best of efficient risk-sharing among liquidity demanders. Ultimately, we will identify the time series properties of price pressures in the data. The structural model allows us to interpret what we find empirically and, in particular, allows for an analysis of social cost. Ho and Stoll (1981) are the first to set up the dynamic program for the risk-averse intermediary and cast it in a finite time horizon where utility is generated only by final wealth. We cannot take this off-the-shelf model as the finite horizon does not mirror the infinite horizon that we find when analyzing supplier inventory and prices in real markets. Rather, we take the model, cast it in recursive form, and solve for the intermediary's optimal policy and the resulting stationary distribution.⁷


Liquidity demand process. Ho and Stoll (1981) use a reduced form model for the transaction rates (where transaction size is fixed at one unit) of public buyers and sellers that was first proposed by Garman (1976). These transaction rates are linear in the ask and bid price, respectively, where the public buy rate decreases in the ask price and public sell rate increases in the bid price.

The social cost analysis that we aim to develop requires a bit more structure on the liquidity demand process. Rather than separating the two sides of the market, we generalize the investor arrival process by considering agents who arrive at the market and depending on the prices quoted decided whether to buy, sell, or not trade. The investors incur idiosyncratic endowment shocks that can be hedged by an offsetting position in a (single) risky security that trades in the market. This hedge value v is stochastic and zero in expectation as it is a private value relative to a security's common value that we fix at zero throughout for ease of notation.⁸ The agent arrives at the market and only trades if v is above the ask (agent buys a unit) or below the bid (agent sells a unit). We parameterize this process as follows. The idiosyncratic shocks happen and agents arrive at the market according to a Poisson process with rate $2\lambda\theta$. Conditional on arrival, the (private) hedge

⁷Other papers also study dynamic inventory control. Madhavan and Smidt (1993) further enriches the Ho and Stoll (1981) setting by adding an informed liquidity demander. Amihud and Mendelson (1980) analyze a dynamic inventory model in recursive form where a risk-neutral intermediary maximizes average profit and stays away from an exogenously given maximum inventory constraint as arriving at that state limits her profit opportunity (she can essentially only exit on one side). We work from the Ho and Stoll (1981) model as we focus on the price risk associated with an suboptimal inventory position.

⁸The common value here is the equilibrium price in the absence of the liquidity friction that is modeled. It is the price at which all idiosyncratic shocks wash in the cross-section of investors. We refer to the investors' idiosyncratic desire to trade as their hedging values, although their private motivations for trade could be more general.

value v is uniformly distributed with support $[-\frac{1}{\theta}, \frac{1}{\theta}]$.⁹ This specification implies the following graphs for the arrival rate per hedge value (left-hand side) and, consequently, public buy and sell rates per ask and bid price, respectively (right-hand side):



where the arrival rate per hedge value is derived as the arrival rate ($2\lambda\theta$) times the hedge value density ($\frac{1}{2}\theta$) and the public buy and sell rates follow by integrating over hedge value rates from the ask to $\frac{1}{\theta}$ and from $\frac{1}{\theta}$ to the bid, respectively. The total hedge value rate (or the “size of the market”) is obtained by setting the bid and ask price equal to zero and is thus represented by the size of the triangle defined by $-\frac{1}{\theta}$ and $\frac{1}{\theta}$ on the price axis and $\lambda\theta$ on the transaction rate axis.¹⁰ The size of this area is λ which therefore represents the size of maximum “consumer surplus.” The realized hedge value rate depends on the bid and ask price set by the intermediary who “produces” liquidity. If she is risk-neutral and therefore operates on a zero cost rate, she charges the monopoly price $p = \frac{1}{2\theta}$ and, for public buys, captures a producer surplus rate equal to the square defined by A in the graph. The consumer surplus rate on this side of the market is the triangle B, and the deadweight loss due to monopoly power is C on this side of the market.

The second parameter θ governs the sensitivity of the net transaction rate (defined as the buy rate minus the sell rate which is the rate at which the intermediary’s inventory changes) standardized by the overall arrival rate. Lowering of both the bid quote and the ask quote by Δp increases the rate of buys minus sells by $2\Delta p\lambda\theta^2$ which, standardized by the arrival rate of $2\lambda\theta$, equals θ per Δp . Comparative statics will show that this is an important parameter in the intermediary’s optimization behavior. In a way, it parallels the price elasticity of demand that a

⁹The use of θ in both the distribution of hedge values and the arrival rate of trades makes investors total hedge value independent of θ . This is useful for examining comparative statics with respect to θ , but is not necessary of any of our results.

¹⁰The analogy is that the expected value of a stochastic variable X can be calculated by integrating $(1-F(X))$ over the domain of X where $F(X)$ is its CDF.

monopolist faces in a standard product market.

The intermediary's dynamic program. We set up the intermediary's optimization in discrete time¹¹ where the Poisson hedge demand arrivals are modeled as discrete probabilities that are linear in prices.¹² The intermediary is risk-averse and maximizes expected utility where her utility is the standard time-separable CRRA utility. She is a monopolist liquidity supplier who produces liquidity supply by issuing firm price quotes that are take-it-or-leave-it prices that liquidity demanders see when arriving at the market.¹³

Let $V(\cdot)$ be the maximum utility that the intermediary can achieve when starting off with an inventory of i_0 shares and a wealth of w_0 :

$$V(i_0, w_0) = \max_{\{c_t, a_t \in [0, \frac{1}{\theta}], b_t \in [-\frac{1}{\theta}, 0]\}_{t=0}^{\infty}} E[\sum_{t=0}^{\infty} \beta^t u(c_t) | i_0, w_0] \quad (1)$$

subject to the budget constraint

$$w_{t+1} = R(w_t - c_t) + i_t \Delta m_{t+1} + a_t \max(-\Delta i_t, 0) - b_t \max(\Delta i_t, 0) \quad (2)$$

and subject to the following transition probabilities across inventory states

$$\begin{aligned} P[i_{t+1} = i_t - 1 | i_t] &= \int_p^{\frac{1}{\theta}} (2\lambda\theta) \left(\frac{1}{2}\theta\right) dv = \lambda\theta(1 - \theta a_t) \\ P[i_{t+1} = i_t + 1 | i_t] &= \lambda\theta(1 + \theta b_t) \end{aligned} \quad (3)$$

where i_t is the intermediary's inventory position at time t , w_t is her wealth, and c_t is her consumption at time t , β is her discount factor, R is the gross riskfree return, Δm_t is her stochastic dividend stream in period $t+1$ (which runs from time t to time $t+1$), and b_t (a_t) is the bid (ask) price she sets at time t . The stochastic dividend stream is a modeling device to minimize the accounting in the model; we emphasize that no sources of uncertainty are ignored (see also Ho and Stoll (1981, p.52)).

We exploit the recursive nature of the problem which leads to the following Bellman equation

¹¹The discrete time setup along with a discretization of the state-space allows us to find an exact numerical solution to the intermediary's dynamic program.

¹²This is a reasonable approximation as the Poisson process itself is the limiting case of event probabilities that are linear in small discrete time increments.

¹³The intermediary can therefore not price discriminate according to the size of the hedge value drawn by the liquidity demander who arrives at the market.

for the inventory state i (using the law of iterated expectations):

$$\begin{aligned}
V(i, w) = & \max_{c_{iw}, a_{iw}} u(c_{iw}) + \beta E_{\Delta m_{t+1}} \left[V(i+1, R(w - c_{iw}) + i\Delta m_{t+1} + b_{iw}) \right] \lambda \theta (1 + \theta b_{iw}) + \\
& \beta E_{\Delta m_{t+1}} \left[V(i-1, R(w - c_{iw}) + i\Delta m_{t+1} + a_{iw}) \right] \lambda \theta (1 - \theta a_{iw}) + \\
& \beta E_{\Delta m_{t+1}} \left[V(i, R(w - c_{iw}) + i\Delta m_{t+1}) \right] (1 - \lambda \theta (1 + \theta b_{iw}) - \lambda \theta (1 - \theta a_{iw})).
\end{aligned} \tag{4}$$

With upperbounds on the (absolute) inventory position and wealth (that the intermediary appears to stay away from endogenously for coefficient of relative risk aversion larger than one and a discount rate that is strictly less than the risk-free rate¹⁴), it is straightforward to show that the functional equations defined by (4) define a contraction in the (i, w) space by verifying Blackwell's pair of sufficient conditions (see Ljungqvist and Sargent (2004, p.1012)). We therefore have a unique fixed point in the space of bounded continuous functions which guarantees the existence of a unique equilibrium.

1.2 Price pressure and social costs: a model-based interpretation

The structure of the dynamic program allows for a calculation of how price pressures relate to the two components of social cost due to inefficient allocation of risk raised earlier. We focus on the price pressure's effect on the ask quotes as the effect on the bid quotes is symmetric.

Price pressure as a measure of the intermediary's cost of risk. The first order condition for the ask price in the dynamic program optimization of equation (4) yields

$$\frac{1}{\theta} - a_{iw} = \frac{E[V(i-1, X_{iw} + a_{iw})] - E[V(i, X_{iw})]}{E[V_2(i-1, X_{iw} + a_{iw})]} \equiv -\Delta_i \tilde{V}_{iw} \tag{5}$$


where V_2 is the derivative of the value function with respect to wealth, $X_{iw} = (R(w - c_{iw}) + i\Delta m_{t+1})$, and the expectation is taken with respect to X_{iw} that has the stochastic dividend as its only source of randomness. The right-hand side expression denoted by $-\Delta_i \tilde{V}_{iw}$ is a measure of how a one unit reduction in inventory affects the intermediary's utility. The numerator captures the differential

¹⁴Huggett (1993) shows that for exogenous nondiversifiable endowment risk these conditions guarantee that the agent does not let wealth go to infinity. The intuition is that with CRRA utility the intermediary becomes more willing to bear absolute risk when wealthier and if the savings rate is strictly smaller than the discount rate she enjoys a benefit out of consuming today rather than tomorrow. Our model extends this setting as the exposure to idiosyncratic risk (inventory) is decided on endogenously. The intermediary, however, does not visit extreme inventory states as her earning potential does not change with inventory—the arrival rate functions are independent of the inventory state—whereas the cost due to a stochastic dividend ($i\Delta m_{t+1}$) does grow with the inventory position. In the numerical solution we set the upperbounds “high enough” so that the intermediary decides never to visit these states. We return to this issue when discussing the numerical solution.

in the value function across inventory states but also across wealth states. The denominator effectively translates the utility difference into wealth units. $\Delta_i \tilde{V}_{iw}$ therefore becomes a measure of the marginal cost of a unit increase in inventory as it represents the amount one has to add to the intermediary's wealth for her to become indifferent between the two.

The ask price difference across inventory states is the change in the conditional price pressure, $a_{i+1,w} - a_{iw}$, and carries the interpretation of the *increase* in marginal cost of adding one inventory unit. If this conditional pressure is constant across inventory state changes (which the numerical solution seems to support) the price pressure linearly increases in the value differential between inventory states: $a_{i+1,w} - a_{iw} = \alpha$ and, therefore, $a_{iw} - a_{0w} = \alpha i$. The average price pressure that we measure empirically as transitory volatility (conditional pressure times the standard deviation of inventory) then becomes a meaningful metric of how costly the idiosyncratic risk is to the intermediary relative to a zero inventory position.

Price pressure as a source of unrealized hedge value. Price pressure not only relates to the social cost of inefficient allocation due to the intermediary bearing risk, it also creates a cost on the liquidity demand side. Even if the size of the bid-ask spread is not changed (so transaction rates are not altered), skewing the quotes so that the midquote does not equal fundamental value (i.e. price is pressured) generates a social cost of unrealized hedge value. The argument is best illustrated graphically:



Here the bid and ask quotes are lowered to mean-revert a positive intermediary inventory position. This makes buys become more likely than sells. The overall transaction rate is not affected as the size of the spread is unchanged. The realized hedge value rate, however, is affected as the hedge value increase due to more buys (light gray area) is smaller than the hedge value lost by fewer sells (dark gray area).

Economically, the expected hedge value of the marginal buyer that the intermediary now accepts is lower than the expected hedge value of marginal seller that she blocks on the other side of the market. This is the source of the loss of hedge value on the demand side as a result of price pressures. If the bid-ask spread is constant across inventory states we can calculate how price pressures affect the realization of investors' hedging gains. The numerical solution to the model find that inventory has little effect on the bid-ask spread.¹⁵

Empirically, the model allows one to retrieve the cost of unrealized hedge value as a function of the time series characteristics of price pressures. In the model, as illustrated by the graph above, the average social cost rate due to unrealized hedge value equals $\lambda\theta^2 E[(\Delta p)^2]$. The empirical equivalent of the right-most factor is the variance of price pressure. The factor by which it needs to be scaled also has an empirical equivalent as, in the model, a price pressure of $-\Delta p$ increases the net transaction rate (buys minus sells) by $2\lambda\theta^2\Delta p$. In the data, a state space model identifies a conditional price pressure parameter α that is the midquote deviation from the fundamental price as a linear function of the intermediary's inventory position, which we denote I in the empirical analysis. The linearity of Δp in inventory implies that $E[(\Delta p)^2] = \alpha^2\text{Var}(I)$. The data allows for estimation of the rate of mean-reversion in inventory by a first order autoregressive process. The sensitivity of the net transaction rate ($\lambda\theta^2$) can therefore be characterized by $(1 - \text{AR1}(I))/\alpha$, where $\text{AR1}(I)$ is the first order autocorrelation in the intermediary's inventory.¹⁶ The empirical measure for the social cost rate due to unrealized hedge value thus becomes:

$$\text{soc_cost_unrealized_hedge_val} = \text{Var}(\Delta p) \frac{(1 - \text{AR1}(I))}{2\alpha} = \frac{1}{2}\alpha\text{Var}(I)(1 - \text{AR1}(I)) \quad (6)$$


which is the price pressure variance multiplied by the rate of inventory mean reversion divided by the conditional price pressure; if price pressure is linear in inventory the reduction in realized hedge value is the conditional price pressure times the variance of inventory times one minus the persistence of inventory.

Is the trade-off of the two social costs efficient? The intermediary trades off the two source of social costs when setting prices. She internalizes the unrealized hedge value loss as less surplus on the demand side implies less surplus for her to capture in the transaction process.

¹⁵Zabel (1981) and Mildestein and Schleef (1983) find the spread is independent of inventory. Ho and Stoll (1981) find inventory has a very small effect of the spread.

¹⁶The numerator captures the net rate per unit of inventory, whereas the denominator is the price pressure per unit of inventory. The linearity of the AR(1) model makes the net transaction rate with the intermediary be proportional to the size of the price pressure which is consistent with the linearity assumed for the hedge value generating process in the economic model.

But, does the intermediary internalize the unrealized hedge value appropriately so that immediacy production is socially efficient? The following graph illustrates how skewing the quotes affects the part of the surplus she gets:



A comparison with the true social cost depicted in the previous graph shows that her private revenue loss is twice the social cost (essentially due to her inability to price discriminate and internalize the total hedge value effect). As her cost of pressuring price/skewing the quotes is larger than the true social cost, she does not pressure prices enough and mean-reversion is slower than what social efficiency would command.¹⁷ A full constrained Pareto optimality solution (to serve as a benchmark for the monopolist intermediary solution) requires one to specify liquidity demander utility so that a social planner can design the optimal policy for the intermediary that maximizes total utility. Such analysis is beyond the scope of the current paper.

1.3 Comparative statics on numerical solution

We numerically solve the intermediary’s dynamic program which is summarized by equation (1). The calibration follows, to the extent possible, the base case parametrization proposed by Ho and Stoll (1981, p.67). We present the calibration details in Appendix I.

[insert Figure 1]

The base case numerical solution. Figure 1 depicts the numerical solution for the standard parametrization. Panel A plots the value function which is defined as the maximum discounted utility the intermediary obtains conditional on starting off in a particular inventory-wealth state. Its

¹⁷This inefficiency critically hinges on the inability to price discriminate, not the monopolistic setting; we do not expect it to disappear in the presence of competition.

concavity in the state is inherited from the *discounting* that occurs in the dynamic program that the intermediary solves (see Ljungqvist and Sargent (2004, §A.2.)). Given this value function and the Bellman equations, it is a straightforward discrete state optimization to establish the intermediary's optimal control policy (i.e., bid price, ask price, and consumption) in each inventory-wealth state which in turn determines the system transition laws. If she starts off with sufficient wealth and zero inventory, she eventually ends up in the stationary distribution over the inventory-wealth states which is also depicted in Panel A. This stationary distribution graph illustrates that she endogenously chooses to stay far from the inventory upperbound and the wealth upperbound. She also steers clear of the absorbing bankruptcy state which therefore gets zero probability in the stationary distribution.

Panel B of Figure 1 illustrates the optimal control policy which consists of the bid price, the ask price, and the consumption in each of the inventory-wealth states. In the graphs, we illustrate the midquote and the bid-ask spread instead of the bid and ask price separately so as to get a better perspective on price pressures. The graph confirms the intuition that the intermediary skews the midquote to generate transactions that mean-reverts her inventory. She lowers the midquote on long positions and raises it on short positions.

The bid-ask spread is relatively stable across the various states and it is close to the spread that a risk-neutral intermediary would charge: $2(2\theta)^{-1} = 0.2$. This is a reassuring result as it confirms the Grossman and Miller (1988) argument that the spread captures the cost of processing orders (or monopoly rents in our case), whereas the distance of the midquote to the efficient price captures the deadweight loss of inefficient risk sharing due to an intermediary who bears idiosyncratic price risk. Consistent with Ho and Stoll (1981), comparing the variability in the bid-ask spread with the variability in the midquote across inventory states suggests that all the action is in the midquote, not in the spread.

The consumption graph illustrates that the intermediary self-insures against adverse dividend shocks by maintaining a wealth buffer. She consumes less than expected earnings on low wealth levels and more than expected earnings on high wealth levels. This is best illustrated by the zero inventory policies as her expected earnings of a potential current period transaction equal a constant $2*(1-0.5*0.23)*10=0.98$ as prices are unchanged across the wealth dimension. She consumes less than these expected earnings on low wealth levels and thus self-insures against adverse future dividend shocks through savings. She consumes more than expected earnings and thus eats out of her large wealth buffer on high wealth levels. This consumption pattern thus makes wealth mean-

reverting. We further notice that she reduces her consumption when on a larger nonzero inventory which reflects the higher need for a wealth buffer given that she is in a more risky state.

[insert Figure 2]

Comparative statics for the numerical solution. Figure 2 illustrates how the numerical solution changes and, more importantly, how the price pressure dynamics change when we increase the liquidity friction by changing the three deep parameters by 25%¹⁸: a higher stochastic dividend risk¹⁹ (σ), a lower hedge value rate (λ), and a less sensitive net transaction rate (θ). The graphs use the intermediary's state contingent bid and ask prices which define the system's transition laws and ultimately its stationary distribution to

- (i) depict expected price pressure conditional on an inventory position and the result of a population regression of the midquote on the inventory position,
- (ii) plot the stationary distribution of the intermediary's inventory,
- (iii) calculate the autocorrelation function for inventory.

The regression of the midquote on the intermediary's position produces a more negative coefficient if we increase the liquidity friction for the intermediary in either of the three dimensions: more volatility, less hedge value, or lower net transaction sensitivity. This conditional price pressure labeled α is larger in all three cases. Interestingly, the conditional expectation of the midquote is close to linear in the various inventory positions, which motivates the linearity we will assume in our econometric model.

The effect of the parameter changes on the average size of price pressure is nontrivial. The size of average price changes is characterized empirically by price volatility. As discussed in the theoretical model, if one assumes linearity for the conditional price pressure, then the average price pressure becomes the product of (the absolute value of) the conditional pressure α and the standard deviation of inventory. It is the latter that makes predicting the net effect of the parameter changes nontrivial. In the case of higher volatility or less hedge value, the intermediary endogenously keeps inventory levels close to zero as entering nonzero inventory becomes more costly. The net effect of these parameter changes on average pressure is therefore ambiguous. In the case of a less sensitive

¹⁸All changes are approximate as we have to work under the constraint that the intermediary stays on the discrete inventory-wealth states.

¹⁹We change the variance of the stochastic dividend by 25% rather than the volatility, which is the standard deviation.

net transaction rate ($\theta = 3.75$), it effectively becomes more expensive for the intermediary to steer inventory across states and she spends more time in nonzero inventory states. Only in this case, the net effect is that more friction generates a larger average price pressure. In other words, only in the case of less sensitive net flow does the social cost of the intermediary bearing risk unambiguously become larger.

Panel A in Table 1 summarizes the numerical results on the conditional and average size of price pressure. It shows that changes in the three deep parameters to increase the liquidity friction raise conditional price pressure—the magnitude of the regression coefficient when regressing the midquote on (signed) inventory, $|\alpha|$. The effect on average price pressure (as summarized by the midquote standard deviation) however is dampened by a lower frequency visiting the nonzero inventory states (as summarized by the standard deviation of inventory) for the cases of more volatility or a lower hedge-value rate. Only in the case that order flow is less sensitive to price pressure do we find an amplifying effect, i.e., a lower θ increases the average price pressure by a greater percentage than the conditional price pressure. The net effect on average price pressure across all parameter permutations differs in economic significance. The numerical solution shows that a 25% change in parameter values implies a small increase of 5% and 7% for higher volatility and less hedge value, respectively, whereas it is more than doubled, i.e., a 64% increase, for a less sensitive net transaction rate.

[insert Table 1]

Panel B provide an analysis of social cost as well as other statistics related to the intermediary’s behavior. Price pressure reduces the realized hedge value by encouraging lower-hedge-value investors on the inventory-mean-reverting side of the market while discouraging high-hedge-value investors on the other side. The numerical solution provides for an exact calculation of this cost and Panel B shows that, relative to the base case scenario, it increases for higher volatility (+27%), it decreases for less hedge value (−5%), and it decreases most for a less sensitive net transaction rate (−32%).²⁰ Intuitively, the increase on higher volatility is not surprising as the intermediary is more aggressively pressuring price to mean-revert inventory more quickly (the autocorrelation of

²⁰The table also calculates the proxy that will be used in the empirical analysis. This proxy uses the population moments (population regression coefficient α , standard deviation of inventory, and the autocorrelation in inventory) to calculate the “empirical statistic” suggested by equation (6). The result shows that it seems to underestimate the true social cost, this is due to the discretization needed to obtain the numerical solution in feasible computational time. When the implicit coarseness is reduced by reducing θ the proxy and exact unrealized hedge value become much closer.

inventory drops). The result that unrealized hedge value declines on lower hedge demand is easily understood as there is less hedge value to begin with. But, it is less clear why less sensitive net transaction rates imply a reduction in this type of social cost given that conditional price pressure increases. The reason is that it is too costly for the intermediary to increase in conditional pressure enough to compensate for the less sensitive transaction rate because. Therefore, she more slowly mean reverts her inventory (autocorrelation of inventory increases). The slow mean reversion reflects lower social costs in terms of a lower rate of substitute of higher-hedge-value investors on the mean-reverting side of the market for lower-hedge-value investors on the other side. Equation (6) shows this as the unrealized hedge value equals the the conditional price pressure times the variance of inventory times one minus the autocorrelation of inventory. In the case of less sensitive net transaction rates the reduction in the last term is greater than the increase in the first two terms, leading to lower social costs of unrealized hedge value.²¹

Interestingly, parameter changes have the exact opposite effect on the social costs due to the intermediary bearing risk. Based on our theoretical model we propose the following measure of social cost due to the intermediary bearing risk. We calculate the deterministic cash flow stream that makes the intermediary indifferent between the value she obtains in her original environment (weighting the various state-depend discounted utilities with the stationary distribution) and this “outside option.” The wedge between this “risk-neutral” cash flow stream and the average cash flow stream in the risky environment expresses the disutility the risk causes the intermediary expressed in monetary terms. The wedge increases for all three parameter changes as does the average price pressure. The size of its increase correlates with the size of the average price pressure increase where, again, the increase for a less price sensitive net transaction rate far dominates the two other dimensions along which the liquidity friction could deteriorate. Because this measure is based on the intermediary’s utility function, it is not possible to identify in the data.

Panel B further provides statistics on the intermediary’s consumption and wealth, the transaction rate, and the bid-ask spread for the three scenarios. It shows that the average consumption drops substantially (-28%) if the hedge value rate is reduced, which is not surprising as less demand reduces her earning power. The average consumption is largely unaffected in the other two states (-3/ + 1%).The average wealth is higher in all three cases which reflects stronger precautionary

²¹In the case of higher volatility and less sensitive arrival rates there is a counterbalancing effect due to a change in the dispersion of the stationary inventory distribution. For higher volatility the intermediary spends less time in nonzero inventory states which dampens the conditional effect. A similar reasoning applies for a less sensitive net transaction rate. This effect seems to be second order as it is appears dominated by the conditional effect.

savings to self-insure against the increased idiosyncratic price risk in the new environment. Its increase correlates with the size of the average price pressure (and the social cost of the intermediary bearing risk) which is perfectly intuitive. The benchmark for both the probability of a transaction and the average bid-ask spread is a risk-neutral monopolist who sets marginal price equal to zero and therefore sets the bid at $-\frac{1}{2\theta}$ and the ask at $+\frac{1}{2\theta}$. The benchmark bid-ask spread is therefore $\frac{1}{\theta}$ and the benchmark transaction rate is $\lambda\theta$. Relative to these benchmarks, the average spread is higher and the average transaction rate is lower which reflects the increased cost that the monopolist incurs in the more adverse market environment.²²

2 An empirical search for conditional and average price pressure and inventory dynamics

In this section we analyze a 12-year balanced panel of 697 NYSE stocks that has end-of-day observations on the midquote (i.e., the average of the bid and ask price) and the NYSE specialist inventory position. We aim to identify the conditional and the average size of price pressure for these stocks both in the cross-section and through time. The structural model suggests an interpretation of the size of the average pressure as a measure of the social cost due to the intermediary bearing risk. In addition, we characterize the inventory dynamics with a first order autoregressive process that, again contingent on the structural model, allows us to calculate the size of the social cost of price pressure due to unrealized hedge value (see equation (6)). Before any estimations, we discuss the data and provide some summary statistics.

2.1 Data and summary statistics

We use CRSP, the NYSE's Trade and Quotes (TAQ), and a proprietary NYSE dataset with specialist positions to prepare data on the end-of-day midquote and NYSE specialist position along with other variables from 1994 through 2005. We construct a balanced panel to make results comparable through time (which controls for stock fixed effects). We start with the sample of all NYSE common stocks that can be matched across TAQ and CRSP and retain the stocks that are present throughout the whole sample period. We then use the stock-split and dividend information from

²²Note that rather than absorbing the cost increase exclusively by reducing average output (expanding the spread and lowering transaction rates) she absorbs most of it by the price pressure instrument in addition with consumption smoothing through precautionary savings. This way she can keep consumption at higher levels than she would otherwise be able to do.

CRSP to remove these effects from the midquote prices in TAQ.²³ Stocks with an average share price of less than \$5 are removed from the sample, as are stocks with an average share price of more than \$1,000. The resulting sample comprises 697 common stocks. Stocks are sorted into quintiles based on market capitalization. Quintile 1 refers to the large-cap stocks and quintile 5 corresponds to the small-cap stocks. We further convert the NYSE specialist position which is in shares in the original database into dollars in order to facilitate comparisons across stocks. We multiply the position with the sample average price so as not to introduce daily price changes in the inventory variable. This would contaminate it as explanatory variable for the transitory price effect in the econometric model.

[insert Table 2]


Summary statistics. Table 2 presents the mean of various trading variables by size quintile. It also presents the within variation which is defined as the standard deviation of the data series after removing stock fixed effects. It thus gives a sense of the variable's variability through time. The statistics lead to a couple of observations. First, the average position of the specialist is positive and economically significant. For the large-cap stocks in Q1, for example, she maintains an almost half a million dollar average inventory position. This position is undoubtedly driven by a cost asymmetry across being in a long and a short position in the course of the trading process (something that is beyond the scope of this study).²⁴ The inventory position for the small-cap stocks in Q5 is \$77,900 and thus is considerably smaller than the average position for the large-caps. Second, the within standard deviation in specialist inventory is \$1.4 million, which seems substantial relative to her average position. It suggests that the specialist is an active intermediary in matching buyers and sellers through time. We will disaggregate this variation by year and by quintile in a later table. Third, the market capitalization is \$34 billion for Q1 stocks and declines to \$290 million for Q5 stocks. The effective half-spread, which is defined as the distance between a the transaction price and the prevailing midquote, is 8 basis points for Q1 stocks and 46 basis points for Q5 stocks. These numbers show that there is considerable heterogeneity across stocks.

²³The empirical analysis uses log prices and the series are adjusted to remove the effects from dividends and stocks splits where the adjustment leaves the sample mean of log price unchanged.

²⁴We interpret the zero position of our model as the deviation from the long-term optimal position of the specialist.

2.2 State space model to distinguish price pressures from efficient price innovations

The challenge to identify price pressures in real markets is that net order imbalance (the dual of the intermediary’s inventory change) might convey information as well as cause pressure which makes prices “overshoot.” The following classic graph based on block trades illustrates such overshooting:



source: Kraus and Stoll (1972, p.575)

To identify the pressure effect in the presence of an information effect, we use the state space approach of Menkveld, Koopman, and Lucas (2007) which models an observed, high-frequency price series as the sum of two unobserved series: a nonstationary efficient price series (“information”) and a stationary series that captures transitory price effects (“pressures”). In its simplest form, the model structure is:

$$p_t = m_t + s_t, \tag{7}$$

$$m_t = m_{t-1} + w_t, \tag{8}$$

where p_t is the observed price, m_t is the unobserved efficient price, s_t is the unobserved transitory price effect, and w_t is the innovation in the efficient price. s_t and w_t are mutually uncorrelated and normally distributed. It is immediate from the structure of the model that only draws on w_t affect the security’s price permanently; any draw on s_t is temporary as it affects prices only at a single point in time. The model is estimated with maximum likelihood where the likelihood is calculated by means of the Kalman filter.²⁵

²⁵The structure of the state space model is similar to the model that Roll (1984) proposes for transaction prices that bounce between the bid and the ask quote. In a way, the price pressures we are after are “bid” and “ask” quotes for net imbalance that the representative investor want the intermediary to take on.

We prefer a state space approach over other approaches (e.g., GMM or an ARIMA model) for several reasons. First, maximum likelihood estimation is asymptotically unbiased and efficient. Second, the model implies that the differenced series is an invertible MA(1) time series model which implies an infinite lag AR model. This is particularly cumbersome if the price series (or the inventory series that we will use as explanatory variable for s_t) have missing values. The Kalman filter ensures maximum efficiency as it deals optimally with missing values.²⁶ Third, after estimation, the Kalman smoother (essentially a backward recursion after a forward recursion with the Kalman filter) facilitates a series decomposition where at any point in time the efficient price and the transitory deviation are estimated using past and future prices. This allows for an in-sample decomposition of prices (as we will do to illustrate our results). We refer to Durbin and Koopman (2001) for an extensive treatment on the use of state space models to analyze economic time series.

In the remainder of this subsection we develop the general state space model to be taken to the data. We first develop the latent efficient price process and then the stationary price deviations that should capture the price pressures that we are after.

Unobserved efficient price process. We use the model to analyze daily midquote series by stock-year. We assume that for a single year the efficient price follows a geometric Brownian motion with a linear trend that captures the required return for the stock. For the remainder of the study we therefore consider log prices and we detrend the original series to remove the linear trend to focus on risk.²⁷ The efficient price series now is a martingale that consists of two components:

$$m_{it} = m_{i,t-1} + \beta_i \hat{\gamma}_t + w_{it}, \quad w_{it} = \kappa_i (I_{it}^{idio} - E_{t-1}[I_{it}^{idio}]) + u_{it} \quad (9)$$

where the subscript i indexes stocks, t indexes days, $\hat{\gamma}_t$ is a common factor innovation²⁸, w_{it} is the idiosyncratic innovation, $E_{t-1}[I_{it}^{idio}]$ is an AR(9) based forecast so that $(I_{it}^{idio} - E_{t-1}[I_{it}^{idio}])$ is a martingale that captures the surprise part of the net order imbalance that might contain information²⁹, and u_{it} is the stock-specific innovation orthogonal to the order imbalance innovation

²⁶It does not lose any information as it operates on the level series rather than the differenced series.

²⁷We prefer not to estimate the mean from the series as it is necessarily noisy in a single year. Instead, we use the year's riskfree rate and the standard compensation for risk by adding a risk premium that is based on the stock's beta and a market risk premium of 6%.

²⁸We find the common factor in midquote returns as a market cap weighted average of standardized stock returns (we standardize to control for heteroskedasticity). We then estimate an AR(4) on this common factor series and take the residual as the common factor innovation.

²⁹In the implementation, the net imbalance innovations are the residuals of a 9th order autoregressive model for inventory.

which is assumed to be a normally distributed white noise process. The decomposition of efficient price innovation in its stochastic components is important as only the idiosyncratic price changes w_{it} represent undiversifiable risk for the specialist.³⁰ The common factor risk is easily hedged by taking opposite position in highly liquid index products.

Unobserved transitory price deviations process. We propose the following process for the stationary price deviations:

$$s_{it} = \varphi_i s_{i,t-1} + \alpha_i I_{it}^{idio} + \beta_i^0 \hat{\gamma}_t + \dots + \beta_i^k \hat{\gamma}_{t-k} + \varepsilon_{it} \quad (10)$$

where I_{it}^{idio} is the end-of-day idiosyncratic inventory of the specialist³¹ and the error term ε_{it} now is assumed to be a normally distributed white noise process that is independent of w_{it} . The general specification nests the two models that we will take to the data. In an exploratory model without the inventory position we set α_i equal to zero. This allows us to test whether the transitory price effects are persistent at the daily frequency. In the main model where we allow these transitory price effects to be explained by inventory we set φ_i equal to zero as the inventory series now captures the persistency. The gamma terms enter the stationary price effect equation to capture a documented lagged adjustment to common factor innovation that is particularly prevalent for relatively inactively traded small-stocks (see, e.g., Campbell, Lo, and MacKinlay (1997)). In the proposed specification the beta coefficient in the efficient price process captures the long-term impact of a common factor shock on the price of the security and any lagged adjustment shows up through negative beta coefficients in the transitory price effects equation.

Observed price process. We close the econometric model with the observation equation:

$$p_{it} = m_{it} + s_{it}. \quad (11)$$

As we intend to analyze price pressures in the cross-section as well as in the time dimension we do all empirical analysis consistently by stock-year. To report the 697*12 stock-year results

³⁰The model explicitly specifies the idiosyncratic innovation as the sum of an imbalance innovation and the orthogonal part to ensure unbiased estimation of the model and, in particular, unbiased estimation of the conditional price pressure as the order imbalance innovation in period t correlates with inventory at time t (which is at the end of period t).

³¹We calculate the common factor as the market cap weighted average of the standardized inventory positions. We consider market-wide inventory shocks irrelevant for our purpose as portfolio shocks keep specialist firms unexposed to the idiosyncratic risk as market risk is easily hedged in highly liquid index securities.

we aggregate stocks into bins according to their size.³² We calculate means for each bin and the number of t -statistics that are outside of the 10% to 90% quantile interval. These t -statistics are available as supplementary material on the authors' websites. In the main tables we report the p -value of the meta statistic of the number of significant t -values in the size-year bin (and in the aggregation across bins).³³

Time series statistics as early evidence on price pressures. Before turning to the model estimates, some time series statistics are useful to test whether the price effects that we are after are borne out by the data.

[insert Table 3]

Table 3 reports first and second order autocorrelation of idiosyncratic midquote returns. The effects of contemporaneous and lagged adjustment to the common factor innovation are removed by regressing the midquote return on the common factor innovation up to four lags. The residuals serve as the idiosyncratic returns. Consistent with the individual stocks autocorrelation results in Campbell, Lo, and MacKinlay (1997) the average first order autocorrelation is negative in 64 of the 70 size-year bins and positive in the other bins. The low p -values indicate that the coefficient estimates are highly significant at conventional significance levels.³⁴ The negative first order autocorrelation is consistent with transitory price effects as the simple state space model of equations (7) and (8) implies a negative first-order correlation in midquote returns. The table further shows that the second order autocorrelation is also significantly negative which is an early indication that the simple model needs to be extended to allow the price deviations s_t to be positively autocorrelated. Price pressures appear to carry over days rather than being just an ultra-high-frequency intra-day phenomenon. This also implies that the unconditional deviations are potentially much larger relative to fundamental volatility than the simple first-order autocorrelations seem to suggest.³⁵

³²We keep the allocation across size quintiles fixed throughout the sample so as to ensure that bins are comparable across years.

³³The meta statistic is binomially distributed where the null probability that a t -value is statistically significant for a stock-year analysis is equal to the significance level of the t -test which is 20% in our case. Commonality across stocks is small in the empirical analysis focuses on idiosyncratic effects by removing a common factor in both the price and in the inventory series.

³⁴Overall, 3525 t -values are significant and 4836 t -values are insignificant. The sign of the significant t -values is primarily negative (2427 negative vs. 1098 positive) which shows that the negative means are statistically significant. We will omit such extensive statistical significance discussion in the remainder of the document and refer to results as being statistically significant if the meta statistic is significant and the sign of the majority of t -value is the same as the sign of the coefficient.

³⁵The simple model with an AR(1) process for the stationary term s_t implies a first-order autocorrelation in midquote returns of

$$ac1 = \frac{\frac{\varphi-1}{1+\varphi}\sigma_\varepsilon^2}{\frac{\sigma_\varepsilon^2}{1-\varphi^2} + \sigma_w^2} \quad (12)$$

Table 3 further reports the standard deviation of specialist inventory, its autocorrelation, and a cross-correlation with subsequent midquote returns as early evidence on the conjectured relation between transitory midquote deviations and inventory. The standard deviation of specialist end-of-day inventory varies is \$1.131 million for the large-cap stocks and monotonically decreases to \$165,000 for the small-cap stocks. It is relatively constant throughout time but tapers off in the last few years in the sample. The cross-sectional variation is undoubtedly due to a higher fundamental volatility and a smaller market (lower hedge value rate) for small-cap stocks. The economic model predicts that the intermediary shies away from frequent and large nonzero inventory position which explains why the inventory volatility is considerably smaller for small-cap stocks.

The inventory volatility is roughly twice the mean inventory across all quintiles. This is evidence of an active intermediation on a day over day level. In other words, inventory management is not a phenomenon that is restricted to the intraday ultra-high frequency level where intermediaries “go home flat.” As a matter of fact, the table further documents a significant first-order autocorrelation in inventories which show that these positions could last for multiple days. Again, there is considerable cross-sectional variation as the average autocorrelation for the large-cap stocks is 0.28 which monotonically increases to an average autocorrelation of 0.72 for small-cap stocks.³⁶ The specialist seems to trade out of most of an end-of-day position in the course of the next day for the large-cap stocks, whereas it take multiple days for the small-cap stocks. Finally, we calculate the correlation between today’s inventory position and tomorrow’s midquote return to verify whether the the two sets of results in the table can be reconciled. The last panel in the table shows that today’s inventory position correlates significantly with tomorrow’s midquote return. The positive signs are consistent with the economic model as the intermediary lowers the midquote on a long position (relative to the long-term average) which elicits an order imbalance that makes inventory mean-revert and, as a consequence, leads to lower midquote discounts. This creates a positive correlation between today’s position and tomorrow’s midquote return.

State space model estimates. We estimate two versions that are nested in the state space model defined by equation (9), (10), and (11). The first version does not yet include the idiosyncratic inventory position as explanatory variable and captures all potential persistence in an AR(1) process on the price deviation, i.e., it sets α_i equal to zero. The second version replaces

where φ is the autoregressive coefficient. The autocorrelation therefore becomes less negative if pressure becomes more persistent, i.e., if φ grows from zero to some positive number. Notice that in the limit ($\varphi \uparrow 1$) the autocorrelation is zero regardless of the magnitude of the transitory price effects.

³⁶These inventory autocorrelations are lower than the puzzlingly large autocorrelations found in NYSE data from the late 1980s and early 1990s in Madhavan and Smidt (1993) and Hasbrouck and Sofianos (1993).

the latent persistence in price deviations with the idiosyncratic inventory position relative to the long-term mean inventory, i.e., it sets φ_i equal to zero.

[insert Table 4]

Table 4 presents the model estimates without inventory as an explanatory variable for price pressure. It shows that these transitory price effects are large and persistent. In terms of size, the standard deviation of the error term in the stationary price pressure equation (10) is statistically significant and is 58 basis points for the large-cap stocks and 67 basis points for the small-cap stocks. These price pressures are persistent as the AR(1) coefficient are significant and range from 0.46 for the large-cap stocks to 0.39 for the small-cap stocks. With these numbers one can calculate the unconditional price pressure which is reported in Panel E and ranges from 83 basis points for small-cap stocks to 79 basis points for large-cap stocks. The table further reports the size of these pressures relative to the size of the innovation in efficient price through variance ratios. These ratios are 0.40 for the large-cap stocks and 0.31 for the small-cap stocks, respectively. They increase to 0.66 and 0.37, respectively, if only idiosyncratic innovations in the efficient price are considered. These ratios illustrate that these price pressures are a nonnegligible part of daily price volatility.

The estimates further confirm two well-documented empirical facts for the cross-section of stocks. The standard deviation of idiosyncratic daily midquote returns is 142 basis points for the large-cap stocks and increases monotonically to 211 basis points for the small-cap stocks. We further find significant evidence of lagged adjustment to common factor innovations predominantly for small-cap stocks. These results are reported in an extended version of the table that is included in the supplementary material.

[insert Figure 3]

Figure 3 illustrates the model's estimates for twenty trading days in a Q5 stock (PERMNO 635030) starting January 8, 2002. It exploits one attractive feature of the state space approach, which is that conditional on the model's parameter estimates the Kalman smoother generates estimates of the unobserved efficient price and price pressure processes conditional on all observations, i.e., it uses past and future observations to estimate the efficient price m_{it} and the temporary price deviation s_{it} at any date t in the sample. The first graph in Panel A plots the observed midquote along with the efficient price estimate and illustrates that the midquote can be far removed from the efficient price. The second graph further decomposes the efficient price into the common factor random walk, the idiosyncratic innovation random walk, and an integrated required return.

Panel B decomposes the stationary price deviation series into a delayed adjustment to the common factor (left) and the nonexplained remainder part (right). It also plots the idiosyncratic inventory as a deviation from its long-term mean on a second y-axis in the graph with the remainder part. The plot provides suggestive evidence that the temporary price deviations might indeed be the price discounts that a representative intermediary applies to mean-revert inventory. We turn to estimation of the model with inventory as explanatory factor to do proper empirical testing.

[insert Table 5]

Table 5 reports the estimates of the state space model with inventory as an explanatory variable for price deviations and their persistence. The key parameter α_i that measures the conditional price pressure has the conjectured sign and is highly significant as 4045 t -values are significantly negative, 4051 are insignificant, and only 265 are significantly positive. Prices are low when the intermediary is on a long position and high when she is on a short position relative to her long-term mean inventory. There is substantial cross-sectional variation in conditional price pressure as α_i is -0.02 for the large-cap stocks and -1.01 for the small-cap stocks. These numbers are economically significant as a \$1.131 million (one standard deviation) position change in specialist inventory creates a price pressure of $1131 \times 0.02 = 23$ basis points (cf. effective half-spread of 8 basis points, see Table 2). A similar position change in the small-cap stocks would create a price pressure of $1131 \times 1.01 = 1142$ basis point or 11.42%! The higher conditional price pressures are undoubtedly the result of higher fundamental volatility and a thinner market (lower hedge value rate) in the context of the economic model.

The table further shows that the average price pressure varies less in the cross-section than the conditional price pressure. By the logic of the economic model and as evidenced by a lower standard deviation of specialist inventory, the intermediary visits large inventory positions less often on higher fundamental volatility or when the market is thin. The average pressure which is measured as the conditional pressure times the standard deviation of inventory is 17 basis points for the large-cap stocks and 120 basis points for the low-cap stocks. The conditional pressure is roughly 7 times higher for the small-cap stocks relative to the large-cap stocks whereas the conditional pressure is 50 times higher. This is consistent with a dampening effect that dynamic inventory control has when fundamental volatility is increased or the market is thinner. The table also reports the size of these average price pressures relative to permanent volatility. The variance ratio is 0.02 for the large-cap stocks which indicates that these price pressures are small relative

to fundamental volatility. This is different for small-cap stocks where the ratio is 1.32: transitory volatility due to price pressures are larger than fundamental volatility.

The social cost due to unrealized hedge value is economically significant and varies considerably in the cross-section of stocks. The structural model suggests a proxy for this social cost based on the conditional price pressure, the variance of inventory, and the first order autocorrelation of inventory (see equation (6)). Panel H of Table 5 presents the cost estimate which is \$9.86 million per year for the largest-cap stocks and monotonically declines to \$0.67 million for the smallest-cap stocks. Cross-sectionally, these results are the exact opposite of the average price pressure results that increased with size. This is undoubtedly caused by the thickness of the market (hedge value rate) as it has an opposite effect on the size of average price pressure and the unrealized hedge value rate (see the change in λ in Table 1).

The aggregate magnitude of the social costs of unrealized hedge value is economically large. The average yearly costs in \$3.60 million per NYSE common stock. Multiplying this by the 697 stocks in the panel is \$2.5 billion per year. On average our balanced panel comprises less than half of the total common stock years. Assuming that social costs in the stocks which are not in the panel is the same as stocks in the panel would double the annual costs to \$5.0 billion.³⁷ Therefore, the total social costs for our 12-year sample period are estimate to be greater than \$50 billion. As noted earlier this is only one component of the inefficiency allocation of risk. The intermediary's cost of bearing risk is the other component which we are not able to measure.

2.3 Price pressure and NYSE market structure

When interpreting our results it is worth discussing the institutional structure of the specialist intermediary at the NYSE. The NYSE grants the specialists a central position in the trading process and imposes obligations upon the specialists. Panayides (2007) shows that the most significant obligation, the Price Continuity Rule which “requires the specialist to smooth transaction prices by providing extra liquidity as necessary to keep transaction price changes small,” is important at the transaction-level horizon. Panayides finds that the rule causes specialists to accumulate inventory to prevent “transaction prices from overshooting beyond their equilibrium levels.” This causes inventories to be positively associated with transitory price effects, the opposite of our

³⁷It is possible that the stocks that are not present for the entire sample period are smaller than stocks present throughout. On the other hand, the number of stocks listed on the NYSE is highest in 1999-2001 when the social costs of price pressure are substantially above the full sample average. Given that the goal is to calculate an approximate number we do not quantify either of these two effects.

relation between intermediary inventory and price pressure. Therefore, if the Price Continuity Rule manifests itself at a daily frequency it causes underestimation of price pressures associated with inventory.

While the NYSE designates a single intermediary, it is possible for other investors to compete with the specialist by placing limit orders to supply liquidity. Such a possibility is especially important given the NYSE's recent market structure changes (after our sample period) which resulted in a reduced role for the specialist (Hendershott and Moulton (2007)). Many liquidity suppliers could reduce the width of the bid-ask spread and could also reduce the social costs of intermediaries bearing risk. The most efficient manner to share risk is for the inventory to be immediately and equally shared across all liquidity suppliers, leading to perfectly correlated positions. How would this affect our estimates of price pressure? First, the conditional price pressure per unit of inventory should be adjusted by the specialist's fraction of inventory, e.g., if the specialist carries one half of the total inventory then the conditional price pressure should be multiplied by one half. The average price pressure is unaffected by additional liquidity suppliers because the standard deviation of inventory is adjusted by the reciprocal of the adjustment to the conditional price pressure. The unrealized hedge value is underestimated by the same magnitude as the underestimation of the standard deviation of inventories (see equation (6)).

2.4 Correlations of price pressure and other variables

[insert Table 6]

Table 6 presents meta correlations on the panel dataset of 12*697 stock-year estimates. It characterizes these correlations separately in the cross-section and in the time dimension by presenting between and within correlations in addition to the overall correlations. We focus on the within correlations as they effectively take out a stock fixed effect. It allows for a focus on fundamental risk and the intermediary's exposure to it which should capture the behavior of the intermediary when confronted different levels of fundamental risk (high risk years vs. low risk years).

The within correlations in the table suggest that at times of elevated volatility the intermediary does not seem to mean-revert quicker, which is what the model predicts she should do. As a matter of fact, she seems to be visiting extreme inventory positions more often rather than less often as suggested by the significantly positive within correlation of 0.13 between inventory and fundamental volatility. In addition, the within correlation of fundamental volatility with conditional

price pressure, 0.03, is lower than its within correlation with the average price pressure, 0.15. This result suggests that dynamic inventory control has an amplifying effect on conditional price pressure. The economic model can only generate these effects if the price elasticity of the net transaction rate drops at times of high fundamental risk. That is, the ability of the intermediary to mean-revert her inventory and quickly share her idiosyncratic risk is reduced which creates a large deadweight loss of inefficient risk-sharing. Order imbalance appears to be less sensitive to prices at times of elevated risk. Also, the social cost of unrealized hedge value is higher at times of elevated fundamental volatility. Perhaps this is the liquidity risk that economists have long been searching for.

3 Conclusion

We construct a theoretical model to understand and characterize the effects of price pressure—the deviation of prices from fundamental values due to the inventory risks born by an intermediary providing liquidity to asynchronously arriving investors with idiosyncratic hedging needs. The structure of the model allows for estimation of the social costs of price pressure due to a reduction in investors’ realized hedging value. Empirically we estimate price pressure using 12 years of NYSE intermediary data. We find:

1. A \$100,000 inventory shock causes price pressure of 1.01% for the small-capitalization stocks and 0.02% for the large-cap stocks. Price pressure conditional on inventory reduces the variability of the intermediaries’ inventory.
2. The daily transitory volatility in stock returns due to price pressure (a measure of average price pressure) is large: 1.20% and 0.17% for small and large stocks, respectively. For small stocks the transitory volatility due to price pressure is greater than the permanent (random-walk or efficient price) volatility.
3. A model-based estimate of price pressure’s annual social cost from lower realized hedging gains is \$9.86 million and \$0.67 million for large and small stocks, respectively. The aggregate social cost is estimated to be greater than \$50 billion for NYSE common stocks for our sample period. The yet unmeasured social costs of inefficient allocation from the intermediary costs of bearing risk make the total social costs related to price pressure even higher.

The significant social costs of price pressure suggest that a goal of financial market regulation should be to mitigate price pressure. One way to do this is by increasing capital for intermediation

as the greater the risk bearing capacity of the intermediaries the smaller the price pressure. Another approach would be to lower costs for investors to monitor the market. This would lead to investor trading being more responsive to price pressures, reducing the duration of price pressure by allowing intermediaries to mean-revert their inventories more quickly.

Appendix I: Details on the calibration and the solution method.

Base case model calibration. We numerically solve the intermediary’s dynamic program which is summarized by equation (1). The calibration follows, to the extent possible, the base case parametrization proposed by Ho and Stoll (1981, p.67).

1. The intermediary’s coefficient of relative risk aversion (ρ) is 2.
2. The time interval length is one day.
3. The arrival rate function is $g(x) = \lambda_0(1 - \theta x)$ where $\lambda_0 = 1$ and $\theta = 5$ so that each day the representative investor arrives and considers a trade. The transaction size is 10 units.³⁸
4. The stochastic dividend risk on one inventory unit equals $\sigma = \sqrt{(0.50/200)} = 0.05$ where one year contains 200 trading days.
5. The intermediary’s discount rate β is 0.9 to capture the two week horizon Stoll proposes for the intermediary.³⁹
6. The intermediary earns a gross daily riskfree rate of $R = 1 + 0.10/200$ on her savings for the next period.

We then discretize the set of inventory-wealth states (i, w) by choosing a grid with carefully selected bounds for both the inventory and the wealth dimension of the state. We choose upperbounds for inventory and wealth so that the state-space become a finite set. We set inventory bounds equal to -50 and 50 and verify that ex-post these inventory states indeed appear to be nonbinding constraints.⁴⁰ The step size for the inventory dimension is governed by the transaction size 10. We choose a wealth upperbound of 10 where we again verify that the limit does not bind endogenously. We set the wealth lower bound equal to 0 which is an absorbing bankruptcy state from which point on consumption is arbitrarily small.⁴¹ The wealth step size is constrained by the stochastic dividend risk which is implemented as a discrete random variable where the outcome is plus or minus σ with equal probability. As the inventory dimension only contains multiples of 10, the stochastic dividend commands a (maximum) wealth step size of $10 \cdot 0.05 = 0.5$.⁴²

We then set up the discretized version of the Bellman equation which is a trivial extension of the system defined by equation (4). The step size on the wealth dimension necessarily discretizes the bid and ask prices that the intermediary can choose, but we can get as close as one might desire to a continuum of prices by

³⁸Ho and Stoll (1981) use the same θ but set $\lambda_0 = 2,000$ as the annual arrival rate. To match their arrival rate we set the transaction size equal to 10 with the interpretation that one successful daily transaction means $2,000/200=10$ intraday transactions of one inventory unit.

³⁹A 0.9 discount rate implies that roughly two-thirds of the total discounted value of a dividend stream of one falls into the first two weeks.

⁴⁰One has to be careful here, as in addition to it being a very risky state, the intermediary also internalizes that she can only exit this state on one side (i.e. less inventory) which limits her earning potential. Ideally, we would like to let the inventory upperbound I tend to infinity to rule out this nonrisk explanation. We “simulate” such limit numerically and verify that the solution does not change by taking larger upperbounds.

⁴¹CRRA utility satisfies the Inada conditions so that intermediaries will steer clear of this state.

⁴²Technically, the intermediary could still end up strictly outside of wealth interval if she gets a large enough draw on the stochastic dividend. In the implementation, we steer her to the lower- or upperbound of wealth in the cases that this happens. This does not affect our solution as the intermediary endogenously chooses to stay away from those states.

shrinking the wealth step size at the cost of computational speed. After discretizing, we solve the dynamic program numerically by iterating on the Bellman equation where we rely on the contraction property of the recursive mapping to achieve convergence.

Bellman equations for the interior of the inventory-wealth state domain. We arrive at the following discrete state dynamic program

$$\begin{aligned}
V(i, w) &= \max_{\tilde{w}_{iw}, a_{iw}, b_{iw}} u(w - R^{-1}\tilde{w}_{iw}) \\
&\quad + \beta E_{\Delta m} [V(i, \tilde{w} + i\Delta m)](1 - g(a_{iw}) - g(b_{iw})) \\
&\quad + \beta E_{\Delta m} [V(i + 1, \tilde{w} + i\Delta m + b_{iw}\Delta i) + \Delta m]g(b_{iw}) \\
&\quad + \beta E_{\Delta m} [V(i - 1, \tilde{w} + i\Delta m - a_{iw}\Delta i) + \Delta m]g(a_{iw}) \\
g(x) &= 1 - \theta x \quad (\text{probability of a transaction}) \\
&\quad i \in \{-50, -40, \dots, 50\} \quad w \in \{0, 0.5, \dots, 10\}
\end{aligned} \tag{13}$$

where the three expected value terms correspond to no-arrival, seller-arrival, and buyer-arrival, respectively. Equation (13) defines the Bellman equations for the nonboundary inventory-wealth states. But, before turning to the boundary Bellman equations, we need to make sure that we stay on the inventory-wealth grid in the iterations. This motivates a binomial distribution for stochastic dividend

$$P[\Delta m = -\sigma] = P[\Delta m = +\sigma] = 0.5 \tag{14}$$

and restricts the set of admissible controls to

$$a_{iw}, b_{iw} \in \{0, 0.05, 0.10, \dots\} \text{ s.t. } g(a_{iw}), g(b_{iw}), 1 - g(a_{iw}) - g(b_{iw}) \geq 0 \tag{15}$$

$$\tilde{w}_{iw} \in \{0, 0.5, \dots, \min(10, 0.5[2Rw])\} \tag{16}$$

The upperbound for end-of-period wealth \tilde{w}_{iw} prevents negative current period consumption.

Bellman equations for the boundary of the inventory-wealth state domain. As for the boundary inventory states, we adjust the Bellman equation defined in equation (13) by reducing the set of admissible controls. For the maximum inventory state, for example, we restrict the bid price to $b_{Iw} = \theta^{-1}$ (which effectively sets the seller arrival rate to zero). As for the boundary wealth states, in the zero wealth state we keep the intermediary on a low enough consumption level relative to her earning power in the nonzero wealth state to make her endogenously choose to stay away from the bankruptcy state. In states close to the maximum wealth state, we redirect her to the maximum wealth state if she enjoys a stochastic dividend that would make her transit to a larger wealth than the maximum wealth. One interpretation is that the government taxes “excessive wealth” away in these high wealth states. We reiterate that these assumptions on the maximum inventory or maximum wealth Bellman equations are inconsequential as we choose the maximum inventory and maximum wealth boundary large enough so that the intermediary endogenously chooses never to get near to these states (see discussion footnote 42).

From this point on we follow Ljungqvist and Sargent (2004, §4.2, p.95), i.e. we iterate on the Bellman equation and the contraction property of the recursive mapping guarantees convergence to the unique solution.

Appendix II: Likelihood optimization

We find the maximum likelihood of the most general state space model in two steps to minimize the probability of ending at a local maximum.

1. We perform an OLS regression of the first difference of the log price series on $\hat{\gamma}_t$ to find the starting value for the various β_i s that capture a potential lagged adjustment to a common factor. We fix β_i throughout. We verify its optimality by running the most general model with all parameters as free parameters.

2. We then optimize the likelihood for a fixed pair $(\sigma(\varepsilon), \varphi)$ which governs the (unobserved stationary component of pricing error). We choose the pair from a ten by ten grid of potential parameter values where φ ranges from -0.8 to 0.8 and $\sigma(\varepsilon)$ ranges from zero to a stock-specific upperbound that is calculated assuming that 80% of a stock's unconditional volatility is pricing error.⁴³

This procedure proves numerically stable as we have strong convergence in the likelihood optimization for all of our stock-year samples.

References

- Acharya, Viral V., Hyun-Song Shin, and Tanju Yorulmazer. 2008. "Fire-Sale FDI." Manuscript, NYU.
- Amihud, Yakov, and Haim Mendelson. 1980. "Dealership Market: Market-Making with Inventory." *Journal of Financial Economics* 8:31–53.
- Brunnermeier, Markus K., and Lasse Heje Pedersen. 2008. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* (forthcoming).
- Campbell, John Y., Sanford J. Grossman, and Jiang Wang. 1993. "Trading Volume and Serial Correlation in Stocks Returns." *The Quarterly Journal of Economics* 108:905–939.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay. 1997. *The Econometrics of Financial Markets*. Princeton, New Jersey: Princeton University Press.
- Coval, Joshua, and Erik Stafford. 2007. "Asset Fire Sales (and Purchases) in Equity Markets." *Journal of Financial Economics* 86:479–512.
- Durbin, Jim, and Siem Jan Koopman. 2001. *Time Series Analysis by State Space Models*. Oxford: Oxford University Press.
- Garman, Mark B. 1976. "Market Microstructure." *Journal of Financial Economics* 3:257–275.
- Greenwood, Robin. 2005. "Short- and Long-Term Demand Curves for Stocks: Theory and Evidence on the Dynamics of Arbitrage." *Journal of Financial Economics* 75:607–649.
- Grossman, Sanford J., and Merton H. Miller. 1988. "Liquidity and Market Structure." *Journal of Finance* 43:617–633.
- Hansch, Oliver, Narayan Y. Naik, and S. "Vish" Viswanathan. 1998. "Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange." *Journal of Finance* 53:1623–1656.
- Harris, Lawrence, and Eitan Gurel. 1986. "Price and Volume Effects Associated with Changes in the S&P 500 List: New Evidence for the Existence of Price Pressures." *Journal of Finance* 41:815–829.
- Hasbrouck, Joel. 1988. "Trades, Quotes, Inventories, and Information." *Journal of Financial Economics* 22:229–252.
- Hasbrouck, Joel, and George Sofianos. 1993. "The Trades of Market Makers: An Empirical Analysis of NYSE Specialists." *Journal of Finance* 48:1565–1593.

⁴³We set the φ upperbound to a value strictly less than one so as to keep the pricing error process stationary. If we use an upperbound close to one the likelihood optimization becomes numerically unstable as the price series potentially gets two random walks to load on.

- Hendershott, Terrence, and Pamela C. Moulton. 2007. "Speed and Stock Market Quality: The NYSEs Hybrid." Manuscript, UC Berkeley.
- Hendershott, Terrence, and Mark S. Seasholes. 2007. "Market Maker Inventories and Stock Prices." *American Economic Review* 97:210–214.
- Ho, Thomas, and Richard Macris. 1984. "Dealer Bid-Ask Quotes and Transaction Prices: An Empirical Study of Some AMEX Options." *Journal of Finance* 39:23–45.
- Ho, Thomas, and Hans R. Stoll. 1981. "Optimal Dealer Pricing under Transaction Cost and Return Uncertainty." *Journal of Financial Economics* 9:47–73.
- Huggett, Mark. 1993. "The Risk-Free Rate in Heterogeneous Agent Incomplete Insurance Economies." *Journal of Economic Dynamics and Control* 17:953–969.
- Kraus, Alan, and Hans Stoll. 1972. "Price Impacts of Block Trading on the New York Stock Exchange." *Journal of Finance* 27:569–588.
- Lehmann, Bruce. 1990. "Fads, Martingales, and Market Efficiency." *The Quarterly Journal of Economics* 105:1–28.
- Ljungqvist, Lars, and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory*. Cambridge Massachusetts: The MIT Press.
- Madhavan, Ananth, and Seymour Smidt. 1991. "A Bayesian Model of Intraday Specialist Pricing." *Journal of Financial Economics* 30:99–134.
- . 1993. "An Analysis of Changes in Specialist Inventories and Quotations." *Journal of Finance* 48:1595–1628.
- Menkveld, Albert J., Siem Jan Koopman, and André Lucas. 2007. "Modelling Round-the-Clock Price Discovery for Cross-Listed Stocks using State Space Methods." *Journal of Business & Economic Statistics* 25:213–225.
- Mildenstein, Eckart, and Harold Schlee. 1983. "The Optimal Pricing Policy of a Monopolistic Market Maker in the Equity Market." *Journal of Finance* 38:218–231.
- Naik, Narayan Y., and Pradeep K. Yadav. 2003. "Do Dealer Firms Manage Inventory on a Stock-by-Stock or a Portfolio Basis?" *Journal of Financial Economics* 69:325–353.
- Panayides, M.A. 2007. "Affirmative Obligations and Market Making with Inventory." *Journal of Financial Economics* 86:513–542.
- Reiss, Peter, and Ingrid Werner. 1998. "Does Risk Sharing Motivate Interdealer Trading?" *Journal of Finance* 53:1657–1704.
- Roll, Richard. 1984. "A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market." *Journal of Finance* 39:1127–1139.
- Stoll, Hans R. 1978. "The Supply of Dealer Services in Securities Markets." *Journal of Finance* 33:1133–1151.
- Townsend, Robert M. 1978. "Intermediation with Costly Bilateral Exchange." *Review of Economic Studies* 45:417–425.
- Zabel, Edward. 1981. "Competitive Price Adjustment without Market Clearing." *Econometrica* 49:1201–1221.

Table 1: Numerical solutions and comparative statics for average and conditional price pressure and social cost rate

This table presents based on the numerical solutions and comparative statics for the dynamic program of a (representative) risk-averse intermediary who supplies liquidity and maximizes utility from consumption. The average price pressure is defined as the standard deviation of the distance between the midquote and a security's fundamental value. The conditional price pressure is defined as the (population) regression coefficient of a univariate regression of the midquote on the intermediary's inventory position. It compares the base case with three alternatives where the friction is increased by either (i) reducing the price-elasticity of arrival rates, (ii) reducing the autonomous arrival rate, or (iii) increasing fundamental volatility.

	base case $\theta=5$ $\lambda=0.2$ $\sigma(w)=0.05$	higher fun- damental volatility $\theta=5$ $\lambda=0.2$ $\sigma(w)=0.0625$	change ^a	lower price- zero arrival rate $\theta=5$ $\lambda=0.15$ $\sigma(w)=0.05$	change ^a	arrival rates less elastic $\theta=3.75$ $\lambda=0.2$ $\sigma(w)=0.05$	change ^a
<i>Panel A: Price pressure and inventory results</i>							
size regr coef midquote on inventory (“ $ \alpha $ ”)	0.0038	0.0043	14%	0.0044	16%	0.0057	53%
st dev inventory	11.2	10.4	-8%	10.3	-8%	12.0	7%
st dev price pressure ^b	0.042	0.045	5%	0.045	7%	0.069	64%
autocorr inventory	0.62	0.57	-8%	0.67	8%	0.79	27%
<i>Panel B: Social costs and intermediary performance measures</i>							
social cost rate unrealized hedge value	0.015	0.019	27%	0.014	-5%	0.010	-32%
proxy social cost rate unrealized hedge value	0.009	0.010	10%	0.008	-15%	0.009	-4%
average cash flow rate	0.809	0.785	-3%	0.582	-28%	0.815	1%
constant cash flow rate that yields same utility	0.707	0.651	-8%	0.434	-39%	0.600	-15%
social cost rate of the intermediary bearing risk	0.102	0.134	32%	0.149	45%	0.215	111%
expected consumption	0.81	0.79	-3%	0.58	-28%	0.81	1%
expected wealth	2.54	3.14	24%	3.42	34%	4.74	87%
probability of a transaction	0.94	0.92	-3%	0.69	-27%	0.59	-37%
average bid ask spread	0.21	0.22	3%	0.22	2%	0.34	60%

^a: Percentage change relative to base case.

^b: The price pressure deviation equals the standard deviation of the midquote in the model as the fundamental value is fixed at zero.

Table 2: Summary statistics

This table presents summary statistics on the dataset, which combines CRSP, NYSE's Trade and Quotes (TAQ) and a proprietary NYSE dataset. It is a balanced panel that contains daily observations on 697 NYSE common stocks from January 1994 through December 2005. Stocks are sorted into quintiles based on market capitalization, where quintile 1 contains large-cap stocks.

variable	description (units)	source	mean Q1	mean Q2	mean Q3	mean Q4	mean Q5	st. dev. wi- thin ^a
<i>midquote_{it}</i>	closing midquote, div/split adjusted ^b (\$)	NYSE	53.76	44.52	36.65	28.58	19.21	22.20
<i>invent_shares_{it}</i>	specialist inventory at the close (1,000 shares)	NYSE	8.19	5.55	4.33	3.23	5.39	34.19
<i>invent_dollar_{it}</i>	specialist inventory at the close ^b (\$1,000)	NYSE/CRSP	412.65	168.95	129.48	75.44	77.90	1,383.43
<i>shares_outst_{it}</i>	shares outstanding (million)	CRSP	729.92	157.74	70.08	36.26	18.73	283.75
<i>market_cap_{it}</i>	shares outstanding times price (\$billion)	CRSP	34.29	5.34	2.06	0.88	0.29	11.57
<i>espread_{it}</i>	share-volume-weighted effective half spread (bps)	TAQ	8.41	12.46	16.50	24.60	46.12	24.20
<i>dollar_volume_{it}</i>	average daily volume (\$million)	TAQ	88.21	23.44	10.13	3.63	0.99	42.31
<i>specialist_particip_{it}</i>	specialist participation rate (%)	NYSE	12.31	12.73	14.10	16.58	20.87	8.30
#observations: 697*3,018 (stock*day)								

^a: Based on the deviations from time means i.e. $x_{it}^* = x_{it} - \bar{x}_i$.

^b: We adjust all price series to account for stock splits and dividends.

Table 3: Specialist inventory mean reversion estimates by year and size quintile

This table implements an estimation strategy to capture dynamics between log price changes and specialist inventory. We explicitly recognize that both series contains a common factor γ_t , which we estimate as the cross-sectional mean each day of the standardized series. We refer to this estimate as $\hat{\gamma}_t$. As an example we describe the procedure for specialist inventory. For each stock, each year we multiply specialist inventory by the average price (to get a dollar amount to facilitate comparison across stocks) and then regress it on $\hat{\gamma}_t$. We save the residuals as the idiosyncratic component of specialist inventory in dollar, i.e. I_{it}^{idio} . We perform the following regressions by size quintile (Q1 contains the largest stocks) and by year:

$$y_{it} = \alpha_i + \beta_i x_{it} + \varepsilon_{it}$$

The table reports p -values in brackets. These p -values are based on a test statistic that counts the number of significant t -values across all stock-year ML estimates in the bin. The test statistic is binomially distributed under the null (we use the 0.1 and 0.9 quantiles in the t -test).

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel A: Autocorrelation 1st lag coef log price change ($x_{it} = y_{i,t-1}$)</i>													
Q1	-0.03 (0.000)	-0.03 (0.000)	-0.05 (0.000)	-0.08 (0.000)	-0.02 (0.000)	-0.00 (0.005)	-0.01 (0.003)	-0.01 (0.001)	-0.05 (0.000)	-0.00 (0.005)	0.00 (0.005)	-0.02 (0.026)	-0.03 (0.000)
Q2	-0.02 (0.000)	-0.02 (0.000)	-0.03 (0.000)	-0.07 (0.000)	-0.01 (0.000)	-0.02 (0.000)	-0.04 (0.000)	-0.00 (0.000)	-0.06 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.03 (0.000)	-0.03 (0.000)
Q3	-0.00 (0.000)	0.00 (0.000)	-0.02 (0.000)	-0.06 (0.000)	-0.01 (0.000)	-0.04 (0.000)	-0.06 (0.000)	-0.02 (0.000)	-0.06 (0.000)	-0.04 (0.000)	-0.03 (0.000)	-0.03 (0.001)	-0.03 (0.000)
Q4	-0.01 (0.000)	-0.02 (0.000)	-0.03 (0.000)	-0.05 (0.000)	-0.00 (0.000)	-0.03 (0.000)	-0.07 (0.000)	-0.04 (0.000)	-0.08 (0.000)	-0.06 (0.000)	-0.08 (0.000)	-0.04 (0.000)	-0.04 (0.000)
Q5	-0.05 (0.000)	-0.04 (0.000)	-0.02 (0.000)	0.00 (0.000)	0.04 (0.000)	0.01 (0.000)	-0.03 (0.000)	0.00 (0.000)	-0.02 (0.000)	-0.01 (0.000)	-0.04 (0.000)	-0.02 (0.000)	-0.02 (0.000)
all	-0.02 (0.000)	-0.02 (0.000)	-0.03 (0.000)	-0.05 (0.000)	-0.00 (0.000)	-0.02 (0.000)	-0.04 (0.000)	-0.01 (0.000)	-0.05 (0.000)	-0.03 (0.000)	-0.03 (0.000)	-0.03 (0.000)	-0.03 (0.000)
<i>Panel B: Autocorrelation 2nd lag coef log price change ($x_{it} = y_{i,t-2}$)</i>													
Q1	-0.03 (0.001)	-0.05 (0.000)	-0.04 (0.000)	-0.02 (0.094)	-0.03 (0.135)	-0.02 (0.001)	-0.06 (0.000)	-0.05 (0.000)	-0.01 (0.026)	-0.01 (0.249)	-0.02 (0.015)	-0.01 (0.135)	-0.03 (0.000)
Q2	-0.03 (0.001)	-0.03 (0.003)	-0.03 (0.009)	-0.01 (0.320)	-0.03 (0.001)	-0.01 (0.249)	-0.04 (0.001)	-0.02 (0.000)	-0.01 (0.015)	-0.02 (0.001)	-0.00 (0.009)	-0.02 (0.649)	-0.02 (0.000)
Q3	-0.02 (0.584)	-0.02 (0.070)	-0.02 (0.045)	-0.01 (0.664)	-0.01 (0.584)	-0.01 (0.029)	-0.04 (0.000)	-0.02 (0.017)	-0.00 (0.010)	-0.00 (0.029)	-0.01 (0.029)	-0.01 (0.001)	-0.02 (0.000)
Q4	-0.01 (0.001)	-0.02 (0.249)	-0.02 (0.041)	-0.03 (0.041)	-0.02 (0.320)	-0.00 (0.063)	-0.03 (0.009)	-0.02 (0.094)	0.00 (0.003)	-0.01 (0.003)	0.01 (0.003)	-0.00 (0.001)	-0.01 (0.000)
Q5	-0.01 (0.029)	-0.01 (0.000)	-0.01 (0.017)	-0.00 (0.017)	-0.00 (0.000)	0.00 (0.010)	-0.01 (0.500)	-0.00 (0.000)	-0.01 (0.017)	-0.01 (0.009)	0.00 (0.087)	0.00 (0.070)	-0.01 (0.000)
all	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.01 (0.010)	-0.02 (0.000)	-0.01 (0.000)	-0.04 (0.000)	-0.02 (0.000)	-0.00 (0.000)	-0.01 (0.000)	-0.00 (0.000)	-0.01 (0.000)	-0.02 (0.000)
<i>Panel C: Standard deviation of idiosyncratic component specialist inventory I_{it}^{idio}</i>													
Q1	691 (0.000)	968 (0.000)	813 (0.000)	964 (0.000)	1126 (0.000)	1336 (0.000)	1344 (0.000)	1489 (0.000)	1472 (0.000)	1122 (0.000)	1119 (0.000)	1128 (0.000)	1131 (0.000)
Q2	472 (0.000)	510 (0.000)	488 (0.000)	524 (0.000)	530 (0.000)	695 (0.000)	819 (0.000)	647 (0.000)	448 (0.000)	391 (0.000)	441 (0.000)	400 (0.000)	530 (0.000)
Q3	374 (0.000)	429 (0.000)	383 (0.000)	372 (0.000)	430 (0.000)	452 (0.000)	668 (0.000)	437 (0.000)	293 (0.000)	242 (0.000)	266 (0.000)	271 (0.000)	385 (0.000)
Q4	226 (0.000)	254 (0.000)	255 (0.000)	261 (0.000)	291 (0.000)	315 (0.000)	320 (0.000)	333 (0.000)	229 (0.000)	163 (0.000)	145 (0.000)	147 (0.000)	245 (0.000)
Q5	167 (0.000)	159 (0.000)	167 (0.000)	234 (0.000)	204 (0.000)	223 (0.000)	210 (0.000)	186 (0.000)	129 (0.000)	108 (0.000)	95 (0.000)	95 (0.000)	165 (0.000)
all	386 (0.000)	464 (0.000)	421 (0.000)	471 (0.000)	516 (0.000)	604 (0.000)	672 (0.000)	619 (0.000)	514 (0.000)	405 (0.000)	413 (0.000)	408 (0.000)	491 (0.000)
<i>Panel D: AR coef estimates idiosyncratic component specialist inventory I_{it}^{idio} ($x_{it} = y_{i,t-1}$)</i>													
Q1	0.28 (0.000)	0.27 (0.000)	0.26 (0.000)	0.22 (0.000)	0.25 (0.000)	0.27 (0.000)	0.28 (0.000)	0.29 (0.000)	0.28 (0.000)	0.34 (0.000)	0.37 (0.000)	0.25 (0.000)	0.28 (0.000)
Q2	0.47 (0.000)	0.46 (0.000)	0.44 (0.000)	0.38 (0.000)	0.35 (0.000)	0.34 (0.000)	0.36 (0.000)	0.32 (0.000)	0.25 (0.000)	0.28 (0.000)	0.33 (0.000)	0.25 (0.000)	0.35 (0.000)
Q3	0.59 (0.000)	0.59 (0.000)	0.56 (0.000)	0.51 (0.000)	0.49 (0.000)	0.45 (0.000)	0.41 (0.000)	0.41 (0.000)	0.30 (0.000)	0.31 (0.000)	0.34 (0.000)	0.24 (0.000)	0.43 (0.000)
Q4	0.74 (0.000)	0.73 (0.000)	0.71 (0.000)	0.66 (0.000)	0.63 (0.000)	0.63 (0.000)	0.59 (0.000)	0.57 (0.000)	0.40 (0.000)	0.38 (0.000)	0.36 (0.000)	0.29 (0.000)	0.56 (0.000)
Q5	0.82 (0.000)	0.80 (0.000)	0.80 (0.000)	0.77 (0.000)	0.78 (0.000)	0.79 (0.000)	0.77 (0.000)	0.76 (0.000)	0.66 (0.000)	0.61 (0.000)	0.57 (0.000)	0.51 (0.000)	0.72 (0.000)
all	0.58 (0.000)	0.57 (0.000)	0.56 (0.000)	0.51 (0.000)	0.50 (0.000)	0.50 (0.000)	0.48 (0.000)	0.47 (0.000)	0.38 (0.000)	0.38 (0.000)	0.39 (0.000)	0.31 (0.000)	0.47 (0.000)

<continued on next page>

<continued from previous page>

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel E: Regression coefficient log price change on lagged idiosyncratic component specialist inventory $I_{i,t-1}^{idio}$</i>													
Q1	0.01 (0.005)	0.01 (0.000)	0.01 (0.009)	0.01 (0.000)	0.01 (0.005)	0.01 (0.135)	0.01 (0.005)	0.01 (0.005)	0.01 (0.001)	0.00 (0.249)	0.00 (0.063)	0.00 (0.135)	0.01 (0.000)
Q2	0.01 (0.041)	0.01 (0.009)	0.01 (0.015)	0.02 (0.000)	0.01 (0.041)	0.02 (0.000)	0.02 (0.001)	0.02 (0.041)	0.02 (0.000)	0.02 (0.135)	0.01 (0.135)	0.01 (0.063)	0.02 (0.000)
Q3	0.02 (0.010)	0.02 (0.001)	0.02 (0.045)	0.02 (0.000)	0.02 (0.145)	0.03 (0.001)	0.05 (0.000)	0.01 (0.000)	0.05 (0.000)	0.03 (0.017)	0.03 (0.102)	0.03 (0.199)	0.03 (0.000)
Q4	0.03 (0.063)	0.02 (0.135)	0.04 (0.000)	0.05 (0.000)	0.04 (0.000)	0.06 (0.000)	0.08 (0.000)	0.05 (0.000)	0.11 (0.000)	0.05 (0.026)	0.08 (0.000)	0.06 (0.026)	0.06 (0.000)
Q5	0.08 (0.000)	0.08 (0.199)	0.03 (0.010)	0.02 (0.017)	0.06 (0.006)	0.06 (0.416)	0.14 (0.000)	0.10 (0.000)	0.15 (0.000)	0.08 (0.009)	0.12 (0.008)	0.13 (0.500)	0.09 (0.000)
all	0.03 (0.000)	0.03 (0.000)	0.02 (0.000)	0.02 (0.000)	0.03 (0.000)	0.03 (0.000)	0.06 (0.000)	0.04 (0.000)	0.07 (0.000)	0.04 (0.000)	0.05 (0.000)	0.05 (0.008)	0.04 (0.000)

*/**: Significant at a 95%/99% level.

Table 4: State space model estimates by year and size quintile ($\alpha_i = \kappa_i = 0$)

This table estimates the following state space model for a latent efficient price and an observed end-of-day midquote:

$$\begin{aligned}
 \text{(observed price)} \quad p_{it} &= m_{it} + s_{it} \\
 \text{(unobserved efficient price)} \quad m_{it} &= m_{i,t-1} + \beta_i \hat{\gamma}_t + w_{it} \quad w_{it} = \kappa_i (I_{it}^{idio} - E_{t-1}[I_{it}^{idio}]) + u_{it} \\
 \text{(unobserved transitory price deviation)} \quad s_{it} &= \varphi_i s_{i,t-1} + \alpha_i I_{it}^{idio} + \beta_i^0 \hat{\gamma}_t + \dots + \beta_i^3 \hat{\gamma}_{t-3} + \varepsilon_{it}
 \end{aligned}$$

where i indexes over stocks and t indexes over days, m_{it} is the end-of-day unobserved efficient price (“state”), $\hat{\gamma}_t$ is a midquote return common factor which is the cross-sectional average of the standardized midquote return series which has been filtered with an AR(4) model to remove intertemporal dynamics, p_{it} is end-of-day observed midquote, I_{it}^{idio} is the idiosyncratic part of the specialist end-of-day USD inventory that remains after removing a common factor across specialist inventories (we use the average price across the entire period to convert inventory in shares to USD), $E_{t-1}[I_{it}^{idio}]$ is an AR(9) based forecast so that $(I_{it}^{idio} - E_{t-1}[I_{it}^{idio}])$ is a martingale that captures the surprise part of the net imbalance, β_i^j captures potential “overreaction” or lagged adjustment to common factor innovations, and u_{it} and ν_{it} are mutually independent i.i.d. error terms. We report maximum likelihood estimates where we assume normality for the error terms. We implement our optimization in ox with ssfpack routines where we use the Kalman filter to evaluate the likelihood (see Koopman, Shephard, and Doornik (1999)). The table reports p -values in brackets. These p -values are based on a test statistic that counts the number of significant t -values across all stock-year ML estimates in the bin. The test statistic is binomially distributed under the null (we use the 0.1 and 0.9 quantiles in the t -test).

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel A: α_i estimates</i>													
Q1	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Q2	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Q3	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Q4	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Q5	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
all	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
<i>Panel B: $\sigma(\varepsilon)_i$ estimates</i>													
Q1	48 (0.000)	44 (0.000)	63 (0.000)	54 (0.000)	70 (0.000)	64 (0.000)	102 (0.000)	83 (0.000)	66 (0.000)	35 (0.000)	31 (0.000)	30 (0.000)	58 (0.000)
Q2	48 (0.000)	43 (0.000)	44 (0.000)	49 (0.000)	62 (0.000)	62 (0.000)	87 (0.000)	72 (0.000)	75 (0.000)	44 (0.000)	39 (0.000)	39 (0.000)	55 (0.000)
Q3	43 (0.000)	34 (0.000)	43 (0.000)	46 (0.000)	61 (0.000)	63 (0.000)	94 (0.000)	70 (0.000)	72 (0.000)	46 (0.000)	46 (0.000)	41 (0.000)	55 (0.000)
Q4	51 (0.000)	51 (0.000)	48 (0.000)	57 (0.000)	52 (0.000)	69 (0.000)	96 (0.000)	75 (0.000)	72 (0.000)	50 (0.000)	56 (0.000)	46 (0.000)	60 (0.000)
Q5	69 (0.000)	69 (0.000)	64 (0.000)	53 (0.000)	63 (0.000)	63 (0.000)	103 (0.000)	76 (0.000)	83 (0.000)	51 (0.000)	57 (0.000)	47 (0.000)	67 (0.000)
all	52 (0.000)	48 (0.000)	52 (0.000)	52 (0.000)	61 (0.000)	64 (0.000)	96 (0.000)	75 (0.000)	74 (0.000)	45 (0.000)	46 (0.000)	40 (0.000)	59 (0.000)
<i>Panel C: φ_i estimates</i>													
Q1	0.46 (0.000)	0.48 (0.000)	0.49 (0.000)	0.32 (0.000)	0.51 (0.000)	0.51 (0.000)	0.55 (0.000)	0.56 (0.000)	0.41 (0.000)	0.39 (0.000)	0.44 (0.000)	0.37 (0.000)	0.46 (0.000)
Q2	0.46 (0.000)	0.41 (0.000)	0.39 (0.000)	0.32 (0.000)	0.46 (0.000)	0.39 (0.000)	0.38 (0.000)	0.54 (0.000)	0.37 (0.000)	0.38 (0.000)	0.40 (0.000)	0.40 (0.000)	0.41 (0.000)
Q3	0.44 (0.000)	0.40 (0.000)	0.34 (0.000)	0.28 (0.000)	0.40 (0.000)	0.30 (0.000)	0.37 (0.000)	0.45 (0.000)	0.35 (0.000)	0.29 (0.000)	0.35 (0.000)	0.37 (0.000)	0.36 (0.000)
Q4	0.37 (0.000)	0.44 (0.000)	0.40 (0.000)	0.33 (0.000)	0.42 (0.000)	0.38 (0.000)	0.37 (0.000)	0.45 (0.000)	0.27 (0.000)	0.28 (0.000)	0.26 (0.000)	0.32 (0.000)	0.36 (0.000)
Q5	0.35 (0.000)	0.40 (0.000)	0.38 (0.000)	0.41 (0.000)	0.48 (0.000)	0.40 (0.000)	0.43 (0.000)	0.43 (0.000)	0.39 (0.000)	0.38 (0.000)	0.33 (0.000)	0.35 (0.000)	0.39 (0.000)
all	0.42 (0.000)	0.43 (0.000)	0.40 (0.000)	0.33 (0.000)	0.45 (0.000)	0.40 (0.000)	0.42 (0.000)	0.48 (0.000)	0.36 (0.000)	0.34 (0.000)	0.36 (0.000)	0.36 (0.000)	0.40 (0.000)
<i>Panel D: $\sigma(w)_i$ estimates</i>													
Q1	116	117	113	128	161	188	220	171	176	121	100	98	142
Q2	136	131	134	140	174	210	247	197	193	144	121	118	162
Q3	147	142	144	154	198	212	245	202	196	149	132	130	171
Q4	162	151	153	157	201	214	236	207	195	159	140	146	177
Q5	176	173	181	185	234	241	268	254	245	204	178	191	211
all	147	143	145	153	194	213	243	206	201	155	134	137	173

<continued on next page>

<continued from previous page>

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel E: transitory volatility</i>													
Q1	64	59	86	65	98	90	145	120	90	46	44	39	79
Q2	64	55	55	59	84	80	107	105	99	57	53	51	72
Q3	56	44	53	55	80	75	114	95	89	56	60	53	69
Q4	64	64	59	67	67	86	115	96	84	58	69	57	74
Q5	84	84	77	67	86	78	133	96	103	64	69	59	83
all	67	61	66	62	83	82	122	103	93	56	59	52	75
<i>Panel F: ratio of transitory and permanent variance</i>													
Q1	0.45	0.43	0.51	0.29	0.44	0.41	0.59	0.52	0.40	0.21	0.30	0.28	0.40
Q2	0.35	0.31	0.26	0.24	0.36	0.26	0.29	0.40	0.38	0.20	0.28	0.33	0.30
Q3	0.32	0.25	0.24	0.25	0.25	0.21	0.33	0.34	0.31	0.16	0.31	0.33	0.27
Q4	0.34	0.35	0.34	0.30	0.25	0.29	0.35	0.35	0.28	0.22	0.33	0.34	0.31
Q5	0.38	0.38	0.36	0.29	0.29	0.24	0.43	0.28	0.32	0.25	0.26	0.26	0.31
all	0.37	0.34	0.34	0.27	0.32	0.28	0.40	0.38	0.34	0.21	0.30	0.31	0.32
<i>Panel G: ratio of transitory and permanent "idiosyncratic" variance</i>													
Q1	0.62	0.51	0.89	0.49	0.85	0.59	0.86	0.95	0.90	0.40	0.46	0.40	0.66
Q2	0.46	0.34	0.39	0.34	0.71	0.32	0.39	0.70	0.89	0.32	0.40	0.49	0.48
Q3	0.41	0.27	0.33	0.32	0.48	0.25	0.43	0.51	0.65	0.25	0.43	0.43	0.40
Q4	0.42	0.36	0.39	0.37	0.36	0.33	0.43	0.51	0.47	0.37	0.48	0.42	0.41
Q5	0.42	0.39	0.39	0.32	0.37	0.25	0.48	0.37	0.44	0.34	0.33	0.32	0.37
all	0.47	0.37	0.48	0.37	0.55	0.35	0.52	0.61	0.67	0.34	0.42	0.41	0.46

Table 5: State space model estimates by year and size quintile ($\varphi_i = 0$)

This table estimates the following state space model for a latent efficient price and an observed end-of-day midquote:

$$\begin{aligned}
 \text{(observed price)} \quad p_{it} &= m_{it} + s_{it} \\
 \text{(unobserved efficient price)} \quad m_{it} &= m_{i,t-1} + \beta_i \hat{\gamma}_t + w_{it} \quad w_{it} = \kappa_i (I_{it}^{idio} - E_{t-1}[I_{it}^{idio}]) + u_{it} \\
 \text{(unobserved transitory price deviation)} \quad s_{it} &= \varphi_i s_{i,t-1} + \alpha_i I_{it}^{idio} + \beta_i^0 \hat{\gamma}_t + \dots + \beta_i^3 \hat{\gamma}_{t-3} + \varepsilon_{it}
 \end{aligned}$$

where i indexes over stocks and t indexes over days, m_{it} is the end-of-day unobserved efficient price (“state”), $\hat{\gamma}_t$ is a midquote return common factor which is the cross-sectional average of the standardized midquote return series which has been filtered with an AR(4) model to remove intertemporal dynamics, p_{it} is end-of-day observed midquote, I_{it}^{idio} is the idiosyncratic part of the specialist end-of-day USD inventory that remains after removing a common factor across specialist inventories (we use the average price across the entire period to convert inventory in shares to USD), $E_{t-1}[I_{it}^{idio}]$ is an AR(9) based forecast so that $(I_{it}^{idio} - E_{t-1}[I_{it}^{idio}])$ is a martingale that captures the surprise part of the net imbalance, β_i^j captures potential “overreaction” or lagged adjustment to common factor innovations, and u_{it} and ν_{it} are mutually independent i.i.d. error terms. We report maximum likelihood estimates where we assume normality for the error terms. We implement our optimization in ox with ssfpack routines where we use the Kalman filter to evaluate the likelihood (see Koopman, Shephard, and Doornik (1999)). The table reports p -values in brackets. These p -values are based on a test statistic that counts the number of significant t -values across all stock-year ML estimates in the bin. The test statistic is binomially distributed under the null (we use the 0.1 and 0.9 quantiles in the t -test).

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel A: α_i estimates</i>													
Q1	-0.03 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.02 (0.000)	-0.01 (0.000)	-0.01 (0.000)	-0.01 (0.000)	-0.02 (0.000)
Q2	-0.06 (0.000)	-0.05 (0.000)	-0.05 (0.000)	-0.04 (0.000)	-0.03 (0.000)	-0.05 (0.000)	-0.05 (0.000)	-0.05 (0.000)	-0.05 (0.000)	-0.03 (0.001)	-0.03 (0.000)	-0.03 (0.000)	-0.04 (0.000)
Q3	-0.11 (0.000)	-0.09 (0.000)	-0.09 (0.000)	-0.07 (0.000)	-0.09 (0.000)	-0.12 (0.000)	-0.11 (0.000)	-0.06 (0.000)	-0.09 (0.000)	-0.06 (0.000)	-0.07 (0.000)	-0.06 (0.000)	-0.09 (0.000)
Q4	-0.33 (0.000)	-0.25 (0.000)	-0.30 (0.000)	-0.27 (0.000)	-0.26 (0.000)	-0.34 (0.000)	-0.27 (0.000)	-0.23 (0.000)	-0.27 (0.000)	-0.14 (0.000)	-0.19 (0.000)	-0.16 (0.000)	-0.25 (0.000)
Q5	-1.07 (0.000)	-1.04 (0.000)	-0.78 (0.000)	-0.74 (0.000)	-1.01 (0.000)	-1.09 (0.000)	-1.19 (0.000)	-1.30 (0.000)	-1.20 (0.000)	-0.86 (0.000)	-0.94 (0.000)	-0.87 (0.000)	-1.01 (0.000)
all	-0.32 (0.000)	-0.29 (0.000)	-0.25 (0.000)	-0.23 (0.000)	-0.28 (0.000)	-0.32 (0.000)	-0.33 (0.000)	-0.33 (0.000)	-0.32 (0.000)	-0.22 (0.000)	-0.25 (0.000)	-0.23 (0.000)	-0.28 (0.000)
<i>Panel B: $\sigma(\varepsilon)_i$ estimates</i>													
Q1	20 (0.000)	17 (0.000)	25 (0.000)	30 (0.000)	24 (0.000)	20 (0.000)	31 (0.000)	22 (0.000)	31 (0.000)	18 (0.000)	12 (0.000)	15 (0.000)	22 (0.000)
Q2	19 (0.000)	19 (0.000)	21 (0.000)	28 (0.000)	24 (0.000)	22 (0.000)	43 (0.000)	21 (0.000)	39 (0.000)	23 (0.000)	17 (0.000)	19 (0.000)	25 (0.000)
Q3	19 (0.000)	17 (0.000)	22 (0.000)	23 (0.000)	27 (0.000)	30 (0.000)	45 (0.000)	27 (0.000)	41 (0.000)	25 (0.000)	20 (0.000)	20 (0.000)	26 (0.000)
Q4	28 (0.000)	25 (0.000)	23 (0.000)	26 (0.000)	18 (0.000)	23 (0.000)	42 (0.000)	30 (0.000)	43 (0.000)	30 (0.000)	30 (0.000)	24 (0.000)	29 (0.000)
Q5	42 (0.000)	42 (0.000)	34 (0.000)	25 (0.000)	21 (0.000)	29 (0.000)	47 (0.000)	38 (0.000)	39 (0.000)	28 (0.000)	28 (0.000)	25 (0.000)	33 (0.000)
all	26 (0.000)	24 (0.000)	25 (0.000)	27 (0.000)	23 (0.000)	25 (0.000)	42 (0.000)	28 (0.000)	39 (0.000)	25 (0.000)	21 (0.000)	21 (0.000)	27 (0.000)
<i>Panel C: φ_i estimates</i>													
Q1	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
Q2	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
Q3	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
Q4	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
Q5	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
all	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)	0.00 (1.000)
<i>Panel D: $\sigma(w)_i$ estimates</i>													
Q1	124	124	128	134	176	198	243	192	187	126	106	103	153
Q2	143	136	138	144	186	213	257	212	206	151	127	125	170
Q3	151	143	147	158	204	213	253	212	205	153	138	136	176
Q4	162	152	153	158	201	211	239	212	199	163	145	151	179
Q5	176	170	181	181	228	232	264	245	244	202	178	190	208
all	151	145	149	155	199	213	251	215	208	159	139	141	177

<continued on next page>

<continued from previous page>

	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	all
<i>Panel E: transitory volatility^a (σ_i)</i>													
Q1	14	13	13	16	19	22	25	30	22	15	14	9	17
Q2	26	24	22	19	20	33	36	30	23	14	15	13	23
Q3	39	40	35	31	40	48	45	35	23	16	20	17	32
Q4	65	64	69	58	63	86	77	63	41	21	24	21	54
Q5	116	109	115	107	152	177	184	155	108	92	71	54	120
all	52	50	51	46	58	73	73	62	43	31	29	23	49
<i>Panel F: ratio of transitory and permanent variance^a</i>													
Q1	0.02	0.02	0.01	0.02	0.01	0.02	0.01	0.03	0.02	0.02	0.03	0.01	0.02
Q2	0.08	0.09	0.04	0.03	0.02	0.04	0.03	0.03	0.02	0.01	0.02	0.02	0.03
Q3	0.28	0.27	0.16	0.09	0.13	0.19	0.05	0.07	0.02	0.01	0.04	0.09	0.12
Q4	0.30	0.47	0.47	0.28	0.23	0.47	0.22	0.16	0.06	0.03	0.04	0.05	0.23
Q5	1.11	1.42	1.14	1.16	1.24	2.55	1.96	1.09	0.79	2.10	0.61	0.62	1.32
all	0.36	0.45	0.36	0.31	0.32	0.65	0.46	0.28	0.18	0.44	0.15	0.16	0.34
<i>Panel G: ratio of transitory and permanent "idiosyncratic" variance^a</i>													
Q1	0.03	0.03	0.02	0.03	0.02	0.02	0.02	0.04	0.03	0.04	0.04	0.01	0.03
Q2	0.10	0.10	0.05	0.04	0.02	0.05	0.04	0.04	0.03	0.02	0.03	0.02	0.04
Q3	0.35	0.28	0.20	0.11	0.18	0.21	0.06	0.10	0.04	0.02	0.05	0.12	0.14
Q4	0.36	0.48	0.54	0.31	0.29	0.51	0.25	0.23	0.08	0.04	0.05	0.05	0.27
Q5	1.28	1.48	1.26	1.28	1.51	2.64	2.16	1.40	0.98	2.31	0.73	0.67	1.47
all	0.42	0.47	0.41	0.35	0.40	0.68	0.51	0.36	0.23	0.48	0.18	0.18	0.39
<i>Panel H: proxy social cost unrealized hedge value (\$ million per year)</i>													
Q1	1.59	2.55	3.23	5.91	10.23	17.20	23.82	20.79	14.34	5.53	6.50	6.66	9.86
Q2	1.66	1.66	2.72	2.88	3.51	8.18	9.79	6.96	3.63	1.75	2.13	2.62	3.96
Q3	0.80	1.30	1.22	1.72	2.69	3.16	4.44	3.82	2.02	1.18	1.55	1.86	2.15
Q4	0.52	0.60	0.75	0.98	1.56	2.39	3.11	2.45	1.43	0.70	1.06	0.86	1.37
Q5	0.29	0.39	0.48	0.82	1.05	0.96	1.23	0.87	0.48	0.41	0.55	0.53	0.67
all	0.97	1.30	1.68	2.46	3.81	6.38	8.48	6.98	4.38	1.91	2.36	2.51	3.60

^a: The transitory volatility/variance is based on the price pressure term only, i.e. $\alpha^2 \sigma^2(\Gamma^{idio})$, as this captures the magnitude of price pressures. It therefore excludes the "residual" stationary term that is captured by ε_{it} .

Table 6: Overall, Between, and Within Correlations

This table presents the overall, between, and within correlations for the Table 5 state space model estimates and trade characteristics that are organized in a 697*12 stockyear panel dataset. We refer to Table 2 for variable definitions.

		<i>dollar_</i> <i>volume</i> <i>e</i> <i>it</i>	$ \alpha_{it} ^*$ $\sigma(I^{Idio})_{it}$	$ \alpha_{it} $	$\sigma(I^{Idio})_{it}$	<i>autocorr</i> $(I^{Idio})_{it}$	<i>soc_cost_</i> <i>unreal_</i> <i>ized_hedge_</i> <i>val</i> <i>it</i>
$\sigma(w)_{it}$	$\rho(\text{overall})$	-0.01	0.18*	0.10*	0.04*	0.04*	0.06*
	$\rho(\text{between})$	-0.08*	0.22*	0.16*	-0.07	0.08*	-0.08*
	$\rho(\text{within})$	0.10*	0.15*	0.03*	0.13*	0.00	0.15*
<i>dollar_volume</i> <i>e</i> <i>it</i>	$\rho(\text{overall})$		-0.15*	-0.15*	0.47*	-0.26*	0.43*
	$\rho(\text{between})$		-0.27*	-0.22*	0.76*	-0.43*	0.66*
	$\rho(\text{within})$		-0.01	-0.01	0.14*	-0.03*	0.27*
$ \alpha_{it} * \sigma(I^{Idio})_{it}$	$\rho(\text{overall})$			0.66*	-0.03*	0.50*	0.02
	$\rho(\text{between})$			0.84*	-0.23*	0.75*	-0.17*
	$\rho(\text{within})$			0.49*	0.11*	0.30*	0.13*
$ \alpha_{it} $	$\rho(\text{overall})$				-0.13*	0.33*	-0.07*
	$\rho(\text{between})$				-0.24*	0.56*	-0.18*
	$\rho(\text{within})$				-0.01	0.06*	0.01
$\sigma(I^{Idio})_{it}$	$\rho(\text{overall})$					0.00	0.52*
	$\rho(\text{between})$					-0.27*	0.84*
	$\rho(\text{within})$					0.22*	0.34*
<i>autocorr</i> $(I^{Idio})_{it}$	$\rho(\text{overall})$						-0.06*
	$\rho(\text{between})$						-0.22*
	$\rho(\text{within})$						0.03*

^a: Based on the time means i.e. $\bar{x}_i = \frac{1}{T} \sum_{t=1}^T x_{it}$.

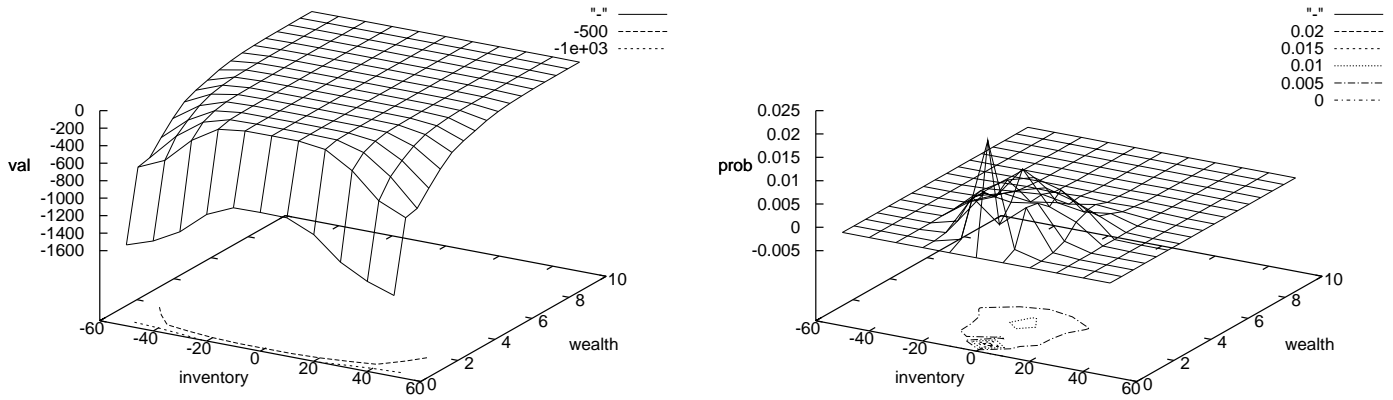
^b: Based on the deviations from time means i.e. $x_{it}^* = x_{it} - \bar{x}_i$.

*: Significant at a 95% level.

Figure 1: Value function, stationary distribution, and optimal control

This figure illustrates the numerical solution to a dynamic programming problem of an intermediary who maximizes expected utility while quoting bid and ask prices to liquidity demanders whose arrival rates are less than perfectly price elastic. Panel A plots the value function (left) and the stationary distribution (right) over all the inventory-wealth states that the intermediary could find herself in. Panel B illustrates her optimal control by plotting the midquote (left), the bid-ask spread (right), and her consumption (bottom) conditional on her inventory-wealth state. The parametrization of the problem is standard and follows Ho and Stoll (1981). The intermediary has CRRA utility with a coefficient of relative risk aversion equal to 2. She operates at a daily frequency and has a two week horizon which we capture by a (daily) discount factor of 0.9. Without loss of generality, we fix the fundamental price at zero and capture inventory price risk through a stochastic dividend with mean zero and daily standard deviation of \$0.05. The probability of a public buyer (seller) arrival depends linearly on the ask (bid) price where the slope θ captures the elasticity of arrival rates to prices (the benchmark is no friction, i.e. $\theta = \infty$).

Panel A: Value function and stationary distribution



Panel B: Optimal controls: midquote, bid-ask spread, and consumption

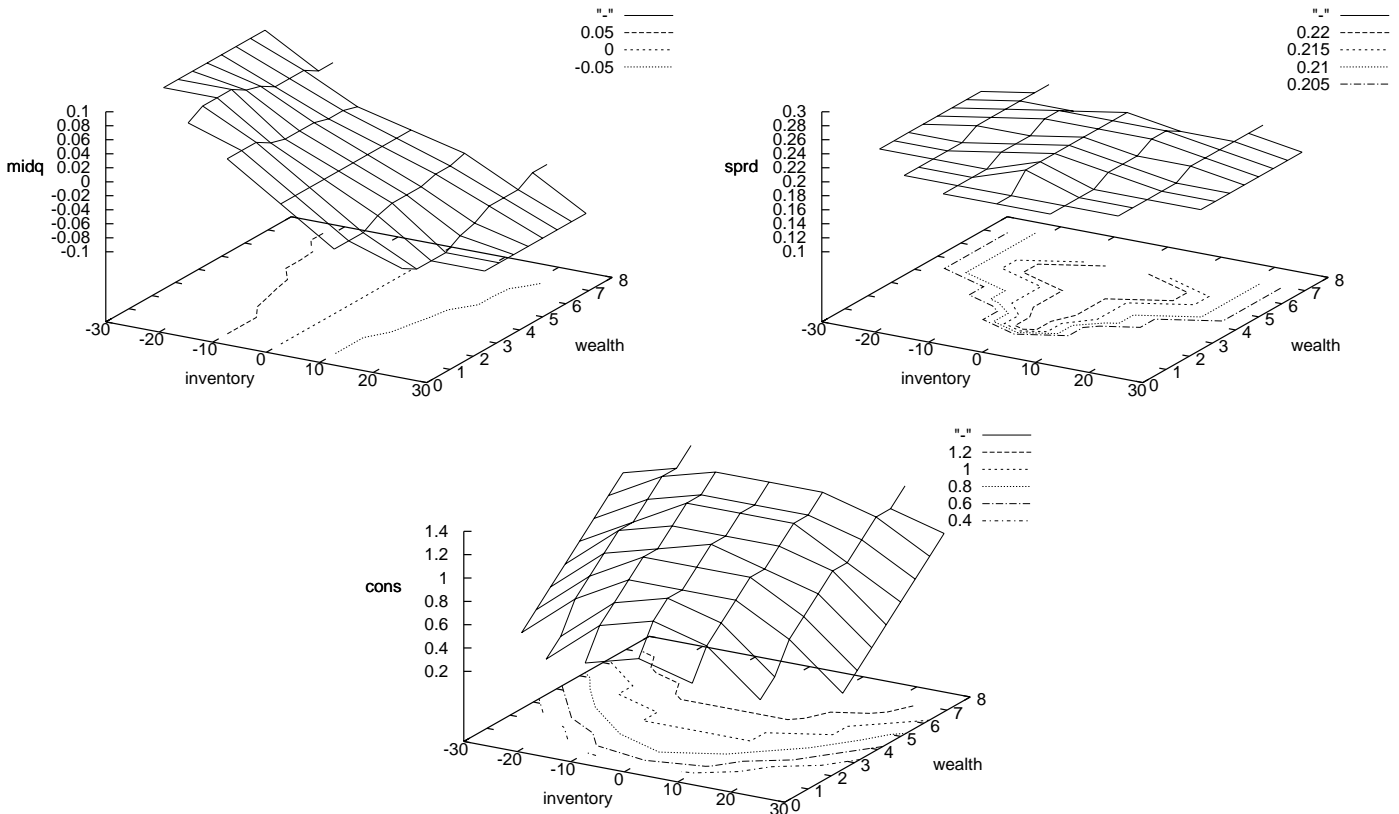
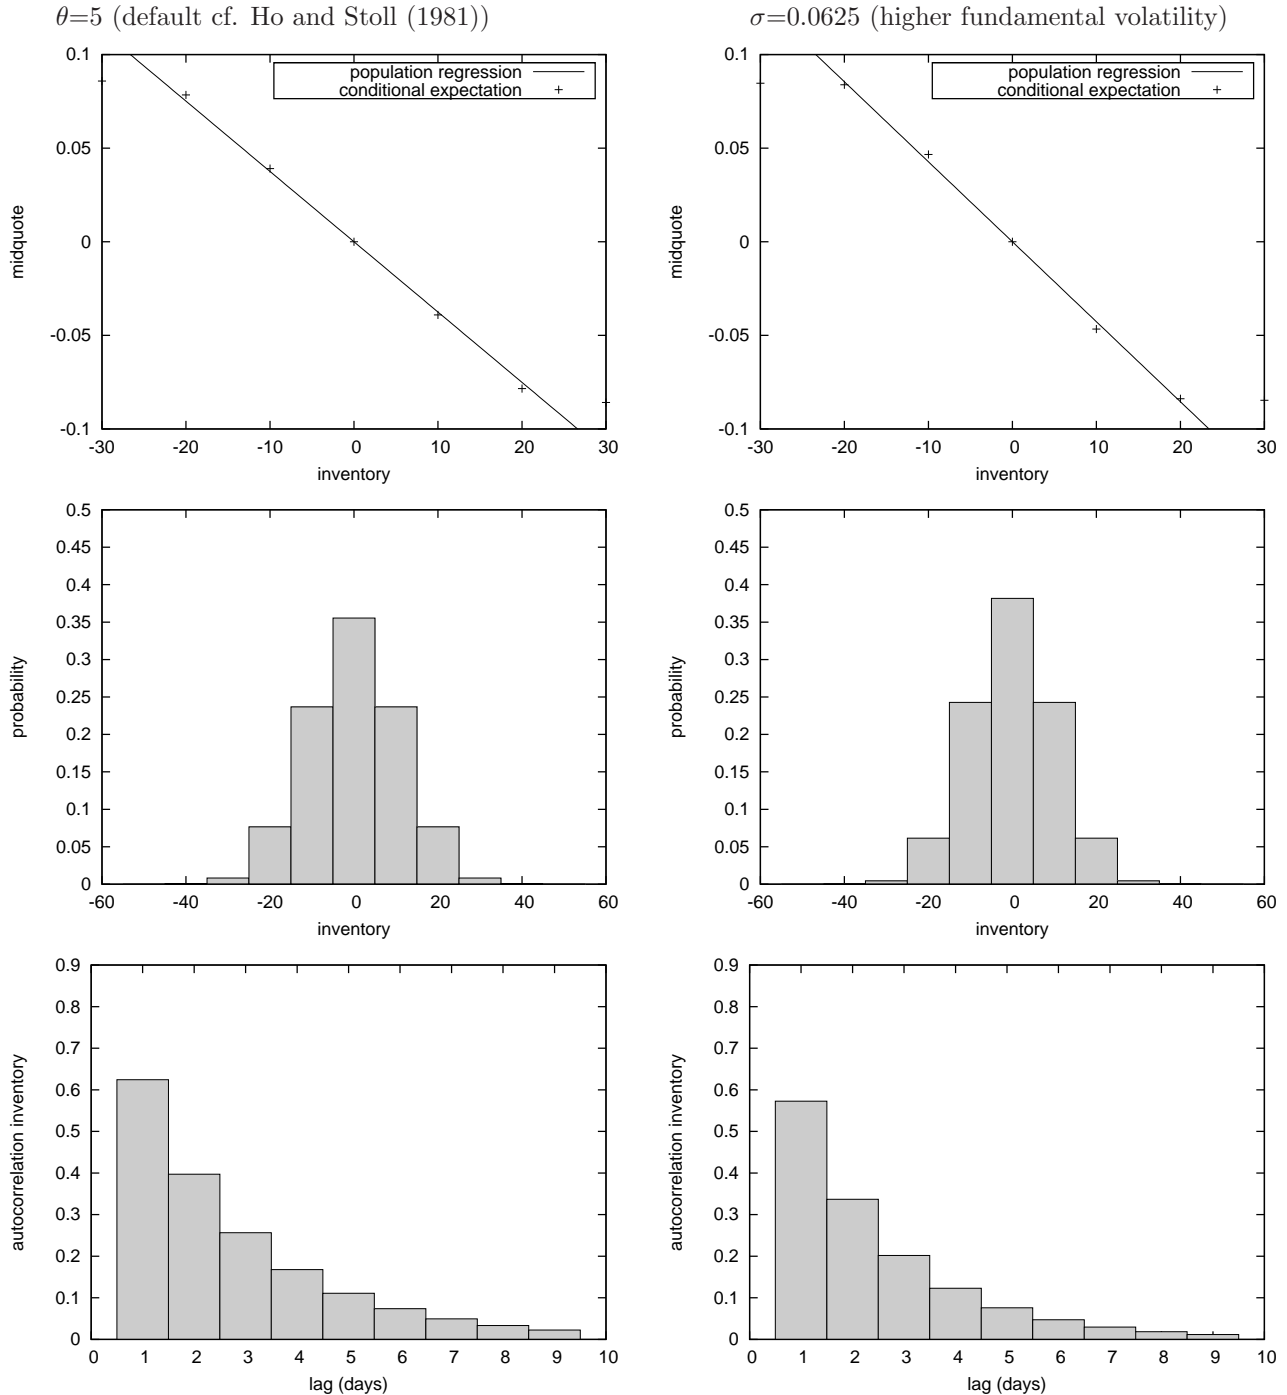


Figure 2: Comparative statistics

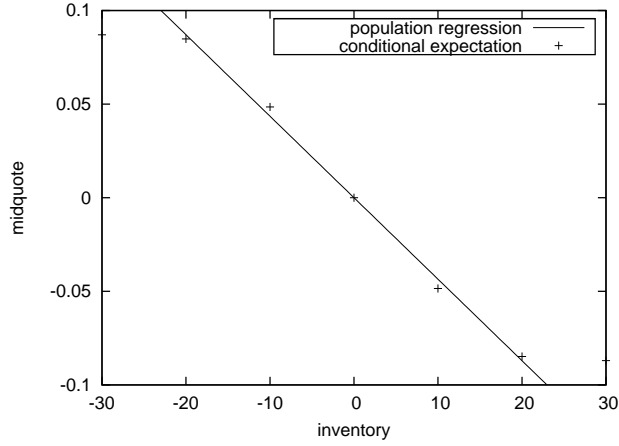
This figure illustrates the numerical solution to a dynamic programming problem of an intermediary who maximizes expected utility while quoting bid and ask prices to liquidity demanders whose arrival rates are less than perfectly price elastic. It depicts the dynamics for the base case elasticity of arrival rates (cf. Ho and Stoll (1981)) in the left column and for more elastic arrival rates in the right column. We plot a histogram to depict the (unconditional) inventory distribution (top), we plot the autocorrelation function to illustrate inventory dynamics (middle), and we plot a population regression to show how the midquote depends on inventory (bottom). This bottom graph shows the conditional expectation of the midquote given inventory. The parametrization of the problem is standard and follows Ho and Stoll (1981). The intermediary has CRRA utility with a coefficient of relative risk aversion equal to 2. She operates at a daily frequency and has a two week horizon which we capture by a (daily) discount factor of 0.9. Without loss of generality, we fix the fundamental price at zero and capture inventory price risk through a stochastic dividend with mean zero and daily standard deviation of \$0.05. The probability of a public buyer (seller) arrival depends linearly on the ask (bid) price where the slope θ captures the elasticity of arrival rates to prices (the benchmark is no friction, i.e. $\theta = \infty$).



(continued on next page)

(continued from previous page)

$\lambda=0.15$ (lower hedge value rate)



$\theta=3.75$ (less elastic net transaction rate)



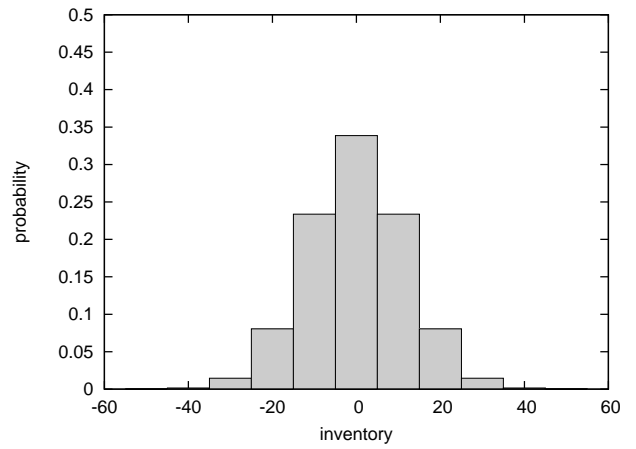
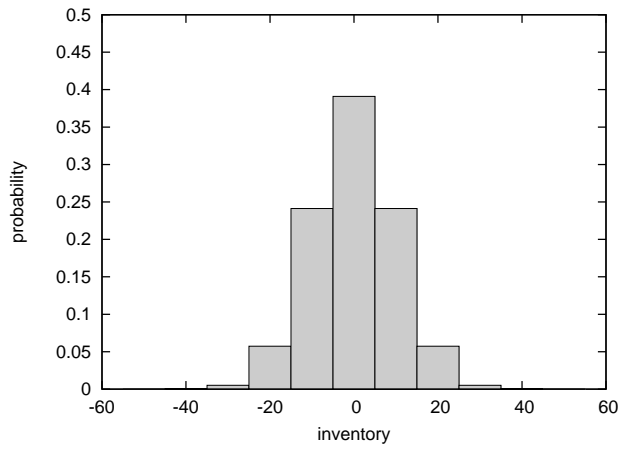
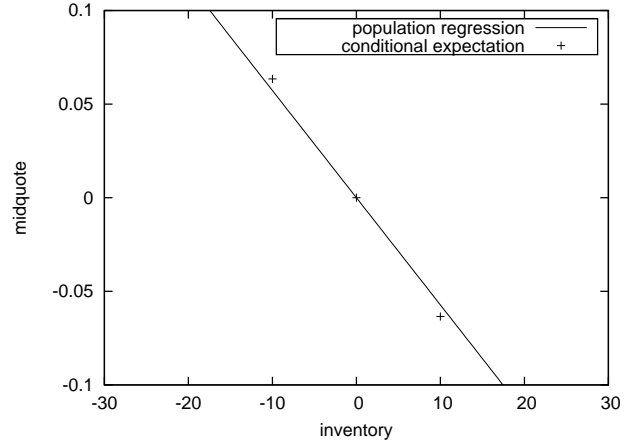



Figure 3: An example of a complete price change decomposition given the state space model estimates ($\alpha_i = 0$)

This figure plots a complete decomposition of the log price series (midquote) of a stock in the lowest-cap size quintile (permno=63503). We start all level series relative to the model's estimate of the efficient price at the first day of the sample.

Panel A: The observed price series (midquote), the (unobserved) efficient price estimate (left) and a decomposition of the efficient price into its two stochastic trend components (right)



Panel B: The wedge between the observed price series and the efficient estimate decomposed into (i) lagged adjustment to common factor (left) and (ii) remaining unobserved stationary term (right). We also plot the idiosyncratic inventory position of the intermediary (right).

