

Stock Returns and the Volatility of Liquidity ^{*}

João Pedro Pereira [†]

Finance Department

ISCTE Business School - Lisbon

Harold H. Zhang [‡]

School of Management

University of Texas at Dallas

December 2, 2008

Abstract

This paper offers a rational explanation for the puzzling empirical fact that stock returns decrease in the volatility of liquidity. We model liquidity as a stochastic price impact process and define the liquidity premium as the additional return necessary to compensate a multi-period investor for the adverse price impact of trading. The model demonstrates that a fully rational, utility maximizing, risk averse investor can take advantage of time-varying liquidity by adapting his trades to the state of liquidity. We provide new empirical evidence supportive of the model.

^{*}We are grateful for the helpful comments of Dong-Hyun Ahn, Evan Anderson, Greg Brown, Jennifer Conrad, Benjamin Croitoru, Bin Gao, Eric Ghysels, Mustafa Gültekin, Kushal Kshirsagar, Hong Liu, Miroslav Misina, Ronnie Sadka, Avanidhar Subrahmanyam, Dragon Tang, the Editor of the JFQA, Hendrik Bessembinder, and seminar participants at the University of North Carolina at Chapel Hill, the 2004 Financial Management Association meeting, the 2004 Equity Premium Puzzle Conference at the University of Exeter in UK, and at the 2005 European Finance Association meeting in Moscow. The paper has benefited greatly from the advice of an anonymous referee. We thank UNC - Chapel Hill for providing computer resources. An earlier version of this paper was titled “The Liquidity Premium in a Portfolio Choice Model with Price Impact”.

[†]Av. Forcas Armadas, 1649-026 Lisboa, Portugal; phone: (+351) 217 903 916; email: joao.pereira@iscte.pt.

[‡]Richardson, TX 75080, USA; phone: (972) 883 4777; email: harold.zhang@utdallas.edu.

I Introduction

There is ample evidence that liquidity affects asset returns. One line of research views liquidity as a characteristic that influences returns beyond trading costs. Investing in illiquid stocks is compensated by higher gross returns.¹ Another line of research emphasizes liquidity as a market-wide risk factor. Stocks with higher sensitivity to innovations in aggregate liquidity have higher expected returns.² However, there is still a considerable debate on the precise definition and role of liquidity.³

In particular, the effect of the *volatility* of liquidity on stock returns is not well understood. Using a sample of monthly returns for NYSE and AMEX stocks for the period from 1966 to 1995, Chordia, Subrahmanyam, and Anshuman (2001) surprisingly find that stocks with higher volatility of liquidity actually have *lower* returns. This relation seems puzzling since it appears to contradict the usual risk-return tradeoff intuition.⁴ Indeed, Hasbrouck (2006, p.31) suggests that their result is so surprising that the problem may even reside in the proxies used for liquidity: “Surprisingly they find that turnover volatility is negatively related to expected returns. This is contrary to the notion that turnover volatility might be acting as proxy for liquidity risk.”

We offer a rational explanation for the puzzling negative relation between stock

¹Several studies have documented that expected returns are decreasing in the level of liquidity, measured by the bid-ask spread (e.g., Amihud and Mendelson (1986)), price impact (e.g., Brennan and Subrahmanyam (1996)), turnover (e.g., Datar, Naik, and Radcliffe (1998)), or trading volume (e.g., Brennan, Chordia, and Subrahmanyam (1998)).

²Pástor and Stambaugh (2003) construct a market-wide liquidity factor and find that stocks whose returns are more correlated with the aggregate liquidity factor have higher expected returns.

³Acharya and Pedersen (2005) detect the effect of liquidity both as a characteristic (return depends on the liquidity level) and as a risk factor (return depends on the covariances between the security’s own return and liquidity with the common liquidity factor). Korajczyk and Sadka (2008), using high frequency data, also find that both liquidity risk and level are priced. However, Hasbrouck (2006) proposes a new way to estimate effective spreads and using a long sample finds only weak support for the effect of liquidity as a characteristic and no support as a risk factor.

⁴In a survey on liquidity and stock returns, Amihud, Mendelson, and Pedersen (2005) state that “because liquidity varies over time, risk-averse investors may require a compensation for being exposed to liquidity risk,” suggesting a positive relation between stock returns and the volatility of liquidity.

returns and the volatility of liquidity and demonstrate that this negative relation is consistent with utility-maximizing investment strategies of risk-averse investors. We examine the relation between expected stock returns and the volatility of liquidity in a dynamic portfolio-choice model with stochastic liquidity. Specifically, a constant relative risk aversion (CRRA) investor allocates his wealth between a risky stock and a risk-free asset. The stock is illiquid in the sense that trading induces adverse price impact.⁵ We assume that price impact follows a stochastic mean-reverting process, thus capturing the important empirical fact that liquidity varies through time. Because trading moves the price against the investor, his expected utility will be lower than in the case of investing in a perfectly liquid stock. The liquidity premium is defined as the extra return that the illiquid stock must earn so that the investor attains the same level of utility as in the case of a perfectly liquid stock.

We calibrate the model to empirically reasonable parameter values and numerically solve for the investor's optimal trading strategy and required liquidity premium. We demonstrate that a rational risk-averse utility-maximizing investor adapts his trading to the state of liquidity and trades large quantities in high liquidity states and small quantities in low liquidity states. A higher liquidity volatility provides more opportunity for the investor to time his trades and leads to a lower required liquidity premium. Therefore, stocks with higher liquidity volatility command a lower return premium. Hence, our analysis offers a rational explanation for the “puzzling” empirical finding on

⁵The liquidity of an asset is a characteristic hard to define and measure. Nonetheless, a commonly accepted definition of liquidity states that an asset is liquid if large quantities can be traded in a short period of time without moving the price too much. Hence, a natural measure of liquidity is the price impact of trading. Studies using price impact as a measure of liquidity include Brennan and Subrahmanyam (1996), Bertsimas and Lo (1998), He and Mamaysky (2005), Amihud (2002), Pástor and Stambaugh (2003), Acharya and Pedersen (2005), and Sadka (2006). The bid-ask spread is also accepted as a measure of liquidity and has been used in earlier studies, starting with Amihud and Mendelson (1986). However, large blocks of shares usually trade outside the bid-ask spread (see, e.g., Chan and Lakonishok (1995) and Keim and Madhavan (1996)).

the negative relation between stock returns and the volatility of liquidity.

The intuition from our model resembles the “tax trading option” of Constantinides and Scholes (1980). They show that stock return volatility leads to higher welfare due to a higher probability of a realized capital loss. Their intuition is that there is a fundamental asymmetry between capital gains and losses, in that gains are deferred and losses realized. Therefore, stocks with higher return volatility should have lower expected returns. Similarly, there is also a fundamental asymmetry between high and low liquidity states, in that the investor can time his trades to avoid bad liquidity states and trade more in good liquidity states.

We provide empirical evidence that the cross sectional negative relation between stock returns and the volatility of liquidity, first reported in Chordia, Subrahmanyam, and Anshuman (2001), still holds in a longer and more recent sample (NYSE and AMEX stocks for 1963–2005), and more important, that it also holds when liquidity is proxied by the price impact measure of Amihud (2002). We also provide new evidence that there is time series Granger causality from price impact to trading activity. In particular, larger (smaller) price impact leads to less (more) trading volume or turnover. This is consistent with existing literature showing that institutional investors time their trades according to the state of liquidity. Hence, the empirical analysis provides supporting evidence for both the main cross sectional implication of the model and the mechanism leading to the implication.

We analyze the relation between returns and the volatility of liquidity in a partial-equilibrium model. This approach is similar to Constantinides (1986), Longstaff (2001), and Jang, Koo, Liu, and Loewenstein (2007), in the sense that we solve for a liquidity premium such that the investor is indifferent between a liquid and an illiquid stock, given an exogenous liquidity process. Naturally, it would be more desirable to study a

general equilibrium model with an endogenous price impact process. However, such a model would be intractable. The relative simplicity of our partial-equilibrium model makes it solvable and allows us to offer a rational explanation for the negative relation between returns and the volatility of liquidity.

Several related papers study the asset pricing effects of transaction costs, including Heaton and Lucas (1996), Vayanos (1998), and Lo, Mamaysky, and Wang (2004). However, none of these papers analyzes the implication of stochastic liquidity on returns. The trading strategies uncovered in our study are broadly consistent with existing research on the best trading strategy for buying or selling a fixed block of shares when there is price impact, including Bertsimas and Lo (1998), Huberman and Stanzl (2005), He and Mamaysky (2005), and Obizhaeva and Wang (2006). However, these papers do not analyze the effect of a stochastic price impact. Another important difference is that we do not assume a target position that the investor has to buy or sell. Instead, our investor optimally chooses the level of shares he desires to hold and the best investment strategy to achieve it.

The next section presents the model. Section III shows numerical solutions. We first consider a simplified model where price impact is constant. We then solve the full model with stochastic liquidity and examine the implication on stock returns. Section IV tests the robustness of the main findings by considering several extensions of the model. Section V presents empirical tests of the model predictions. Section VI concludes.

II The Model

Consider an investor who maximizes the expected utility of terminal wealth, $E[u(W_T)]$, where W_t is the investor's wealth at time t and T is the terminal date. Assume that the

investor has the constant relative risk aversion (CRRA) utility function given by

$$(1) \quad u(W) = \begin{cases} W^{1-\gamma}/(1-\gamma) & \text{if } \gamma > 1 \\ \ln(W) & \text{if } \gamma = 1 \end{cases}$$

where γ represents the investor's coefficient of relative risk aversion. At each point in time he invests his wealth in N_t units of a risky stock and M_t units of a risk-free bond (or money market account). The bond price, B_t , follows the discrete process

$$(2) \quad B_{t+1} - B_t = rB_t$$

where r is the risk-free rate. The price of a perfectly liquid stock follows the process

$$(3) \quad S_{t+1} - S_t = S_t(\mu + \sigma\varepsilon_{t+1}).$$

where the random shock ε_{t+1} is distributed $\varepsilon_{t+1|t} \sim N(0, 1)$.

We now consider two departures from the standard portfolio choice problem with perfect liquidity by introducing price impact and the irrelevance of paper wealth.

(1) *Price impact.* We assume that the investor has to make a price concession when selling the stock and must pay a higher price when buying the stock. Hence, the stock is illiquid in the sense that trading moves its price. Specifically, we consider a price impact function similar to He and Mamaysky (2005) and Breen, Hodrick, and Korajczyk (2002). In the absence of any other factors, the price change caused by trading is given by

$$(4) \quad (S_{t+1} - S_t)/S_t = \psi_{t+1}(N_{t+1} - N_t)$$

where $N_{t+1} - N_t$ represents the shares traded and ψ_{t+1} is a positive value. Hence,

the stock return is proportional to (signed) volume. If the investor buys the stock ($N_{t+1} > N_t$), he pays a higher price. If he sells ($N_{t+1} < N_t$), he receives a lower price.

The price impact assumed here is permanent. Theoretically, this form of price impact can be motivated by the existence of a market maker who extracts information from the order flow and adjusts the price accordingly (as in Kyle (1985)). The empirical evidence supports a permanent effect for large trades. For example, Holthausen, Leftwich, and Mayers (1990) show that most of the price effect associated with block trades is permanent. Seppi (1992) shows that large-block trades reveal information about earnings. Sias, Starks, and Titman (2001) argue that the price pressure caused by institutional trading tends to be permanent (information related) rather than transitory (inventory related). Furthermore, we allow the permanent price impact to vary through time. In other words, as the underlying market conditions change, traders and market makers may infer different information from otherwise identical trades. This assumption is consistent with the theoretical model of Saar (2001), where permanent price impacts vary according to the underlying economic environment, and with the empirical evidence in Chiyachantana, Jain, Jiang, and Wood (2004). Hence, our model captures the most important aspect of trading for large investors.⁶

In the presence of a permanent price impact, the price of an illiquid stock thus moves according to the sum of the two components in (3) and (4):

$$(5) \quad S_{t+1} - S_t = S_t [\mu + \lambda + \sigma \varepsilon_{t+1} + \psi_{t+1} (N_{t+1} - N_t)]$$

⁶Nevertheless, in reality trading may also induce an additional transitory price effect due to inventory risk. For example, Sadka (2006), extending the methodology of Glosten and Harris (1988), decomposes the price impact into an (informational) permanent component and a (non-informational) transitory component. Hence, section IV extends the model to include a temporary price effect. An alternative and also economically sensible modeling approach would be to build our benchmark model around a time-varying temporary price impact process. Even though the computational details would change, the economic intuition derived from the model would remain the same: randomness in (the temporary component of) liquidity would still allow the investor to time his trades and thus decrease the required expected return.

The parameter λ in equation (5) represents the liquidity premium, that is, the extra expected return that the stock must earn to compensate the investor for its illiquidity.

A natural specification for the evolution of liquidity is to assume that it fluctuates randomly around a long-term mean. Thus, we define the following mean-reverting process for the price impact coefficient:

$$(6) \quad \psi_{t+1} = \bar{\psi} + \rho(\psi_t - \bar{\psi}) + \varphi\varepsilon_{t+1}^\psi$$

where $\bar{\psi}$ is the long-term mean, ρ is the first-order autocorrelation, φ is the volatility, and ε_{t+1}^ψ is white noise. Amihud (2002) also assumes a first-order autoregressive process for a closely related price impact measure. Acharya and Pedersen (2005) and Pástor and Stambaugh (2003) use second-order autoregressive processes.

The timing is as follows. One instant before choosing N_{t+1} , the investor observes price $\bar{S}_{t+1} \equiv S_t[1 + \mu + \lambda + \sigma\varepsilon_{t+1}]$, and the price impact coefficient, ψ_{t+1} . The investor then chooses N_{t+1} . This forms the new market price S_{t+1} . Note that the price impact coefficient is known at the time the investor chooses the action, i.e., N_{t+1} is chosen after observing both \bar{S}_{t+1} and ψ_{t+1} . The following scheme represents the timing.

t		$t + 1$		$t + 2$
S_t	→	1) Observe \bar{S}_{t+1}		
		2) Choose N_{t+1}		
		3) S_{t+1} is created	→	\bar{S}_{t+2}
		4) W_{t+1} known		

This can be interpreted as the investor seeing the midpoint of the bid-ask, \bar{S}_{t+1} , as well as the whole demand-supply schedules (both sides of the order book) at the time he chooses N_{t+1} . When the investor submits his order of $N_{t+1} - N_t$ shares, he hits an order in the book, trading takes place, and a transaction price S_{t+1} is recorded.

The investor's wealth is $W_t \equiv M_t B_t + N_t S_t$. Imposing a self-financing constraint, we arrive at

$$(7) \quad W_{t+1} - W_t = M_t(B_{t+1} - B_t) + N_t(S_{t+1} - S_t).$$

Note that N_{t+1} does not influence W_{t+1} directly; it only does so indirectly through S_{t+1} . Equation (7) implies that the trade is executed at the post-impact price S_{t+1} . The dollar value traded in stocks, $S_{t+1}(N_{t+1} - N_t)$, is exactly absorbed by changes in the money-market account, $B_{t+1}(M_{t+1} - M_t)$. Therefore, N_{t+1} influences W_{t+1} only by changing the price, S_{t+1} , of the N_t shares already owned.

We replace $M_t = (W_t - N_t S_t)/B_t$ and use (2) to get $W_{t+1} - W_t = (W_t - N_t S_t)r + N_t(S_{t+1} - S_t)$, and then (5) to arrive at

$$(8) \quad W_{t+1} = W_t(1 + r) + N_t S_t [\mu + \lambda - r + \sigma \varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)].$$

With price impact, we create an important departure from the standard model: trading by itself changes the investor's wealth. For example, even if we set $r = 0$, $\mu + \lambda = 0$, and $\sigma = 0$, simply buying more shares ($N_{t+1} - N_t > 0$) at the post-impact price of S_{t+1} increases the wealth by $W_{t+1} - W_t = N_t(S_{t+1} - S_t) = N_t S_t \psi_{t+1}(N_{t+1} - N_t)$. However, this increase in wealth is only "on paper". It can be reversed when the investor sells the stock. Hence, we must also introduce a second restriction.

(2) *Paper wealth is irrelevant.* One important feature of investing in illiquid stocks is that both the initial accumulation and the final unloading of the stock induce adverse price movements. Therefore, we model an investor that starts without any holdings of the stock (his initial wealth is in the form of cash): $N_0 = 0$. Similarly, the stock holding must be liquidated by the end of the investment horizon (T): $N_T = 0$. The stock will

have to be sold because the investor only derives utility from the wealth that can be used to pay for consumption. Hence, the investor must formulate an optimal trading strategy constrained to starting and ending with zero shares.

The irrelevance of paper wealth is a concept also present in Bertsimas and Lo (1998) and He and Mamaysky (2005). They define a liquidity cost as the difference between the theoretical market value of a block of shares and the actual cash that can be realized after accounting for the price impact of selling that block. Brunnermeier and Pedersen (2005) further distinguish between an asset's paper value, its orderly liquidation value, and its distressed liquidation value.⁷

To summarize, the investor's maximization problem is formulated as follows:

$$\begin{aligned}
(9) \quad & \underset{\{N_t\}_{t=0}^T}{\text{maximize}} && E_0[u(W_T)] \\
& \text{subject to} && W_{t+1} = W_t(1+r) + N_t S_t [\mu + \lambda - r + \sigma \varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)] \\
& && S_{t+1} = S_t [1 + \mu + \lambda + \sigma \varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)] \\
& && \psi_{t+1} = \bar{\psi} + \rho(\psi_t - \bar{\psi}) + \varphi \varepsilon_{t+1}^\psi \\
& && N_0 = N_T = 0
\end{aligned}$$

The liquidity premium, λ , is such that the maximized expected utility in (9) is the

⁷Constantinides (1986), Heaton and Lucas (1996), and Vayanos (1998) analyze the asset pricing effects of transaction costs. The standard result in those papers is that while transaction costs cause the investor to reduce the trading frequency, they induce only a negligible utility loss, i.e., transaction costs have only a second-order effect on assets prices. However, those models produce counterfactual low trading volume. When investors desire to trade large amounts very frequently, transaction costs can have a significant effect. For example, Longstaff (2001) models a stock with stochastic volatility of returns, which induces more desired trading than in the standard portfolio-choice model due to the necessary portfolio rebalancing, thus making trading frictions relevant. Lo, Mamaysky, and Wang (2004) demonstrate that the impact of transaction costs is very large when heterogenous investors trade to hedge their exposure to an exogenous nontradable endowment risk. Jang, Koo, Liu, and Loewenstein (2007) show that investors have higher trading needs under a stochastic investment opportunity set and that trading costs can have a first-order effect. In our model, investors trade more than in the standard portfolio-choice model because they must start and end only with cash.

same as that in the standard perfectly liquid case.

III Numerical Results

Given that solving the full model with stochastic price impact is a challenging optimization problem, we first consider a simplified model with a constant price impact. We then solve the full model with stochastic price impact and discuss the implications on the relation between stock returns and volatility of liquidity.

A A Simplified Model with Constant Liquidity

We begin with the basic case of a constant price impact process, i.e., $\psi_t = \psi, \forall t$. The analysis of this problem allows us to gain intuition for the properties of the model and to verify that it produces sensible results. This analysis serves as the foundation for the main results in section B.

1 Calibration and Solution Method

For the baseline case we calibrate the drift and volatility of the stock price process to reflect the characteristics of very liquid stocks. During the period 1926–2002 a portfolio of the largest (top decile) NYSE stocks had an average annual return of 11% with an annual standard deviation of 18%. We allow the investor to trade every month and hence set $\mu = 0.11/12$ and $\sigma = 0.18\sqrt{1/12}$. We also set the monthly risk-free rate to $r = 0.05/12$. The baseline coefficient of relative risk aversion is set at $\gamma = 3$. We consider several levels of initial wealth W_0 : 10^4 , 10^5 , and 10^6 . We set $S_0 = 1$, thus the different W_0 can be interpreted as multiples of the initial stock price.⁸ Following

⁸For example, if the initial stock price is \$10, $W_0 = 10^5$ corresponds to an initial wealth of \$1 million.

Longstaff (2001), we consider investment horizons of either one or two years.

Several studies estimate price impact coefficients. Breen, Hodrick, and Korajczyk (2002) report that on average a 0.1% increase in net turnover during a 5-minute interval induces a 2.65% price increase for NYSE and AMEX listed firms and a 1.85% increase for NASDAQ firms. Given the average shares outstanding of 10 million for NYSE and AMEX listed firms, this corresponds to an average price impact coefficient of $\psi = 2.65 \times 10^{-6}$. In other words, a sale of 10,000 shares in a single block (representing 0.1% of shares outstanding in the average firm) moves the price down by 2.65%. Hasbrouck (2006) estimates the price impact of signed dollar volume aggregated over 5-minute intervals for a sample of 300 NYSE/Amex and Nasdaq firms. He finds that on average a \$10,000 buy order moves the price up by 28 basis points. In terms of our model, this corresponds to an average price impact coefficient of $\psi = 2.8 \times 10^{-6}$ for a trade of 1,000 shares in a stock priced at \$10. Çetin, Jarrow, Protter, and Warachka (2006) estimate the price impact coefficient in a regression very similar to our equation (5). Using a sample of five liquid NYSE stocks (each stock has an option trading on the CBOE), and considering only small transactions (they exclude trades larger than 1000 shares), they find price impact coefficients ranging from $\psi = 0.2 \times 10^{-6}$ to $\psi = 1.3 \times 10^{-6}$. These can be seen as lower limits to the magnitude of price impact. Hence, we report both the liquidity premium and optimal trading strategies for different degrees of price impact, ranging from $\psi = 1 \times 10^{-6}$ to $\psi = 5 \times 10^{-6}$ (which moves the price down by 1% to 5% for a sale of 10,000 shares in a single block). These values of the price impact coefficient cover a broad range of liquidity levels, from very liquid to very illiquid stocks.

Define x_t to be the vector of state variables known at time t . The optimal solution to problem (9) consists of: (1) a trading policy with $T + 1$ decision rules, $\{N_t(x_t)\}_{t=0}^T$, i.e., a sequence of *functions* mapping all future possible states (x_t) to the possible ac-

tions (number of shares to hold); and (2) the value function at time 0, representing the maximum expected utility given the state at time 0. This solution can be obtained through dynamic programming and is sometimes called a *closed-loop control* (see Bertsekas (2000)). However, the problem can also be solved by a suboptimal method known as the *open-loop feedback control* (see Bertsekas (2000)). With this method, after observing the state at time 0, the investor selects a sequence of actions as if no further information about the state will be received in the future. Hence, the open loop is suboptimal because it does not use the information about the state that will be available in the future, i.e., the policy is a single sequence of *numbers*: $N_t(x_t) = n_t$, for all states at time t . Nevertheless, this method usually provides a good approximation to the value function at time 0.⁹ Due to the specific characteristics of our problem, we use the open-loop solution as an initial approximation and then use the optimal closed-loop control to pin down the solution.¹⁰

2 Optimal Trading Strategy and Liquidity Premium

Let $\omega_t \equiv N_t S_t / W_t$ denote the optimal proportion of wealth invested in the stock. Figure 1 compares three different situations: (1) the standard trading strategy for a perfectly liquid stock (ω_t^* corresponding to the optimal solution N_t^* for $\psi = 0$ and $\lambda = 0$); (2) the trading strategy when there is price impact but no liquidity premium (ω_t^{**} for $\psi > 0$ and $\lambda = 0$); and (3) the strategy when the stock is illiquid and earns a premium (ω_t^{***} for $\psi > 0$ and $\lambda > 0$). We plot the optimal strategies along a representative path, namely the path where the disturbance is always at its expected value, $\varepsilon_t = 0, \forall t \in [0, T]$.

⁹Bertsekas (2000) states that the open-loop feedback control is a fairly satisfactory mode of control for many problems. Carlin, Lobo, and Viswanathan (2007) use an open-loop method to solve a trading game and study the effect of cooperation on liquidity. They argue that a closed-loop solution to their problem is not substantially different from their main open-loop solution.

¹⁰A separate appendix with a detailed discussion of our solution procedure is available upon request.

When the stock is perfectly liquid, the optimal equity proportion is the well-known Merton (1969) solution, $\omega^* = (\mu - r)/(\gamma\sigma^2)$. The investor immediately jumps to the optimal level and stays at that level until the last period when he liquidates his entire stockholding. When there is price impact, it is optimal to break a trade into several partial orders to obtain a lower average buying price or a higher average selling price. Consequently, the strategies ω_t^{**} and ω_t^{***} show that the optimal stock holdings slowly increase in the beginning of the investment period and then slowly decrease to zero by the terminal date. Strategy ω_t^{**} shows that with price impact (but no liquidity premium yet), the expected proportion of wealth invested in an illiquid stock is less than the standard Merton (1969) solution. Intuitively, with price impact the net drift of the stock price was smaller, making it less attractive relative to the bond. Finally, when the stock earns a higher expected return due to a liquidity premium, the investor chooses to hold a higher equity proportion than in the previous case: $\omega_t^{***} > \omega_t^{**}, \forall t$. The combination of the optimal trading strategy ω_t^{***} and the liquidity premium gives the investor the same expected utility as in the perfectly liquid case.

Table 1 presents the liquidity premium for different levels of price impact. The investment horizon (T) is either 12 months (panel A) or 24 months (panel B). In both cases, the investor may trade once each month. For an initial wealth of $W_0 = 10^5$ and an investment horizon of one year, the liquidity premium ranges from 3.11% (for $\psi = 1 \times 10^{-6}$) to 8.89% (for $\psi = 5 \times 10^{-6}$) per annum. In particular, for a price impact of $\psi = 3 \times 10^{-6}$, similar to the average level found in Breen, Hodrick, and Korajczyk (2002), the required liquidity premium is 6.41% per annum. For NYSE size sorted portfolios for 1926–2002, the average annual return on the lowest decile (18%) is approximately 7 percentage points higher than the return on the highest decile (11%). Since the price impact in small stocks is substantially higher than in large stocks (e.g.,

Ghysels and Pereira (2008)), these numbers give us a rough guideline for a maximum liquidity premium around 7%. Acharya and Pedersen (2005) decompose a total liquidity premium of 4.6% into a liquidity risk premium (1.1%) and a liquidity *level* premium of 3.5%.

The liquidity premium displays several features. First, it increases with the degree of illiquidity (or price impact, ψ). If the investor must pay a larger premium (concede a larger discount) when buying (selling) the stock, he demands a higher expected return.

Second, the liquidity premium (λ) is a concave function of the price impact coefficient (ψ). This concavity is consistent with the theoretical and empirical findings of Amihud and Mendelson (1986). However, while in the model of Amihud and Mendelson (1986) the concavity is generated by a clientele effect (investors with longer horizons require a smaller increase in the premium of illiquid assets), the concavity uncovered here is caused by a different reason. In our model, the investor is not constrained in the choice of the optimal number of shares, N_t . He can respond to an increase in price impact by demanding a higher liquidity premium, by reducing his optimal stock holdings, or both. As figure 2 illustrates, the optimal N_t decreases with ψ . By trading less, the investor reduces the adverse price impact of his trades. Hence, an increase in price impact leads to a less than linear increase in the liquidity premium because the investor reduces his total holdings of shares. The effect in our model is similar to that of proportional transaction costs on equilibrium asset returns uncovered by Constantinides (1986) and to the effect of fixed transaction costs in Lo, Mamaysky, and Wang (2004).

Third, an investor with more wealth demands a higher liquidity premium. Again, there are two simultaneous effects at work (though more complex in this case). As the top panel in figure 3 shows, increasing the initial wealth (W_0) reduces the equity proportion (ω_t). Still, the reduction in ω_t is not strong enough to compensate the

increase in wealth, and therefore the actual number of shares held by the investor, N_t , increases with wealth. Since the magnitude of price impact is proportional to changes in N_t , a larger investor induces larger (unfavorable) price changes and requires a higher liquidity premium. Therefore, our model suggests that large investors will tend to prefer liquid stocks, leaving the illiquid stocks to smaller investors who require a lower liquidity premium. This prediction has indeed been verified empirically by Falkenstein (1996). He finds that mutual funds show an aversion to small firms and that their demand is increasing in liquidity (measured by turnover).

Finally, a longer investment horizon leads to a smaller liquidity premium. The bottom panel in figure 3 shows that the optimal equity proportion increases with the investment horizon. The existence of price impact makes it optimal to trade a given quantity through several smaller partial orders. With a longer time horizon, the investor can achieve a higher peak holding, while still trading small blocks and therefore minimizing the adverse price impact. Furthermore, the investor enjoys the additional stock drift for a longer period of time. The investor thus requires a smaller liquidity premium.¹¹

B The Model with Stochastic Liquidity

This section analyzes the full model in (9) with stochastic price impact. Our focus is on the relation between the liquidity premium and the variability of price impact.

We demonstrate that our model offers a rational explanation for the negative relation

¹¹As a robustness check, we also study the effect of a *random* investment horizon: instead of the fixed known investment horizon, investors may face the possibility of an “emergency” liquidation of the stock. The results (not shown and available upon request) demonstrate that the liquidity premium increases strongly with the probability of an emergency liquidation. Uncertainty about the horizon induces a higher required liquidity premium because a quick unplanned sale of illiquid assets can induce large wealth losses (through the necessary price discounts the investor concedes to unwind his position in a short period of time). Hence, our simplified model with a random horizon is consistent with the results in Koren and Szeidl (2002) and Huang (2003).

between expected returns and the volatility of liquidity.

1 Calibration of the Model

The price impact coefficient follows the AR(1) process in (6): $\psi_{t+1} = \bar{\psi} + \rho(\psi_t - \bar{\psi}) + \varphi \varepsilon_{t+1}^\psi$. We consider several parameter values that make the results directly comparable with the basic model without stochastic liquidity. Specifically, we set the mean long-term price impact at $\bar{\psi} = 3 \times 10^{-6}$, one of the constant price impact coefficients used in section A. We explore several values for the volatility of liquidity, representing a range from relatively stable ($\varphi = 0.1 \times \bar{\psi}$) to very volatile ($\varphi = \bar{\psi}$) liquidity processes. We consider processes with high persistence (large ρ) and low persistence (small ρ). The risk aversion is set at $\gamma = 3$ and the initial wealth at $W_0 = 10^5$. The investment horizon is one year and the investor may trade once per month. We extend the state space for constant price impact to include the random price impact process. The details are described in a separate appendix available upon request.

2 Optimal Trading Strategy and Liquidity Premium

Table 2 presents the liquidity premium for different values of the volatility of liquidity. Our results indicate that introducing a stochastic price impact reduces the liquidity premium. The magnitude of the effect depends on the values of the conditional volatility (φ) and autocorrelation (ρ) parameters. For the case where ψ_t follows a simple white-noise process, i.e., $\rho = 0$, we find that the liquidity premium decreases quite substantially with the variance of liquidity: from the constant-liquidity benchmark case of 6.41% per year, to 5.69% for $\varphi = 0.5\bar{\psi}$, and further to 4.46% for $\varphi = \bar{\psi}$. When ψ_t is strongly autocorrelated, the reduction in the premium is smaller. For example, when $\rho = 0.9$ the liquidity premium is 6.41% for $\varphi = 0.1\bar{\psi}$, 6.29% for $\varphi = 0.5\bar{\psi}$, and 6.06%

for $\varphi = \bar{\psi}$. Further, for a given liquidity volatility, the liquidity premium increases with the autocorrelation of price impact. The effect is stronger at high levels of volatility.

The finding may seem contradictory to the usual risk-return tradeoff intuition and therefore deserves a careful analysis. Indeed, if the trading strategy $\{N_t\}_{t=0}^T$ was fixed, increasing the variance of ψ_t would increase the variance of the stock price, thus increasing the variance of terminal wealth, leading to a higher expected return to hold the stock. However, the trading strategy is not predetermined. $N_t(x_t)$ is a state-contingent function, namely contingent on ψ_t . The investor optimally chooses the number of shares to trade at time t knowing the price impact coefficient at that moment, ψ_t . Hence, the investor can adapt his trading to take advantage of periods of high liquidity (low ψ) and to ameliorate the adverse effects of periods of low liquidity (high ψ).

Figure 4 illustrates this “value of information” by plotting representative optimal trading strategies under two different scenarios: (1) liquidity suddenly increases (top panel); and (2) liquidity suddenly decreases (bottom panel). The top panel shows the trading strategy for the representative path where $\varepsilon_t = 0, \forall t \in [0, 12]$, with the stock becoming more liquid at time $t = 9$, that is, $\psi_t = \bar{\psi}, \forall t \in [0, 12] \setminus \{9\}$, and ψ_9 takes a smaller value of 0.53×10^{-6} , which is the minimum value in the discrete grid specified for the price impact coefficient. In this case, the investor takes advantage of the sudden increase in liquidity at $t = 9$ by selling a large block of shares at a “good price”, that is, suffering only a small price concession (N_9 is now below the corresponding share holding for the constant-liquidity case). Further, the more likely this sudden increase in liquidity is temporary (the lower the price impact persistence ρ) and ψ_t will soon revert to its mean, the more shares the investor decides to unload (see the strategy for $\rho = 0.2$). On the other hand, if the high-liquidity state is likely to persist (e.g., $\rho = 0.9$), we have two countervailing effects. While it is still advantageous to sell immediately with a low price

concession, it is now also advantageous to hold the shares longer, enjoying a positive stock drift longer and still being able to sell without incurring too much adverse price impact. Hence, when ψ_t is highly persistent, $\rho = 0.9$, the optimal action is to sell some shares at $t = 9$, but not as much as in the $\rho = 0.2$ case.

The bottom panel in figure 4 shows the opposite case: a sudden decrease in liquidity at time $t = 9$. We set $\psi_9 = 5.47 \times 10^{-6}$, which is the maximum value in the discrete grid specified for the price impact coefficient. The investor responds to this drop in liquidity by delaying the sale of shares. If the spike in ψ_9 is likely to be short-lived, i.e., if the autocorrelation is low, the investor may not sell any shares at $t = 9$. For the case where $\rho = 0.2$, the investor even tries to “manipulate” the price, that is, he buys more shares at $t = 9$, pushing the price up, thus increasing the value of the shares he currently holds. Since ψ_t is expected to subsequently revert back down to its mean, the investor expects to be able to sell all the shares later without having to offer a significant price discount. However, if the price impact coefficient is highly persistent (large ρ), thus expected to remain high in the future, the investor delays selling the shares somewhat, but not too much. Otherwise, he might have to sell a big block later on while still facing a high price impact coefficient (see the strategy for $\rho = 0.9$).¹²

To summarize, our model shows that a fully rational utility maximizing investor can take advantage of the volatility in liquidity. He is willing to hold more shares than when liquidity is constant; equivalently, he requires a lower liquidity premium. This result comes from the investor being able to time his trades to take advantage of periods of

¹²In a separate appendix (available upon request), we provide further evidence that it is the ability to adapt the trading strategy to the liquidity state that reduces the liquidity premium. Specifically, we solve the model using a suboptimal open-loop method, where information that will be available in the future is not taken into consideration. This leads to a liquidity premium that *increases* with the volatility of liquidity. In this case, the investor does not formulate contingent decision rules. His trading strategy formulated at time 0 is only optimal “on average.” The investor does not take advantage of timing his trades according to the state of liquidity. Hence, it is the fact that the investor considers all possible future information that causes the premium to decrease in volatility.

high liquidity and reduce the effects of periods of low liquidity.

IV Robustness and Extensions of the Model

This section checks the robustness of our main finding by examining the following extensions of the model: different degrees of risk aversion, higher trading frequency, the existence of temporary price impact, and correlation between stock returns and liquidity.

A Degree of Risk Aversion

The coefficient of relative risk aversion was set at $\gamma = 3$ in the baseline model above, as this is typically regarded as a reasonable value for the CRRA utility function. This section examines the effect of different levels of risk aversion ranging from 1 (representing the frequently used log utility) to 10 (usually considered an upper bound for reasonable values of γ). We consider different values for the volatility parameter (φ) while fixing the long-term mean and correlation of the price impact process.

Table 3 presents the resulting liquidity premium. For each level of the volatility of liquidity, the liquidity premium decreases in risk aversion. More risk-averse investors want to invest less in the stock and thus need to trade less. Therefore, they suffer less adverse price impact and hence require a smaller liquidity premium.

More important, for every level of risk aversion, we observe that the liquidity premium decreases with the volatility of price impact. Furthermore, the proportional reduction in the premium is very similar across different levels of risk aversion. Hence, we conclude that our main result of a negative relation between the volatility of liquidity and the liquidity premium is robust to the level of risk aversion of the investor.

B Analysis of the Trading Frequency

While in our baseline case investors trade only once a month, in reality they can trade much more frequently. For example, Chan and Lakonishok (1995) study the trading of large blocks of shares by institutional investors and find that only 20% of the trades are completed in one day. More than half of the dollar value traded by institutions takes four or more days of execution. Investors divide their trades into smaller partial orders to reduce the price impact of trading. Hence, it is interesting to assess the impact of different trading frequencies on the liquidity premium.

We modify our model to accommodate the trading patterns described in Chan and Lakonishok (1995). Specifically, we compute a new liquidity premium for three different trading patterns. In case 1, denoted “24 evenly spaced”, the investor can trade twice per month, for a total of 24 evenly spaced trades during the year. This case isolates the effect of increasing the number of trades within a given horizon. With more opportunities to trade, the required liquidity premium should decrease. In case 2, denoted “6 beg + 6 end”, we consider daily trading, i.e., the investor can trade in the first six and last six days of the one-year investment period, with no trading in between, thus totalling 12 trades. The investor trades once in each day, that is, his order for the day is executed as a single trade. Case 3, denoted “12 beg + 12 end”, compounds the effects of daily trading with more opportunities to trade: the investor can trade in the first and last 12 days, totalling 24 trades.

Table 4 shows the resulting liquidity premiums for different specifications of the price impact process. Panel A shows the simplified model with constant liquidity. Case 1, “24 evenly spaced”, shows a decrease in the required liquidity premium. For example, for the benchmark case of a price impact coefficient of $\psi = 3 \times 10^{-6}$, the liquidity premium is $\lambda = 4.58\%$. This is lower than the $\lambda = 6.41\%$ for the monthly trading setup reported

in table 1 because the investor is able to split his trading into smaller packets and thus suffer less price impact. Case 2, “6 beg + 6 end”, shows a more interesting decrease in the liquidity premium. For example, for $\psi = 3 \times 10^{-6}$, the liquidity premium decreases from the benchmark $\lambda = 6.41\%$ to $\lambda = 3.01\%$. In other words, allowing the investor to concentrate a given number of trades in the beginning and ending of the investment horizon, as opposed to having those trades equally spaced over the investment horizon, reduces the liquidity premium roughly in half. Intuitively, this reduction is due to the investor being able to build its target stock holding faster and enjoy the stock drift longer. Case 3, “12 beg + 12 end”, shows a further decrease in the liquidity premium: $\lambda = 2.18\%$. This is lower than in case 2 because the investor can trade for twice as many days and thus reduce the price impact of trading. Furthermore, it is also lower than in case 1, again due to the fact that the portfolio is built faster.

From these examples it might be tempting to conclude that there is a simple way to eliminate the liquidity costs altogether. In theory, the investor could follow a strategy of breaking his desired quantities into very small orders of a few shares traded every few minutes, which would induce an almost negligible price impact. However, in reality this is not the case for two reasons. First, one must take into consideration the costs of trading, such as brokerage fees. The costs can quickly overcome the benefits of frequent trading. Second, our analysis here assumes the same price impact coefficient regardless of the trading frequency. As documented in Dufour and Engle (2000), the price impact of trading increases as the time duration between transactions decreases. Trading the same quantity in the course of six days would induce a much larger adverse price impact than trading it in six months. A more appropriate comparison would therefore require a higher price impact for the daily trading setup. The resulting new liquidity premiums in table 4 would be higher.

Panel B in table 4 shows the full model with stochastic price impact. We fix $\bar{\psi} = 3 \times 10^{-6}$ and $\rho = 0.5$, and allow the volatility to vary from $\varphi = 0.1\bar{\psi}$ to $\varphi = \bar{\psi}$. The most important observation is that the liquidity premium still decreases with the volatility of liquidity. At low volatility, $\varphi = 0.1\bar{\psi}$, the liquidity premiums for the three trading cases are very close to the corresponding values in Panel A; the premiums then decrease as the volatility increases to $\varphi = \bar{\psi}$. The three cases show some interesting differences regarding the speed at which the premium decreases. Specifically, the schemes with 24 trading periods (case 1 and case 3) show a faster reduction than the one with 12 trading periods (case 2). With more discretion over when to trade, the investor is better able to move his trades to periods of high liquidity. To summarize, we conclude that the importance of the volatility of liquidity is robust to different trading frequencies.

C Temporary Price Impact

This section extends the baseline model to include temporary price impact. The temporary price effect of a trade consists of an adverse price movement that lasts only for that trade. Subsequently, the stock price reverts back to the level associated only with the permanent effect.

We model both the permanent and temporary price impacts by the following equation for the trade price (\hat{S}_{t+1}):

$$\hat{S}_{t+1} = S_t[1 + \mu + \lambda + \sigma\varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t) + \delta(N_{t+1} - N_t)]$$

where the parameter $\delta > 0$ measures the temporary price impact. The effect is proportional to the quantity traded: buying increases the trade price, while selling decreases it. We assume that market-making costs are relatively stable through time and therefore δ

does not depend on t .¹³

The timing is as follows. As in the baseline model, one instant before choosing N_{t+1} , the investor observes a middle price $\bar{S}_{t+1} \equiv S_t[1 + \mu + \lambda + \sigma\varepsilon_{t+1}]$. The price impact coefficients, ψ_{t+1} and δ , are also known now. The investor then chooses N_{t+1} . The difference to the baseline case is that trading occurs at \hat{S}_{t+1} . Immediately after the trade, the market price reverts back to S_{t+1} (as in equation 5).

Since the price reverts back to S_{t+1} , trading induces an immediate wealth loss corresponding to the block bought or sold at a, respectively, “too high” or “too low” price. For a buy trade the wealth reduces to $W_{t+1} = M_t B_{t+1} + N_t S_{t+1} - (N_{t+1} - N_t)(\hat{S}_{t+1} - S_{t+1})$, while for a sell trade it reduces to $W_{t+1} = M_t B_{t+1} + N_t S_{t+1} - (N_t - N_{t+1})(S_{t+1} - \hat{S}_{t+1})$. Both cases are nested in $W_{t+1} = M_t B_{t+1} + N_t S_{t+1} - |N_{t+1} - N_t| \cdot |\hat{S}_{t+1} - S_{t+1}|$. Using the specification for \hat{S}_{t+1} , subtracting $W_t \equiv M_t B_t + N_t S_t$, and replacing $M_t = (W_t - N_t S_t)/B_t$, we get the wealth process for the case of both permanent and temporary price impact:

$$W_{t+1} = W_t(1 + r) + N_t S_t [\mu + \lambda - r + \sigma\varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)] - \delta S_t (N_{t+1} - N_t)^2$$

Hence, the investor’s problem for the case of both permanent and temporary price

¹³Almgren and Chriss (2000), Huberman and Stanzl (2005), and Obizhaeva and Wang (2006) consider similar temporary price impact functions.

impacts is the following:

(10)

$$\text{maximize}_{\{N_i\}_{i=0}^T} E_0[u(W_T)]$$

subject to

$$W_{t+1} = W_t(1 + r) + N_t S_t [\mu + \lambda - r + \sigma \varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)] - \delta S_t (N_{t+1} - N_t)^2$$

$$S_{t+1} = S_t [1 + \mu + \lambda + \sigma \varepsilon_{t+1} + \psi_{t+1}(N_{t+1} - N_t)]$$

$$\psi_{t+1} = \bar{\psi} + \rho(\psi_t - \bar{\psi}) + \varphi \varepsilon_{t+1}^\psi$$

$$N_0 = N_T = 0$$

Table 5 shows the liquidity premium for different combinations of permanent and temporary price impacts. We start by considering the simplified model with constant permanent price impact. When the price impact is entirely temporary ($\psi = 0, \delta = 3 \times 10^{-6}$) rather than permanent as in the baseline model ($\psi = 3 \times 10^{-6}, \delta = 0$), the liquidity premium increases to 10.11%, which compares with 6.41% in the baseline model. Even when the total benchmark price impact is split evenly between the permanent and temporary components ($\psi = 1.5 \times 10^{-6}, \delta = 1.5 \times 10^{-6}$), the premium increases to 8.37%. These results are expected: the temporary impact increases the required liquidity premium because it causes the investor to lose wealth in each single trade.

The last three columns in Table 5 show the liquidity premium for the full model with stochastic liquidity defined in (10). Again, the existence of temporary price impact increases the premium. Most important, an increase in the volatility of the permanent component (φ) still decreases the total required liquidity premium. The reduction in the premium is nonetheless lower than in the benchmark case of permanent price impact only. This is to be expected as the total price impact is now split between two

components: the temporary value (δ) that is fixed through time and thus cannot be avoided; the permanent value (ψ_t) that the investor can still time, but now only represents a fraction of the total cost of trading. Hence, we conclude that our main result of a negative relation between the volatility of liquidity and stock returns is robust to the existence of both permanent and temporary price impacts.

D Correlation between Stock Returns and Liquidity

Our baseline model assumes that liquidity is uncorrelated with stock returns. But for some stocks, illiquidity may be negatively correlated with returns. For example, small stocks are typically harder to sell during market downturns. Furthermore, since liquidity is persistent, a positive shock to illiquidity predicts future high illiquidity, causing the contemporaneous stock price to decrease in order to raise expected future returns, thus generating a negative correlation between illiquidity and contemporaneous returns.¹⁴ This section analyzes the effect of such correlation.

Let \bar{R}_{t+1} denote the stock return due to the arrival of new public information, $\bar{R}_{t+1} \equiv \mu + \lambda + \sigma\varepsilon_{t+1}$. Table 6 shows the liquidity premium for different values of $corr_t(\bar{R}_{t+1}, \psi_{t+1})$.¹⁵ First, we observe that the liquidity premium is relatively insensitive to this correlation. As the correlation changes from the baseline case of zero to a large value of $corr_t(\bar{R}_{t+1}, \psi_{t+1}) = -0.3$, the liquidity premium changes only by a few basis points (less than 10bp for all the cases considered in the table). The reason for this result is the following. Consider first the later time periods when the investor is unwinding the portfolio. A negative correlation has two opposite effects: on the one hand, it means that low returns will be accompanied by high price impact, which hurts the

¹⁴See, for example, Acharya and Pedersen (2005), Pástor and Stambaugh (2003), and Amihud (2002).

¹⁵We impose the correlation between \bar{R}_{t+1} and ψ_{t+1} using the method described in a separate appendix available upon request. For the stocks in our data, the correlation ranges between -0.37 and +0.35, with a median of -0.07.

investor during this selling phase; on the other hand, it also means that high returns will be accompanied by low price impact, which benefits the investor. Our results suggest that the net effect is very small. Again, this is because the investor is able to time his trades, that is, he is able to move part of his sell orders from the bad low- \bar{R} and high- ψ days to the good high- \bar{R} and low- ψ days. Consider now the initial time periods when the investor is building up the portfolio. Here, the investor actually benefits from price impact as his trades increase the paper value of his current holdings. In this case, the negative correlation means that low price impact scenarios are compensated by higher stock returns, which benefits the investor. Therefore, we observe that a negative $corr_t(\bar{R}_{t+1}, \psi_{t+1})$ actually reduces the liquidity premium, even if by just a few basis points.

Finally, and most important, by comparing the three columns in table 6, we conclude that the correlation between returns and liquidity, be it negative or positive, does not change our main finding: the liquidity premium decreases with the volatility of liquidity. The reduction in the premium is similar for all levels of correlation.

V Empirical Tests of the Model

Our model has two main testable implications. First, in the cross section, expected stock returns should be negatively related to the volatility of liquidity. Hence, we follow Chordia, Subrahmanyam, and Anshuman (2001)(CSA) and start by showing that stock returns are negatively related to the volatility of price impact. Second, in the time series, our model suggests that price impact should be negatively related to trading activity. The fact that investors trade less (more) when there is less (more) liquidity is actually the mechanism that generates the first cross sectional result. Therefore, the second part of this section shows that higher price impact causes lower trading volume.

A Cross-sectional Relation between Stock Returns and the Volatility of Liquidity

1 Testing Methodology and Data

We follow the methodology of CSA, initially developed by Brennan, Chordia, and Subrahmanyam (1998). It allows us to test whether many different stock characteristics are related to stock returns, using the full cross section of all stocks. Specifically, we perform the following cross-section regression each month:

$$(11) \quad R_{jt}^* = \sum_{i=1}^I c_{it} Z_{ijt} + e_{jt},$$

where Z_{ijt} represents characteristic i for stock j in month t . We consider two alternatives for the dependent variable, R_{jt}^* . First, the simple excess return, $R_{jt}^* = R_{jt} - R_{ft}$, where R_{jt} is the return for stock j during month t and R_{ft} is the risk-free rate. Second, the risk-adjusted return using the Fama and French three-factor model, $R_{jt}^* = R_{jt} - R_{ft} - [\beta_{jM}(R_{Mt} - R_{ft}) + \beta_{js}\text{SMB}_t + \beta_{jh}\text{HML}_t]$, where the factor loadings $(\beta_{jM}, \beta_{js}, \beta_{jh})$ are estimated with the time-series regression $R_{jt} - R_{ft} = a_j + \beta_{jM}(r_M - r_f) + \beta_{js}\text{SMB} + \beta_{jh}\text{HML} + \varepsilon_j$ using stock returns and Fama and French factors for the prior 60 months.

As in CSA, we consider the following list of characteristics:

SIZE – the log of market capitalization (in \$ billions) at month $t - 2$.

BM – the log of the book-to-market ratio, using the previous year data.

PRICE – $\ln(1/P_{t-2})$, where P is the share price.

YLD – the dividend yield, defined as $\sum_{s=1}^{12} d_{t-1-s}/P_{t-2}$, where d_t is the dividend paid in month t .

RET2-3 – the cumulative return $\prod_{s=2}^3(1 + R_{t-s}) - 1$.

RET4-6 – the cumulative return $\prod_{s=4}^6(1 + R_{t-s}) - 1$.

RET7-12 – the cumulative return $\prod_{s=7}^{12}(1 + R_{t-s}) - 1$.

The basic data consists of monthly information for NYSE and AMEX common stocks for the period from 1963 to 2005. Returns, volume, prices, dividends, and shares outstanding are from CRSP and book values are from Compustat. To be included in the sample for a given month, a stock had to satisfy the following filters as in Pástor and Stambaugh (2003) and Acharya and Pedersen (2005). First, its price had to be between \$5 and \$1000. Second, it had to have at least 15 days of returns and volume in that month. Also, we discarded the first and last month of trading for each stock. To correct for outliers in each month, we exclude all stocks whose SIZE, BM, PRICE, or YLD is smaller than the 0.5 percentile or greater than the 99.5 percentile. This yields an average number of 1558 stocks per month. Panel A in Table 7 provides summary statistics on firm characteristics. The log transformations described above correct the considerable skewness of the raw data. The key variables are similar to those used in CSA.

2 Trading Activity

We first follow CSA and examine the relation between average stock returns and trading activity measured by the following variables:

DVOL – the log of dollar trading volume during month $t - 2$ (in \$ millions).

CV(DVOL) – the log of the coefficient of variation (ratio of standard deviation to the mean) of dollar volume computed over $t - 37$ to $t - 2$.

TURN – the log of share turnover during month $t - 2$.

CV(TURN) – the log of the coefficient of variation of turnover computed over $t - 37$ to $t - 2$.

While DVOL and TURN capture the effect of the level of liquidity on expected stock returns, CV(DVOL) and CV(TURN) are included to uncover the relation between stock returns and the volatility of liquidity. Table 8 shows the results. Consistent with the findings reported in CSA, the level of dollar volume and turnover have a significant negative effect on both excess and risk-adjusted returns. The estimated coefficients for the key variables of interest — CV(DVOL) and CV(TURN) — remain negative and highly significant both statistically and economically. This suggests that the negative relation between expected stock returns and the volatility of liquidity (based on measures of trading activity) is not specific to CSA’s sample period (1963-1995). Instead, it is a persistent feature of the data, robust to a longer and more recent sample period.

3 Price Impact

This section tests whether expected stock returns are also negatively related to the volatility of liquidity when liquidity is measured by price impact. We use the price impact measure introduced by Amihud (2002) — the ratio of absolute return to dollar volume. For each stock, we compute the monthly averages of the daily ratios of absolute return to dollar volume (times 10^6). We then use two different ways to characterize the level and volatility of this price impact series.

First, in line with CSA, we compute the following characteristics:¹⁶

PI-level – the log of price impact at month $t - 2$.

¹⁶Our model does not provide guidance on econometric issues, namely on the frequency and window over which the volatility of price impact should be measured. Our choice here follows CSA and Amihud (2002). Nonetheless, our findings are robust to different measuring frequencies and even to assuming that investors have perfect foresight and know the future volatility of price impact (results available upon request).

CV(PI-level) – the log of the coefficient of variation of price impact computed over $t - 37$ to $t - 2$.

Second, we fit an AR(1) specification as in (6) to the price impact series. However, we are faced with the following econometric issue: while an AR(1) is a stationary process, Amihud’s measure has dollar volume in the denominator and is thus nonstationary (due to inflation, at least). Therefore, following Pástor and Stambaugh (2003) and Acharya and Pedersen (2005), before fitting the AR(1) process, we first scale the price impact series through $(m_{t-1}/m_1) \times PI_t$, where PI is Amihud’s measure and m is the total market value of all stock included in the sample. While this transformation corrects to some extent the nominal increase in the denominator of Amihud’s measure, it does not necessarily correct for the secular increase in share volume over the sample period. As documented extensively in Lo and Wang (2000), even turnover is not stationary during our sample period. Following these authors, we abstain from further transformations of the price impact series, and instead use rolling windows of five years to perform our analysis. Specifically, for each stock and at each month t , we estimate an AR(1) over a 60-month rolling window (from $t - 61$ to $t - 2$). We then use the estimated mean ($\bar{\psi}$) as a proxy for the price impact level at month t and use the estimated conditional standard deviation (φ) as a proxy for the volatility of liquidity at month t . Hence, we include the following characteristics in equation (11):

PI- $\bar{\psi}$ – the log of the unconditional mean of price impact ($\bar{\psi}$ in equation (6)).

PI- φ – the log of the conditional volatility of price impact (φ in equation (6)).

PI- ρ – the autocorrelation of price impact (ρ in equation (6)).

Panel C in Table 7 reports summary statistics for the price impact variables. Both measures of the level of price impact (PI-level and PI- $\bar{\psi}$) share similar moment charac-

teristics. The measures of volatility (CV(PI-level) and PI- φ) also exhibit similar cross-sectional distributions. The price impact has an average first-order autocorrelation of 0.65 with a standard deviation of 0.15.

Table 9 reports the Fama-MacBeth estimates for regression (11) using price impact as a proxy for liquidity. We find that liquidity related variables, as well as all other firm characteristics, have the same signs and similar magnitudes as in the case based on trading activity. As expected, both measures of price impact (PI-level and PI- $\bar{\psi}$) have a positive effect on expected stock returns. More important, both measures of the volatility of liquidity (CV(PI-level) and PI- φ) have a negative effect on expected stock returns, implying that a higher volatility of liquidity is associated with a lower expected stock return. The estimated coefficients suggest that the negative effect of the volatility of liquidity on risk-adjusted returns is statistically and economically significant. The evidence on excess returns is slightly weaker. While the conditional standard deviation (PI- φ) remains statistically significant, the t-statistic on the coefficient of variation CV(PI-level) is slightly below conventional test levels.

In addition to the negative relation between the liquidity premium and the volatility of liquidity, our numerical analysis in section III also implies that the liquidity premium increases in the autocorrelation of price impact (ρ). The results in Table 9 suggest that the effect of the autocorrelation on returns is statistically insignificant. There are at least two possible reasons. First, the empirical estimate of the autocorrelation varies in a narrow range between 0.5 to 0.7 within one standard deviation. For small changes in ρ , the increase in the liquidity premium is small. This makes it difficult to empirically detect the true effect of the autocorrelation on the liquidity premium. Second, there are measurement errors in estimating the persistence of price impact. This measurement error emerges as the result of the econometric issues in estimating the AR(1) process

for the Amihud’s measure as discussed above.

Overall, our empirical analysis based on measures of price impact provides supporting evidence on the negative relation between expected stock returns and volatility of liquidity. This finding is consistent with the results based on using measures of trading activity as proxies for liquidity.

B Time-series Relation on Liquidity and Trading Activity

In our model, the investor is able to time his trades according to the state of liquidity (this is the crucial feature that generates a negative cross-section relation between expected returns and the volatility of liquidity). Empirically, we thus expect to observe time-series causality from liquidity to trading activity. In particular, we expect higher illiquidity (higher price impact) to lead to less trading.

One potential approach to test this hypothesis would be to analyze detailed information on individual trades by large investors. Some studies have used proprietary databases on institutional investors to investigate related questions. We therefore start by reviewing some of the existing literature on institutional investors. While these papers do not directly test our hypothesis, they do find supporting evidence. An alternative approach is to test the hypothesis using the time series of liquidity and trading activity observed in the market. We follow this approach and provide new evidence that price impact Granger causes trading activity.

1 Existing Literature on Institutional Investors

Some empirical evidence that investors adapt their trading to the state of liquidity comes from the survey by Economides and Schwartz (1995). They assess asset managers demand for immediacy by surveying traders of US equity funds. The vast majority (77%)

of respondents say they delay trades in hope of finding a better price for 25 to 75% of their trades (their table 3). Further, at least 51% of traders sometimes do not adjust their portfolios because the market is too illiquid (their table 14).

Keim and Madhavan (1997) find that transaction costs incurred by institutional investors depend on their investment style: “value traders have lower costs than index traders, who in turn have lower costs than technical traders”. They argue that indexers and technical traders have a strong demand for immediacy, while value traders “may incur lower costs because of more patient trading strategies involving limit or working orders”. Christoffersen, Keim, and Musto (2006) reach similar conclusions: “active managers, despite the presumably higher information content of their trades, trade at lower cost than index managers.” They conclude that “active managers add value both through patient trading and informed stock-picking.”

Foster, Gallagher, and Looi (2005) study transactions from Australian equity managers and conclude that “institutional investors are aware of liquidity in the market” and “choose to trade at times where the market is more accepting of unusually large flows”. Lipson and Puckett (2007) analyze proprietary institutional trading data and also find that “institutions who already wished to sell decide to sell more actively as markets are rising and those who already wished to buy decide to buy more actively as markets are falling”. They thus infer that institutions have “trading strategies based on long term price movements that also seek to minimize implementation costs by selling (buying) when there is increased demand (supply)”.

Chordia, Roll, and Subrahmanyam (2001) find that market liquidity declines and trading activity slows on Fridays, while Tuesdays display the opposite pattern. They rationalize the persistence of such day-of-the-week regularities with the adage that “liquidity begets liquidity”: “although a return anomaly is subject to arbitrage forces, a

liquidity anomaly is self-perpetuating; that is, as agents find out about such an anomaly, they will avoid trading in illiquid periods, which will further reduce liquidity in those periods.” In other words, it is rational for investors to move their trades, at least to some extent, from less liquid to more liquid days.

Hence, existing research on institutional investors suggests that they do adjust their trades to the state of liquidity in order to minimize transaction costs.

2 Vector Autoregression Results

We now directly test for causality between liquidity and trading activity. While our model provides one rationale for causality from price impact to trading activity, one should actually expect the data to show bidirectional causality. Several theoretical models predict a negative correlation between volume and illiquidity. In Kyle (1985), informed trading increases with uninformed trading, while illiquidity (Kyle’s λ) decreases with uninformed trading. Hence, higher volume is associated with lower illiquidity. Admati and Pfleiderer (1988) and Foster and Viswanathan (1990) extend Kyle’s model by allowing uninformed liquidity traders to have some discretion over when they trade. Both papers predict that trading costs are low when trading volume is high. Therefore, to allow bidirectional causal effects, we estimate a Vector Autoregression (VAR) for liquidity and trading activity, controlling for other relevant variables.¹⁷

We analyze time series data on the same sample of stocks as in the previous section, but with a daily frequency since we are interested in detecting very short term relations. At each day t , we aggregate all stocks into a “market” portfolio and compute the following quantities:

¹⁷Several papers, including Hasbrouck (1991), Hiemstra and Jones (1994), Chordia, Roll, and Subrahmanyam (2001), Chordia, Sarkar, and Subrahmanyam (2005a), and Chordia, Sarkar, and Subrahmanyam (2005b), have studied the time-series relation between several measures of liquidity, trading activity, and returns. However, they do not analyze the causality from liquidity to trading volume.

$\text{PrcImpact}_t := \frac{1}{N} \sum_{i=1}^N PI_{it}$, is the measure of portfolio price impact, where N is the total number of stocks trading on day t and PI_{it} is the price impact (ratio of absolute return to dollar volume) of stock i on that day.

$\text{DVolume}_t := \sum_{i=1}^N V_{it}$, is the portfolio total dollar volume, where V_{it} represents the dollar volume of stock i .

$\text{Turnover}_t := \sum_{i=1}^N w_{it} V_{it}$, is the portfolio weighted-average turnover, where w_{it} the weight of firm i based on its market capitalization.

$\text{RetVolat}_t := r_{pt}^2$, measures current market volatility, where r_{pt} is the portfolio value-weighted return.

$\text{RetPos}_t := \max(r_{pt}, 0)$ is the positive market returns.

$\text{RetNeg}_t := \min(r_{pt}, 0)$ is the negative market returns.

This definition of price impact for a portfolio (PrcImpact) follows Amihud (2002). Return volatility typically leads to more volume (e.g., Gallant, Rossi, and Tauchen (1992)). The decomposition of current returns into their positive (RetPos) and negative (RetNeg) components allows for a nonlinear relation between volume and returns, i.e., volume can be high both in positive and negative return days (see Chordia, Roll, and Subrahmanyam (2001)). Since a VAR model requires stationary series, we transform all the variables by taking logarithms and subtracting a four-week moving average (e.g., Ghysels and Pereira (2008)). Standard Augmented Dickey-Fuller tests reject a unit root in the transformed series.

The VAR model is thus

$$y_t = C + \sum_{l=1}^L A_l y_{t-l} + \varepsilon_t,$$

where y is a vector consisting of the five variables defined above: PreImpact, Trading (either DVolume or Turnover), RetVolat, RetPos, and RetNeg. To account for documented seasonal regularities in trading (e.g., Chordia, Sarkar, and Subrahmanyam (2005a)), each row in the matrix C includes a constant plus 4 dummies for day-of-the-week and 11 dummies for calendar month. The matrices A_l represent the coefficients to be estimated. There are 10,780 daily observations. We determine the number of lags in the VAR (L) using the Akaike Information Criterion and the Schwarz Bayesian Information Criterion, choosing the smaller value when the two criteria give different results. In all cases reported below, we use five lags ($L = 5$).¹⁸

Table 10 shows Granger causality tests for four different specifications. The table shows the F-statistics and corresponding P-values for a test on whether the five coefficients associated with the five lags of a given variable (y_{t-l}^i) are jointly zero in a particular equation of the VAR (y^j). If the null is rejected, we may infer that there is Granger causality from variable y^i to y^j . On the left panel, trading activity is measured by dollar volume. The first two columns are for a simple VAR with only two equations: price impact and dollar volume. There is evidence of strong bidirectional causality. Volume Granger causes price impact ($F = 34.4, P < 0.01$) and price impact also Granger causes volume ($F = 5.2, P < 0.01$). To correct potential biases due to missing variables, the next two columns present the full model with five equations. The evidence for bidirectional causality remains very strong. In particular, we find that price impact Granger causes dollar trading volume ($F = 4.1, P < 0.01$). The other three variables (RetVolat, RetPos, and RetNeg) also help to forecast volume and liquidity (all F tests are strongly significant). The right side panel of the table measures trading activity by turnover. The results are similar to dollar volume. Again, we find evidence of bidirec-

¹⁸We confirm that with five lags the estimated residuals in each VAR equation are not autocorrelated.

tional causality at conventional test levels. In particular, price impact Granger causes turnover ($F = 2.2, P = 0.05$, in the VAR with five equations.)

To better analyze the effect of liquidity on trading activity, figure 5 shows Impulse Response Functions (IRF) for those two variables. The IRF shows the response of one variable to a one-time, one standard-deviation, positive impulse (“shock”) to another variable. Even though the figure shows the response of just the two variables of interest (liquidity and trading), those responses are computed using the full VAR with five variables. The top panels in the figure show IRFs when trading activity is measured by dollar volume. The top left panel shows that price impact decreases when there is a positive shock to volume. More important, consistent with the implication from our model, the right panel shows that volume decreases when there is a positive innovation to price impact (solid thick line). The increase in price impact leads to a decrease in dollar volume over the next 5 days. The two bottom panels show the IRFs when trading is measured by turnover. The results are very similar, though slightly weaker. In particular, the bottom right panel shows that a positive innovation to price impact leads to less turnover over the next 4 or 5 days.¹⁹

Overall, these results show that there is Granger causality from price impact to trading activity. Together with the existing empirical evidence on institutional investors reviewed above, they support the hypothesis that investors trade less (more) when the market is more (less) illiquid.

¹⁹The IRF uses orthogonal innovations and thus its results are sensitive to the ordering of the variables. As noted in Chordia, Sarkar, and Subrahmanyam (2005a), one approach is to place the variables according to the order in which they influence the other variables. In our case, given the results from the model and from the literature reviewed above, the presumption is that liquidity influences trading, which suggests that price impact should be placed before trading activity.

VI Conclusion

We offer a rational explanation for the negative relation between expected stock returns and the volatility of liquidity first documented in Chordia, Subrahmanyam, and Anshuman (2001). While this finding may seem puzzling given the standard risk-return tradeoff intuition, our model demonstrates that it is consistent with the optimal behavior of a risk-averse utility-maximizing investor. The key feature of the model is that the investor can adapt his trading strategy to take advantage of periods of high liquidity and to ameliorate the effects of periods of low liquidity. Consequently, for a given level of average liquidity, investors actually benefit from volatility around that level of liquidity. This result is robust to several different parameterizations and model extensions.

Furthermore, we provide new empirical evidence that price impact Granger causes trading volume in the time series. This evidence supports the main prediction of the model — the liquidity premium decreases with the volatility of liquidity because investors adapt their trading strategy to the state of liquidity.

References

- Acharya, V. V. and L. H. Pedersen. “Asset Pricing with Liquidity Risk.” *Journal of Financial Economics*, 77 (2005), 375–410.
- Admati, A. and P. Pfleiderer. “A Theory of Intraday Patterns: Volume and Price Variability.” *Review of Financial Studies*, 1 (1988), 3–40.
- Almgren, R. and N. Chriss. “Optimal Execution of Portfolio Transactions.” *Journal of Risk*, 3 (2000), 5–39.
- Amihud, Y. “Illiquidity and Stock Returns: Cross-Section and Time-Series effects.” *Journal of Financial Markets*, 5 (2002), 31–56.
- Amihud, Y. and H. Mendelson. “Asset Pricing and the Bid-Ask Spread.” *Journal of Financial Economics*, 17 (1986), 223–249.
- Amihud, Y.; H. Mendelson; and L. H. Pedersen. “Liquidity and Asset Prices.” *Foundations and Trends in Finance*, 1 (2005), 269–364.
- Bertsekas, D. P. *Dynamic Programming and Optimal Control*, vol. 1. Belmont, Massachusetts: Athena Scientific, second edn. (2000).
- Bertsimas, D. and A. W. Lo. “Optimal Control of Execution Costs.” *Journal of Financial Markets*, 1 (1998), 1–50.
- Breen, W. J.; L. S. Hodrick; and R. A. Korajczyk. “Predicting Equity Liquidity.” *Management Science*, 48 (2002), 470–483.
- Brennan, M. J.; T. Chordia; and A. Subrahmanyam. “Alternative Factor Specifications, Security Characteristics, and the Cross-Section of Expected Stock Returns.” *Journal of Financial Economics*, 49 (1998), 345–373.

- Brennan, M. J. and A. Subrahmanyam. “Market Microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns.” *Journal of Financial Economics*, 41 (1996), 441–464.
- Brunnermeier, M. K. and L. H. Pedersen. “Predatory Trading.” *Journal of Finance*, 60 (2005), 1825–1863.
- Carlin, B. I.; M. S. Lobo; and S. Viswanathan. “Episodic Liquidity Crises: Cooperative and Predatory Trading.” *Journal of Finance*, 62 (2007), 2235–2274.
- Çetin, U.; R. Jarrow; P. Protter; and M. Warachka. “Pricing Options in an Extended Black Scholes Economy with Illiquidity: Theory and Empirical Evidence.” *Review of Financial Studies*, 19 (2006), 493–529.
- Chan, L. K. C. and J. Lakonishok. “The Behavior of Stock Prices Around Institutional Trades.” *Journal of Finance*, 50 (1995), 1147–1174.
- Chiyachantana, C.; P. Jain; C. Jiang; and R. Wood. “International Evidence on Institutional Trading Behavior and Price Impact.” *The Journal of Finance*, 59 (2004), 869–898.
- Chordia, T.; R. Roll; and A. Subrahmanyam. “Market Liquidity and Trading Activity.” *Journal of Finance*, 56 (2001), 501–530.
- Chordia, T.; A. Sarkar; and A. Subrahmanyam. “An Empirical Analysis of Stock and Bond Market Liquidity.” *Review of Financial Studies*, 18 (2005a), 85–129.
- Chordia, T.; A. Sarkar; and A. Subrahmanyam. “Liquidity Dynamics and Cross-Autocorrelations.” Working Paper, SSRN eLibrary (2005b).

- Chordia, T.; A. Subrahmanyam; and V. R. Anshuman. “Trading Activity and Expected Stock Returns.” *Journal of Financial Economics*, 59 (2001), 3–32.
- Christoffersen, S. K.; D. B. Keim; and D. K. Musto. “Valuable Information and Costly Liquidity: Evidence from Individual Mutual Fund Trades.” Working Paper, SSRN eLibrary (2006).
- Constantinides, G. M. “Capital Market Equilibrium with Transaction Costs.” *Journal of Political Economy*, 94 (1986), 842–862.
- Constantinides, G. M. and M. S. Scholes. “Optimal Liquidation of Assets in the Presence of Personal Taxes: Implications for Asset Pricing.” *Journal of Finance*, 35 (1980), 439–449.
- Datar, V. T.; N. Y. Naik; and R. Radcliffe. “Liquidity and Stock Returns: An Alternative Test.” *Journal of Financial Markets*, 1 (1998), 203–219.
- Dufour, A. and R. F. Engle. “Time and the Price Impact of a Trade.” *Journal of Finance*, 55 (2000), 2467–2498.
- Economides, N. and R. Schwartz. “Equity Trading Practices and Market Structure: Assessing Asset Managers’ Demand for Immediacy.” *Financial Markets, Institutions and Instruments*, 4.
- Falkenstein, E. G. “Preferences for Stock Characteristics As Revealed by Mutual Fund Portfolio Holdings.” *Journal of Finance*, 51 (1996), 111–135.
- Foster, F. and S. Viswanathan. “A Theory of the Interday Variations in Volume, Variance, and Trading Costs in Securities Markets.” *Review of Financial Studies*, 3 (1990), 593–624.

- Foster, F. D.; D. R. Gallagher; and A. Looi. “Institutional Trading and Share Returns.” Working Paper, SSRN eLibrary (2005).
- Gallant, A. R.; P. E. Rossi; and G. Tauchen. “Stock Prices and Volume.” *Review of Financial Studies*, 5 (1992), 199–242.
- Ghysels, E. and J. P. Pereira. “Liquidity and Conditional Portfolio Choice: A Nonparametric Investigation.” *Journal of Empirical Finance*, 15 (2008), 679–699.
- Glosten, L. R. and L. E. Harris. “Estimating the Components of the Bid/Ask Spread.” *Journal of Financial Economics*, 21 (1988), 123–142.
- Hasbrouck, J. “Measuring the Information Content of Stock Trades.” *Journal of Finance*, 46 (1991), 179–207.
- Hasbrouck, J. “Trading Costs and Returns for US Equities: Estimating Effective Costs from Daily Data.” Working Paper, New York University (2006).
- He, H. and H. Mamaysky. “Dynamic Trading Policies with Price Impact.” *Journal of Economic Dynamics and Control*, 29 (2005), 891–930.
- Heaton, J. and D. J. Lucas. “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing.” *Journal of Political Economy*, 104 (1996), 443–487.
- Hiemstra, C. and J. D. Jones. “Testing for Linear and Nonlinear Granger Causality in the Stock Price-Volume Relation.” *Journal of Finance*, 49 (1994), 1639–1664.
- Holthausen, R. W.; R. W. Leftwich; and D. Mayers. “Large-Block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects.” *Journal of Financial Economics*, 26 (1990), 71–95.

- Huang, M. “Liquidity Shocks and Equilibrium Liquidity Premia.” *Journal of Economic Theory*, 109 (2003), 104–129.
- Huberman, G. and W. Stanzl. “Optimal Liquidity Trading.” *Review of Finance*, 9 (2005), 165–200.
- Jang, B.-G.; H. K. Koo; H. Liu; and M. Loewenstein. “Liquidity Premia and Transaction Costs.” *Journal of Finance*, 62 (2007), 2329 – 2366.
- Keim, D. B. and A. Madhavan. “The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects.” *Review of Financial Studies*, 9 (1996), 1–36.
- Keim, D. B. and A. Madhavan. “Transactions Costs and Investment Style: An Inter-Exchange Analysis of Institutional Equity Trades.” *Journal of Financial Economics*, 46 (1997), 265–292.
- Korajczyk, R. A. and R. Sadka. “Pricing the commonality across alternative measures of liquidity.” *Journal of Financial Economics*, 87 (2008), 45–72.
- Koren, M. and A. Szeidl. “Portfolio Choice with Illiquid Assets.” Working Paper, Harvard University (2002).
- Kyle, A. S. “Continuous Auctions and Insider Trading.” *Econometrica*, 53 (1985), 1315–1335.
- Lipson, M. L. and A. Puckett. “Institutional Trading During Extreme Market Movements.” Working Paper, SSRN eLibrary (2007).
- Lo, A. W.; H. Mamaysky; and J. Wang. “Asset Prices and Trading Volume under Fixed Transactions Costs.” *Journal of Political Economy*, 112 (2004), 1054–1090.

- Lo, A. W. and J. Wang. “Trading Volume: Definitions, Data Analysis, and Implications Of Portfolio Theory.” *Review of Financial Studies*, 13 (2000), 257–300.
- Longstaff, F. A. “Optimal Portfolio Choice and the Valuation of Illiquid Securities.” *Review of Financial Studies*, 14 (2001), 407–431.
- Merton, R. C. “Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case.” *Review of Economics and Statistics*, 51 (1969), 247–257.
- Obizhaeva, A. and J. Wang. “Optimal Trading Strategy and Supply/Demand Dynamics.” Working Paper, MIT (2006).
- Pástor, Ľ. and R. F. Stambaugh. “Liquidity Risk and Expected Stock Returns.” *Journal of Political Economy*, 111 (2003), 642–685.
- Saar, G. “Price Impact Asymmetry of Block Trades: An Institutional Trading Explanation.” *Review of Financial Studies*, 14 (2001), 1153–1181.
- Sadka, R. “Momentum and Post-Earnings-Announcement Drift Anomalies: The Role of Liquidity Risk.” *Journal of Financial Economics*, 80 (2006), 309–349.
- Seppi, D. J. “Block Trading and Information Revelation Around Quarterly Earnings Announcements.” *Review of Financial Studies*, 5 (1992), 281–305.
- Sias, R. W.; L. T. Starks; and S. Titman. “The Price Impact of Institutional Trading.” Working Paper, Washington State University (2001).
- Vayanos, D. “Transaction Costs and Asset Prices: A Dynamic Equilibrium Model.” *Review of Financial Studies*, 11 (1998), 1–58.

Table 1: Liquidity Premium with Constant Liquidity

Liquidity premiums (λ) for different combinations of price impact (ψ) and initial wealth (W_0). The liquidity premium is the annual additional return (drift), in percentage points, the investor requires to hold the illiquid stock. In panel A, the investment horizon is one year with monthly rebalancing ($T = 12$ months); in panel B, two years ($T = 24$ months). All other parameters are set at the baseline values in section III.

Panel A: Horizon = 1 year					
W_0	Price impact ψ ($\times 10^{-6}$)				
	1	2	3	4	5
10^4	1.02	1.34	1.61	1.85	2.08
10^5	3.11	4.91	6.41	7.71	8.89
10^6	13.66	20.41	25.54	29.78	33.47

Panel B: Horizon = 2 years					
W_0	Price impact ψ ($\times 10^{-6}$)				
	1	2	3	4	5
10^4	0.68	0.80	0.89	0.97	1.05
10^5	1.42	1.99	2.52	2.97	3.47
10^6	5.33	8.18	10.5	12.43	14.08

Table 2: Liquidity Premium with Stochastic Liquidity

Liquidity premium, in percentage points per year, when the price impact follows the mean-reverting process in (6). The long-term mean is $\bar{\psi} = 3 \times 10^{-6}$, the initial wealth is $W_0 = 10^5$, and the horizon is one year with monthly rebalancing ($T = 12$ months). All other parameters are set at the baseline values in section III.

Correlation	Volatility			
	$\varphi = 0.1\bar{\psi}$	$\varphi = 0.5\bar{\psi}$	$\varphi = 0.75\bar{\psi}$	$\varphi = \bar{\psi}$
$\rho = 0.00$	6.39	5.69	5.07	4.46
$\rho = 0.20$	6.39	5.79	5.26	4.74
$\rho = 0.50$	6.40	5.94	5.54	5.15
$\rho = 0.90$	6.41	6.29	6.23	6.06

Table 3: Liquidity Premium for Different Degrees of Risk Aversion

This table shows the liquidity premium for different degrees of risk aversion (γ) for the full model with stochastic liquidity, with fixed $\bar{\psi} = 3 \times 10^{-6}$ and $\rho = 0.5$, and φ taking different values. The initial wealth is $W_0 = 10^5$, the horizon is one year with monthly rebalancing ($T = 12$ months), and all other parameters are set at the baseline values in section III.

Risk Aversion	Volatility		
	$\varphi = 0.1\bar{\psi}$	$\varphi = 0.5\bar{\psi}$	$\varphi = \bar{\psi}$
$\gamma = 1$	12.81	12.08	10.94
$\gamma = 2$	8.30	7.74	6.77
$\gamma = 3$	6.40	5.94	5.15
$\gamma = 5$	4.59	4.29	3.69
$\gamma = 10$	2.98	2.81	2.44

Table 4: Liquidity Premium for Alternative Trading Patterns

Liquidity premiums (λ) for different trading schemes. Panel A shows the simplified model with constant liquidity, for several price impact coefficients (ψ). Panel B shows the full model with stochastic liquidity, with fixed $\bar{\psi} = 3 \times 10^{-6}$ and $\rho = 0.5$, and with φ taking different values. Each panel shows three different trading patterns: (case 1) “24 evenly spaced”, the investor trades twice per month, for a total of 24 evenly spaced trades; (case 2) “6 beg + 6 end”, he trades in the first 6 days and last 6 days of the year, totalling 12 trades; (case 3) “12 beg + 12 end”, he trades in the first and last 12 days, totalling 24 trades. In all cases the horizon is one year, the initial wealth is $W_0 = 10^5$, and all other parameters are set at the baseline values in section III.

Panel A: Model with constant liquidity			
Trades	Price impact ψ ($\times 10^{-6}$)		
	1	3	5
(Case 1) 24 evenly spaced	2.37	4.58	6.17
(Case 2) 6 beg + 6 end	1.42	3.01	4.36
(Case 3) 12 beg + 12 end	0.95	2.18	3.08

Panel B: Model with random liquidity			
Trades	Volatility		
	$\varphi = 0.1\bar{\psi}$	$\varphi = 0.5\bar{\psi}$	$\varphi = \bar{\psi}$
(Case 1) 24 evenly spaced	4.45	2.53	0.49
(Case 2) 6 beg + 6 end	3.00	2.83	2.46
(Case 3) 12 beg + 12 end	2.16	1.50	0.37

Table 5: Liquidity Premium with Permanent and Temporary Price Impact

Liquidity premium, in percentage points per year, when there is both permanent and temporary price impact according to the model in (10). Each row shows the premium for different combinations of the permanent (ψ) and temporary (δ) effects. For the simplified model with constant liquidity in column 3, the constant ψ is set at the long-term mean ($\bar{\psi}$) for that row. For the full model with stochastic liquidity in columns 4–6, we fix $\rho = 0.5$ and allow $\bar{\psi}$ and φ to vary as indicated in the table. The initial wealth is $W_0 = 10^5$, the horizon is one year with monthly rebalancing ($T = 12$ months), and all other parameters are set at the baseline values in section III.

Price impact ($\times 10^{-6}$)		Model with constant liquidity	Model with random liquidity Volatility		
Permanent	Temporary	($\psi = \bar{\psi}$)	$\varphi = 0.1\bar{\psi}$	$\varphi = 0.5\bar{\psi}$	$\varphi = \bar{\psi}$
$\bar{\psi} = 0$	$\delta = 3$	10.11	—	—	—
$\bar{\psi} = 1.5$	$\delta = 1.5$	8.37	8.36	8.27	8.01
$\bar{\psi} = 3$	$\delta = 0$	6.41	6.40	5.94	5.15
$\bar{\psi} = 3$	$\delta = 1.5$	11.50	9.97	9.75	9.18

Table 6: Liquidity Premium with Correlation between Stock Returns and Liquidity

Liquidity premium, in percentage points per year, when the stock return due to public information (\bar{R}) is correlated with price impact. The premiums are for the full model with stochastic liquidity with $\bar{\psi} = 3 \times 10^{-6}$, $\rho = 0.5$, and different φ values as indicated in the table. The initial wealth is $W_0 = 10^5$, the horizon is one year with monthly rebalancing ($T = 12$ months), and all other parameters are set at the baseline values in section III.

$corr_t(\bar{R}_{t+1}, \psi_{t+1})$	Volatility		
	$\varphi = 0.1\bar{\psi}$	$\varphi = 0.5\bar{\psi}$	$\varphi = \bar{\psi}$
0.1	6.40	5.97	5.15
0.0	6.40	5.94	5.15
-0.1	6.38	5.91	5.11
-0.3	6.35	5.88	5.06

Table 7: Descriptive Statistics

Times-series averages of cross-sectional statistics for a monthly average of 1558 NYSE and AMEX stocks. The sample is from Jan/1966 to Dec/2005. The variables are as described in the text.

	Mean	St.Dev.	Median
A: Firm Characteristics			
SIZE	-1.770	1.887	-1.790
BM	-0.277	0.689	-0.212
PRICE	-2.753	0.935	-2.917
YLD	2.562	3.052	1.926
RET2-3	0.027	0.163	0.014
RET4-6	0.038	0.197	0.022
RET7-12	0.077	0.295	0.046
B: Measures of Trading Activity			
DVOL	1.602	2.298	1.727
CV(DVOL)	-0.401	0.449	-0.417
TURN	-3.535	0.992	-3.462
CV(TURN)	-0.570	0.410	-0.586
C: Measures of Liquidity			
PI-level	-1.584	1.954	-1.615
CV(PI-level)	-0.575	0.353	-0.590
PI- ψ	-1.681	1.849	-1.636
PI- φ	-1.610	0.450	-1.606
PI- ρ	0.655	0.150	0.665

Table 8: Fama-MacBeth Regression Estimates for Volume and Turnover

Fama-MacBeth regressions of stock returns on firm characteristics. For each panel, the dependent variable is excess returns in the first column and Fama-French risk-adjusted returns in the second column. The independent variables are as described in the text. The sample is from Jan/1966 to Dec/2005. All coefficients are multiplied by 100. T-statistics are in parenthesis.

	A: Volume		B: Turnover	
	Excess Ret.	Risk-adj. Ret.	Excess Ret.	Risk-adj. Ret.
interc	0.972 [2.47]	-0.011 [-0.04]	-0.033 [-0.06]	-1.439 [-4.35]
SIZE	0.039 [0.72]	0.069 [1.78]	-0.092 [-3.05]	-0.116 [-4.56]
BM	0.232 [4.39]	0.106 [2.44]	0.236 [4.44]	0.11 [2.51]
PRC	0.083 [0.85]	-0.052 [-0.61]	0.071 [0.73]	-0.065 [-0.76]
YLD	-0.007 [-0.44]	0.001 [0.13]	-0.005 [-0.33]	0.005 [0.44]
RET2-3	0.837 [3.13]	0.754 [2.73]	0.838 [3.13]	0.761 [2.75]
RET4-6	1.129 [4.65]	1.14 [4.74]	1.113 [4.58]	1.109 [4.60]
RET7-12	1.072 [7.70]	0.829 [5.06]	1.052 [7.56]	0.79 [4.83]
DVOL	-0.124 [-2.58]	-0.184 [-5.19]	—	—
CV(DVOL)	-0.269 [-3.89]	-0.375 [-5.38]	—	—
TURN	—	—	-0.135 [-2.78]	-0.199 [-5.53]
CV(TURN)	—	—	-0.309 [-4.94]	-0.372 [-5.63]

Table 9: Fama-MacBeth Regression Estimates for Price Impact

Fama-MacBeth regressions of stock returns on firm characteristics. For each panel, the dependent variable is excess returns in the first column and Fama-French risk-adjusted returns in the second column. The independent variables are as described in the text. The sample is from Jan/1966 to Dec/2005. All coefficients are multiplied by 100. T-statistics are in parenthesis.

	A: Simple PI		B: AR(1) parameters	
	Excess Ret.	Risk-adj. Ret.	Excess Ret.	Risk-adj. Ret.
interc	0.6 [1.51]	-0.537 [-2.11]	0.583 [1.67]	-0.459 [-1.61]
SIZE	0.021 [0.40]	0.055 [1.27]	0.006 [0.09]	0.029 [0.59]
BM	0.223 [3.91]	0.085 [1.82]	0.207 [3.46]	0.066 [1.33]
PRC	-0.016 [-0.21]	-0.164 [-2.29]	0.06 [0.79]	-0.044 [-0.60]
YLD	0.013 [0.73]	0.027 [2.10]	0.014 [0.77]	0.025 [1.92]
RET2-3	0.719 [2.32]	0.543 [1.81]	0.485 [1.51]	0.346 [1.11]
RET4-6	0.970 [3.57]	0.885 [3.35]	0.830 [2.95]	0.796 [2.93]
RET7-12	1.018 [6.25]	0.794 [4.62]	1.205 [6.71]	1.021 [5.67]
PI-level	0.104 [2.42]	0.136 [3.54]	—	—
CV(PI-level)	-0.129 [-1.45]	-0.227 [-2.63]	—	—
PI- $\bar{\psi}$	—	—	0.093 [1.97]	0.114 [2.50]
PI- φ	—	—	-0.147 [-2.16]	-0.269 [-3.90]
PI- ρ	—	—	0.032 [0.18]	-0.182 [-1.03]

Table 10: Granger Causality Tests from VAR Model

We estimate a VAR(5) with up to five variables: PrcImpact, Trading (either DVolume or Turnover), RetVolat, RetPos, and RetNeg. Each pair of numbers in the table (F-statistic and corresponding P-value) tests the null hypothesis that the five coefficients associated with the row variable are jointly zero in the VAR equation denoted in the top of the block. On the left (right) panel, trading activity is measured by dollar volume (turnover). The first two columns in each panel are for a simpler VAR with only two variables: price impact and trading activity. We compute “market” time series by aggregating NYSE and AMEX common stocks. The sample has 10,780 daily observations from 1963 to 2005.

	Trading = Dollar Volume				Trading = Turnover			
	F-stat	P-val	F-stat	P-val	F-stat	P-val	F-stat	P-val
Eqn: PrcImpact								
PrcImpact	227.7	0.00	186.1	0.00	272.5	0.00	221.3	0.00
Trading	34.4	0.00	32.0	0.00	23.3	0.00	22.9	0.00
RetVolat			4.1	0.00			3.6	0.00
RetPos			4.7	0.00			4.8	0.00
RetNeg			39.1	0.00			41.4	0.00
Eqn: Trading								
PrcImpact	5.2	0.00	4.1	0.00	2.8	0.02	2.2	0.05
Trading	636.1	0.00	535.4	0.00	628.7	0.00	528.7	0.00
RetVolat			7.0	0.00			7.6	0.00
RetPos			24.9	0.00			24.1	0.00
RetNeg			21.9	0.00			14.3	0.00
Eqn: RetVolat								
PrcImpact			3.0	0.01			4.4	0.00
Trading			1.1	0.36			1.2	0.33
RetVolat			1.7	0.13			1.4	0.22
RetPos			12.3	0.00			13.2	0.00
RetNeg			0.8	0.54			0.7	0.59
Eqn: RetPos								
PrcImpact			9.0	0.00			14.5	0.00
Trading			3.0	0.01			1.5	0.19
RetVolat			9.2	0.00			10.4	0.00
RetPos			91.1	0.00			86.3	0.00
RetNeg			118.3	0.00			115.9	0.00
Eqn: RetNeg								
PrcImpact			3.8	0.00			5.0	0.00
Trading			2.0	0.08			2.5	0.03
RetVolat			12.7	0.00			11.7	0.00
RetPos			24.7	0.00			25.6	0.00
RetNeg			67.0	0.00			66.8	0.00

Figure 1: Optimal Trading Strategy Along a Representative Path

Optimal percentage of wealth invested in an illiquid stock. The perfectly liquid case ($\psi = 0$) is given by $\omega^* = (\mu - r)/(\gamma\sigma^2)$. The strategy ω^{**} is the optimal solution when there is price impact ($\psi > 0$), but the liquidity premium is set to zero ($\lambda = 0$). The third strategy ω^{***} , determined jointly with a positive liquidity premium λ , allows the investor to attain the same level of expected utility as in the first case, even though now there is price impact ($\psi > 0$). These are closed-loop solutions for the representative path where $\varepsilon_t = 0, \forall t \in [0, T]$. The initial wealth is $W_0 = 10^5$, the price impact coefficient is $\psi = 3 \times 10^{-6}$, the investment horizon is one year with monthly rebalancing ($T = 12$ months), and all other parameters are set at the baseline values in section III.




Figure 2: Optimal Trading Strategy for Different Levels of Price Impact

Each line represents the optimal number of shares N_t^{***} for a different coefficient of price impact ψ as noted in the legend. They are the closed-loop solutions for the representative path where $\varepsilon_t = 0, \forall t \in [0, T]$. The initial wealth is $W_0 = 10^5$ and the investment horizon is 1 year with monthly rebalancing ($T = 12$ months). All other parameters are set at the baseline values in section III.




Figure 3: Optimal Trading Strategies for Different Parameters

Each line represents the optimal percentage of wealth invested in the illiquid stock when the stock is earning the optimal liquidity premium (strategy ω^{***}). These are closed-loop solutions for the representative path where $\varepsilon_t = 0, \forall t \in [0, T]$. In the first plot, each line assumes a different initial wealth W_0 , as noted in the legend. The investment horizon is 1 year with monthly rebalancing ($T = 12$). The second plot compares the optimal strategies for two different investment horizons, 1 year ($T = 12$) and 2 years ($T = 24$), both with monthly rebalancing. The initial wealth is $W_0 = 10^5$. In both panels the price impact is $\psi = 3 \times 10^{-6}$ and all other parameters are set at the baseline values in section III.

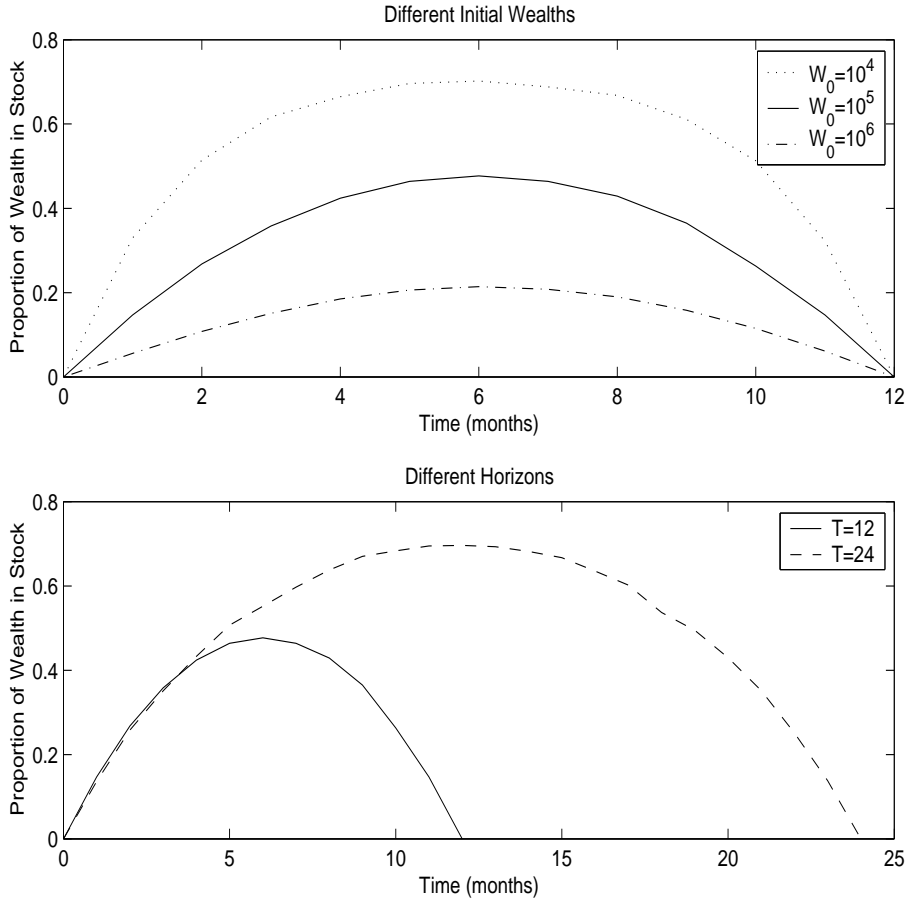


Figure 4: Examples of Trading Strategies with Liquidity Shocks

Each line represents the optimal closed-loop trading strategy under a representative path where $\varepsilon_t = 0$. In the first panel, the stock experiences a sudden increase in liquidity at time 9: $\psi_9 = 0.53 \times 10^{-6}$ and $\psi_t = \bar{\psi}, \forall t \in [0, 12] \setminus \{9\}$. The second panel displays the reverse, a sudden decrease in liquidity: $\psi_9 = 5.47 \times 10^{-6}$. The price impact coefficient ψ_t follows the mean-reverting process in (6) with $\bar{\psi} = 3 \times 10^{-6}$, $\varphi = \bar{\psi}$, and ρ as in the legend. In the “base” case liquidity is constant, $\psi_t \equiv \bar{\psi}$. In all cases, the stock is not earning a liquidity premium, $\lambda = 0$. The initial wealth is $W_0 = 10^5$, the investment horizon is one year with monthly rebalancing ($T = 12$ months), and all other parameters are set at the baseline values in section III.




Figure 5: Impulse Response Functions

The top (bottom) panels show Impulse Response Functions (IRF) when trading is measured by dollar volume (turnover). Even though the figure shows the response of just the two variables of interest (price impact and trading), those responses are computed with the full VAR with 5 variables, as detailed in the text.

