Predicting Equity Liquidity

William J. Breen • Laurie Simon Hodrick • Robert A. Korajczyk
Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208-2001
Graduate School of Business, Columbia University, 809 Uris Hall, 3022 Broadway, New York, New York 10027
Kellogg School of Management, Northwestern University, 2001 Sheridan Road, Evanston, Illinois 60208-2001
wilbreen@aol.com • *laurie.hodrick@columbia.edu* • *r-korajczyk@northwestern.edu*

In this paper we develop a measure of liquidity, price impact, which quantifies the change in a firm's stock price associated with its observed net trading volume. For a large set of institutional trades we compare out-of-sample, characteristic-based estimates of price impact to actual price impacts. Predictive predetermined firm characteristics, chosen to proxy for the severity of adverse selection in the equity market, the non-information-based costs of making a market in the stock, and the extent of shareholder heterogeneity, include relative size, historical relative trading volume, institutional holdings, and the inverse of the stock price. We find numerous aspects of trade execution which are significantly related to the price impact forecast error in economically plausible ways: For example, the predicted price impact overestimates the actual price impact for very large trades, for trades executed in a more patient manner, and for trades where the institution pays higher commissions.

(*Liquidity; Price Impact; Transactions Costs*)

In this paper we develop a measure of liquidity, price impact, which quantifies the change in a firm's stock price associated with its observed net trading volume. The systematic patterns of price impact across a large data set of firms allow us to predict equity liquidity out of sample. Price impact captures the extent to which trade execution influences the stock price: A perfectly liquid asset trades without any price impact while a perfectly illiquid asset cannot be traded at any price. Therefore, our measure encompasses important aspects of liquidity which are not captured by existing measures like the bid-ask spread and quoted depth.

Liquidity is a vital aspect of trading equities. While smaller equity trades are often executed inside the quoted prices (Petersen and Fialkowski 1994, Knez and Ready 1996), larger trades often face prices far inferior to those quoted. Such costs of trading securities are important for many reasons. In asset pricing, recent research suggests that asset prices are affected by expected trading costs. For example, Brennan and Subrahmanyam (1996) find that measures of liquidity are positively related to equity returns. Further, the inclusion of all relevant trading costs may erase apparent asset-pricing anomalies. For example, Knez and Ready (1996) find that trading costs eliminate the apparent profitability of a trading strategy based on autocorrelations. In portfolio management, better measures of expected trading costs improve the ability to effectively implement portfolio strategies and to monitor the quality of trade execution.

We find that increasing the magnitude of net turnover during a five-minute interval by 0.1% of the shares outstanding produces, on average, an incremental price effect of 2.65% for NYSE and AMEX-listed firms and 1.85% for NASDAQ firms. These averages, however, mask considerable cross-sectional variation. We relate the measured price impact to a set of predetermined firm characteristics, expecting liquidity to be influenced by (i) the severity of the adverse selection problem faced by uninformed traders, (ii) the non-information-based costs of market making, and (iii) the extent of shareholder heterogeneity. We find that liquidity varies cross-sectionally as predicted by these theories, and that the cross-sectional relation between price impact and these characteristics is stable over the sample period. This

Management Science © 2002 INFORMS
Vol. 48, No. 4, April 2002 pp. 470–483

0025-1909/02/4804/0470$5.00
1526-5501 electronic ISSN

cross-sectional relation is useful in predicting price impacts for other periods or for equities not used in the estimation stage. We compare the actual and predicted impact for a large sample of trades made by institutional traders. We find that numerous aspects of trade execution are significantly related to the price impact forecast error in economically reasonable ways. For example, the predicted price impact overestimates the actual price impact for very large trades, for trades executed in a more patient manner, and for trades where the institution pays higher commissions. Our cross-sectional results allow researchers back-testing portfolio-trading strategies to incorporate our estimate of the price impact. For example, Mitchell and Pulvino (2001) find that incorporating our price impact has an important effect on the profitability of a merger-arbitrage portfolio strategy.

1. Data and Methodology

The sample used to estimate price impact covers the period from January 1993 through May 1997 and consists of all firms included in the COMPUSTAT PST and OTC files. For each of these firms, we use the New York Stock Exchange's (NYSE) TAQ database to obtain net volume as well as price and quote information. We classify every trade in the sample as buyer initiated or seller initiated on the basis of whether the trade price is greater than or less than the midpoint of the prevailing best bid and ask quotes. Trades executed exactly at the midpoint are classified as neither buyer nor seller initiated and contribute zero to net turnover. As in Lee and Ready (1991), when a transaction occurs within five seconds of a quote revision we use the quote in the TAQ database as of five seconds prior to the trade. Odders-White (2000) finds that the "quote method" of trade classification used here has the lowest frequency of misclassification of the three methods she studies.

Once all of the trades are classified, we calculate net turnover for each 5- and 30-minute trading interval during the trading day.[1] Net turnover for period t (NTO_t) is defined as buyer-initiated volume less seller-initiated volume (times 1,000) as a fraction of shares outstanding. That is, $\mathrm{NTO}_t = 1$ corresponds to net turnover of 0.1% of shares outstanding. $\mathrm{P}_{i,\tau}$ is the price at which the last trade occurred within the time period τ and $\mathrm{Q}_{\mathrm{i}\tau}$ is the quote midpoint prevailing at the end of period τ. We calculate returns using both the percentage change in the last traded price and the percentage change in the end-of-period quote midpoints. We exclude intervals without any trades.

For each month, t, where $t = 1, \ldots, 53$, we have $\tau_{\mathrm{i}}(t)$ 5- or 30-minute observations on firm i, $i = 1, \ldots, N(t)$. We estimate, for each month, the regression of firm i's $\tau_i(t)$ equity returns on its corresponding net turnover.

$$(\mathrm{P}_{i,\tau} - \mathrm{P}_{i,\tau-1})/\mathrm{P}_{i,\tau-1} = \alpha_{it}^{\mathrm{P}} + \beta_{it}^{\mathrm{P}} \mathrm{NTO}_{i,\tau} + \epsilon_{i,\tau} \quad (1)$$

$$(\mathrm{Q}_{i,\tau} - \mathrm{Q}_{i,\tau-1})/\mathrm{Q}_{i,\tau-1} = \alpha_{it}^{\mathrm{Q}} + \beta_{it}^{\mathrm{Q}} \mathrm{NTO}_{i,\tau} + \epsilon_{i,\tau} \quad (2)$$

for $\tau = 1, 2, \ldots, \tau_i(t)$ and $i = 1, \ldots, N(t)$. This generates four time series of monthly price impact coefficients, $\hat{\beta}_{i,t}$, for every firm: two series (5- and 30-minute intervals) using price returns and two series using quote returns. The cross-sectional sample consists of 6,513 firms from January 1993 through May 1997, with a typical month having data for 3,699 firms. While the extant literature has investigated liquidity cross-sectionally, our sample includes a much larger cross-section of firms, a longer time series, and a larger array of explanatory variables. Hasbrouck (1991a, 1991b) investigates the relation between firm size and a measure of price impact for samples of 80 and 177 firms, respectively, during the first quarter of 1989. Glosten and Harris (1988) investigate the cross-sectional relation between temporary and permanent price effects and proxies for adverse selection- and nonadverse selection-based trading costs for a sample of 250 firms over the period from December 1981 through January 1983.

The specification of (1) and (2) is motivated by the linear pricing rule of Kyle (1985), which expresses price changes as a linear function of net volume:

$$\mathrm{P}_{i,\tau} - \mathrm{P}_{i,\tau-1} = \lambda_{i,t} \mathrm{NVOL}_{i,\tau} + \epsilon_{i,\tau}.$$

Notice that our specification scales net volume by shares outstanding to get net turnover and scales price changes by the beginning-of-period price to get returns. Using scaled measures provides more meaningful cross-sectional and intertemporal comparisons.

---

[1] Our sample does not include any overnight returns.

Table 1 **Time-Series Mean of the Cross-Sectional Distributional Statistics of the Price Impact Coefficient, 5-Minute Intervals, January 1993–May 1997, $T = 53$**

	NYSE/AMEX		NASDAQ	
	Beta Using Price Returns from (1)	Beta Using Quote Returns from (2)	Beta Using Price Returns from (1)	Beta Using Quote Returns from (2)
Mean, $\bar{\beta} = (\bar{\beta}_1 + \bar{\beta}_2 + \cdots + \bar{\beta}_T)/T$	3.09	2.15	1.96	0.25
Average Cross-Sectional Standard Deviation	26.69	14.08	6.16	0.77
Average Percentiles 15%	0.32	0.23	0.20	0.00
25%	0.48	0.36	0.36	0.01
75%	2.13	1.61	2.03	0.31
85%	3.25	2.45	3.14	0.49
Mean, $\bar{\beta}$ with High/Low Truncation	2.65	1.95	1.85	0.24
Average Cross-Sectional Standard Deviation with High/Low Truncation	10.94	8.07	3.67	0.54

For example, for firms engaging in stock splits, a given net volume corresponds to a larger fraction of the firm before the split than after the split. If the price reaction is solely a function of the *fraction* traded, then the coefficient λ would change around a stock split while β would not. The same would be true for cross-sectional comparisons of otherwise identical firms that have a different number of shares outstanding.

For similar reasons we use returns in (1) and (2) rather than price changes. The use of returns makes price or quote changes nonlinear functions of NTO across multiple periods.[2] Huberman and Stanzl (2000) have argued that the potential for profitable market manipulation exists if the permanent component of the price effect of trade is not linear in net turnover. Our impact measure is meant to measure both permanent and temporary price impacts; hence, nonlinearity does not imply arbitrage opportunities. From an empirical standpoint, using returns rather than price changes does not seem to make a significant difference in the results of Hasbrouck (1991b).

We expect that β^{P}_{it} will be larger than β^{Q}_{it} if fixed costs of trading are reflected in price returns. This is consistent with microstructure models which imply both permanent and transitory effects of trade on prices (Glosten 1987, Easley and O'Hara 1987). For example, if the bid and ask prices were constant through time with all transactions occurring at these quotes, quote returns would be identically zero, implying that $\beta^{\mathrm{Q}}_{it} = 0$, while β^{P}_{it} would be positive.

Table 1 contains summary statistics on the cross-sectional distribution of the price impact coefficient from Equations (1) and (2) using 5-minute intervals. Given the differences in the extent of inter-dealer trading between the exchanges and NASDAQ (Gould and Kleidon 1994), we report separate statistics for NYSE/AMEX versus NASDAQ firms. The average coefficients are 3.09% (exchanges) and 1.96% (NASDAQ) using price returns, and are 2.15% and 0.25% using quote returns. There is substantial cross-sectional dispersion, as can be seen in the average monthly standard deviations of 26.69% and 6.16% using price returns, and 14.08% and 0.77% using quote returns. Moreover, in all cases, the 85th percentile average coefficient is more than 10 times that of the 15th percentile value. Since coefficient outliers exist, we also report the mean price impact after truncating the sample.[3] This truncation leads to slightly

---

[2] Using (1) and the definition of returns we get that $\mathrm{P}_{i,\tau} - P_{i,\tau-1} = \mathrm{P}_{i,\tau-1}[\alpha^{\mathrm{P}}_{it} + \beta^{\mathrm{P}}_{it}\mathrm{NTO}_{i,\tau} + \epsilon_{i,\tau}]$. This implies that trading a given NTO over two periods has a different price effect than trading the same NTO over one period.

[3] In our truncation, we remove the high and low observations of $\beta_{i,t}$ each month. This eliminates two observations out of an average of 3,699 each month. Similar results are obtained using alternative

Figure 1 **Time Series of Average Price Impact Coefficient, Beta Net Turnover, High/Low Truncation: NYSE/AMEX Versus NASDAQ**

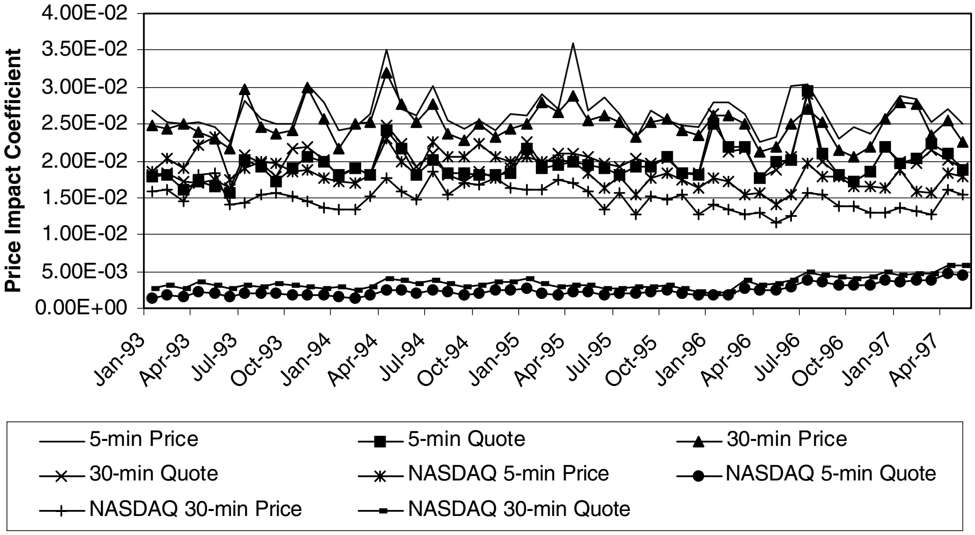

smaller means: 2.65% and 1.85% for Equation (1) and 1.95% and 0.24% for Equation (2). The cross-sectional standard deviation drops dramatically with the truncation. For Equation (1) the standard deviation drops from 26.69% to 10.94% and from 6.16% to 3.67%. Similarly, for Equation (2), the standard deviation drops from 14.08% to 8.07% and from 0.77% to 0.54%. All subsequent analyses are performed on the truncated sample.

The time series of average monthly coefficients for all four specifications are plotted in Figure 1. The average coefficients are relatively stable over time. Also, the price impact coefficients estimated using the last traded prices in the period are consistently larger than those estimated using the quote midpoints. This is consistent with microstructure models implying both permanent and transitory effects of trade on prices. This difference is found to be more pronounced for NASDAQ firms.

It is instructive to compare the magnitudes of our price impact measures to similar variables estimated in previous work. Hasbrouck (1991a, 1991b) models the reaction of quote midpoints to classified trading volume using exchange-traded firms, making his results most directly comparable to our quote specification in (2) for NYSE/AMEX firms. His average price impacts per 1,000-share trade are 0.299% and 0.255% (1991a, Table IV and 1991b, Table 1). In our sample, a 1,000-share trade corresponds to an average net turnover of 0.00919% for exchange-traded firms, so our average coefficient of 1.95% for turnover of 0.1% corresponds to a predicted impact of a 1,000-share trade equal to 0.179% [= (1.95%/0.1%) × 0.00919%], slightly lower than that in Hasbrouck (1991a, 1991b).

Glosten and Harris (1988, Table 2) report a permanent price effect per 1,000-share trade between \$0.0102 and \$0.0133 and a total (permanent plus transitory) price effect between \$0.0375 and \$0.0567 for firms traded on the NYSE. Given their average price of \$20.00 (Table 1), these dollar price impacts correspond roughly to a percentage impact between 0.05% and 0.07% (permanent) and between 0.19% and 0.28% (permanent plus transitory). While their estimate of the permanent price impact is considerably smaller than both our estimates and those of Hasbrouck (1991a, 1991b), their total price impact is similar to our NYSE firms' estimate of 0.244% using price returns.

2. Cross-Sectional Results

Equity liquidity can differ across firms at a point in time and across time for the same firm. To predict

---

truncations rules. Using a 20 standard-deviation truncation eliminates an average of 1.7 firms per month.

Table 2 Time-Series Mean of the Cross-Sectional Average and Cross-Sectional Standard Deviation of Firm Characteristics and Estimates of the Average Cross-Sectional Relation Between $\hat{\beta}_{i,t}$, Estimated Using Net Turnover (Equations (1) and (2)), and the Firm-Specific Predetermined Variables, Using 5-Minute Intervals, with Returns Defined Using Transactions Prices and Quote Midpoints: January 1993–May 1997

	NYSE/AMEX		NASDAQ		NYSE/AMEX				NASDAQ			
					$\bar{\Gamma}$, Prices: 5-Minute		$\bar{\Gamma}$, Quotes: 5-Minute		$\bar{\Gamma}$, Prices: 5-Minute		$\bar{\Gamma}$, Quotes: 5-Minute	
Explanatory Variable	$\bar{X}$	Std. Dev.	$\bar{X}$	Std. Dev.	$\bar{\gamma}_i$ ×100	*t*-Stat.	$\bar{\gamma}_i$ ×100	*t*-Stat.	$\bar{\gamma}_i$ ×100	*t*-Stat.	$\bar{\gamma}_i$ ×100	*t*-Stat.
Intercept	1.00	0.00	1.00	0.00	3.90	31.00	3.14	34.20	2.36	31.10	0.26	18.50
Relative Market Cap.	2.18	9.55	−0.24	3.27	0.24	8.69	0.20	7.18	0.18	9.34	0.03	12.60
Relative Trading Volume	0.10	2.37	−0.07	3.10	−0.58	−9.91	−0.47	−8.97	−0.14	−11.40	−0.02	−8.75
Price Appreciation	0.09	0.71	0.10	0.33	−0.65	−3.91	−0.84	−5.32	−0.63	−7.96	−0.04	−2.51
Price Movement	0.22	0.66	0.23	0.26	0.16	0.95	0.49	2.80	0.14	1.79	0.06	3.47
S&P Inclusion Dummy	0.19	0.39	0.03	0.17	−0.70	−7.70	−0.57	−6.97	−0.34	−7.48	−0.01	−1.50
Dividend Yield	0.02	0.04	0.01	0.02	1.27	1.35	−3.26	−4.24	14.60	9.15	0.29	1.46
R^2 Returns vs. NYSE	0.15	0.13	0.13	0.10	−2.34	−4.91	−2.17	−4.79	−1.92	−11.30	−0.03	−1.18
NYSE Inclusion Dummy	0.56	0.49			0.28	3.83	0.02	0.39				
Earnings Release Dummy	0.42	0.46	0.40	0.45	0.72	4.52	0.55	4.22	0.05	1.32	−0.01	−1.12
Percentage Institutional	33.30	24.73	29.70	21.70	−0.06	−19.90	−0.04	−17.40	−0.02	−27.50	−0.00	−8.22
Option Traded Dummy	0.30	0.46	0.14	0.34	−0.03	−0.46	−0.03	−0.47	0.09	4.05	0.14	11.30
Inverse Price	0.14	0.40	0.11	0.19	2.68	13.60	0.86	9.26	3.15	9.84	0.06	1.62

liquidity, we consider a set of predetermined firm characteristics that proxy for (i) the severity of the adverse selection costs faced by uninformed traders transacting with informed traders (Glosten and Milgrom 1985), (ii) the non-information-based costs of market making (Stoll 1978), and (iii) the extent of shareholder heterogeneity (Bagwell 1992). The first four columns of Table 2 provide the time-series means of the cross-sectional average and standard deviation for each of the firm characteristics. The variables included in the cross-sectional analysis are as follows.

(1) The first variable is a constant.

(2) The second variable is the relative market capitalization of the firm, measured as the market capitalization of the firm's common equity at the end of the previous month divided by the average market capitalization of firms in the CRSP (NYSE/AMEX/NASDAQ) index, minus one. A firm with a market capitalization larger (smaller) than the typical CRSP firm would have a positive (negative) relative size.

(3) The third variable is the historical relative trading volume of the firm's equity, measured as the total trading volume from the previous three months divided by the trading volume, over the same three months, of the average firm traded on the NYSE, minus one. Equities trading with high volume should pose less of an inventory risk to dealers since their expected holding period is shorter, and thus these equities should be more liquid. Alternatively, causality might run in the opposite direction: Assets for which there is greater liquidity should be traded more frequently. These arguments suggest that liquidity should be positively related to volume.

(4) The fourth variable is the recent price appreciation of the stock, measured as the firm's stock price at the end of the previous month divided by the price six months prior, minus one. Bagwell (1991) develops a model wherein cost basis value heterogeneity influences liquidity through the extent of "locked-in" capital gains. In her model, holders of appreciated assets have a disincentive to sell those shares because of locked-in capital gains, while holders of depreciated assets have an incentive to sell and realize capital losses. The supply of tax-motivated traders is expected to be smaller in the former case and larger in the latter, predicting that liquidity should be negatively related to appreciation.

(5) The fifth variable is recent price movement, measured as the absolute value of the previous

variable, $|(P_{t-1}/P_{t-7}) - 1|$. Both gains and losses might induce non-information-based trading due to portfolio rebalancing toward target proportions (Constantinides 1986). This suggests that liquidity and price movement should be positively related.

(6) The sixth variable is a dummy variable equal to unity if the firm is included in the S&P 500 Index and equal to zero otherwise. The S&P 500 portfolio is a popular benchmark for passive index funds. Harris and Gurel (1986) find permanent increases in trading volume for firms added to the Index driven by increased institutional demand. Harford and Kaul (1998) also find evidence suggesting an increase in firm liquidity following addition to the Index.

(7) The seventh variable is the dividend yield on the firm's stock, measured as the most recent indicated annual dividend (Quarterly COMPUSTAT item number 20) divided by the share price at the end of the previous month. Dividend capture, a non-information-based motive for trade, is less costly to implement when the yield is high.[4] If there is less adverse selection risk when the fraction of uninformed traders in the stock is large, then a negative relation between the price impact and dividend yield is expected.

(8) The eighth variable is the percentage of the firm's return variance explained by the return on an aggregate stock market portfolio. We use the coefficient of determination, R^2, from a regression of the firm's monthly percentage stock price change on the monthly percentage change in the NYSE index, estimated over the previous 36 months. Some of the price impact may be compensation to market makers for bearing inventory risks, as argued in Scholes (1972), because firm-specific (idiosyncratic) risks should be more difficult to hedge. Also, insiders are more likely to have firm-specific private information, making the severity of the adverse selection risk increasing in the level of firm-specific risk. Finally, the dividend capture discussed above is more easily hedged when firm-specific risk is lower. We predict a smaller price impact when R^2 is larger.

(9) The ninth variable is a dummy variable equal to unity if the firm is traded on the NYSE and equal to zero if traded on the AMEX (this variable is not relevant for NASDAQ firms). Liquidity differences across exchanges could be due either to actual differences in liquidity provision or to different listing requirements which are correlated with liquidity. There is a large literature investigating the effects of different market structures on liquidity (Marsh and Rock 1986, Hasbrouck and Schwartz 1988, Huang and Stoll 1996).

(10) The tenth variable is equal to unity if the firm's last earnings release was more than two months ago (and less than 40% of the other firms have at least a two-month reporting lag) or if the firm's last earnings release was more than three months ago and equal to zero otherwise. Korajczyk et al. (1992) derive a model in which the precision of the managers' private information increases between regularly scheduled information releases such as earnings. Immediately after a release, the adverse selection problem is small. As time passes since an information release, the asymmetry of information gets larger, implying a larger price impact.

(11) The eleventh variable is the percentage of the firm's equity held by institutional investors at the end of the previous quarter, in percentage points (i.e., the institutional holding data for March are used in the cross-sectional regressions for April, May, and June). Hodrick (1999) argues that institutional holdings proxy negatively for both taxation- and information-induced illiquidity.

(12) The twelfth variable is a dummy variable equal to unity if options are traded on the firm's equity and equal to zero otherwise. While dealers may demand a smaller price concession if firm-specific risk can be hedged through option transactions, it may instead be the case that price impact is larger for stocks with traded options if order fragmentation keeps a market maker from determining whether the initiator of the trade is also trading in the options market (Bernhardt and Hughson 1997).

(13) The thirteenth variable is the inverse of the stock price from the previous month. Some of the

---

[4] Tax-motivated dividend capture strategies, designed to take advantage of the dividend-received deduction, need to be held for a minimum of 46 days during the 90-day period beginning 45 days before the stock goes ex-dividend (U.S. Department of the Treasury 1999). We therefore expect non-information-based trading to be high in both ex-dividend and non-ex-dividend months for high-yield firms.

measured price impact of trades using price returns might be due to bid-ask-bounce, and we seek to avoid having this effect attributed to the dividend yield variable. Given the discreteness of bid-ask spreads, price impact would be more pronounced for lower-priced stocks.

For each month, t, we estimate the cross-sectional relation between the estimates, $\hat{\beta}_{i,t}$, and the firm-level characteristics:

$$\hat{\beta}_t = X_t\Gamma_t + \nu_t \tag{3}$$

where $\hat{\beta}_t$ is the $N(t)\times 1$ vector of estimated coefficients for month t and X_t is the $N(t)\times 13(N(t)\times 12)$ matrix of predetermined firm characteristics used to explain differences in $\hat{\beta}_{i,t}$ for NYSE/AMEX (NASDAQ) firms. As in Fama and MacBeth (1973), we use the time series of estimates $\widehat{\Gamma}_1, \widehat{\Gamma}_2, \ldots, \widehat{\Gamma}_T$ $(T=53)$ to estimate the average $\overline{\Gamma} = (\widehat{\Gamma}_1 + \widehat{\Gamma}_2 + \cdots + \widehat{\Gamma}_T)/T$ and to get standard errors for the elements of $\overline{\Gamma}$. We estimate Equation (3) each month using ordinary least squares (OLS).[5]

Let $\hat{\gamma}_{j,t}$ denote the jth element of $\widehat{\Gamma}_t$. The point estimates, $\bar{\gamma}_j$, and associated t-statistics are presented in the last eight columns of Table 2. Since autocorrelation-consistent standard errors (Newey and West 1987) are close to noncorrected standard errors, we report t-statistics in Table 2 that are not adjusted for autocorrelation in the estimates. For NYSE/AMEX firms, all of the cross-sectional coefficients are statistically significant at the 5% level except for the dividend yield, option dummy, and price movement variables. For NASDAQ firms, all variables are significant except for the earnings release and price movement variables. The average monthly R^2 is 9.2% for the NYSE/AMEX subsample and 10.5% for the NASDAQ subsample.

The magnitudes of the coefficient estimates suggest significant cross-sectional variation in liquidity. For exchange firms, the coefficient of 0.24 for the relative market capitalization implies that if a firm's relative size were to increase by one (cross-sectional) standard deviation, then the coefficient would increase by 2.28%. This positive relation stems from measuring the coefficient per unit of net turnover: As firm size increases, the size of a trade which represents a given percent turnover also increases. If we instead define volume as a fixed number of shares, we find that the relation between size and the coefficient is significantly negative as found in Hasbrouck (1991a, 1991b).

The coefficient of −0.58 for relative trading volume implies that if a firm's trading volume in the previous quarter were to increase by one standard deviation, then the coefficient would decline by 1.38%. These findings are consistent with the conventional wisdom that the stocks of firms with greater trading volume tend to trade with more liquidity.

The coefficient of −0.65 for price appreciation implies that if a firm's price appreciation over the previous six months were to increase by one standard deviation, then the coefficient would decline by 0.46%. This seems consistent with the finding in Brown and Ryngaert (1992) that tendering rates in repurchase offers are greater following a larger run-up in share prices and inconsistent with the hypothesis that firms with greater price appreciation are less liquid due to the "lock-in" induced by capital gains taxation. Odean (1998) finds evidence that investors have a greater tendency to sell assets that have appreciated rather than those that have depreciated. If this behavior results in a sufficiently large source of non-information-based traders following appreciation, it would be consistent with our observed increase in liquidity for those firms.

Firms included in the S&P 500 Index have significantly lower price coefficients, consistent with the documented increase in institutional demand following inclusion in the Index. The coefficient of −2.34 on the R^2 variable implies that were it to increase by one standard deviation, then the coefficient would decline by 0.29%. This is consistent with both adverse selection and nonadverse selection costs of making a market. Firms trading on the New York Stock Exchange have higher price impact coefficients than do American Stock Exchange firms. The coefficient of 0.72 on the earnings release dummy variable implies that firms that have not released earnings recently have a higher price impact, consistent with the severity of adverse selection increasing with the time since an earnings release.

---

[5] We also estimate (3) using weighted least squares, where the weights are inversely proportional to the standard error of the residuals from the first-stage regression (1). We find that the results are qualitatively similar and do not report them here.

The coefficient of -0.06 on the percentage of stock held by institutions implies that if a firm's institutional holdings were to increase by one standard deviation, then the price impact would decline by 1.41%. As in Bagwell (1992), this might either reflect that institutional holders choose to hold more liquid stocks or that institutional holdings affect liquidity. The coefficient of 2.68 on the inverse price variable implies that a one standard deviation increase would lead to an increase in the price impact of 1.06%.

For firms traded on NASDAQ, the coefficients have the same sign as for the exchange-traded firms with one exception: The coefficient for the option dummy is positive for NASDAQ. The late earnings release variable becomes insignificant (at the 5% level) for NASDAQ firms while the dividend yield and option-traded variables become significant.

We perform many tests of the robustness of the results.[6] The cross-sectional results are virtually identical whether we use price impact coefficients calculated from transactions price returns or quote midpoint returns, and whether we use 5- or 30-minute time intervals. We also find intraday patterns in the price impact: The coefficient is larger during the first and last half hour of trading versus the rest of the day. Consistent with the literature, we find that there are small but statistically significant nonlinearities in the relation between returns and NTO, and we find that net buyer-initiated volume has a larger price impact, on average, than net seller-initiated volume. Finally, we find that the price impact coefficient is positively related to bid-ask spreads but that spreads are not a sufficient statistic for the cross-sectional variation in price impact.

3. Out-of-Sample Impact Prediction

There are many important applications for which an estimate of the expected price impact would be extremely useful. While the direct costs of commissions are the easiest component of transactions costs to observe, the price impact is likely to be much larger for sizable trades (Loeb 1983). For many documented "anomalies" where profitable trading strategies seem to exist, it is vital to quantify the price impact that would be incurred by the trades. Moreover, portfolio managers may wish to evaluate the trade execution provided by a broker. An estimate of the expected price impact, accounting for cross-sectional variation in liquidity, would provide a benchmark against which actual trade execution can be compared.

For a given hypothetical trade of size NTO in asset i, price impact can either be estimated directly, $\hat{\beta}_{i,t} \times \text{NTO}$, or out of sample from the predicted value obtained from the cross-sectional Regressions (3), $X_{i,t}\overline{\Gamma} \times \text{NTO}$. While the direct estimate has the advantage of incorporating any firm-specific factors not explained by the cross-sectional regression, the predicted approach can be applied to data not covered by the original sample. For example, Mitchell and Pulvino (2001) use our cross-sectional regression coefficients to estimate the price impact incurred in a merger arbitrage trading strategy over a period far longer than that covered by our sample.

There are many reasons why the predicted and actual price impacts may differ. Our price impact measure implicitly assumes that a trade of size NTO is executed during the observation interval, though large orders are often split up and not executed in one trade (Kyle 1985, Back 1992, Bertsimas and Lo 1998). For orders that are broken up, the actual price impact should be less than that predicted. Similarly, while the predicted impact implicitly assumes that market orders are used to execute the trade, patient traders may choose to execute the trade with an alternative order, such as a limit order, in which case the actual price impact should be less than that predicted.

We compare the price impact predicted out of sample, $X_{i,t}\overline{\Gamma} \times \text{NTO}$, to the actual price impact for a sample of trades executed by 21 institutional traders over the period from January 1991 to March 1993. The data, collected by the Plexus Group and including detailed information on the equity trades by these institutions, are those from Keim and Madhavan

---

[6] The detailed results are not reported here to conserve space but are available at ⟨http://papers.ssrn.com/sol3/papers.cfm?abstract_id = 255131⟩.

(1995, 1997).[7] The comparison is out of sample in that the actual equity trades occur over a period with very little overlap with the period used to estimate $\hat{\beta}_{i,t}$ and $\overline{\Gamma}$. The actual price impacts provide a useful benchmark against which our out-of-sample impact predictions can be compared, and we investigate how various aspects of the transaction relate to the forecast error.

While the Plexus data set does not provide the identity of the participating institutions, Plexus does classify the institutions by investment style: value, technical (momentum), and indexing. We know (a) the type of order placed including an index of the urgency of the order; (b) the price of the stock both prior to the decision to trade and prior to sending the order to the trading desk; (c) the number of shares sought and the direction of trade (i.e., a buy vs. a sell); (d) the average trade price and the actual number of shares traded; (e) the dates at which the trading decision was made, the order was sent to the trading desk, trading began, and the order was completed. While we can determine the number of days that the broker took to execute the trade, we cannot determine the number of separate trades used to execute the trade, because the institutions only receive information on the total shares traded and the average price within a day. We exclude trades for which the time between either the decision date and the desk date, the desk date and the beginning of trade, or the beginning and completion of trade exceeds three weeks. The results are not sensitive to this 21-day screen. After matching the Plexus data with our sources for $X_{i,t}$, we have data from 56,211 trades for a total of 1.9 billion shares representing \$62.4 billion in value.

On average there are 1.9 days between making the decision to trade and sending the order to the trading desk. The time between the receipt of the order by the trading desk and the initiation of trade averages 1.21 days. The average number of days between the beginning and end of trading is 0.37. Most orders (92.5%) are filled on the same day that trading begins.

We calculate two measures of actual price impact. The first is the average trade price minus the closing price the day before the trade is sent to the trading desk, expressed as a fraction of the prior closing price:

$$I_{\text{Desk}} = 100 \times \frac{\overline{\mathrm{P}} - \mathrm{P}_{\text{Desk}}}{\mathrm{P}_{\text{Desk}}}.$$

The second is the average trade price minus the closing price the day before the decision to trade:

$$I_{\text{Decide}} = 100 \times \frac{\overline{\mathrm{P}} - \mathrm{P}_{\text{Decide}}}{\mathrm{P}_{\text{Decide}}}.$$

If there is no leakage of information about the trade between the decision date and the date the order is sent to the trading desk, then I_{Desk} would be the better of the two measures since $\mathrm{I}_{\text{Decision}}$ would equal I_{Desk} plus the noise due to other information released between the decision and desk dates. Conversely, if information leaks between the decision date and the date the order is sent to the trading desk, then I_{Desk} would not include the full price impact, while $\mathrm{I}_{\text{Decision}}$ would.

These two actual price impact measures are compared to two predicted price impact measures. The first, which assumes that the order is executed within one trading interval (5 or 30 minutes), comes directly from our cross-sectional regression:

$$\widehat{I}^{u}_{i,t} = X_{i,t}\overline{\Gamma} \times \mathrm{NTO}_{i,t}. \tag{4}$$

The second measure is adjusted for the amount of time taken to execute the order:

$$\widehat{I}^{a}_{i,t} = X_{i,t}\overline{\Gamma} \times \mathrm{NTO}_{i,t} \times 0.5 \times (1 + 1/n), \tag{5}$$

where n is the number of days over which trading of the order extends. The superscripts u and a denote unadjusted and adjusted: The second measure is meant to adjust for the effects of breaking up the order (Bertsimas and Lo 1998). For example, if the order is split in half, half of the order would be executed at half the price impact in Equation (4) and half would be executed at the full impact in Equation (4). This leads to an average price effect of Equation (4) times $(0.5 + 0.25)$ as in Equation (5). Since the data do not allow us to determine the actual manner in which the order is broken up, we use the simple approximation of one trade per day.

---

[7] We thank Don Keim and Ananth Madhavan for providing us with an extensively cleaned version of the data and Wayne Wagner of the Plexus Group for granting us permission to use the data.

Table 3 Actual Versus Predicted Price Impacts by Market, Investment Style, and Order Type

				Actual Impact		Predicted Impact: Price Returns Regressions		Predicted Impact: Quote Returns Regressions	
Order Type	Market	Style	Sample Size	I_{Desk}	I_{Decision}	Unadjusted (4)	Adjusted (5)	Unadjusted (4)	Adjusted (5)
Sell	NYSE/AMEX	Value	7,066	−0.047	0.100	−1.046	−0.909	−0.864	−0.752
Sell	NYSE/AMEX	Technical	11,779	−0.391	−0.594	−1.209	−1.193	−1.054	−1.043
Sell	NYSE/AMEX	Index	2,198	−0.155	−0.239	−0.474	−0.446	−0.399	−0.376
Sell	NASDAQ	Value	512	−0.302	−0.081	−1.719	−1.540	−0.387	−0.341
Sell	NASDAQ	Technical	2,150	−0.892	−0.915	−3.025	−3.023	−0.767	−0.766
Sell	NASDAQ	Index	203	−0.634	−0.730	−1.032	−0.993	−0.228	−0.217
Buy	NYSE/AMEX	Value	6,230	0.179	0.139	1.241	1.056	1.021	0.870
Buy	NYSE/AMEX	Technical	11,268	0.254	0.443	1.372	1.361	1.174	1.166
Buy	NYSE/AMEX	Index	10,457	0.106	0.201	0.179	0.165	0.144	0.133
Buy	NASDAQ	Value	605	0.256	0.103	1.991	1.737	0.435	0.364
Buy	NASDAQ	Technical	2,331	0.533	0.847	2.554	2.549	0.657	0.656
Buy	NASDAQ	Index	1,412	0.524	0.676	0.198	0.178	0.042	0.039

Note. I_{Desk} is the actual impact measured relative to the price prior to sending the order to the trading desk and I_{Decision} is the actual impact measured relative to the price prior to deciding to trade.

Table 3 provides averages of the actual and predicted impacts by investment style, by buy/sell order type, and by exchange-traded versus NASDAQ. The predicted price impact is generally larger (in absolute value) than the actual price impact. To aggregate across buy and sell orders we calculate the signed forecast error:

$$\delta_{i,t} \times (I_{z,i,t} - \hat{I}^{y}_{i,t}), \tag{6}$$

where $z =$ Decision or Desk, $y = a$ or u, and $\delta_{i,t} = 1$ for buys and -1 for sells. A positive (negative) average signed forecast error indicates that the predicted price impact was smaller (larger) than the actual price impact.

We first compare the actual impacts to those predicted using price returns. The consistently negative average forecast errors indicate that the actual price impacts tend to be smaller than the predicted price impacts. Across style categories, the predicted values are closest to the actual values for index funds, with an average (weighted by number of trades) signed forecast error of −6 basis points. The average signed forecast error is −91 basis points for value investors and −113 basis points for technical investors. When sorted by exchanges versus NASDAQ, we find an average signed forecast error of −71 basis points for exchange firms and −144 basis points for NASDAQ. When sorted by buy versus sell orders, we find an average signed forecast error of −74 basis points for buy orders and −98 basis points for sell orders.

When quote returns estimate impact, the predicted values are closer to the actual values. Again, the predicted values are closest to actual price impacts for index funds, with an average (weighted by number of trades) signed forecast error of −0.04 basis points. The average signed forecast error is −65 basis points for value and technical investors. When sorted by exchanges versus NASDAQ, we find average signed forecast errors of −57 basis points (exchange) and 9 basis points (NASDAQ). When sorted by buy versus sell orders, we find an average signed forecast error of −45 basis points for buy orders and −54 basis points for sell orders.

To investigate potential explanations for the observed differences, we regress the forecast errors in (6) on characteristics of the trade and a constant. The characteristics are:

Trade-size-related variables: (1) NTO (net turn over); and (2) NTO^2.

Style-related variables: (3) A value dummy variable equal to unity if the institution is a value-style investor and equal to zero otherwise. (4) A technical dummy variable equal to unity if the institution is a technical investor and equal to zero otherwise.

Exchange-related variable: (5) A NASDAQ dummy variable equal to unity if the stock is listed on NASDAQ and equal to zero otherwise.

Variables measuring the time taken to complete the trade: (6) The number of days between the decision date and the date the order is sent to the trading desk. (7) The number of days between the date the order is sent to the trading desk and the date trading begins. (8) The number of days between the beginning and end of trading.

Variables measuring urgency of trade and the type of order: (9) An urgency code (ranging from 1 for very urgent to 5 for not urgent). (10) A working-order dummy variable equal to unity if the broker is instructed to work the order and equal to zero otherwise. (11) A limit order dummy variable equal to unity if the order is a limit order and equal to zero otherwise. (12) A market-not-held dummy variable equal to unity if the order is a market-not-held order and equal to zero otherwise. (13) A cross dummy variable equal to unity if the order was executed using a crossing network and equal to zero otherwise. (14) A principal dummy variable equal to unity if the order was executed through a principal trade and equal to zero otherwise.

Nonprice impact costs: (15) The commissions paid to the broker, expressed in percent of the trade value (determined by the stock price prior to sending the order to the trading desk). (16) A trade shortfall variable equal to the difference between the desired trade size and the actual trade size (measured in 100,000 share units).

We expect that very large executed trades in the Plexus sample will have actual price impacts smaller than predicted (i.e., a more negative forecast error) since the choice to execute the trade is endogenous. This could be manifested either as a negative coefficient on NTO^2 or as a negative coefficient on NTO with a nonpositive coefficient on NTO^2.

Since we have incorporated dummy variables for both value and technical investing, the base-case investment style is index fund investing. Since value investors have less demand for immediacy, we would expect the coefficient on the value dummy to be negative. Keim and Madhavan (1995) find evidence that technical traders are willing to trade slightly more patiently than index traders but less patiently than value traders, which leads us to expect a coefficient for the technical dummy variable between zero and the coefficient for the value dummy variable.

A number of cross-exchange differences could be picked up by the NASDAQ dummy variable. Given that volume is "double counted" on NASDAQ, we would expect the measured values of price impact to be downward biased, generating a positive forecast error.

The times between the date the decision is made to trade, the date the order is sent to the trading desk, and the date trading begins may reflect the portfolio manager's and the trading desk's judgement about the times at which market conditions are most favorable. If so, we would expect actual price impacts to be lower for longer waiting periods, leading to negative coefficients on these variables. However, the causality could also run in the opposite direction if institutions must wait longer for difficult, low-liquidity trades. If the time between the beginning and ending of trade is a proxy for breaking up the trade, then we expect the forecast error to be more negative the longer the trading time when we use the unadjusted predicted impact. For the adjusted price impact, we would expect to find no relation between the forecast error and the trading interval if our approximation of one trade per day is reasonable. As above, an alternative interpretation would be that trades taking a long time are difficult, high-impact trades.

We predict a negative coefficient on the urgency variable if a very urgent trade has a higher actual price impact. We predict that trades executed through limit orders will have a lower price impact, implying a negative coefficient on the limit order dummy variable. All else equal, a working order or a market-not-held order should have lower price impact than a market order, predicting that the coefficients on these two dummy variables would be negative. However, all else may not be equal in that these types of orders are more often used for more difficult trades.

Commissions and price impact may be substitutes as transactions costs: One may pay higher commission rates to get better execution. Therefore, we

Table 4 Regression of Signed Price Impact Forecast Errors on Trade and Institutional Characteristics

	$\delta_{i,t} \times (I_{\mathrm{Desk}:i,t} - \widehat{I}^{u}_{i,t})$		$\delta_{i,t} \times (I_{\mathrm{Decision}:i,t} - \widehat{I}^{u}_{i,t})$		$\delta_{i,t} \times (I_{\mathrm{Desk}:i,t} - \widehat{I}^{a}_{i,t})$		$\delta_{i,t} \times (I_{\mathrm{Decision}:i,t} - \widehat{I}^{a}_{i,t})$	
	Coefficient	*t*-Stat	Coefficient	*t*-Stat	Coefficient	*t*-Stat	Coefficient	*t*-Stat
Intercept	1.89	2.21	3.73	4.02	1.63	1.96	3.47	3.82
NTO (0.1%)	0.00	0.56	0.00	0.57	0.02	2.40	0.02	2.25
NTO^2 (0.1%^{2})	−0.04	−132.19	−0.04	−119.35	−0.03	−130.43	−0.03	−117.27
Value dummy	−0.71	−11.86	−1.09	−16.81	−0.54	−9.35	−0.93	−14.60
Technical dummy	−0.55	−11.17	−0.64	−11.85	−0.51	−10.49	−0.59	−11.23
NASDAQ dummy	−0.49	−8.50	−0.44	−7.09	−0.47	−8.31	−0.42	−6.88
Decision to desk (days)	−0.11	−13.20	−0.10	−11.15	−0.12	−13.99	−0.11	−11.81
Desk to begin (days)	−0.09	−13.76	−0.05	−7.13	−0.09	−14.60	−0.05	−7.75
Begin to end (units = days)	−0.09	−8.55	−0.10	−9.02	0.00	0.08	−0.01	−1.10
Urgency (1 = urgent to 5 = not urgent)	−0.45	−1.60	−1.07	−3.47	−0.39	−1.41	−1.01	−3.33
Work order dummy	−0.10	−1.33	−0.04	−0.53	−0.07	−0.99	−0.02	−0.20
Limit order dummy	−0.42	−3.19	−0.48	−3.31	−0.41	−3.19	−0.47	−3.31
Not held order dummy	0.09	1.77	0.52	9.12	0.06	1.18	0.48	8.72
Cross dummy	0.14	1.15	0.07	0.53	0.09	0.78	0.02	0.18
Principal trade dummy	−0.50	−2.88	−0.82	−4.33	−0.41	−2.42	−0.73	−3.95
Commissions (%)	−1.71	−17.08	−1.21	−11.08	−1.54	−15.73	−1.03	−9.69
Shortfall (100,000 shares)	−0.10	−2.68	0.01	0.21	−0.09	−2.34	0.02	0.60
R^2	0.27		0.23		0.26		0.22	

predict that the coefficient on commissions will be negative. The shortfall variable is related to the opportunity costs associated with the failure to trade. There might also be a mechanical link between the forecast error and the shortfall: An institution might specify a desired trade size and a maximal (minimal) price at which to buy (sell). The closer that price boundary is to the current market price, the smaller is the maximal price impact and the larger is the probability of a shortfall.

The results for the regressions are reported in Table 4. The signs of the coefficients and incidence of statistical significance (at the 5% level) are generally consistent across specifications. As predicted, the coefficients on NTO are positive, but only significant for the adjusted measures, while the coefficients on NTO^2 are consistently significantly negative. Thus, the predicted price impact will overestimate the actual price impact for large trades that are executed. Both the value and technical dummy variables are negative and statistically significant, suggesting that these institutions demand less liquidity than index funds. The coefficient for the technical dummy variable is smaller in absolute value than the value dummy variable, consistent with the evidence found in Keim and Madhavan (1995). The NASDAQ dummy variable is significantly negative.

The time between the decision to trade and sending the order to the trading desk has a significantly negative coefficient, as does the time between sending the order to the trading desk and the beginning of trading. This is consistent with portfolio managers waiting for favorable market liquidity before handing the order to the trading desk and with the trading desks waiting for favorable market liquidity before beginning to trade. The time it takes to execute the trade has a significantly negative coefficient when the unadjusted predicted impact is used, and an insignificant coefficient (that changes sign) when we use the adjusted predicted impact. As discussed above, this is what one would expect if both the trading interval is related to the manner in which the order is broken up and the approximation of one trade per day is reasonable. Less urgent trades yield lower actual price impact, though the effect is only significant when calculating the actual impact using the price prior to the decision to trade.

As expected, the prediction errors are more negative for limit orders than for working orders, consistent with limit orders providing rather than demanding liquidity. The coefficients on the not-held-order dummy variable are positive and significant only when we calculate the actual price impact using the price prior to the decision to trade. This may be due to portfolio managers choosing to submit not-held orders when the stock price has moved against them during the period prior to sending the order to the desk. The coefficients on the crossing network dummy are insignificantly positive and the coefficients on the principal trade dummy are significantly negative.

The coefficients on commissions are significantly negative, consistent with the argument that higher commissions are paid to execute trades with less price impact. That the coefficients are greater than one in absolute value indicates a greater than one-for-one trade-off between extra commissions and lower price concessions. The coefficient on the trade size shortfall is significantly negative only when impact is measured relative to the trading desk date.

All of the above results are for signed forecast errors obtained using price-return regressions to predict the price impact. We obtain similar results for signed forecast errors using the quote-return regressions to predict impact, so we only provide a summary of the main difference for brevity. With quote returns, the coefficient on NTO becomes significantly negative and the coefficient on the NASDAQ dummy variable becomes significantly positive.

The institutional data collected by the Plexus Group provides an important out-of-sample benchmark against which our predicted price impacts may be compared. The results suggest that the magnitude of the predicted price impact tends to be larger than the actual price impact, especially when price-returns are used to estimate the predicted price impact. The prediction errors are significantly related to the investment style of the institution and the characteristics of the order: More patient trading leads to a more negative price impact forecast error.

4. Summary and Conclusions

We measure equity liquidity using a measure of price impact, the change in a firm's stock price associated with its observed net trading volume for a large cross-section of firms. There is considerable cross-sectional variation in price impact. We study the relation between our measure of price impact and a set of predetermined firm characteristics that proxy for the severity of adverse selection, non-information-based costs of making a market, and the extent of shareholder heterogeneity. We employ the fitted cross-sectional relation between the price impact coefficient and firm characteristics to generate out-of-sample predicted price impacts for a sample of institutional equity trades. The trades represent a total of 1.9 billion shares and $62.4 billion in value. The predicted impact overstates the actual impact, on average, with the difference being the smallest for trades by index funds. Numerous trade characteristics are significantly related to the signed price impact prediction error in economically reasonable ways. For example, the predicted price impact overestimates the actual price impact for very large trades, for trades executed in a more patient manner, and for trades where the institution pays higher commissions.

This study can be extended in many directions. One important question about asset pricing is whether price impact has incremental explanatory power over other measures of liquidity (such as the spread in Amihud and Mendelson 1986) for the purposes of explaining asset pricing (Brennan and Subrahmanyam 1996). Further, price impact might help explain why various asset pricing anomalies are not exploited as well as the cross-sectional differences in mutual fund performance.

Acknowledgments

The authors thank Margaret Forster, Thomas George, Joel Hasbrouck, Ravi Jagannathan (the editor), Don Keim, Hans Stoll, an anonymous referee, and seminar participants at Columbia University, the Federal Reserve Bank of New York, Harvard University, University of Iowa, Michigan State University, Vanderbilt University, and the University of Virginia for helpful comments. We also thank Valery Polkovnichenko for excellent research assistance. Hodrick gratefully acknowledges the generosity of the Lynde and Harry Bradley Foundation and the National Science Foundation through a Presidential Young Investigator Award. Previous drafts of this paper were circulated under the title "Equity Price Concessions."

References

Amihud, Y., H. Mendelson. 1986. Asset pricing and the bid-ask spread. *J. Financial Econom.* **15** 223–249.

Back, K. 1992. Insider trading in continuous time. *Rev. Financial Stud.* **5** 387–409.

Bagwell, L. S. 1991. Share repurchase and takeover deterrence. *RAND J. Econom.* **22** 72–88.

——. 1992. Dutch auction repurchases: An analysis of shareholder heterogeneity. *J. Finance* **47** 71–105.

Bernhardt, D., E. Hughson. 1997. Splitting orders. *Rev. Financial Stud.* **10** 69–101.

Bertsimas, D., A. W. Lo. 1998. Optimal control of execution costs. *J. Financial Markets* **1** 1–50.

Brennan, M. J., A. Subrahmanyam. 1996. Market microstructure and asset pricing: On the compensation for illiquidity in stock returns. *J. Financial Econom.* **41** 441–464.

Brown, D. T., M. Ryngaert. 1992. The determinants of tendering rates in interfirm and self-tender offers. *J. Bus.* **65** 529–556.

Constantinides, G. 1986. Capital market equilibrium with transaction costs. *J. Political Econom.* **94** 842–862.

Easley, D., M. O'Hara. 1987. Price, trade size, and information in securities markets. *J. Financial Econom.* **19** 69–90.

Fama, E. F., J. MacBeth. 1973. Risk, return and equilibrium: Empirical tests. *J. Political Econom.* **71** 607–636.

Glosten, L. R. 1987. Components of the bid/ask spread and the statistical properties of transaction prices. *J. Finance* **42** 1293–1307.

Glosten, L. R., L. E. Harris. 1988. Estimating the components of the bid/ask spread. *J. Financial Econom.* **21** 123–142.

——, P. R. Milgrom. 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *J. Financial Econom.* **14** 71–100.

Gould, J. F., A. W. Kleidon. 1994. Market maker activity on Nasdaq: Implications for trading volume. *Stanford J. Law, Bus. Finance* **1** 1–17.

Harford, J., A. Kaul. 1998. Do concentrated trading equilibria exist? The migration of informed trading following index addition. Working paper, University of Oregon, Eugene, OR.

Harris, L., E. Gurel. 1986. Price and volume effects associated with changes in the S&P 500 list: New evidence for the existence of price pressures. *J. Finance* **41** 815–829.

Hasbrouck, J. 1991a. Measuring the information content of stock trades. *J. Finance* **46** 179–207.

——. 1991b. The summary informativeness of stock trades: An econometric analysis. *Rev. Financial Stud.* **4** 571–595.

——, R. A. Schwartz. 1988. Liquidity and execution costs in equity markets. *J. Portfolio Management* **14** 10–16.

Hodrick, L. S. 1999. Does stock price elasticity affect corporate financial decisions? *J. Financial Econom.* **52** 225–256.

Huang, R. D., H. R. Stoll, 1996. Competitive trading of NYSE listed stocks: Measurement and interpretation of trading costs. *Financial Markets, Institutions & Instruments* **5** 1–55.

Huberman, G., W. Stanzl. 2000. Arbitrage-free price-update and price-impact functions. Working paper, Columbia University, New York.

Keim, D. B., A. Madhavan. 1995. Anatomy of the trading process: Empirical evidence on the behavior of institutional traders. *J. Financial Econom.* **37** 371–398.

——, ——. 1997. Transactions costs and investment style: An inter-exchange analysis of institutional equity trades. *J. Financial Econom.* **46** 265–292.

Knez, P. J., M. J. Ready. 1996. Estimating the profits from trading strategies. *Rev. Financial Stud.* **9** 1121–1163.

Korajczyk, R. A., D. J. Lucas, R. L. McDonald. 1992. Equity issues with time-varying asymmetric information. *J. Financial Quant. Anal.* **27** 397–417.

Kyle, A. S. 1985. Continuous auctions and insider trading. *Econometrica* **53** 1315–1335.

Lee, C. M. C., M. J. Ready. 1991. Inferring trade direction from intraday data. *J. Finance* **46** 733–754.

Loeb, T. F. 1983. Trading costs: The critical link between investment, information, and results. *Financial Analysts J.* **39** 39–44.

Marsh, T., K. Rock. 1986. Exchange listing and liquidity: A comparison of the American Stock Exchange with the NASDAQ national market system. Report #2, American Stock Exchange Transactions Data Research Project, New York.

Mitchell, M., T. Pulvino. 2001. Characteristics of risk and return in risk arbitrage. *J. Finance* **56** 2135–2175.

Newey, W. K., K. D. West. 1987. A simple, positive semi-definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica* **55** 703–708.

Odders-White, E. 2000. On the occurrence and consequences of inaccurate trade classification. *J. Financial Markets* **3** 259–286.

Odean, T. 1998. Are investors reluctant to realize their losses? *J. Finance* **53** 1775–1798.

Petersen, M. A., D. Fialkowski. 1994. Posted versus effective spreads: Good prices or bad quotes? *J. Financial Econom.* **35** 269–292.

Scholes, M. 1972. The market for securities: Substitution vs. price pressure and the effects of information on share prices. *J. Bus.* **45** 179–211.

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

U.S. Department of the Treasury. 1999. Corporations. Publication 542, Internal Revenue Service, Washington, D.C.

Accepted by Ravi Jagannathan; received November 5, 2001. This paper did not require revision.

Copyright 2002, by INFORMS, all rights reserved. Copyright of Management Science is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.