Limited attention and the allocation of effort in securities trading

Shane A. Corwin[a,*], Jay F. Coughenour[b]

[a]*Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556*
[b]*Lerner College of Business and Economics, University of Delaware, Newark, DE 19716*

June 2005

Abstract

The influence of limited attention on decision making has been analyzed in a variety of economic and psychological settings, but its impact on financial markets is not well understood. In this paper, we test whether the limited attention of human market makers influences trading in securities markets. We conduct pooled time-series and cross-sectional tests on individual NYSE specialist portfolios and find that transaction costs increase with the total trading activity of other stocks handled by the same specialist. The results indicate that specialists face significant processing limits and that they allocate effort toward their most active stocks during periods of increased trading activity. Concurrently, specialists reduce their attention to the other stocks in their portfolio, resulting in less frequent price improvement and increased transaction costs. Our evidence suggests that limited attention, and the resulting allocation of effort across stocks, have a significant impact on liquidity provision in securities markets.

JEL classification: G10, G14, D23, J22

Keywords: Limited Attention; Microstructure; Specialists; Transaction Costs

---

We thank Rob Battalio, Tom Cosimano, Glen Dowell, Alok Kumar, Jeff Harris, Brian Hatch, Paul Laux, Marc Lipson, Lin Peng, Paul Schultz, Kumar Venkataraman, Wei Xiong, and seminar participants at the University of Delaware for helpful comments. Any errors are the responsibility of the authors.

* Corresponding author. Tel.: + 1-574-631-6026

E-mail addresses: scorwin@nd.edu (S. Corwin), coughenj@lerner.udel.edu (J. Coughenour).

1. Introduction

Substantial evidence suggests that humans are limited in their ability to process information and to perform multiple tasks simultaneously. Kahneman (1973) argues that this type of "limited attention" requires individuals to allocate their cognitive resources across tasks, so that attention spent on one task must reduce attention available for other tasks.[1] Recently, several theoretical and empirical studies have applied this concept to the analysis of financial markets and investor behavior. Peng and Xiong (2005) show that investors with limited attention will resort to simple decision rules, such as categorization, and that the use of these rules can explain well-documented patterns in asset return co-variation. Further, Peng (2005) illustrates that investors will optimally allocate their limited attention across sources of uncertainty to minimize total portfolio uncertainty.[2] Consistent with these theories, Huberman (2001), Huberman and Regev (2001), and Barber and Odean (2005) provide evidence that investors tend to focus on familiar or attention-grabbing stocks and that information may not be incorporated into prices until it attracts the attention of investors. While this empirical research provides indirect evidence consistent with limited attention, direct tests are scarce because human attention and its allocation across tasks are difficult to measure in financial market settings.

In this paper, we study limited attention in the context of market making on the New York Stock Exchange (NYSE). This setting is ideal for analyzing the effects of limited attention for several reasons. First, the NYSE features individual specialists who are obligated to provide liquidity for a well-defined set of securities. As a result, we can directly identify the set of securities across which the specialist must divide their attention. Second, we can measure factors that necessitate the allocation of attention across securities. We use the trading activity of all stocks in the specialist's portfolio to measure intertemporal

---

[1] See Kahneman (1973) and Pashler (1998) for detailed discussions of attention limits and attention allocation. Limited attention is also related to the concept of 'bounded rationality,' which has been applied extensively to other economic questions. Unlike economic theories where agents are assumed to optimize over all possible alternatives and to fully understand the consequences of all choices, bounded rationality assumes that decision-makers are rational only within limits. These limits may result from the decision environment or from the computational capacities of the agent. For a general discussion of bounded rationality, see Simon (1955, 1979) and Sargent (1993). See Gifford (2001) for a discussion of limited attention as a bound on rationality.

[2] In related work, Hirshleifer and Teoh (2003) and Hirshleifer, Lim, and Teoh (2004) show that firms may alter their financial reporting decisions in settings where investors are subject to limited attention. Limited attention has also been applied in models of behavior within organizations. A common framework in these models is that of a manager who must allocate limited effort across multiple projects (e.g., Radner and Rothschilde (1975) and Gifford (2001)).

variation in the attention a specialist can provide to any given stock. Finally, because specialists provide an important source of liquidity through their participation in trading, the effects of limited attention can be measured in the form of transaction costs. If limited attention leads to the allocation of specialist effort across stocks, we expect to observe a positive relation between transaction costs and the trading activity of other stocks handled by the same specialist, all else constant. We refer to this as the *Limited Attention Hypothesis*.

The Limited Attention Hypothesis is based on the assumption that individual specialists face time and processing constraints that limit their ability to monitor and execute orders, particularly during busy trading periods. While the specialist's limited attention can affect all stocks, we expect the effects to be most evident for small, inactive stocks for several reasons. First, specialists participate in a larger fraction of trades and provide a greater proportion of liquidity for inactive securities (see Madhavan and Sofianos (1998)). As a result, changes in specialist participation rates should be most apparent for these securities. Second, cost-benefit models of attention allocation suggest that agents will allocate attention in a manner that maximizes their total utility.[3] Since specialists put more capital at risk when trading the most active stocks and also derive a large fraction of their profits from these stocks (see Sofianos (1995)), we argue that they are less likely to divert attention from these actively-traded securities.

We test the Limited Attention Hypothesis using intraday data on individual NYSE specialist portfolios. Results from pooled time-series and cross-sectional regressions indicate that bid-ask spreads are wider and price improvement is less frequent when the combined trade frequency of the other stocks handled by the same specialist increases. These results hold after controlling for variation in the stock's own trading activity, for firm fixed effects, for time-of-day effects, and for other variables known to affect transaction costs. The findings are also robust to alternative definitions of busy trading periods and to alternative econometric techniques. To account for potential commonality between individual stock trading activity and panel trading activity, we repeat the analysis using only those periods in which own-

[3] Gabaix and Laibson (2004) and Gabaix, Laibson, Moloche, and Weinberg (2003) develop a general cost-benefit model of endogenous attention allocation and provide experimental evidence in support of this directed cognition effect. See Peng (2005) for an application to financial markets.

stock trading activity is normal. The results confirm that our findings are driven by variation in panel activity rather than own-stock trading activity. Consistent with our hypothesis, the effects of limited attention are most evident for the least active stocks. While the positive relation between panel activity and bid-ask spreads is evident for all categories of stocks, the negative relation between panel activity and the rate of price improvement (a proxy for specialist involvement in trading) is significant for only the less active securities. Together, our results suggest that market makers allocate limited attention across the securities in their portfolio, and that this allocation has a significant impact on transaction costs.

Our evidence is particularly notable given that several NYSE characteristics may reduce the effects of limited attention. NYSE specialists are intensely regulated and their performance with respect to liquidity provision is closely monitored. During unusually busy periods, specialists can increase capacity by calling on 'relief specialists' or additional clerks. In addition, specialist firms appear to allocate stocks to panels in a manner that reflects attention limits, with the most active stocks trading apart from one another and at smaller panels. Together, these factors may mitigate the potential effects of an individual specialist's limited attention. Nevertheless, we document a significant relation between limited attention and liquidity provision.

Our empirical work is closely related to three concurrent studies that focus on individual specialist trading. Battalio, Ellul, and Jennings (2004) examine time-series changes in transaction costs, but focus specifically on changes in floor location. They find that specialists form cost-reducing relationships with floor brokers and that these relationships take time and attention to develop following a reorganization of the trading floor. Our results provide additional evidence that the location of a security on the trading floor can influence transaction costs. In cross-sectional analyses, Huang and Liu (2004) find that NYSE specialists subsidize the illiquid stocks in their portfolio and Hatch, Johnson, and Lei (2004) find that quote adjustment speeds depend upon the prominence of the stock within the specialist's portfolio. While our evidence is generally consistent with these two studies, we note that cross-sectional analyses of limited attention are difficult to interpret given the endogenous relation between stock characteristics and specialist portfolio assignments. In contrast, our study focuses on time-series variation in trading activity within the specialist's portfolio and on the relation between this panel activity and transaction costs. This allows us to minimize the aforementioned endogeneity problem and to directly test whether variation in market maker attention affects transaction costs.

Overall, our evidence indicates that transaction costs are significantly affected by the limited attention of market makers and the resulting allocation of effort across securities. These findings point to a potential but unexplored benefit of recent NYSE initiatives to automate a larger fraction of trading. Specifically, increased automation may reduce capacity constraints and allow specialists to focus on those trades for which they add the most value. However, our analysis does not permit us to draw conclusions about the optimality of alternative market structures or to determine whether a reduction in capacity constraints would result in lower transaction costs for the overall market. We argue only that market maker attention is not unlimited and that the resulting effects are significant enough to be considered along with other costs and benefits of market design. While our tests are based on data from the NYSE, our findings may be applicable to any market where the effort of intermediaries must be allocated across multiple securities.

The remainder of the paper is organized as follows. In Section 2 we discuss related literature and develop our main hypotheses. Section 3 describes the data and sample characteristics. In Section 4 we provide the main empirical tests of the Limited Attention Hypothesis. Section 5 describes several robustness tests and a brief summary concludes the paper in Section 6.

2. Background and motivation

2.1. The NYSE trading floor and the role of the specialist

Each security traded on the NYSE is handled by a single specialist who is responsible for providing two-sided quotes and making a "fair and orderly market" in the security. However, individual specialists are typically responsible for making markets in multiple securities.[4] The decision of assigning

[4] As of August 2002, there were seven active specialist firms on the NYSE trading at 17 trading posts and 357 trading panels. The number of securities traded at an individual specialist panel (including preferred stocks, warrants, trusts, and other structured

a particular security to an individual specialist involves input from the listing firm, the specialist firm, and the Exchange. Initially, stocks are allocated to specialist firms in accordance with the Exchange's Allocation Policy and Procedures (see Corwin (2004)). The specialist firm then identifies the individual specialist who is to provide liquidity for the stock. Once allocated, reassignments of stocks across specialist firms are rare.[5] However, reorganizations of stocks within a specialist firm are relatively common and specialist firms have some flexibility in how they organize stocks across trading panels.

Corwin (2004) finds that stock allocations to NYSE specialist firms reflect both performance and non-performance variables. For example, extremely active and extremely inactive stocks tend to be distributed across specialist firms. Preference is also given to specialist firms that already trade securities in the same industry. These findings suggest that the NYSE takes both potential profitability and risk into consideration when assigning stocks to specialist firms. However, prior research does not address the factors that affect stock allocations within specialist firms. Notably, the role of performance in the NYSE allocation process suggests that there is a cost to the specialist of providing less than adequate market quality for a stock. As a result, specialists may be unwilling to set unusually wide quotes or avoid participation in the trading process for extended periods of time.

2.2. Market making and the Limited Attention Hypothesis

In microstructure models, market maker behavior is typically determined by two types of risk. First, market makers face inventory risk as their participation in trading moves them away from their desired inventory position (see, for example, Stoll (1978), Ho and Stoll (1981), and O'Hara and Oldfield (1986)). Second, market makers face the risk of trading with someone who possesses superior information (see, for example, Copeland and Galai (1983), Glosten and Milgrom (1985), and Easley and O'Hara

products) ranged from one to 63. Throughout the paper, we refer to the stocks at a single panel as an individual specialist portfolio.

[5] Schwartz (1993) notes that 11 stocks were reallocated across specialist firms due to poor specialist performance following the market crash in October 1987. Stocks are also reassigned in association with specialist firm mergers. Corwin (2004) provides a description of specialist firm mergers and associated stock assignments from 1992 through 1998. Hatch and Johnson (2002) provide an analysis of the impact of specialist firm mergers on market quality.

(1987)).[6] Prior research also suggests that the NYSE specialist can play a unique role in managing adverse selection risks. For example, Glosten (1989) shows that a monopolist specialist can improve market liquidity in periods of extreme information asymmetry because of his ability to average profits across trades. In addition, Benveniste, Marcus, and Wilhelm (1992) show that the specialist can manage adverse selection risk by forming relationships with floor traders, and Leach and Madhavan (1993) show that the specialist's ability to average profits across time allows them to experiment with prices to induce more informed order flow. Together this research suggests that an unconstrained specialist will set bid and ask prices for all stocks conditional on their inventory risk and the probability of informed trade. However, these models do not account for the possibility that individual specialists may be subject to limited attention.[7]

If specialists face attention limits, they may not be able to continuously incorporate all relevant information for all securities. During periods of unusually busy trading, specialists may be forced to allocate effort across the securities in their portfolio. What factors determine how a specialist will allocate their effort? Given the intuition from Peng (2005), we expect a constrained specialist to allocate effort toward those stocks that have the greatest impact on his utility. In particular, we expect specialists to focus on those stocks that have the largest influence on their portfolio profits and risk. Given that specialists place the most capital at risk when trading the largest, most active securities, it is reasonable to expect that these securities have the greatest impact on specialist risk. In addition, the largest, most active securities account for the vast majority of specialist profits. For example, Coughenour and Harris (2005) find that roughly 82% of specialist revenues are derived from the 100 most active stocks. Thus, both

[6] Empirical studies of NYSE specialist behavior generally support both inventory and adverse selection effects. For example, Hasbrouck and Sofianos (1993), Madhavan and Smidt (1993), and Madhavan and Sofianos (1998) find evidence of active, though possibly slow, inventory adjustment by specialists through the timing and size of trades. Kavajecz (1999) finds that specialists also manage quoted depth to control both inventory and adverse selection risks.

[7] It is worth noting that individual specialists share market-making risks and rewards by forming specialist firms. Prior research finds that liquidity provision by individual specialist differs across specialist firms. Cao, Choe, and Hatheway (1997) and Corwin (1999) show that transaction costs differ across specialist firms. Coughenour and Deli (2002) and Coughenour and Saad (2004) find that certain specialist firm characteristics may partially explain differences in transaction costs across specialist firms and co-variation in transaction costs across stocks, respectively. Although it is possible that a specialist's decision to allocate attention across stocks is influenced by specialist firm characteristics, we focus on whether individual specialists allocate attention across their set of assigned stocks.

profit and risk considerations suggest that a specialist will allocate their attention toward the largest, most active stocks in their market making portfolio.

If specialists allocate attention toward the largest, most active securities, the amount of attention devoted to the less active securities in their portfolio must be reduced. This suggests that the effects of limited attention should be most evident in the transaction costs of small, inactively-traded securities. The effects of limited attention may also have a greater impact on inactive securities because specialists provide a larger fraction of liquidity for these securities. Madhavan and Sofianos (1998), for example, report that specialist's participate in 54% of share volume in the least active decile of NYSE stocks, compared to only 15% in the most active decile.[8] They conclude that the importance of the specialist as a source of liquidity decreases as the trading activity of the stock increases. Therefore, since small, inactive securities rely more on the specialist to supply liquidity, transaction costs for these securities will be tied more directly to specialist actions. Conversely, for more active securities, the specialist faces substantial competition from other traders and transaction costs are determined in large part by the limit order book.

3. Data and sample characteristics

To analyze the effects of limited attention within specialist portfolios, we must identify the trading location of each NYSE security on each trading day. To accomplish this, we obtained daily NYSE Specialist Directories for the period from August 1, 2002 through October 31, 2002. For every NYSE-listed security, the Specialist Directory identifies the specialist firm assigned to the security, as well as the post and panel where the security trades on the NYSE floor. The number of securities listed in the specialist directory (including all preferred stocks, warrants, and structured products) ranges from 3,599 on August 1, 2002 to 3,609 on October 31, 2002.

To refine the sample, we combine the Specialist Directory data with additional information from the Center for Research in Security Prices (CRSP). We start by identifying the sample of securities

[8] Coughenour and Harris (2005) find that specialists participate in 30% of trades for the top 100 large-cap securities and 44% for their group of 1,311 small-cap securities. In related work, Kavajecz (1999) shows that specialists selectively participate in quoted depth. He reports that 25 to 50% of posted quotes on the NYSE are either outside the best limit order price (hidden limit orders) or include no specialist participation.

included in both CRSP and the NYSE Specialist Directory for the full sample period from August 1 through October 31, 2002. This provides an initial sample of 2,515 securities. We then restrict the sample to common stocks and ADRs (CRSP Share code equal to 10, 11, 12, 30, or 31). This reduces the sample to 1,920 securities. Throughout the rest of the paper, we refer to this as the "full sample" and we estimate all panel size and panel activity variables using this sample.

For all analyses of transaction costs, we focus on the subset of stocks that meet an additional set of price and trading restrictions. We remove stocks that experience a stock split during the sample period, stocks with an average daily closing price of less than \$3 or greater than \$200, stocks with a minimum transaction price during the sample period of less than \$2, and stocks that trade in fewer than 800 of the 840 30-minute trading intervals during the sample period. These restrictions reduce the sample by 20, 125, 15, and 498 securities, respectively.[9] The restricted sample used in the transaction cost analysis includes 1,262 NYSE-listed common stocks and ADRs.

For each security in the restricted sample, we estimate transaction cost measures at 30-minute intervals based on intraday trade and quote data from the NYSE's TAQ database. For each transaction on the NYSE, let t denote transaction time, P denote transaction price, a denote the ask price, b denote the bid price, and m denote the bid-ask midpoint. For each trade, we define the quoted spread (qs_t) as a_t-b_t, the percentage quoted spread (pqs_t) as $100 \cdot qs_t/m_t$, the effective spread (es_t) as $2|P_t-m_t|$, and the percentage effective spread (pes_t) as $100 \cdot es_t/m_t$. Trade-weighted quoted and effective spreads are then estimated using all trades that occur during each 30-minute interval. As a proxy for the rate of specialist involvement in trades, we also define the rate of price improvement as the proportion of all trades that occur inside the bid and ask quotes.[10]

[9] To the extent that limited attention effects are most pronounced for the least active securities, these restrictions should bias our tests against the Limited Attention Hypothesis.

[10] We apply several standard filters to the trade and quote data. We utilize only those trades that occur on the NYSE during regular exchange hours, have a positive price and quantity traded, and have a trade correction code less than 3. We include only NYSE quotes with 'mode' = 0, 1, 2, 6, 10, or 12. Bid-ask spreads and price improvement are defined based on the most current NYSE quote at the time of each trade. To measure these variables, we exclude quotes if either the bid or offer depth is non-positive, if the bid-ask spread is greater than \$5, or if either the bid or ask differs by more than 25% from the preceding quote (unless it is the first trade of the day). We also exclude the first trade of the day and any trade for which the effective spread is more than \$5.

The choice of a 30-minute aggregation period is largely driven by the nature of our hypothesis. This interval should be fine enough to capture periods during which the specialist becomes time-constrained and also allows us to capture well-known intraday variation in trading activity and transaction costs (e.g. McInish and Wood (1992)). Aggregating over higher frequency time periods (such as 5-minutes) would result in many zero trade observations for the less-actively traded stocks and noisier estimates of transaction costs. Lower frequency aggregation periods (such as daily) would reduce the power of our tests by smoothing over intraday periods of intense trading.

The Limited Attention Hypothesis suggests that specialist will allocate effort away from least active securities and toward the most active securities. Therefore, following Coughenour and Harris (2005), we sort our sample into three groups based on total trade frequency; the 100 most active; the next 400 mid-active; and the remaining 762 least active securities. Throughout the paper we conduct tests on the complete sample, as well as these trade-activity subsamples.

Table 1 provides summary statistics for the final sample of 1,262 NYSE-listed securities, as well as the three trade-activity subsamples. Across the full sample (Panel A), the average transaction price ranges from $3.15 for the lowest-priced stock to $121.72 for the highest-priced stock, with a mean of $26.51. The average firm in the sample trades at least once in 837 of the 840 30-minute trading intervals, with an average volume of 66,000 shares or $1.9 million per period and an average trade size of 852 shares. Trading activity ranges from an average of 5.9 trades and 1600 shares per period for the least active stock in the sample to 387 trades and 1.6 million shares per period for the most active stock.

Turning to transaction cost measures, the average firm has a quoted bid-ask spread of 5.3 cents and an effective spread of 3.8 cents. The average percentage quoted spread is 27.4 basis points (bps) and the average percentage effective spread is 19.5 bps. Percentage quoted spreads range from 6.6 bps to 163 bps, while percentage effective spreads range from 4.3 bps to 114 bps. On average, 39.9% of NYSE trades occur inside the quotes and the frequency of price improvement ranges from 18.1 to 60.4%.

The results for the three trade activity subsamples (Table 1, Panel B) illustrate the substantial cross-sectional variation in trading activity. The most active stocks average 179 trades per period, while the mid- and low- trade activity stocks average 79 and 23 trades per period, respectively. Similarly, dollar volume drops from $11.8 million per period for the most active stocks, to $2.4 million for the mid-activity stocks and $0.32 million for the least active stocks. Transaction costs also vary widely across trade-activity groups. Percentage quoted spreads range from 13.7 bps for the most active stocks to 34.4 bps for the least active stocks and percentage effective spreads range from 9.2 bps for the most active stocks to 24.7 bps for the least active stocks. The rate of price improvement ranges from 39.4% for the least active stocks to 42.6% for the most active group.[11] The restriction of equal means across trade activity categories is easily rejected for all variables (p-values < 0.0001).

To provide a clearer picture of the NYSE trading floor, Table 2 provides a description of NYSE post and panel composition as of August 1, 2002. On this date, the trading floor included 17 active trading posts and 357 active trading panels. Using the full specialist directory, the mean and median panel sizes are 10.1 and 9.0 securities, respectively, and panel size ranges from one to 63 securities. The four securities that trade alone at a panel on August 1 are Nortel Networks, CIT Group, and the SPY and QQQ exchange traded funds (ETFs).[12] Of the 3,599 securities in the directory at the start of the sample period, 15% change trading posts and 27.8% change trading panels at some point during the sample period. This highlights the importance of identifying specialist portfolios on a daily basis.

After excluding funds, REITs, Units, Trusts, and other structured products, we see that common stocks and ADRs from the final sample trade at 331 different panels, with an average panel size of 5.8 stocks. Notably, the reduction in number of panels relative to the full specialist directory suggests that 26 panels trade no common stocks or ADRs. The largest panel now includes 21 stocks and there are 11 securities traded at panels with no other common stocks. Of the 1,920 common stocks and ADRs in the sample, 17.6% change posts and 31.5% change panels at some point during the sample period.

[11] Although the price improvement rate is highest for the most active stocks, price improvement is more likely to reflect specialist involvement for the least active stocks for which specialists play a more prominent role in providing liquidity. We also conducted all tests using "net price improvement," defined as the proportion of trades inside the quotes minus the proportion of trades outside the quotes. The results are qualitatively similar.

[12] Both Nortel Networks and CIT Group trade at larger panels for most of the sample period. As of October 31, only the three largest ETFs (SPY, QQQ, and DIA) and one common stock (OM Group) are traded alone at a panel. OM Group was moved to a single-stock panel on the last day of the sample period.

The full distribution of panel size is illustrated in more detail in Figure 1. As suggested in Table 2, the distribution of panel size in the full specialist directory (Panel A) is substantially skewed. The mode of panel size is nine and there are approximately 25 to 35 panels at each panel size from five to eleven. However, there are also numerous panels with more than 20 securities. Restricting the sample to common stocks and ADRs (Panel B) substantially reduces the skewness in panel size. The mode of panel size is reduced to five and there is only one panel with more than 20 securities.

For completeness, the last row of Table 2 describes panel size in the price and trade restricted sample of 1,262 stocks. Of the 357 original panels, 323 trade at least one security from the restricted sample. The average panel includes four stocks from this sample and the largest panel includes 12 sample stocks. Of the 1,262 stocks in the restricted sample, 17.3% change posts and 32.4% change panels during the sample period.

If limited attention and allocation of effort are important market making considerations, we expect NYSE specialist firms to allocate their most active stocks to panels with few other securities. As an initial test of our hypothesis, we examine the relation between individual stock trade activity and panel assignments on the NYSE trading floor. For each stock in our restricted sample (N=1,262), we calculate the average number of securities at the same panel, the stock's average rank at the panel based on number of trades, and the stock's average proportion of trades and dollar volume at the panel. Panel size is measured based on both the complete specialist directory and the sample of 1,920 common stocks. Panel ranks and trade shares are defined based on the common stock sample only. To account for changes in panel characteristics over time, we estimate the rank and percentage of total trade activity during each 30 minute period, and then calculate an average for each stock across all periods.

Table 3 reports the cross-sectional means of panel characteristics for each of our three trade activity categories. Considering all types of securities, the most active stocks trade on panels with 7.2 securities, on average, while the least active stocks trade on panels with 12.5 securities. Counting only common stocks and ADRs, the most active stocks have an average panel size of 4.0, while the least active stocks have an average panel size of 7.7. Results for panel ranks are similar. The most active stocks have an average panel rank of 1.1, while the least active stocks have an average panel rank of 3.8. We also find that the most active stocks account for a significantly larger proportion of trading volume at the panel, representing 72.8% of dollar volume and 61.5% of transactions, on average. In contrast, the least active stocks represent only 11.2% of panel dollar volume and 13.6% of transactions, on average. The restriction of equal means across trade activity groups is easily rejected for all variables (p-values < 0.0001).

The evidence in Table 3 is consistent with the hypothesis that NYSE specialist firms tend to place their most actively traded securities at smaller panels. Similar results are reported in Huang and Liu (2004). We conclude that the observed allocation of stocks to panels provides *prima facie* evidence that NYSE specialist firms recognize the marginal costs associated with limited attention and allocation of effort. In the following sections, we consider the significance of these costs in light of the fact that they may be reduced by the stock allocation decisions of specialist firms.

4. Empirical tests for limited attention and effort allocation

In this section we present our empirical tests for limited attention and effort allocation in securities trading. For comparability to other research, we start by examining the cross-sectional relation between transaction costs and specialist portfolio characteristics. We then move to our primary tests based on pooled time-series and cross-sectional analysis.

4.1. Preliminary cross-sectional analysis

Prior research suggests that NYSE securities may benefit from subsidization or diversification benefits within market making portfolios (see, for example, Cao, Choe, and Hatheway (1997), Gehrig and Jackson (1998), and Huang and Liu (2004)). While subsidization and diversification effects may exist for some securities, the Limited Attention Hypothesis suggests that these benefits will be diminished as trading at the panel increases.

As a preliminary experiment, we examine the cross-sectional relation between transaction costs and panel characteristics. The cross-sectional regressions include several control variables that are known to explain the cross-section of bid-ask spreads and take the following general form:

$$ExecutionCost_i = \alpha + \beta_1 \cdot InvPrice_i + \beta_2 \cdot LogTrades_i + \beta_3 \cdot LogTradeSize_i + \beta_4 \cdot StDevRet_i + \sum_{g=1}^{3} \beta_{5g} \cdot LogPanelActivity_i * T_g + \varepsilon_i , \quad (1)$$

where for each stock *i*, $StDevRet_i$ is the standard deviation of 30-minute quote midpoint returns, $InvPrice_i$ is the inverse of average trade price, $LogTrades_i$ is the natural log of average trade frequency, and $LogTradeSize_i$ is the natural log of average trade size. Our key variable is $LogPanelActivity_i$, defined as the natural log of average panel activity, where panel activity is the total number of trades for all other stocks traded at the same specialist panel. For each security, price, trade frequency, trade size, and panel activity are estimated during each 30-minute trading period and then averaged across all 840 periods. To analyze subsidization effects, we interact $LogPanelActivity_i$ with three dummy variables (T_g) that identify the trade activity category to which the stock belongs. To the extent that subsidization effects are reflected in the spreads of the least active stocks, we expect a negative relation between transaction costs and panel activity for these securities.

The cross-sectional regressions are reported in Table 4. The results suggest that the effects of panel activity differ significantly by trade activity category. For the most active securities, the coefficient on panel activity is positive and significant for all measures of dollar and percentage spreads. In contrast, for the least active stocks, the coefficient on panel activity is always negative and is significant for percentage spreads. These results suggest that inactive stocks may benefit from being traded at a panel with active securities, consistent with conclusions of Huang and Liu (2004). Although subsidization or diversification benefits may exist, the Limited Attention Hypothesis suggests that these potential benefits will decrease as the specialist allocates attention away from the inactive stocks.[13]

[13] To examine this possibility in the cross-sectional setting, we reestimated the regressions with an additional interaction term to identify those inactive securities that are traded at a panel with one of the 100 most active stocks in the sample. We expect the specialists at these panels to face greater attention constraints than other specialists. The results are consistent with this

We note that the coefficients from these cross-sectional tests must be interpreted with caution given the endogenous relation between trading activity and stock allocation decisions on the NYSE.[14] In particular, we expect specialist firms to allocate stocks so that expected trading activity is somewhat balanced across panels. Due to this endogenous relation, we believe that cross-sectional tests are not particularly well-suited to test the Limited Attention Hypothesis.

4.2. Pooled time-series and cross-sectional analysis

The Limited Attention Hypothesis implies a positive relation between transaction costs and the degree to which a specialist faces time and processing constraints through time. In other words, transaction costs should increase during periods when the specialist is forced to allocate their effort away from a specific stock. In this section, we test this hypothesis directly using a pooled time-series and cross-sectional analysis. As a proxy for the time-variation in specialist attention to each stock, we measure the combined trading activity of all other stocks at the same floor panel. Because these tests focus on changes in transaction costs and panel trading activity over time, they are less susceptible to the endogeneity concerns that plague cross-sectional tests.

One important consideration in our estimation is the potential correlation in order flow across stocks. For example, Hasbrouck and Seppi (2001) report strong positive co-variation in order flow across the 30 stocks in the Dow Jones Industrial Average. To isolate the effects of limited specialist attention, we control for own-stock trading activity in two ways. First, our regressions control for own-stock trading activity directly. This allows us to interpret separately the effects of panel trading activity and own-stock trading activity. Second, we repeat all of our analysis using only those periods when own-stock trading

---

hypothesis. While inactive stocks appear to have lower percentage spreads when traded on an active panel, this effect is significantly reduced when the same specialist trades one of the 100 most active NYSE stocks. These results are available upon request.

[14] Indeed, two important factors suggest that specialists would find it difficult or costly to set wide spreads on active stocks in order to subsidize inactive securities. First, for actively traded securities, the specialist faces substantial competition from other traders and transaction costs are determined in large part by the limit order book. This competition limits the specialist's ability to increase spreads on these securities. Second, specialist performance ratings may be heavily influenced by their performance for the most visible securities. This suggests that specialists would be unwilling to widen spreads on these stocks for the purpose of subsidization.

activity is within normal bounds. This allows us to test the robustness of our results to own-stock trading effects.

Before presenting the regression results, we provide univariate evidence on the relation between transaction costs and the combined trade activity of other stocks at the panel. We provide results for five measures of transaction costs: dollar quoted and effective spreads, percentage quoted and effective spreads, and the rate of price improvement. To control for the effects of correlated order flow, as discussed above, we present summary statistics for periods when own trading activity is low, normal and high, relative to periods when panel trading activity is low, normal, and high. For each stock, we estimate the mean and standard deviation of the number of trades across all 30-minute trading periods. Periods of high and low trading activity are defined as periods with activity more than one standard deviation above and below the mean, respectively. All other periods are defined as normal. Periods of low, normal, and high panel trading activity are defined similarly, where panel activity for a given stock equals the total number of trades for all other stocks traded at the same floor panel.

The univariate results are presented in Table 5, with normal own-trading in Panel A, low own-trading in Panel B, and high own-trading in Panel C. In each panel, results are provided separately for the low, mid, and high trading activity subsamples. In addition, we provide a test of whether means are equal across periods of low, normal, and high panel activity. The evidence for normal own-trading periods (Panel A) indicates that transaction costs differ significantly based on the level of panel trading activity. For low activity stocks, percentage effective spreads range from 23.2 basis points (bps) for low panel activity periods to 25.5 bps for high panel activity periods and percentage quoted spreads range from 32.5 bps to 35.2 bps across these categories. The rate of price improvement ranges from 38.1% for high panel activity periods to 40.0% for low panel activity periods. Results for dollar spreads are qualitatively similar, though equality of means across panel activity groups cannot be rejected for these variables.

The results for mid-activity stocks are largely similar to those reported for low-activity stocks. For the most active stocks, however, spreads and price improvement rates do not vary significantly with variation in panel activity. Overall, the results in Panel A are generally consistent with the Limited

Attention Hypothesis and suggest that transaction costs increase as panel trading activity increases, especially for the least active securities. However, these results do not control for other stock characteristics and time-of-day effects that can influence transaction costs. These issues are addressed in more detail in the regression analysis below.

For comparison, univariate results for low and high own-trading periods are presented in Panels B and C, respectively. Results for periods of low own-trading (Panel B) are generally consistent with the normal own-trading evidence, though somewhat stronger. During periods of low own-trading activity, all three categories of stocks exhibit significantly higher dollar and percentage spreads during high panel activity periods than during low panel activity periods. However, the rate of price improvement during low own-trading does not differ significantly across panel trading activity levels for any of the three categories of stocks.

The results for high own-trading provide a striking contrast to the results presented in Panels A and B. For both low and mid activity stocks, bid-ask spreads actually decrease as panel activity increases. At the same time, the rate of price improvement continues to decrease as panel activity increases. One potential explanation for these results is that the interaction of own-stock trading activity and panel trading activity reveals something about the trading environment. For example, periods when own-stock trading activity is high but panel activity is low (or normal) may reflect significant firm-specific events. Consistent with this explanation, bid-ask spreads are significantly wider during these periods than when both panel and own-stock trading activity are high, at least for the less-active stocks. These results highlight the importance of controlling for own-stock trading activity when analyzing the effects of panel activity on transaction costs.

As a more direct test of the Limited Attention Hypothesis, we provide pooled time-series and cross-section regressions that take the following general form:

$$ExecutionCost_{it} = \alpha + \sum_{p=2}^{13} \alpha_i \cdot HH_p + \beta_1 \cdot InvPrice_{it} + \beta_2 \cdot LogTrades_{it} + \beta_3 \cdot LogTradeSize_{it} + \beta_4 \cdot StDevRet_{it} + \beta_5 \cdot PanelActivity_{it} + \varepsilon_{it}, \quad (2)$$

where for each stock i, $StDevRet_{it}$ is the standard deviation of trade-to-trade quote midpoint returns during period t, $InvPrice_{it}$ is the inverse of average trade price during period t, $LogTrades_{it}$ is the natural log of the number of trades during period t, and $LogTradeSize_{it}$ is the natural log of average trade size during period t. To control for intraday patterns in transaction costs we also include dummy variables (HH_p) to identify 30-minute trading periods from 9:30 a.m. to 4:00 p.m.

We define two alternative measures of $PanelActivity_{it}$ based on the total number of trades during period t for all other stocks traded at the same specialist panel. First, we include the natural log of Panel Activity (*LogPanelActivity*) during period t, where Panel Activity is defined as the total number of trades for all other stocks at the same specialist panel.[15] Second, we define dummy variables to identify periods of unusually low and unusually high panel trading activity. Specifically, $Busy_t$ ($Slow_t$) is equal to one if both (i) panel activity during period t is more than one standard deviation above (below) its own panel-specific mean, and (ii) panel activity is higher (lower) than the 75th (25th) percentile of panel activity observations across all stocks and periods. The first characteristic identifies periods that are busy relative to the normal level of trading at that panel, while the second factor ensures a high absolute level of trading activity.[16] If transaction costs increase during periods of high panel activity, we expect the coefficients on *Busy* and *LogPanelActivity* to be positive. In addition, if the specialist is able to provide additional liquidity during periods of unusually low panel activity, we expect the coefficient on *Slow* to be negative.

Given the pooled nature of the data, we expect the error terms in the model to include a security-specific component. To account for this, the model is estimated including one-way fixed effects. We assume the following structure for the error terms:

---

[15] We believe trade frequency provides a more appropriate measure of specialist attention to a stock than share volume. Nevertheless, the results are similar if panel activity is defined based on share volume rather than number of trades.

[16] We include the latter restriction to account for the fact that some panels may always be relatively active or relatively inactive due to the nature of the stocks at the panel. However, results are qualitatively similar if we define busy and slow periods using only the panel-specific definition. Based on the panel-specific restriction, 14.5% of firm-periods are classified as slow and 15.4% are classified as busy. When the absolute trading restriction is also applied, 7.9% of firm-periods are classified as slow and 9.6% are classified as busy.

$$\varepsilon_{it} = \nu_i + u_{it} , \quad (3)$$

where v_i is a firm-specific fixed effect and u_{it} is the classical error term with zero mean.[17] As in the cross-sectional regressions, the model is estimated on dollar quoted and effective spreads and percentage quoted and effective spreads. In addition, as a proxy for participation by the specialist, we also estimate the model for the rate of price improvement. The Limited Attention Hypothesis suggests that the specialist will decrease their participation in trading when other stocks at the panel become unusually busy. As a result, we expect the coefficients on *LogPanelActivity* and the *Busy* period dummy variable to be negative in the price improvement regression. Following prior research, the quoted bid-ask spread is included as an explanatory variable in the price improvement regression.

The results for the full sample of 1,262 stocks are provided in Table 6. Models including *Busy* and *Slow* panel activity dummy variables are described in Panel A. In general, all control variables have the expected signs and significance levels. In addition, the restrictions that firm fixed effects and time-of-day effects are equal to zero are easily rejected in all specifications. Overall, the results in Panel A provide strong evidence that transaction costs increase with panel activity. The coefficient on the *Busy* period dummy is positive and significant at the one percent level for all four bid-ask spread measures, and significantly negative for the rate of price improvement. These results indicate that the limited attention of individual specialists significantly influences transaction costs for NYSE securities.

The coefficient on the *Slow* period dummy variable is insignificant in all four bid-ask spread models, suggesting that these results are driven primarily by busy trading periods. However, both the *Busy* and *Slow* period dummy variables are significant in the price improvement model. The coefficients on these variables suggest that the rate of price improvement decreases during unusually busy trading at the panel and increases when panel trading activity is light. Together with the results for bid-ask spreads, these findings suggest that the specialist's ability to provide liquidity is significantly affected by the trading activity of other stocks at the panel, even after controlling for own-stock trading activity.

[17] For completeness, we also estimated all of the models using random effects. A Hausman test easily rejects the random effects specification for all models.

Panel B of Table 6 presents results based on the continuous measure of panel activity (*LogPanelActivity*). Consistent with the results in Panel A, these findings suggest that transaction costs on NYSE stocks increase as the trading activity at the panel increases. The coefficient on *LogPanelActivity* is positive and significant at the one percent level for all four measures of bid-ask spread and significantly negative for price improvement. In addition, when we interact the panel activity variable with the *Busy* panel dummy variable, the coefficient is significant in all cases. This suggests that the effects of limited attention are most pronounced during extremely busy periods. However, the coefficient on *LogPanelActivity* remains significant, indicating that the allocation of attention across stocks influences transactions costs even in periods when the overall panel activity is normal or low. In sum, we interpret the evidence from Table 6 as supportive of the Limited Attention Hypothesis.

To examine the economic significance of the results, we estimate the predicted value of transaction costs based on the regression coefficients in Table 6 and the mean values of all explanatory variables. The results based on the busy period dummy variable (Table 6, Panel A) suggests that effective spreads are approximately one percent higher during periods of busy panel trading activity than during normal periods. Similar results are obtained based on the continuous panel activity variable (Table 6, Panel B). Here, the results suggest that effective spreads increase by approximately 0.8% if panel trading activity goes from the 25th to the 75th percentile, and by approximately 1.5% if panel trading activity goes from the 10th to the 90th percentile.[18]

To put these costs in perspective, it is useful to characterize total NYSE trading activity. The NYSE Fact Book reports that trading volume on the NYSE averaged approximately $40 billion per day during 2002. Our data suggest that approximately 9.5% of this volume, or $3.8 billion, was executed during periods of busy panel activity. Based on the 19 basis point average effective spread during our sample period, a one percent increase would imply a floor-wide increase in trading costs of approximately

[18] Based on the busy period dummy variable, the effective spread estimate increases from 5.98 cents and 31.05 bps during normal trading periods to 6.04 cents and 31.32 bps during busy trading periods, respectively. Based on the continuous panel activity variable, effective spreads increase from 5.50 cents and 30.64 bps at the 10th percentile of panel activity to 5.59 cents and 31.07 bps at the 90th percentile of panel activity.

$72.2 thousand per day or $18.2 million per year. While these cost estimates are only approximate, they suggest that the potential floor-wide benefits from reductions in capacity constraints (perhaps through the use of automatic execution systems) could be large.[19]

5. Robustness tests

5.1. Individual time-series regressions

Although the pooled analysis should minimize the potential endogeneity bias associated with stock allocation decisions on the NYSE, these effects may not be completely eliminated. As a robustness check, we also estimated individual time-series regressions for each stock based on ordinary least squares. These regressions completely eliminate cross-security effects and focus solely on the time-series relation between panel activity and transaction costs. The models follow the specification in equation (2).

To test our hypothesis, we examined the cross-sectional distribution of the coefficients on all explanatory variables. The results from these tests confirm the findings based on the pooled time-series and cross-sectional analysis. Transaction costs are significantly positively related to the trading activity of other stocks traded by the same specialist, and these effects are most pronounced during periods of unusually busy panel activity. These findings suggest that our results are robust to potential endogeneity concerns and to alternative estimation methods. To conserve space, these results are not reported.

5.2. Normal own-trading periods

As noted above, covariation in order flow across stocks is a significant concern in our analysis. We wish to interpret the effects of panel trading activity in isolation from variation in own stock trading. In the pooled regressions above, we account for this by including own-stock trading activity as an explanatory variable. As an additional test, we repeat our pooled analysis including only those periods

---

[19] To the extent that the most active stocks are unaffected by limited attention, we should ignore these stocks in our cost estimates. If we exclude the most active stock at each panel, the resulting cost estimates are roughly $28 thousand per day or $7.1 million per year. We also note that estimates based on normal own-trading periods lead to higher cost estimates. These tests, reported in Table 7 and discussed below, suggest that effective spreads increase by as much as two percent during periods of busy panel activity. To be conservative, we report cost estimates based on the full-sample estimation.

during which each stock's own trading activity is within one standard deviation of its own time-series mean. The results are presented in Table 7.[20]

The results based on only normal own-trading periods are even stronger than for the full sample and confirm that our findings are robust to own-stock trading considerations. As in the full sample, the coefficient on the *Busy* panel dummy variable (Table 7, Panel A) is positive and significant for all four measures of bid-ask spread. However, the coefficients tend to be approximately 50% larger when estimated during normal own-trading periods than during all trading periods. For example, the coefficient on the busy period dummy variable in the dollar effective spread regressions is 0.084 based on normal own-trading periods, compared to 0.059 based on all trading periods. A similar comparison can be made for all four spread measures. This suggests that the estimates of economic significance discussed in the previous section may be understated (see footnote 19). Turning to price improvement, we see that busy panels are associated with significantly lower price improvement rates and slow panels are associated with significantly higher price improvement rates.

Conclusions based on the continuous panel activity variable during normal own trading (Table 7, Panel B) are similar. *LogPanelActivity* remains positively related to all four measures of bid-ask spread and negatively related to price improvement rates. As in Table 6, panel activity effects are evident during all trading periods, but are most pronounced during periods of unusually busy panel activity. As in Panel A, the magnitude of the coefficients is substantially larger when we focus on only normal own-trading periods. Coefficients on the four spread measures increase by 50 to 90% compared to Table 6.

5.3. Differences across individual stock trade activity levels

The Limited Attention Hypothesis suggests that the effects of panel activity should be most evident for inactive securities. To examine this hypothesis, we re-estimate the pooled regressions separately for each trade-activity subsample. As in Table 7, we focus only on those periods where own-stock trading is "normal." In Table 8 we report results for percentage effective spreads and the rate of

[20] As an alternative specification, we estimated the pooled regressions using all observations with interaction terms to separate the effects of panel activity during periods of normal own-stock trading activity. The results are qualitatively similar.

price improvement for the active and inactive stock subsamples. The results for other spread measures are similar and are not reported. Consistent with the full sample regressions, we find a significant positive relation between percentage effective spreads and panel activity for both active and inactive stocks. However, consistent with our arguments that inactive stocks should be more sensitive to variation in specialist attention, the magnitude of the panel activity coefficients is greater for the inactive securities.[21]

The results for price improvement also differ substantially across groups. For the less active stocks, the price improvement rate is significantly lower during periods of busy panel activity and is significantly negatively related to *LogPanelActivity*. For the most active stocks, however, price improvement rates are not significantly related to any measure of panel activity. To the extent that price improvement reflects the actions of the specialist, these results suggest that specialists are less likely to divert attention from their more active stocks. The results may also reflect that specialist actions are more likely to affect price improvement rates for inactive securities than for active securities. Together, the results in Table 8 indicate that the effects of limited attention are most pronounced for the least active securities.

6. Summary and conclusions

It is well known that human beings are limited in their ability to process information and to perform multiple tasks simultaneously (see Kahneman (1973) and Pashler (1998) for reviews). Despite the documented importance of "limited attention" in other settings, its impact on financial markets has only recently attracted attention. Recent research provides evidence that, when faced with complicated information, investors resort to simplified decision rules such as categorization. Prior studies also suggest that limited attention on the part of investors can have significant effects on asset price behavior and financial reporting decisions by firms. To date, however, empirical studies have been limited by the absence of direct measures of attention and its allocation across securities.

[21] In unreported regressions, we reestimated the full sample analysis including interaction terms between panel activity and trade activity category. The results confirm that the coefficient on *LogPanelActivity* differs significantly across trade activity categories. These results are available upon request.

This study provides the first direct evidence that limited attention influences the provision of liquidity in financial markets. Since individual NYSE specialists are assigned a well-defined set of securities, this setting provides an ideal framework for analyzing the allocation of limited attention across securities. We show that transaction costs for NYSE securities are significantly affected by the level of trading activity in other securities traded by the same specialist. Consistent with the Limited Attention Hypothesis, our evidence indicates that market makers face attention limits and that they allocate their effort across multiple securities. In the process, their ability to act as an important source of liquidity is reduced for at least a subset of the securities in their market making portfolio. Therefore, while the design of the NYSE may yield diversification and/or subsidization benefits by allowing specialists to handle a portfolio of stocks, our paper identifies potential costs associated with this organizational arrangement.

We do not argue that specialists actively impair liquidity during periods of busy panel activity. Instead, it is likely that the specialist is unable to act as an additional source of liquidity during these periods, leaving spreads to be determined primarily by the actions of other traders. This interpretation is consistent with prior evidence that the NYSE tends to provide lower transaction costs than other trading systems. During normal conditions, the specialist steps in to facilitate trade and act as an important source of liquidity. During busy periods, however, the specialist's ability to provide this service is limited.

This research has important implications for future theoretical work considering the influence of limited attention on trading. In particular, attention limits could influence prices through the supply of liquidity as well as through the demand for liquidity. Our paper also has implications for the allocation of stocks on the trading floor and contributes to the ongoing debate over the merits of floor-based trading versus electronic trading. In particular, increased automation, such as that currently being considered at the NYSE, may relieve specialist capacity constraints and reduce the necessity to allocate effort across stocks. While our tests are based on data from the NYSE our findings may be applicable to any market where effort must be allocated across multiple securities.

References

Barber, B. M., Odean, T., 2005. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors, Unpublished Working Paper, University of California, Berkeley.

Battalio, R., Ellul, A., Jennings, R., 2004. Reputation effects in trading on the New York Stock Exchange, Unpublished Working Paper, University of Notre Dame.

Benveniste, L. M., Marcus, A. J. Wilhelm, W. J., 1992. What's special about the specialist?. Journal of Financial Economics 32, 61-86.

Cao, C., Choe, H., Hatheway, F., 1997. Does the specialist matter? Differential execution costs and inter-security subsidization on the NYSE, Journal of Finance 52, 1615-1640.

Copeland, T. E., Galai, D. 1993. Information effects and the bid-ask spread. Journal of Finance 38, 1457-1469.

Corwin, S. A., 1999. Differences in trading behavior across NYSE specialist firms. Journal of Finance 54, 721-745.

Corwin, S. A., 2004. Specialist performance and new listing allocations on the NYSE: An empirical analysis. Journal of Financial Markets 7, 27-51.

Coughenour, J., Deli, D., 2002. Liquidity provision and the organizational form of NYSE specialist firms. Journal of Finance 56, 841-869.

Coughenour, J., Harris, L, 2005. Specialist profits and the minimum price increment, Unpublished Working Paper, University of Delaware.

Coughenour, J., Saad, M., 2004. Common market makers and commonality in liquidity. Journal of Financial Economics, 73, 37-70.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69-90.

Gabaix, X., Laibson, D., 2004. Bounded rationality and directed cognition. Unpublished Working Paper, MIT and NBER.

Gabaix, X., D. Laibson, G. Moloche, Weinberg, S., 2003. The allocation of attention: Theory and evidence. Unpublished Working Paper, MIT.

Gehrig, T., Jackson, M., 1998. Bid-ask spreads with indirect competition among specialists. Journal of Financial Markets 1, 89-119.

Gifford, S., 2001. Limited attention as the bound on rationality. Unpublished Working Paper, Rutgers University.

Glosten, L. R., 1989. Insider trading, liquidity, and the role of the monopolist specialist. Journal of Business 62, 211-236.

Glosten, L. R., Milgrom, P.R., 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders, Journal of Financial Economics 14, 71-100.

Hasbrouck, J., Seppi, D. J., 2001. Common factors in prices, order flows, and liquidity. Journal of Financial Economics 59, 383-411.

Hasbrouck, J., Sofianos, G., 1993. The trades of market-makers: an analysis of NYSE specialists. Journal of Finance 48, 1565-1594.

Hatch, B. C., Johnson, S. A., 2002. The impact of specialist firm acquisitions on market quality. Journal of Financial Economics 66, 139-167.

Hatch, B. C., Johnson, S. A., Lei, A. Y. C., 2004. Net dealer revenue and the speed of quote adjustments on the NYSE, Unpublished Working Paper, Texas A&M University.

Hirshleifer, D., Teoh, S. H., 2003. Limited attention, information disclosure, and financial reporting. Journal of Accounting and Economics 36, 337-386.

Hirshleifer, D., Lim, S. S., Teoh, S. H., 2004. Disclosure to an audience with limited attention. Unpublished Working Paper, The Ohio State University.

Huang, R. D., Liu, J. W., 2004. Do individual NYSE specialists subsidize illiquid stocks? Unpublished Working Paper, University of Notre Dame.

Ho, T., Stoll, H. R., 1981. Optimal dealer pricing under transactions and return uncertainty. Journal of Financial Economics 9, 47-73.

Huberman, G., 2001. Familiarity breeds investment. Review of Financial Studies 14, 659-680.

Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: A nonevent that made market prices soar. Journal of Finance 56, 387-396.

Kahneman, D., 1973. Attention and Effort (Prentice Hall, New Jersey).

Kavajecz, K. A., 1999. A specialist's quoted depth and the limit order book. Journal of Finance 54, 747-771.

Leach, J. C., Madhavan, A., 1993. Price experimentation and security market structure. Review of Financial Studies 6, 375-404.

Madhavan A., Smidt, S., 1993. An analysis of daily changes in specialists' inventories and quotations. Journal of Finance 48, 1595-1628.

Madhavan, A., Sofianos, G., 1998. An empirical analysis of NYSE specialist trading. Journal of Financial Economics 48, 189-210.

McInish, T. H., Wood, R. A., 1992. An analysis of intraday patterns in bid-ask spreads for NYSE stocks. Journal of Finance 47, 753-764.

O'Hara, M., Oldfield, G. S., 1986. The microeconomics of market making. Journal of Financial and Quantitative Analysis 21, 361-376.

Pashler, H. E., 1998. The Psychology of Attention (MIT Press).

Peng, L., 2005. Learning with Information Capacity Constraints. Journal of Financial and Quantitative Analysis, *forthcoming*.

Peng, L., Xiong, W., 2005. Limited attention and asset prices. Journal of Financial Economics, *forthcoming*.

Radner, R., Rothschild, M., 1975. On the allocation of effort. Journal of Economic Theory 10, 358-376.

Sargent, T. J., 1993. Bounded Rationality in Macroeconomics (Oxford University Press, 1993).

Schwartz, R. A., 1993. Reshaping the Equity Markets: A Guide for the 1990s (Business One Irwin, Homewood, IL).

Simon, H. A., 1955. A behavioral model of rational choice. Quarterly Journal of Economics 69, 99-118.

Simon, H. A., 1979. Rational decision making in business organizations. The American Economic Review 69(4), 493-513.

Sofianos, G., 1995. Specialist gross trading revenues at the New York Stock Exchange. Unpublished Working Paper, New York Stock Exchange (#95-01).

Stoll, H. R., 1978. The supply of dealer services in securities markets. Journal of Finance 33, 1133-1151.

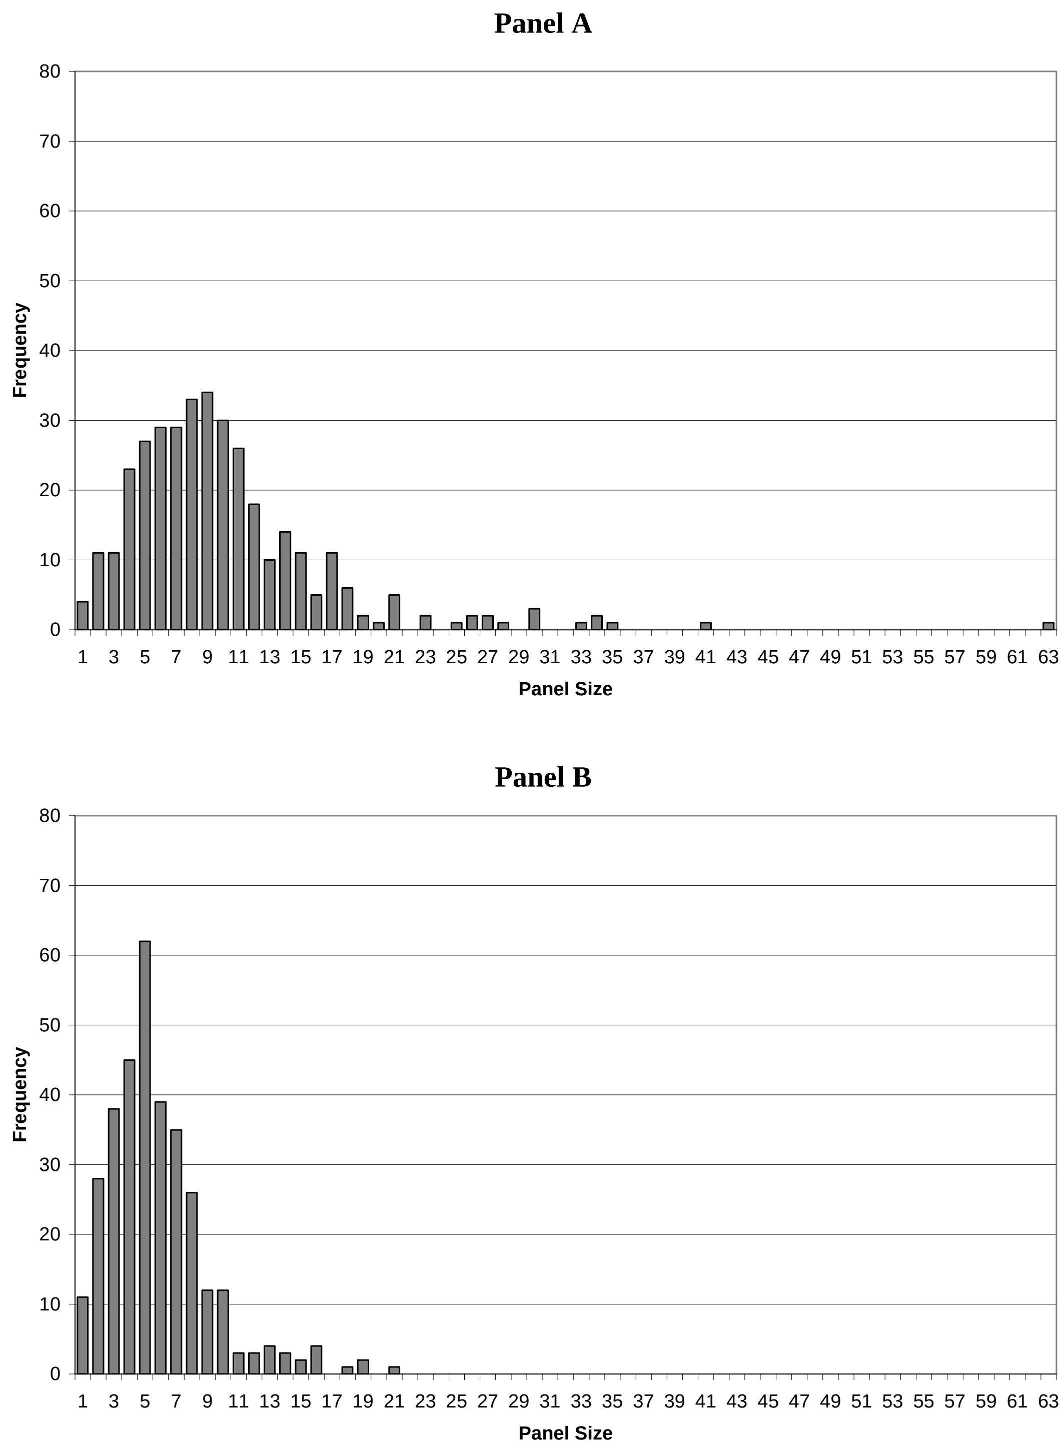

Figure 1 – Distribution of Panel Size

The figure shows the distribution of panel size on the New York Stock Exchange as of August 1, 2002. Panel A reflects panel size based on the full sample of 3,599 securities in the August 1st NYSE Specialist Directory. Panel B reflects panel size after restricting the sample to only common stocks and ADRs as identified in CRSP (N=1,920). We use this set of 1,920 stocks to compute all panel size and panel activity variables.

Table 1
Summary statistics

The sample includes 1,262 common stocks and ADRs listed on the NYSE from August 1 through October 31, 2002. All securities meet the following restrictions: (1) identified in both CRSP and the NYSE Specialist Directory, (2) no stock splits during the sample period, (3) a minimum transaction price greater than $2 and an average transaction price of less than $200, and (4) at least one trade in 800 of the 840 available 30-minute trading periods. All variables are averaged across 30-minute trading intervals for each security. The table then reports the cross-sectional mean of each variable. Quoted and effective bid-ask spreads are trade weighted. Price improvement is the proportion of all trades that occur at prices inside the bid-ask quotes. Stocks are divided into three categories based on the average daily number of trades: the 100 most-active stocks, the next 400 most-active stocks, and the least active 762 stocks. Means are reported separately for these three trading activity categories. The *p*-value is from a test of the restriction that means are equal across trade activity categories, based on analysis of variance.

Panel A – Full Sample				
	Mean	Median	Min	Max
Half-hours traded	837.06	840.00	800.00	840.00
Price ($)	26.51	23.40	3.15	121.72
Dollar volume per half-hour ($000)	1883.18	519.62	18.32	43219.18
Share volume per half-hour (000)	66.63	23.70	1.62	1578.20
Trades per half-hour	53.38	36.13	5.94	387.48
Trade size	852.59	637.82	637.57	6546.23
Std deviation of midpoint returns (%)	0.88	0.79	0.11	3.04
Quoted spread (¢)	5.30	4.92	1.84	23.82
Effective spread (¢)	3.78	3.49	1.23	20.08
% Quoted spread (bps)	27.37	21.16	6.57	162.66
% Effective spread (bps)	19.46	14.96	4.33	114.04
Price improvement (% of trades)	39.89	39.82	18.10	60.35

Panel B – Trade Activity Subsamples				
	High Activity (N=100)	Mid Activity (N=400)	Low Activity (N=762)	*p*-value
Half-hours traded	839.87	839.94	835.17	0.000
Price ($)	39.74	31.92	21.93	0.000
Dollar volume per half-hour ($000)	11785.18	2385.25	320.15	0.000
Share volume per half-hour (000)	362.51	88.29	16.42	0.000
Trades per half-hour	179.18	78.96	23.45	0.000
Trade size	1980.81	1036.58	607.95	0.000
Std deviation of midpoint returns (%)	0.88	0.89	0.92	0.000
Quoted spread (¢)	4.48	4.59	5.78	0.000
Effective spread (¢)	3.07	3.22	4.16	0.000
% Quoted spread (bps)	13.70	17.37	34.42	0.000
% Effective spread (bps)	9.21	12.06	24.69	0.000
Price improvement (% of trades)	42.57	40.11	39.43	0.000

Table 2
Panel characteristics

The table provides specialist panel characteristics from the NYSE trading floor as of August 1, 2002. In the first row of the table, panel size is defined based on the full set of 3,599 securities listed in the August 1st NYSE Specialist Directory. In row two, panel size is defined based on the set of 1,920 common stocks and ADRs with available CRSP data. In row three, panel size is defined based on the final restricted sample of 1,262 stocks. All securities in the final sample meet the following restrictions: (1) identified in both CRSP and the NYSE Specialist Directory, (2) no stock splits during the sample period, (3) a minimum transaction price greater than $2 and an average transaction price of less than $200, and (4) at least one trade in 800 of the 840 available 30-minute trading periods. % Post Change and % Panel Change are the proportions of securities in each sample that change post or panel locations, respectively, during the sample period.

Sample	Panels	Panel Size				Securities	% Post Changes	% Panel Changes
		Mean	Min	Median	Max			
Full Specialist Directory	357	10.1	1.0	9.0	63.0	3,599	15.0	27.8
Common Stocks	331	5.8	1.0	5.0	21.0	1,920	17.6	31.5
Final Restricted Sample	323	3.9	1.0	4.0	12.0	1,262	17.3	32.4

Table 3
Panel size and trading characteristics by trade activity category

The table provides average panel size and trading characteristics for the sample of 1,262 common stocks and ADRs listed on the NYSE from August 1 through October 31, 2002. All securities meet the following restrictions: (1) identified in both CRSP and the NYSE Specialist Directory, (2) no stock splits during the sample period, (3) a minimum transaction price greater than $2 and an average transaction price of less than $200, and (4) at least one trade in 800 of the 840 available 30-minute trading periods. For each security, panel and trading characteristics are averaged across all intraday trading periods, and we report the cross-sectional average of these individual stock means. "All Securities at Panel" is the panel size based on all securities listed in the NYSE Specialist Directory. "All Common Stocks at Panel" is the panel size based on the full sample of 1,920 common stocks and ADRs listed on the NYSE during the sample period (prior to applying price and trading restrictions). Panel ranks and market shares for dollar volume and trades are also defined based on the unrestricted sample of 1,920 stocks. Sample stocks are divided into three categories based on the average daily number of trades: the 100 most-active stocks, the next 400 most-active stocks, and the least active 762 stocks. Means are reported separately for each category. The *p*-value is from a test of the restriction that means are equal across trade activity categories, based on analysis of variance.

	Trade Activity Category:			
	High Activity (N=100)	Mid Activity (N=400)	Low Activity (N=762)	*p*-value
All Securities at Panel	7.16	9.35	12.45	0.000
All Common Stocks at Panel	3.95	5.77	7.74	0.000
Rank among Common Stocks at Panel	1.14	1.82	3.84	0.000
Share of Dollar Volume at Panel (%)	72.79	37.37	11.19	0.000
Share of Trades at Panel (%)	61.57	34.20	13.55	0.000

Table 4
Cross sectional relation between transaction costs and panel activity

The table lists coefficient estimates from OLS regressions of transaction costs on measures of panel activity and a set of control variables. The model includes 1,262 firms and all variables are firm-specific averages across 840 30-minute trading periods from August 1 through October 31, 2002. The dependent variables include dollar and percentage quoted and effective spreads, as well as the rate of price improvement. Spreads are transaction-weighted averages executed within the 30-minute interval. The rate of price improvement is the percentage of trades during the 30-minute interval that are executed inside the bid-ask quotes. The independent variables include the natural log of panel activity (*LogPanelActivity*), the inverse of the average trade price (*InvPrice*), the natural log of the average number of trades (*LogTrades*), the natural log of the average trade size (*LogTradeSize*), and the standard deviation of midpoint-to-midpoint returns (*StDevRet*). Individual specialist portfolios are identified using daily NYSE specialist directories. For each firm and each trading period, we define panel activity as the total number of trades during the period for all other stocks traded at the same specialist panel. For each firm, this variable is then averaged across all trading periods. *HighActive* is set equal to one for the 100 most active stocks, and equals zero otherwise; similarly *MidActive* (N=400) and *LowActive* (N=762) are dummy variables used to identify stocks in the other trade-frequency categories. *p*-values are reported in parentheses below the coefficients.

	Quoted Spread (¢)	Effective Spread (¢)	Quoted Spread (%)	Effective Spread (%)
Intercept	16.880	12.253	62.340	45.863
	(0.000)	(0.000)	(0.000)	(0.000)
InvPrice	-29.716	-22.447	285.962	195.877
	(0.000)	(0.000)	(0.000)	(0.000)
LogTrades	-1.999	-1.527	-12.653	-9.646
	(0.000)	(0.000)	(0.000)	(0.000)
LogTradeSize	-0.706	-0.479	-2.566	-1.642
	(0.000)	(0.000)	(0.000)	(0.000)
StDevRet	184.239	137.660	1265.553	939.609
	(0.000)	(0.000)	(0.000)	(0.000)
*LogPanelActivity*HighActive*	0.531	0.378	2.826	1.975
	(0.000)	(0.000)	(0.000)	(0.000)
*LogPanelActivity*MidActive*	0.161	0.117	0.685	0.432
	(0.051)	(0.056)	(0.080)	(0.136)
*LogPanelActivity*LowActive*	-0.054	-0.037	-0.675	-0.549
	(0.514)	(0.549)	(0.087)	(0.060)
Adjusted R^2	0.480	0.485	0.868	0.861

Table 5
Spreads and price improvement rates by panel activity and own-trading activity

The table lists mean values for five measures of transaction costs categorized by panel trading activity and own-stock trading activity. For each stock, we estimate the mean and standard deviation of the number of trades across all 30-minute trading periods. Periods of low, normal, and high own trading activity are defined as periods more than one standard deviation below the mean, within one standard deviation of the mean, and greater than one standard deviation above the mean, respectively. Periods of low, normal, and high panel trading activity are defined similarly, where panel activity for a given stock equals the total number of trades for all other stocks traded at the same floor panel. The transaction cost measures are dollar quoted and effective spreads, percentage quoted and effective spreads, and the rate of price improvement. Results are provided separately for the 100 most active NYSE stocks (High Activity Stocks), the next 400 most active stocks (Mid), and the 762 least active stocks (Low). *p*-values are from a test of the restriction that means are equal across these trade activity categories, based on an analysis of variance.

Panel A: Own trading activity is normal

	Panel Activity			
	Low	Normal	High	*p*-value
Low activity stocks (N=762)				
Effective spread (¢)	4.03	4.16	4.19	0.182
Quoted spread (¢)	5.62	5.77	5.78	0.344
Effective spread (b.p.)	23.23	24.72	25.45	0.029
Quoted spread (b.p.)	32.53	34.43	35.15	0.072
Rate of price improvement (%)	39.96	39.50	38.09	0.000
Mid activity stocks (N=400)				
Effective spread (¢)	3.12	3.22	3.23	0.243
Quoted spread (¢)	4.64	4.60	4.59	0.345
Effective spread (b.p.)	11.31	12.07	12.28	0.015
Quoted spread (b.p.)	16.38	17.42	17.63	0.043
Rate of price improvement (%)	40.62	40.44	39.27	0.001
High activity stocks (N=100)				
Effective spread (¢)	2.89	3.06	3.18	0.153
Quoted spread (¢)	4.25	4.47	4.64	0.154
Effective spread (b.p.)	8.59	9.14	9.61	0.208
Quoted spread (b.p.)	12.86	13.64	14.33	0.291
Rate of price improvement (%)	42.52	42.86	42.70	0.911

Table 5
continued

Panel B: Own trading activity is low				
	Panel Activity			
	Low	Normal	High	*p*-value
Low activity stocks (N=762)				
Effective spread (¢)	4.00	4.24	4.41	0.002
Quoted spread (¢)	5.61	5.88	6.16	0.002
Effective spread (b.p.)	22.85	24.45	26.64	0.001
Quoted spread (b.p.)	32.28	34.18	37.46	0.002
Rate of price improvement (%)	39.83	39.39	38.78	0.266
Mid activity stocks (N=400)				
Effective spread (¢)	3.18	3.23	3.48	0.006
Quoted spread (¢)	4.46	4.54	4.89	0.002
Effective spread (b.p.)	11.45	11.76	12.76	0.011
Quoted spread (b.p.)	16.16	16.76	18.09	0.005
Rate of price improvement (%)	39.11	39.38	39.60	0.605
High activity stocks (N=100)				
Effective spread (¢)	2.93	3.11	3.51	0.007
Quoted spread (¢)	4.26	4.51	5.06	0.004
Effective spread (b.p.)	8.45	9.22	10.66	0.002
Quoted spread (b.p.)	12.56	13.63	15.61	0.008
Rate of price improvement (%)	42.39	42.72	43.56	0.466
Panel C: Own trading activity is high				
Low activity stocks (N=762)				
Effective spread (¢)	4.37	4.28	4.00	0.000
Quoted spread (¢)	6.16	5.99	5.57	0.000
Effective spread (b.p.)	25.95	25.69	24.26	0.108
Quoted spread (b.p.)	36.75	36.00	33.77	0.036
Rate of price improvement (%)	42.38	40.45	37.78	0.000
Mid activity stocks (N=400)				
Effective spread (¢)	3.53	3.38	3.06	0.000
Quoted spread (¢)	5.03	4.84	4.37	0.000
Effective spread (b.p.)	13.24	12.95	11.90	0.003
Quoted spread (b.p.)	19.15	18.76	17.18	0.005
Rate of price improvement (%)	41.14	40.60	38.21	0.000
High activity stocks (N=100)				
Effective spread (¢)	3.07	3.18	3.02	0.593
Quoted spread (¢)	4.46	4.61	4.41	0.632
Effective spread (b.p.)	9.34	9.81	9.37	0.692
Quoted spread (b.p.)	13.85	14.40	13.88	0.814
Rate of price improvement (%)	42.10	41.87	40.75	0.233

Table 6

Pooled time-series analysis of transaction costs and panel activity

The table reports coefficient estimates from pooled time-series and cross-section regressions relating transaction costs to measures of panel activity. The model includes 1,262 firms and 840 30-minute trading periods from August 1 through October 31, 2002. The dependent variables include dollar and percentage quoted and effective spreads, and the rate of price improvement. Spreads are defined as a trade-weighted average within each 30-minute interval. The rate of price improvement is the percentage of trades during the 30-minute interval that are executed inside the bid-ask quotes. The independent variables include measures of panel activity, the inverse of the average trade price (*InvPrice*), the natural log of the average number of trades (*LogTrades*), the natural log of the average trade size (*LogTradeSize*), and the standard deviation of trade-by-trade midpoint returns (*StDevRet*). Individual specialist portfolios are identified using daily NYSE specialist directories. For each stock and each trading period, *Panel Activity* is defined as the total number of trades during the period for all other stocks traded at the same specialist panel. Busy (slow) periods are then defined as periods for which panel activity is more than one standard deviation above (below) the mean, based on panel-specific means and standard deviations. In Panel A, the model is estimated using dummy variables to identify *Busy* and *Slow* periods based on panel activity. In Panel B, the model is estimated including panel activity and the interaction between panel activity and the busy period dummy variable. All models also include firm fixed effects and dummy variables for time-of-day effects. *p*-values are reported in parentheses below the coefficients.

	Quoted Spread (¢)		Effective Spread (¢)		Quoted Spread (%)		Effective Spread (%)		Rate of Price Improvement	
	Panel A – Panel Activity Dummy Variables									
Intercept	7.510	7.510	5.278	5.279	18.967	18.966	12.872	12.875	0.313	0.313
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
InvPrice	-8.977	-8.985	-6.617	-6.628	311.157	311.172	220.776	220.726	-0.035	-0.032
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.003)	(0.000)
LogTrades	0.019	0.018	-0.035	-0.036	0.043	0.043	-0.253	-0.254	0.013	0.014
	(0.001)	(0.000)	(0.000)	(0.000)	(0.218)	(0.216)	(0.000)	(0.000)	(0.000)	(0.000)
LogTradeSize	0.184	0.184	0.190	0.190	0.974	0.974	1.021	1.022	-0.017	-0.017
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
StDevRet	0.173	0.173	0.022	0.022	0.543	0.543	-0.013	-0.012	-0.005	-0.005
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.859)	(0.858)	(0.000)	(0.000)
Quoted Spread ($)	-	-	-	-	-	-	-	-	0.024	0.024
									(0.000)	(0.000)
Busy Panel	0.067	0.067	0.059	0.059	0.274	0.274	0.273	0.274	-0.006	-0.003
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Slow Panel	-	-0.007	-	-0.011	-	0.014	-	-0.046	-	0.003
		(0.503)		(0.182)		(0.824)		(0.343)		(0.000)
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-Day Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.433	0.433	0.407	0.407	0.616	0.616	0.576	0.576	0.264	0.264

Table 6
continued

	Quoted Spread (¢)		Effective Spread (¢)		Quoted Spread (%)		Effective Spread (%)		Rate of Price Improvement	
					Panel B – Continuous Panel Activity					
Intercept	7.213 (0.000)	7.278 (0.000)	4.956 (0.000)	5.007 (0.000)	17.612 (0.000)	17.868 (0.000)	11.366 (0.000)	11.604 (0.000)	0.345 (0.000)	0.344 (0.000)
InvPrice	-9.039 (0.000)	-9.051 (0.000)	-6.688 (0.000)	-6.697 (0.000)	310.584 (0.000)	310.536 (0.000)	220.345 (0.000)	220.301 (0.000)	-0.028 (0.023)	-0.028 (0.024)
LogTrades	0.018 (0.001)	0.017 (0.001)	-0.036 (0.000)	-0.037 (0.000)	0.036 (0.297)	0.034 (0.324)	-0.259 (0.000)	-0.261 (0.000)	0.014 (0.000)	0.014 (0.000)
LogTradeSize	0.184 (0.000)	0.185 (0.000)	0.191 (0.000)	0.191 (0.000)	0.969 (0.000)	0.971 (0.000)	1.022 (0.000)	1.023 (0.000)	-0.017 (0.000)	-0.017 (0.000)
StDevRet	0.173 (0.000)	0.173 (0.000)	0.022 (0.045)	0.022 (0.045)	0.541 (0.000)	0.541 (0.000)	-0.013 (0.847)	-0.013 (0.848)	-0.005 (0.000)	-0.005 (0.000)
Quoted Spread ($)	-	-	-	-	-	-	-	-	0.024 (0.000)	0.024 (0.000)
LogPanelActivity	0.054 (0.000)	0.041 (0.000)	0.057 (0.000)	0.047 (0.000)	0.253 (0.000)	0.203 (0.000)	0.273 (0.000)	0.227 (0.000)	-0.006 (0.000)	-0.005 (0.000)
*LogPanelActivity*Busy Panel*	-	0.009 (0.000)	-	0.007 (0.000)	-	0.033 (0.002)	-	0.031 (0.000)	-	-0.000 (0.131)
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-Day Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.433	0.433	0.407	0.407	0.616	0.616	0.576	0.576	0.264	0.264

Table 7
Pooled time-series analysis of transaction costs and panel activity during normal own-trading periods

The table reports coefficient estimates from pooled time-series and cross-section regressions of transaction costs on measures of panel activity, including only those observations where the number of trades in the stock is within one standard deviation of its time-series mean. The model includes 1,262 firms and 840 30-minute trading periods from August 1 through October 31, 2002. The dependent variables include dollar and percentage quoted and effective spreads, as well as the rate of price improvement. Spreads are defined as a trade-weighted average within each 30-minute interval. The rate of price improvement is the percentage of trades during the 30-minute interval that are executed inside the bid-ask quotes. The independent variables include measures of panel activity, the inverse of the average trade price (*InvPrice*), the natural log of the average number of trades (*LogTrades*), the natural log of the average trade size (*LogTradeSize*), and the standard deviation of trade-by-trade midpoint returns (*StDevRet*). Individual specialist portfolios are identified using daily NYSE specialist directories. For each stock and each trading period, *Panel Activity* is defined as the total number of trades during the period for all other stocks traded at the same specialist panel. Busy (slow) periods are then defined as periods for which panel activity is more than one standard deviation above (below) the mean, based on panel-specific means and standard deviations. In Panel A, the model is estimated using dummy variables to identify *Busy* and *Slow* periods based on panel activity. In Panel B, the model is estimated including panel activity and the interaction between panel activity and the busy period dummy variable. All models also include firm fixed effects and dummy variables for time-of-day effects. *p*-values are reported in parentheses below the coefficients.

	Quoted Spread (¢)		Effective Spread (¢)		Quoted Spread (%)		Effective Spread (%)		Rate of Price Improvement	
	Panel A – Panel Activity Dummy Variables									
Intercept	7.960	7.960	5.519	5.520	22.090	22.096	14.541	14.549	0.313	0.313
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
InvPrice	-9.202	-9.217	-6.671	-6.686	308.971	308.845	221.371	221.232	-0.075	-0.071
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LogTrades	-0.029	-0.030	-0.057	-0.057	-0.345	-0.348	-0.466	-0.468	0.015	-0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LogTradeSize	0.140	0.140	0.160	0.160	0.728	0.728	0.869	0.869	-0.017	0.024
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
StDevRet	0.173	0.173	0.034	0.034	0.491	0.490	-0.003	-0.003	-0.004	-0.004
	(0.000)	(0.000)	(0.007)	(0.008)	(0.000)	(0.000)	(0.970)	(0.967)	(0.000)	(0.000)
Quoted Spread ($)	-	-	-	-	-	-	-	-	0.024	0.024
									(0.000)	(0.000)
Busy Panel	0.096	0.096	0.084	0.084	0.457	0.457	0.416	0.417	-0.004	-0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Slow Panel	-	-0.014	-	-0.014	-	-0.113	-	-0.125	-	0.004
		(0.237)		(0.138)		(0.137)		(0.036)		(0.000)
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-Day Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.436	0.437	0.411	0.411	0.619	0.619	0.579	0.579	0.268	0.268

Table 7
continued

	Quoted Spread (¢)		Effective Spread (¢)		Quoted Spread (%)		Effective Spread (%)		Rate of Price Improvement	
	Panel B – Continuous Panel Activity									
Intercept	7.523	7.617	5.082	5.158	19.504	19.900	12.157	12.511	0.350	0.349
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
InvPrice	-9.299	-9.319	-6.677	-6.794	308.445	308.364	220.858	220.786	-0.063	-0.063
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LogTrades	-0.031	-0.032	-0.059	-0.060	-0.361	-0.364	-0.478	-0.481	0.015	0.015
	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LogTradeSize	0.140	0.141	0.161	0.161	0.730	0.732	0.872	0.874	-0.017	-0.017
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
StDevRet	0.173	0.173	0.034	0.034	0.489	0.489	-0.004	-0.004	-0.004	-0.004
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.957)	(0.959)	(0.000)	(0.000)
Quoted Spread ($)	-	-	-	-	-	-	-	-	0.024	0.024
									(0.000)	(0.000)
LogPanelActivity	0.079	0.061	0.078	0.063	0.465	0.390	0.426	0.359	-0.007	-0.006
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
*LogPanelActivity*Busy Panel*	-	0.012	-	0.009	-	0.051	-	0.045	-	-0.000
		(0.000)		(0.000)		(0.000)		(0.011)		(0.270)
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-Day Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.437	0.437	0.411	0.411	0.619	0.619	0.579	0.579	0.268	0.268

Table 8

Pooled time-series analysis of transaction costs and panel activity by stock trading categories

The table reports coefficient estimates from pooled time-series and cross-section regressions of percentage spreads and price improvement rates on measures of panel activity, including only those observations where the number of trades in the stock is within one standard deviation of its time-series mean. Results are provided separately for 762 Inactive firms and 100 Active firms and the data cover 840 30-minute trading periods from August 1 through October 31, 2002. Spreads are defined as a trade-weighted average within each 30-minute interval. The rate of price improvement is the percentage of trades during the 30-minute interval that are executed inside the bid-ask quotes. The independent variables include measures of panel activity, the inverse of the average trade price (*InvPrice*), the natural log of the average number of trades (*LogTrades*), the natural log of the average trade size (*LogTradeSize*), and the standard deviation of trade-by-trade midpoint returns (*StDevRet*). Individual specialist portfolios are identified using daily NYSE specialist directories. For each stock and each trading period, *Panel Activity* is defined as the total number of trades during the period for all other stocks traded at the same specialist panel. Busy (slow) periods are then defined as periods for which panel activity is more than one standard deviation above (below) the mean, based on panel-specific means and standard deviations. All models also include firm fixed effects and dummy variables for time-of-day effects. *p*-values are reported in parentheses below the coefficients.

	Inactive Firms (N=762)				Active Firms (N=100)			
	Effective Spread (%)		Rate of Price Improvement		Effective Spread (%)		Rate of Price Improvement	
Intercept	13.547	10.512	0.342	0.399	1.522	0.476	0.453	0.455
	(0.000)	(0.000)	(0.000)	(0.000)	(0.002)	(0.350)	(0.000)	(0.000)
InvPrice	253.616	252.893	-0.084	-0.065	114.067	113.83	0.150	0.142
	(0.000)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
LogTrades	-0.604	-0.621	0.016	0.016	0.645	0.628	-0.013	-0.012
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LogTradeSize	1.156	1.162	-0.022	-0.022	0.852	0.842	-0.006	-0.006
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
StDevRet	-0.125	-0.127	-0.003	-0.003	-0.243	-0.249	-0.006	-0.006
	(0.242)	(0.237)	(0.002)	(0.002)	(0.012)	(0.010)	(0.016)	(0.017)
Quoted Spread ($)	-	-	0.023	0.023	-	-	0.028	0.028
			(0.000)	(0.000)			(0.000)	(0.000)
Busy Panel	0.435	-	-0.004	-	0.136	-	-0.000	-
	(0.000)		(0.000)		(0.022)		(0.857)	
LogPanelActivity	-	0.539	-	-0.010	-	0.221	-	0.000
		(0.000)		(0.000)		(0.000)		(0.654)
Firm Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time-of-Day Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.531	0.531	0.244	0.245	0.681	0.682	0.456	0.457