Is There Systematic Variation in Market Efficiency?

Dominik M. Rösch, Avanidhar Subrahmanyam, and Mathijs A. van Dijk*

March 2013

Abstract

Market efficiency remains central to the study of financial markets. We examine how the degree of efficiency varies across stocks and over time. We find: (i) systematic variation in efficiency (measured using short-horizon return predictability and put-call parity deviations) across individual stocks, and (ii) systematic variation in aggregate efficiency across different efficiency measures (short-horizon measures as well as longer-horizon measures based on monthly reversals and momentum). The systematic component of market efficiency varies through time with aggregate funding liquidity, frictions that impede arbitrage, and variables that affect market-making efficacy. Our results indicate that microstructural efficiency measures share a common factor with broader measures of market efficiency and quality, and imply that policies that impact funding liquidity can systematically impact variations in this factor.

* Rösch and van Dijk are at the Rotterdam School of Management, Erasmus University; Subrahmanyam is at the UCLA Anderson School. E-mail addresses: drosch@rsm.nl, asubrahm@anderson.ucla.edu, and madijk@rsm.nl. We thank Yakov Amihud, Tarun Chordia, Carole Comerton-Forde, Thierry Foucault, Amit Goyal, Terry Hendershott, Ed Lin, Christophe Pérignon, Vikas Raman, Matti Suominen, Avi Wohl, and participants at the 2012 Brazilian Finance Conference (Sao Paulo), the 2012 EFMA meetings (Barcelona), the 2012 Frontiers of Finance Conference (Warwick Business School), the 2013 Campus for Finance conference (WHU Otto Beisheim School of Management), and at seminars at Deakin University, Erasmus University, Goethe University Frankfurt, and UCLA Anderson for valuable comments. This work was sponsored by the BiG Grid project for the use of computing and storage facilities, with financial support from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (Netherlands Organisation for Scientific Research, NWO). We thank SARA, and in particular Lykle Voort, for technical support on computing and storage, and OneMarketData for the use of their OneTick software. Van Dijk gratefully acknowledges financial support from the Vereniging Trustfonds Erasmus Universiteit Rotterdam and from the NWO through a "Vidi" grant.

Is There Systematic Variation in Market Efficiency?

March 2013

Abstract

Market efficiency remains central to the study of financial markets. We examine how the degree of efficiency varies across stocks and over time. We find: (i) systematic variation in efficiency (measured using short-horizon return predictability and put-call parity deviations) across individual stocks, and (ii) systematic variation in aggregate efficiency across different efficiency measures (short-horizon measures as well as longer-horizon measures based on monthly reversals and momentum). The systematic component of market efficiency varies through time with aggregate funding liquidity, frictions that impede arbitrage, and variables that affect market-making efficacy. Our results indicate that microstructural efficiency measures share a common factor with broader measures of market efficiency and quality, and imply that policies that impact funding liquidity can systematically impact variations in this factor.

Financial markets serve many roles. First, they allow individuals to reallocate consumption across time by investing for the future at an appropriate expected return. Second, they allow people to optimize their reward to risk ratios based on their preferences. Third, financial market prices serve as proxies for aggregate investor beliefs and thus convey important information to financial managers and policy makers. It goes almost without saying that the efficacy of all of these roles depend on the financial market being relatively free of frictions and of high quality. In such an ideal market, prices accurately reflect fundamentals, and, in doing so, obey the law of one price that assets with identical cash flows sell for the same price. Such a market is commonly referred to as an “efficient market” (Fama, 1970).

There are a number of ways to capture the efficiency of a market. For example, the microstructure literature measures efficiency via metrics such as intraday return predictability and variance ratios (Bessembinder, 2003; Chordia, Roll, and Subrahmanyam, 2008). One may also consider the extent to which markets obey the law of one price (such as put-call parity deviations—viz. Finucane, 1991; Cremers and Weinbaum, 2010), or measure the extent to which longer-term return anomalies prevail (for example, the well-known reversal and momentum effects of Jegadeesh, 1990, and Jegadeesh and Titman, 1993). In a perfect market, prices would be efficient at both short and long horizons, and also obey the law of one price.

However, the enforcement of these attributes of a perfect financial market requires arbitrage. But, any arbitrage activity – whether risky or comparatively risk-free – requires capital, and because of frictions, arbitrage may not be perfect at all times or across all securities, as pointed out by Shleifer and Vishny (1997). These frictions can take the form of market illiquidity, limited capital, short-sale constraints, and volatility, and the severity of such frictions may vary considerably over time.[1] For example, earlier literature suggests that secular and non-secular changes in liquidity influence short-horizon market efficiency, where inefficiency is defined by the extent of short-horizon return predictability from past order flow (Boehmer and Kelley, 2009; Chordia, Roll, and Subrahmanyam, 2005, 2008). Other work demonstrates that monthly equity return predictability and the magnitude of mispricings exploited by hedge funds in relative-value trades vary over time with constraints on arbitrage capital (Akbas, Armstrong, Sorescu, and Subrahmanyam, 2010; Mitchell and Pulvino, 2012).

The efficiency of price formation is also likely to differ across different securities. It is well-known that there is substantial cross-sectional variation in the liquidity of individual stocks (e.g., Benston and Hagerman, 1974; Brennan and Subrahmanyam 1995), and in other frictions that impede arbitrage such as short-sales constraints (Nagel, 2005). Thus, there are sound reasons to believe that efficiency metrics such as short-horizon return predictability and put-call parity deviations also exhibit considerable variation across individual securities.

The idea that efficiency varies both over time and across securities raises the question to what extent time-variation in the price efficiency of individual securities is systematic. While market efficiency depends on liquidity, the latter has firm-specific as well as systematic components (Chordia, Roll, and Subrahmanyam, 2000; Hasbrouck and Seppi, 2000). In addition, time-variation in liquidity depends on macroeconomic funding constraints (Brunnermeier and Pedersen, 2009) as well as variables that influence market making behavior, such as return volatility and net order imbalances (Chordia, Roll, and Subrahmanyam, 2002). More generally, systematic fluctuations in the severity of frictions that hamper arbitrage may lead to systematic variation in efficiency across individual securities. Furthermore, since such fluctuations in

[1] See, for example, Shleifer and Vishny (1997), Mitchell, Pulvino, and Stafford (2002), and D'Avolio (2002) for theoretical and empirical explorations of how limits to arbitrage can cause market inefficiencies to persist.

frictions may similarly affect microstructural market quality measures (such as short-horizon return predictability), efficiency measures related to the law of one price (such as put-call parity deviations), and measures of longer-term return anomalies (such as reversal and momentum), they could also give rise to systematic variation in efficiency across different efficiency measures.

Motivated by these observations, in this paper, we do the following. We first compute daily market efficiency measures for individual stocks based on three metrics: short-horizon (intraday) return predictability based on past order flow, intraday return autocorrelation, and put-call parity deviations using a substantial sample of NYSE stocks over a long sample period of fifteen years.[2] We then construct market-wide measures of efficiency from the individual stock measures and estimate the extent of systematic variation in market efficiency (or "commonality in efficiency") across individual stocks, as the R^2s of a "market model" regression of individual stock measures on the market-wide measure. These analyses show that time-variation in market efficiency, as measured by each of the three short-horizon metrics, has a material common component. This common component indicates that market efficiency is prone to systematic variation, providing pointers that not only might liquidity be affected across the board in a financial crisis, but so might pricing efficiency.

Our next goal is to analyze the degree of systematic variation in aggregate market efficiency *across* different efficiency metrics (or "commonality in efficiency" across efficiency metrics). We relate the short-horizon measures of market-wide efficiency to a variance-ratio-

[2] Busse and Green (2002) find that news reports about individual stocks on the financial television network CNBC are incorporated into stock prices within one to two minutes. Epps (1979) studies price formation for firms in one industry (automobiles). He finds rapid (but not instantaneous) adjustments across firms to common industry news. Note that our measure of market efficiency is not mechanically related to liquidity. For example, illiquidity is caused by the presence of informed agents in Kyle (1985) but prices are not predictable from order flow.

based measure of aggregate market quality,[3] and to measures of market-wide efficiency based on longer-horizon anomalies, namely monthly return reversals (Jegadeesh, 1990) and momentum (Jegadeesh and Titman, 1993).[4] We find that both short- and long-horizon measures of market-wide efficiency are mostly positively cross-correlated and the majority of these correlations are significant, suggesting a pervasive market efficiency factor.[5] We extract the factor via principal component analysis from the monthly time-series of our market-wide efficiency measures and show that this first component explains about one-third of the joint variation in our market efficiency measures.

We next propose that this systematic component varies across time for three reasons: (i) variation in funding liquidity proxied by variables such as hedge fund flows and short rates, (ii) variation in frictions impeding arbitrage such as short-horizon volatility and bid-ask spreads, and (iii) variation in microstructural aspects that affect market making efficacy such as trading activity. Our analysis indicates that the systematic component of market efficiency depends on all three classes of determinants. For example, we document that the systematic component is strongly related to hedge fund flows, the TED spread (a common indicator of funding liquidity), returns to the banking sector, short-horizon volatility, and the aggregate number of transactions

---

[3] This efficiency metric is based on the observation that for a random walk price process, the variance of long-horizon returns is n times the variance of short-horizon returns, where n is the number of short-horizon intervals in the longer horizon; see, for example, Bessembinder (2003).

[4] Fama and French (1996) point out that momentum is left unexplained by traditional risk-based factor models, pointing to the notion that it is a form of inefficiency; Fama (1998) argues that momentum is among very few anomalies that survives closer scrutiny. Cooper (1999) and Mase (1999) suggest that investor overreaction can account for monthly return reversals, suggesting that this anomaly also is a form of market inefficiency. We do not include very long-term reversals (DeBondt and Thaler), because of controversy about the power of statistical tests documenting this phenomenon (Fama, 1998).

[5] Griffin, Kelly, and Nardari (2010) examine market efficiency measures across countries, but do not examine common variation in the measures over time. In another significant paper, Pasquariello (2012) considers deviations of arbitrage bounds across countries in foreign exchange markets and American Depositary Receipts (or ADRs), but does not relate these bounds to microstructural efficiency measures or funding constraints.

in the stock market. We find that almost two-thirds of the time-variation in the systematic component of efficiency can be explained by our regressors.

Our demonstration that there is systematic variation in market efficiency is important for several reasons. First, we show that market efficiency is not a static concept: it exhibits significant time-variation and also has a systematic component, across individual stocks and across short- and long-horizon metrics. Second, our results link the microstructure literature, that has mostly concentrated on short-horizon market quality measures of efficiency, to the asset pricing literature that addresses longer-horizon reversals and momentum. Third, we go beyond the well-known link between market liquidity and funding liquidity and demonstrate a further connection between funding liquidity and common variation in the efficiency of price formation. The latter result suggests that policy attempts to increase funding liquidity not only may have a direct impact on trading costs but also systematically affect the efficiency of stock market prices and, in turn, the efficacy with which resources are allocated. Fourth, our results suggest a new source of uncertainty in asset markets, that there are pervasive dynamic shifts to market efficiency.

This paper is organized as follows. In Section 1, we discuss the estimation of microstructural efficiency measures. Sections 2 and 3 respectively estimate these efficiency measures and document systematic variation in these measures, as well as in daily put-call parity deviations, across individual stocks. Section 4 demonstrates that short- and long-horizon measures of efficiency have a common component; and analyzes determinants of time-series variation in this component. Section 5 concludes.

1. Microstructural efficiency estimation: Data and methodology

We begin by estimating efficiency measures obtained from intraday data for individual stocks: the predictability of returns from order flow and from past returns (i.e., return autocorrelation). Following Chan and Fong (2000), Boehmer and Wu (2007), and Chordia, Roll, and Subrahmanyam (2005, 2008), we use short-horizon return predictability from past order flow and short-horizon return autocorrelation as inverse indicators of microstructural efficiency. One immediate issue that arises is whether predictability of intraday returns from order flow and returns is synonymous with market illiquidity. But, we note that illiquidity does not necessarily imply any return predictability from order flow or past returns. In Kyle (1985), for example, markets are illiquid but prices are martingales because market makers are risk-neutral. Return predictability from order flow will arise, however, if market makers have limited risk-bearing capacity.[6] In this case agents such as floor brokers and floor traders will be able to conduct arbitrage trades, which will act to mitigate the predictability. So, in a sense, our measure of predictability, or lack thereof, is a metric of the efficacy of such short-horizon arbitrage.

To estimate microstructural efficiency, we obtain data on all trades and quotes (time-stamped with microsecond precision) as well as their respective sizes for individual U.S. stocks from the Thomson Reuters Tick History (TRTH) database. Our data start in 1996, which is the earliest year available in the TRTH database, and run till the end of 2010. Our sample consists of all NYSE stocks that were an S&P 100 constituent at any time during 1996-2010. We drop one stock (Capital Cities Inc.) since it has only one month of data over this period.

We discard trades that fall outside the continuous trading session (9:30 am till 4:00 pm EST) on the NYSE (in total 2,572,874 trades). We also discard trades with a negative price

[6] See Stoll (1978). Chordia and Subrahmanyam (2004) show that positively autocorrelated imbalances induce a positive relation between returns and lagged imbalances.

(2,395 trades), or a price that is more than 10% different from the trade price of the ten surrounding trades (137 trades). We further drop trades of more than 100,000 shares (543,767 trades), since large trades are often negotiated before they get reported (Glosten and Harris, 1988). Similarly, we discard quotes outside the continuous trading session (3,089,002 quotes), quotes with a non-positive bid or ask price (1,409 quotes), quotes of which the bid price exceeds the ask price (1 quote), and a number of outliers (1,654,494 quotes). We consider quotes to be an outlier when the bid (ask) price is more than 10% different from the average bid (ask) price of the ten surrounding quotes. We also classify quotes as outliers when the ask price is more than US$5 higher than the bid price, or when the proportional quoted spread is greater than 25%. We note that while the absolute numbers of trades and quotes discarded because of these data screens are large, they are small relative to the total number of trades and quotes in the sample. For example, our data screens lead us to discard a mere 0.2% of all trades in the sample.

Our final sample consists of 156 stocks and 1,541,744,764 trades. To prevent survivorship bias, we use data for these stocks over the entire period for which we have data during 1996-2010, and not only during the period over which they were an S&P 100 constituent. We sign trades using the Lee and Ready (1991) algorithm.[7] Because of a decrease in reporting errors since 1998 (Madhavan, Richardson, and Roomans, 2002), we do not use a delay between a trade and its associated quote. We are able to sign 1,541,265,557 trades, which corresponds to 99.97% of all trades in our final sample.

[7] The Lee/Ready algorithm classifies a trade as buyer- (seller-) initiated if it is closer to the ask (bid) of the prevailing quote. If the trade is exactly at the midpoint of the quote, the trade is classified as buyer- (seller-) initiated if the last price change prior to the trade is positive (negative.) Of course, there is inevitably some assignment error, so the resulting order imbalances are imperfect estimates. Lee and Radhakrishna (2000) and Odders-White (2000) indicate that the Lee/Ready algorithm is quite accurate for NYSE stocks, suggesting that assignment errors should have minimal impact on the results.

We estimate the short-horizon efficiency of each individual stock for each day in the sample based on regressions of the returns over short intervals within the day on order imbalance (or on returns) in the previous interval. When estimating efficiency this way, we face a trade-off in choosing the length of the intraday intervals over which we measure returns and order imbalances. On the one hand, the intervals need to be long enough to contain trades to calculate the order imbalance. On the other hand, the intervals need to be short enough to capture predictability in returns. As shown by Chordia, Roll, and Subrahmanyam (2005), "convergence to market efficiency" takes less than 30 minutes in 1996 and around five minutes in 2002. Since our sample period lasts till 2010, we have to use intervals shorter than five minutes to still capture meaningful predictability in the later part of the sample period. However, estimating predictive regressions over much shorter horizons (say, 30-second intervals) would lead us to discard a substantial part of the stock-day observations in our sample (especially in the early years of our sample) due to an abundance of intervals without trades.

In light of these considerations, we estimate predictability based on intraday returns and order imbalances measured over one-minute intervals (with a robustness check based on two-minute intervals). This leaves us with a sufficient number of observations even in the first few years of the sample period (for example, the stocks in our sample have on average 155 one-minute intervals per day with at least one trade in 1996), while we take into account the increase in the speed of convergence to market efficiency over the sample period as a result of, e.g., improved liquidity and the advent of algorithmic trading.

We estimate the extent of short-horizon return predictability from order flow for each stock i and day d in the sample based on the following regression estimated using intraday data aggregated over one-minute intervals:

$$R_{i,d,t} = a_{i,d} + b_{i,d} OIB_{i,d,t-1} + \varepsilon_{i,d} \, , \tag{1}$$

where $R_{i,d,t}$ is the return of stock i in one-minute interval t on day d based on the mid-quote associated with the last trade to the mid-quote of the first trade in the interval (to avoid the bid-ask bounce), and $OIB_{i,d,t-1}$ is the order imbalance for the same stock and day in the previous interval t-1, computed as the difference between the total dollar amount of trades initiated by buyers and sellers (*OIB$*). We refer to the efficiency measure based on this regression specification as the *OIB predictability* measure. For robustness, we also report results for three alternative measures of intraday return predictability from order flow, each named after the single feature that distinguishes them from the *OIB predictability* measure. The *allquotes* measure is based on returns computed using all quotes within each interval rather than only using quotes associated with trades; the *2minutes* measure is based on two-minute instead of one-minute intervals; and the *oib#* measure is based on order imbalance expressed in number of trades rather than dollars (*OIB#*). We also present and discuss the results using one-minute return predictability from past returns instead of past order flow and label this the *autocorrelation* measure.

We require at least one signed trade in both the interval over which we calculate the return as well as the previous interval. This leads us to drop a non-negligible fraction of the intraday intervals in the early years of the sample period, but since 2000 almost all stocks have at least one trade in almost all of the intraday intervals. We also discard stock-days for which we have fewer than 20 one-minute intervals with valid data on the stock return within that interval and on the OIB or return in the preceding interval (in total 9,082 stock-day observations), and days for which TRTH reports a data gap that overlaps with the continuous trading session (in total 56 days).

We use the R^2s of the efficiency regressions in equation (1) as measures of the efficiency of individual stocks, where a higher R^2 indicates a lower degree of efficiency. We then estimate the degree of systematic variation in efficiency for each stock i as the R^2 of the following regression (Karolyi, Lee, and van Dijk, 2012):

$$\Delta Eff_{i,m,d} = \alpha_{i,m} + \beta_{i,m}\Delta MktEff_{i,m,d} + \gamma_{i,m}\Delta MktEff_{i,m,d-1} + \delta_{i,m}\Delta MktEff_{i,m,d+1} + \eta_{i,m}, \qquad (2)$$

where $\Delta Eff_{i,m,d}$ is the change in the efficiency of stock i on day d of month m (based on the R^2 of the daily efficiency regression in equation (1)), and $\Delta MktEff_{i,m,d}$ is the change in market-wide efficiency (defined as the equally-weighted average efficiency across all stocks in our sample excluding stock i). We estimate equation (2) monthly based on daily efficiency estimates within the month. Based on the five different short-horizon efficiency measures, we also obtain five different "commonality in efficiency" measures. For these regressions, we drop stock-months with fewer than 15 days with valid individual stock efficiency estimates within the month (in total 644 months for 9 different stocks). In unreported robustness tests, we estimate equation (2) based on efficiency levels rather than changes and based on contemporaneous market efficiency as the only independent variable (that is, no lead and lagged market-wide efficiency), and obtain similar results.

Table 1 presents summary statistics of the variables constructed based on intraday trade and quote data from TRTH that serve as inputs to our regressions. The table reports cross-sectional summary statistics (the mean, standard deviation, as well as the 25th, 50th (median), and 75th percentile across the 156 stocks in our sample) of the time-series averages by stock of these variables. The average number of trades per day is around 3,000. We note that there is considerable variation in the number of months individual stocks are in the sample. The last row of Table 1 shows that, on average, stocks are in the sample for 144 months (that is, we are able to estimate the commonality in efficiency measure for these months). First Interstate BNCP is in the sample for only three months, while 58 stocks are part of the sample during the entire sample period 1996:02-2010:12 (179 months). In light of the dramatic increase in trading activity over this period (e.g., Angel, Harris, and Spatt, 2011; Chordia, Roll, and Subrahmanyam, 2011), the average number of trades per day is likely to be substantially lower for stocks that are only in the sample during the earlier years. In other words, some of the cross-sectional variation highlighted by the summary statistics in Table 1 is due to variation across stocks in the calendar period over which data are available. The average daily trading volume is an average of 0.097 or US$97m, with an interquartile range from 0.039 to 0.130. The mean and median mid-quote returns are both equal to 0.01 basis point, which corresponds to four basis points per day. There is a slight positive average order imbalance over the one-minute intervals in our sample. The average proportional effective bid-ask spreads in our sample is 0.13%, with a median of 0.09%.

The final three rows of Table 1 provide information on the number of observations included in our analysis and on the impact of our data filters. We require at least 20 one-minute intervals per day with valid data on the stock return within that interval and on the OIB (or return) in the preceding interval. The average number of days we discard per stock is relatively modest at 61, and most of these days are concentrated in the early years of our sample period. Our data filters allow us to estimate the efficiency of the 156 stocks in our sample for on average 2,972 days over the period 1996-2010.

2. Estimating the microstructural efficiency of individual stocks over time

Table 2 presents the results of our regressions to estimate the microstructural efficiency of individual stocks at the daily frequency. We measure the degree of efficiency using regressions to predict short-interval individual stock returns based on lagged order imbalance or lagged returns. As described in Section 1, we run these regressions by stock-day based on intraday data. Table 2 shows the results of four alternative ways of measuring return predictability from order flow (*OIB predictability, allquotes, 2minutes,* and *oib#*) as well as the *autocorrelation* measure which is based on regressions of one-minute mid-quote returns on lagged mid-quote returns.

We find that OIB positively predicts future returns over short intervals. The average coefficient on lagged OIB$ (expressed in billions of dollars) across the approximately 460,000 stock-day regressions is quite similar across the first three short-horizon efficiency measures in Table 2, ranging from 0.511 to 0.997. The coefficient on lagged OIB for the fourth efficiency measure cannot be directly compared to these numbers since OIB in these regressions is expressed in number of trades instead of dollars, but it is also positive. Its magnitude of 0.664 is in fact similar to that of the coefficients on lagged OIB in the first three columns of Table 2, since *OIB$* in the first three columns is scaled by 10^9 and *OIB#* in the fourth column is scaled by 10^4 and since Table 1 shows that *OIB#* is roughly 10^5 as large as *OIB$*. We also note that the return autocorrelation estimate for the full sample is positive at 0.025.

The first number in parentheses below the average coefficient ("*t*-stat avg") is the average *t*-statistic across all stock-day regressions, which is equal to 1.510 for *OIB predictability*, 0.821 for *allquotes*, 0.850 for *2minutes*, 2.189 for *oib#*, and 0.412 for *autocorrelation*. Although for all but one measure the simple average *t*-statistic does not exceed critical values associated with conventional confidence levels, these average *t*-statistics are nonetheless suggestive of non-trivial return predictability over our sample period – especially considering that the *t*-statistics of the individual stock-day regressions can be based on as few as 20 intraday observations. This assertion is confirmed by the second number in parentheses in each column ("*t*-stat cross"), which is the *t*-statistic computed from the cross-sectional distribution of estimated coefficients by stock-day, and which thus exploits the power obtained from the large cross-section of predictability regressions. These *t*-statistics are highly significant and range from 26.9 for *allquotes* to 424.0 for *oib#*.[8] Table 2 also presents the fraction of stock-day regressions which yields a positive ("% positive") and a positive and significant ("% positive significant") coefficient on lagged OIB and on lagged returns. Depending on the efficiency measure, around 70 to 90% of the estimated OIB coefficients are positive, and around 35 to 65% are positive and significant based on the individual *t*-statistics. For the return autocorrelation, about 60% of the coefficients are positive, and 30% are positive and significant.

Overall, Table 2 provides evidence of considerable intraday return predictability in our sample over 1996-2010. The finding that our predictive regressions over one- and two-minute horizons show significant return predictability suggests that these regressions yield meaningful short-horizon efficiency measures that we can use to examine to what extent the efficiency of individual stocks varies over time and to what extent there is systematic variation in efficiency across different stocks.

Figure 1 plots the monthly time-series development of the five market-wide microstructural efficiency measures, computed as the equally-weighted average across stocks of the equally-weighted average R^2 across days within the month for each of the five daily efficiency regressions reported in Table 2. The figure shows that the efficiency of NYSE-listed S&P100 stocks has improved considerably over the sample period 1996-2010. For our *OIB predictability* measure, the market-wide R^2 has declined from around 6% at the beginning of the

[8] These cross-sectional *t*-statistics (pooled *t*-statistics to be more precise) are based on the assumption that the estimation errors in the estimated coefficients are independent across the stock-day regressions for each efficiency measure. If, instead, these errors are positively correlated, the *t*-statistics need to be corrected downwards. Unreported results based on a random selection of calendar days and stocks from our sample suggest that the corrections are nowhere near large enough to make the cross-sectional *t*-statistics insignificant.

sample period to around 1% towards the end. These results indicate that the declining trend in short-term return predictability over the period 1993-2002 documented by Chordia, Roll, and Subrahmanyam (2008) has persisted since then, although the biggest improvement in efficiency during our sample period stems from the 1996-2003 period. The levels of the market-wide R^2s differ across the five efficiency measures, but the time-series patterns are remarkably similar. The average correlation across the five monthly efficiency measures is 0.92. Although the dominant feature of time-variation in market-wide efficiency is the long-run downward trend, there is also substantial variation in market efficiency at a higher frequency. For example, there is a marked increase in the market-wide R^2 for each of the five efficiency measures in July of 2002 followed by a marked drop in the second half of 2002.

To illustrate the degree of cross-sectional variation in short-horizon efficiency across the stocks in our sample, Figure 2 presents yearly boxplots of the *OIB predictability* measure. For each stock, we compute the yearly efficiency as the average R^2 of the daily *OIB predictability* regressions from Table 2 for that stock across days within the year. The boxes reflect the 25th, 50th (median), and 75th percentile of the yearly average R^2s across the 156 stocks in our sample over the period 1996-2010. The top (bottom) of the line above (below) the boxes represent the highest (lowest) individual stock R^2 below (above) the 75th (25th) percentile plus (minus) 1.5 times the interquartile range. Clearly, efficiency does not only vary considerably over time, but also across stocks. For example, in 1996 the average *OIB predictability* R^2 ranges from 2.4% for Mobil Corporation to 11% for Unisys Corporation and the interquartile range is 1.5%. Furthermore, the extent of dispersion in efficiency across stocks also changes over time. For example, the dispersion is considerably larger in 1999 than in 2000, even though the median R^2 is about the same in both years. The dispersion tends to be smaller in later years, but so is the median R^2. In relative terms, there is little indication that the dispersion in efficiency across stocks is lower in the second half of the sample period compared to the first. For example, the ratio between the interquartile range and the median R^2 is 0.25 in 1996 and 0.42 in 2010. Figure 2 also indicates that the time-series development of the market-wide efficiency measures in Figure 1 is not driven by outliers, as the boxplots follow roughly the same development as the market-wide *OIB predictability* graph in Figure 1.

In sum, Figures 1 and 2 suggest that there is considerable variation in the degree of price efficiency over time and across stocks. In the next section, we investigate to what extent there are common components in the time-varying efficiency of individual stocks.

3. Estimating systematic variation in efficiency across individual stocks

We now set out to examine whether there is systematic variation in market efficiency across individual stocks. To further enrich this analysis, we supplement the *OIB predictability* and *autocorrelation* measures of the previous section with a law of one price measure derived from the stock and options markets for individual stocks. The measure is based on implied volatility discrepancies between put and call option prices on a daily basis. The use of this measure deepens our understanding of commonality in efficiency by extending the notion to derivatives markets for individual stocks.

This *put-call parity* measure is estimated using the OptionMetrics database as the absolute difference between the implied volatilities of a call and a put option of the same series (i.e., pairs of options on the same underlying stock with the same strike price and the same expiration date), averaged across all option pairs for each stock. We use end-of-day quotes from all option series with positive implied volatilities, expiring in two weeks to one year, and with a call delta between 0.3 and 0.7.[9] We are able to estimate the *put-call parity* measure for 123 of the 156 stocks in our sample, for on average 151 of the 180 months in our sample period. The average number of underlying (put and call) option series is nine, with a minimum of four and a maximum of 22. We discard stock-days with fewer than three valid underlying option series. The mean absolute put-call parity deviation (expressed in terms of implied volatility) across stock-days in the sample is 1.2%.

To estimate the extent of systematic variation in efficiency across stocks, we run time-series regressions of changes in the efficiency of individual stocks on contemporaneous, lead, and lagged changes in market-wide efficiency (computed as the equally-weighted average efficiency across the stocks in our sample, excluding the stock of interest), see equation (2).[10] We estimate these regressions by stock-month based on daily efficiency estimates for individual stocks within the month. We perform the commonality regressions for the *OIB predictability* measure, the *autocorrelation* measure, and the *put-call parity* measure.

The results are in Table 3. For each of the three efficiency measures, the table reports average coefficients across all stock-month regressions, average *t*-statistics, cross-sectional *t*-statistics, the fraction of coefficients across all stock-month regressions that is positive and that is positive and significant, and the average R^2 as well as the average adjusted R^2 across all stock-month regressions.

Table 3 reveals evidence of significant systematic variation in efficiency across stocks. The average coefficient on contemporaneous changes in market-wide efficiency is positive and economically substantial for all efficiency measures, ranging from 0.389 for the *OIB*

[9] This measure is also used in Cremers and Weinbaum (2010). These authors note that while, strictly speaking, put-call parity does not hold as an equality for the American call options on individual stocks, a discrepancy in implied volatilities from binomial models nonetheless is indicative of an inefficiency in the stock and options markets.

[10] Since our sample consists of 156 stocks, it is unfortunately not feasible to investigate commonality at a finer level, e.g., at that of the industry.

predictability measure to 0.556 for the *put-call parity* measure. The average *t*-statistic of this coefficient is not significant in any of the cases, but the cross-sectional *t*-statistics (computed from the cross-sectional distribution of estimated coefficients by stock-month) range from 7.4 for the *put-call parity* measure to 30.2 for the *autocorrelation* measure, which is clear evidence of significant commonality in efficiency. We note that the cross-sectional *t*-statistics are calculated under the assumption that the estimation errors in the estimated coefficients are independent across the stock-month regressions, a presumption we examine in more detail below. The fraction of individual coefficients on contemporaneous changes in market-wide efficiency that is positive is 57% or more for all three efficiency measures; at least around 8% (*OIB predictability*) and up to 17% (*put-call parity*) of the coefficients are positive and significant on an individual basis. Although lead and lagged changes in market-wide efficiency have coefficients that are positive on average, there is little evidence that they are important in explaining time-variation in the efficiency of individual stocks.

The average (adjusted) R^2s of the commonality regressions range from 19.3% (4.4%) for *OIB predictability* to 33.6% (26.0%) for *put-call parity*. Hence, although Table 3 shows evidence of significant common components in the efficiency of individual stocks, these components leave a substantial fraction of the time-variation in the efficiency of individual stocks unexplained. This suggests that there are also important idiosyncratic sources of time-variation in efficiency. Of course, there may also be non-negligible noise in our short-horizon efficiency estimates for individual stocks. Nonetheless, we note that the R^2s in Table 3 for *OIB predictability* and *autocorrelation* are of the same order of magnitude as the R^2s of similar regressions to estimate commonality in liquidity reported by Karolyi, Lee, and van Dijk (2012), who find that the monthly commonality in liquidity R^2 in the U.S. averaged around 23% over the period 1995-2009. In unreported robustness tests, we rerun the commonality regressions in equation (2) by stock-year (instead of stock-month) based on daily efficiency estimates within the year and obtain similar adjusted R^2s for *OIB predictability* and *autocorrelation* as Chordia, Roll, and Subrahmanyam (2000) report for their yearly regressions to estimate commonality in liquidity. Commonality in microstructural measures of efficiency is thus roughly an equally strong phenomenon as commonality in liquidity. The adjusted R^2 for the *put-call parity* regressions is much larger than the adjusted R^2 for the two microstructural measures, indicating much greater commonality in put-call parity deviations than in intraday predictability of returns from order flow and past returns.

In Table 4, we explore the extent to which the cross-sectional *t*-statistics reported in Table 3 need to be corrected for the effect of cross-equation dependence in estimation error. Footnote 8 of Chordia, Roll, and Subrahmanyam (2000) reports that – under some simplifying assumptions – the ratio of the true standard error to the standard error used when assuming independence can be expressed as $[1 + (N\text{-}1)\times\rho]^{½}$, where ρ is the correlation between each pair of residuals. (This implies that for negative ρ, the standard errors used in Table 3 are too large and the reported *t*-statistics are too small.) Table 4 reports sample statistics of the correlations between residuals of the stock-month commonality regressions across stocks, estimated on an annual basis. The table reveals little evidence of cross-equation dependence. For each of the three short-horizon efficiency measures and for each year, the average correlation is very close to zero and the *p*-values are never significant.

Figure 3 shows the time variation in the degree of commonality in efficiency across stocks. The figure plots the equally-weighted average across stocks of the R^2 of the monthly commonality regressions. Each line in the figure represents commonality in efficiency based on one of the three short-horizon efficiency measures in Table 3 (*OIB predictability*, *autocorrelation*, and *put-call parity*). The figure shows considerable variation for each of the three commonality series. This suggests that the common components in the time-varying efficiency of individual stocks are stronger in some periods than in others. The time-series pattern is at least somewhat similar across the three measures. For example, each measure shows a peak in mid to late 2007, which could be related to the breakdown of arbitrage as a result of the "quant crisis" (Khandani and Lo, 2011), and towards the end of 2008, possibly related to the onset of the global financial crisis around the Lehman Brothers collapse. The correlation across the three monthly commonality in efficiency measures is positive at 0.28.

4. Estimating systematic variation across aggregate efficiency measures

The previous section looked for evidence on systematic variation in efficiency across individual stocks. In this section, we consider the hitherto unexplored issue of whether there is systematic variation across market-wide efficiency measures. In other words, we examine commonality across measures of market efficiency at the aggregate market level. The underpinning argument is that if the efficacy of arbitrage varies over time owing to variation in the availability of capital or funding constraints, both intraday and longer-horizon versions of efficiency will tend to vary systematically over time. The key innovation here is that we link high-frequency, microstructural measures of market quality to broader, longer-horizon market efficiency metrics.

4.1 Common variation in short- and long-horizon measures of efficiency

We extend the efficiency analyses of the previous section in two ways. First, we add a measure of efficiency used in Bessembinder (2003): namely, a variance ratio that examines how closely the market adheres to a random walk benchmark. Second, we add two measures based on longer-horizon return anomalies that are prominent in the asset pricing literature: namely, monthly reversals (Jegadeesh, 1990) and momentum (Jegadeesh and Titman, 1993). While there are a large number of cross-sectional anomalies (Fama and French, 2008; McLean and Pontiff, 2012), for parsimony, we focus on the two well known departures from weak-form efficiency in the cross-section; this accords with our usage of the short-horizon predictability of returns from past returns and order flow.

Specifically, we consider the following measures:

- *Variance ratio*: |1 - 13×VAR(30-min) / VAR(OC)|, where VAR(30-min) is the return variance estimated from 30-minutes, market-wide (equal-weighted) mid-quote returns within a day and VAR(OC) is the return variance estimated from open-to-close, market-wide mid-quote returns. The scaling factor 13 is based on the number of 30-minute intervals within the 6.5 hour trading day. We discard stock-days with fewer than four non-zero 30-minute returns. The variance ratio is estimated each month and tends to unity as serial dependence in asset returns tends to zero as per Bessembinder (2003); therefore, it measures how closely the market adheres to a random walk.[11]
- *Reversal*: returns on a portfolio that is long losers and short winners over the past month.
- *Momentum*: returns on a portfolio that is long winners and short losers over the past twelve months, skipping the first month (i.e., months *m*-12 up to and including *m*-2, where *m* is the current month);

[11] Variance ratios are computed from equal-weighted mid-quote returns and do not utilize traded prices, mitigating the problem of non-synchronous trading. We use variance ratios at the aggregate market level, because our exploratory analyses and those of Andersen, Bollerslev, and Das (2001) suggest that return outliers at high frequencies render intraday variances unreliable at the individual stock level. Our use of 30-minute-to-daily variance ratios is similar to the hourly-daily measure computed by Bessembinder (2003) and Chordia, Roll, and Subrahmanyam (2011); and the use of hourly-daily variance ratios does not substantially alter our conclusions.

The monthly time-series of returns on the reversal and momentum factors are computed as the monthly average of the daily returns on those factors as obtained from Ken French's website (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/).[12] We aggregate the three stock-specific efficiency measures (*OIB predictability*, *autocorrelation*, and *put-call parity*) to monthly, market-level measures by first averaging across stocks each day, and then averaging across days within the month.

Table 5 presents the correlation matrix between the six different monthly market-wide efficiency measures. All six measures are inverse indicators of the degree of market efficiency. We present the Pearson and Spearman correlations in Panels A and B, respectively. Of the 15 correlations in Panel A, 13 are positive, and seven are significant at the 10% level or less. In Panel B, 14 of 15 correlations are positive, with nine being significant at the 10% level or less. This is indication of non-trivial systematic variation in aggregate market efficiency across various efficiency measures. We note that the *put-call parity* and *reversal* measures (and to a lesser extent the *variance ratio* measure) are positively and significantly related to the microstructural efficiency measures (i.e., *OIB predictability* and *autocorrelation*); however, the correlations between *momentum* and these measures, while positive, are much smaller. In addition, the correlations of *momentum* with *put-call parity* and *variance ratio* are indistinguishable from zero (and positive in only one of four cases corresponding to the two methods of computing correlations). *Momentum* is, however, positively correlated with *reversal*,

---

[12] As per the website, the *momentum* factor is formed as follows: "We use six value-weight portfolios formed on size and prior (2-12) returns to construct Mom. The portfolios, which are formed daily, are the intersections of 2 portfolios formed on size (market equity, ME) and 3 portfolios formed on prior (2-12) return. The daily size breakpoint is the median NYSE market equity. The daily prior (2-12) return breakpoints are the 30th and 70th NYSE percentiles. [The *momentum* factor] is the average return on the two high prior return portfolios minus the average return on the two low prior return portfolios." The *reversal* factor is formed similarly, except that the prior month's return replaces the prior 2-12 months' return.

and significantly so in Panel A. The correlation between the two microstructural measures is close to unity, and highly significant in both Panels A and B.

Given the evidence of significant common variation in aggregate market efficiency across different measures in Table 5, we now seek to extract a systematic efficiency factor via principal component analysis (PCA). We follow Hasbrouck and Seppi (2000) and extract the principal components using the correlation matrix of Panel A in Table 5. We omit the autocorrelation-based microstructural measure from the PCA because, as observed above, it is a near-perfect substitute for the OIB-based measure. We find that the first principal component explains 32.9% of the total variation in the five monthly time-series of market-wide efficiency measures. Moreover, the loadings of the five different efficiency measures on the first principal component all have the same sign, with the greatest loadings for *OIB predictability*, *put-call parity*, and *reversal*. The second to fifth components account for 23%, 20%, 17%, and 8% of the variance, respectively. Overall, this evidence points at the existence of strong common factors in both microstructural and longer-horizon measures of market efficiency.

Given that the first component explains almost a full one-third of the total variation, explains ten percent more variation than the next component, and exhibits same-sign factor loadings across all five efficiency measures, we use this component as representative of systematic variation in aggregate market efficiency. To get a time-series of the first principal component, we standardize each efficiency measure to have zero mean and unit standard deviation, and multiply the matrix of standardized efficiency measures by the vector of the loadings of each measure on the component.

4.2 Determinants of systematic variation in aggregate market efficiency

We now turn to an analysis of the drivers of systematic variation in aggregate market efficiency, defined by the first principal component. First, we propose that the systematic component of aggregate efficiency can be related to variables that measure funding liquidity, such as hedge fund flows and short rates. To the extent that fluctuations in the funding liquidity of the financial system have pervasive effects on arbitrage activity (e.g., Akbas, Armstrong, Sorescu, and Subrahmanyam, 2010), the systematic component of efficiency can be affected by changes in funding liquidity. Second, arbitrage activity is hampered more generally by frictions such as idiosyncratic volatility and bid-ask spreads. Variation in these frictions can result in systematic variation in efficiency. Third, common variation in efficiency can be induced by microstructural aspects that affect market making efficacy (such as market-wide trading activity) and/or by irrational investor behavior.

We collect data on three different proxies for funding liquidity. *Hedge fund flow* is the monthly percentage in money inflow into hedge funds.[14] Greater hedge fund inflows should spur arbitrage activity. We note that hedge fund flow data are available for only part of the sample period. *TED spread* is the difference between the three-month LIBOR and the three-month T-bill rate from the FRED database of the Federal Reserve Bank of St. Louis and is a widely used indicator of funding liquidity (Brunnermeier, Nagel, and Pedersen, 2008; Brunnermeier, 2009; Brennan, Chordia, Subrahmanyam, and Tong, 2012).[15] *Bank returns* is the monthly total return on the Dow Jones U.S. financial industry index taken from TRTH. Following Hameed, Kang, and Viswanathan (2010) and Karolyi, Lee, and van Dijk (2012), we argue that a rise in the

[14] We thank Matti Suominen and LIPPER-TASS for data on hedge fund flows. The sample includes all hedge funds that report their returns in U.S. dollars and have a minimum of 36 monthly return observations over our sample period. The sample period spans 1996 to 2008.

[15] The notion is that the TED spread may proxy for counterparty risk, which, when elevated, can lead to funding illiquidity.

market value of the financial industry is likely to be associated with a stronger aggregate balance sheet of the funding sector.

We include two proxies for market frictions that can hamper arbitrage. *5-minute volatility* is the average across stocks of the monthly return volatility estimated by stock based on 5-minute intervals within the month. Prior research (e.g., de Jong, Rosenthal, and van Dijk, 2009) suggests that idiosyncratic volatility impedes arbitrage. *PESPR* is the average across stocks of the monthly time-series average of the daily proportional effective bid-ask spread of each stock, computed from TRTH data.

Regarding the third category of determinants (sources of mispricing), we argue that fundamental uncertainty can make deviations from efficient pricing more likely across the board and we use the forward-looking *VIX* volatility index obtained from TRTH as a proxy for market-wide uncertainty. *#trades* is the total number of trades per month across all the stocks in our sample. We include this variable to account for the marked increase in turnover and decrease in average trade size over our sample period which has been related to the advent of algorithmic trading (Hendershott, Jones, and Menkveld, 2011; Chordia, Roll, and Subrahmanyam, 2011). Finally, Baker and Wurgler (2006) show that waves of investor sentiment affect many stocks at the same time. To test the hypothesis that fluctuations in market-wide investor sentiment affect systematic variation in efficiency, we include their U.S. investor sentiment index (*sentiment*).

Table 6 presents summary statistics of the five (non-standardized) efficiency measures (Panel A) and of the variables in each of the three categories of potential determinants of the systematic component of aggregate efficiency (Panel B). The mean R^2 of the *OIB predictability* regressions is around 2% (like in Table 2) and the mean absolute *put-call parity* deviation is about 1.2%. The mean absolute deviation of the *variance ratio* from unity is 0.24, which is broadly in line with the findings of Chordia, Roll, and Subrahmanyam (2011). The mean daily rewards for *reversal* and *momentum* are 8.6 and 2.6 basis points, respectively.

We are interested in what factors explain time-variation in the systematic component of efficiency. We run simple time-series regressions of the first principal component on the contemporaneous determinants in the three categories. Prior to usage as a dependent variable, we detrend the efficiency measure with linear and quadratic trend terms.[16] Table 7 presents the estimated coefficients and their associated *t*-statistics in each of these six regressions. The results indicate that the systematic component of efficiency depends on all three categories of determinants. The regression reported in the first column includes only the funding liquidity measures. We find that the coefficient of *hedge fund flow* is negative and significant at the 1% level, whereas the coefficients on *TED spread* and *bank returns* are positive and negative, respectively, and significant at the 10% and 1% levels. Thus, an increase in hedge fund inflows and returns to the banking sector or a decrease in the TED spread all improve the systematic component of market efficiency, which is consistent with intuition. The economic magnitudes of these effects are considerable. A one standard deviation increase in *hedge fund flow* is associated with a 0.27 standard deviation decrease in the systematic component of efficiency (which is an inverse measure of aggregate market efficiency).[17] The corresponding economic effects of *TED spread* and *bank returns* are equal to 0.24 and 0.26, respectively.

The regression reported in the second column of Table 7 includes the microstructural frictions variables. The coefficient on *5-minute volatility* is strongly positive; a one standard

[16] Detrending addresses the possible concern that trends in dependent and explanatory variables could lead to spurious conclusions, and follows Roll, Schwartz, and Subrahmanyam (2007). The results below are largely unchanged whether we detrend the principal component, or whether we detrend the individual efficiency measures and then extract the principal component.

[17] This economic magnitude is computed by multiplying the coefficient on *hedge fund flow* of -0.103 in the first column of Table 7 by the time-series standard deviation of this variable of 1.533 (from Table 6) and then dividing by the time-series standard deviation of the detrended first principal component of the five aggregate market efficiency measures, which is equal to 0.584 (not tabulated).

deviation decrease in *5-minute volatility* is associated with a 0.77 standard deviation improvement in systematic market efficiency. Idiosyncratic volatility can deter arbitrage in individual stocks, and our result suggests that a reduction in the average total volatility across stocks is associated with increased efficiency. The coefficient on PESPR is not significant, somewhat surprisingly. We revisit this issue below.

We present coefficients for the last category of determinants (sources of mispricing) in the third column of Table 7. The coefficient on *VIX* is positive and significant, indicating that increased ex ante uncertainty leads to decreased efficiency, as is intuitive. A one standard deviation reduction in *VIX* is associated with a 0.70 standard deviation improvement in market efficiency. However, *#trades* or *sentiment* are not related to common variation in efficiency in this specification.

To isolate the role of the bid-ask spread, in the fourth column we include *PESPR* as the sole variable in the regression. We find that the coefficient on this variable is positive and significant at the 10% level, which is in line with the view that illiquidity hampers arbitrage and thus harms efficiency. Combining the evidence in the second and fourth columns, this suggests that short-horizon volatility supplants liquidity as a determinant of systematic variation in market efficiency.

The fifth column of Table 7 includes all of the determinants of efficiency. The coefficients on the funding liquidity measures continue to have the expected sign and preserve their statistical and economic significance, as does the coefficient on *5-minute volatility*. In the presence of this variable, however, the coefficient on *VIX* loses significance. Further, the coefficient on *#trades* is negative and significant in the presence of the other variables. The negative sign implies that, controlling for other determinants of efficiency, an increase in the number of trades implies greater market efficiency, which is consistent with the notion that increased market activity leads to greater camouflage for arbitrageurs and increases their efficacy. Further, an increase in the number of trades could also be associated with increased algorithmic trading in later years of our sample (Hendershott, Jones, and Menkveld, 2011), and under the presumption that some of this trading is arbitrage-driven, it would enhance market efficiency. The economic significance of the effect of *#trades* is substantial. A one standard deviation increase in this variable is associated with a 0.35 standard deviation improvement in market efficiency.

We note that the coefficient on the proportional effective spread is negative and significant in the fifth column and thus switches sign relative to the fourth column. To ascertain whether the sign flip is due to the collinearity of *PESPR* with *#trades* (we expect these variables to be intimately related), we use the *#trades* measure orthogonalized with respect to *PESPR* as a separate variable in the sixth and last column of Table 7 (*#trades* $\perp$ *PESPR*). We find that the coefficient on this orthogonalized measure remains negative and significant while the coefficient on the spread is no longer significant, as in the second column. Thus, the negative sign of the *PESPR* coefficient appears to arise from its interaction with *#trades*, and while the spread has a positive and marginally significant effect as a standalone determinant of common variation in efficiency, it is supplanted by the other variables in the regression.

Taken together, the results in Table 7 indicate that the systematic component of aggregate market efficiency varies over time with aggregate funding liquidity, frictions that impede arbitrage, and microstructural aspects that affect market making efficacy. As indicated by the raw and adjusted R^2s reported in Table 7, the determinants in these three categories jointly explain about two-thirds of the time-variation in the systematic component of market efficiency.

5. Conclusions

We examine variation over time and across stocks in various short-horizon market efficiency metrics (specifically, intraday return predictability from order flow or returns and daily put-call parity deviations) for a substantial sample of NYSE stocks over a period of fifteen years. We show that each of the efficiency metrics demonstrates considerable time-variation and also exhibits a strong systematic component across stocks. We correlate the market-wide equivalents of these short-horizon efficiency metrics with market-wide variance ratios and with the well-known reversal and momentum anomalies, and find that almost all of these short- and long-horizon market efficiency measures are positively cross-correlated, with the majority of correlations being statistically significant. This finding links microstructural measures of market quality to measures of the law of one price, and to broader, longer-horizon measures of market efficiency. We then consider the pervasive component of efficiency via the first principal component across all our aggregate market efficiency metrics. We show that time-variation in the common efficiency component depends on funding liquidity, frictions that impede arbitrage, and traditional variables that measure the efficacy of market making such as trading activity.

Recognizing that market efficiency has a common component opens new vistas for research. First, it would be worth exploring whether there is a global component to variation in market efficiency. This would allow us to ascertain whether global financial crises might lead to systematic deterioration in the quality of price formation in markets across the world. Second, it would be worth investigating whether commonality in market efficiency extends to other asset classes such as fixed income securities, foreign exchange, and derivatives. Third, we have not attempted to explore asset pricing implications of our results given the relatively short sample period of fifteen years. Nevertheless, it seems reasonable to conjecture that if agents prefer efficient markets and dislike unexpected deteriorations in market efficiency (due to, for example, unexpected shocks in arbitrage funding constraints), then both the level and the volatility of efficiency may be priced in the cross-section of asset returns. These and other issues are left for future research.

References

Akbas, F., W. Armstrong, S. Sorescu, and A. Subrahmanyam, 2010, Funding constraints and market efficiency, working paper, Texas A&M University.

Andersen, T., T. Bollerslev, and A. Das, 2001, Variance-ratio statistics and high-frequency data: Testing for changes in intraday volatility patterns, *Journal of Finance* 56, 305-327.

Angel, J, L. Harris, and C. Spatt, 2011, Equity trading in the 21st century, *Quarterly Journal of Finance* 1, 1-53.

Baker, M., and J. Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645-1680.

Benston, G., and R. Hagerman, 1974, Determinants of bid-asked spreads in the over-the-counter market, *Journal of Financial Economics* 1, 353-364.

Bessembinder, H., 2003, Trade execution costs and market quality after decimalization, *Journal of Financial and Quantitative Analysis* 38, 747-777.

Boehmer, E., and E. Kelley, 2009, Institutional investors and the informational efficiency of prices, *Review of Financial Studies* 22, 3563-3594.

Boehmer, E., and J. Wu, 2007, Order flow and prices, working paper, University of Georgia.

Brennan, M, T. Chordia, A. Subrahmanyam, and Q. Tong, 2012, Sell-order liquidity and the cross-section of expected stock returns, *Journal of Financial Economics*, forthcoming.

Brennan, M., and A. Subrahmanyam, 1995, Investment analysis and price formation in securities markets, *Journal of Financial Economics* 38, 361-382.

Brunnermeier, M., 2009, Deciphering the liquidity and credit crunch 2007-2008, *Journal of Economic Perspectives* 23, 77-100.

Brunnermeier, M., S. Nagel, and L. Pedersen, 2008, Carry trades and currency crashes, working paper, Princeton University.

Brunnermeier, M., and L. Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201-2238.

Busse, J., and T. Green, 2002, Market efficiency in real time, *Journal of Financial Economics* 65, 415-437.

Chan, K., and W. Fong, 2000, Trade size, order imbalance, and the volatility-volume relation, *Journal of Financial Economics* 57, 247-273.

Chordia, T., R. Roll, and A. Subrahmanyam, 2000, Commonality in liquidity, *Journal of Financial Economics* 56, 3-28.

Chordia, T., R. Roll, and A. Subrahmanyam, 2002, Order imbalance, liquidity, and market returns, *Journal of Financial Economics* 65, 111-130.

Chordia, T., R. Roll, and A. Subrahmanyam, 2005, Evidence on the speed of convergence to market efficiency, *Journal of Financial Economics* 76, 271-292.

Chordia, T., R. Roll, and A. Subrahmanyam, 2008, Liquidity and market efficiency, *Journal of Financial Economics* 87, 249-268.

Chordia, T., R. Roll, and A. Subrahmanyam, 2011, Recent trends in trading activity, *Journal of Financial Economics* 101, 243-263.

Chordia, T., and A. Subrahmanyam, 2004, Order imbalance and individual stock returns: Theory and evidence, *Journal of Financial Economics* 72, 485-518.

Cooper, M., 1999, Filter rules based on price and volume in individual security overreaction, *Review of Financial Studies* 12, 901-935.

Cremers, M., and D. Weinbaum, 2010, Deviations from put-call parity and stock return predictability, *Journal of Financial and Quantitative Analysis* 45, 335-367.

D'Avolio, G, 2002, The market for borrowing stock, *Journal of Financial Economics* 66, 271-306.

De Jong, A., L. Rosenthal, and M. van Dijk, 2009, The risk and return of arbitrage in dual-listed companies, *Review of Finance* 13, 495-520.

DeBondt, W., and R. Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 43, 793-808.

Epps, T., 1979, Comovements in stock prices in the very short run, *Journal of the American Statistical Association* 74, 291-298.

Fama, E., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–417.

Fama, E., 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics*, 49, 283-306.

Fama, E., and K. French, 1996, Multifactor explanations of asset pricing anomalies, *Journal of Finance* 51, 55-84.

Fama, E., and K. French, 2008, Dissecting anomalies, *Journal of Finance* 63, 1653-1678.

Finucane, T., 1991, Put-call parity and expected returns, *Journal of Financial and Quantitative Analysis* 26, 445-457.

Glosten, L., and L. Harris, 1988, Estimating the components of the bid/ask spread, *Journal of Financial Economics* 21, 123-142.

Griffin, J., P. Kelley, and F. Nardari, 2010, Do market efficiency measures yield correct inferences? A comparison of developed and emerging markets, *Review of Financial Studies* 23, 3225-3277.

Hameed, A., W. Kang, and S. Viswanathan, 2010, Stock market declines and liquidity, *Journal of Finance* 65, 257-293.

Hasbrouck, J., and D. Seppi, 2000, Common factors in prices, order flows and liquidity, *Journal of Financial Economics* 59, 383-411.

Hendershott, T., C. Jones, and A. Menkveld, 2011, Does algorithmic trading improve liquidity?, *Journal of Finance* 66, 1-33.

Jegadeesh, N., 1990, Evidence of predictable behavior in security returns, *Journal of Finance* 45, 881-898.

Jegadeesh, N., and S. Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *Journal of Finance* 48, 65-92.

Karolyi, G., K.-H. Lee, and M. van Dijk, 2012, Understanding commonality in liquidity around the world, *Journal of Financial Economics* 105, 82-112.

Khandani, A., and A. Lo, 2011, What happened to the quants in August 2007? Evidence from factors and transactions data, *Journal of Financial Markets* 14, 1-46.

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

Lee, C., and B. Radhakrishna, 2000, Inferring investor behavior: Evidence from TORQ data, *Journal of Financial Markets* 3, 83-111.

Lee, C., and M. Ready, 1991, Inferring trade direction from intra-day data, *Journal of Finance* 46, 733-747.

Madhavan, A., M. Richardson, and M. Roomans, 2002, Why do security prices change? A transaction-level analysis of NYSE stocks, *Review of Financial Studies* 10, 1035-1064.

Mase, B., 1999, The predictability of short-horizon stock returns, *European Finance Review* 3, 161-173.

McLean, D., and J. Pontiff, 2012, Does academic research destroy stock return predictability?, working paper, Boston College.

Mitchell, M., and T. Pulvino, 2012, Arbitrage crashes and the speed of capital, *Journal of Financial Economics* 104, 469-490.

Mitchell, M., T. Pulvino, and E. Stafford, 2002, Limited arbitrage in equity markets, *Journal of Finance* 57, 551-584.

Nagel, S., 2005, Short-sales, institutional investors, and the cross-section of stock returns, *Journal of Financial Economics* 78, 277-309.

Odders-White, E., 2000, On the occurrence and consequences of inaccurate trade classification, *Journal of Financial Markets* 3, 205-332.

Pasquariello, P., 2012, Financial market dislocations, working paper, University of Michigan.

Roll, R., E. Schwartz, and A. Subrahmanyam, 2007, Liquidity and the law of one price: The case of the futures/cash basis, *Journal of Finance* 62, 2201-2234.

Shleifer, A., and R. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 52, 35-55.

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

Table 1: Cross-sectional summary statistics of time-series averages, 1996-2010

This table reports the cross-sectional (across the 156 S&P100 stocks in the sample) mean, standard deviation, as well as the 25th, 50th (median), and 75th percentile of the time-series average by stock of the daily number of trades *(#trades*), trading volume in US$b. (*dollar volume*), one-minute mid-quote returns in basis points (*1-min mid-quote return*), percentage annualized volatility of one-minute mid-quote returns (*1-min return volatility*), percentage annualized volatility of 15-minute mid-quote returns (*15-min return volatility*), difference between the total amount of trades initiated by buyers and sellers over one-minute intervals (*1-min OIB#*), difference between the total dollar amount of trades initiated by buyers and sellers over one-minute intervals (*1-min OIB$*), proportional effective spread (*PESPR*), the number of days with fewer than 20 one-minute intervals during the day with a valid return observation for that interval and a valid OIB observation in the previous interval (*filters*), as well as the number of days *(# days in sample*) and the number of months *(# months in sample*) the stock is in the sample. The first column of the table indicates the number of stocks over which the summary statistics are computed. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute all variables in the table are from TRTH.

	# stocks	mean	st.dev.	25th	median	75th
#trades	156	2,958	1,918	1,527	2,880	4,014
dollar volume	156	0.097	0.081	0.039	0.073	0.130
1-min mid-quote return	156	0.01	0.04	0	0.01	0.03
1-min return volatility	156	28.70	7.64	23.78	26.57	32.18
15-min return volatility	156	26.62	6.36	22.05	25.41	29.61
1-min OIB#	156	0.185	0.149	0.089	0.170	0.251
1-min OIB$	156	18,840	15,701	8,428	14,995	26,708
PESPR (%)	156	0.13	0.12	0.08	0.09	0.14
filters	156	61	179	0	3	45
# days in sample	156	2,972	1,012	2,585	3,563	3,667
# months in sample	156	144	50	119	170	179

Table 2: Predictive regressions of one-minute returns on lagged order imbalance or on lagged returns, 1996-2010

This table reports the average of the efficiency regression results from equation (1) estimated by stock-day for all 156 S&P100 stocks in the sample. Each of the first four columns in the table presents the results of a different way to estimate the efficiency regression. For *OIB predictability*, the dependent variable *return* is computed as the return of the mid-quote associated with the last trade to the mid-quote of the first trade in each one-minute interval and the independent variable *OIB* is the difference between the total dollar amount of trades initiated by buyers and sellers in the previous one-minute interval. The other methods are named after the single feature that distinguishes them from the *OIB predictability* regressions. The *allquotes* regressions are based on returns computed using any quotes within each interval rather than only using quotes associated with trades; the *2minutes* regressions are based on two-minute instead of one-minute intervals; and the *oib#* regressions are based on order imbalance expressed in number of trades rather than dollars. The final column (*autocorrelation*) presents the results of similar regressions estimated by stock-day in which the independent variable is $return_{t-1}$ instead of OIB_{t-1}. The first number in each column is the average slope coefficient in the efficiency regressions. The average ("*t*-stat avg") and cross-sectional *t*-statistics ("*t*-stat cross") are in parentheses below the coefficients. "% positive" is the percentage of positive coefficients, and "% positive significant" is the percentage of coefficients with *t*-statistics greater than 1.645 (the 5% critical level in a one-tailed test). Intercepts have been suppressed to conserve space. For readability, the *OIB* coefficient has been scaled by 10^9 for the *OIB predictability*, *allquotes*, and *2minutes* regressions and by 10^4 for the *oib#* regressions. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute returns and order imbalances are from TRTH.

Dependent variable: $return_t$					
Efficiency measure:	*OIB predictability*	*allquotes*	*2minutes*	*oib#*	*autocorrelation*
OIB_{t-1}	0.997	0.511	0.849	0.664	
$return_{t-1}$					0.025
t-stat avg	(1.510)	(0.821)	(0.850)	(2.189)	(0.412)
t-stat cross	(48.375)	(26.852)	(37.848)	(423.985)	(151.250)
% positive	83.25	72.57	72.32	88.88	58.80
% positive significant	53.11	34.52	36.12	67.36	30.89
R^2	2.24	1.19	2.02	3.37	1.37
adj. R^2	1.68	0.62	1.29	2.81	0.70
# regressions	463,469	463,476	466,634	463,191	451,958

Table 3: Monthly regressions of daily individual stock efficiency on contemporaneous, lagged, and lead market-wide efficiency, 1996-2010

This table reports the average of the monthly commonality regression results from equation (2) estimated for all 156 S&P100 stocks in the sample. The dependent variable $\Delta Eff_{i,d}$ is the change in the efficiency of stock *i* on day *d*. The independent variable $\Delta MktEff_d$ is the change in market-wide efficiency on day *d*, computed as the equally-weighted average efficiency of all individual stocks on day *d*, excluding stock *i*. Each commonality regression also includes a one-day lead and lag of changes in market-wide efficiency. Each of the three columns in the table presents the results of the commonality regressions based on a different efficiency measure. The *OIB predictability* efficiency measure is equal to the R^2 of the daily regressions of the mid-quote return in each one-minute interval on order imbalance expressed in dollars in the previous one-minute interval (from Table 2). The *autocorrelation* measure is the R^2 of the daily regressions of the mid-quote return in each one-minute interval on the mid-quote return in the previous one-minute interval (from Table 2). The *put-call parity* measure is the absolute difference between the implied volatilities of a call and a put option of the same series (i.e., pairs of options on the same underlying stock with the same strike price and the same expiration date), averaged across all option pairs for each stock (based on end-of-day quotes from all option series with positive implied volatilities, expiring in two weeks to one year, and with a call delta between 0.3 and 0.7). Each column presents the average slope coefficients in the commonality regressions. The average ("*t*-stat avg") and cross-sectional *t*-statistics ("*t*-stat cross") are in parentheses below the coefficients. "% positive" is the percentage of positive coefficients, and "% positive significant" is the percentage of coefficients with *t*-statistics greater than 1.645 (the 5% critical level in a one-tailed test). Intercepts have been suppressed to conserve space. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH and OptionMetrics.

Dependent variable: $\Delta Eff_{i,d}$			
Efficiency measure:	*OIB predictability*	*autocorrelation*	*put-call parity*
$\Delta MktEff_d$	0.389	0.509	0.556
t-stat avg	(0.219)	(0.329)	(0.543)
t-stat cross	(20.657)	(30.215)	(7.429)
% positive	57.45	60.42	64.05
% positive significant	7.94	10.37	17.33
$\Delta MktEff_{d-1}$	0.025	0.019	0.067
t-stat avg	(0.017)	(0.025)	(0.044)
t-stat cross	(1.359)	(1.148)	(1.555)
% positive	49.92	49.61	49.94
% positive significant	7.04	7.29	8.59
$\Delta MktEff_{d+1}$	0.025	0.019	0.141
t-stat avg	(0.024)	(0.017)	(0.024)
t-stat cross	(1.303)	(1.159)	(1.995)
% positive	50.54	49.76	49.58
% positive significant	7.30	7.16	8.06
R^2	19.32	19.95	33.64
adj. R^2	4.38	5.09	25.96
# regressions	22,430	21,660	18,681

Table 4: Check for cross-equation dependence in estimation error, 1996-2010

This table assesses to what extent the reported cross-sectional *t*-statistics in Table 3 need to be corrected for cross-equation dependence in residuals. Estimating equation (2) monthly for each individual stock yields daily residuals for each stock for each month. The table reports the average pairwise correlation between these residuals for every stock for each year in the sample for each of the three stock-level efficiency measures in Table 3 (*OIB predictability*, *autocorrelation*, and *put-call parity*) as well as the associated average *p*-values in parentheses. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH.

Efficiency measure:	*OIB predictability*	*autocorrelation*	*put-call parity*
1996	-0.001	-0.001	0.003
	(0.430)	(0.456)	(0.428)
1997	-0.001	-0.004	0.002
	(0.412)	(0.427)	(0.432)
1998	0.000	-0.003	-0.002
	(0.413)	(0.414)	(0.418)
1999	0.001	-0.002	-0.001
	(0.412)	(0.415)	(0.435)
2000	0.000	-0.002	-0.002
	(0.411)	(0.420)	(0.402)
2001	-0.001	-0.003	-0.003
	(0.405)	(0.416)	(0.414)
2002	-0.003	-0.006	0.001
	(0.416)	(0.408)	(0.404)
2003	-0.003	-0.002	-0.001
	(0.409)	(0.412)	(0.409)
2004	-0.002	-0.002	0.000
	(0.423)	(0.415)	(0.409)
2005	-0.002	-0.003	-0.004
	(0.412)	(0.417)	(0.409)
2006	-0.002	-0.002	-0.002
	(0.411)	(0.411)	(0.414)
2007	-0.002	-0.002	-0.005
	(0.396)	(0.409)	(0.330)
2008	-0.001	-0.003	0.001
	(0.405)	(0.393)	(0.231)
2009	-0.002	-0.004	0.014
	(0.417)	(0.397)	(0.351)
2010	-0.003	-0.005	0.002
	(0.416)	(0.401)	(0.354)

Table 5: Correlations between aggregate market efficiency measures, 1996-2010

This table reports correlation coefficients between six different monthly aggregate market efficiency measures. Three of the efficiency measures (*OIB predictability*, *autocorrelation*, and *put-call parity*) are aggregated from daily stock-level efficiency measures by first averaging across stocks each day, and then averaging across days within the month. *Variance ratio* is the monthly absolute difference between one and the scaled ratio of the 30-minute mid-quote return variance to the open-to-close mid-quote return variance. *Reversal* and *momentum* are the average daily returns each month on portfolios that are long losers and short winners over the past month and long winners and short losers over the past twelve months skipping the first month, respectively. All six measures are inverse indicators of the degree of market efficiency. Panels A and B present Pearson and Spearman correlations, respectively. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH, OptionMetrics, and the website of Ken French. *p*-values are in parentheses. Significance at the 1%, 5%, and 10% level is indicated by ***, **, and *, respectively.

	OIB predictability	*auto-correlation*	*put-call parity*	*variance ratio*	*reversal*	*momentum*
Panel A: Pearson correlations						
OIB predictability	1					
autocorrelation	0.915***	1				
	(0.00)					
put-call parity	0.591***	0.454***	1			
	(0.00)	(0.00)				
variance ratio	0.112	0.145*	0.035	1		
	(0.14)	(0.05)	(0.64)			
reversal	0.124*	0.125*	0.068	0.002	1	
	(0.10)	(0.09)	(0.36)	(0.98)		
momentum	0.087	0.104	-0.05	-0.015	0.154**	1
	(0.24)	(0.17)	(0.51)	(0.84)	(0.04)	
Panel B: Spearman correlations						
OIB predictability	1					
autocorrelation	0.933***	1				
	(0.00)					
put-call parity	0.566***	0.490***	1			
	(0.00)	(0.00)				
variance ratio	0.097	0.139*	0.125*	1		
	(0.20)	(0.06)	(0.09)			
reversal	0.226***	0.236***	0.220***	0.133*	1	
	(0.00)	(0.00)	(0.00)	(0.08)		
momentum	0.091	0.063	0.017	-0.061	0.016	1
	(0.22)	(0.40)	(0.82)	(0.42)	(0.83)	

Table 6: Summary statistics of aggregate market efficiency and its potential determinants, 1996-2010

This table reports the time-series mean, standard deviation, as well as the 25th, 50th (median), and 75th percentile of the five monthly aggregate market efficiency measures included in the principal component analysis (Panel A) and three categories of potential determinants of time-series variation in aggregate market efficiency (Panel B). Three of the efficiency measures (*OIB predictability*, *autocorrelation*, and *put-call parity*) are aggregated from daily stock-level efficiency measures by first averaging across stocks each day, and then averaging across days within the month. *Variance ratio* is the monthly absolute difference between one and the scaled ratio of the 30-minute mid-quote return variance to the open-to-close mid-quote return variances. *Reversal* and *momentum* are the average daily returns each month on portfolios that are long losers and short winners over the past month and long winners and short losers over the past twelve months skipping the first month, respectively. All five measures are inverse indicators of the degree of market efficiency. Panel B distinguishes between three categories of potential determinants. (i) Funding liquidity: percentage money inflows into hedge funds (*hedge fund flow*), *TED spread* (FRED ID: USD3MTD156N minus TB3MS), and monthly returns on the Dow Jones U.S. financial industry index (*bank returns*; RIC: .DJUSFN). (ii) Frictions impeding arbitrage: aggregate (i.e., equally-weighted cross-sectional average) volatility based on five-minute intervals (averaged over all observations within the month) (*5-min volatility*) and aggregate proportional effective spread (averaged over daily observations within the month) (*PESPR*). (iii) Sources of mispricing: VIX index (*VIX*), aggregate number of trades in millions *(#trades*), and the sentiment index from Baker and Wurgler (2006) (*sentiment*).

	# obs.	mean	st.dev.	25th	median	75th
Panel A: Measures of aggregate market efficiency						
OIB predictability (% R^2)	180	2.27	1.84	0.80	1.13	3.43
put-call parity (%)	180	0.24	0.23	0.10	0.19	0.33
variance ratio	180	1.17	0.54	0.71	1.01	1.59
reversal (bp per day)	180	2.56	28.58	-7.38	3.25	16.66
momentum (bp per day)	180	8.58	24.74	-2.97	5.82	16.95
Panel B: Determinants of aggregate market efficiency						
(i) Funding liquidity						
hedge fund low (%)	168	0.781	1.533	0.113	1.006	1.671
TED spread (%)	180	0.574	0.440	0.242	0.482	0.722
bank returns (%)	179	0.265	5.797	-1.831	1.010	3.415
(ii) Frictions impeding arbitrage						
5-min volatility (% per annum)	180	19.32	7.58	13.78	17.37	22.17
PESPR (%)	180	0.112	0.075	0.044	0.094	0.170
(iii) Sources of mispricing (fundamental uncertainty, trading, sentiment)						
VIX	180	22.22	8.38	16.61	20.98	25.48
#trades (millions)	180	8.570	7.746	2.251	6.208	12.747
sentiment	180	0.212	0.608	-0.140	0.098	0.435

Table 7: Regressions to explain systematic variation in aggregate market efficiency, 1996-2010

This table presents the results of time-series regressions to explain systematic variation in aggregate market efficiency, defined as the first principal component of the five monthly aggregate market efficiency measures from Table 6. To get a time-series of the first principal component, we standardize each efficiency measure to have zero mean and unit standard deviation, and multiply the matrix of standardized efficiency measures by the vector of the loadings of each measure on the component. We then detrend the first principal component with linear and quadratic trend terms. The resulting dependent variable in the regressions is an inverse indicator of the degree of market efficiency. The (contemporaneous) independent variables are described in Table 6 and are grouped into three categories: (i) funding liquidity, (ii) frictions impeding arbitrage, and (iii) sources of mispricing (*#trades* ⊥ *PESPR* is the *#trades* variable orthogonalized with respect to the *PESPR* variable). The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Intercepts have been suppressed to conserve space; *t*-statistics are in parentheses. Significance at the 1%, 5%, and 10% level is indicated by ***, **, and *, respectively.

Dependent variable: *detrended first principal component of five monthly aggregate market efficiency measures*

	(1)	(2)	(3)	(4)	(5)	(6)
(i) Funding liquidity						
hedge fund flow	-0.103***				-0.051**	-0.051**
	(-3.16)				(-2.32)	(-2.32)
TED spread	0.324*				0.336***	0.336***
	(1.70)				(3.15)	(3.15)
bank returns	-0.026***				-0.012*	-0.012*
	(-3.63)				(-1.79)	(-1.79)
(ii) Frictions impeding arbitrage						
5-min volatility		0.059***			0.046***	0.046***
		(8.04)			(3.81)	(3.81)
PESPR		0.253		2.071*	-1.324*	0.408
		(0.48)		(1.79)	(-1.97)	(0.80)
(iii) Sources of mispricing (fundamental uncertainty, trading, sentiment)						
VIX			0.049***		-0.003	-0.003
			(7.01)		(-0.36)	(-0.36)
#trades			0.002		-0.026***	
			(0.20)		(-3.73)	
#trades ⊥ *PESPR*						-0.026***
						(-3.73)
sentiment			0.106		0.068	0.068
			(1.38)		(1.55)	(1.55)
R^2	40.64	59.61	51.20	7.16	66.03	66.03
adj. R^2	39.46	59.15	50.37	6.64	64.17	64.17
# obs.	155	180	180	180	155	155

Figure 1: Time-variation in microstructural market efficiency measures, 1996-2010

This figure shows the monthly time-variation in five aggregate market efficiency measures estimated from microstructure data. Each of the five microstructural efficiency measures is estimated at the individual stock-level as the R^2 of daily regressions predicting short-interval mid-quote returns from order flow or returns, and is then aggregated to a monthly, aggregate market efficiency measure by first averaging across stocks each day, and then averaging across days within the month. A high R^2 indicates low efficiency. For *OIB predictability*, the dependent variable is computed as the return of the mid-quote associated with the last trade to the mid-quote of the first trade in each one-minute interval and the independent variable *OIB* is the difference between the total dollar amount of trades initiated by buyers and sellers in the previous one-minute interval. The *allquotes* regressions are based on returns computed using any quotes within each interval rather than only using quotes associated with trades; the *2minutes* regressions are based on two-minute instead of one-minute intervals; and the *oib#* regressions are based on order imbalance expressed in number of trades rather than dollars. The *autocorrelation* measure is based on similar regressions estimated by stock-day in which the independent variable is the return (instead of the OIB) in the previous one-minute interval. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH.

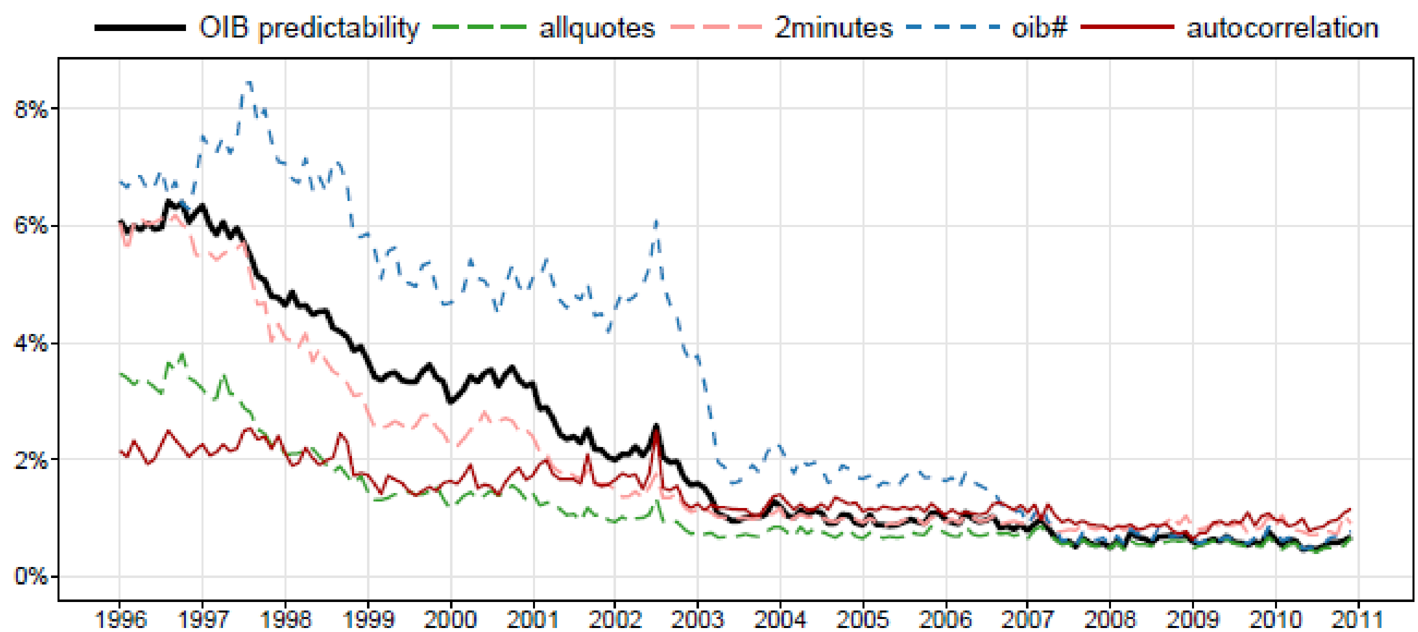

Figure 2: Cross-sectional variation in microstructural efficiency across stocks, yearly boxplots, 1996-2010

This figure shows yearly boxplots of the *OIB predictability* efficiency measure, defined as the average R^2 for all 156 S&P100 stocks in the sample based on daily regressions of mid-quote returns on lagged order imbalance measured over one-minute intervals within the day, averaged by stock across days within the year. A high R^2 indicates low efficiency. The bottom, middle, and top line of the boxes reflect the 25th, 50th (median), and 75th percentile of the yearly average R^2s across the 156 stocks in our sample, respectively. The top (bottom) of the line above (below) the boxes represents the highest (lowest) individual stock R^2 below (above) the 75th (25th) percentile plus (minus) 1.5 times the interquartile range, called the end of the whiskers. Observations above the end of the top whisker and below the end of the bottom whisker are not shown. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH.

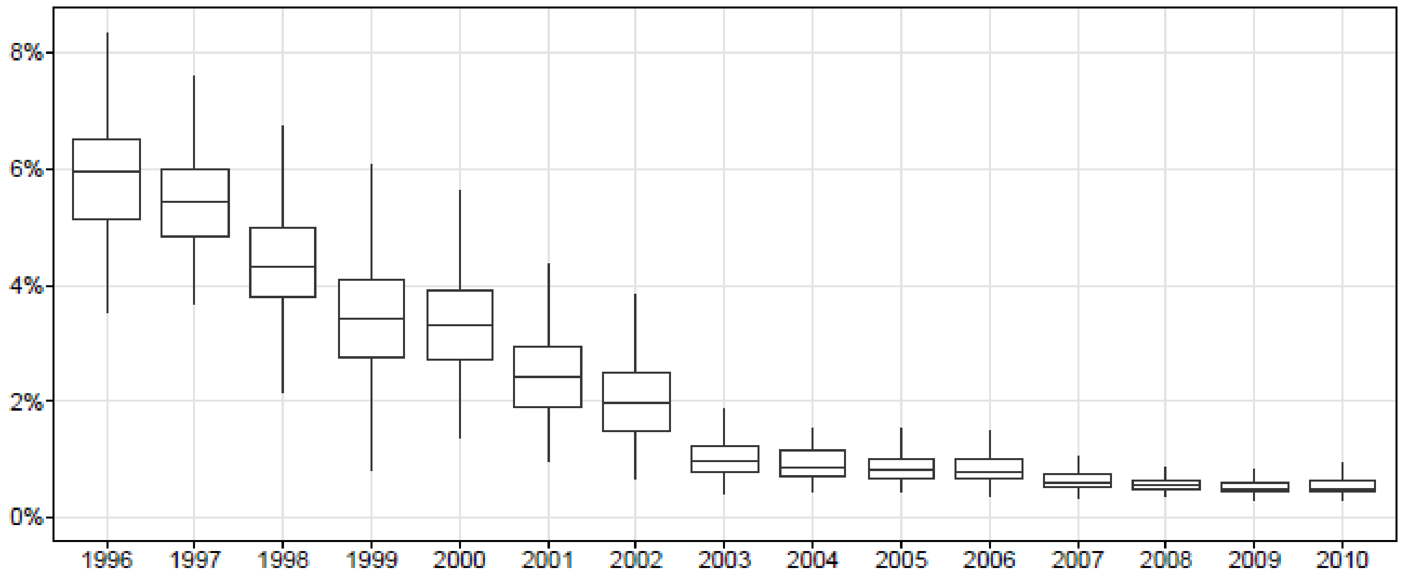

Figure 3: Time-variation in commonality in efficiency, 1996-2010

This figure shows the monthly time-variation in the degree of commonality in efficiency across stocks for three short-horizon efficiency measures. The *OIB predictability* efficiency measure is equal to the R^2 of the daily regressions of the mid-quote return in each one-minute interval on order imbalance expressed in dollars in the previous one-minute interval (from Table 2). The *autocorrelation* measure is the R^2 of the daily regressions of the mid-quote return in each one-minute interval on the mid-quote return in the previous one-minute interval (from Table 2). The *put-call parity* measure is the absolute difference between the implied volatilities of a call and a put option of the same series (i.e., pairs of options on the same underlying stock with the same strike price and the same expiration date), averaged across all option pairs for each stock (based on end-of-day quotes from all option series with positive implied volatilities, expiring in two weeks to one year, and with a call delta between 0.3 and 0.7). All three measures are inverse indicators of the degree of efficiency. For each efficiency measure, a monthly measure of commonality in efficiency across stocks is computed as the average R^2 across all stocks in the sample of monthly regressions of daily changes in individual stock efficiency on contemporaneous, lead, and lagged changes in market-wide efficiency. The sample includes all 156 NYSE-listed stocks that were part of the S&P100 at any time during 1996-2010. Data to compute the efficiency measures are from TRTH.