10-14-2016

Does Price Efficiency Increase with Trading Volume? Evidence of Nonlinearity and Power Laws in ETFs

Gunduz Caginalp
University of Pittsburgh

Mark DeSantis
Chapman University, desantis@chapman.edu

Recommended Citation

Caginalp, G., & DeSantis, M. (2016). Does price efficiency increase with trading volume? Evidence of nonlinearity and power laws in ETFs, *Physica A, 467*: 436-452. http://dx.doi.org/10.1016/j.physa.2016.10.039

Accepted Manuscript

Does price efficiency increase with trading volume? Evidence of nonlinearity and power laws in ETFs

Gunduz Caginalp, Mark DeSantis

PII: S0378-4371(16)30726-9
DOI: http://dx.doi.org/10.1016/j.physa.2016.10.039
Reference: PHYSA 17597

To appear in: *Physica A*

Received date: 4 May 2016
Revised date: 6 August 2016

***Highlights (for review)**

Highlights for the manuscript PHYSA-16760 entitled "Does price efficiency increase with trading volume? Evidence of nonlinearity and power laws in ETFs" are as follows:

1. The relationship between price efficiency and volume tests two competing theories.
2. Low volume levels obey a power law that supports efficient market theories.
3. Very high volume levels correspond to a power law increase in inefficiency.
4. Cross-over behavior is established between the two regimes.
5. The model accounts for contemporaneous correlation and fund heterogeneity.

Does price efficiency increase with trading volume? Evidence of nonlinearity and power laws in ETFs

Gunduz Caginalp
University of Pittsburgh
Department of Mathematics
301 Thackeray Hall
Pittsburgh, PA 15260
Phone: 412.624.8375
E-mail: caginalp@pitt.edu

Mark DeSantis
(corresponding author)
Chapman University
Argyros School of Business and Economics
One University Drive, Orange, CA 92866
Phone: 714.997.6957
E-mail: desantis@chapman.edu

Abstract

Whether efficiency increases with increasing volume is an important issue that may illuminate trader strategies and distinguish between market theories. This relationship is tested using 124,236 daily observations comprising 68 large and liquid U.S. equity exchange traded funds (ETFs). ETFs have the advantage that efficiency can be measured in terms of the deviation between the trading price and the underlying net asset value that is reported each day. Our findings support the hypothesis that the relationship between volume and efficiency is nonlinear. Indeed, efficiency increases as volume increases from low to moderately high levels, but then decreases as volume increases further. The first part tends to support the idea that higher volume simply facilitates transactions and maintains efficiency, while the latter part, i.e., even higher volumes, supports the ansatz that increased volume is associated with increased speculation that ignores valuation and decreases efficiency. The results are consistent with the hypothesis that valuation is only part of the motivation for traders. Our methodology accounts for fund heterogeneity and contemporaneous correlations. Similar results are obtained when daily price volatility is introduced as an additional independent variable.

Key words: Volume; Nonlinearity; Scaling laws; Cross-over behavior; Price efficiency; Volatility

1. Introduction

The empirical relationship between volume and efficiency is an important issue that has not been previously investigated. Thus there is no data-driven basis for even a phenomenological theory such as a power law relationship between the two. Classical finance hints that efficiency should improve with increased volume, but does not specify the micro-economic mechanism whereby this would be realized. On the other hand, behavioral finance generally suggests that high volume is often a consequence of increased speculation that by definition corresponds to less focus on value than price action. So it is not surprising to have lower efficiency with higher volume. However, it is clear that without a careful analysis of data, it would be impossible to determine which line of reasoning prevails in the various regimes of the financial markets.

With the recent rise in computerized (algorithmic) and high frequency trading, the typical volume of trading has dramatically escalated along with the controversy surrounding it. Proponents claim that the increased trading results in higher volume that increases efficiency. Meanwhile, many consumer advocates feel that high-frequency trading is effectively fleecing the retail trader and money manager. Once again, without a detailed analysis, the discussion is more philosophical and political. Thus, it is an important open question as to whether a dramatic increase in trading volume increases efficiency by providing more information relating to price discovery and greater liquidity, or whether the increased volume diminishes efficiency due to the introduction of traders who are oblivious to valuation in their pursuit of very short term trends.

A number of studies have confirmed that stocks that are characterized by low liquidity (including low volume) have been mispriced more than those that are traded with high volume (see Bodie et al. 2008). To the best of our knowledge, however, there has not been a prior study that establishes that the *same* stock is less efficient when the volume is lower, as one might expect from classical efficient market theory (as discussed below). There is a significant difference when one is comparing different stocks, as noted in Bodie et al. (2008), in that fewer analysts cover the lower volume stocks, for example. We find that efficiency increases with rising volume, but only to a point after which efficiency decreases. Thus, our results offer support for the two hypotheses (described below) that efficiency should either improve or deteriorate with increasing volume.

The case for a positive correlation between high volume and high efficiency can be formulated based on classical finance as follows. Suppose there are two financial instruments that represent the same asset. An example may be a basket of individual stocks (that are highly liquid with high volume) and a fund that invests in them (e.g., an exchange traded fund that is also highly liquid). If the latter is trading at a discount compared to the underlying stocks, then an investor seeking to buy these stocks would profit from buying the undervalued instrument. Even for an investor who does not seek a position, there would be a profit from buying the undervalued and selling short the overvalued instrument.

However, in practice there is the issue of not only how much of a discount but what is the dollar amount that can be purchased at that discount. For example, if there is a 2% discount and the volume is such that only $1,000 worth of the asset can be bought per day, it will not be attractive to an arbitrageur. If there is a 2% discount with high volume so that $1 million can be bought, then the potential profit of $20,000 will likely draw buyers. Thus, a 2% discount that can easily persist when the dollar volume is in the thousands will not likely persist when it is in the millions.

In terms of market microstructure and consistent with this line of reasoning, it is not much of a surprise if stocks with a modest volume exhibit clear deviations between price and value. Indeed, these stocks usually coincide with those with very few market makers. The market makers are aware that as more of them enter a small market (i.e., involving a stock with small volume) the potential profit becomes negligible. On the other hand, for stocks with a huge volume, even small deviations from what they perceive to be realistic value will result in large absolute (arbitrage[1]) profits, so many agents will be functioning as market makers for the stock, thereby increasing the efficiency.

On the other hand, there is also the behavioral perspective which suggests that increased volume and trading may also be associated with speculative activity that tends to distort prices.[2] This has been observed with the high-tech and internet bubble and with the initial public offerings of the late 1990's. The idea is that as public participation increases dramatically, more money enters the market with less attention to value. Indeed, increased volume in a particular asset tends to push up prices and induce an uptrend, which in turn attracts momentum traders who are less concerned with the fundamental value of an asset. This perspective suggests that very high volume should be associated with larger deviations, i.e., greater inefficiency than associated with moderately high volume.

In particular, consider for a particular stock the distribution of its daily dollar volume. The market makers specializing in this stock will have a supply of capital that is adequate to exploit any inefficiency in pricing under most circumstances. In the event that there is increased speculative activity that boosts the volume to unusual levels, for example to three standard deviations above the mean, then it is quite possible that the capital of the market makers will be exhausted in the face of the increased speculative activity. Meanwhile the speculators, by definition, are relatively unconcerned with valuation and deviations from a realistic value and will buy or sell at levels that are not necessarily close to the valuation. Thus, one can argue, based on behavioral finance arguments, that increased volume is often a result of a behavioral bias, e.g., trend following, over-reaction to a news event, etc., so that one would expect large volume to result in greater inefficiency. As noted above, the arbitrage capital of the market makers would tend to be exhausted during these episodes, so the forces that lead to efficient markets would take a subordinate role.

While the two sets of arguments are antithetical, they are not mutually exclusive. In statistical

[1] Friedman (1953) and Fama (1965) argue that sophisticated investors will arbitrage any deviations from fundamental value due to the actions of irrational or "noise" traders; thereby keeping prices in line with valuations.

[2] This reasoning is consistent with the investor sentiment literature. Beginning with Keynes (1936), researchers have studied the idea that investor sentiment can cause prices to deviate from valuations. Consider, for example, Baker and Wurgler (2006, 2007) and Stambaugh et al. (2012). In addition, Huddart et al. (2009) find evidence of higher stock trading volumes when prices exceed their 52-week highs.

mechanics, one often observes different regimes in which different phenomena are evident. Each regime may have its own power law. The two regimes are often connected through cross-over behavior. In this case there may be different power laws for the efficiency argument and the behavioral argument. Classical finance would stipulate $C_1 \times Volume^{p_1}$ where p_1 is negative, while the behavioral argument would stipulate $C_2 \times Volume^{p_2}$ where p_2 is positive, and both C_1 and C_2 are positive constants. Under these conditions, the sum of these two effects would be manifested in a function with a critical value (i.e., a minimum).

Hence, the first step in understanding the role of volume in efficiency is to perform an empirical study in order to determine whether inefficiency decreases monotonically with volume, or whether it has a minimum value, suggesting a role for the behavioral argument. For the latter possibility, the goal is to determine the minimum in terms of the number of standard deviations of volume from its mean.

Since there is no unique way to assess the value of a stock, testing these alternative hypotheses decisively is difficult. However, one important class of assets, namely exchange traded funds (ETFs) provides an opportunity to evaluate these theories. Since ETFs have a net asset value (NAV) that is announced at the end of each trading day, the difference between the trading price and NAV can be examined in the context of trading volume. ETFs also have "authorized participants" who can request (subject to a minimum number of shares) the underlying stocks in an ETF. This issue was discussed in Caginalp et al. (2014), which concluded that the deviations from NAV within this arbitrage range were significant in terms of determining underlying trader behavior.

In essence, we treat the collection of ETFs as a laboratory analogous to the controlled experimental environments of, for example, Smith et al. (1988). As noted by Frydman et al. (2014), “The advantage of experiments is that they give researchers a large degree of control over the trading and information environment, which can make it easier to tease theories apart.” Indeed, in asset market experiments the fundamental value of the asset can be known with certainty. While the intrinsic value of a stock remains elusive and debatable, the NAV of an ETF is a suitable proxy for this value. Thus, the ETF universe serves as a large laboratory in which we may effectively test hypotheses related to the price’s deviation from fundamental value.

Indeed, we use a large data set consisting of 124,236 daily observations of ETF prices to test whether volume and efficiency are correlated and to determine the nature of their relationship (if one exists). The main idea is to perform a regression of the daily deviation between price and NAV against the standardized dollar volume and its square. We utilize the higher power of the volume in order to extract any nonlinearity from the data. We include both fund and time fixed effects to account for fund heterogeneity and contemporaneous correlations, respectively. We therefore obtain coefficients of the quadratic polynomial of dollar volume. We can then plot these to examine the relationship between the relative deviation of price and NAV and consider whether this deviation is a decreasing function of the volume.

We find that the deviation does indeed decrease with increasing volume so long as volume remains below approximately 0.91 standard deviations. In particular as we move from relatively low[3] volume to average volume, there is a significant decrease in deviation, so efficiency improves in this range. The situation is different, however, once the volume reaches the 0.91 standard deviation level. At this point the deviations begin to increase significantly.

It is important to note that we are interested in the impact of volume changes on the pricing efficiency of ETFs. The goal is not to determine whether funds with larger average trading volumes have smaller deviations from NAV than those with smaller average trading volumes (which would entail a cross-sectional analysis). Rather we consider whether, for example, an increase in volume has, on average, a positive/negative effect on the deviation between price and NAV. This explains our standardization procedure in which the independent variables are standardized on a per fund basis rather than standardizing across all funds and time periods. The daily volume in terms of the number of standard deviations from the mean daily volume for that fund is a measure of the additional market participation.

2. Utilizing ETF data and inefficiency levels to test a key concept in markets

[3] Note that the ETFs utilized in this study are quite liquid. The median firm's average daily dollar trading volume is 16.56 million. Thus, funds with a "relatively" low volume on a particular day are, on average, still very liquid instruments.

In this section we discuss the characteristics of ETFs and the features of the premium/discount that will be instrumental in our analysis of the relationship between volume and efficiency. We also review the literature on volume's expected impact on price changes/efficiency.

During the past two decades, ETFs have become a major segment of the stock markets. While they have been designed to closely track the underlying assets as discussed below, the deviation from their net asset value is not negligible (as documented by several works cited at the end of this section). Our interest, however, is not to add to this list of works, but rather to use this large set of funds as a laboratory to explore a key idea that is intrinsic to markets. In particular, as discussed above, one clearly expects that price efficiency should increase with rising volume. For ETF's, unlike most stocks, the efficiency can be measured by comparison with the NAV. If one can show that efficiency decreases with increasing volume beyond a critical value of volume, then it would be consistent with the hypothesis that traders are influenced by additional factors beyond the valuation of the underlying assets. To develop this further, we need to understand the basic structure of ETF's.

ETF Characteristics

ETFs are similar to closed-end funds in that they have a publicly available NAV and are actively traded on the exchanges. However, while closed-end funds tend to trade at persistent premiums or discounts (see Anderson and Born 2002 for a survey of the literature); ETFs tend to trade close to their NAVs. This difference is explained by the creation/redemption mechanism afforded to authorized participants of ETFs (as explained in the ETF prospectuses). If a fund is trading at a large discount to NAV, then an authorized participant will redeem a large block of ETF shares (typically 50,000) for an equivalent number of shares of the underlying assets. Similarly, if the fund is trading at a large premium to NAV, then the authorized participant will exchange a portfolio of the underlying asset shares for a block of ETF shares (i.e., new ETF shares will be created). Thus, if authorized participants behave as expected, then ETF prices should closely track the NAV.

As discussed in Caginalp, et al. (2014), the authorized participants only act when their gains exceed their costs. That is, there exist arbitrage bounds within which these authorized participants do not trade. However, in addition to the authorized participants, all other traders are also attempting to maximize their profits by buying/selling due to under/over-valuation. Thus, on balance, one expects that the daily ETF price movements (relative to NAV) largely reside within these "arbitrage bounds." Moreover, one might still expect price efficiency to increase with volume even within the arbitrage bounds of the authorized participants.

Price efficiency

Practitioners have long acknowledged a relationship between volume and price as described by adages like "Volume is validity" and "It takes volume to make prices move." Thus, it is not surprising that there is an extensive literature pertaining to trading volumes and prices. Note, however, that these studies focus on volume's effect on prices or price changes – not price efficiency or mispricings (i.e., the deviation between price and fundamental value). This, of course, is to be expected as the fundamental value of most assets is not easily observed or calculated.

Bris et al. (2007) note that "As the voluminous literature on efficient market theory suggests, there is no universal test for relative market efficiency…" Indeed, since the intrinsic value of an asset is not observable, there has been a need to develop sophisticated measures for the efficiency of stocks.[4] ETFs, on the other hand, provide a rather unique opportunity to circumvent this issue (or at least diminish its effect) as they provide a publicly reported (daily) net asset value, which we utilize as a proxy for the intrinsic value.[5] Thus, smaller deviations between price and NAV correspond to more efficient prices, while larger deviations correspond to less efficient prices.

[4] See, for example, studies by Bris et al. (2007), Hou and Moskowitz (2005), and Saffi and Sigurdson (2011).

[5] We follow the example set forth by Klibanoff et al. (1998), who utilize the NAV as a proxy for the fundamental value of closed-end funds. Indeed, they contend that NAV is a better proxy for the fundamental value of a closed-end fund than its price. In a similar fashion, we utilize the NAV as a proxy for the intrinsic value of the ETF.

The NAV incorporates the market price of the underlying securities, so it is already the best possible adjustment for the impact of basic economic and financial indicators such as GDP growth, employment, factory orders, interest rates, etc. that are commonly used as additional independent variables in studying the price dynamics of stocks. Furthermore, the NAV presents a more accurate adjustment since different groups of stocks are impacted differently by these variables. For example, an ETF consisting of bank stocks would be impacted much more strongly by a rise in interest rates than ETFs in a software sector that does not need to borrow. Thus, the use of these macro indicators as independent variables may be important for examining ordinary stocks, but in the context of examining the efficiency of ETFs it would amount to double inclusion of these macro indicators.

Literature on ETF premiums/discounts

As noted above, it is generally assumed that ETFs trade close to their NAVs. For example, Engle and Sarkar (2006) find that the deviation between price and NAV for domestic ETFs is very small and, on average, only lasts a few minutes. However, other studies have found empirical evidence suggesting ETF prices can deviate, significantly in some situations, from the underlying NAVs. See, for example, Delcoure and Zhong (2007), Tse and Martinez (2007), Ackert and Tian (2008), and Caginalp et al. (2014). We also find a significant deviation (in magnitude) between NAV and price. Indeed, Table 1 notes that one-quarter of the funds in our study have a deviation with a magnitude of at least 0.056%.

Petajisto (2013) analyzes the price efficiency of ETFs and finds that funds holding domestic, liquid assets are typically more efficient than those holding illiquid (foreign) assets. However, “even U.S. sector funds holding liquid domestic stocks … exhibit nontrivial premiums.” While Petajisto reports volume characteristics of ETFs, he does not consider the role volume plays with respect to price efficiency. Delcoure and Zhong (2007) consider 20 iShares country ETFs and find that volume, along with other factors such as institutional ownership and the bid-ask spread, affects the magnitude of the premium (discount); however, they note that these factors do not fully explain the deviations, leaving room for behavioral considerations.

Using a data set comprised of 28 ETFs, Ackert and Tian (2008) find an inverted "u-shaped" relationship between ETF premium (deviation between NAV and price) and liquidity. However, this nonlinearity is insignificant when they restrict their focus to U.S. ETFs. They also have a different focus, the "cross-market liquidity effects for country ETFs," than the present manuscript in which the focus is on how the efficiency of a U.S. equity ETF depends on its own volume.

Our work is in a different direction as we seek a relationship between efficiency and volume, rather than studying the magnitude of inefficiency. The finding that price efficiency does not continue to increase as volume increases is interesting from a fundamental perspective, as it is difficult to envision a classical explanation for why efficiency should begin to decrease at some point as volume continues to increase. As volume increases, the limitations to efficient markets diminish, at least from a classical perspective. In examining this issue, the addition of a plethora of independent variables beyond volume might be useful in a practical context but counterproductive in the analysis of whether classical factors continue to dominate at the very high volume end of the trading spectrum.

3. Data

The data for this study were obtained from the Bloomberg terminal. We restrict our focus to the 68 active large- and mid-cap U.S. equity ETFs with observations for each day of the time period ranging from April 1, 2009 to July 1, 2016. Thus, we have a time series cross-sectional (or panel) data set. This data set is balanced as each fund has 1,827 daily observations. We select this date range subsequent to the Great Recession crisis period to exclude a very tumultuous trading period. Furthermore, straddling a major crisis might introduce spurious artifacts of regime change. Nevertheless, it is interesting to examine whether the efficiency-volume relationship is different in distinct time periods. We include an analysis of the post-crisis, crisis, and pre-crisis time periods in Appendix A along with an examination of the data set that encompasses all three time periods, January 3, 2006 through July 1, 2016.

We treat DIA and SPY separately as they are substantively different from the other ETFs. Indeed, an obvious distinction is that they cover the major indexes that are benchmarks for many managers. Also, there are many mutual funds and pensions that track them. So, these can be expected to have a level of efficiency that is much higher than the typical ETF.

We consider large U.S. equity ETFs. That is, at least 70% of each fund's holdings are large- or mid-cap equities falling within the U.S. None of the funds are leveraged, actively managed, or derivatives-based. Moreover, none of the funds hedge exchange rate risk or utilize swaps as a primary means of tracking an index. The funds' underlying indexes are weighted via market cap (59 funds), dividend (3 funds), fundamentals (2 funds), equally (1 fund), and multi-factor (3 funds). The funds' average inception date was June 26, 2002.

Table 1 provides statistics for this data set. The average daily trading volume for these ETFs is comparable to that of the most actively traded U.S. stocks. Indeed, the mean (per fund) daily trading volume is approximately 3.43 million shares with one-half of the funds having an average volume of at least 250,000 shares, and one-quarter with an average volume of at least 1.13 million shares. The minimum average volume is 19,000 shares and maximum average volume for any fund is 68.1 million shares. Thus, the funds are liquid and the data set provides a robust cross-section of funds by daily trading volume.

The primary objective of this study is to assess the impact of dollar trading volume on price efficiency, which we define as the relative deviation between net asset value (NAV) and price, while the dollar trading volume represents the number of shares traded multiplied by the price of each trade. The valuation spread ranges from -3.33% to 3.25% with an average (by fund) of -0.0035%. Note that the average of a list containing both positive and negative numbers can be misleading due to cancellations of positive with negative deviations. As such, we report statistics for the absolute value of the valuation spread. It has an average value of 0.053%. Further, one-half of the funds have a valuation spread magnitude of at least 0.046%, and one-quarter of the funds have a valuation spread greater than 0.0.056%. As noted above, these figures are consistent with previous literature. For example Krause and Tse (2013) find an average deviation between NAV and price of 0.168%. Thus, we find evidence of significant deviations between NAV and price, though the deviation is about 1/3 of the prior study and may be due to the less turbulent time period of our data set.[6]

A possible explanation for an ETF price's deviation from NAV is that the underlying securities are either illiquid or trade during a different time period than the ETF itself. In these situations the NAV might reflect the "stale" prices of the underlying securities and not accurately represent the fundamental value of the ETF. Petajisto (2013) finds that "stale" prices do not completely explain the mispricings, and, as noted above, he finds that ETFs with liquid underlying securities are typically "priced relatively efficiently." We avoid the concern of stale prices by considering (as discussed in the next section) a data set consisting of 68 large, U.S. equity ETFs not country-specific funds.[7] These ETFs are comprised of highly liquid stocks with large trading volumes. Indeed, even the ETF with the lowest average daily volume (JKD, iShares Morningstar Large Core Index) has very liquid holdings, including Microsoft Corporation (MSFT), Johnson & Johnson (JNJ), Procter & Gamble Company (PG), and Bank of America Corporation (BAC). The prior studies cited above show that the large U.S. equity ETFs have, on average, smaller deviations between price and NAV than country funds, for example. The deviations from NAV might also be attributed to large bid-ask spreads which serve to limit arbitrage activity. However, the bid-ask spread for the ETFs in this study is generally on the order of a penny as noted in Caginalp et al. (2014). Thus, our findings are surprising in that

[6] The Krause and Tse (2013) figure of 0.168% is for the time period April 3, 2003 through September 30, 2010 and therefore includes the Great Recession crisis period. For comparison purposes, the mean absolute valuation spread for our 55 fund data set (see Appendix A) for the time period January 3, 2006 (the earliest date in our data set) through July 1, 2016 is 0.116%. If we restrict the date range to the crisis period, September 1, 2008 through March 31, 2009, then the mean absolute valuation spread is 0.28% (see Panel B in Table 3).
[7] Petajisto (2013) notes that the "ETF market has increased almost 3,000%" between the end of 2000 and the end of 2010. Thus, these earlier papers were limited by the number of funds that could be studied. Indeed, according to the Investment Company Institute's report on March 28, 2014, there were 1,327 ETFs with over $1.7 trillion in net assets as of February 2014. Per the Investment Company Institute's 2013 Fact Book, 53rd edition, there were only 152 ETFs with $228 million in net assets at the end of 2004, which roughly corresponds to the ending date of the Ackert and Tian (2008) data set.

we note a convex relationship between trading volume and the absolute relative deviation between price and NAV. [8]

The significance of the magnitude of mispricing varies with investment strategies. For an investor with a ten year horizon, a 0.056% inefficiency may be trivial. However, for active traders, repeated profits of 0.056% or more would represent a substantial profit. Moreover, from our perspective, the nature of the deviations illuminates key aspects of trader motivations that will be useful in constructing a phenomenological theory.

An appropriate measure for volatility is the relative difference between the high and low prices for the day (refer to Section 4 for the detailed definition). The average (per fund) volatility of prices is 1.27% with a minimum value of 0.96% and a maximum of 2.23%. The average (per fund) correlation between volatility and dollar trading volume is 0.19. A positive correlation between these two factors would suggest that an influx of trading interest may be boosting both volume and volatility. This could be attributed to changes that bring some uncertainty as well as news. But it may also be due to an influx of less informed traders, explaining the increased volume, volatility and inefficiency. In our sample the observed correlation is positive, but small, at 19%, indicating that the volume and volatility are almost independent.

[8] Tetlock (2008) also finds a positive relationship between liquidity (proxied by trading volume as well as the bid-ask spread and market depth) and efficiency by examining prediction markets.

Table 1. Descriptive statistics for data set. Mean values are calculated on a per fund basis and then averaged over all funds. Note that the standard deviation is not listed as a percentage.

68 fund data set	Mean	St. Dev.	Min	First Quartile	Median	Third Quartile	Max
Mean Market Capitalization (Millions)	4,988	7,778	246	556	1,645	5,645	41,399
Mean Volume (# of shares in Millions)	3.43	10.78	0.019	0.076	0.250	1.13	68.10
Mean Dollar Volume (Millions)	162.23	463.51	0.88	3.93	16.56	75.43	3,430.48
Mean Volatility (%)	1.27	0.0025	0.96	1.10	1.20	1.38	2.23
Mean Valuation Spread (%)	-0.0035	0.00012	-0.041	-0.012	0.0014	0.0043	0.019
Minimum Valuation Spread (%)	-	-	-3.33	-	-	-	-
Maximum Valuation Spread (%)	-	-	-	-	-	-	3.25
Mean Absolute Valuation Spread (%)	0.053	0.00025	0.030	0.038	0.046	0.056	0.170
Expense Ratio	0.25	0.15	0.05	0.135	0.2	0.4	0.59
Mean Correlation* (Dollar Volume and Volatility)	0.19 (0.05)	-	-	-	-	-	-

*- Reported correlation is calculated on a per fund basis and then averaged across funds. The value in parentheses represents the correlation across the entire data set.

4. Variable Definitions

For an ETF, a key measure of fundamental value is the NAV. Clearly, if the price and NAV of an ETF differ, then they cannot both be efficient. As such, we define a valuation spread variable, V_t, as the relative deviation between NAV and price. Indeed, we have

$$V_t = \frac{NAV_t - P_t}{NAV_t}$$

where P_t corresponds to the close price[9] on day t. As we are interested in the mispricing, we utilize the absolute value of this variable, $|V_t|$, as the dependent variable in our regressions (see Section 5).

[9] The CEF_PCT_PREM field in Bloomberg is multiplied by -1/100 and utilized for the Valuation Spread variable.

We focus our efforts on analyzing the relationship between dollar trading volume and efficiency.[10] Dollar trading volume measures the total dollar amount of trading in a particular security during the day. This is more significant than share volume since a trader seeking to profit from a mispricing needs to examine the total profit in dollars that is possible. Dollar volume is also often more informative than, for example, other more vague measures of the ease of buying or selling. In this case one might, for example, consider the bid/ask spread. However, from the perspective of a market maker, or any short term trader, it is only a small part of the picture. For example, the bid/ask spread might be two cents, but the number of shares available at that spread might be only 1,000. Thus, the total dollar volume is valuable because it is a measure of how many dollars of the stock one can purchase without altering the price. Suppose there are two stocks that both feature a penny bid/ask spread. However, stock A might have a dollar volume of $10 million, while stock B has a volume of $10 billion. If one would like to buy $5 million of stock, then one can buy stock B at the same price as someone buying a thousand dollars of the stock.[11] However, with Stock A, it would be impossible to buy $5 million of stock (i.e., half of the daily volume) at the same price (i.e., without moving the price up). If one is a market maker or a trader trying to make money on either bid/ask spreads or stochastic fluctuations, then it is the dollar volume that is a key factor in determining whether and how to trade that ETF. The total dollar volume represents the volume in a way that one can compare the opportunity (relative to other stocks) for the market makers.

The daily dollar trading volume variable, Vol_t, corresponds to the "Turnover" variable available via Bloomberg. This represents the number of shares traded multiplied by the price at which the transaction occurred.

An additional issue is whether high daily volume is associated with high daily volatility. In other

[10] Trading volume is frequently associated with liquidity, though a consensus on the appropriate measure for liquidity does not exist. The literature uses several proxies for liquidity including volume, turnover (ratio of volume to shares outstanding), price, and the bid-ask spread. We do not take a stance on the appropriate measure for liquidity. Rather, we focus on the fundamental relationship between price efficiency and volume.

[11] Note that $5 million is 1/200th of $1 billion. So, this purchase would be equivalent to buying an amount that is traded approximately 200 times per day. Put another way; consider that the market is open for 6.5 hours or 390 minutes. This $5 million purchase is likely to be completed in 390/200 or about 2 minutes. Thus, if one has an order to buy at the asking price there is not much doubt that it would be filled by the sell orders that come in every two minutes.

words, we consider the possibility that the larger NAV-Price deviations are caused by high volatility that is associated with high volume. To resolve this issue, we make an appropriate definition of volatility for our purposes, namely, the absolute value of the relative difference between the day's high and low prices.[12,13] That is, we define the volatility variable as the relative deviation between the high, H_t, and low, L_t, price for an ETF on a specific day, t. That is, we define

$$Volatility_t = \frac{H_t - L_t}{L_t}.$$

The introduction of the volatility variable is important in terms of eliminating the possibility that volatility is responsible for both high volume and inefficiency. In other words we examine whether inefficiency is caused by volatility since it makes arbitrage more difficult.

To determine the existence of a nonlinear relationship between the independent variables and the dependent variable, we include the square of both independent variables and their interaction term in the regressions.

5. Methodology

Our data set may be characterized as a panel or time-series cross-sectional data set with big N and big T (i.e., a large number of cross-sections, funds, and observations, days). We employ regressions with fixed effects for both fund and time to account for heterogeneity across ETFs and contemporaneous correlations, respectively. Indeed, if a fund's price is inefficient due to an inherent, time-invariant fund characteristic, then the fund fixed effect will account for this. Similarly, if there is an exogenous shock to the market that, say, increases trading volume for all funds on a particular day, then the time fixed effect will allow for this.

---

12 The daily high and low prices correspond to the Bloomberg PX_HIGH and PX_LOW fields, respectively.

13 This definition is in the spirit of the Parkinson (1980) measure, which is defined as $[ln(H_t) - ln(L_t)]/\sqrt{4ln(2)}$ for daily data. As the results are essentially unchanged with the use of the Parkinson measure in place of our Volatility variable, they are not included. However, they are available upon request.

The data set is grouped by fund, and for each fund we have 1,827 daily observations. To test the existence of a (potentially nonlinear) relationship between the magnitude of the valuation spread and volume, we perform a regression analysis of the form[14]:

$$|V_t| = \beta_0 + \beta_1 Vol_t + \beta_2 Vol_t^2 + \beta_3 Volatility_t + \beta_4 Volatility_t^2 + \beta_5 Vol_t * Volatility_t + u_{i,t}$$

where

$$u_{i,t} = \mu_i + \gamma_t + \varepsilon_{i,t}.$$

Let i=1,…, 68 denote funds, while t=1,…, 1,827 represents days. The fund-specific effect and time-specific effect correspond to μ_i and γ_t, respectively, while the idiosyncratic error term is given by $\varepsilon_{i,t}$.

In addition to controlling for correlations across time for a given fund and across funds for a specific date[15], we also consider the possibility that the regression residuals might be correlated by fund and/or time, which could lead to biased standard error estimates. One method to account for this correlation is to cluster the standard errors by fund, time, or both fund and time (see Petersen 2009 and Thompson 2011). Thompson (2011) notes that “all else equal, it is more important to cluster along the dimension with fewer observations…” and that clustering by both fund and time is most effective when the number of observations of both fund and time are similar. As we have 68 funds each with 1,827 observations, we cluster by fund. Thus, we account for any within-fund (across time) correlations. In addition, robust standard errors are employed to account for heteroscedasticity[16] and to ensure we do not overstate the significance of our findings.

[14] Note the independent variable is restricted to nonnegative values, while the Volume variable may be negative. We used each regression reported in Table 2 to estimate the predicted values for $|V_t|$. For each regression, less than 1% of the estimates were negative. Thus, it is reasonable to utilize this model rather than investigate potential generalized linear models.

[15] A detailed description of how the time fixed effect accounts for contemporaneous correlations is provided in Caginalp, et al. 2014.

[16] The SAS PANEL procedure provides different options (HCCME) for producing a heteroscedasticity-consistent covariance matrix. These options are detailed in both the SAS user's guide and Davidson and MacKinnon (1993). HCCME=0 produces the robust standard errors due to White. However, as noted in Davidson and MacKinnon (1993),

To reduce the impact of any outlier data points we winsorize the independent variables at the 5th and 95th percentiles. That is, if the variable's value is in the 5th (95th) percentile, then that value is replaced by the value corresponding to the 5th (95th) percentile. We winsorize by percentile rather than standard deviation to ensure equal weights are given to the observations in both tails of the distribution.

To facilitate comparison of the effect each (winsorized) independent variable has on the magnitude of the valuation spread, we standardize these variables (dollar trading volume and volatility). That is, for each fund i we compute the mean, m_i, and standard deviation, σ_i, for each variable, x. We then define the standardized variable as $(x - m_i)/(2\sigma_i)$. Gelman (2008) argues that standardizing by two standard deviations is better than one and suggests division by two as the default approach. This is essentially a scaling issue, and we follow Gelman's recommendation.[17]

To produce Figures 2 and 3, the plotted intercept term is calculated as follows. Let $\hat{\beta}_i^{(cs)}$ and $\hat{\beta}_j^{(t)}$ represent the estimated coefficients for the fund (cross-section) and time fixed effects, respectively. Then, for each fund *i* we compute the corresponding intercept, i.e. mean $\left\{\sum_{j=1}^{T}\left(INT + \hat{\beta}_i^{(cs)} + \hat{\beta}_j^{(t)}\right)\right\}$ where *INT* is the SAS-reported intercept and *T* represents the number of days (e.g., *T*=1,827 for the 68 fund data set). Finally, we average this fund-specific intercept over all funds and utilize the resulting value as the intercept for the plots.

6. Results

6.1 Main results

---

there are alternative methods for correcting for heteroscedasticity. Our results are robust to these different options (HCCME=0,1,2,3,4). We report results for HCCME=3 which utilizes the squared residuals divided by $(1 + h_t)^2$ to estimate the diagonal entries of the error covariance matrix. The variable, h_t, is defined as $X_t(X^TX)^{-1}X_t^T$ where X is the regressor matrix with t^{th} row denoted by X_t.

[17] To facilitate interpretation of our results the discussion in the proceeding section reflects standardization by one standard deviation. Consistent with this approach, we present Figures 2 and 3 with the axes corresponding to the independent variables that reflect standardization by one standard deviation. That is, when plotting the magnitude of the valuation spread using the estimated coefficients, $\hat{\beta}$, from the corresponding regressions, we divide the volume and volatility values by two. Thus, a one unit change in one independent variable corresponds to a one standard deviation change in that variable.

We test for a relationship between inefficiency and dollar volume. If such a relationship does exist, then we examine the nature of this association. That is, does inefficiency decrease with increasing volume (see Figure 1A)? Does inefficiency increase with volume (see Figure 1B)? Do both relations hold? If so, can we identify the transition from one regime to the other (see Figure 1C)? We begin by considering a simple model of inefficiency versus volume for which we run a two-way fixed effects regression with robust standard errors clustered by fund. In Model 1 (Table 2) we use volume as the sole independent variable and the magnitude of the valuation spread, $|V_t|$, as the dependent variable. While the coefficient of the Volume variable is negative, consistent with the theory that prices become more efficient with increasing volume, the coefficient is not statistically significant from zero with a t-value of -0.87. Thus, there is little to support the thesis that increased volume should be accompanied by increased efficiency. However, linear regressions can obscure a fundamentally nonlinear relationship, and do not address our key question of whether there exists a convex relation.

Figure 1. Plot of price inefficiency versus dollar volume. The curve in Figure 1A represents the notion that price efficiency should improve with higher volumes. Figure 1B displays the speculative viewpoint that efficiency should decrease with higher volumes. Figure 1C reflects the hypothesized existence of a nonlinear relationship between price and value and marks the transition between the two regimes shown in Figures 1A and 1B. That is, prices become more efficient with increasing volume, but only up to a point after which efficiency deteriorates.

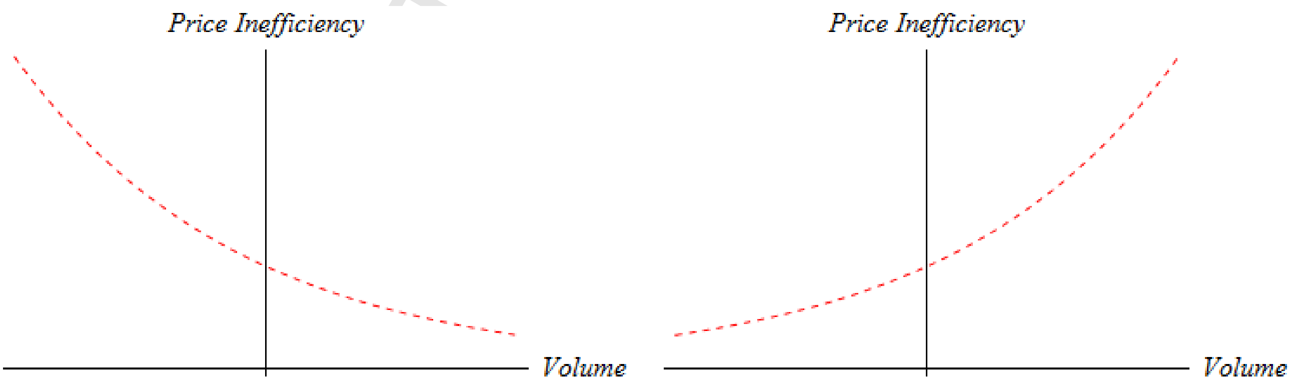

(A) Perspective: Efficiency improves with volume (B) Perspective: Efficiency decreases with volume

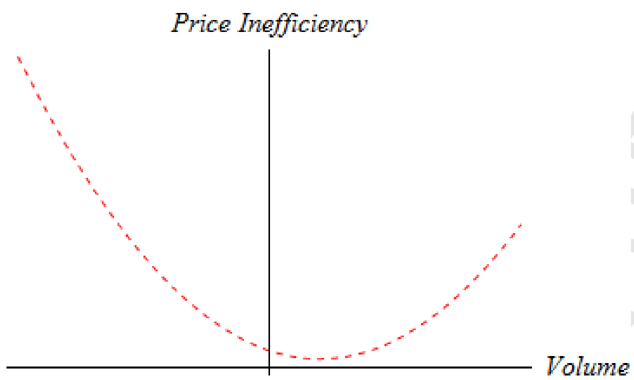

(C) Hypothesis: Transition between two regimes

Table 2. Regression Results. Results for the two-way fixed effects model with robust standard errors clustered by fund. Volume is the daily dollar trading volume, while volatility is the relative difference between the daily high and low prices. Models 1 and 2 may be considered baseline cases, while Models 3 and 4 augment these baselines with the Volatility variables and an interaction term. Models 5 and 6 serve as robustness checks. Model 5 augments Model 4 with proxies for illiquidity (square root of the Amihud Illiquidity measure) and fund size, while Model 6 includes the lagged dependent variable.

Term	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Volume	-0.02 (-0.87)	-0.07** (-2.16)	-0.02 (-1.03)	-0.07** (-2.31)	-0.04 (-1.5)	-0.03 (-1.46)
Volume^2		0.073*** (3.16)		0.077*** (3.22)	0.064*** (2.79)	0.058*** (2.78)
Volatility			0.018 (0.53)	-0.01 (-0.34)	-0.02 (-0.62)	-0.02 (-0.64)
Volatility^2				0.09*** (3.49)	0.091*** (3.49)	0.085*** (3.38)
Volume*Volatility				-0.04 (-1.39)	-0.04 (-1.32)	-0.03 (-1.27)
Market Capitalization					-0.08** (-2.47)	-0.08** (-2.51)
Amihud Illiquidity					0.022 (1.39)	0.019 (1.28)
Lagged |Valuation Spread|						98.1*** (8.37)
R-Square	0.2527	0.2535	0.2527	0.2541	0.2549	0.2612
No. Observations	124,236	124,236	124,236	124,236	124,236	124,168
No. Groups/Funds	68	68	68	68	68	68
No. Days (per fund)	1,827	1,827	1,827	1,827	1,827	1,826
F Test for no fixed effects	21.84***	21.55***	18.84***	18.39***	16.87***	12.69***

Notes: (a) t-values are reported in parentheses; (b) *** indicates P < 0.01, ** indicates P<0.05, and * indicates P<0.1; (c) Coefficient values have been multiplied by 1,000 for exposition; (d) The reported R-Square value corresponds to the R-Square measure developed by Theil (1961); (e) To ensure multi-collinearity is not an issue we compute Variance Inflation Factors, which are all less than 3.52.

We next consider whether the relationship between efficiency and volume is nonlinear. Indeed, we augment Model 1 with the square of the Volume variable and report the results in Model 2 (Table 2).[18] The coefficient of the linear Volume variable remains negative but is now statistically significant, while the coefficient of the quadratic Volume term is statistically significant with a positive sign. To better understand

[18] We perform a joint hypothesis test to determine if both we can reject the hypothesis that the coefficients of both the linear and quadratic volume terms are simultaneously zero. The Wald (10.69***), Lagrange Multiplier (153.52***), and Likelihood Ratio (153.62***) test statistics are all significant and enable us to reject the hypothesis that both coefficients are zero.

the relationship suggested by the results from Model 2, we plot the magnitude of the valuation spread as a function of volume only. Consider the dotted (black) curve in Figure 3. Note that prices become more efficient as volume increases to approximately 0.96 standard deviations above the mean. However, as volume increases beyond this point, prices become more inefficient.

This result suggests that increases in volume will usually lead to more efficient prices. However, if volumes are sufficiently large, which occurs approximately 16% of the time in our data set, then prices become increasingly less efficient. Indeed, 19,428 out of 124,236 dollar trading volume values are greater than the volume, 0.48 standard deviations, corresponding to the minimum inefficiency value for the data set. This is a non-trivial and economically significant result, particularly in light of the fact that this 16% of the data set has considerably greater volume than a typical day. Thus, the fraction of trades within this regime will be significantly higher than 16%. For market makers, who make their profit from the bid/ask spread, the significance is even greater.

We examine the role of volatility in price efficiency by including this variable along with the volume in linear and nonlinear regressions. Hence, we augment Models 1 and 2 with the Volatility variable(s) described in Section 4. Model 3 includes linear Volume and Volatility terms as independent variables. Note that neither variable has a statistically significant coefficient. However, the signs of both variables are consistent with the theory that efficiency should increase (decrease) with increasing volume (volatility).

The next question we seek to address is whether the relationship between inefficiency, volume, and volatility is nonlinear. Indeed, a linear regression may be utilized to explore a nonlinear relationship between the dependent and independent variable(s) by including higher powers of the dependent variable(s).[19] To this end, we augment Model 3 with the squared Volume and Volatility variables as well as an interaction term. The results, which are consistent with those from Models 2 and 3, are reported in Model 4 (Table 2). In fact,

[19] Consider $f(x) = x(x-1)$. A linear regression with $f(x)$ as the dependent variable and x as the sole independent variable will lead to a small and most likely insignificant coefficient for x. Indeed, a horizontal line would be the best linear approximation. Thus, a higher order term, in this case x^2, must also be included to understand the true relationship between the variables.

there is a positive correlation between the volatility and the volume per fund as noted in Section 3. This is consistent with the explanation that with increased trading the role of the market makers of a particular stock are subordinated to the activities of the additional traders. In other words, the resources of the market makers are finite, so the additional trading tends to reduce the capital of the market makers relative to the total trading volume. The simultaneous increase in volatility and inefficiency tends to suggest that the increased trading involves less informed investors.

The Volume variable has a statistically significant negative coefficient in Model 4 which is consistent with the notion that higher volume leads to lower deviations between price and intrinsic value. However, there is a different aspect to the relationship between volume and price efficiency. Indeed, the square of the Volume variable is also significant, but with a positive coefficient.[20] This suggests the deviation between price and NAV decreases with increasing volume to a point (approximately 0.91 standard deviations above the mean) beyond which the deviation actually increases. The linear Volatility variable's coefficient is negative, but statistically insignificant; while the quadratic Volatility variable has a statistically significant positive coefficient. Note the interaction term is not statistically significant.[21] Figure 2 depicts the three-dimensional relationship graphically through the plot of the function:

$$V(Vol, Volatility) = 0.201 - 0.07Vol + 0.077Vol^2 - 0.01Volatility + 0.09Volatility^2 - 0.04Vol * Volatility$$

where $V(.,.)$ represents the absolute value of the valuation spread variable described in Section 4.

[20] We again test the joint hypothesis that both volume coefficients are zero. We are able to reject this hypothesis as the Wald (12.02***), Lagrange Multiplier (156.54***), and Likelihood Ratio (156.54***) test statistics are all significant.
[21] An interaction term was also added to the linear Volume and Volatility variables in Model 3. The results mirror those of Model 3. Similar to Model 4, this interaction term is not statistically significant. For brevity, this regression is not included though it is available upon request from the authors.

Figure 2. Three-dimensional plot of inefficiency versus Volume and Volatility. Model 4 is used to produce a plot of the magnitude of the valuation spread, i.e. the relative mispricing, versus volume and volatility.

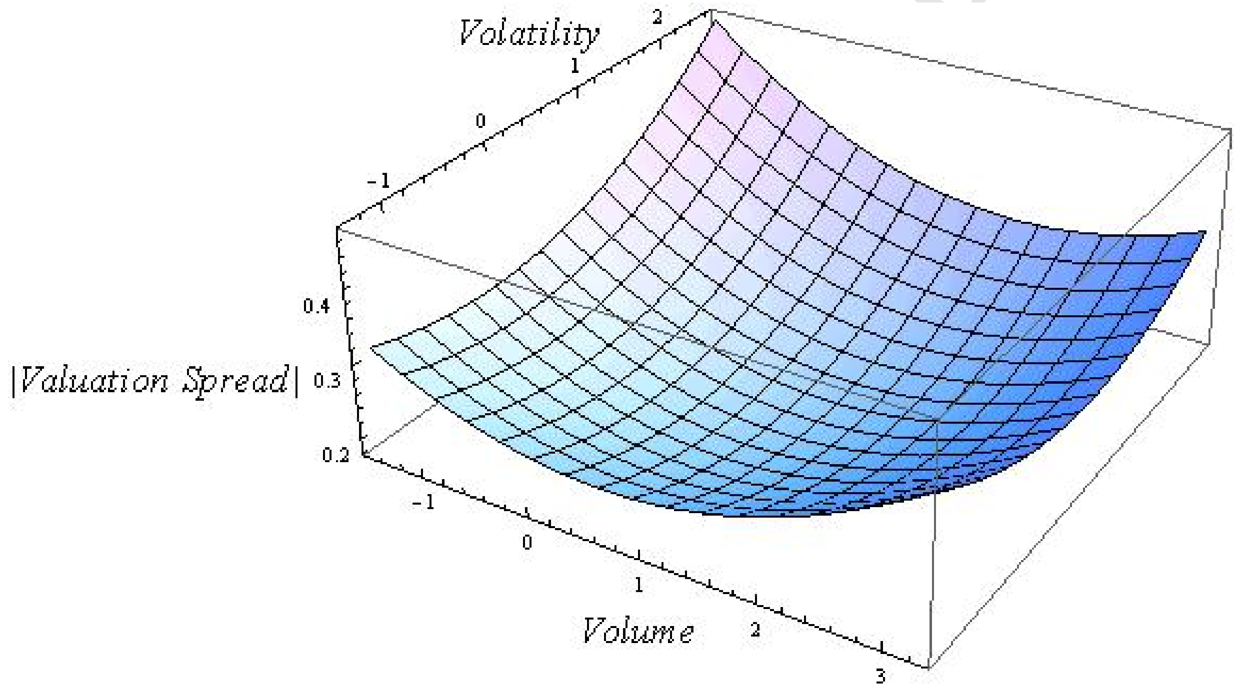

As previously noted we rescale the independent axes so that a one unit change in a variable corresponds to a one standard deviation change in that variable, i.e. we plot $V(Vol/2, Volatility/2)$. Also, note the quadratic relationship between inefficiency and both Volume and Volatility.

To gain a better understanding of the relationship between the valuation spread and volume, we consider the cross-section of Figure 2 in which the Volatility is zero, i.e., the Volatility corresponds to the average. The solid (blue) curve in Figure 3 reflects the function $V(Vol/2,0)$ for Model 4 and represents the plot of the valuation spread's magnitude versus dollar trading volume. Notice the curve decreases for volumes between approximately -1.5 and 0.91 standard deviations (above the mean) at which it attains a minimum value of 0.000185. For volumes greater than 0.91 standard deviations, the valuation spread magnitude increases. It is interesting to note that the inefficiency, as measured by the deviation between price and value, in a low volume environment is similar to that of a high volume environment. Indeed, the dashed (red) curve in Figure 3 shows that the inefficiency corresponding to a Volume of approximately -1.2

standard deviations is equivalent to the inefficiency associated with a Volume of approximately 3 standard deviations.

Figure 3. Price efficiency versus Volume. We use Model 2 to plot the absolute value of the valuation spread versus daily dollar trading volume (dotted, black curve). The solid (blue) curve corresponds to Model 4 with the standardized Volatility assumed to be zero. With respect to Model 4 the dashed (red) line shows that a three standard deviation increase in volume has, on average, an effect on price efficiency that is similar to a -1.2 standard deviation decrease in volume. Also note the mispricing decreases to approximately 0.185 (0.000185 as values have been multiplied by 1,000) at a volume of 0.91 standard deviations above the mean. Beyond this point greater volumes lead to larger deviations between price and value.

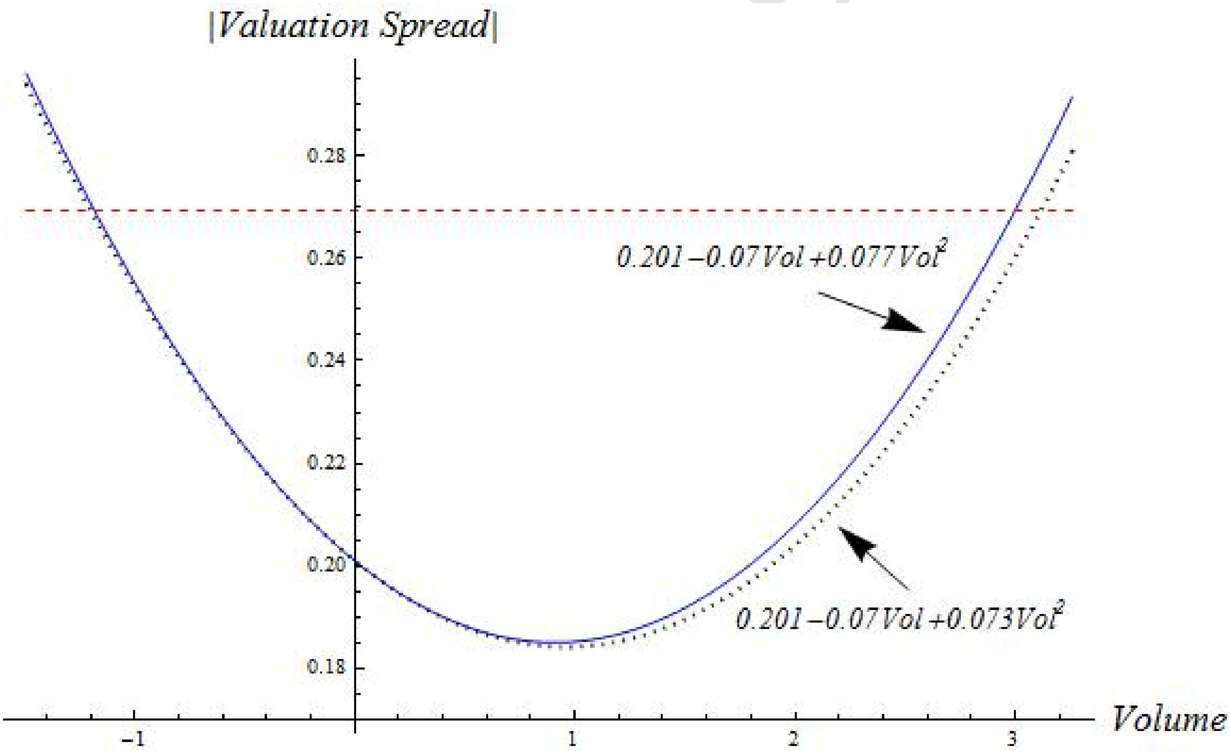

Model 4 was augmented with cubic Volume and Volatility variables as well as interaction terms up to second order. While the coefficient of the cubic Volume term is not significant in this regression, the quadratic Volume term's coefficient is both significant and positive. This provides further support for the hypothesis that the relationship between Volume and price efficiency is quadratic in nature. The results from this regression are not reported, but are available from the authors upon request.

6.2 Robustness tests

The nonlinearity observed in Figures 2 and 3 is robust to alternate model specifications. The fund and day fixed effects account for any omitted variable bias provided the excluded variable does not vary across time (e.g., an indicator variable denoting whether the fund is actively managed) and/or does not vary across funds (e.g., a macroeconomic variable like GDP would be constant across funds for a specific day). Two variables that do vary across these dimensions are fund size and the Amihud Illiquidity measure.[22] Fund size has been hypothesized to impact pricing efficiency. To test the robustness of our results (specifically the nonlinear relationship between efficiency and volume) to this factor, we consider a standardized Market Capitalization variable. In addition, we utilize the square root of Amihud's Illiquidity measure as a proxy for liquidity. We augment Model 4 with these additional control variables and report the results in Table 2 as Model 5. Note the standard errors are still clustered by fund.

While the coefficient of the square root of the Amihud Illiquidity measure is not statistically significant, the coefficient of the Market Capitalization variable is significant and negative. This does not suggest that prices of larger funds are more efficient. Rather, it indicates that if on a particular day a fund's market capitalization is greater than average, then its price tends to be more efficient. Since the number of shares outstanding for a particular fund tends to be fairly constant, the periods of higher prices for the stock are associated with higher efficiency. Note that the coefficient of the quadratic volume term remains statistically significant and positive.

To test the impact of prior price deviations from NAV, we include the lagged absolute valuation spread as an independent variable in our models. Model 5 is augmented with the lagged dependent variable as another regressor, and the results are reported in Model 6 (see Table 2).[23] As expected, the coefficient of this variable is highly significant with a positive coefficient. The presence of this variable, however, does not impact our results. In particular, the square of the Volume variable remains statistically significant and

[22] The Amihud Illiquidity measure (Amihud 2002) is defined as the ratio of the absolute value of the asset's return to its dollar trading volume.
[23] Inclusion of the lagged dependent variable reduces the number of observations in each regression by the number of cross-sections (funds).

positive.[24] Note that inclusion of this lagged variable leads to a dynamic panel model which may produce inconsistent coefficient estimates if both the number of funds and the number of time periods do not approach infinity (see, for example, Nickel 1981). Fortunately, we have a large number of both cross-sections (68 funds) and time periods (1,827 days). To ensure these numbers are large enough, we perform several Monte Carlo simulations. We generate data via the formula

$$y_{i,t} = \rho y_{i,t-1} + \sqrt{\delta_i}\varepsilon_{i,t}$$

where $\varepsilon_{i,t} \sim N(0,1)$ and δ_i is a random draw from a uniform distribution over the interval [0,1000]. We allow the variance to differ by cross-section so that our simulated data set more closely resembles our actual data set. Further to that end, we set the autocorrelation coefficient, ρ, to be equal to the autocorrelation of the absolute valuation spread in our data set, 0.288. We run 500 simulations and compute the absolute relative deviation between ρ and its estimate, $\hat{\rho}$. That is, for each simulation we calculate $|\hat{\rho} - \rho|/\rho$. These values are averaged across all simulations producing an average absolute relative deviation of approximately 0.92%. Thus, we have a sufficient number of both time periods and cross-sections to run the dynamic panel model.

We perform two further robustness tests in Appendices A and B. First, Model 10 in Appendix A is analogous to Model 5 but applied to a balanced 55 fund data set for the same time period (April 1, 2009 through July 1, 2016). Descriptive statistics for this data set are included in Appendix A (Panel D of Table 3). Note that our primary result, efficiency increases with increasing volume but only to a point after which efficiency declines as volume continues to rise, still holds.

Second, to enhance our intuitive understanding of the results, we discretize the region of standardized volume values into 50 subintervals. We then compute the average magnitude of the valuation spread variable within each of these subintervals and plot the results. This provides a straightforward check of our results. Indeed, consider Figure 4 in Appendix B. Most of the records generating this curve are close

[24] For Models 5 and 6 we again test the joint hypothesis that both volume coefficients are zero. We are able to reject this hypothesis as the Wald (14.56***/15.49***), Lagrange Multiplier (96.74***/80.44***), and Likelihood Ratio (96.78***/80.47***) test statistics are all significant for Model 5/6.

to zero where the curve is fairly flat. And, while this curve does not necessarily appear quadratic nature, it does appear that less efficient prices correspond to larger volumes.

7. Conclusion

We test the hypothesis that for a particular exchange traded fund the days of (relatively) larger volumes correspond to greater price efficiency, i.e. smaller relative deviations between price and fundamental value. Examining this deviation is difficult for ordinary corporate stocks because their fundamental values are not reported and/or easily calculated on a daily basis. To circumvent this issue we restrict our focus to ETFs, which are a large and important asset class, and utilize their NAVs as a proxy for the fundamental value. The specific funds we consider are large and liquid US Equity ETFs with large and liquid underlying equities. If there is a significant difference in the pricing of the ETF and the underlying assets, then they cannot both be efficient. Furthermore, if there is a pattern to the inefficiency, then it must reflect a non-random factor in trader motivations or strategies.

We find that larger volumes correspond to smaller deviation magnitudes, but only to a point (approximately 0.91 standard deviations above the mean). Volumes in excess of this point actually correspond to greater deviations, i.e. larger price inefficiencies. Thus, our results imply the existence of a nonlinear relationship between volume and price efficiency.

The first part of the inefficiency-volume relationship is intuitively well understood. When volume is small, the absolute profit for potential arbitrage agents is small, so that an inefficiency is more likely to persist. However, the reason behind the rise in inefficiency with rising volume is not innately obvious. This phenomenon is consistent with the hypothesis that some participants are motivated to trade based on factors beyond valuation. As in the Smith et al. (1988) experiments, when there is no doubt about valuation, the only uncertainty involves the anticipated actions of other traders. These motivations may include a spectrum of possibilities such as momentum trading, the heuristic affect (trading based on the attractiveness of the concept or theme), fear generated by news, etc. The results suggest that two regimes are at work and quantify the cross-over from one to the other. The minimum in the Volume-Inefficiency graphs is quantified and represents the point at which the transition is made from one dominant behavior to the other. In terms of the general theory, this minimum specifies the point at which further volume appears to be due to speculation that diminishes efficiency.

From a market microstructure perspective, one can also understand the empirical results in terms of the motivations of the shortest term traders who are essentially functioning as market makers. When the volume is not very high, the traders who exploit the bid/ask spread restore the market to efficiency. These traders are equipped with a level of capital that is commensurate with the volume under most conditions. When the volume is significantly higher than the mean, however, there is relatively little capital compared to those of the less informed traders (in terms of short term price movements) who have entered the market. Often price trend (or momentum) is a key factor for these less informed traders (see experimental results in Smith et al. 1988). Thus, as the trading volume in a particular stock moves above 0.91 standard deviations (above the mean), the market makers play a smaller role, contrary to their purported function.

In some extreme cases, one can see how this effect will be dominant. During the height of a financial crisis (e.g., in 2008), trading intensifies and volume expands, but market makers and short term traders are still limited to the same amount of capital. Hence, one can expect larger deviations from fundamental value simply because of the shift in balance between market makers (and short term traders) and investors who are eager to sell without much regard to value. Aggravating this situation is that investors focused on fundamentals are aware that they can obtain better bargains simply by waiting until panic drives prices lower.

Another aspect of the study is that one can use this methodology to understand the differences in trader motivations between different eras or asset classes. The question of whether market efficiency or trader motivations differed during the Great Recession crisis period can be addressed by examining this as a distinct time period. The results presented in Appendix A show, for example, that the linear coefficient of volume is positive for the pre-crisis period, suggesting the additional volume entering the market leading up to the crisis tended to result from the actions of traders/investors who were less concerned with valuation.

Acknowledgment

The authors thank Professors Jeffrey Wurgler and Fadel Lawandy as well as two anonymous referees for helpful comments.

References

Ackert, L.F., Tian, Y.S., 2008. Arbitrage, liquidity, and the valuation of exchange traded funds. Financial Markets Institutions and Instruments 17, 331-362.

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31-56.

Anderson, S.C., Born, J.A., 2002. Closed-end fund pricing: Theories and evidence. Norwell, MA: Kluwer Academic Publishers.

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

--, 2007. Investor sentiment in the stock market. Journal of Economic Perspectives 21, 129-151.

Bodie, Z., Kane, A., Marcus, A.J., 2008. Investments. Boston, MA: McGraw-Hill Education.

Bris, A., Goetzmann, W.N., Zhu, N., 2007. Efficiency and the Bear: Short Sales and Markets around the World. Journal of Finance 62, 1029-79.

Caginalp, G., DeSantis, M., Sayrak, A., 2014. The nonlinear price dynamics of U.S. equity ETFs. Journal of Econometrics 183, 193-201.

Davidson, R., MacKinnon, J.G., 1993. Estimation and inference in econometrics. New York: Oxford University Press.

Delcoure, N., Zhong, M., 2007. On the premiums of iShares. Journal of Empirical Finance 14, 168-195.

Engle, R., Sarkar, D., 2006. Premiums-discounts and exchange traded funds. Journal of Derivatives (Summer) 27-45.

Fama, E., 1965. The behavior of stock market prices. Journal of Business 38, 34-105.

Friedman, M., 1953. The case for flexible exchange rates, In Essays in Positive Economics. Chicago: University of Chicago Press.

Frydman, C., Barberis, N., Camerer, C., Bossaerts, P., Rangel, A., 2014. Using neural data to test a theory of investor behavior: An application to realization utility. Journal of Finance 69, 907-946.

Gelman, A., 2008. Scaling regression inputs by dividing by two standard deviations. Statistics in Medicine 27, 2865–2873.

Hou, K., Moskowitz, T.J., 2005. Market Frictions, Price Delay, and the Cross-section of Expected Returns. Review of Financial Studies 18, 981-1020.

Huddart, S., Lang, M., Yetman, M., 2009. Volume and price patterns around a stock's 52-week highs and lows: Theory and evidence. Management Science 55, 16-31.

Keynes, J.M., 1936. The general theory of employment, interest and money. Boston: MacMillan.

Klibanoff, P., Lamont, O., Wizman, T.A., 1998. Investor reaction to salient news in closed-end country funds. Journal of Finance 53, 673−699.

Krause, T., Tse, Y., 2013. Volatility and return spillovers in Canadian and U.S. industry ETFs. International Review of Economics and Finance 25, 244-259.

Nickell, S., 1981. Biases in dynamic models with fixed effects. Econometrica 49, 1417-1426.

Parkinson, M., 1980. The extreme value method for estimating the variance of the rate of return. Journal of Business 53, 61-65.

Petajisto, A., September 2013. Inefficiencies in the pricing of Exchange-Traded Funds. Working paper, BlackRock, Inc.

Petersen, M., 2009. Estimating standard errors in finance panel data sets: comparing approaches. Review of Financial Studies 22, 435-480.

Saffi, P., Sigurdsson, K., 2011. Price Efficiency and Short Selling. Review of Financial Studies 24, 821-852.

Smith, V.L., Suchanek, G.L., Williams, A.W., 1988. Bubbles, crashes and endogenous expectations in experimental spot asset markets. Econometrica 56, 1119-1151.

Stambaugh, R.F., Yu, J., Yuan, Y., 2012. The short of it: Investor sentiment and anomalies. Journal of Financial Economics 104, 288-302.

Tetlock, P., May 2008. Liquidity and prediction market efficiency. Working paper, Columbia Business School, Columbia University, New York.

Theil, H., 1961. Economic Forecasts and Policy, 2nd edition. Amsterdam: North-Holland.

Thompson, S., 2011. Simple formulas for standard errors that cluster by both firm and time. Journal of Financial Economics 99, 1-10.

Tse, Y., Martinez, V., 2007. Price discovery and informational efficiency of international iShares funds. Global Finance Journal 18, 1-15.

Appendix A

The choice of the primary data period was made in order to study an era that was free of historic turbulence, namely, the 2008 Great Recession and crisis. However, the question of whether or not the relationship between efficiency and volume is different during crisis periods from non-crisis periods is an interesting one, and we consider it in this appendix.

Of the 68 ETFs in our primary analysis, only 55 have daily data prior to the Great Recession. Thus, our investigation in this appendix will focus on a balanced data set of 55 funds each with 2,643 daily observations[25] spanning the time period January 3, 2006 through July 1, 2016. The features of this data set (noted in Panel D of Table 3) are similar to those of the 68 fund data set described in Table 1. Indeed, at least 70% of each fund's holdings are large- or mid-cap equities falling within the U.S. None of the funds are leveraged, actively managed, or derivatives-based. Moreover, none of the funds hedge exchange rate risk or utilize swaps as a primary means of tracking an index. The funds' underlying indexes are weighted via market cap (50 funds), fundamentals (1 fund), equally (1 fund), and multi-factor (3 funds).

Note this longer time series includes three distinct time periods: pre-crisis (January 3, 2006 – August 31, 2008), crisis (September 1, 2008 – March 31, 2009), and post-crisis (April 1, 2009 – July 1, 2016).[26] The crisis period ranges from shortly before the collapse of Lehman Brothers (September 2008) through the market reaching its lowest point in March 2009. Each distinct sub-data set as well as the entire 55 fund data set is considered in this Appendix. Descriptive statistics for the three sub-data sets are included in Table 3. Not surprisingly, the volume, dollar volume, and volatility metrics are higher in the crisis period than the other two periods. Moreover, the mean absolute valuation spread (i.e., the magnitude of the relative deviation between price and NAV) is also larger with a higher standard deviation in the crisis period than the other periods. Consider our two primary variables of interest: efficiency (proxied via the mean absolute

[25] While the 55 funds considered in this Appendix are a subset of the 68 funds studied in the main analysis, the time period analyzed with the 55 fund data set encompasses that considered with the 68 fund data set.
[26] Although the data set considered in the main analysis of the present study covers the same date range as the post-crisis data set, they differ in the number of funds considered (68 vs. 55 funds, respectively).

valuation spread) and dollar trading volume. The median fund during the crisis period had a mean absolute valuation of 0.24% compared to 0.045% and 0.086% in the post-crisis and pre-crisis periods, respectively. The median fund's dollar volume during crisis period was 31.17 million compared to 23.76 million and 11.97 million in the post-crisis and pre-crisis periods, respectively. Given the differences in these values, it would not be surprising to discover that the relationship between these two variables is different during these different time periods.

Table 3. Descriptive statistics for the pre-crisis, crisis, and post-crisis data sets. Panels A, B, and C provide statistics for the post-crisis, crisis, and pre-crisis data sets, respectively. Each data set represents a distinct subset of the 55 fund data set described in Panel D. Mean values are calculated on a per fund basis and then averaged over all funds. Note that the standard deviation is not listed as a percentage.

A. Post-crisis	**Mean**	**St. Dev.**	**Min**	**First Quartile**	**Median**	**Third Quartile**	**Max**
Mean Market Capitalization (Millions)	5,665	8,374	246	603	2,172	7,513	41,399
Mean Volume (# of shares in Millions)	4.18	11.88	0.019	0.088	0.285	2.18	68.09
Mean Dollar Volume (Millions)	197.01	509.91	0.88	5.20	23.76	157.10	3,430.48
Mean Volatility (%)	1.29	0.0022	0.96	1.16	1.24	1.40	1.81
Mean Valuation Spread (%)	-0.0003	0.000088	-0.027	-0.0052	0.0026	0.0052	0.019
Minimum Valuation Spread (%)	-	-		-	-	-	-
Maximum Valuation Spread (%)	-	-	-	-	-	-	6.34
Mean Absolute Valuation Spread (%)	0.052	0.00025	0.030	0.038	0.045	0.053	0.17
Mean Correlation* (Dollar Volume and Volatility)	0.20 (0.05)	-	-	-	-	-	-
Number of observations per fund	1,827						

B. Crisis	**Mean**	**St. Dev.**	**Min**	**First Quartile**	**Median**	**Third Quartile**	**Max**
Mean Market Capitalization (Millions)	2,206	3,235	77	331	693	2,440	14,764
Mean Volume (# of shares in Millions)	11.25	40.53	0.053	0.17	0.94	6.59	228.78
Mean Dollar Volume (Millions)	323.03	992.87	0.68	6.27	31.17	175.16	6,472.10
Mean Volatility (%)	4.67	0.011	2.99	4.11	4.44	4.89	7.93
Mean Valuation Spread (%)	-0.0028	0.00069	-0.18	-0.043	-0.0068	0.022	0.22
Minimum Valuation Spread (%)	-	-	-9.59	-	-	-	-
Maximum Valuation Spread (%)	-	-	-	-	-	-	5.26
Mean Absolute Valuation Spread (%)	0.28	0.0011	0.15	0.19	0.24	0.35	0.53
Mean Correlation* (Dollar Volume and	0.26 (0.10)	-	-	-	-	-	-

Volatility)							
Number of observations per fund	146	-	-	-	-	-	-
C. Pre-crisis	**Mean**	**St. Dev.**	**Min**	**First Quartile**	**Median**	**Third Quartile**	**Max**
Mean Market Capitalization (Millions)	2,420	3,859	53	329	825	2,571	18,367
Mean Volume (# of shares in Millions)	4.67	19.21	0.014	0.063	0.17	1.37	131.23
Mean Dollar Volume (Millions)	212.96	841.23	0.42	3.44	11.97	64.35	5,927.30
Mean Volatility (%)	1.40	0.0034	0.89	1.17	1.30	1.55	2.29
Mean Valuation Spread (%)	-0.0086	0.00025	-0.14	-0.016	-0.0018	0.0057	0.038
Minimum Valuation Spread (%)	-	-	-3.46	-	-	-	-
Maximum Valuation Spread (%)	-	-	-	-	-	-	6.34
Mean Absolute Valuation Spread (%)	0.10	0.00042	0.064	0.077	0.086	0.12	0.31
Mean Correlation* (Dollar Volume and Volatility)	0.37 (0.16)	-	-	-	-	-	-
Number of observations per fund	670	-	-	-	-	-	-
D. Entire Time Period	**Mean**	**St. Dev.**	**Min**	**First Quartile**	**Median**	**Third Quartile**	**Max**
Mean Market Capitalization (Millions)	4,652	6,880	192	493	1,741	6,043	33,940
Mean Volume (# of shares in Millions)	4.697	14.86	0.021	0.093	0.288	2.18	82.14
Mean Dollar Volume (Millions)	208.01	615.97	1.05	5.74	19.66	152.32	4,231.44
Mean Volatility (%)	1.51	0.0028	1.07	1.33	1.42	1.60	2.16
Mean Valuation Spread (%)	-0.0025	0.00011	-0.043	-0.0091	-0.0004	0.0049	0.021
Minimum Valuation Spread (%)	-	-	-9.59	-	-	-	-
Maximum Valuation Spread (%)	-	-	-	-	-	-	6.34
Mean Absolute Valuation Spread (%)	0.077	0.00030	0.050	0.057	0.068	0.090	0.23
Expense Ratio	0.25	0.15	0.05	0.14	0.2	0.43	0.59
Mean Correlation* (Dollar Volume and Volatility)	0.20 (0.08)	-	-	-	-	-	-

Number of observations per fund	2,643	-	-	-	-	-	-

*- Reported correlation is calculated on a per fund basis and then averaged across funds. The value in parentheses represents the correlation across the entire data set.

We run the same regressions on these data sets as executed on the 68 fund data set and report the results in Table 4. For space considerations we only include the regressions corresponding to Model 5 in our main analysis. The key features of the other regressions (i.e., those analogous to Models 1 through 4), namely the significance and sign of the coefficient of the quadratic Volume term, are consistent with the reported regressions. The excluded regressions are available upon request.

Models 7, 8, and 9 were run on the post-crisis, crisis, and pre-crisis data sets, respectively, while Model 10 was run against the entire data set (January 3, 2006 through July 1, 2016).[27] Consistent with the analysis in the main text, the coefficient of the squared Volume variable is both positive and statistically significant for the post-crisis period. Note that during the crisis period the linear and quadratic Volume terms have larger (in magnitude) coefficients than in the post-crisis time period. While the coefficient of the squared Volume term is positive, neither Volume coefficient is statistically significant.[28] Interestingly, the Amihud Illiquidity measure is significant and negative during this time period suggesting that if returns are held constant, then an increase in dollar trading volume corresponds to a decrease in efficiency. It is also interesting to note that the coefficient of the Volatility variable is highly significant and large in magnitude relative to the other coefficients. This indicates that inefficiency rose as a result of increased volatility during the crisis period.

Both the linear and the quadratic Volume variables' coefficients are statistically significant during the pre-crisis period. While the linear Volume term is positive, the quadratic term is negative. The former is consistent with the behavioral theory of increasing volumes corresponding to lower efficiency, and the latter

[27] As we analyze each time period individually, we winsorize and standardize each sub-data set separately. We also winsorize and standardize the entire data set for Model 10.
[28] Note that there are only 146 days per fund in the crisis period compared to 670 and 1,827 in the pre-crisis and post-crisis time periods, respectively.

suggests there exists a point (approximately 0.94 standard deviations) beyond which increasing volumes correspond to more efficient prices. Note, however, that only approximately 8% of the observations in the pre-crisis data set satisfy this criterion. Thus, on average during this time period an increase in volume accompanied a decrease in efficiency.

Table 4. Regression Results. Results correspond to the 55 fund data set discussed in Appendix A. The two-way fixed effects model with robust standard errors clustered by fund is utilized to generate these results. Volume is the daily dollar trading volume, while volatility is the relative difference between the daily high and low prices. Models 7, 8, and 9 correspond to the post-crisis, crisis, and pre-crisis data sets, respectively. Model 10 corresponds to the entire 55 fund data set ranging from January 3, 2006 through July 1, 2016. All models are analogous to Model 5 in Table 2.

Term	**Model 7** (Post-Crisis)	**Model 8** (Crisis)	**Model 9** (Pre-Crisis)	**Model 10** (Entire period)
Volume	-0.04 (-1.56)	-0.21 (-1.15)	0.262*** (3.72)	-0.07 (-1.05)
Volume^2	0.063*** (2.71)	0.245 (1.37)	-0.14** (-2.4)	0.094** (2.39)
Volatility	-0.008 (-0.19)	1.02*** (3.91)	0.228*** (3.7)	0.134** (2.52)
Volatility^2	0.09*** (3.2)	0.564 (1.6)	0.303*** (3.77)	0.281*** (4.39)
Volume*Volatility	-0.02 (-0.74)	0.012 (0.03)	0.104* (1.68)	-0.05 (-0.6)
Market Capitalization	-0.08** (-2.11)	0.026 (0.11)	0.045 (0.92)	-0.18*** (-3.59)
Amihud Illiquidity	0.016 (1.02)	-0.57*** (-3.15)	0.049 (1.57)	-0.01 (-0.34)
R-Square	0.2649	0.3563	0.2155	0.3634
No. Observations	100,485	8,030	36,850	145,365
No. Groups/Funds	55	55	55	55
No. Days (per fund)	1,827	146	670	2,643
F Test for no fixed effects	14.92***	12.99***	10.12***	19.14***

Notes: (a) t-values are reported in parentheses; (b) *** indicates P < 0.01, ** indicates P<0.05, and * indicates P<0.1; (c) Coefficient values have been multiplied by 1,000 for exposition; (d) The reported R-Square value corresponds to the R-Square measure developed by Theil (1961); (e) To ensure multicollinearity is not an issue we compute Variance Inflation Factors, which are all less than 4.7; (f) We reject the hypothesis that the coefficients of both Volume variables are zero in Models 7/9/10 by performing a joint hypothesis test. The Wald (10.7***/18.65***/14.38***), Lagrange Multiplier (76.28***/77.12**/55.57***), and Likelihood Ratio (76.31***/77.2***/55.58***) test statistics are all significant.

Model 10 displays results of a regression run on all days included within the 55 fund data set. The descriptive characteristics of this data set are included in Panel D of Table 3. Consistent with the analysis in the main text the coefficient of the squared Volume variable is both significant and positive. As this data set includes three distinct time periods, we test the robustness of the results presented in Table 4 by interacting our independent regressors with a dummy variable indicating the specific time period (pre-crisis, crisis, or post-crisis). This essentially results in a block diagonal structure for the matrix of regressors with three blocks determined by the time periods. Thus, the coefficient of each variable is estimated for each time period simultaneously. That is, each variable appears three times in the vector of dependent variables (e.g., Post-crisis Volume, Crisis Volume, and Pre-crisis Volume). Focusing on our primary variable of interest, the quadratic Volume term, notice that the results of Model 11 are similar to those of Models 7, 8, and 9. Although the sign of the quadratic term differs for the crisis period, the term is not statistically different from zero.

Table 5. Regression Results. Results correspond to the 55 fund data set discussed in Appendix A. The two-way fixed effects model with robust standard errors clustered by fund is utilized to generate these results. A dummy variable indicating the specific time period (pre-crisis, crisis, and post-crisis) is interacted with the regressors to produce the regression results. Volume is the daily dollar trading volume, while volatility is the relative difference between the daily high and low prices.

Term	**Model 11** (Entire Period)
Post-Crisis Variables	
Volume	-0.03 (-0.95)
Volume^2	0.055** (2.5)
Volatility	-0.0088 (-0.17)
Volatility^2	0.085** (1.96)
Volume*Volatility	-0.0078 (-0.22)
Market Capitalization	-0.07* (-1.8)
Amihud Illiquidity	0.034* (1.69)
Crisis Variables	
Volume	0.03 (0.09)
Volume^2	-0.18 (-0.77)
Volatility	-0.75** (-2.12)
Volatility^2	0.797*** (3.35)
Volume*Volatility	-0.1 (-0.05)
Market Capitalization	-0.1 (-0.17)
Amihud Illiquidity	-0.65*** (-4.08)
Pre-Crisis Variables	
Volume	0.267*** (2.77)
Volume^2	-0.13** (-2.02)
Volatility	0.23*** (3.13)
Volatility^2	0.408*** (4.49)
Volume*Volatility	-0.04 (-0.5)

Market Capitalization	-0.14* (-1.76)
Amihud Illiquidity	0.017 (0.58)
R-Square	0.3941
No. Observations	436,095 (=145,365*3)
No. Groups/Funds	165 (=55*3)
No. Days (per fund)	2,643
F Test for no fixed effects	18.15***

Appendix B

We include a straightforward check to determine if our results are reasonable. Indeed, we discretize the range of volume values[29] for each data set into 50 subintervals. We then compute the average magnitude of the valuation spread variable for each subinterval. These average values are displayed in Figure 4. The majority of the volume values are close to zero where, it seems, the curves are fairly flat or slightly positively sloped. However, a fair number of values still lie in the tails. Indeed, Figure 3 indicates that volumes greater than approximately 0.91 standard deviations should, on average, correspond to less efficient prices. There are 20,247 volume values (approximately 16.3% of the total number of observations) greater than 0.91 standard deviations. While the convex nature of the relationship evidenced in Figure 3 is somewhat apparent in Figure 4, it is clear that less efficient prices correspond to extreme volumes.

Figure 4. Efficiency versus Volume. Plot of the average magnitude of the valuation spread versus the average daily dollar trading volume for the 68 fund data set. The volume is discretized into 50 subintervals, and the average magnitude of the valuation spread is calculated across all funds for each subinterval.

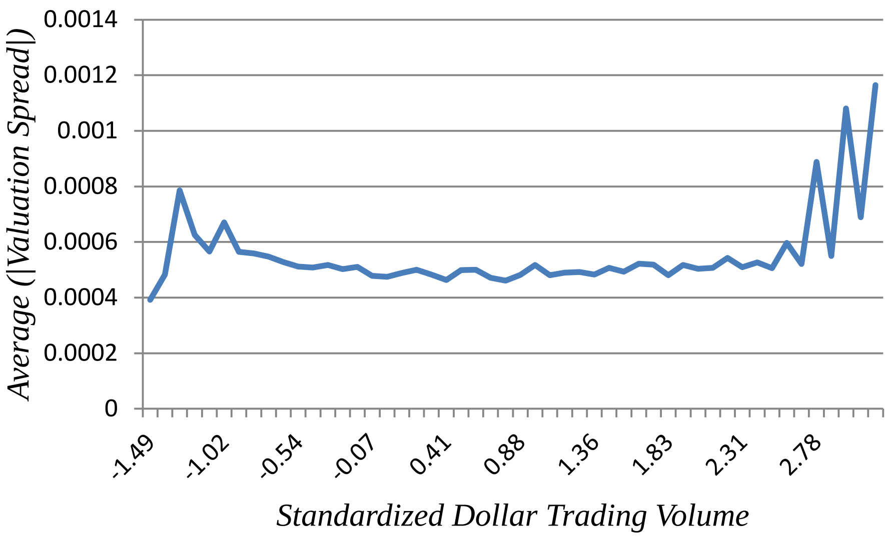

[29] The volume values are winsorized at the 5th and 95th percentiles and then standardized (by fund) by one standard deviation (to facilitate comparisons with Figure 3).