Do ETFs Increase Volatility?

Itzhak Ben-David
Fisher College of Business, The Ohio State University, and NBER

Francesco Franzoni
University of Lugano (USI) and the Swiss Finance Institute

Rabih Moussawi
Villanova School of Business, Villanova University
and WRDS, The Wharton School, University of Pennsylvania

October 2015

Abstract

Due to their low trading costs, ETFs are likely to be a catalyst for liquidity traders. The liquidity shocks can propagate to the underlying securities through the arbitrage channel. Therefore, we explore whether ETFs increase the volatility of the securities in their baskets. We exploit exogenous changes in index membership, and find that stocks with higher ETF ownership display significantly higher volatility. ETF ownership is also related to significant departures of stock prices from a random walk at the intraday and daily frequencies. Additional time-series evidence suggests that ETFs introduce a new layer of liquidity trading into the market, as opposed to just reshuffling existing liquidity shocks across securities.

Keywords: ETFs, volatility, arbitrage, fund flows
JEL Classification: G12, G14, G15

We are especially grateful to Andrew Ellul, Marco Di Maggio, Robin Greenwood (AFA discussant), and Martin Oehmke (NBER discussant). We thank Pierre Collin-Dufresne, Chris Downing, Vincent Fardeau, Thierry Foucault, Rik Frehen, Denys Glushkov, Jungsuk Han, Johan Hombert, Augustin Landier, David Mann, Rodolfo Martell, Massimo Massa, Albert Menkveld, Robert Nestor, Marco Pagano, Ludovic Phalippou, Anton Tonev, Tugkan Tuzun, Dimitri Vayanos, Scott Williamson, Hongjun Yan, and participants at seminars and conferences at the NBER Summer Institute (Asset Pricing), Toulouse School of Economics, Insead, HEC Paris, the Cambridge Judge Business School, Villanova University, the University of Lugano (USI), the 4th Paris Hedge Funds Conference, the 5th Paul Woolley Conference (London School of Economics), the 8th Csef-IGIER Symposium (Capri), the 5th Erasmus Liquidity Conference (Rotterdam), the 1st Luxembourg Asset Pricing Summit, the Center for Financial Policy Conference at the University of Maryland, Jacobs Levy's Quantitative Financial Research Conference at the Wharton School, the Geneva Conference on Liquidity and Arbitrage, the 20th Annual Conference of the Multinational Finance Society, the 7th Rothschild Caesarea Conference, the Swedish House of Finance, the FIRS conference (Toronto), and SAC Capital Advisors for helpful comments and suggestions. Ben-David acknowledges support from the Neil Klatskin Chair in Finance and Real Estate and from the Dice Center at the Fisher College of Business. An earlier version of this paper circulated under the title "ETFs, Arbitrage, and Shock Propagation."

Disclosure Statement by Itzhak Ben-David

I declare that I have no relevant or material financial interests that relate to the current study.

Itzhak Ben-David

Disclosure Statement by Francesco Franzoni

I declare that I have no relevant or material financial interests that relate to the current study.

Francesco Franzoni

Disclosure Statement by Rabih Moussawi

I declare that I have no relevant or material financial interests that relate to the current study.

Rabih Moussawi

1 Introduction

Passive investing is gaining popularity in the asset management industry. While almost all mutual funds followed active strategies in 1980, by the end of 2014, 30% of assets were in passive allocations (Morningstar, 2015). Some authors believe that the shift to passive investing is welfare improving, given the drop in intermediation fees and the improvement in portfolio diversification that index funds provide (French, 2008). Furthermore, Stambaugh (2014) argues that the rise in passive investing is symptomatic of improved market efficiency, as profit opportunities for active managers are shrinking.

Exchange traded funds (ETFs) are at the forefront of this trend in the United States, as well as globally (Cheng, Massa, and Zhang, 2015).[1] However, because of their peculiar characteristics, ETFs do not conform to the traditional view of passive funds as buy and hold investors. First, ETFs provide intraday liquidity to their investors. As a result, they attract high-frequency demand, which can translate into price pressure on the underlying securities, due to the arbitrage relation between the ETF and its basket. To compound this effect, trading strategies that were previously too expensive suddenly become affordable for retail investors thanks to ETFs. The liquidity needs of these new investors contribute an additional layer of shocks to the prices of the underlying securities. In sum, if ETFs foster additional liquidity trading, the volatility of the securities in their basket can increase.

Despite the ways in which ETFs differ from traditional passive funds, and despite their prominent role in today's investment space, there has been virtually no large sample study exploring the causal effect of ETFs on the volatility of the underlying securities' prices.[2] This paper aspires to fill this gap.

---

[1] The first ETF started trading in the United States in 1993. At the end of 2014, exchange traded passive vehicles had a market capitalization of $2 trillion .ETFs, along with other exchange traded products (ETPs), had reached $2.8 trillion of assets under management (AUM) globally as of December 2014 (BlackRock, December 2014). Also important, ETPs are involved in an increasing share of transactions in equity markets. For example, in August 2010, ETPs accounted for about 40% of all trading volume in U.S. markets.

[2] A few papers test whether ETFs have an impact on the underlying securities, but most prior research focuses on specific types of ETFs or specific events. Cheng and Madhavan (2009) and Trainor (2010) investigate whether the daily rebalancing of leveraged and inverse ETFs increases stock volatility; they find mixed evidence. Bradley and Litan (2010) voice concerns that ETFs may drain the liquidity of already illiquid stocks and commodities. Madhavan (2012) relates market fragmentation in ETF trading to the Flash Crash of 2010. Da and Shive (2015) and Israeli, Lee, and Sridharan (2015) are large sample studies, like our paper. The results in these papers support our main claim but differ in the identification strategy. A discussion follows below.

We conjecture that ETFs affect the volatility of the securities in their baskets through the arbitrage channel. Specifically, the demand shocks in the ETF market put pressure on ETF prices. If the ETF price deviates from the net asset value (NAV) of the portfolio holdings, arbitrageurs trade the underlying securities in the same direction as the initial shock. Thus, the underlying securities inherit the shocks that occur in the ETF market. Consequently, their volatility increases.

This effect is similar to that of mutual fund flows on the prices of the portfolio holdings (Coval and Stafford, 2007; Lou, 2012; Cella, Ellul, and Giannetti, 2013; Hombert and Thesmar, 2014; and in the general context of large trades: Ellul, 2006). However, transactions in ETFs, as well as arbitrage activity, take place continuously throughout the day. Thus, ETFs are potentially a more rapid conduit for the propagation of demand shocks than other managed portfolios.

With regard to the counterfactual, one could conjecture that in the absence of ETFs, the same liquidity shocks could hit directly the underlying securities. If this were the counterfactual, the net effect of ETFs on volatility would be zero. Hence, for ETFs to have a positive effect on the basket stocks' volatility, a necessary condition is that ETFs foster additional liquidity trading.

To verify this condition, our empirical analysis starts by showing that ETFs attract short-term investors. ETFs are, on average, significantly more liquid than the basket of underlying securities in terms of bid-ask spread, price impact, and turnover. Consistent with theories positing that short-horizon clienteles self-select into assets with lower trading costs (Amihud and Mendelson, 1986), we find that the institutions holding ETFs have a significantly shorter horizon than those holding the underlying securities. We take this evidence as satisfying the necessary condition for the argument that ETFs attract a new layer of layer of liquidity trading, which would not otherwise hit the securities in their baskets.

To test the effect of ETFs on volatility, we start by providing evidence on the correlation between daily stock volatility and ETF ownership, using OLS regressions. ETF ownership is the total fraction of a stock's capitalization that is held by ETFs. For the universe of S&P 500 stocks, we find that a one-standard-deviation increase in ETF ownership is associated with a statistically significant increase in daily volatility that ranges between 7% and 13% of a standard deviation. The effect is, therefore, economically significant.

Although the OLS regressions control for observable stock characteristics and include stock fixed effects, there is a concrete possibility that ETF ownership is an endogenous variable.[3] To address this concern, we rely on the natural experiment provided by the annual reconstitution of the Russell indexes (Chang, Hong, and Liskovich, 2015). We closely follow the approach in Appel, Gormley, and Keim (2015), who run an instrumental variable (IV) regression exploiting the mechanical rule allocating stocks between the Russell 1000 (top 1000 stocks by market capitalization) and the Russell 2000 (next 2000 stocks by market capitalization) indexes in June of each year. Due to the large difference in index weights, the top stocks in the Russell 2000 receive significantly larger amounts of passive money than do the bottom stocks in the Russell 1000. The identifying assumption is that a switch to either index generates exogenous variation in ETF ownership, after controlling for market capitalization, which determines index assignment. Hence, we use the index switch as an instrument to identify the effect of ETF ownership on volatility. In addition, we control for lagged volatility, which is positively correlated with the probability of switching and, therefore, could cause us to find a spurious effect.

This methodology confirms that the impact of ETF ownership on volatility is positive and statistically significant. The IV estimates exceed those from the OLS regressions, averaging around 32% of a standard deviation, which suggests a negative omitted variable bias in the OLS specifications. To make sense of the larger IV coefficients, we also note that the IV slopes measure a local average treatment effect (LATE), which is the average effect across the units in the sample that receive treatment (the 'index switchers', in our context). Because the treated stocks experience a drastic change in the extent to which they are owned by ETFs, we should not expect such a large effect for the average stock in the sample, experiencing only marginal variation in ETF ownership.

To corroborate the IV analysis, we assess the validity of the exclusion restriction. In particular, we ask whether switching index can affect stock volatility in other ways than through a change in ETF ownership. We support the validity of our identification by showing that

[3] For example, new ETFs might track investment themes that have gained popularity among investors. The stocks in these segments of the market might be more volatile because of the attention they already receive, not because ETFs attract liquidity trading. This mechanism would generate a positive bias in the OLS estimates. Alternatively, higher ETF ownership may signal companies that belong to multiple indexes, which have less volatile stocks because they are more established companies. This fact would lead to a negative omitted variable bias.

switching index has a stronger impact in months when ETF ownership is larger. Furthermore, the effect of interest is larger for stocks with a higher ratio of ETF ownership to other mutual fund ownership (either index or active funds). This finding suggests that ETFs, as opposed to other institutional investors, are the drivers of the effects that we measure.

Rather than increased liquidity trading, the observed effect could reflect higher investor attention, which makes prices react more strongly to information (e.g., cash flow news), as shown by Andrei and Hasler (2015). To separate these two alternatives, we rely on the premise that liquidity shocks subsequently revert, while cash flow news causes permanent price changes. Then, we construct the absolute difference from one of intraday and daily variance ratios of stock returns. We find that the deviation in the variance ratios of stock returns from unity increases with ETF ownership, suggesting a link between the presence of ETFs and greater deviations from a random walk for the underlying security prices. Moreover, we estimate predictive regressions of stock returns as a function of ETF flows at the stock level and daily frequency. We find that almost half of the contemporaneous positive impact of flows reverts over the next 20 days, confirming that the presence of ETFs is significantly related to the mean-reverting component of prices. This evidence suggests that at least part of the measured increase in volatility reflects transitory shocks, consistent with the effect of liquidity trading.

To provide evidence on the arbitrage channel of liquidity shock propagation, we ask whether the impact of ETF arbitrage activity on stock prices is weaker for securities that display higher arbitrage costs. Indeed, we find that a proxy for arbitrage activity (the difference between the ETF price and the NAV, commonly labeled 'mispricing') has a smaller effect on volatility for stocks in the top half of the distribution of the bid-ask spread and of share-lending fees, i.e., for stocks with higher limits to arbitrage. Moreover, strongly supporting the arbitrage channel, the coefficient on share-lending fees is significant only in the subsample for which the arbitrage trades involve shorting the stock (that is, when mispricing is negative).

Our evidence corroborates the view that ETFs have a positive causal effect on the arrival of liquidity shocks in the prices of the stocks that they own. We find this result interesting, as it establishes a novel dimension along which institutional trading can affect prices. However, because of the cross-sectional identification, the evidence could signal a migration of liquidity traders from securities with low ETF ownership to those with high ownership as well as an overall increase in liquidity trading in the stock market. In the last part of the analysis, we try to separate these two scenarios in a time-series setting. We provide suggestive evidence that the average share of ETF ownership in the market is positively associated with average stock volatility. The effects persist when we control for aggregate ownership by other mutual funds and include a time trend, which helps to absorb omitted factors. Although the time-series identification does not allow us to make unambiguous causal inferences, this finding is fully consistent with the view that ETFs attract a new layer of liquidity trading that would not otherwise reach the stock market if ETFs did not exist.

Our study relates to several strands of the literature. There is mounting evidence on the role of institutions in impounding non-fundamental demand shocks into asset prices as a result of flows from their investors (Brunnermeier and Nagel, 2004; Coval and Stafford, 2007; Lou, 2012; Cella, Ellul, and Giannetti, 2013; Vayanos and Woolley, 2013; Hombert and Thesmar, 2014). We highlight a previously unexplored channel: arbitrage activity between ETFs and the underlying baskets. By documenting an effect of institutional ownership on volatility, we join a body of work that focuses on the impact of institutions on the second moments of returns (Greenwood and Thesmar, 2011; Anton and Polk, 2014; Lou and Polk, 2014, Basak and Pavlova, 2013a, 2013b).[4]

A few recent studies also focus on the effect of ETFs on second moments, but use different empirical approaches. Da and Shive (2015) find that ETF ownership is associated with higher comovement of the underlying securities. This idea is subsumed by our results: ETFs impound the same shocks into all the stocks in their basket and, therefore, make them comove. Simply showing a relation between ETF ownership and the second moment of stock returns is not by itself sufficient because of the potential endogeneity of ETF ownership. Hence, the empirical challenge that we take on in this paper is to show that ETFs indeed *cause* price variation in the underlying stocks. If the causal link exists, then comovement is a by-product of this effect. Similarly, the recent paper by Israeli, Lee, and Sridharan (2015) makes the claim that increased ETF ownership can lead to higher trading costs and lower benefits from information

[4] Our paper indirectly relates to the rich literature on the effect of indexing (Shleifer, 1986; Barberis, Shleifer, and Wurgler, 2005; Greenwood, 2005; Wurgler, 2011; Chang, Hong, and Liskovich, 2015). The trigger for the effect that we measure is trading in ETFs, as opposed to index reconstitution. Index membership matters only in defining the stocks that are affected by ETFs.

acquisition for the basket securities, a combination that results in less informative stock prices. While their results supports our conclusions, once again, our analysis differs in its effort to identify truly exogenous variation in ETF ownership. Leippold, Su, Ziegler (2015) build a model in which the impact of ETFs on return correlations exceeds the effect of futures and they provide consistent empirical evidence in time-series tests. Finally, Malamoud (2015) develops a theory that accurately describes the functioning of the ETF market. In his model, ETFs can affect volatility through the liquidity shock transmission channel that we describe in our paper, but also through the time-varying risk premia that investors require as compensation for exposure to these shocks.

Another theme in the literature to which our study relates is the long-running debate on the effect of derivatives on the quality of the underlying securities' prices. On one side of the debate is the concern that liquidity shocks in derivatives markets can trickle down to the cash market, making prices less informative. For example, Stein (1987) shows that imperfectly informed speculators in futures markets can destabilize spot prices. Among the supporters of the alternative view, Grossman (1989) argues that the existence of futures provides additional market-making power to absorb the impact of liquidity shocks. As a result, volatility in the spot market is reduced (see also Danthine, 1978; Turnovsky, 1983).[5] We contribute to this literature by providing systematic evidence from ETFs, an asset class that has a similar flavor to futures but is potentially more attractive to liquidity traders due to the lack of margin requirements and the absence of roll-over risk. In December 2014, the assets under management (AUM) in ETFs tracking the S&P 500 surpassed the open interest in futures on the same index, suggesting that ETFs are becoming the security of choice to achieve exposure to the stock market (Amery, 2015).

The paper proceeds as follows. Section 2 provides institutional details on ETFs and develops the testable hypotheses, and Section 3 describes the data. Section 4 presents the main

[5] Earlier studies that examine the impact of derivatives on volatility focused on futures. The proposed economic channel in this literature is the same as the one that we test in this paper. In a cross-sectional analysis, Bessembinder and Seguin (1992) find that high trading volume in the futures market is associated with lower equity volatility. However, consistent with the idea that non-fundamental shocks in the futures market are passed down to the equity market, they find that unexpected futures trading volume is positively correlated with equity volatility. Chang, Cheng, and Pinegar (1999) document that the introduction of futures trading increased the volatility of stocks in the Nikkei index stocks. Roll, Schwartz, and Subrahmanyam (2007) find evidence of Granger causality between prices in the futures and equity markets: price shocks are transmitted from the futures market to the equity market and vice versa.

evidence of the effect of ETF ownership on stock volatility and variance ratios. In Section 5, we provide evidence on the role of arbitrage in driving the main effects. Section 6 addresses the question of whether ETFs attract a "new layer" of volatility to the stock market. Section 7 concludes.

2 Institutional Details and Hypotheses Development

2.1 Mechanics of Arbitrage

ETFs are investment companies whose objective is to replicate the performance of an index, very much like index mutual funds. Unlike index funds, ETFs are listed on an exchange and trade throughout the day. ETFs were first introduced in the late 1980s and became popular with the issuance in January 1993 of the Standard & Poor's Depository Receipts (SPDR) tracking the S&P 500 (ticker: SPY). SPY is currently the largest ETF in the world, with about $181 billion of assets (December 2014). The number of ETFs subsequently exploded to more than 1,600 (including other exchange products, i.e., ETPs) by the end of 2014, spanning various asset classes.

To illustrate the growing importance of ETFs in the ownership of common stocks, we present descriptive statistics for the S&P 500 and Russell 3000 universes in Table 1, Panel A. ETF ownership of individual stocks has increased almost twenty times over our sample period. For S&P 500 stocks, the average fraction of a stock's capitalization held by ETFs has risen from 0.22% in 2000 to 3.90% in 2012. The table shows that the number of ETFs holding the average stock in the S&P500 universe grew from about 2 to about 49 during the same period. The total AUM for ETFs holding S&P 500 stocks was, in 2012, about $414 billion. The statistics for the Russell 3000 universe paint a similar picture. The growth of ETFs appears as even more drastic when compared to trend for index funds (Panel B), which are the closest substitute, and active funds (Panel C). Index fund ownership of stocks, although on average higher, increased only by a factor of two in this sample. At the end of the period, while still lagging behind, ETFs had closed much of the initial gap (ownership of 3.90% vs. 4.69%, for the average Russell 3000 stock). The ownership by active funds, while on average higher than for the passive vehicles, remained more or less stable.

Futures are similar to ETFs in that they track an index. Unlike futures, ETFs do not involve a rollover of the expiring contract. Rollover can erode performance for investors with horizons spanning beyond the short maturity of a futures contract. According to BlackRock, the annualized rollover cost of a futures position in large cap stocks (S&P 500, Euro Stoxx 50, FTSE 100) ranges from 0.9% to 1.4%. The total expense ratio for an ETF on the same indexes can be as low as 0.05% (e.g., the Vanguard S&P 500 ETF). Hence, ETFs provide a more cost efficient way to track an index, especially to investors with uncertain trading horizons. This cost efficiency can explain the fact that in December 2014 the assets in ETFs tracking the S&P 500 surpassed the open interest in futures contracts for the first time (see Amery, 2015).

In our analysis, we focus on ETFs that are listed on U.S. exchanges and whose baskets contain U.S. stocks. Further, we restrict to "plain vanilla" products that engage in physical replication, that is, they hold the securities of the basket that they aim to track. We omit from our sample leveraged and inverse leveraged ETFs that use derivatives to deliver the performance of the index, which represent at most 2% of the assets in the sector according to BlackRock (December 2014). These more complex products are studied by Cheng and Madhavan (2009), among others. We also omit active ETFs that are still below 1% of AUM in the sector.

ETFs are traded in the secondary market by retail and institutional investors, in a fashion similar to closed-end funds. However, unlike closed-end funds, new ETF shares can be created and redeemed. Because the price of ETF shares is determined by the demand and supply in the secondary market, it can diverge from the value of the underlying securities (the NAV).[6] Some institutional investors called "authorized participants" (APs), who are dealers that have signed an agreement with the ETF provider, can trade bundles of ETF shares (called "creation units," typically 50,000 shares) with the ETF sponsor. An AP can create new ETF shares by transferring the securities underlying the ETF to the ETF sponsor. These transactions constitute the primary market for ETFs. Similarly, the AP can redeem ETF shares and receive the underlying securities in exchange. For some funds, ETF shares can be created and redeemed in cash.[7]

---

[6] Unlike premia and discounts in closed-end funds (e.g., Lee, Shleifer, and Thaler, 1991; Pontiff, 1996), price divergence between the ETF and the NAV can be more easily arbitraged away thanks to the possibility of continuously creating and redeeming ETF shares. As a result, ETF premia/discounts are orders of magnitude smaller than for closed-end funds.

[7] Creation and redemption in cash is especially common with ETFs on foreign assets or for illiquid assets, e.g., fixed income ETFs.

To illustrate the arbitrage process through the creation/redemption of ETF shares, we start from an ETF trading at a premium relative to the NAV. In this case, APs have an incentive to buy the underlying securities, submit them to the ETF sponsor, and ask for newly created ETF shares in exchange. Then the AP sells the new supply of ETF shares on the secondary market. This process puts downward pressure on the ETF price and potentially leads to an increase in the NAV, reducing the premium. In the case of a discount, APs buy ETF units in the market and redeem them for the basket of underlying securities from the ETF sponsor. Then the APs can sell the securities in the market. This generates positive price pressure on the ETF and possibly negative pressure on the NAV, which reduces the discount.

Creating/redeeming ETF shares has limited costs in most cases, especially for equity-focused funds. These costs include the fixed creation/redemption fee plus the costs of trading the underlying securities. Petajisto (2013) describes the fixed creation/redemption costs as ranging in absolute terms from $500 to $3,000 per creation/redemption transaction, irrespective of the number of units involved.

ETF arbitrage also takes place continuously throughout the day as a result of hedge fund and high-frequency trader activity.[8] These investors do not need to engage in primary market trades. On the secondary market, they can buy the inexpensive asset and short sell the more expensive one between the ETF and the basket of underlying securities. They hold the positions until prices converge, at which point they close down the positions to realize the profit. ETF sponsors facilitate arbitrage activity by disseminating NAV values at a 15-second frequency throughout the trading day. They do so because the smooth functioning of arbitrage is what brings about the low tracking error of these instruments. As a result of the low trading costs and availability of information, arbitraging ETFs against the NAV has become a very popular trading strategy in recent years. According to some industry participants, statistical arbitrage accounts for 50% of the volume in the S&P 500 SPDR, which is the most traded security in the United States with $26 billion average daily volume (last 3 months of 2014).

[8] To be precise, although these trading strategies involve claims on the same cash flows, they may not be arbitrages in the strict sense because they can involve some amount of risk. In particular, market frictions can introduce uncertainty into the process (e.g., execution may not be immediate, shares may not be available for short selling, or mispricing can persist for longer than the arbitrageurs' planned horizon for the trade). In the remainder of the paper, when referring to arbitrage, we imply the broader definition of "risky arbitrage."

Both the creation/redemption activity by APs, which takes place at the daily frequency, and the intraday arbitrage by high-frequency traders have the potential to move the prices of the underlying securities. We provide evidence consistent with the effects of ETF arbitrage playing out at both the daily and the intraday frequencies.

These institutional details, with some modifications, also apply to synthetic ETFs, which replicate the performance of the index using total return swaps and other derivatives, and for which creation and redemption accour in cash. The secondary market arbitrage still involves transactions in the underlying securities. Thus, the potential for the propagation of demand shocks from the ETF market to the underlying securities via arbitrage is also present among synthetic ETFs. Similarly, the arbitrage process is an inherent characteristic of all types of ETFs, beyond the equity-based ones that are studied here. Hence, one should expect the effects that we describe in this paper to hold for all types of underlying assets.

2.2 Hypotheses Development

The main testable hypothesis of the paper is that ETFs are a catalyst for liquidity trading and that the ensuing price pressure propagates to the underlying securities via arbitrage. According to this hypothesis, stocks with higher ETF ownership should display higher volatility, all else being equal.

The mechanism for the propagation of demand shocks from the institutional portfolios to the underlying securities is similar in nature to the demand pressure originating from institutional investors that experience flows (e.g., Coval and Stafford, 2007; Lou, 2012; Vayanos and Woolley, 2013). Analogously to the effect of mutual fund flows on stock prices, the demand for ETF shares translates into price pressure on the underlying securities because of the arbitrage activity. What makes ETFs special is that ETFs allow investors to access the market continuously and at a low trading cost. Hence, ETFs attract potentially more liquidity trading than standard mutual funds do.

To illustrate the channel for the propagation of liquidity shocks, we imagine a situation in which the ETF price and the NAV of its portfolio are aligned at the level of the fundamental value, as in Figure 1a. Then, a liquidity shock, i.e., one that is unrelated to information about future cash flows, hits the ETF market. Arbitrageurs absorb the liquidity demand by shorting the

ETF. Because they are risk averse, arbitrageurs require compensation for the (negative) inventory in the ETF that they are taking on. Hence, the ETF price has to rise (Figure 1b). At the same time, to hedge their short ETF position, arbitrageurs take a long position in the securities in the ETF basket. Again, to compensate the arbitrageurs for the risk that they take, the prices of the basket securities have to rise, as in Figure 1c. Eventually, when other sources of liquidity materialize, prices revert to fundamentals (Figure 1d).

To make a concrete example of this channel, consider hedge funds' trading practices in ETFs. Some hedge funds that specialize in high-frequency strategies carry out arbitrage trades of ETFs against the underlying baskets. These trades conform to the mechanism described in Figure 1. In addition, hedge funds can impound mispricing indirectly through their use of ETFs in statistical arbitrage. Suppose hedge funds short-sell an overpriced stock and hedge the industry risk by going long in the corresponding sector ETF. This trade puts upward pressure on the ETF price (as in Figure 1b). Then, cross-market arbitrageurs transfer the price pressure to the securities in the ETF basket (as in Figure 1c). This argument suggests that ETFs can propagate mispricing to the underlying securities not only because they are traded directly by uninformed investors, but also because they are traded indirectly through their participation in long-short strategies that involve other mispriced securities.

We note that the sequence of events in Figure 1 generates predictions that partly overlap with those from an alternative scenario positing gradual price discovery after a fundamental shock, defined as news about future cash flows. If price discovery occurs first in the ETF market, ETF prices adjust immediately to the new information, while the underlying securities' prices remain temporarily fixed ("stale pricing"). We illustrate this scenario in Figure 2. The initial equilibrium (Figure 2a) is perturbed by a shock to the fundamental value of the ETF components (Figure 2b). The ETF price moves first because of price discovery (Figure 2c), and the prices of the underlying securities move with a delay because of stale pricing (Figure 2d). In this alternative situation, ETFs improve price discovery and the arbitrage activity facilitates the adjustment of prices to fundamentals. As a result, there could be a positive link between ETF ownership and "good" volatility (i.e., fundamental volatility). To disentangle the two scenarios, it is not sufficient to show that stocks with higher ETF ownership display higher volatility. We also need to show that ETFs are associated with increased mean reversion in prices, which follows from the propagation of liquidity shocks (as per Figure 1).

The testable hypothesis spelled out above actually posits an increase in liquidity trading in the underlying securities because of ETF ownership. This conjecture compels us to specify the counterfactual. If the same amount of liquidity traders merely shifted from trading a given stock to trading the ETFs holding that security, the volatility in the stock's price would not increase. To observe an increase in volatility, one also needs the additional assumption that liquidity traders prefer ETFs to stocks as their habitat. Under this assumption, the creation of ETFs entails a migration of liquidity shocks from stocks with low ETF ownership to stocks with high ETF ownership. The sort of liquidity traders that affect volatility at high frequency are likely to be short-horizon investors.

Furthermore, the paper addresses an additional conjecture. Rather than simply redistributing existing liquidity shocks from securities with low ETF ownership to those with high ownership, this new asset class could also cause a new layer of liquidity trading, and therefore of volatility, to materialize in the stock market. The effect could follow from the enhanced trading opportunities that come with ETFs. For example, relative to standard mutual funds (including index funds), ETFs allow intraday trading and shorting at a low cost in a wide variety of market themes. Hence, liquidity traders can gain access to previously unavailable opportunities to express their views. The possibility that ETFs attract a new layer of volatility is our second testable hypothesis.

To test this conjecture, we look for a significant positive relation between the average stock volatility in the market and the average ETF ownership of stocks in the time series. To measure the effect of interest more closely, we will need to control for other contemporaneous developments in the market. Admittedly, the time-series tests can never completely rule out omitted factors, so that the evidence in favor of this second testable hypothesis is only suggestive.

2.3 ETFs vs. Stocks: Liquidity, Investor Types, and Trading Horizon

The main testable hypothesis of the paper posits that ETFs increase the volatility of the security in their baskets because of the propagation of liquidity shocks. This conjecture requires the assumption that, in a world without ETFs (i.e. the counterfactual world), liquidity trading in the underlying securities would occur to a lower extent. To verify this conjecture, we start by contrasting the liquidity of ETFs to that of their portfolio constituents. Ultimately, we want to verify whether ETFs are more appealing to high turnover investors thanks to their higher liquidity.

The bid-ask spreads on ETFs are on average low, potentially due to a lack of information asymmetry. For a few representative ETFs, Madhavan and Sobczyk (2014) provide evidence that the bid-ask spread is lower than the average spread in the corresponding basket. These authors put forward a convincing argument for the higher liquidity of ETFs. Investors with stock-level private information are more likely to trade individual securities, and market makers impose higher bid-ask spreads to overcome adverse selection. In contrast, investors who place uninformed directional bets or trade for hedging purposes are more likely to trade entire baskets, such as ETFs. As a result, ETF spreads are less likely to contain an adverse selection premium.

We carry out a similar analysis in our sample covering all U.S.-equity-based ETFs listed on U.S. exchanges (660 different products; see Section 3 for details on sample construction). Table 2, Panel A, presents systematic evidence on the difference in liquidity between ETFs and the underlying portfolios along three dimensions: the percentage bid-ask spread, the Amihud (2002) measure of price impact, and daily turnover. For all the ETFs in our sample, we compute the average of each liquidity measure across all the stocks in the basket in a given quarter. Then, to replicate the strategy of an investor that allocates funds to all ETFs according to their market capitalization, we take the value-weighted mean of these measures across all ETFs in a given quarter. The table reports the time-series average of these means in the 52 quarters of the sample (2000:Q1–2012:Q4), along with the results of tests for the statistical significance of their difference. Along all three dimensions, the average ETF is significantly more liquid than its basket stocks. The bid-ask spread is lower by about 20 bps. Price impact, as measured by the Amihud (2002) ratio, is also significantly lower for ETFs. Finally, ETF turnover is about 8.3% higher.

A corollary of the conjecture that ETFs are more liquid than the underlying baskets is that ETF investors should display higher turnover. This prediction stems, for example, from Amihud and Mendelson's (1986) clientele effect, whereby short-horizon investors choose to trade in more liquid securities. The bottom of Table 2, Panel A supports this conjecture. We compare ETFs to their underlying baskets in terms of two measures of the investor churn ratio. The first measure comes from Cella, Ellul, and Giannetti (2013), who compute an institutional investor–level churn ratio as the sum of quarterly absolute changes in dollar holdings over average AUM, using institutional holdings in the 13-F filings. This measure is then averaged across institutions at the stock level using the fraction of a company held by each institution as a weight. The second measure differs only in that the investor-level churn ratio is computed as the minimum between the absolute value of buys and sells, divided by prior quarter holdings.[9] In Table 2, Panel A, we note that the average ETF has a significantly higher investor churn ratio than its underling basket by about 6.7% per quarter for the first measure and by 2.9% for the second measure. These differences are economically significant, as the average churn ratio for the basket of stocks is 24% and 12.5%, respectively, for the two measures. This evidence confirms that ETFs, rather than stocks, are the preferred habitat of investors with a short trading horizon.

Next, we compare classes of ownership for all ETFs in our sample and all common stocks in CRSP. Panel B of Table 2 uses Thomson-Reuters' classification of institutional owners filing 13-F forms. Internet Appendix Table A2 contains detailed definitions of the various investor classes. The panel reports shares held by each group as a fraction of total shares outstanding. The first striking fact is that the institutional ownership of ETFs is by far smaller (at 47.4% on average) than the institutional ownership of stocks (at 62.1% on average) throughout the entire sample period.[10] One can roughly infer retail ownership as one minus the fraction of institutional ownership.[11] Based on Stambaugh's (2014) argument that uninformed traders are mostly present among retail investors, this evidence suggests a higher density of liquidity traders among ETF clients. Second, research firms, which include broker-dealers, have greater ETF

[9] Buys (sells) are the sum of the dollar value of the quarterly positive (negative) changes in stock holdings for a given institutional portfolio, as reported in SEC 13-F form. Values are computed using beginning-of-quarter prices.
[10] It is worth noting that part of the direct institutional ownership in stocks is through ETFs.
[11] This way of computing retail ownership is an approximation due to two elements. First, small institutions managing less than $100 million and professional investors managing solely their own proprietary accounts are not required to file a 13-F form. Second, reported shares include shares that are short-sold. Because we compute ownership as a fraction of shares outstanding, total institutional ownership for a firm could exceed one. This issue is especially relevant for ETFs, as short interest for some ETFs can be very large (even exceeding the total shares outstanding). However, expressing ownership as a fraction of the shares outstanding plus shares short sold would give an even higher estimate of retail ownership: 1 – Institutional Shares / (Shares Outstanding + Shares Sold Short) > 1 – Institutional Shares / Shares Outstanding.

ownership (5.8%) than stock ownership (0.6%). This class of owners, along with hedge funds, corresponds to ETF arbitrageurs and market makers (including APs).[12]

In sum, Table 2 reveals that ETFs are more liquid than the underlying securities. Moreover, ETFs attract investors with higher turnover. From this analysis, we can conclude that a necessary condition for ETFs to have a positive net effect on stock volatility is verified. Specifically, ETFs have the potential to attract a new layer of demand shocks, which would not otherwise hit the underlying securities.

3 Data

We use Center for Research in Security Prices (CRSP), Compustat, Bloomberg, and OptionMetrics data to identify ETFs traded on the major U.S. exchanges and to extract returns, prices, and shares outstanding. We first draw information from CRSP for all securities that have a historical share code of 73, which exclusively defines ETFs in this data set. We then screen all U.S.-traded securities in the Compustat XpressFeed and OptionMetrics data, identifying ETFs using the security-type variables, and merge this sample with the CRSP ETF sample.[13] Our initial sample consists of 1,673 ETFs between 1993 and 2012.

Because very few ETFs traded during the 1990s, we restrict the sample to the 2000–2012 period. We further restrict our sample to ETFs that invest primarily in U.S. domestic equity stocks, because they are not plagued with stale pricing issues (global equity or bond ETFs) or other issues affecting the ease of replication (short bias, volatility, and futures-based ETFs, commodities, etc.). Therefore, we exclude leveraged ETFs, short equity ETFs, and all ETFs that invest in international or non-equity securities, or in futures and physical commodities. We also eliminate active and long/short ETFs as well as dedicated short bias funds and focus on plain vanilla U.S. domestic long equity ETFs. To do so, we use both the CRSP Style Codes and Lipper prospectus objective codes in the CRSP Mutual Fund Database and restrict our sample to the fund objectives that span broad-based U.S. Diversified Equity funds and U.S. sector ETFs that

[12] From Table 2, Panel B, we also note that investment companies, which mostly coincide with mutual funds, have minimal investments (1.7%) in ETFs compared to stocks (16.3%). Mutual funds only use ETFs to temporarily park their cash and avoid accumulating tracking errors with respect to their benchmark.

[13] Note that in 2011, the time of the first draft of this paper, the CRSP-Compustat merged product did not correctly link ETF securities in the CRSP and Compustat universes. For this reason, we use historical CUSIP and ticker information to map securities in the CRSP, Compustat, and OptionMetrics databases.

invest in equities (e.g., U.S. companies investing in oil and natural resources vs. those investing in oil or commodity futures).[14] We end up with 660 distinct equity ETF securities.

We obtain quarterly holdings information using the Thomson-Reuters Mutual Fund holdings database. ETFs are subject to Investment Company Act reporting requirements, and similar to mutual funds, they have to disclose their portfolio holdings at the end of each fiscal quarter.[15] We use these data to align ETF ownership every month using the most recently reported holdings. Then, for every stock, we sum the total ownership by various ETFs to construct our ETF ownership measure. We also use the Thomson-Reuters Mutual Fund holdings database to compute the ownership by mutual funds other than ETFs, that is, index funds and active funds. To do so, we use the index fund flag in the CRSP Mutual Fund Database, and merge it with Thomson-Reuters holdings data using WRDS MFLinks. Similar to how ETF ownership is calculated, we compute monthly index and active fund ownership by using the most recently reported holdings.

We use total shares outstanding at day-end to compute the daily market capitalization of each ETF and to measure the net share creations/redemptions (i.e., flows) for each ETF daily. Bloomberg is our primary source for shares outstanding and the related net flow measures. We use Compustat and OptionMetrics to complement the ETF series when there are gaps in the Bloomberg data.[16]

[14] The Lipper Asset Code is not sufficient to accurately filter for U.S. domestic equity funds, because the Equity Funds code comprises a wide array of U.S. and global funds that implement various direct investment or alternative/inverse strategies. Instead, we use the Lipper Objective Code classifications that are assigned by Lipper to a specific population of equity funds and that are based on how the fund invests by looking at the actual holdings of the fund to determine market cap and style versus a benchmark. We restrict our sample to the following Lipper Objective Codes: Broad Based U.S. Equity: S&P 500 Index Objective Funds, Mid-Cap Funds, Small-Cap Funds, Micro-Cap Funds, Capital Appreciation Funds, Growth Funds, Growth and Income Funds, and Equity Income Funds (CA, EI, G, GI, MC, MR, SG, and SP, respectively). We also include Sector Funds that invest in U.S. companies: Basic Materials, Consumer Goods, Consumer Services, Financial Services, Health/Biotechnology, Industrials, Natural Resources, Real Estate, Science and Technology, Telecommunications, Specialty/Miscellaneous Funds, and Utilities (BM, CG, CS, FS, H, ID, NR, RE, TK, TL, S, and UT, respectively).

[15] We find that until mid-2010, Thomson Mutual Fund Ownership data are more reliable and more complete than CRSP Mutual Fund Holdings.

[16] Because CRSP shares outstanding figures are stale during the month, we assess the accuracy of three databases that provide data on shares outstanding at a daily frequency: Bloomberg, Compustat, and OptionMetrics. Thanks to direct validation by BlackRock, we concluded that Bloomberg is more accurate and timely in updating ETF shares outstanding when newly created or redeemed shares are cleared with the Depository Trust & Clearing Corporation (DTCC). On many occasions, Compustat and OptionMetrics shares-outstanding data lag Bloomberg by up to three and sometimes as many as five days.

As a dependent variable of our main tests, we compute daily stock volatility at the monthly frequency as the standard deviation of daily returns within a month. For some of our tests, we compute volatility at a daily frequency using second-by-second data from the Trade and Quote database (TAQ). For each stock, we compute a return in each second during the day using the last trade price at the end of each second during market hours (between 9:30 am and 4:00 pm). Then, we compute the standard deviation of those second-by-second returns as the intraday volatility measure.[17]

We extract stock lending fees from the Markit Securities Finance (formerly Data Explorers) database. We use the variable that reports the average lending fee over the prior seven days.

Table 3 reports summary statistics for the variables that we use in the analysis. Panel A presents summary statistics for the monthly-stock-level sample of our main regressions; Panel B reports the correlations for the same variables. Panel C presents summary statistics for the variables that are used in the return regressions at the daily frequency. Panel D presents statistics for the stock-day-level sample. We further describe these variables in later sections and provide definitions in Appendix Table A1.

4 The Effect of ETF Ownership on Volatility

4.1 ETF Ownership and Volatility: OLS Regressions

We start by asking whether ETF ownership correlates with higher volatility of the underlying securities. To this purpose, we exploit variation in ETF ownership across stocks and over time in a simple OLS framework.

ETF ownership of stock *i* in month *t* is defined as the sum of the dollar value of holdings by all ETFs investing in the stock, divided by the stock's capitalization at the end of the month:

$$ETF\ ownership_{i,t} = \frac{\sum_{j=1}^{J} w_{i,j,t} AUM_{j,t}}{Mkt\ Cap_{i,t}}, \tag{1}$$

[17] We also compute intraday volatility using intraday returns based on National Best Bid and Offer (NBBO) midpoints; the results are similar.

where J is the set of ETFs holding stock i; $w_{i,j,t}$ is the weight of the stock in the portfolio of ETF j, which is extracted from the most recent quarterly report; and $AUM_{j,t}$ is the assets under management (AUM) of ETF j at the end of the month.

Based on Equation (1), variation in ETF ownership comes from three sources. First, stocks are typically part of multiple indices (e.g., a stock might be part of the S&P 500, the S&P 500 Value, the Russell 3000, and a sector index). Second, there is variation in ETFs' AUM over time and across products. Third, there is variation in weighting schemes. For example, the S&P 500 and the Russell 2000 are capitalization-weighted, but the Dow Jones is price-weighted; also, our sample contains 17 products that explicitly mention equal-weighting in their names.

The three sources of variation in ETF ownership present different types of potentially spurious correlation with the dependent variable of interest, stock volatility. The portfolio weights follow the weighting scheme of the index. One caveat is that if the weights do not grow at the same rate as the market capitalization at the denominator (e.g., for equal-weighted indexes), there could be a spurious link between ETF ownership and volatility resulting from the correlation between stock size and volatility. To avoid this issue, we include market capitalization (in logarithm) as a control in our regressions. ETF's AUM as well as the number of ETFs covering a stock also present some potential endogeneity issues. For example, investors' demand for existing or new ETFs may relate to how popular a given sector or asset class is at a particular point in time. This popularity also affects the amount of trading intensity and the volatility of the underlying securities. This argument can generate a positive relation between ETF ownership and volatility that confounds the causal effect we are trying to identify. On the other hand, the number of ETFs tracking a given stock depends on the number of indexes in which a stock appears. If more established, less volatile firms are more likely to be members of an index, then there can be a negative bias in the relation between ETF ownership and volatility.

In our tests, we take several steps to guard against potentially omitted variables. First, we include stock and month fixed effects. In addition, we control for stock size and liquidity, which is measured by the inverse of the stock price, the Amihud (2002) illiquidity measure of price impact, and the bid-ask spread. Also, we include standard predictors of returns that could also relate to volatility, such as book-to-market, past-12-month returns, and gross profitability (gross income scaled by total assets, as in Novy-Marx, 2013). Yet, we cannot entirely avoid the concern that ETF ownership is an endogenous variable within this framework. For this reason, in the next subsection, we provide additional analysis that identifies exogenous variation of ETF ownership using the annual reconstitution of the Russell indexes.

With this caveat in mind, we start by reporting the results of OLS regressions of daily volatility in a given month on ETF ownership at the end of the prior month. In Table 4, we present separate regressions for S&P 500 stocks and for the broader sample of Russell 3000 stocks. The goal is to assess how the effect of interest varies with firm size. All the controls date from the end of the prior month. We also include stock and month fixed effects in all regressions. Standard errors are double-clustered at the stock and month levels.

The results of the analysis are presented in Table 4. To ease interpretation, we standardize volatility and ETF ownership by subtracting the sample mean and dividing by the sample standard deviation. From Column (1) of Table 4, we infer that the relationship between ETF ownership and volatility is positive and strongly statistically significant. The economic magnitude is also large, as a one-standard-deviation move in ownership is associated with 13.2% of a standard deviation change in daily volatility. This result is consistent with the hypothesis that ETFs impound liquidity shocks in the underlying securities' prices.

Next, we test whether ETF ownership captures a different effect from the ownership of other institutional investors. Among these, open-end mutual funds are the most similar to ETFs because they also receive daily flows. ETFs are, however, different from other open-end funds in that they allow intraday trading. In Column (2) of Table 4, we include lagged ownership by active and index mutual funds, measured in the same way as ETF ownership (and standardized). The coefficients on both mutual fund ownership variables are positive and significant. However, the point estimates of both mutual fund ownership variables are significantly smaller in magnitude than the slope on ETF ownership, which remains intact. Thus, it appears that ETF ownership has an independent and stronger tie to volatility, which, according to the main hypothesis, depends on the fact that ETFs attract high-turnover investors.

In Column (3) of Table 4, we include three lags of the dependent variable to address the concern that the persistence in volatility could introduce reverse causality. The coefficient on ETF ownership remains large and significant at 7.3% of a standard deviation.

Extending the universe to smaller stocks (Columns (4) to (6)), the relationship between ETF ownership on volatility is weaker, amounting to about 3.3% to 5.2% of a standard deviation. Moreover, the slope is no longer statistically distinguishable from the coefficients for other mutual funds (Columns (5) and (6)).

The lower sensitivity of volatility to ETF ownership in a sample that is dominated by small stocks is consistent with the conjecture channel. The arbitrage activity that occurs at high frequency throughout the day does not require the creation or redemption of ETF shares. Hence, arbitrageurs can choose to minimize transaction costs by concentrating on the larger stocks in the ETF baskets when constructing the replicating portfolio. Such behavior, called 'optimized replication' can explain why smaller stocks inherit fewer shocks from the ETF market.

Given that ETF trading and the arbitrage activity involving the underlying securities occur intraday, the effect that we identify should also be visible at higher frequencies. In Internet Appendix Table A3, we replicate the analysis of Table 4 using intraday volatility as the dependent variable, computed from second-by-second returns within a day. ETF ownership is updated daily using the daily market capitalization of the stock and daily ETF flows. The results from these daily stock-level regressions confirm the sign and significance from the monthly sample. The economic magnitude is also in the same ballpark (the variables of interest are standardized): a one-standard-deviation increase in ETF ownership is associated with an increase of 10.6% in intraday volatility for S&P 500 stocks and with an increase of 2.4% in intraday volatility for Russell 3000 stocks. We give more emphasis to the results using daily volatility (Table 4) to stress the fact that we are not merely identifying a microstructure effect that washes out at lower frequencies.

4.2 Identification Using a Quasi-Natural Experiment

An identification based on cross-sectional and time-series variation in ETF ownership, which underlies the OLS results in Table 4, can be questionable if the stock-level controls fail to capture characteristics that co-determine ETF ownership and volatility. For this reason, we next present results from a more robust identification approach.

Chang, Hong, and Liskovich (2015) devise an identification strategy that exploits the exogenous variation in membership to the Russell 1000 and the Russell 2000 indexes. In our implementation, we closely follow Appel, Gormley, and Keim (2015), who cast this experiment within an instrumental variable (IV) framework.[18]

The Russell 1000 index comprises the top 1000 stocks by market capitalization, while the Russell 2000 includes the next 2000 stocks. Russell Inc. reconstitutes the indexes on the last Friday of June every year, based only on end-of-May stock capitalization; hence, no discretion is involved in index assignment.[19] Index composition remains constant for the rest of the year. For stocks in a close neighborhood of the cutoff, changes in index membership are random events, once controlling for the assignment variable, i.e., market capitalization, as they result from random variation in stock prices at the end of May.

Chang, Hong, and Liskovich (2015) show that, although the amount of passive assets benchmarked to the Russell 1000 is 2 to 3.5 times larger than those tracking the Russell 2000, the weights of the top stocks in the Russell 2000 are about 10 times larger than those of the bottom stocks in the Russell 1000. Consequently, a significantly larger amount of passive money tracks the top Russell 2000 stocks.

Figure 3a provides evidence that is consistent with the latter claim in the context of ETFs. The figure plots average ETF ownership as a function of market capitalization rankings for the Russell 3000 universe, in bins of 10 stocks, for 500 stocks to the right and 500 to the left of the cutoff (the 1000th position). We note that around the cutoff position, there is a discontinuity in ownership. Stocks immediately after the cutoff appear to display higher ownership than stocks immediately to the left. We note, further, that such discontinuity is not as evident when considering index fund (Figure 3b) and active fund (Figure 3c) ownership. This fact reassures us on the validity of the Russell in identifying the effect of ETF ownership on volatility, as opposed to the effect of institutional ownership in general.

Spurred by this evidence, we focus on stocks that move between the two indexes, and we use the switch event as an instrument for ETF ownership. Then, we regress our outcome variable, daily stock volatility, on instrumented ETF ownership. To identify the effect of interest, we rely on Appel, Gormley, and Keim's (2015) insight that variation in institutional ownership

---

[18] Other papers that exploit the Russell reconstitution are Cao, Gustafson, and Velthuis (2014); Crane, Michenaud, and Weston (2014); Fich, Harford, and Tran (2015); Lu (2014); and Mullins (2014).

[19] When the last Friday falls on the 29th or 30th day of the month, the two indexes are reconstituted on the preceding Friday. For more details, on the index formation process, see Appel, Gormley, and Keim (2015), and Russell Investments (2013).

around the cutoff is exogenous, after controlling for the market-capitalization-based ranking of the stock.[20] Unlike these authors, we also need to control for lagged volatility, because it affects the probability of switching indexes and it is correlated with our dependent variable.

The additional identifying assumption that needs to be satisfied is the exclusion restriction, that is, the requirement that the event affects the outcome variable only through the treatment variable. In our context, this translates into the condition that a switch in index membership only affects volatility through ETF ownership. Later in this section, we discuss reasons why this assumption might fail and conclude that this concern does not appear to be relevant in our context.

While Appel, Gormley, and Keim (2015) are constrained to use annual data by the availability of their governance measures, we cast our analysis at the monthly frequency because every month we have a different observation on the dependent variable (i.e., daily volatility). We note that the exogenous variation in ETF ownership only comes from the June switch. However, this exogenous component of ETF ownership is contained in all the monthly observations of this variable through May of the next year. Therefore, the 12 monthly observations of ETF ownership provide relevant explanatory power for the different observations of the dependent variable. Using one observation per year would entail a loss of power.

The first index reconstitution in our sample occurs in May 2000. Mullins (2014) and Appel, Gormley, and Keim (2015) report that the classification method of stocks to the Russell indices was modified after the reconstitution of June 2006. Until the June 2006 reconstitution, the cutoff for reclassification was simply the 1000th position in terms of market capitalization. Thus, in our main tests, we include end-of-month data between June 2000 and May 2007. As in Appel, Gormley, and Keim (2015), we consider several bandwidths: 100, 200, 300, 400, and 500 stocks on each side of the cutoff. Another difference in our implementation is that we consider switches in both directions, i.e., to the Russell 1000 from the Russell 2000 and to the Russell 2000 from the Russell 1000, while these authors only consider moves to the Russell 2000. The

[20] Indeed, Russell, Inc. uses a proprietary methodology to compute market capitalization (see Mullins, 2014). This fact implies that we cannot perfectly control for the ranking variable if we use market capitalization from CRSP. However, Appel, Gormley, and Keim (2015) show that the procedure is robust to substituting the CRSP measure with the Russell proprietary measure of market capitalization, for the years between 2002 and 2006, in which it is available.

drop in ETF ownership that comes from switching to the Russell 1000 is informative in our context.

The validity of the Russell experiment becomes questionable after 2006, as market capitalization is no longer strictly related to the index switch. Specifically, starting with the 2007 reconstitution, Russell Inc. adopted a banding rule whereby stocks only switch from their current index if they move beyond a 5% range around the market capitalization percentile of the 1000th stock. As expected, switches are more frequent before the introduction of the banding rule (Internet Appendix Table A4, Panel A). In Internet Appendix Table A4, we show, however, that our results hold even in the longer sample period.

We carry out a two-stage least squares estimation. In each stage, we run our regressions on two separate groups of stocks: those that in May, before index reconstitution, are in the Russell 1000 and those that are in the Russell 2000. The sample composition remains constant for all the months between June, the first month after index reconstitution, and May of the next year. The first stage consists of a regression of ETF ownership on an indicator variable for whether the stock switches index membership in June. For the Russell 1000 sample, the indicator variable flags stocks that switch to the Russell 2000. Vice versa for the Russell 2000 sample, the dummy captures a switch to the Russell 1000. In regression form, the first stage is

$$ETF\ Ownership_{it} = \alpha + \beta * I(Switched\ to\ other\ index\ )_{it} + Controls + Fixed\ effects + \varepsilon_{it}. \tag{2}$$

In the second stage, for the same two separate groups of stocks, we regress volatility on the fitted value of ETF ownership from the first stage. In regression form, this reads as

$$Volatility_{it} = \alpha + \beta * ETF\ \widehat{Ownership}_{it} + Controls + Fixed\ effects + \varepsilon_{it}. \tag{3}$$

In addition to the control for market capitalization, i.e., the assignment variable, we include the same set of controls as in the OLS regressions. While time fixed effects are part of the regression, we do not include stock fixed effects because identification in this experiment is inherently cross sectional; that is, it results from comparing switchers to non-switchers in a given time period. Standard errors are double-clustered at the stock and month levels. We standardize the ownership variables and volatility in the relevant samples to ease interpretation. Finally, following Appel, Gormley, and Keim (2015), we include different polynomials of the ranking variable: first (Panels A and B of Table 5), second, and third degree (Internet Appendix, Table 5, Panels C1-D2). Here, we report only the estimates on the main variable, to save space. The full results are in the Internet Appendix.[21]

Table 5, Panel A shows the first-stage regressions. We separately consider stocks that belong to the Russell 1000 before index reconstitution (Columns (1)–(5)) and stocks belonging to the Russell 2000 before index reconstitution (Columns (6)–(10)). The instrument is an indicator for whether the stock switches to the other index. The dependent variable (ETF ownership) is measured in each month following the index reconstitution. To illustrate the setting, consider Column (1). The sample includes stocks that are in the Russell 1000 index in May (prior to the reconstitution). We use the end-of-May cutoff to determine the stocks that are included in the sample (±100 stocks around the cutoff). The stocks remain in the sample in all months between June (after reconstitution) and May of next year, unless they drop out of the sample due to exogenous events (e.g., a merger). The indicator variable flags the stocks that switch to the Russell 2000 after reconstitution.

The results of this test show that switching indexes has a strong impact on ETF ownership. The slope on the switch indicator in Column (1) suggests that ETF ownership in the 12 months after reconstitution increases for the stocks switching to the Russell 2000 by about 24.4% of a standard deviation. Across Columns (1) through (5), the average effect is larger at about 41%.

Column (6) focuses on stocks that start out in the Russell 2000 in May prior to reconstitution, with the same bandwidth (±100 stocks around the cutoff). For the stocks that switch to the Russell 1000 after reconstitution, ETF ownership decreases by about 14.5% of a standard deviation. Across Columns (6) through (10), the average estimate is about -35%, which is of similar magnitude to the effect of switching to the Russell 2000. The strong statistical significance of the first-stage regressions reassures us about the validity of the instrument.

Table 5, Panel B, reports the second-stage estimates of the effect of ETF ownership on volatility in the next month. Analogous to the layout in Panel A, the instruments are indicators for a switch to either index, and the sample is restricted to members of either index before

[21] The Internet Appendix is at http://www.people.usi.ch/franzonf/ETFs_Internet_Appendix.pdf or http://fisher.osu.edu/fin/faculty/Ben-David/articles/ETFs_Internet_Appendix.pdf.

reconstitution. The effect of ETF ownership on volatility is significant across all samples and bandwidths. The magnitudes in Table 5, Panel B, are considerably larger than the OLS estimates in Table 4. The coefficients range between 16.8% and 80.3%, averaging about 32% of a standard deviation.

At first sight, the magnitudes of the results in Tables 4 and 5 may appear contradictory. On the one hand, in Table 4, we find stronger effects among the large S&P 500 stocks than among the broader universe of Russell 3000 companies. On the other hand, the effects in Table 5 are more economically significant than the ones in Table 4, although Table 5 focuses on relatively small stocks. This apparent inconsistency arises because the two tables use different estimation techniques, applied to different samples. First, the larger IV estimates from Table 5 may be revealing that the endogeneity of ETF ownership induces a negative bias in the OLS estimates in Table 4. A negative bias can occur if, for example, higher ETF ownership signals companies that belong to multiple indexes, which have less volatile stocks because they are more established companies. Second, Table 4 reports estimates for the average effects across all the stocks in the Russell 3000 index. In Table 5, instead, we examine the effects in a neighborhood around the cutoff between the Russell 1000 and Russell 2000. Specifically, the IV estimates measure the local average treatment effect (LATE), where the weights in the average are the ex-ante probability that a unit receives treatment (see, e.g., Lee and Lemieux, 2010). In other words, the IV estimates overweight stocks that are highly likely to switch indexes. These stocks can experience a drastic change from not being included in arbitrageurs' strategies to having top weights in these strategies, and vice versa. Arguably, we should expect that changes in ETF ownership have a bigger effect on these stocks. Given these considerations, we conclude that the IV estimates likely represent an upper bound, while the OLS coefficients are the lower bound, for the effect of ETF ownership on volatility.

A sign of a well-specified experiment is the fact that the estimates are stable when different degrees of the polynomials of the ranking variable are included (Lee and Lemieux, 2010). The Internet Appendix presents additional results using quadratic and cubic polynomials, (the first stages are adjusted accordingly). Reassuringly, the estimates are in the same ballpark as in Panel B of Table 5, which uses a first degree polynomial.

Finally, we come back to assessing the validity of the exclusion restriction in our context. The exclusion restriction is not satisfied if there is a correlated omitted variable that varies with index switches and affects volatility. ETF ownership could merely be a proxy for this omitted factor. For example, a violation occurs if, after appearing among the top stocks in the Russell 2000, a firm becomes more visible to investors. It is then possible that prices react more quickly to fundamental information and returns become more volatile, as shown by Andrei and Hasler (2015). Appel, Gormley, and Keim (2015) find that analyst coverage is largely unaffected after inclusion in the Russell 2000. Similarly, Crane, Michenaud, and Weston (2014) show that a switch to the Russell 2000 does not lead to increased media coverage. These results are important for our study as they suggest that the increase in volatility associated with the increase in ETF ownership is not likely to be caused by a rise in information production following switching. Moreover, in the next subsection, we show that mean reversion of prices increases with ETF ownership, suggesting that the driver of the increase in volatility is not only improved price discovery. Therefore, the evidence seems to rule out this specific violation of the exclusion restriction.

More generally, we obtain further corroboration of the validity of the exclusion restriction by combining the cross-sectional identification from the index switching experiment with time-series variation in ETF ownership. In particular, if the IV exercise is truly measuring the causal effect of ETF ownership, we should observe a stronger impact of index switching on volatility at times when aggregate ETF ownership is larger. Following this argument, we regress stock-level volatility on the interaction between the index-switching indicator and the (equally weighted) average of ETF ownership across Russell 2000 stocks. If the exclusion restriction is satisfied, then we would expect switching to the Russell 2000 (Russell 1000) to have a more positive (negative) effect on volatility at times when ETF ownership is overall higher.[22] The regressions include additional interactions to control for aggregate ownership by other mutual funds (passive and active) and for a time trend, given that ETF ownership increases over time. We also include the uninteracted variables as well as the usual stock-level controls and month fixed effects. The estimates in Panel C of Table 5 broadly support the validity of the exclusion

[22] We remind the reader that we expect a switch to the Russell 2000 (Russell 1000) to increase (decrease) volatility because that switch increases (decreases) ETF ownership, based on the evidence in Panel A of Table 5.

restriction. In Columns (1)–(5), the addition to the Russell 2000 has a larger impact on volatility at times of higher average ETF ownership. In two cases, the effect is statistically significant. In all five specifications focusing on the switch to the Russell 1000 (Columns (6)–(10)), the decrease in volatility is significantly larger in months when ETF ownership is higher.

While Panel C shows that the effect of ETF ownership on volatility is larger and more significant than the effects of other mutual funds, one may still wonder to what extent ETF ownership captures an independent effect from that of other mutual funds. In particular, the concern exists that index-switching stocks experience an overall migration of institutional owners, which in turn causes the change in volatility. The evidence in Figure 3, suggesting that only ETF ownership experiences a drastic change around the cutoff for index transition, partly alleviates this concern. While active funds have leeway in the way they track a benchmark, the evidence that ETF and index fund ownerships are not perfectly correlated may be surprising, as both types of funds do passive replication of the index. Internet Appendix Table A5 provides some insight into why we should not expect perfect correlation. Focusing on data from December 2014, we group ETFs and index funds into those that follow the index strictly using perfect replication and those that adopt alternative weighting schemes (e.g., growth-tilt, value-tilt, equal weights, low-volatility-tilt, etc.). We omit leveraged and short ETFs because they are not in our sample. The table shows that the percentages of funds that are strict followers differ between ETFs and index funds. For example, while 17% of Russell-1000 ETFs' market capitalization perfectly replicates the index, and 83% of ETFs follow alternative weighting schemes, the percentages are different for index funds: 29% are strict followers of the Russell 1000 and only 71% follow alternative schemes. The percentages are closer for the Russell 2000, but still not perfectly correlated. Even within the group of funds with alternative weighting schemes, the allocation of assets differs across ETFs and index funds because of the different allocation of assets to weighting schemes. Given the variety of weighting schemes, we conclude that ETFs and index funds are bound to have different demand for transitioning stocks following index reconstitution.

To conclude, in Panel D of Table 5, we provide a formal test of the importance of ETF ownership relative to ownership by other funds in driving the volatility effect. We study whether stocks switching to the Russell 1000, for which ETF ownership goes down, experience a larger decrease in volatility if in May before reconstitution they have a higher ratio of ETF ownership to ownership by other funds (either index or active funds). Reassuringly, the tests in Panel F reveal that the ownership ratio, with either index fund or active fund ownership at the denominator, magnifies the decrease in volatility for stocks quitting the Russell 2000.[23]

Given the evidence in Table 5, we feel more confident in giving a causal interpretation to the positive relation between ETF ownership and stock-level volatility. We next study whether the observed increase in volatility is related to an increase in transitory shocks to stock prices.

4.3 Identifying the Impact on Non-Fundamental Volatility

4.3.1 Variance Ratios

The finding that higher ETF ownership is associated with increased volatility does not necessarily confirm our hypothesis that ETFs impound liquidity shocks into the prices of the underlying securities. For example, Amihud and Mendelson (1987) provide a simple model in which the volatility of trading prices is positively related to the speed at which prices adjust to fundamentals. In addition, Andrei and Hasler (2015) prove theoretically and empirically that investor attention increases the sensitivity of prices to fundamentals and, therefore, volatility. If ETF arbitrage makes prices adjust more promptly to fundamentals, or if stocks in ETFs are exposed to higher investor attention, then the *fundamental* volatility of the underlying securities could go up. In this context, fundamental volatility coincides with price variation resulting from cash flow news that impounds a permanent shock in stock prices. This increase in volatility differs from our hypothesis, which instead focuses on *non-fundamental* volatility, defined as the price variation resulting from liquidity shocks to prices. These transitory shocks induce mean-reversion in prices.

O'Hara and Ye (2011) use variance ratios to measure the transitory component of stock prices. At time t, stock i's variance ratio is defined as:

$$VR_{i,t} = \left| \frac{Var(r_{k,i,t})}{k \cdot Var(r_{1,i,t})} - 1 \right|, \tag{4}$$

[23] We do not expect the symmetric effect to hold in the transition from the Russell 1000 to the Russell 2000, because we have no prior on how ETF ownership before switching, keeping ownership by other funds constant, affects volatility of stocks that experience an *increase* in ETF ownership after reconstitution.

where the numerator is the variance of *k*-period returns in the estimation window corresponding to time *t*, and the denominator is *k* times the variance of the single-period log returns in the same window *t* (also see Lo and MacKinlay, 1988).[24] When prices follow a random walk, the ratio of variances should be closer to one and the quantity in Equation (4) approaches zero. This simple device provides a non-parametric test of the impact of ETFs on non-fundamental volatility. If ETFs impound transitory shocks into prices, *VR* should increase with ETF ownership.

In our application, we construct the variance ratio using two different horizons because ETF arbitrage can affect stock prices at two different frequencies. As discussed in Section 2, arbitrageurs can affect volatility intraday, through the activity of hedge funds and high-frequency traders. We test this channel by constructing variance ratios from intraday returns. Moreover, APs can impact stock prices at the daily frequency through their daily creation and redemption of ETF shares. To find evidence of this channel, we construct the variance ratio using daily returns.

First, to construct intraday variance ratios (VR 15), we measure single-period returns from transaction prices at 5-second intervals and choose $k = 3$, so that multiperiod returns are measured over 15-second intervals. To estimate both variances, we use all the returns within a day. Then, we average the daily estimates over the month to obtain monthly observations. The choice of 15-second time intervals follows from the observation that ETF sponsors disseminate information about the portfolio NAV at 15-second intervals to facilitate high-frequency arbitrage. This frequency is therefore necessary to capture the intraday effect of arbitrageurs on the underlying stock prices. Second, to capture the lower frequency at which the activity of APs takes place, we also compute the ratio of the 5-day return variance to five times the 1-day return variance (VR 5). To have sufficient observations to estimate these variances, we use all the returns within a quarter. Hence, the frequency of this sample is quarterly.

Table 6 reports estimates from regressions of the standardized values of the stock-level variance ratio on standardized ETF ownership in the prior period. Panel A shows OLS regressions. The results point unambiguously to a positive and significant relation between ETF ownership and variance ratios. At both frequencies, the evidence suggests that the prices of stocks with higher ETF ownership are farther away from a random walk and therefore contain a

[24] Strictly speaking, only the first element in the absolute value in Equation (4) is a ratio of variances. We label the whole expression a "variance ratio" for convenience.

larger mean-reverting component. The effect is twice as large at the intraday frequency as it is at the 5-day frequency (about 11% of a standard deviation of VR 15 for a one-standard-deviation change in ETF ownership, for S&P 500 stocks). It is, however, statistically and economically significant for the 5-day frequency as well. Consistent with the pattern in Table 4, the effect is reduced but still significant in the intraday setting, when the universe is extended to smaller stocks (Russell 3000).

Given the concerns about the potential endogeneity of ETF ownership in the OLS regressions, we implement the IV estimation based on the Russell indexes reconstitution, using the variance ratio as the dependent variable in the second-stage regressions. (The first stage is identical to Table 5, Panel A). As before, we restrict the sample to the months following the reconstitutions between 2000 and 2006. Panel B of Table 6 reports the results for VR 15 from the monthly sample, while Panel C has the results for VR 5 from the quarterly sample. In all specifications, the IV confirms the positive slope on ETF ownership. Statistical significance is present in the majority of cases. Finally, the larger magnitude of the IV slopes than the OLS slopes mirrors the previous evidence regarding the effect on total volatility and can be explained in the same way. The results from the IV analysis give us more confidence in our causal interpretation of the positive link between ETF ownership and the mean-reverting component in stock prices.

4.3.2 Price Reversals

An alternative way to test whether ETFs impound transitory shocks into the underlying securities is to look for direct evidence of the sequence of events that appear in Figure 1. Following a demand shock in the ETF market, the prices of the underlying securities should move in the same direction as the initial shock. Then, prices revert to the initial level as the effect of the shock vanishes. Finding evidence of mean reversion in prices would also help us rule out the story that ETFs merely improve price discovery (as in Figure 2), which is an alternative explanation for the observed increase in volatility.

For this analysis, we focus on the daily frequency because the demand shocks in the ETF market can be clearly identified by measuring daily flows in ETFs. As explained above, ETF flows (redemptions and creations) are the result of APs' activity. Stock-level flows are defined as the weighted average of the daily flows in the ETFs that own the stock. The weights are the fraction of ownership in the stock held by each ETF. Dollar daily ETF flows are then expressed as a fraction of prior-day stock capitalization.

On the day when flows occur, we expect a price move in the same direction as the flows, irrespective of whether the motive for trade is new information or liquidity demand. To the extent that at least part of the originating shock is due to liquidity demand, a reversal should occur in the next days. To capture this behavior, we regress returns at different horizons (5, 10, and 20 days) on stock-level flows, using overlapping daily observations. We include the usual stock-level controls and time fixed effects, in addition to order imbalance, calculated as the dollar value of buy minus sell trades from TAQ, divided by market capitalization. Order imbalance is a natural control in this context because daily flows in ETFs could merely be a proxy for aggregate demand in the underlying securities, which induces negative autocorrelation in returns (Chordia and Subrahmanyam, 2004). The standard errors are clustered at the day level and we correct for the autocorrelation of residuals induced by overlapping observations for multiday returns using the Newey and West (1987) estimate of variance.

In Table 7, returns are expressed in percentages, whereas net flows are standardized. From Column (1), we note that on the same day, ETF flows and returns move in the same direction. The contemporaneous price move is 16.7 bps for a one-standard-deviation change in net flows for S&P 500 stocks. The high significance is not surprising, as flows and returns are measured on the same day (hence, this is *not* a predictive regression). In addition, we note that the magnitude of the change in prices exceeds the half-spread, which is about 8.5 bps for the sample of large stocks. This magnitude rules out the possibility that flows cause a simple bid-ask bounce.

More relevant to identifying the transmission of liquidity shocks, ETF flows predict a reversal of the underlying stocks' prices in the next 20 days (Columns (2)–(4)). This evidence is consistent with the conjecture that the demand shocks in the ETF market add a mean-reverting component to stock prices. From Column (4), we can infer that almost half of the initial price impact is reversed (1.00167*0.00072/0.00167 = 0.43). Extending the horizon farther out to 40 days does not increase the magnitude of reversals (not reported). As in the prior tables, the absolute effects are smaller in the extended universe of Russell 3000 stocks.[25]

In sum, the evidence in this subsection suggests that the positive link between ETF ownership and volatility, which we report in Tables 4 and 5, is consistent with ETFs being a catalyst for liquidity trading, which then propagates to the underlying securities. Specifically, ETFs appear to add a mean-reverting component to stock prices both intraday and at the daily frequency.

5 Exploring the Arbitrage Channel

Having established a link between ETF ownership and stock volatility, we next look for evidence that the transmission of shocks between ETFs and the underlying securities occurs because of arbitrage activity. To this end, we first define a proxy for arbitrage activity. The difference between the ETF price and the net asset value of the underlying basket (NAV), commonly labeled ETF mispricing, is a signal for the profitability of ETF arbitrage. Hence, we expect a stock's involvement in arbitrage trades to be a positive function of the mispricing of the ETFs that hold the stock. Using this proxy, we study whether arbitrage activity has an incremental impact on volatility and variance ratio for a given level of ETF ownership.[26]

Then, we conjecture that the proxy for expected arbitrage activity should have a weaker effect on prices for stocks that are harder to arbitrage. In other words, we seek evidence that limits to arbitrage play a role in attenuating the propagation of shocks to the underlying securities. This evidence would indirectly testify to the importance of the arbitrage channel.

[25] An alternative explanation for our findings in Table 7 could be that APs create ETF units (and generate positive ETF flows) to lend to clients wishing to short the ETF. If these shorts were informed, ETF prices as well as those of basket securities would subsequently fall. To rule out this alternative, in Internet Appendix Table A6, we repeat the test, controlling for lending fees, which proxy for the tightness of the share lending market. Because this variable is only available for a subset of securities, including it reduces our sample size. Reassuringly, the results in the S&P 500 sample are mostly unaffected.

[26] It could actually be the case that ETF mispricing signals a *lack of* arbitrage activity. That is, more mispricing is present when arbitrageurs refrain from entering the market. This could be an issue for our tests if the reason why arbitrageurs abstain from their trades is volatility in the underlying securities, which is the dependent variable in our tests. In such a case, the endogeneity of mispricing could bring a positive spurious correlation with volatility. To address this concern, we control for the lagged value of the dependent variable, so that we study the impact of mispricing on innovations in volatility. This helps to attenuate the endogeneity concern because arbitrage trades are not likely to condition on innovations in volatility in the next period. Further, in the tests in which we interact mispricing with measures of the limits to arbitrage, this potential endogeneity would lead to the opposite sign of the coefficient relative to what we find. See the discussion below.

We use two proxies for limits to arbitrage: the stock-level bid-ask spread and share-lending fees. First, because ETF arbitrage involves a roundtrip transaction in the stock, a large stock-level bid-ask spread reduces the profitability of arbitrage trades and therefore the incidence of arbitrage trading in a given stock. Second, when the arbitrage transaction involves shorting the stock (i.e., the NAV is above the ETF price), higher stock lending fees discourage arbitrageurs. In addition, a high share-lending fee can reflect a shortage of shares for lending, meaning that some arbitrageurs may simply not be able to carry out the trade (Cohen, Diether, and Malloy, 2007).

Given the high-frequency fluctuations in arbitrage activity, we carry out our tests at the daily frequency, which allows us to measure the variables of interest in a more timely way. Thus, the dependent variables for these tests are intraday volatility, which is estimated from second-by-second returns within a day, and the daily variance ratio resulting from the comparison of 15-second returns to three times 5-second returns within a day (VR 15). The main explanatory variable is the stock-level measure of absolute ETF mispricing in the prior day. This variable is calculated by summing the absolute dollar mispricing (i.e., the difference between the ETF price and NAV, as a fraction of the ETF price, multiplied by the dollar holdings in the stock) across all ETFs holding stock i, and expressing this quantity as a fraction of a stock's capitalization:

$$Abs\left(mispricing_{i,t}\right) = \frac{\sum_{j=1}^{J} w_{i,j,t} * AUM_{j,t} * \left|mispricing_{j,t}\right|}{Mkt\ Cap_{i,t}}. \tag{5}$$

Within this variable, we interact the effect of the ETF mispricing—a signal for the attractiveness of the stock for arbitrage trades—with the ownership of each ETF in the stock's capital base, which measures the relative importance of each ETF for the given stock. We take the absolute value of mispricing because arbitrage activity is triggered by both positive and negative discrepancies between ETF prices and the NAV. It is therefore important to avoid netting out these deviations across ETFs. In a second set of tests, in which we condition on the direction of the arbitrage trades, we use net mispricing. Net mispricing differs from absolute mispricing for the omission of the absolute value in Equation (5).

In Panel A of Table 8, the dependent variable is intraday volatility. The sample consists of S&P 500 stocks, where, according to our prior results, most of the effect of ETFs occurs. In Column (1), we test whether absolute mispricing at the close of day t – 1, which is a proxy for arbitrage activity on day *t*, has an incremental effect on volatility for a given level of ETF ownership. In addition to the usual controls, we include the mispricing on day *t* – 2 and the lagged dependent variable. The goal is to capture the effect of the innovation in mispricing on the innovation in volatility, given that mispricing on day *t* – 1 could itself depend on volatility (i.e., ETFs holding stocks that are more volatile are more likely to be mispriced, as discussed in footnote 26). We also include the return on the stock on day *t* – 1 to capture variation in mispricing that is exogenous to movements in the stock price itself. That is, we identify variation in mispricing resulting from movements in the ETF price or in the prices of the other stocks in the basket, but not in the stock's own price. We note that the effect of absolute mispricing is positive and significant, amounting to about 2.3% of a standard deviation of the dependent variable for a one-standard-deviation change in mispricing (both variables are standardized). Further, the effect of ETF ownership drops in magnitude relative to the specification without mispricing (compare with Internet Appendix Table A3). This evidence supports the view that arbitrage activity, as proxied by mispricing, is the transmission channel for the effect of ETF ownership on volatility.

Next, we report specifications that include interactions of absolute mispricing with the proxies for arbitrage costs. For each measure of limits to arbitrage, we define a dummy variable for stocks that are in the top of half of the distribution of the variable in the prior period.[27] We leave out stock fixed effects, because we wish to achieve identification from the cross-sectional variation in the proxies. From Table 8, Panel A, Column (2), we infer that the effect of arbitrage on volatility, as proxied by absolute mispricing, is significantly weaker for stocks with a high bid-ask spread. This evidence suggests that limits to arbitrage are playing a role in the transmission of shocks to the underlying securities.

Next, we break up the sample by the sign of net mispricing. A priori, we do not expect the sign of mispricing to matter for the interaction with the bid-ask spread, because the arbitrage trade involves a roundtrip transaction in the underlying stock in any case. The results in Columns (3) and (4) confirm this conjecture.

[27] Information on share-lending fees is sparse, especially in the initial part of the sample. Therefore, we use the average fee in the month.

In Column (5), Panel A, Table 8, share-lending fees have a marginally significant impact in attenuating the effect of arbitrage on volatility. More importantly, we now expect this effect to differ based on the sign of net mispricing. Only when mispricing is negative (i.e., the ETF price is below the NAV) does the arbitrage trade involve a short sale of the underlying stocks. The estimates in Columns (6) and (7) square nicely with this prediction and provide strong evidence for the role of arbitrage activity in generating the effect of interest.

We also note that the sign of the interactions with the proxies for arbitrage costs tends to rule out concerns about the endogeneity of mispricing (see footnote 26). Indeed, if mispricing were capturing the fact that arbitrageurs abstain from trading because volatility discourages them, we would expect this effect to be even stronger for illiquid stocks or for stocks that are hard to locate, given that these characteristics correlate positively with volatility. That is, the sign on the interactions with arbitrage costs should be positive. Instead, contrary to this view, the interactions have negative and significant coefficients.

Panel B of Table 8 replicates the analysis using the variance ratio from intraday returns as dependent variable. Mostly, the results mirror the evidence in Panel A, further supporting the role of arbitrage in transmitting liquidity shocks to the prices of the securities in the ETF baskets.

For completeness, in Internet Appendix Table A7, we report the analysis for the Russell 3000 universe. As expected, in this sample, the effects are weaker or nonexistent. These results confirm our prior belief that arbitrageurs tend to focus on the larger stocks in the ETF baskets when doing optimized replication. Therefore, the effect of interest is located mostly among large stocks.

6 Are ETFs Attracting a "New Layer" of Volatility?

The second testable hypothesis spelled out in Section 2 is that the liquidity shocks hitting ETF-owned stocks represents a *new layer* of demand that would not be present in the market if ETFs did not exist. The argument is that ETFs provide previously unavailable trading opportunities, at low cost and high frequency that cause new traders and/or new trading strategies to materialize. For example, ETFs might make the hedging of industry risk in statistical arbitrage strategies that involve mispriced securities significantly cheaper, so that the volume of these trades might increase. In this sense, the introduction of ETFs is analogous to a decrease in trading costs that enables traders to operate at higher frequency.

The alternative view to this hypothesis is that ETFs merely provide a convenient conduit for existing investors who wish to trade the underlying securities. According to this view, liquidity trading is reshuffled from stocks with low ETF ownership to stocks with high ETF ownership. We label this argument the "reshuffling hypothesis." To stay with the previous example, the same statistical arbitrageurs that currently employ ETFs for hedging purposes would have previously constructed a hedging portfolio using stocks in the same industry as the mispriced security.

To disentangle the two alternative we need to resort to time-series evidence. The ultimate prediction of the new-layer hypothesis is that the growth in the ETF market attracts liquidity trading to the stock market. Hence, we should observe higher volatility at times of higher ETF stock ownership. The reshuffling hypothesis, instead, predicts that aggregate volatility should not change because of ETFs.

In Panel A of Table 9, we report the estimates from a regression of average daily volatility on lagged average ETF ownership across all stocks in CRSP. The frequency is monthly, and volatility is computed using the daily returns in a month. We include lagged volatility to set up the regression as a test for Granger causality. To mitigate the concern that ETF ownership proxies for omitted factors related to institutional ownership, we include lagged average ownership by index and active funds. Importantly, we add a time trend as a catch-all control for developments in aggregate conditions (e.g., a protracted reduction in trading costs). We find that ETF ownership significantly predicts volatility, with a positive sign. The economic magnitude falls between that reported in Tables 4 and Table 5. In Column (2), we replicate the analysis in first differences and find consistent results.

We note that this analysis does not imply that volatility has increased over the sample period as a result of the positive trend in ETF ownership. Indeed, the positive association between ETF average ownership and volatility in Table 9 does not hold without the inclusion of the time trend. Therefore, our results point to a significant relation between *de-trended* ETF ownership and aggregate volatility. For this reason, our findings are not in contradiction with

Brandt, Brav, Graham, and Kumar (2010), who show a decrease in aggregate stock volatility after 2003.

While this analysis does not support the extreme version of the reshuffling hypothesis for the market-wide effect of ETF ownership on volatility, it is still interesting to ask whether some stocks experience a decrease in volatility at the expense of others as ETF ownership increases. In other words, we ask whether some partial reshuffling of liquidity trading is taking place. To this purpose, we test whether, as aggregate ETF ownership increases, volatility declines for some groups of stocks and rises for others. We sort the universe of stocks into five quintiles by ETF ownership. The average ETF ownership in the bottom quintile is 0.70% of a stock's capitalization, while in the top quintile it is about 4%. For each quintile, at the monthly frequency, we regress the average volatility for that group of stocks on the (lagged) average ETF ownership across the entire market and include controls. The explanatory variables are the same across quintiles, because the goal is to test whether the same aggregate developments are related to different changes in the volatility of different groups of stocks. From Panel B of Table 9, we note that all groups of stocks experience a significant increase in volatility as aggregate ETF ownership increases. This evidence does not support a partial reshuffling of liquidity trading across stocks. Quite relevantly, the effect of interest is strongest in the quintile with the top ETF ownership, which strengthens the case for a causal interpretation of the time-series association between ETF ownership and volatility. Finally, running the same regressions in first differences confirms our results (Table 9, Panel C).

To conclude, while the time-series setting of this analysis prevents us from drawing unambiguous causal inferences, the evidence in this section is consistent with the view that ETFs attract a new layer of demand shocks to the market as opposed to causing a reshuffling of existing demand across stocks.

7 Conclusion

With $3 trillion of assets under management globally (August 2015), ETFs are rising steadily among the big players in the asset management industry. This asset class is also capturing an increasing share of transactions in financial markets. For example, in August 2010, ETFs and other exchange traded products accounted for about 40% of all trading volume in U.S. markets.

The success of ETFs is justified by the fact that these investment vehicles offer an unprecedented source of diversification at low cost and high liquidity. However, providing investors with new ways to express their liquidity needs can raise the volatility of the securities in the ETF baskets. The evidence in this paper appears consistent with this previously unexplored effect of ETFs.

We present results showing that the stocks in ETF baskets display higher volatility than otherwise similar securities. Through a quasi-natural experiment based on the reconstitution of the Russell indexes, we are able to attach a causal interpretation to this finding. The presence of ETFs also appears to cause the underlying securities' prices to diverge from random walks, both intraday and daily. These effects are significantly related to proxies for the intensity of arbitrage activity between the ETFs and their baskets.

This evidence paints a picture in which liquidity shocks in the ETF market are passed down to the prices of the underlying securities by the transmission chain of arbitrage trades. Moreover, because of their ease of trade and cost effectiveness, ETFs attract higher turnover investors than the average stock in their baskets. Consequently, liquidity trading and volatility in stock prices increase with ETF ownership.

In addition, we find that aggregate volatility varies significantly over time with aggregate ETF ownership in the stock market, controlling for ownership by other mutual funds and for a time trend. With the caveat that the time-series identification of this effect does not allow for a conclusive causal inference, the evidence suggests that ETFs bring a new layer of volatility to the market, as opposed to just causing a migration of existing liquidity traders across securities. We explain this finding as a result of the new trading opportunities, at low cost and high frequency, made possible by ETFs.

To conclude, a new theoretical framework seems necessary to gauge the tradeoff between the decreased transaction costs and the improved access to diversification that ETFs bring about and the increase in volatility, which can lead to higher perceived risk and higher risk premia. The general equilibrium and welfare implications of this important wave of financial innovation therefore deserve further investigation.

References

Amery, Paul, 2015, ETFs Giants Bet on Futures for Flows, *Financial Times*, February 1, 2015.

Amihud, Yakov, 2002, Illiquidity and Stock Returns: Cross-Section and Time-Series Effects, *Journal of Financial Markets* 5, 31–56.

Amihud, Yakov, and Haim Mendelson, 1986, Asset Pricing and the Bid-Ask Spread, *Journal of Financial Economics* 17, 223–249.

Amihud, Yakov, and Haim Mendelson, 1987, Trading Mechanisms and Stock Returns: An empirical investigation, *Journal of Finance* 42(3), 533–553.

Andrei, Daniel, and Michael Hasler, 2015, Investor Attention and Stock Market Volatility, *Review of Financial Studies* 28(1), 33-72.

Anton, Miguel, and Christopher Polk, 2014, Connected Stocks, *Journal of Finance* 69(3), 1099–1128.

Appel, Ian, Todd A. Gormley, and Donald B. Keim, 2015, Passive Investors, Not Passive Owners, University of Pennsylvania, Working Paper.

Barberis, Nicholas, Andrei Shleifer, and Jeffrey Wurgler, 2005, Comovement, *Journal of Financial Economics* 75(2), 283–317.

Basak, Suleyman, and Anna Pavlova, 2013a, Asset Prices and Institutional Investors, *American Economic Review* 103(5), 1728–1758.

Basak, Suleyman, and Anna Pavlova, 2013b, A Model of Financialization of Commodities, Working paper, London Business School.

Bessembinder, Hendrik, and Paul J. Seguin, 1992, Futures-Trading Activity and Stock Price Volatility, *Journal of Finance* 47 (5), 2015–2034.

Blackrock, 2014, ETP Landscape, Industry Highlights, December.

Bradley, Harold, and Robert E. Litan, 2010, Choking the Recovery: Why New Growth Firms Aren't Going Public and Unrecognized Risks of Future Market Disruptions, Working paper, Kauffman Foundation.

Brandt, Michael W., Alon Brav, John R. Graham and Alok Kumar, 2010, The Idiosyncratic Volatility Puzzle: Time Trend or Speculative Episodes? *Review of Financial Studies* 23(2), 863–899.

Brunnermeier, Markus K., and Stefan Nagel, 2004, Hedge Funds and the Technology Bubble, *Journal of Finance* 59(5), 2013–2040.

Cao, Charles, Matthew Gustafson, and Raisa Velthuis, 2014, Capital Market Consequences of Index Membership: Evidence from Small Firms, Working Paper Pennsylvania State University.

Cella, Cristina, Andrew Ellul, and Mariassunta Giannetti, 2013, Investors' Horizons and the Amplification of Market Shocks, *Review of Financial Studies* 26(7), 1607–1648.

Chang, Eric C., Joseph W. Cheng, and J. Michael Pinegar, 1999, Does Futures Trading Increase Stock Market Volatility? The Case of the Nikkei Stock Index Futures Markets, *Journal of Banking & Finance* 23, 727–753.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, 2015, Regression Discontinuity and the Price Effects of Stock Market Indexing, *Review of Financial Studies*, forthcoming.

Cheng, Minder, and Ananth Madhavan, 2009, The Dynamics of Leveraged and Inverse Exchange-Traded Funds, *Journal of Investment Management* 7(4), XXX–XXX.

Cheng, Si, Massimo Massa, and Hong Zhang, 2015, The Multiple Facets of ETF Investing: A World-Wide Analysis, INSEAD, Working Paper.

Chordia, Tarun, and Avanidhar Subrahmanyam, 2004, Order Imbalance and Individual Stock Returns: Theory and Evidence, *Journal of Financial Economics* 72(3), 485–518.

Cohen, Lauren, Karl B. Diether, and Christopher J. Malloy, 2007, Supply and Demand Shifts in the Shorting Market, *Journal of Finance* 62(5), 2061–2096.

Coval, Joshua, and Erik Stafford, 2007, Asset Fire Sales (and Purchases) in Equity Markets, *Journal of Financial Economics* 86(2), 479–512.

Crane, Alan D., Sebastien Michenaud, and James Weston, 2014, The Effect of Institutional Ownership on Payout Policy: A Regression Discontinuity Design Approach, Rice University, Working Paper Rice University.

Da, Zhi, and Sophie Shive, 2015, When the Bellwether Dances to Noise: Evidence from Exchange-Traded Funds, Notre Dame University, Working Paper.

Danthine, Jean-Pierre, 1978, Information, Futures Prices, and Stabilizing Speculation, *Journal of Economic Theory* 17, 79–98.

Ellul, Andrew, 2006, Ripples Through Markets: Inter-Market Impacts Generated by Large Trades, *Journal of Financial Economics* 82(1), 173–196.

Fich, Eliezer M., Jarrad Hartford, and Anh L. Tran, 2015, Motivated Monitors: The Importance of Institutional Investors’ Portfolio Weights, *Journal of Financial Economics*, forthcoming.

French, Kenneth R., 2008, The Cost of Active Investing, *Journal of Finance* 63(4), 1537–1573.

Greenwood, Robin, 2005, Short- and Long-Term Demand Curves for Stocks: Theory and Evidence on the Dynamics of Arbitrage. *Journal of Financial Economics* 75(3), 607–649.

Greenwood, Robin, and David Thesmar, 2011, Stock Price Fragility, *Journal of Financial Economics* 102(3), 471–490.

Grossman, Sanford Jay, 1989, An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies, NBER Working paper 2357.

Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and Market Structure, *Journal of Finance*, 43(3), 617-633.

Hombert, Johan, and David Thesmar, 2014, Overcoming Limits of Arbitrage: Theory and Evidence, *Journal of Financial Economics* 111(1), 26–44.

Israeli, Doron, Charles M. C. Lee, and Suhas Sridharan, 2015, Is there a Dark Side to Exchange Traded Funds (ETFs)? An Information Perspective, Working Paper Stanford University.

Lee, Charles M. C., Andrei Shleifer, and Richard H. Thaler, 1991, Investor Sentiment and the Closed-End Fund Puzzle, *Journal of Finance* 46(1), 75–109.

Lee, David S., and Thomas Lemieux, 2010, Regression Discontinuity Designs in Economics, *Journal of Economic Literature* 48(2), 281–355.

Leippold, Markus, Lujing Su, and Alexander Ziegler, 2015, Do Index Futures and ETFs Affect Stock Return Correlations? Working Paper University of Zurich.

Lo, Andrew W., and Archie Craig MacKinlay, 1988, Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test, *Review of Financial Studies* 1(1), 41–66.

Lou, Dong, 2012, A Flow-Based Explanation of Returns Predictability, *Review of Financial Studies* 25(12), 3457–3489.

Lou, Dong, and Christopher Polk, 2014, Comomentum: Inferring Arbitrage Activity from Return Correlations, Working paper, London School of Economics.

Lu, Ruichang, 2014, How Does Institutional Ownership Affect Bank Loan Pricing: Evidence from a Regression Discontinuity Design, Working paper.

Madhavan, Ananth, 2012, Exchange-Traded Funds, Market Structure and the Flash Crash, Working paper, BlackRock Inc.

Madhavan, Ananth and Aleksander Sobczyk, 2014, Price Dynamics and Liquidity of Exchange-Traded Funds, SSRN working paper.

Malamud, Semyon, 2015, A Dynamic Equilibrium Model of ETFs, Swiss Finance Institute Working Paper.

Marshall, Ben R., Nhut H. Nguyen, and Nuttawat Visaltanachoti, 2013, ETF Arbitrage: Intraday Evidence, *Journal of Banking and Finance* 37(9), 3486–3498.

Morningstar, 2015, Morningstar Direct: U.S. Asset Flows Update, February 15.

Mullins, William, 2014, The Governance Impact of Index Funds: Evidence from Regression Discontinuity, Working Paper, MIT.

Newey, Whitney K. and Kenneth D. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55(3), 703–708.

Novy-Marx, Robert, 2013, The Other Side of Value: The Gross Profitability Premium, *Journal of Financial Economics* 108(1), 1–28.

O'Hara, Maureen, and Mao Ye, 2011, Is Market Fragmentation Harming Market Quality? *Journal of Financial Economics* 100, 459–474.

Petajisto, Antti, 2013, Inefficiencies in the Pricing of Exchange-Traded Funds, NYU Stern School of Business, Working paper.

Pontiff, Jeffrey, 1996, Costly Arbitrage: Evidence from Closed-End Funds, *Quarterly Journal of Economics* XXX, 1135–1151.

Ramaswamy, Srichander, 2011, Market Structures and Systemic Risks of Exchange-Traded Funds, Working paper, Bank of International Settlements.

Richie, Nivine, Robert Daigler, and Kimberly C. Gleason, 2008, The Limits to Stock Index Arbitrage: Examining S&P 500 Futures and SPDRs, *Journal of Futures Markets* 28(12), 1182–1205.

Roll, Richard, Eduardo Schwartz, and Avanidhar Subrahmanyam, 2007, Liquidity and the Law of One Price: The Case of the Futures-Cash Basis, *Journal of Finance* 62 (5), 2201–2234.

Russell Investments, 2013, Russell U.S. Equity Indexes Construction and Methodology.

Shleifer, Andrei, 1986, Do Demand Curves for Stocks Slope Down? *Journal of Finance* 41, 579–590.

Stambaugh, Robert F., 2014, Investment Noise and Trends, *Journal of Finance* 69 (4), 1415–1453.

Stein, Jeremy C., 1987, Informational Externalities and Welfare-Reducing Speculation, *Journal of Political Economy* 95, 1123–1145.

Trainor, Jr., William J., 2010, Do Leveraged ETFs Increase Volatility? *Technology and Investment* 1(3), 215–220.

Turnovsky, Stephen J., 1983, The Determination of Spot and Futures Prices with Storable Commodities, *Econometrica* 51, 1363–1387.

Vayanos, Dimitri, and Paul Woolley, 2013, An Institutional Theory of Momentum and Reversal, *Review of Financial Studies* 26 (5), 1087–1145.

Wurgler, Jeffrey, 2011, On the Economic Consequences of Index-Linked Investing, in Gerald Rosenfeld, Jay W. Lorsch, and Rakesh Khurana, eds.: Challenges to Business in the Twenty–First Century (American Academy of Arts and Sciences, Cambridge, MA).

Table 1. ETF Ownership Statistics

The table presents descriptive statistics for ETF (Panel A), Index Fund (IF, Panel B), and Active Fund (AF, Panel C) ownership of stocks. For each year, across months and stocks, we average the number of institutions and the percentage of each stock owned by an institution type. We also report cumulative assets owned by these institutions in the stocks in our sample. We present statistics for S&P 500 stocks (left columns) and for Russell 3000 stocks (right columns).

Panel A: ETFs

	S&P 500			Russell 3000		
Year	Average Number of ETFs in firm	Average ownership of ETF in firm (%)	Cumulative assets in ETFs ($mil)	Average Number of ETFs in firm	Average ownership of ETF in firm (%)	Cumulative assets in ETFs ($mil)
2000	2.45	0.22	32103.81	2.41	0.25	38070.46
2001	13.45	0.48	51863.61	8.91	0.36	60868.03
2002	15.47	0.90	64222.06	10.18	0.83	79973.03
2003	15.95	1.05	79002.90	10.42	0.95	98669.69
2004	21.40	1.22	97156.25	14.30	1.26	131744.95
2005	24.75	1.51	130211.35	15.73	1.55	184462.52
2006	25.80	1.67	173305.69	16.81	1.84	236819.09
2007	36.04	2.00	224374.08	22.60	2.21	302123.01
2008	50.61	2.76	245628.88	30.26	2.87	340147.47
2009	53.19	3.27	230367.72	31.30	3.53	344573.84
2010	52.08	3.30	286060.19	30.08	3.74	451925.78
2011	52.77	3.61	357624.80	28.87	3.81	552138.06
2012	49.25	3.90	414072.51	27.24	3.91	614854.72
Average	30.69	2.09	179547.12	20.13	2.37	271977.64

Panel B: Index Funds

	S&P 500			Russell 3000		
Year	Average Number of IF in firm	Average ownership of IF in firm (%)	Cumulative assets in IF ($mil)	Average Number of IF in firm	Average ownership of IF in firm (%)	Cumulative assets in IF ($mil)
2000	121	3.69	441378.69	38	2.50	510947.95
2001	137	4.59	447082.06	44	3.33	511055.06
2002	147	5.26	419124.27	49	3.71	484953.78
2003	151	5.16	442165.27	53	3.55	518196.64
2004	151	6.04	591421.69	53	5.06	727197.21
2005	147	6.01	609892.09	52	4.43	745496.10
2006	137	6.01	670264.68	51	4.17	803490.86
2007	139	6.46	791656.16	53	4.36	954011.74
2008	146	7.01	668895.96	58	4.81	806636.84
2009	141	7.17	595467.90	58	4.82	705905.00
2010	126	7.32	730526.92	52	4.92	877263.49
2011	119	6.34	692858.38	51	4.42	833349.16
2012	115	6.81	833837.99	50	4.69	992960.63
Average	137	5.99	609898.90	51	4.23	736336.70

Panel C: Active Funds

	S&P 500			Russell 3000		
Year	Average Number of AF in firm	Average ownership of AF in firm (%)	Cumulative assets in AF ($mil)	Average Number of AF in firm	Average ownership of AF in firm (%)	Cumulative assets in AF ($mil)
2000	591	16.36	1723372.63	84	11.85	2234091.17
2001	609	17.63	1477374.71	96	12.42	1811560.24
2002	601	18.29	1281444.04	101	13.44	1583896.67
2003	611	17.47	1301381.68	103	13.36	1640344.79
2004	576	17.67	1587185.42	108	14.66	2053690.46
2005	557	17.49	1692922.11	111	14.83	2234306.76
2006	519	17.81	1905312.35	110	15.25	2472396.64
2007	552	17.96	2166531.29	128	15.70	2837458.34
2008	605	18.32	1693863.00	144	15.53	2213057.71
2009	620	18.58	1461597.13	148	14.76	1867356.53
2010	564	18.50	1727851.31	140	15.05	2251251.62
2011	534	17.88	1856191.87	134	15.03	2429274.78
2012	525	16.95	1960550.80	128	14.15	2538780.18
Average	574	17.76	1679068.27	118	14.31	2178583.89

Table 2. ETFs versus Stocks: Liquidity, Institutional Ownership, Churn Ratio

The table reports statistics for ETF- and stock-level liquidity, investor turnover, and ownership. Panel A shows the security-level liquidity measures (bid-ask spread, Amihud (2002) ratio, and daily turnover) as well as the churn ratio measures of the investors in the securities (churn ratios 1 and 2). For all ETFs in our sample, we compute the average measure of liquidity or churn ratio across the stocks in the basket in a given quarter. Then, we value-weight the ETF-level and basket-level measures across all ETFs at the quarter level using ETF market capitalization (thus having 52 quarters in our sample). Churn ratio 1 is from Cella, Ellul, and Giannetti (2013), who compute an institutional investor–level churn ratio as the sum of quarterly absolute changes in dollar holdings over average assets under management (the data come from SEC 13-F filings). This measure is then averaged across institutions at the stock level using the fraction of a company held by each institution as a weight. Churn ratio 2 differs only in that the investor-level churn ratio is computed as the minimum between the absolute value of buys and sells, divided by prior quarter holdings. Panel B presents information about institutional ownership: averaged across all 52 quarters, in the first quarter of the sample, and in the last quarter of the sample. Ownership averages are presented for the ETFs in our sample and all the stocks in the CRSP data set. Variable descriptions are provided in Appendix Table A1. *t-statistics* for the test of the null hypothesis that the difference is equal to zero are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between 2000:Q1 and 2012:Q4.

Panel A: Liquidity and Investors' Churn Ratio Measures

Liquidity measures	Variable	Quarters	ETFs	Stocks	Difference	t-stat
Security-level	Bid-Ask Spread	52	0.003	0.005	-0.002***	(-3.518)
	Amihud ratio	52	0.002	0.008	-0.006***	(-9.702)
	Daily turnover	52	0.093	0.011	0.083***	(13.462)
Investor-level	Churn Ratio 1	52	0.307	0.240	0.067***	(10.195)
	Churn Ratio 2	52	0.154	0.125	0.029***	(7.493)

Panel B: Types of Institutional Ownership

	Ownership averaged across the sample			Ownership at 2000:Q1		Ownership at 2012:Q4	
Type of Institution	Quarters	ETFs	Stocks	ETFs	Stocks	ETFs	Stocks
All Institutions	52	0.474	0.621	0.280	0.511	0.492	0.651
Banks	52	0.131	0.137	0.052	0.114	0.202	0.116
Endowments	52	0.006	0.001	0.000	0.003	0.001	0.001
Hedge Funds	52	0.033	0.030	0.022	0.019	0.028	0.036
Insurance	52	0.014	0.033	0.007	0.034	0.011	0.026
Investment Advisors	52	0.198	0.211	0.125	0.166	0.167	0.231
Investment Companies	52	0.017	0.163	0.010	0.139	0.023	0.196
Pension Funds	52	0.009	0.035	0.001	0.031	0.008	0.026
Individual Investor (in 13F)	52	0.001	0.000	0.000	0.000	0.000	0.001
Research Firms	52	0.058	0.006	0.061	0.003	0.028	0.008
Corporations	52	0.003	0.001	0.000	0.001	0.017	0.003
Venture Capital	52	0.000	0.000	0.000	0.000	0.000	0.001
Private Equity	52	0.000	0.001	0.000	0.000	0.000	0.001
Sovereign Funds	19	0.000	0.000	0.000	0.000	0.000	0.001

Table 3. Summary Statistics

The table presents summary statistics for the variables used in the study. Panels A shows summary statistics for the stock-month sample. Panel B reports correlations for the same sample, and Panel C shows summary statistics for the variables used in the return regressions (stock-day sample). Panel D shows summary statistics for the stock-day sample. Panels A, C, and D present separate statistics for the S&P 500 and the Russell 3000 universes. The samples range between January 2000 and December 2012.

Panel A: Monthly Sample

S&P 500

	N	Mean	Std Dev	Min	Median	Max
Daily stock volatility (%)	67,261	2.180	1.470	0.612	1.770	10.800
Log(abs(VR 5 days/(5*1 day)))	67,261	-1.810	0.900	-3.880	-1.590	-0.588
ETF ownership (%)	67,261	2.060	1.580	0.012	1.730	9.610
Index Fund ownership (%)	67,261	5.970	2.080	0.440	5.800	12.300
Active Fund ownership (%)	67,261	17.900	6.430	0.803	17.600	36.600
log(Mktcap ($m))	67,261	9.250	1.070	4.760	9.180	11.300
1/Price	67,261	0.040	0.037	0.006	0.030	0.578
Amihud	67,261	0.000	0.001	0.000	0.000	0.026
Bid-ask spread (%)	67,261	0.287	0.528	0.022	0.080	3.190
Book-to-Market	67,261	0.466	0.394	0.029	0.361	2.560
Past 12-month Return	67,261	0.090	0.378	-0.773	0.067	2.230
Gross Profitability	67,261	0.308	0.224	-0.242	0.271	1.080

Russell 3000

	N	Mean	Std Dev	Min	Median	Max
Daily stock volatility (%)	289,563	2.620	1.660	0.612	2.170	10.800
Log(abs(VR 5 days/(5*1 day)))	289,563	-1.740	0.906	-3.880	-1.530	-0.588
ETF ownership (%)	289,563	2.360	1.980	0.012	1.850	9.610
Index Fund ownership (%)	289,563	4.660	2.480	0.315	4.440	12.300
Active Fund ownership (%)	289,563	16.500	8.280	0.500	16.500	36.600
log(Mktcap ($m))	289,563	7.350	1.410	4.170	7.100	11.300
1/Price	289,563	0.061	0.060	0.006	0.044	0.578
Amihud	289,563	0.011	0.031	0.000	0.002	0.335
Bid-ask spread (%)	289,563	0.333	0.484	0.022	0.152	3.190
Book-to-Market	289,563	0.518	0.423	0.029	0.414	2.560
Past 12-month Return	289,563	0.144	0.479	-0.773	0.087	2.230
Gross Profitability	289,563	0.301	0.244	-0.242	0.268	1.080

Panel B: Correlations

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Daily stock volatility (%)	(1)	1.00										
Log(abs(VR 5 days/(5*1 day)))	(2)	0.01	1.00									
ETF ownership (%)	(3)	0.00	-0.02	1.00								
Index Fund ownership (%)	(4)	-0.03	-0.06	0.34	1.00							
Active Fund ownership (%)	(5)	-0.04	-0.10	0.21	0.41	1.00						
log(Mktcap ($m))	(6)	-0.31	-0.07	-0.03	0.29	0.27	1.00					
1/Price	(7)	0.35	0.02	-0.05	-0.12	-0.23	-0.39	1.00				
Amihud	(8)	0.18	0.09	-0.19	-0.26	-0.35	-0.39	0.30	1.00			
Bid-ask spread (%)	(9)	0.23	0.06	-0.39	-0.26	-0.27	-0.25	0.28	0.49	1.00		
Book-to-Market	(10)	0.21	0.03	0.12	0.08	-0.15	-0.24	0.34	0.16	0.19	1.00	
Past 12-month Return	(11)	-0.11	-0.03	-0.05	-0.03	0.02	0.05	-0.17	-0.10	-0.12	-0.31	1.00
Gross Profitability	(12)	0.01	-0.03	-0.01	0.02	0.20	0.00	-0.05	-0.06	-0.03	-0.26	0.04

Table 3. Summary Statistics (Cont.)

Panel C: Variables Used in Return Regressions (Daily Frequency)

S&P 500

	N	Mean	Std Dev	Min	Median	Max
Ret(t) (%)	1,123,157	0.062	2.114	-9.459	0.023	10.403
Ret(t+1,t+5) (%)	1,123,157	0.244	4.520	-19.922	0.244	21.330
Ret(t+1,t+10) (%)	1,123,157	0.465	6.146	-23.823	0.514	25.242
Ret(t+1,t+20) (%)	1,123,157	0.898	8.605	-31.350	1.071	33.629
net(ETF Flows) (%)	1,123,157	0.000	0.000	-0.001	0.000	0.001

Russell 3000

	N	Mean	Std Dev	Min	Median	Max
Ret(t) (%)	5,014,804	0.070	2.373	-9.459	0.000	10.405
Ret(t+1,t+5) (%)	5,014,804	0.228	5.054	-19.923	0.200	21.333
Ret(t+1,t+10) (%)	5,014,804	0.452	6.829	-23.825	0.456	25.245
Ret(t+1,t+20) (%)	5,014,804	0.873	9.626	-31.355	0.950	33.641
net(ETF Flows) (%)	5,014,804	0.000	0.000	-0.001	0.000	0.001

Panel D: Daily Sample

	N	Mean	Std Dev	Min	Median	Max
ETF ownership (%)	1,029,618	2.320	1.360	0.030	2.190	9.770
Abs(mispricing) (bps)	1,024,398	0.256	0.287	0.028	0.161	2.390
Net(mispricing) (bps)	1,002,866	-0.049	0.237	-1.160	-0.004	0.466
Intraday volatility (%)	1,029,618	0.019	0.015	0.005	0.014	0.123
Variance Ratio (VR 15)	1,000,903	-2.040	1.040	-8.230	-1.820	-0.546
Share lending fee (%)	1,029,618	0.213	1.110	0.000	0.099	72.400

Russell 3000

	N	Mean	Std Dev	Min	Median	Max
ETF ownership (%)	4,554,404	2.660	1.780	0.030	2.320	9.770
Abs(mispricing) (%)	4,357,317	0.324	0.349	0.028	0.211	2.390
Net(mispricing) (%)	4,260,036	-0.108	0.334	-1.160	-0.023	0.466
Intraday volatility (%)	4,554,404	0.021	0.015	0.005	0.017	0.123
Variance Ratio (VR 15)	4,485,442	-2.680	1.290	-8.230	-2.440	-0.546
Share lending fee (%)	4,554,404	0.444	2.240	0.000	0.132	132.000

Table 4. ETF Ownership and Stock Volatility

The table reports estimates from OLS regressions of daily volatility on ETF ownership and controls. In Columns (1) to (3), the sample consists of S&P 500 stocks, and in Columns (4) to (6) the sample consists of Russell 3000 stocks. The frequency of the observations is monthly, and volatility is computed using all daily returns within the month. The dependent variable and the ownership variables have been standardized by subtracting the mean and dividing by the standard deviation. Variable descriptions are provided in Appendix Table A1. Standard errors are double-clustered at the stock and month levels. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between January 2000 and December 2012.

Dependent variable:	Daily stock volatility					
Sample:	S&P 500			Russell 3000		
	(1)	(2)	(3)	(4)	(5)	(6)
ETF ownership	0.132***	0.127***	0.073***	0.052***	0.042***	0.033***
	(4.828)	(4.700)	(4.488)	(4.606)	(3.784)	(4.661)
log(Mktcap (t-1))	0.048	0.038	-0.035*	-0.096***	-0.116***	-0.104***
	(1.271)	(1.010)	(-1.691)	(-3.799)	(-4.550)	(-5.630)
1/Price (t-1)	1.574**	1.502**	0.073	0.814***	0.905***	0.172
	(2.446)	(2.343)	(0.209)	(2.954)	(3.291)	(0.918)
Amihud (t-1)	-9.242	-3.037	-8.665	0.195	0.305	-0.127
	(-0.604)	(-0.206)	(-1.238)	(0.704)	(1.087)	(-0.690)
Bid-ask spread (t-1)	-7.346**	-7.215**	-4.544***	-8.894***	-8.168***	-5.513***
	(-2.085)	(-2.041)	(-2.826)	(-2.989)	(-2.723)	(-3.337)
Book-to-Market (t-1)	0.531***	0.530***	0.190***	0.332***	0.325***	0.158***
	(9.162)	(9.172)	(7.110)	(9.539)	(9.381)	(8.032)
Past 12-month Return (t-1)	0.025	0.016	0.055***	0.056***	0.061***	0.050***
	(0.668)	(0.430)	(2.714)	(2.791)	(3.051)	(4.294)
Gross Profitability (t-1)	0.454***	0.493***	0.189***	0.032	0.042	0.041
	(4.143)	(4.400)	(4.213)	(0.607)	(0.788)	(1.484)
Index Fund Ownership		0.028**	0.008		0.021***	0.010***
		(2.287)	(1.503)		(3.537)	(2.982)
Active Fund Ownership		0.062***	0.030***		0.060***	0.037***
		(3.884)	(4.227)		(6.509)	(6.994)
Volatility (t-1)			0.295***			0.217***
			(17.911)			(17.608)
Volatility (t-2)			0.170***			0.155***
			(9.312)			(18.587)
Volatility (t-3)			0.197***			0.186***
			(11.930)			(21.707)
Stock fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	67,261	67,261	65,866	289,563	289,563	275,962
Adjusted R^2	0.645	0.647	0.743	0.593	0.595	0.664

Table 5. Quasi-Natural Experiment Based on the Russell Index Reconstitution

The table reports estimates from a quasi-natural experiment relying on the reconstitution of the Russell 1000 and Russell 2000 indexes. The frequency of the data is monthly at the stock level. In Panel A, the dependent variable is ETF ownership. The explanatory variables are a dummy for inclusion in the Russell 2000, for stocks in the Russell 1000 before index reconstitution (Columns (1)–(5)); and a dummy for inclusion in the Russell 1000, for stocks in the Russell 2000 before index reconstitution (Columns (6)–(10)). Stocks are ranked in terms of market capitalization in May of each year. Different ranges of this rank around the cutoff are used for inclusion in the sample: 100 stocks on each side (Columns (1) and (6)), 200 stocks on each side (Columns (2) and (7)), 300 stocks on each side (Columns (3) and (8)), 400 stocks on each side (Columns (4) and (9)), and 500 stocks on each side (Columns (5) and (10)). The same stocks enter the sample in the June after index reconstitution and remain in the sample until May of the next year, except if delistings occur. The controls in all panels include logged market capitalization, the lagged inverse share price ratio, the lagged Amihud (2002) ratio, the lagged average bid-ask spread, the lagged book-to-market ratio, lagged past 12 months' returns, lagged gross profitability (as in Novy-Marx, 2013), lagged volatility, index fund ownership, and active fund ownership. In Panels B, C, and D, the dependent variable is daily stock volatility (computed using all daily returns within a month). The main explanatory variable is instrumented ETF ownership. The instruments are either a dummy for inclusion in the Russell 2000 for stocks in the Russell 1000 before reconstitution (Columns (1)–(5)) or a dummy for inclusion in the Russell 1000 for stocks in the Russell 2000 before reconstitution (Columns (7)–(10)). The same bandwidths around the cutoff are used to restrict the sample as in Panel A. The regressions in Panel B, as well as in Panel A, include a linear specification of the ranking variable (not reported). For the two-stage estimation whose results are reported in Panel B, ETF ownership, as well as ownership by index and active funds, is standardized by subtracting the mean and dividing by the standard deviation in the estimation sample. In Panel C, the dependent variable is daily volatility. The main explanatory variable is an interaction between the dummy variables for index inclusion and average ETF ownership in the sample of Russell 2000 stocks. Other explanatory variables include interactions between the index inclusion dummies and ownership by index and active funds, a time trend. Panel C also includes controls for average ETF ownership as well as average index funds ownership and mutual fund ownership in the Russell 2000 sample. In Panel D, the dependent variable is daily volatility. The main explanatory variable is an interaction between the dummy variable for Russell 1000 inclusion and the stock-level ratio of ETF ownership to other fund (either Index or Active fund) ownership in May before reconstitution. The dependent variable and the ownership variables have been standardized by subtracting the mean and dividing by the standard deviation. Variable descriptions are provided in Appendix Table A1. Standard errors are double-clustered at the stock and month level. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between June 2000 and May 2007.

Panel A: First-Stage Regressions, First Degree Polynomial

Dependent variable:	ETF ownership									
Instrument:	Switch to the Russell 2000					Switch to the Russell 1000				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Switching index	0.244***	0.345***	0.451***	0.518***	0.490***	-0.145***	-0.404***	-0.415***	-0.383***	-0.396***
	(6.397)	(7.100)	(7.644)	(8.825)	(8.991)	(-2.842)	(-8.461)	(-9.822)	(-8.115)	(-8.528)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,885	10,152	16,162	22,403	28,742	6,528	12,798	18,360	23,744	29,186
Adjusted R^2	0.375	0.390	0.378	0.346	0.336	0.410	0.421	0.455	0.471	0.461

Panel B: Second-Stage Regressions, First Degree Polynomial

Dependent variable:	Daily stock volatility									
Instrument:	Switch to the Russell 2000					Switch to the Russell 1000				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ETF ownership (instrumented)	0.498**	0.362**	0.263***	0.171**	0.240***	0.803**	0.258***	0.202***	0.220***	0.168***
	(2.206)	(2.603)	(3.229)	(2.617)	(3.632)	(2.232)	(3.660)	(3.357)	(3.569)	(3.517)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,885	10,152	16,162	22,403	28,742	6,528	12,798	18,360	23,744	29,186

Table 5. Quasi-Natural Experiment Based on the Russell Index Reconstitution (Cont.)

Panel C: Interaction with Average ETF Ownership

Dependent variable:	Daily stock volatility									
Instrument:	Switch to the Russell 2000					Switch to the Russell 1000				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ETF ownership × Switch	0.294**	0.284**	0.137	0.051	0.029	-0.262	-0.398***	-0.483***	-0.536***	-0.625***
	(2.040)	(2.118)	(1.049)	(0.432)	(0.228)	(-1.590)	(-3.592)	(-3.848)	(-3.469)	(-3.321)
Index funds ownership × Switch	-0.082***	-0.052***	-0.028**	-0.026**	-0.023**	0.069***	0.014	-0.005	-0.011	-0.016
	(-5.246)	(-3.993)	(-2.425)	(-2.443)	(-2.025)	(3.409)	(0.906)	(-0.377)	(-0.867)	(-0.975)
Active funds ownership × Switch	0.280***	0.052	0.021	0.025	0.018	-0.146**	0.005	0.077	0.093	0.155*
	(3.983)	(0.916)	(0.365)	(0.445)	(0.332)	(-2.103)	(0.086)	(1.053)	(1.369)	(1.859)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Aggregate ownership controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time trend, interacted with switch	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,887	10,157	16,173	22,420	28,769	6,532	12,808	18,372	23,761	29,204
Adjusted R^2	0.631	0.606	0.602	0.592	0.593	0.495	0.473	0.476	0.484	0.488

Panel D: Ratio of ETF to Other Fund (Index Fund or Active Fund) Ownership

Dependent variable:	Daily stock volatility									
Instrument:	Switch to the Russell 1000									
Ratio:	ETF / Index Fund					ETF / Active Fund				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Switch	-0.074*	-0.027	-0.021	-0.006	0.009	-0.082**	-0.061*	-0.058**	-0.043*	-0.029
	(-1.905)	(-0.969)	(-0.861)	(-0.276)	(0.439)	(-2.093)	(-1.898)	(-2.203)	(-1.804)	(-1.387)
Ratio × Switch	-0.000	-0.001***	-0.001***	-0.001***	-0.002***	-0.004	-0.003	-0.003**	-0.003**	-0.003**
	(-0.328)	(-3.045)	(-3.388)	(-4.010)	(-4.409)	(-1.557)	(-1.666)	(-2.389)	(-2.311)	(-2.256)
Ratio	-0.003***	-0.002***	-0.002***	-0.002***	-0.002***	-0.006**	-0.008***	-0.008***	-0.008***	-0.008***
	(-3.899)	(-5.000)	(-5.811)	(-8.427)	(-8.700)	(-2.621)	(-7.319)	(-9.607)	(-11.074)	(-11.513)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	5,629	10,934	15,684	20,200	24,643	5,558	10,775	15,376	19,897	24,363
Adjusted R^2	0.441	0.396	0.400	0.402	0.413	0.434	0.394	0.396	0.397	0.405

Table 6. ETF Ownership and Price Efficiency: Variance Ratios

The table reports estimates from OLS regressions of variance ratios on ETF ownership and controls (Panel A), as well as IV regressions (Panels B and C). In Panel A, Columns (1) and (2), the sample consists of S&P 500 stocks, and in Columns (3) and (4), the sample consists of Russell 3000 stocks. The frequency of the observations is monthly for VR 15 seconds and quarterly for VR 5 days. *VR 15 seconds* is the absolute value of the ratio of the variance of 15-second log returns on day *t* and 3 times the variance of 5-second log returns on day *t* – 1, using data from the TAQ database and averaging the numerator and denominator within a month. *VR 5 days* is the absolute value of the ratio of the variance of 5-day returns in a given quarter and 5 times the variance of one-day returns in the same quarter. Panels B and C show IV regressions for 15-second variance ratios (Panel B) and for 5-day variance ratios (Panel C) based on the Russell 1000/Russell 2000 inclusion experiment. The dependent variable and the ownership variables have been standardized by subtracting the mean and dividing by the standard deviation. The controls in all panels include logged market capitalization, the lagged inverse share price ratio, the lagged Amihud (2002) ratio, the lagged average bid-ask spread, the lagged book-to-market ratio, lagged past 12 months' returns, lagged gross profitability (as in Novy-Marx, 2013), lagged volatility, index fund ownership, and active fund ownership. Variable descriptions are provided in Appendix Table A1. Standard errors are double-clustered at the stock and time level. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between January 2000 and December 2012 in Panel A, and June 2000 and May 2007 in Panels B and C.

Panel A: Variance Ratios (OLS)

Sample:	S&P 500		Russell 3000	
Dependent Variable:	VR 15 seconds	VR 5 days	VR 15 seconds	VR 5 days
	(1)	(2)	(3)	(4)
ETF ownership	0.109***	0.049*	0.061***	0.013
	(4.485)	(1.809)	(7.620)	(1.188)
Controls	Yes	Yes	Yes	Yes
Stock fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Observations	56,623	22,887	268,418	100,836
Adjusted R^2	0.473	0.032	0.488	0.041

Panel B: Variance Ratios – 15 Seconds (IV)

Dependent variable:	VR 15 seconds									
Instrument:	Switch to the Russell 2000					Switch to the Russell 1000				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ETF ownership (instrumented)	-0.074	0.316*	0.238**	0.152	0.283***	0.436	0.016	0.055	-0.007	-0.057
	(-0.258)	(1.686)	(2.076)	(1.597)	(2.666)	(1.382)	(0.193)	(0.812)	(-0.099)	(-1.064)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,809	10,008	15,947	22,114	28,387	6,469	12,673	18,175	23,493	28,863

Panel C: Variance Ratios – 5 Days (IV)

Dependent variable:	VR 5 days									
Instrument:	Switch to the Russell 2000					Switch to the Russell 1000				
Bandwidth:	± 100	± 200	± 300	± 400	± 500	± 100	± 200	± 300	± 400	± 500
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
ETF ownership (instrumented)	0.248	0.232	0.241*	0.189*	0.190*	0.584	0.470**	0.490**	0.537**	0.515**
	(1.108)	(1.556)	(1.886)	(1.828)	(1.806)	(1.284)	(2.343)	(2.720)	(2.476)	(2.505)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Linear polynomials of rank	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,540	3,214	5,124	7,118	9,148	2,071	4,048	5,798	7,498	9,217

Table 7. Price Reversals

The table reports estimates from OLS regressions of one- and multiday returns on ETF flows and controls. The specifications also include the *k*-period lagged dependent variable, where *k* is set to have the return-measurement horizon end in *t*-1. In Columns (1) to (4), the sample consists of S&P 500 stocks, and in Columns (5) to (8), the sample consists of Russell 3000 stocks. The frequency of the observations is daily. Returns are in percentages. Flows have been standardized by subtracting the mean and dividing by the standard deviation. The controls in all panels include logged market capitalization, the lagged inverse share price ratio, the lagged Amihud (2002) ratio, the lagged average bid-ask spread, the lagged book-to-market ratio, lagged past 12 months' returns, lagged gross profitability (as in Novy-Marx, 2013), order imbalance, and the lagged dependent variable. Variable descriptions are provided in Appendix Table A1. Standard errors are clustered at the day level and are computed using the Newey and West (1987) estimator. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between January 2000 and December 2012.

Sample:	S&P 500				Russell 3000			
Dependent variable:	Ret(t)	Ret(t+1,t+5)	Ret(t+1,t+10)	Ret(t+1,t+20)	Ret(t)	Ret(t+1,t+5)	Ret(t+1,t+10)	Ret(t+1,t+20)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
net(ETF Flows)	0.167***	-0.024	-0.054**	-0.072**	0.062***	-0.011	0.000	-0.032**
	(17.420)	(-1.328)	(-2.248)	(-2.278)	(13.176)	(-1.136)	(0.027)	(-2.104)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,120,058	1,120,058	1,120,058	1,120,058	4,988,503	4,988,503	4,988,503	4,988,503
Adjusted R^2	0.361	0.311	0.289	0.287	0.332	0.271	0.243	0.240

Table 8. Evidence on the Arbitrage Channel

The table reports estimates from OLS regressions of intraday volatility (Panel A) and intraday variance ratios (Panel B) on absolute stock-level mispricing in the prior period, interacted with measures of arbitrage costs. The frequency is daily, and the observations are at the stock level. The sample includes S&P 500 stocks. In Columns (2)–(4), arbitrage cost is captured by the bid-ask spread in the prior day, and in Columns (5)–(7), by the average share-lending fee in the month. For both measures of arbitrage costs, we construct dummy variables denoting whether the stock is in the top half of the distribution of that measure in the relevant period. In Columns (3) and (6), we restrict the sample to observations for which the stock-level mispricing is positive. In Columns (4) and (7), we restrict the sample to observations for which the stock-level mispricing is negative. The controls in all panels include logged market capitalization, the lagged inverse share price ratio, the lagged Amihud (2002) ratio, the lagged average bid-ask spread, the lagged book-to-market ratio, lagged past 12 months' returns, lagged gross profitability (as in Novy-Marx, 2013), lagged returns, the lagged dependent variable, and the absolute mispricing in period *t* - 2. Variable descriptions are provided in Appendix Table A1. Standard errors are double-clustered at the stock and day level. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between January 2000 and December 2012.

Panel A: Intraday Volatility

Dependent variable:	Intraday stock volatility						
	All	All	Misp > 0	Misp < 0	All	Misp > 0	Misp < 0
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Abs(Mispricing) (t-1)	0.023***	0.043***	0.076***	0.035***	0.020***	0.048***	0.018***
	(6.897)	(8.025)	(10.300)	(6.988)	(4.642)	(7.530)	(4.302)
× I(High bid-ask spread)		-0.053***	-0.055***	-0.043***			
		(-6.554)	(-4.257)	(-6.392)			
× I(High lending fee)					-0.009*	0.003	-0.010**
					(-1.956)	(0.444)	(-2.548)
High bid-ask spread		0.042***	0.038***	0.045***			
		(7.186)	(6.318)	(7.130)			
High lending fee					-0.005	-0.003	-0.005
					(-1.565)	(-0.957)	(-1.433)
ETF ownership (t-1)	0.022***	0.022***	-0.004	0.031***	0.021***	-0.006	0.031***
	(3.611)	(4.217)	(-0.692)	(5.531)	(4.042)	(-1.227)	(5.417)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stock fixed effects	Yes	No	No	No	No	No	No
Observations	1,022,548	1,022,548	509,240	513,308	1,022,548	509,240	513,308
Adjusted R^2	0.549	0.500	0.505	0.498	0.499	0.504	0.497

Table 8. Evidence on the Arbitrage Channel (Cont.)

Panel B: Intraday Variance Ratio

Dependent variable:	Intraday variance ratio (VR 15)						
	All	All	Misp > 0	Misp < 0	All	Misp > 0	Misp < 0
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Abs(Mispricing) (t-1)	0.002	0.030***	0.055***	0.024***	0.007	0.026***	0.006
	(0.741)	(4.658)	(6.920)	(3.531)	(1.440)	(4.012)	(1.276)
× I(High bid-ask spread)		-0.056***	-0.062***	-0.046***			
		(-5.781)	(-4.739)	(-4.765)			
× I(High lending fee)					-0.011**	-0.001	-0.012**
					(-2.083)	(-0.125)	(-2.127)
High bid-ask spread		0.120***	0.119***	0.120***			
		(12.401)	(11.907)	(11.991)			
High lending fee					-0.009**	-0.006	-0.010**
					(-2.115)	(-1.341)	(-2.298)
ETF ownership (t-1)	0.037***	0.030***	0.016**	0.032***	0.027***	0.010	0.030***
	(4.864)	(4.401)	(2.123)	(4.599)	(3.825)	(1.297)	(4.161)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Stock fixed effects	Yes	No	No	No	No	No	No
Observations	983,625	983,625	490,321	493,304	983,625	490,321	493,304
Adjusted R^2	0.245	0.186	0.188	0.185	0.180	0.183	0.179

Table 9. Volatility and ETF Ownership in the Time-Series

Panel A reports estimates from a time-series regression at the monthly frequency of average daily volatility in a given month across the stocks in the CRSP universe on lagged average ETF ownership for the same universe. The controls include lagged average volatility, lagged average index and active fund ownership (IF and AF variables, respectively), and a time trend. The same regression is performed in first differences, excluding the time trend. Panel B reports estimates from time-series regressions of average volatility in each quintile of ETF ownership on lagged average ETF ownership across all stocks in the CRSP universe. Panel C reports estimates from time-series regressions of the changes in the average volatility in each quintile of ETF ownership on lagged changes in the average ETF ownership across all stocks in the CRSP universe. The other ownership variables are also computed as averages across all stocks. The dependent as well as the ownership variables have been standardized by subtracting the mean and dividing by the standard deviation. Variable descriptions are provided in Appendix Table A1. *t-statistics* are presented in parentheses. ***, **, and * represent statistical significance at the 1%, 5%, and 10% levels, respectively. The sample ranges between January 2000 and December 2012.

Panel A: Regression of Volatility on ETF Ownership

Dependent Variable:	Volatility (t+1)	ΔVolatility (t+1)
	(1)	(2)
ETF ownership (t)	0.216***	
	(3.938)	
IF ownership (t)	-0.066	
	(-1.106)	
AF ownership (t)	0.082	
	(1.629)	
Volatility (t)	0.701***	
	(12.939)	
Trend	-0.001	
	(-0.483)	
ΔETF ownership (t)		0.344***
		(4.388)
ΔIF ownership (t)		-0.052
		(-0.632)
ΔAF ownership (t)		0.030
		(0.358)
ΔVolatility (t)		-0.150*
		(-1.934)
Observations	149	148
Adjusted R^2	0.721	0.151

Table 9. Volatility and ETF Ownership in the Time-Series (Cont.)

Panel B: Regression of Volatility and ETF Ownership, by Quintiles of ETF Ownership

Dependent Variable:	Volatility (t+1)				
Quintile of ETF ownership:	Smallest	2	3	4	Largest
ETF ownership (t)	0.189***	0.213***	0.206***	0.216***	0.247***
	(3.532)	(4.075)	(3.890)	(4.154)	(3.856)
Index mutual funds ownership (t)	-0.063	-0.068	-0.049	-0.049	-0.093
	(-1.062)	(-1.144)	(-0.823)	(-0.863)	(-1.381)
Active mutual funds ownership (t)	0.025	0.060	0.091*	0.089*	0.150**
	(0.497)	(1.193)	(1.799)	(1.823)	(2.560)
Volatility (t)	0.760***	0.688***	0.656***	0.638***	0.667***
	(15.050)	(12.508)	(11.261)	(10.779)	(11.548)
Trend	-0.000	-0.000	-0.001	-0.000	-0.002
	(-0.027)	(-0.238)	(-0.606)	(-0.183)	(-1.216)
Observations	149	149	149	149	149
Adjusted R^2	0.759	0.691	0.648	0.662	0.726

Panel C: Regression of Changes in Volatility and ETF Ownership, by Quintiles of ETF Ownership

Dependent Variable:	ΔVolatility (t+1)				
Quintile of ETF ownership:	Smallest	2	3	4	Largest
ΔETF ownership (t)	0.297***	0.289***	0.326***	0.312***	0.384***
	(3.886)	(3.654)	(4.161)	(4.189)	(4.482)
ΔIndex mutual funds ownership (t)	-0.049	-0.064	-0.040	-0.061	-0.030
	(-0.607)	(-0.777)	(-0.479)	(-0.764)	(-0.326)
ΔActive mutual funds ownership (t)	-0.016	0.009	0.053	0.044	0.051
	(-0.195)	(0.114)	(0.638)	(0.549)	(0.551)
ΔVolatility (t)	-0.098	-0.149*	-0.184**	-0.200**	-0.197**
	(-1.219)	(-1.859)	(-2.356)	(-2.596)	(-2.540)
Observations	148	148	148	148	148
Adjusted R^2	0.120	0.128	0.157	0.159	0.150

Figure 1: Illustration of the Propagation of Liquidity Shocks via Arbitrage

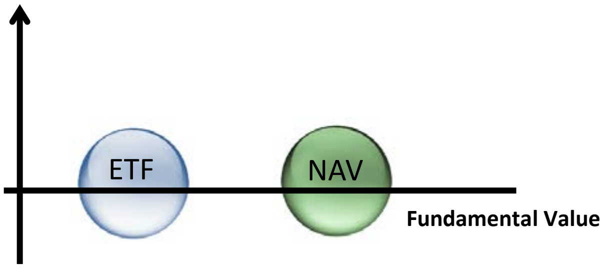

Figure 1a. Initial equilibrium

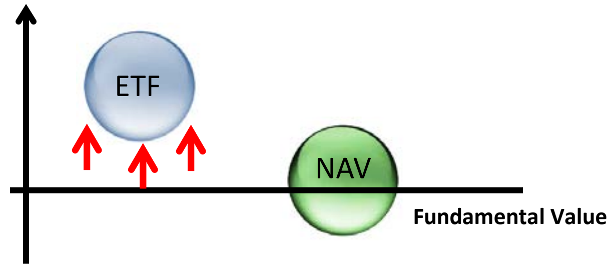

Figure 1b. Liquidity shock to ETF

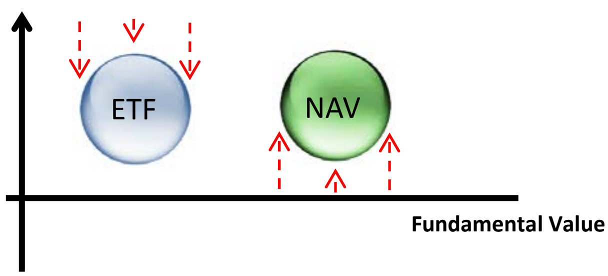

Figure 1c. Initial outcome of arbitrage: the shock is propagated to the NAV, and the ETF price starts reverting to the fundamental value.

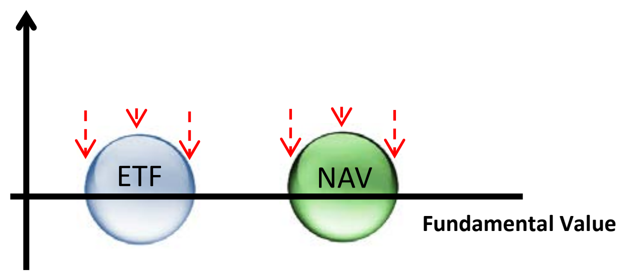

Figure 1d. Equilibrium reestablished: after some time, both the ETF price and the NAV revert to the fundamental value.

Figure 2: Illustration of the Propagation of a Fundamental Shock with Price Discovery Occurring in the ETF Market

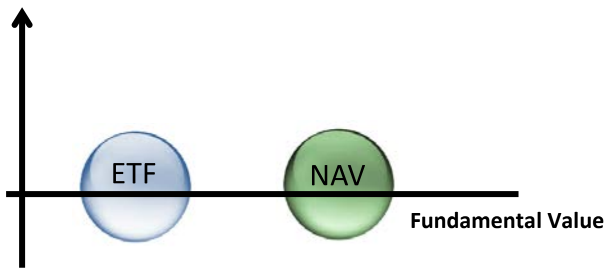

Figure 2a. Initial equilibrium

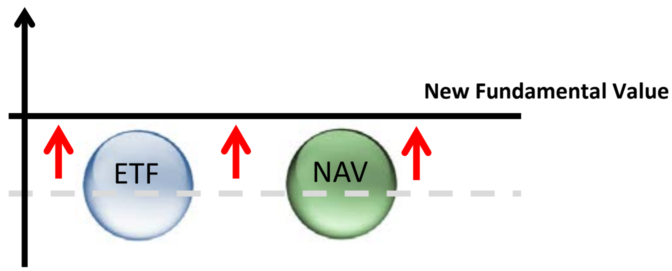

Figure 2b. Shock to fundamental value

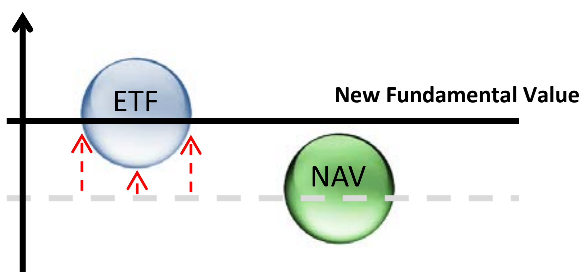

Figure 2c. Price discovery takes place in the ETF market. The ETF price moves to the new fundamental value.

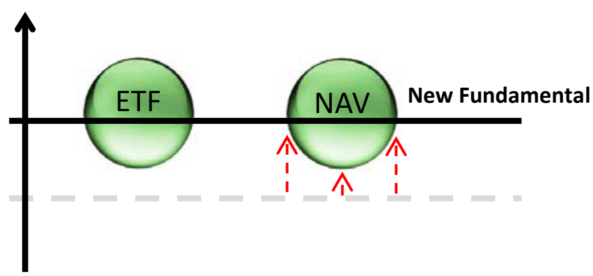

Figure 2d. After a delay, the NAV catches up with the new fundamental.

Figure 3: Fund Ownership around the Russell Cutoff

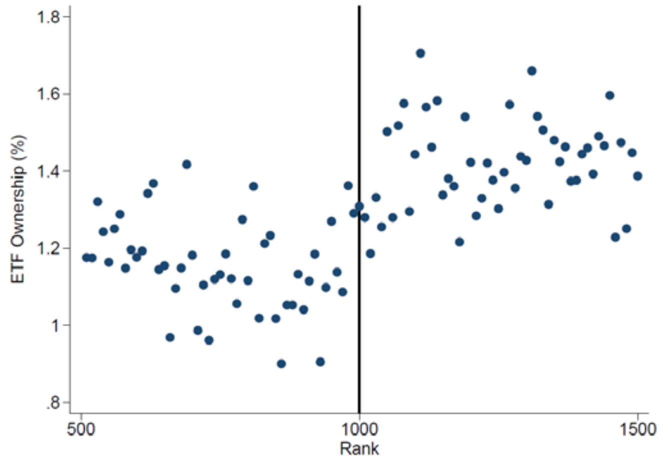

Figure 3a. ETF Ownership

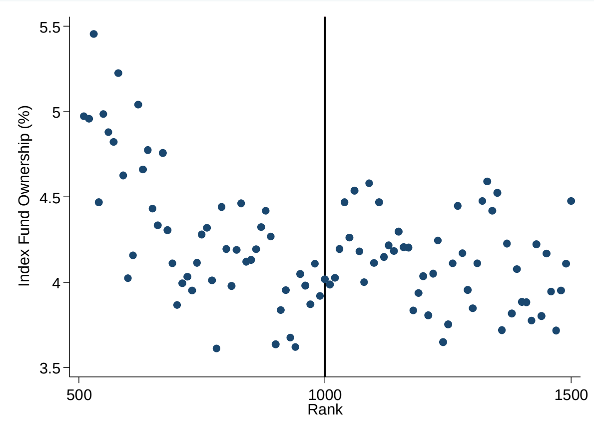

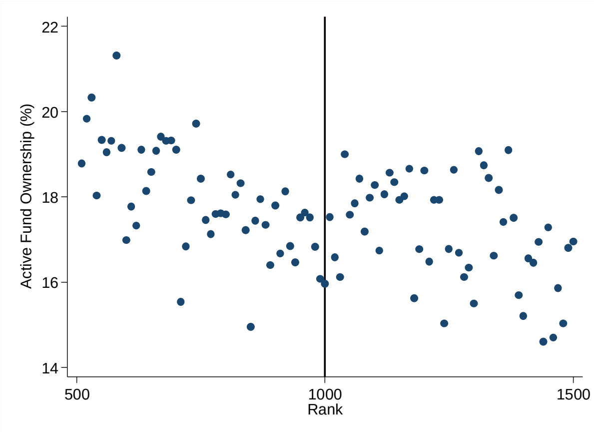

Figure 3b. Index Fund Ownership

Figure 3c. Active Fund Ownership

The figure reports average ownership (in %) by ETFs (3a), Index Funds (3b), and Active Funds (3c) f`or stocks ranked by market capitalization and included in the Russell 3000. The average is computed first by ranking over time, then across the ranking in bins of 10 stocks. The vertical line denotes the 1000th rank. The sample ranges between January 2000 and May 2007.

Appendix Table A1. Variable Definitions

Variable	Description	Source
ETF ownership	The sum of the ownership of all ETFs holding the stock, using the most recent quarterly investment company reports for equity ETFs. The lagged quarterly portfolio weights are interacted with daily ETF AUM and daily stock capitalization to compute daily ownership. The monthly variable is defined accordingly.	Thomson-Reuters, CRSP, Bloomberg
Index (or active) mutual fund ownership	The sum of the ownership by all index (or active) mutual funds holding the stock, using the most recent quarterly investment company reports.	Thomson-Reuters, CRSP Mutual Fund, and MFLinks
Daily volatility	Standard deviation of daily stock returns within a month.	CRSP
Intraday volatility	Standard deviation of second-by-second intraday returns.	TAQ
Variance ratio 15 seconds	The ratio of 15-second log return variance divided by 3 times the 5-second log return variance minus 1. The numerator and denominator are computed using returns within a day and averaged over a month. The dependent variable in the regressions is the logarithm of the absolute value of this difference.	TAQ
Variance ratio 5 days	The ratio of 5-day return variance divided by 5 times the 1-day return variance minus 1. The numerator and denominator are computed using daily and 5-day returns within a quarter. The dependent variable in the regressions is the logarithm of the absolute value of this difference.	CRSP
Net(ETF flows)	Stock-day-level measure. Weighted average of the percentage change in ETF shares outstanding across the ETFs holding the stock. The weight is ETF ownership of the stock.	Bloomberg, Compustat
Ret(t_1, t_2)	The total return of the stock between the close of t_1 and the close of t_2.	CRSP
Abs(mispricing)	Sum of absolute dollar mispricing across all the ETFs holding the stock divided by stock capitalization (Equation (5)). Dollar mispricing is the product of ETF mispricing (i.e., the difference between the ETF price and its NAV, as a fraction of the ETF price) times dollar holdings of an ETF in the stock.	Thomson-Reuters, CRSP, Bloomberg
Net(mispricing)	Similar construction to abs(mispricing). The only difference is that the ETF-level mispricing is not in absolute value.	Thomson-Reuters, CRSP, Bloomberg
Lending fees	Share-lending fee at the security level, 7-day average. Average within the month.	Markit
log(Mktcap)	The logged market capitalization of the stock (in $ millions) at the end of the month.	CRSP
1/Price	The inverse of the nominal share price at the end of the month.	CRSP
Amihud ratio	Absolute return scaled by dollar volume in $million, average within the month. Based on Amihud (2002).	CRSP
Bid-ask spread	The quoted spread divided by the bid-ask midpoint. End-of-month value.	CRSP
Book-to-Market	Book value of assets / market value of assets.	CRSP, Compustat
Past 12-month Return	Cumulative returns in the previous 12 months.	CRSP
Gross Profitability	(Revenue – cost of goods sold) / total assets. Following Novy-Marx (2013)	Compustat
Churn ratio 1	This measure follows Cella, Ellul, and Giannetti (2013) in computing the investor-level churn ratio, which is then aggregated at the stock level using ownership weights.	CRSP, 13-F
Churn ratio 2	This measure uses an investor-level churn ratio that is computed as the minimum between the absolute value of buys and sells divided by prior quarter holdings. Buys and sells use prior quarter prices.	CRSP, 13-F