Working Paper Series 651
(ISSN 1211-3298)

Does Index Arbitrage Distort the Market Reaction to Shocks?

Stanislav Anatolyev
Sergei Seleznev
Veronika Selezneva

CERGE-EI
Prague, December 2019

ISBN 978-80-7343-458-8 (Univerzita Karlova, Centrum pro ekonomický výzkum a doktorské studium)
ISBN 978-80-7344-515-7 (Národohospodářský ústav AV ČR, v. v. i.)

Does Index Arbitrage Distort the Market Reaction to Shocks?*

Stanislav Anatolyev
CERGE-EI and NES

Sergei Seleznev
INECO Capital Ltd

Veronika Selezneva†
CERGE-EI

Abstract

We show that ETF arbitrage distorts the market reaction to fundamental shocks. We confirm this hypothesis by creating a new measure of the intensity of arbitrage transactions at the individual stock level and using an event study analysis to estimate the market reaction to economic shocks. Our measure of the intensity of arbitrage is the probability of simultaneous trading of ETF shares with shares of underlying stocks estimated using high frequency data. Our approach is direct, and it accounts for statistical arbitrage, passive investment strategies, and netting of arbitrage positions over the day, which the existing measures cannot do. We conduct several empirical tests, including the use of a quasi-natural experiment, to confirm that our measure captures fluctuations in the intensity of arbitrage transactions. We focus on oil shocks because they contain a large idiosyncratic component which facilitates identification of our mechanism and interpretation of the results. Oil shocks are identified using weekly oil inventory announcements.

Keywords: high-frequency data, stock market, ETF, arbitrage intensity, oil shock, market efficiency

JEL classifications: G12, G14, G23, Q43

*We thank Lubos Pastor, Carol Alexander, and seminar participants at the University of Sussex, and conference participants of NASMES-2019, CEMA-2019, and ESEM-2019 for helpful comments and suggestions. Grant GACR 17-27567S from the Czech Science Foundation is gratefully acknowledged.

†Corresponding author, veronika.selezneva@cerge-ei.cz. CERGE-EI, a joint workplace of Charles University and the Economics Institute of the Czech Academy of Sciences, Politickych vězňů 7, 11121 Prague 1

1 Introduction

The recent growth of exchange traded funds (ETFs) has intensified index trading by opening access to the market for a broader set of market participants. Although the benefits of increased participation are apparent, the side effects of the dramatic change in the market structure are yet to be fully assessed.

Recent research on equity ETFs examines whether ETF arbitrage can damage the underlying markets. The literature faces two main challenges: how to measure the intensity of arbitrage transactions and how to assess distortions in the underlying markets. In this paper, we innovate along both dimensions. We show that ETF arbitrage distorts the market reaction to fundamental shocks. In order to do that, we create a new measure of arbitrage intensity at the individual stock level.

A standard way to measure the intensity of ETF arbitrage at the stock level is to use changes in ETF holdings. Ben-David et al. (2018) construct an ETF ownership measure by multiplying assets under management by the index weight of the security[1]. A similar approach is taken in Brown et al. (2018), Israeli et al. (2017), Agarwal et al. (2018), Saglam et al. (2019), and Brogaard et al. (2019). Currently ETF redemptions and creations data are available even at daily frequency. However, changes in ETF positions may significantly underestimate arbitrage intensity. If positive and negative price deviations are equally likely to occur, the net accumulated position of arbitrageurs over the day can be rather small despite intensive arbitrage[2]. Moreover, ETF arbitrage can also be performed by statistical arbitrageurs who close their positions at market prices and thus do not participate in the creation/redemption process. To the best of our knowledge, only Da and Shive (2018) take a different approach and use ETF turnover to measure arbitrage intensity, but only at the ETF level.

We propose a direct measure of arbitrage intensity. We propose to estimate the probability of simultaneous trading of ETF shares with shares of underlying stocks using high frequency data. We are motivated by the mechanics of ETF arbitrage. All arbitrageurs, irrespective of their type, respond to a price deviation by simultaneously trading ETF shares and the shares of underlying securities[3]. Simultaneity is crucial to avoid taking the unnecessary risks of adverse market movements. Our measure can be computed over any interval of the day and can account for statistical arbitrage. Thus, our approach delivers

[1] This approach might not work for a sizable proportion of funds that replicate their target index by investing in a representative basket of securities rather than the entire index basket. Brogaard et al. (2019) document that at least 22% of ETFs replicate their target index, including six of the ten largest ETFs as of 2018.

[2] Evans et al. (2019) also argue that the time period between the original mispricing/arbitrage event and the corresponding creation event may be as large as six days, as authorized participants have incentives to delay ETF creation.

[3] An offsetting operation can be performed later, once the price difference has been eliminated. Authorized participants can exchange the baskets for units of ETF shares in an in-kind transaction via a creation/redemption mechanism.

a more precise measure of the intensity of arbitrage transactions compared to existing approaches in the literature.

Another challenge is to document distortions in the underlying markets induced by ETFs. The most explored channel is the propagation of new liquidity shocks, brought to the market by ETFs' clientele and transmitted to underlying markets by ETF arbitrage. As a result, the underlying markets can be characterized by excess volatility as shown by Ben-David et al. (2018), and excess comovement (see Glosten et al. (2016), Israeli et al. (2017), Staer and Sottile (2018), Da and Shive (2018)); finally, Brown et al. (2018) show that ETF creations and redemptions predict returns for both the underlying securities and the ETFs themselves. The main difficulty in this line of research is to show that the volatility and/or comovement of the underlying returns are not driven by the market reaction to market-wide news. A few papers measure the impact of ETFs on price efficiency on the underlying markets by investigating the market response to earnings announcements. Israeli et al. (2017) document that an increase in ETF ownership is associated with lower return responsiveness to earnings and higher return synchronicity. In contrast, Glosten et al. (2016) find that an increase in ETF ownership brings quarterly stock returns closer to earnings realized over the same period.

We aim to document the distortion in the market reaction to fundamental shocks due to ETF arbitrage. ETFs facilitate price discovery, hence when new information arrives it can be incorporated into the ETF price first. As the ETF price deviates from the price of the underlying basket, it triggers arbitrage transactions. By purchasing or selling the underlying securities, arbitrageurs transmit the initial shock to the individual securities. This mechanical pricing pressure can be substantial if arbitrage is intense while the liquidity of the underlying market is low. As a result, the stock price response to the shock no longer reflects the true fundamental exposure.

We use event study methodology to identify exogenous fundamental oil shocks. Oil shocks have an important advantage over standard macroeconomic news shocks by containing large *idiosyncratic* component. An individual firm's exposure to oil shocks depends mainly on the nature of its business. Oil betas – betas from regressing a stock's return on the oil return – of most non-oil related companies are naturally expected to be *negative* (at least in the period before the shale boom). Intuitively, an increase in oil prices will raise input costs for most businesses, as well as force consumers to spend more money on gasoline and less on everything else. Although a systematic component may be present in the oil shocks, the idiosyncrasy is likely to dominate. We use oil inventory announcements to identify oil shocks. Oil inventory reports are released by the Department of Energy on a weekly basis at a pre-specified time and strongly affect the price of oil, thus providing us with a strong and frequent signal[4].

[4]See Anatolyev et al. (2018) for institutional details and for references to other papers that use oil inventory announcements to identify oil shocks.

We use high frequency data on stocks and oil futures to estimate stock market sensitivity to oil shocks. To measure arbitrage intensity, we sample data at a one-second frequency, and estimate the probability of simultaneous trading of ETF shares with shares of underlying stocks. Finally, we estimate a linear panel regression of stocks' sensitivity to oil shocks on the probability of simultaneous trading. In our regression analysis, we include stock-level fixed effects, sector by year fixed effects, various measures of the intensity of overall trading, measures of liquidity and price impact, volatility, and variables that capture time-varying firm value and size. Our sample of firms consists of all non-energy firms comprising the S&P 500 index and covers the period from 2010 to 2016. We consider the SPDR S&P 500 ETF (SPY) that tracks the S&P 500 index and represents the largest ETF in the world currently managing $280 billion in assets under management.

Our results confirm the main hypothesis that ETF arbitrage distorts the market reaction to shocks. The firms that are traded more often with SPY display a stronger reaction to oil shocks, regardless of their fundamental exposure. The effect is economically and statistically significant; an increase in the probability of simultaneous trading from the first to the third quantile is associated with an increase in the oil beta by 40% of the median oil beta.

One way to further test our mechanism is to show that the distortion is greater in exactly those firms where theory would predict that the effect of arbitrage transactions on prices is likely to be stronger. A natural characteristic of a stock to explore is, of course, liquidity. However, the presence of statistical arbitrage complicates the problem, as the same characteristic can have opposite effects on the intensity of arbitrage transactions and on the magnitude of distortion that the arbitrage transactions cause. Our direct measure of the intensity of arbitrage allows to separately measure these two effects. We show that our results are consistent with the arbitrage mechanism: the less liquid the stock is, the less likely it will be included in the arbitrage portfolio, but if included, the distortion will be larger.

Although the percent increase in oil betas due to ETF arbitrage activity is large, one can argue that the absolute increase in oil betas may not be large enough to be considered meaningful for the overall market and have serious implications. It should be noted that our empirical setup delivers the most *conservative* estimate of the effects of ETF arbitrage on the market reaction to shocks, large part of the effect is picked up by our controls, especially by the sector by year fixed effects. Moreover, there is nothing specific about oil shocks used in our analysis, we focus on oil inventory news for identification purposes only. While a median exposure to oil shocks is relatively low, a median exposure to other fundamental shocks can be much larger, and so can be the distortion.

There are two main challenges with our approach that open up our results to alternative interpretations. First, there is a clear omitted variable in our regressions which is the true oil beta. Second, our ability to clean our measure of arbitrage intensity from the intensity of overall trading may be considered insufficient. As a result, the positive correlation between the estimated oil betas and the probability of simultaneous trading that we document can potentially be driven by a positive relationship between the *true oil betas* and *the intensity of overall trading.* Although in our regression analysis we do include a large number of controls to proxy for the true oil beta and its variation over time, and we use a battery of different measures to control for the overall intensity of trading, one may not be fully convinced, especially considering that the positive link between the volume of trade and stock returns around public announcements can easily be generated by an asset pricing model with heterogeneous prior beliefs. Intuitively, the larger is the announcement return, the larger must be the revisions of agents beliefs, and thus the larger is the trading volume triggered by rebalancing activity. We carefully address this important alternative interpretation of our findings and run a number of tests to show that it is unlikely to drive our results. One particular advantage of our approach is our usage of oil shocks and oil betas (along with a clean identification of oil shocks) rather than macroeconomic shocks and market betas. Oil exposure of most non-energy companies is likely to be negative (or zero). The presence of negative oil betas allows us to directly test our mechanism against this alternative explanation. If there is a link between *true oil betas* and the intensity of overall trading it should be a *negative* relationship for a sizeable subset of the data, while our arbitrage mechanism predicts a positive relationship uniformly[5]. It should also be noted that we calculate the probability of simultaneous trading over the entire year, and not just over the short windows of time following oil inventory announcements.

Our findings are robust and can stand against against a wide range of reasonable alternatives, we list and carefully discuss the most plausible alternative interpretations for our results.

Our findings imply that ETFs affect the very basic role of asset prices, which is to reflect fundamental information. The distortion in the market reaction to oil shocks that we document represents a new source of an increase in the volatility of underlying assets, which is unrelated to liquidity shocks introduced by ETFs' clientele, and thus our results can be viewed as complimentary to Ben-David et al. (2018). We also show that the distortion is not short-lived. The long-lived distortions may seem to contradict our mechanism, as one would expect fundamental traders to (eventually) correct the mispricings. Although it is not unusual in the finance literature to document prolonged mispricings as arbitrage can be limited and arbitrage capital can be slow moving, recent theoretical research on ETFs offers yet another explanation. Bhattacharya and O'Hara (2018) study informational effects of ETFs in a factor model of prices. When ETFs are available for trading, informed agents optimally choose to herd on systematic part of

[5]This test would not be possible, if we were to use market betas instead of oil betas.

information. As a result, asset prices do not reflect idiosyncratic component and exhibit lower informational efficiency. This is consistent with our findings. Indeed, our results show that the non-energy companies in the S&P 500 index overreact to oil shocks. Now, if the stock market index correctly responds to oil shocks, then it mechanically implies underreaction of the energy firms in the index to oil shocks. As a result, idiosyncratic variation in the betas gets partially lost due to the presence of ETFs and ETF arbitrage, in line with Bhattacharya and O'Hara (2018). Distortion of the market reaction to fundamental news represents an important destabilizing mechanism of financial trading on asset price dynamics.

Our mechanism requires that ETFs play at least partial role in the price discovery process. That is, we need new information to be incorporated into the ETF price first, at least part of the time, to create a mispricing and trigger arbitrage activity. Is that a reasonable assumption? In a recent paper, Box et al. (2019) identify intraday ETF mispricings and investigate how trading in the ETF and the underlying stocks resolves these mispricings. According to their findings, mispricing is typically preceded by a permanent shock in the underlying portfolio and is subsequently corrected by quote adjustment and not by arbitrage transactions. Stale ETF pricing and absence of arbitrage transactions could raise a serious concern against plausibility of our mechanism. However, we show that SPY and the underlying portfolio react simultaneously to our oil inventory announcements even when we use one-second frequency data, while Box et al. (2019) work with a one-minute frequency. Moreover, we show that SPY overreacts to news relative to the underlying portfolio, and thus potentially opens up an arbitrage opportunity. Finally, the direction of these mispricings is such that arbitrageurs would push the prices of the underlying securities in the direction of the initial shock, consistent with our mechanism. We also get different results for the behavior of the directional volume. Thus, even though some smaller ETFs may indeed have negligible effects on the underlying markets, that is clearly not the case for SPY.

As a final robustness check, we perform an ETF-level analysis using non-energy U.S. equity ETFs. Our results indicate that more liquid ETFs, which are more likely to be involved in intensive arbitrage transactions, display a stronger reaction to oil shocks. Hence, our ETF-level analysis confirms our previous results that index trading intensifies the market reaction to fundamental oil shocks.

In a recent paper, Shim (2018) aims to document the distortion in the market reaction to factor information caused by ETF arbitrage. However, a few serious identification issues are present in the paper. Shim (2018) uses ETF betas – betas from regressing a constituent stock's return on the ETF's return using daily data (basically using market betas as ETF betas are likely to be highly correlated with market betas). Usage of ETF betas makes it challenging to identify fundamental news arrivals and separate them from liquidity shocks that could drive ETF returns. Therefore, the relationship between the intensity of arbitrage and ETF betas cannot be easily interpreted as the distortion in the market reaction to fundamental information[6]. Moreover, the ETF price reaction to news itself can be distorted, and thus, relative responses become even harder to interpret. Finally, Shim (2018) uses a measure of arbitrage sensitivity, which represents a combination of index weight and price impact sensitivity, and thus similar to other existing measures it does not account for statistical arbitrage. It is also impossible to separately measure the effect of liquidity on the intensity of arbitrage from its effect on the magnitude of distortion, while our approach allows us to do that.

Relatively little empirical evidence exists on whether fundamental shocks are amplified or distorted by arbitrage activity. One exception is Hong et al. (2012), who show that arbitrage activity amplifies the market reaction to earnings surprises, but they consider a very different economic mechanism: that stock prices overreact to good news due to short covering.

From a practical perspective, our results have important implications for risk management. We document distortions in the market reaction to oil shocks. As there is nothing specific about oil shocks in our analysis, we can expect that the stock market reaction to other fundamental shocks is also distorted by the presence of index arbitrage. Hence, standard measures of risk exposures have to be adjusted to remove the bias induced by the index arbitrage. Similarly, from an academic perspective, any event study analysis has to remove the bias in the market reaction to shocks induced by the index arbitrage before the results can be interpreted as reflecting true fundamental exposures. Hence, our findings caution against estimation of discontinuous betas recently put forward by Todorov and Bollerslev (2010), Bollerslev et al. (2013), Bollerslev et al. (2016). Intuitively, large stock market movements either triggered by arrivals of news or related to other sources of market discontinuity, should provide more information about changes in fundamentals and should be less subject to noise in the price formation process. However, we argue that such betas are likely to be significantly distorted by index arbitrage.

We also contribute to the literature by developing a direct measure of the intensity of arbitrage. Our measure is model free and it can be easily computed for an arbitrary set of assets linked by arbitrage. In contrast to other popular measures based on a cointegration assumption, our approach does not require long data series needed to estimate a long-term relationship between the assets. To provide additional empirical evidence, we perform several tests of our measure. Our approach is to identify times when arbitrage intensity is known to change and to test if our measure captures these changes. First, we analyze the S&P 500 index additions and deletions. We show that the probability of simultaneous trading increases when a firms is added to the index, and decreases when it is deleted,

[6]To better identify arrivals of fundamental information, Shim (2018) identifies days with large-in-magnitude ETF returns. However, large returns do not necessarily reflect news, as suggested by the Flash Crash of May 6, 2010 and similar events. We believe that our identification of fundamental shocks is more precise.

confirming our interpretation that the probability of simultaneous trading reflects changes in index trading intensity. Second, we analyze SPY creations and redemptions and show a positive link between changes in holdings and our measure. Finally, we use mutual fund fire-sales to identify exogenous variations in prices and document a positive link with our measure.

This paper is organized as follows. Section 2 illustrates the main mechanism and builds the foundation for our empirical framework. Section 3 outlines our empirical methodology, while section 4 describes our data. Our main results are presented in Section 5, where we also discuss other alternative explanations for our findings and discuss the magnitude of our effect. Subsequently, in Section 6 we provide empirical evidence that our measure captures the intensity of arbitrage transactions. We also carefully compare our measure of the intensity of arbitrage transactions with existing measures in the literature. We perform an additional exercise using ETF-level data in Section 7. Finally, in Section 8 we discuss our results and conclude.

2 Hypothesis Development

The main testable hypothesis of this paper is that ETF arbitrage distorts the market reaction to fundamental shocks. This section illustrates the mechanism and builds a foundation for our empirical framework. We start by briefly outlining the mechanics of ETF arbitrage.

Two types of ETF arbitrage

ETF arbitrage can be conducted in two ways: via redemption/creation of ETF shares by Authorized Participants or via statistical arbitrage. Both strategies assume the same initial response, but differ in the way the offsetting operation is performed.

Authorized Participants can profit from riskless arbitrage. Consider a case in which an ETF price exceeds the weighted price of the basket of underlying securities. An AP can simultaneously purchase the constituents of the index that the ETF holds in the exact same weights as the index, and sell the ETF shares. Once the market closes, an in-kind transaction occurs: The AP delivers the stocks to the sponsor of the fund and in return he receives a block of ETF shares, which covers his previous short position. This creation or redemption mechanism allows the AP to conduct riskless arbitrage. Typically, only large banks like Goldman Sachs and Bank of America can become APs. Moreover, only the elimination of sufficiently large price deviations should be profitable given the large transaction costs associated with a purchase of the entire basket of securities.

However, if an arbitrageur is eager to bear a certain amount of risk, arbitrage can take a different form. Statistical arbitrage also prescribes a simultaneous purchase(sale) of a subset of underlying securities and a sale(purchase) of ETF shares in response to a price deviation. In contrast to the previous strategy, however, an offsetting transaction is conducted as a typical market transaction once the price deviation has been eliminated. Although statistical arbitrage involves risk, it allows one to economize on transaction costs: An arbitrageur does not have to acquire the entire basket of underlying securities. A subset of the stocks can do the job as long as it tracks the index sufficiently well.

ETF arbitrage and the market reaction to shocks

How could ETF arbitrage distort the market reaction to shocks? We provide a simple illustration of the mechanism. Once introduced, ETFs facilitate price discovery. Imagine that new information arrives and becomes incorporated into the ETF price first[7]. When a shock arrives and the price of the ETF deviates from the price of the underlying basket, arbitrageurs start trading to eliminate the discrepancy. By purchasing or selling the underlying securities, arbitrageurs transmit the initial shock to the individual securities. This mechanical pricing pressure can be substantial if arbitrage is intensive while the liquidity of the underlying security is low. As a result, the stock price response to the shock no longer reflects the true fundamental exposure. The distortion, however, may be transitory if, over time, fundamental traders eliminate the induced mispricing.

Our goal in this paper is to create a measure of the intensity of arbitrage transactions and document a link between arbitrage transactions and distortions in the market reaction to fundamental news.

Stocks characteristics and the magnitude of distortion

Can we predict which stocks will experience the largest distortions associated with ETF arbitrage? In the presence of statistical arbitrage, we have to consider two channels.

First, a certain stock characteristic can influence arbitrage intensity. Not all stocks should necessarily be included into the arbitrage transaction. The portfolio of chosen firms must track the index well enough to avoid the risk of substantial price deviations at the time of offsetting the transaction. However, some firms can be omitted to economize on costs. Firms with larger index weight, volatility and liquidity are more likely to be included in arbitrage transactions. Indeed, when index weight is large, any variation in the stock price leads to a substantial variation in the index price. Similarly, the more volatile the stock price is, the more likely such a large deviation occurs. Finally, liquidity or the price impact function determines to what extent the price of the security will be affected by arbitrageurs' transactions, and thus affects the profitability of arbitrage.

[7] We would like to emphasize that we do not assume that *all information* is first incorporated in the ETF price. However, it is reasonable to assume that the ETF market plays at least a partial role in the price discovery process.

At the same time, for a given intensity of arbitrage transactions, the same stock characteristic can determine the magnitude of the market distortion. The less liquid the stock is, the larger the affect of arbitrage transactions on its price will be, and thus we can expect a larger distortion in the stock's response to shocks.

The total effect of some stock characteristics on the size of market distortion can be ambiguous. Indeed, the less liquid the stock is, the less likely it will be included in the portfolio. Hence, arbitrage will be less intensive, but the distortion caused by arbitrage transactions will be larger. Without a direct measure of arbitrage intensity, it is impossible to separate these two opposing effects of liquidity. However, because we directly measure the intensity of arbitrage transactions, we can separately measure the effect of a stock characteristic on the intensity of arbitrage transactions from its effect on the magnitude of distortion that these arbitrage transactions cause.

3 Methodology

In this section we outline our empirical strategy. We aim to estimate the effect of ETF arbitrage on the market response to fundamental shocks. Our approach involves three steps. First, we construct a measure of the intensity of ETF arbitrage transactions. Second, we use event study analysis to identify oil shocks and estimate the stock market reaction to these shocks. Finally, we use regression analysis to document the relationship between the intensity of arbitrage transactions and the market reaction to shocks. In what follows, we describe each step in detail.

3.1 Measuring the intensity of arbitrage transactions

Our measure is based on the following observation. Regardless of whether arbitrage is conducted by Authorized Participants or by statistical arbitrageurs, the initial response to a price deviation is the same: arbitrage prescribes the simultaneous trading of ETF shares and its constituents. Simultaneity is crucial to avoid taking the unnecessary risks of adverse market movements. Hence, we propose to use the estimated probability of simultaneous trading of ETF shares with shares of underlying stocks as a measure of the intensity of arbitrage transactions at the individual stock level.

Formally, consider an index and an ETF that tracks this index. For each firm j in the index, we use high-frequency data to estimate the probability of simultaneous trading of firm j and ETF shares over day t:

$$\pi_{j,t} = \frac{\sum_{\tau=1}^{T} \mathbb{I}_{V_{ETF,\tau}>0} \mathbb{I}_{V_{j,\tau}>0}}{T},$$

where $V_{j,\tau}$ is the volume of stock j traded over interval $\tau = 1, .., T$ of day t[8]. We propose to sample data on 1 second frequency.

3.2 Measuring the market response to shocks

In the second step, we identify fundamental shocks and measure stock market sensitivity to these shocks. We follow on our previous research, Anatolyev et al. (2018), and use oil inventory announcements to identify fundamental oil shocks.

Weekly estimates of crude oil inventories in the U.S. are provided by the U.S. Energy Information Administration (EIA), a statistical and analytical agency within the U.S. Department of Energy. A summary report is released in the form of an EIA publication, the *Weekly Petroleum Status Report*[9]. The report becomes available to the public at 10:30am Eastern time. The oil price typically reacts quite strongly to the announcements, thus providing us with a necessary signal.

We use an event study approach to estimate the stock market sensitivity to oil shocks. Our main identification assumption is the absence of any market-wide shocks except for the oil inventory news in the announcement window. To determine the size of the window, we face the standard tradeoff: we need the window to be narrow enough to justify our identification assumption, but we would like to avoid the complications of modeling microstructure noise. Hence, in our main exercise, we use a 1 minute window. Formally, we use the most recent transaction price as of 10:31:00 am (or one minute after the release) and the most recent transaction price 5 seconds before the announcement[10]. For robustness, we repeat our analysis with announcement returns calculated using a 5-minute window to account for potentially delayed information processing[11].

For each year and each firm j we estimate the following linear regression:

$$r_{j,t} = \alpha_j + \beta_j r_{oil,t} + \varepsilon_{j,t} \tag{1}$$

where t represents an announcement day, $r_{j,t}$ is the firm $j's$ announcement return on announcement day t, $r_{oil,t}$ is the first oil futures contract announcement return on announcement day t, and $E(\varepsilon_{j,t}|r_{j,t}) = 0$. We estimate oil betas separately for each firm over each year using OLS.

[8]For robustness we also use an estimate of conditional probability (see appendix A.4). By conditioning on firm j's trading, we can further filter our intensity of overall trading from our measure of the intensity of arbitrage transactions.

[9]For some weeks which include holidays, releases are delayed by one day.

[10]For each firm we exclude weeks during which no trading is recorded in the 5-minute interval before the announcement. Only a negligible number of weeks were excluded.

[11]We do not expect the lack of trading to be a problem in our dataset. Inventory announcements move oil prices significantly, and trading in the oil futures market intensifies around EIA announcements. Moreover, the S&P 500 index is comprised of actively traded firms, and our data typically shows multiple transactions recorded in the minute both before and after the announcement.

3.3 Regression analysis

To measure the effect of ETF arbitrage on the market reaction to shocks, we estimate the following fixed effect panel regression:

$$\beta_{j,t} = c_j + \alpha_{J,t} + \gamma_t \pi_{j,t} + \delta_t X_{j,t} + \varepsilon_{j,t} \tag{2}$$

where $\beta_{j,t}$ is the estimated oil beta of stock j in year t, $\pi_{j,t}$ is the estimated probability of simultaneous trading of stock j and SPY over the same period, and $X_{j,t}$ is a vector of control variables. Controls include variables describing the size and value of the firms such as the logarithms of market capitalization and book-to-market ratio. We also include the logarithms of the dollar volume and turnover to capture the overall intensity of trading of $j's$ stock. It is critical to control for the overall intensity of trading, as our measure cannot distinguish arbitrage transactions from coincidental trades, which are more likely for actively traded firms. Finally, we include the bid-ask spread and the price impact ratio calculated using high frequency data to capture overall liquidity and price impact of trades. The coefficient c_j captures the unobserved time-invariant stock fixed effect, which can be interpreted as stock $j's$ true sensitivity to oil shocks. We also include year by sector effects, $\alpha_{J,t}$. Time variability can be captured by allowing the coefficient of interest, γ_t, and the effects of controls δ_t to vary over time.[12]

4 Data and descriptive statistics

Before we proceed to our main results, we describe our data and characterize the behavior of the estimated intensity of arbitrage transactions and oil betas.

4.1 Data

In our empirical exercise we focus on SPY, an ETF that tracks the S&P 500 index. SPY is the largest ETF in the world, currently managing $280 billion in assets under management. We consider the firms comprising the S&P 500 index as of November 2018. Although the composition of the index has changed significantly over the last decade, we fix the set of firms to facilitate the estimation.

We use the closest to maturity WTI oil futures traded at NYMEX (CME group) to measure oil betas. We use a standard rolling procedure, replacing a soon-to-expire futures contract with the next one on the 5th day of the month. At about that time, the liquidity on the market moves from the first to the second contract markets.

The high-frequency data on stocks are obtained from the Trade and Quote Database

[12] On a technical note, we do not consider the estimated probability as a generated regressor. We treat the estimated probability itself as a measure of the intensity of arbitrage transactions.

Figure 1: Probability of simultaneous trading.

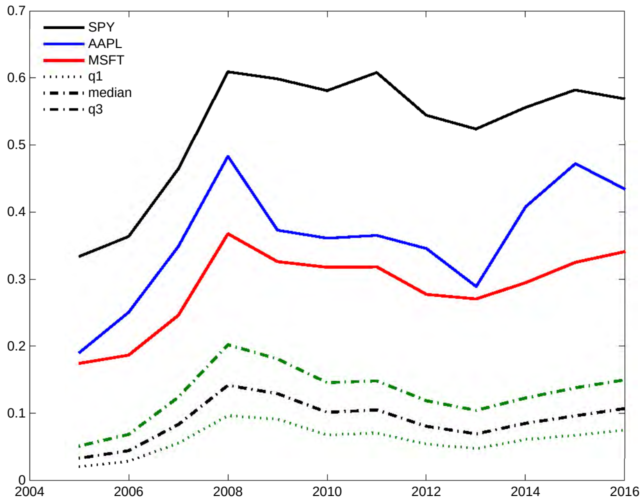

(TAQ) and on oil futures from TickData[13]. Our sample covers the period from 2005 to 2016. To calculate an estimate of the probability of simultaneous trading, we consider the time period from 9:30 am to 4:00 pm, and sample the data on 1 second frequency.

We use CRSP to calculate bid-ask spreads, and Compustat for market capitalization, turnover, and book-to-market ratio.

4.2 Descriptive statistics

In this subsection, we describe the time and cross-sectional behavior of our measure of the intensity of arbitrage transactions and oil betas.

Intensity of arbitrage transactions

Figure 1 displays the average probability of simultaneous trading with the SPY. The black line corresponds to SPY itself, and thus displays the probability of trading over a 1 second interval. The blue and red lines correspond to Apple and Microsoft, respectively, two firms with the largest weights in the index. Finally, the three lower lines correspond

[13] Prior to December 9, 2013 TAQ does not report odd lot trades (seeO'Hara et al. (2014) for the general discussion of this issue). The omission of the odd lots in the earlier subsample can lead to an underestimation of the probability of simultaneous trading and thus underestimation of the intensity of arbitrage. However, we believe that it is unlikely to bias our main results. To bias our results, the underestimation of the probability of simultaneous trading must correlate with the underestimation of oil betas. However, our sample of the firms consists of the most liquid firms. As we observe trading immediately following the announcements, the bias in the estimates of oil betas should be minimal (if any). Moreover, as we discuss below our results become even more pronounced as we increase the estimation window from 1 minutes to 5 minutes after the oil announcement. To provide further support, we repeat our main estimations including measures of algorithmic trading, in particular the odd lot volume ratio. We show that the results remain the same (see appendix A.11). It is also worth noticing, that we do not observe any significant changes in the probability of simultaneous trading around the date of the change in reporting requirements.

Figure 2: Evolution of oil betas.

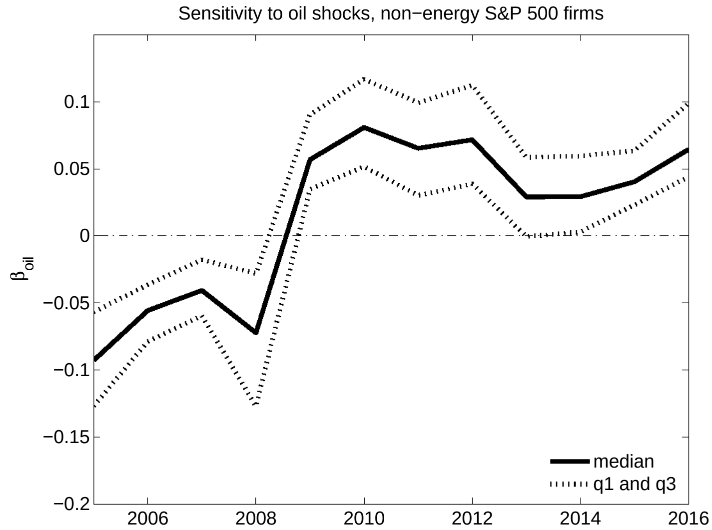

to the first, second and third quantiles of the distribution of probabilities across all firms in the S&P 500 index.

We see substantial time variation in trading intensity and in the estimated probability of simultaneous trading. Trading clearly intensified over the first 4 years of the sample; by 2008 SPY was traded about 60% of the time. The median probability of simultaneous trading also increased by 2008, but continued fluctuating afterwards.

More importantly, we observe significant cross-sectional variation in the estimated probability of simultaneous trading with SPY[14]. In particular, Apple and Microsoft, two firms with the highest index weight, clearly have much larger probability to be traded synchronously with SPY. This is a more general result, and partially reflects active trading of these stocks. We further discuss the determinants of the probability of simultaneous trading in section 6.

Oil betas

Figure 2 displays the evolution of oil betas for the sample of non-energy firms in the S&P 500 index over our sample period. Before 2009, all oil betas are negative[15], in line with historical evidence of a negative relationship between oil prices and the macroeconomy. However, the relationship suddenly becomes positive in 2009 and remains positive subsequently. The timing of the structural shift might not be surprising since in times of financial crises, fire sales and alleviated systematic risk, all assets typically move together. What is more surprising is that the positive link has not been reversed subsequently.

[14] The breakdown by industry is provided in appendix in Table 11. We do not see any particular patterns across industries, except for somewhat lower numbers in the real estate sector, and higher in the energy sector.

[15] See the breakdown of oil betas by industry and years in the appendix in Table 11.

It is beyond the scope of this paper to investigate the drivers of the shift, i.e. whether it was the development of unconventional oil, or geopolitical changes, or something else. However, in our regression analysis we assume time-invariant firm fixed effects. Hence, we need to restrict our sample to start in 2010 to avoid this dramatic shift in betas. Moreover, Figure 1 documents a period of rising overall intensity of trading until 2008. As we cannot fully wash out the intensity of overall trading from our measure of the intensity of arbitrage transactions, we have an additional reason to restrict the sample to exclude earlier years.

5 Main Results

5.1 Arbitrage intensity and market reaction to oil shocks

In this section, we document the relationship between sensitivities to oil shocks and our measure of index trading intensity. Table 1 reports our main results. Each column corresponds to a different specification, depending on whether we include controls and allow the regression coefficient to vary over time. For robustness, we repeat estimation for 1 minute and 5 minute returns.

Regarding the results for 1 minute returns, the probability of trading has a large and positive coefficient, and is robust to the inclusion of controls and time effects. The firms that are traded more often with SPY display a stronger reaction to oil shocks. To gauge the magnitude of the effect, consider an increase in the probability of simultaneous trading by 0.07, which is the difference between the third and the first quantiles. Such an increase is associated with an increase in oil beta by 0.02. In relative terms, the effect is large, as the median oil beta is 0.05 for most industries (see Figure 2 or Table 11). We further argue that our approach delivers a *conservative* estimate of the magnitude of the ETF distortion effect in section 5.5.

Importantly, inclusion of the measures of the overall intensity of trading, such as dollar volume and turnover, as well as the inclusion of market capitalization, does not affect our results. Hence, our measure of the intensity of trading contains additional information orthogonal to the intensity of overall trading, and we show that this information has the power to explain firms' response to oil shocks. In appendix A.4, we repeat the estimation using the conditional probability of simultaneous trading, which helps to filter our overall intensity of trading. We obtain even stronger results. In appendix A.11 we also show that inclusion additional controls for the intensity of algorithmic trading following Weller (2018) and Lee and Watts (2018) does not change our results[16].

[16] We follow the paper and use the SEC's MIDAS data to construct two measures of the intensity of algorithmic trading: the cancel-to-trade ratio, or the number of cancellations divided by the number of trades; and the trade-to-order volume ratio, or the total volume divided by the total volume across all orders placed. The data is available since 2012.

Table 1: Impact of probability of simultaneous trading with SPY on oil betas in the panel of S&P 500 firms.

The table displays the estimates of the following panel fixed effect regression: $\beta_{j,t} = c_j + \alpha_{J,t} + \gamma_t \pi_{j,t} + \delta_t X_{j,t} + \varepsilon_{j,t}$, where $\beta_{j,t}$ is the estimated oil beta of stock j in year t, $\pi_{j,t}$ is the average probability of simultaneous trading of stock j and SPY over the same period, $\alpha_{J,t}$ are year by sector dummies, and $X_{j,t}$ is a vector of control variables. Controls include logarithms of market capitalization, book-to-market ratio, dollar volume, turnover, volatility, and three measures of liquidity and price impact: bid-ask spread, the price impact coefficient estimated using high frequency data, and the Amihud ratio. St.err are clustered at the firm level, t-statistics in parenthesis. The sample period covers 2010-2016.

	1 min returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ProbTrad	0.49	0.47	0.28	0.31	0.28	0.29	0.20	
	(7.53)	(4.81)	(2.53)	(2.31)	(2.37)	(1.87)	(1.94)	
Prob -								
2010								0.29 (1.23)
2011								0.31 (1.86)
2012								0.27 (1.72)
2013								0.12 (0.61)
2014								0.25 (1.94)
2015								0.23 (2.06)
2016								0.36 (2.84)
Controls(mcap, book-to-market, turnover, dollar volume)		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bid-ask spread				Yes				
Price impact					Yes			
Amihud						Yes		
Volatility							Yes	
Year sector effects			Yes	Yes	Yes	Yes	Yes	Yes
Time varying γ_t, δ_t								Yes
R-sq within	0.037	0.096	0.254	0.266	0.254	0.254	0.269	0.292
R-sq between	0.003	0.207	0.022	0.020	0.022	0.018	0.099	0.000
N	3,058	2,877	2,877	2,876	2,877	2,877	2,877	2,877
N groups	461	444	444	444	444	444	444	444

Table 1(continue)

	5 min returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
ProbTrad	1.25	1.42	0.61	0.65	0.61	0.65	0.49	
	(10.15)	(8.06)	(3.43)	(3.08)	(3.26)	(2.73)	(3.00)	
Prob -								
2010								0.60 (1.59)
2011								0.56 (2.10)
2012								0.36 (1.34)
2013								0.55 (1.87)
2014								0.51 (2.41)
2015								0.51 (2.76)
2016								0.56 (3.33)
Controls(mcap, book-to-market, turnover, dollar volume)		Yes	Yes	Yes	Yes	Yes	Yes	Yes
Bid-ask spread				Yes				
Price impact					Yes			
Amihud						Yes		
Volatility							Yes	
Year sector effects			Yes	Yes	Yes	Yes	Yes	Yes
Time varying γ_t, δ_t								Yes
R-sq within	0.069	0.268	0.489	0.490	0.489	0.490	0.497	0.514
R-sq between	0.015	0.139	0.072	0.068	0.071	0.052	0.140	0.001
N	3,058	2,877	2,877	2,876	2,877	2,877	2,877	2,877
N groups	461	444	444	444	444	444	444	444

The effect of probability on oil betas also survives the inclusion of various measure of liquidity and price impact: the point estimate remains unchanged and significant. The effect of probability weakens when we include volatility, but remains significant at the 10% level. One reason could be that ETF arbitrage increases stock volatility itself, as shown by Ben-David et al. (2018), as new noise and liquidity shocks arise in the ETF market and propagate to the underlying markets via ETF arbitrage. We confirm this basic finding in our data using our measure of arbitrage intensity in appendix A.6.

The results are similar when we allow all coefficients to vary over time, see specification IV. With the exception of 2013, the estimates are extremely close to the value of 0.31 in the time-invariant specifications. The coefficients are significant in the 2014-16 period, although lacking significance in the earlier years. The lack of significance can be due to the lack of a strong effect of oil inventory announcements on oil prices, especially over the 2012-13 period as we document in Anatolyev et al. (2018). When the oil price responses to news are small, the signal-to-noise ratio decreases and our estimates of oil betas are likely to become less precise.

We also see that the results become even stronger when announcement returns are measured over 5 minutes after the release.

Overall, our results indicate a significant link between probabilities and oil betas confirming that ETF arbitrage distorts the market reaction to oil shocks.

5.2 Differential effects

One way to provide further support for our mechanism is to show that the distortion is greater in exactly those firms where theory would predict that the effect of arbitrage transactions on prices is likely to be stronger. A natural characteristic of a stock to explore is, of course, liquidity. However, the presence of statistical arbitrage complicates the problem, as the same characteristic can have opposite effects on the intensity of arbitrage transactions and on the magnitude of distortion that these arbitrage transactions cause. Our direct measure of the intensity of arbitrage allows to separately measure these two effects. In this section we aim to test if for the same intensity of arbitrage transactions, a stock characterized by a higher price impact displays a stronger reaction to shocks.

We use three different measures of liquidity and price impact. First, we use the bid-ask spread calculated as the quoted spread divided by the price from CRSP. Second, we follow Amihud (2002) and measure price impact by computing the absolute daily return divided by the total dollar daily volume. Finally, we calculate price impact using high frequency data[17].

For brevity, we only consider one specification that assumes the time invariant effects of controls and probabilities, but contains time by sector fixed effects (specifications (4-5)

[17] See appendix A.1 for details.

Table 2: Estimates of the differential effect of probability on oil betas for high and low liquidity stocks

The table displays the estimates of the following panel fixed effect regression: $\beta_{j,t} = c_j + \alpha_{J,t} + (\gamma_0 + \gamma_1 Z_{j,t})\pi_{j,t} + \delta_0 Z_{j,t} + \delta' X_{j,t} + \varepsilon_{j,t}$, where $\beta_{j,t}$ is the estimated oil beta of stock j in year t, $\pi_{j,t}$ is the average probability of simultaneous trading of stock j and SPY over the same period, $\alpha_{J,t}$ are year by sector dummies, and $X_{j,t}$ is a vector of control variables. Controls include logarithms of market capitalization, book-to-market ratio, dollar volume, and turnover. Variables $Z_{j,t}$ represent three measures of liquidity and price impact: bid-ask spread, price impact calculated using intraday data, and Amihud ratio calculated using monthly returns and volumes. St.err are clustered at the firm level, t-statistics in parenthesis.

Panel A: The entire period 2010-2016

	1 min returns			5 min returns		
	(1)	(2)	(3)	(1)	(2)	(3)
ProbTrad	2.12	4.89	7.98	2.77	6.01	12.84
	(1.24)	(1.36)	(1.81)	(1.04)	(1.10)	(2.04)
ProbTrad × Bid ask	0.22			0.26		
	(1.13)			(0.85)		
ProbTrad × Price impact		0.33			0.39	
		(1.31)			(1.02)	
ProbTrad × Amihud			0.31			0.49
			(1.79)			(1.99)

Panel B: The period from 2014-2016.

	1 min returns			5 min returns		
	(1)	(2)	(3)	(1)	(2)	(3)
ProbTrad	3.82	10.85	8.83	4.31	11.63	9.56
	(2.45)	(4.19)	(3.42)	(2.70)	(4.05)	(3.24)
ProbTrad × Bid ask	0.41			0.41		
	(2.37)			(2.33)		
ProbTrad × Price impact		0.74			0.76	
		(4.11)			(3.82)	
ProbTrad × Amihud			0.34			0.36
			(3.39)			(3.05)

in Table 1). We consider each of our three measures of liquidity and price impact separately, and we interact each of them with the probability of simultaneous trading. Panel A in Table 2 shows the estimates for the entire sample period. All the interaction terms are positive, implying that conditional on the level of intensity of arbitrage transactions, firms that are characterized by lower liquidity and a larger price impact, display a larger reaction to oil shocks, which is consistent with our mechanism. However, we only find significance when we use the Amihud measure. Panel B shows the estimates for the 2014-2016 period when oil price fluctuations provide a stronger signal. Now we see that all the coefficients for the interaction of probability and measures of liquidity and price impact are positive and significant, no matter which measure we use and how we measure announcement returns.

In section6 we also show that less liquid stocks are associated with less intense arbitrage. Hence, our results are consistent with the arbitrage mechanism: the less liquid the stock is, the less likely it will be included in the arbitrage portfolio, but if included, the distortion will be larger.

5.3 Long-term effects

One concern may be that our results document only a temporary distortion of the market reaction to shocks by ETF arbitrage. Eventually, fundamental traders may correct the mispricing and thus, the detrimental effect of ETFs is short-lived.

Even if that is the case, our results imply that ETFs introduce yet another source of noise to the market. There is nothing specific about oil shocks in our analysis; our choice of oil shocks was driven by our ability to identify these fundamental shocks using oil inventory announcements. Naturally, we may expect that the market reaction to all fundamental shocks becomes distorted by ETF arbitrage. Hence, ETF arbitrage introduces a continuous cycle of mispricings and corrections and leads to excessive fluctuations of asset prices. We identify another source of the increased volatility of underlying securities due to the presence of ETFs. Our paper can be viewed as complimentary to the findings of Ben-David et al. (2018).

However, it is likely that initial distortions never become fully corrected. In appendix A.7, we repeat our analysis using oil betas calculated over the entire day. We calculate the cumulative return on each stock from the time of announcement until the end of the day. The oil beta is defined as the regression coefficient of these cumulative returns on 1 minute oil announcement returns (to keep clean identification of oil shocks). We show that our effect is still present and significant. Even by the end of the day, the mispricing associated with the intensity of ETF arbitrage is not corrected.

We find that ETF distortions are not short-lived. Hence, ETFs not only introduce a new source of noise to the market, but actually affect the very basic role of asset prices to reflect fundamental information. Hence, ETFs have a detrimental impact on price efficiency.

5.4 Alternative explanations

Our preferred interpretation of the results in Table 1 is the distortion of the market reaction to oil shocks induced by ETF arbitrage. In this section, we discuss other potential explanations for our findings. We also discuss our assumptions and our approach in general.

There are two main challenges with our approach that open up our results to alternative interpretations. First, there is a clear omitted variable in our regressions which is the true oil beta. Second, our ability to clean our measure of arbitrage intensity from the intensity of overall trading can be limited. Hence, the positive correlation between the estimated oil betas and the probability of simultaneous trading that we document can potentially be driven by a relationship between the *true oil betas* and *the intensity of overall trading.* Below we discuss one potentially important driver of this relationship, however, first we would like to describe our general efforts to solve this identification issue. First, in our regression analysis we include a number of controls to proxy for the true oil beta and its variation over time. In particular, we include stock fixed effects and sector by year fixed effects, in addition to other standard controls. It should also be noted that we carefully control for the intensity of overall trading in our regressions. We use the dollar volume and turnover as our preferred measures, and in addition, in appendix A.4 we repeat estimation using the conditional probability of simultaneous trading, which further filters out the intensity of trading of the individual stock. We obtain even stronger results. To provide even more empirical evidence that our measure indeed measures the intensity of arbitrage transactions and not simply reflect the intensity of overall trading, we devote Section 6 to various empirical tests of our measure. Finally, we would like to emphasize that we use *oil shocks and oil betas*, and not macroeconomic shocks and market betas. Below we argue how this alleviates the identification problem, because true oil betas are more likely to be orthogonal to both the intensity of arbitrage transactions and the intensity of overall trading, than market betas.

We start by addressing the most plausible alternative interpretation of our results, which is based on the relationship between true oil betas and the intensity of overall trading. The reverse causality argument states that the stocks with larger true oil betas are traded more intensely when oil news arrive. This prediction can be generated by an asset pricing model with heterogeneous prior beliefs. The larger is the announcement return, the larger must be the revisions of agents beliefs, and thus the larger is the trading volume that we observe as the agents rebalance their portfolios. Although plausible, this alternative reverse causality argument is unlikely to drive our results for three reasons. First, we use *oil shocks and oil betas*, and not macroeconomic shocks and market betas. True oil exposure of most non-energy companies is likely to be negative (or zero). Intuitively, an increase in oil prices will raise input costs for most business, as well as force consumers to spend more money on gasoline and less on everything else. The idiosyncrasy in firms' sensitivities to oil shocks dominates any systematic component of oil price fluctuations. Thus, even if there is a link between *true oil betas* and the intensity of overall trading, it actually should be a *negative* relationship for a sizable subset of the data. The presence of negative oil betas allows us to test our mechanism against the alternative explanation. The test can be performed by repeating our main regressions using the *absolute value* of the beta. If the alternative mechanism drives our results, then taking the absolute value of the beta should make the results even stronger, as stocks with the greater exposure (i.e. more negative betas, larger in absolute value) are associated with more intensive trading. In contrast, we argue that more intensive arbitrage increases oil betas uniformly, in particular, it makes negative oil betas *less* negative, hence the results should become weaker. As shown in the appendix A.10, the results do become weaker when we use the absolute value of the oil beta thus supporting our story. Note that, if we were to use market betas instead such a test would be impossible and we would have a more serious identification issue, as the true market betas are more likely to be related to the intensity of overall trading [18,19].

Second, we calculate the probability of simultaneous trading over the entire trading period using all trading days, not just over a short window immediately following an oil inventory announcement. Oil news constitute only a negligible fraction of overall news and thus can hardly drive the average intensity of trading, and thus our results.

Finally, it should be noted, that the reverse causality mechanism inadvertently makes a very restrictive assumption on the cross-sectional distribution of the true oil betas. The companies that are traded intensely and have large estimated oil betas, are presumed to have large true oil betas. In the data, we find that firms that are traded more intensely, are typically the firms with the largest weight in the index. However, it is not clear why such companies, including, say, Amazon and Facebook, should necessarily have larger positive true exposures to oil shocks relative to the firms with lower index weight. In contrast, our story does not impose any restrictions on the distribution of true oil betas. It is the arbitrage activity that creates a link between the intensity of overall trading and

[18] Although there is little evidence of a strong relationship even between true market betas and the intensity of overall trading. To illustrate it, we use two measures of the intensity of trading, the dollar volume and the probability of trading (estimated as a fraction of one second intervals with non-zero trading volumes), to calculate the cross-sectional correlations with estimated market betas (by year). In 7 out of 12 years in our sample, the correlations of market betas with either of the two measures of the intensity of trading were lower than 0.15, although positive, and only slightly higher in other years.

[19] We also calculate the cross-sectional correlation of oil and market betas for each year and find that this correlation was typically very low (see appendix A.5 for details and additional discussion, including the discussion of the shale boom and its potential effects on the systematic component of oil price fluctuations).

observed oil betas. Amazon and Facebook are more likely to be involved in statistical arbitrage, and thus are more exposed to arbitrage distortions and have more distorted oil betas. The true oil betas can be uniformly zero or even negative. Thus, our mechanism generates a link between observed betas and trading volumes under much less restrictive assumptions. Hence, it is unlikely that our results are driven by the reverse causality argument.

Our results can also be partially driven by infrequently traded firms. Such firms have low probability of simultaneous trading with SPY due to low trading activity. Moreover, these same firms could have zero or close to zero estimated oil betas, as the lack of trading in a short window around the announcement prevents new information from being incorporated into the price. However, we believe that our results are not driven by infrequently traded firms. First, we consider the S&P 500 firms, which are large and actively traded companies. Second, we observe no trading in the event window only for a negligibly small number of firm-event pairs. More importantly, as Table 1 reports, our results become stronger for 5 minute returns, which goes against this alternative explanation. As we increase the event window from 1 minute to 5 minutes, at least some trading is more likely to occur and thus oil betas become more precise and no longer zero. Hence, the effect should have weakened if the alternative explanation was true.

The effect of arbitrage on stock's oil betas that we interpret as a distortion, may alternatively be interpreted as an enhancement, a positive effect of ETFs on price efficiency. According to this view, by attracting new clientele and thus improving liquidity of the underlying stocks, the ETFs make trading based on fundamental information more attractive and thus improve stock responsiveness to news. Although this could be a plausible alternative for small and illiquid stocks, it is hardly relevant for the most liquid and frequently traded S&P 500 stocks that we consider. Indeed, this idea is further studied by Glosten et al. (forthcoming) who investigate the effect of ETF ownership on stock responsiveness to earnings announcements and only find a positive and significant effect in small stocks and stocks with low analyst coverage. Moreover, this alternative view also suggests that the effect should be stronger when the absolute values of betas are used, however that is not the case.

Clearly, we cannot rule out all possible alternative explanations for a positive relationship between oil betas and the probability of simultaneous trading. However, if one takes into account stock fixed effects and year by industry effects, as well as a large set of controls, our findings can stand against a wide range of reasonable alternatives. It should also be noted that any alternative explanation should also account for the differential effects of liquidity, that we document.

One might also question our approach to the estimation of oil betas. We use oil inventory surprises during each year to estimate stocks' sensitivity to oil shocks. In doing so, we do not distinguish the sources of shocks that drive inventory changes. However, Kilian and Park (2009) argue that the stock market reaction to oil shocks differs depending on whether the change in the price of oil is driven by demand or supply shocks in the oil market. This is not necessarily an issue for us. As long as the distribution of shocks is relatively over time, we do not have a problem at all. However, even if the composition of shocks changes over time, under a mild monotonicity assumption our results remain valid. Indeed, in our main regression specification we include individual stock fixed effects and sector by year effects. As long as individual firms that belong to one sector experience a uniform shift in betas over the years due to changes in the composition of shocks, the resulting changes in oil betas will be picked up by sector-year effects and thus would not affect our main results. The only situation in which our results would be affected is if we were to have individual changes in oil betas positively correlated with changes in the probability of simultaneous trading with SPY due to changes in the composition of shocks. We find this combination of effects extremely unlikely, and thus we believe our results are robust to changes in the shocks composition.

Another concern has been raised recently by Box et al. (2019) who identify intraday ETF mispricings and investigate how trading in the ETF and the underlying stocks resolves these mispricings. According to their findings, mispricing is typically preceded by a permanent shock in the underlying portfolio and is subsequently corrected by quote adjustment and not by arbitrage transactions. Although we also consider an arrival of fundamental information as a source of mispricing, stale ETF pricing and absence of arbitrage transactions could raise a serious argument against our mechanism. In appendix A.9, we follow their approach and investigate the behavior of quotes and volumes around our oil inventory announcements. In contrast to Box et al. (2019), we show that SPY and the underlying portfolio react simultaneously to our oil inventory announcements even when we use 1 second frequency data, while Box et al. (2019) work with one minute frequency. Moreover, we show that SPY reacts stronger to news and thus potentially opens up an arbitrage opportunity in the direction that is consistent with our mechanism. We also get different results for the behavior of the directional volume. Thus, even though some smaller ETFs may indeed have negligible effects on the underlying markets, that is clearly not the case for SPY[20].

5.5 Magnitude of our effect and discussion

We also would like to comment on the magnitude of our main results. Although the percent increase in oil betas is non-trivial (40% in our benchmark case), one can argue that going from a very small exposure, 0.05, to a slightly larger exposure, 0.07, does not lend confidence that this effect is meaningful for the overall market and has serious

[20]In a different setting and for a different time period, Foucault et al. (2016) identify arbitrage opportunities due to asynchronous price adjustments. They find out that these opportunities terminate with an arbitrageurs' trade in about two-thirds of the case.

implications if the change is driven by distortions. First of all, while a median exposure to oil shocks is relatively low, a median exposure to other fundamental shocks can be much larger, and so can be the distortion. There is nothing specific to oil shocks in our analysis, our choice of oil shocks is driven solely by our desire for a clean identification.

We also would like to emphasize that our empirical setup delivers the most *conservative* estimate of the effects of ETF arbitrage on the market reaction to shocks. Our effect is partially picked up by numerous controls that we use in our regression analysis. For example, fluctuations of the overall intensity of trading are partially driven by changes in the intensity of arbitrage transactions. When we remove the controls, the effect of ETF arbitrage on betas becomes two and a half times larger.

Finally, we maintain the assumption that any changes in average industry sensitivities to oil shocks are not related to ETF arbitrage and potentially reflect fundamental changes in the U.S. market (perhaps due to the development of unconventional oil). This changes are removed from the analysis by sector by year indicators. However, instead, these changes, including the dramatic change in oil betas from uniformly negative to uniformly positive in 2008, may be yet another consequence of the presence of ETFs.

Interestingly, the relative dynamics of the average industry oil betas provides yet another argument in favor of our mechanism. Table 11 shows that all non-energy sectors experienced a dramatic increase in oil betas from uniformly negative to uniformly positive. Surprisingly, the average beta of the energy sector, actually decreased from 0.42-0.45 in the four years before 2008 to just 0.31 in the period of 2009-2013. Thus, while non-oil sectors became positively related to oil shocks, the sensitivity of the oil sector itself to oil shocks decreased. The unification of betas, an increase in the non-energy oil betas and a decrease in the energy oil betas, although puzzling from a standard macro perspective, is consistent with our arbitrage mechanism. Imagine, that oil news are processed at the ETF market first[21]. Once the market responds, an arbitrage opportunity opens up, and ETF arbitrage spreads out the shock to the individual markets almost uniformly. As long as two stocks have similar weights in the index and have similar liquidity, we will observe the same price pressure due to arbitrage transactions. The nature of the firm becomes irrelevant, it does not matter whether the firm belongs to the energy sector or not. Hence, ETF arbitrage has a tendency to unify individual market responses to news. Of course, in reality fundamental traders at least partially correct the mispricings, potentially eliminating or reducing the price response of some non-energy companies and increasing the price response of energy companies. However, if this fundamental trading is limited or slow, the unification of betas can be observed over a sufficient period of time to be detected[22], in line with our findings.

[21]The market response consists of the direct effect of the revaluaton of the energy firms included in the index, but also of the general equilibrium effects of higher or lower oil prices relevant for all firms.

[22]We conduct our main analysis using only stocks of the non-energy firms for two reasons. First, arbitrage pressure is more likely to dominate fundamental trading in individual non energy securities as

One can argue that long-lasting permanent effects are less consistent with market distortions and more consistent with permanent changes in prices. However, a number of recent theoretical papers point out several limitations that fundamental traders face to identify and correct the mispricings when ETFs are present on the market. Some papers suggest existence of equilibria with permanently incorrect prices. Bhattacharya and O'Hara (2018) study the informational effects of ETFs in a factor model of prices. A newly introduced ETF attracts new clientele with additional information about the fundamental values of the underlying securities, both the systemic factor and individual shocks. The market makers on the underlying markets can now extract information from the ETF price. However, if informed traders use short-term trading strategies, the introduction of an ETF facilitates a herding equilibrium (in the spirit of Froot et al. (1992)). In a herding equilibrium, informed traders rationally choose to trade based on information on the systematic factor only and completely disregard all idiosyncratic information. Each informed trader on the market for each individual security, can accurately anticipate the information flow into their own market once market makers learn from observing other markets and the ETF price, and thus can foresee the price bump and profit from it. As a result, in equilibrium asset prices exhibit lower informational efficiency. In other words, the *market reaction to news is distorted*, as all individual information is disregarded and never becomes reflected in prices. It should also be noted that endogenous information acquisition and processing further reinforce the herding. Over time, any incentives to acquire and process idiosyncratic information cease if informed traders use short-term strategies and trade based on a systematic piece of information only. Hence, our results can be seen as an empirical evidence of such herding.

Other recent papers on algorithmic trading focus on how fast arbitrage can alter incentives for fundamental traders to acquire information. This literature is particularly relevant for us, as liquid and cheap ETFs are perfectly suited to attract high frequency traders. Weller (2018) documents that algorithmic trading deter information acquisition. Intuitively, improved screening of informed order flow erodes information rents and reduces profitability of fundamental trading. Similarly, Lee and Watts (2018) exploit SEC's randomized tick size experiment to provide more direct evidence of the causal effect of tick size (and thus intensity of algorithmic trading) on fundamental information acquisition. Among other results, they find that with a larger tick size treatment firms experience a significant increase in EDGAR web traffic in the days leading up to each earnings announcement. Foucault et al. (2016) argue that arbitrageurs are often able to detect arbitrage opportunities due to asynchronous adjustments in asset prices following information arrival. By arbitraging away such a price discrepancy, arbitrageur cannibalize on stale quotes and expose dealers and other market participants to the risk

oil news have only marginal importance for these firms, relative to energy firms. Second, we do not have a sufficiently large number of energy stocks in the S&P 500 index to conduct meaningful inference.

of trading against more informed counter party. Foucault et al. (2016) document an increase in bid-ask spreads in the affected markets as the dealers seek to insure themselves against this risk. However, some market participants may also respond by pulling out from the trading altogether, or at least at times of known public announcements, further deteriorating fundamental trading in the individual markets.

Overall, the presence of ETFs can negatively affect expected profitability of fundamental trading and thus allow for mispricings to persist in the markets for individual stocks over a long period of time consistent with our findings.

6 Empirical tests of our measure and comparison with other measures

In this section, we investigate our measure and provide empirical evidence that the estimated probability of simultaneous trading with SPY measures the intensity of arbitrage transactions. Our approach is to identify times when arbitrage intensity is known to increase, and analyze whether our measure captures this intensification. We analyze i) algorithmic market making activity; ii) additions and deletions to the S&P 500 index; iii) days of intensive creations and redemptions of SPY units; and iv) times of mutual fund fire sales. Finally, we compare our measure to existing measures in the literature.

Determinants of the probability of simultaneous trading

We start by analyzing the determinants of our measure. We transform our estimate of probability using the logistic function and estimate the following fixed effects regression:

$$y_{j,t} = c_j + \alpha_{J,t} + \delta_t X_{i,t} + \varepsilon_{i,t},$$

where $y_{j,t} = \ln\frac{\pi_{j,t}}{1-\pi_{j,t}}$ is the logarithm of the odds, $\pi_{j,t}$ is the estimated probability of simultaneous trading of stock j and SPY in year t, and $X_{j,t}$ is a vector of our explanatory variables, including measures of liquidity and price impact and also dollar volume, turnover, market capitalization and book-to-market ratio. Time-invariant stock effects are captured by c_j. We also include year by sector fixed effects. Time variability is captured by allowing the coefficients δ_t to vary over time.

Panel A in Table 3 shows the results for the time invariant specification. Not surprisingly, we find a strong relationship between the probability of simultaneous trading and dollar volume (and turnover). As SPY is traded frequently, when an individual stock is traded more often, a simultaneous transaction becomes more likely to occur. As we cannot fully filter out these coincidental trades, it is critical to account for the overall intensity of trading in our regressions.

Panel A in Table 3 also shows that liquidity measures have an insignificant effect on the probability. To further explore this relationship, Panel B allows for the effects of liquidity measures (as well as other controls) to vary over time. We find that, in the early years, all coefficients are negative and mostly significant. The high frequency measure of price impact also indicates a negative and significant relationship in 2015. The negative link between price impact and the intensity of arbitrage transactions partially reflects the reluctance of arbitrageurs to include an illiquid stock in the arbitrage portfolio. By conducting statistical arbitrage, the arbitrageurs can avoid purchasing the entire basket of individual stocks to economize on transaction and market impact costs. A sufficiently large portfolio would suffice, as long as it tracks the index well enough. Thus, our results (weakly) confirm the presence of statistical arbitrage.

Test 0: Comparison to other measures of algorithmic trading

We start by exploiting the tight relationship between different types of algorithmic trading activities. This should be viewed as suggestive evidence, rather than a precise test.

The statistical arbitrage is just one of many applications of algorithmic trading, along with optimal order execution and market making. While representing separate activities, arbitrage and market making are not independent and often compete for the same order flow. Indeed, a transitory shock triggered by a large demand order can be traded away be arbitrageurs or can be absorbed by limit orders supplied by liquidity providers. Thus, active market making reduces arbitrage profits, eventually ceasing arbitrage trading activity. In line with this argument, Chaboud et al. (2014) and Brogaard et al. (2014) find that HFTs[23] tend to consume liquidity when eliminating temporary pricing errors; at the same time the frequency of arbitrage opportunities decreases as algotrading intensifies. Thus, we can use currently developed measures of algorithmic trading activity that pick up variation in market making, to show the ability our measure to pick up variation in the intensity of statistical arbitrage.

Earlier papers on the effects of algorithmic trading were confined to use regulatory and proprietary data. Recently, however, huge regulatory interest led to a creation of the United States Securities and Exchange Commissions's Market Information Data Analytics (MIDAS) system which makes some of the stock-level statistics publicly available. We follow the SEC's suggestions and Weller (2018) and construct two types of daily measures of the intensity of algotrading. First two measures calculate the relative number of order cancellations. We use the cancel-to-trade ratio, or the number of cancellations divided by the number of trades and the trade-to-order volume ratio, or the total volume divided by the total volume across all orders placed. Next two measures calculate the amount of small trades. We use the total volume executed in quantities smaller than 100 shares divided

[23]It is worth noting, that many papers specifically focus on the speed of transactions aim to identify high frequency trading, but do not try to separate market making and arbitrage activities.

Table 3: Determinants of the probability of simultaneous trading with SPY.

The table displays the estimates of the following panel regression: $y_{j,t} = c_j + \alpha_{J,t} + \delta X_{j,t} + \varepsilon_{j,t}$, where $y_{j,t} = \ln \frac{\pi_{j,t}}{1-\pi_{j,t}}$ is the logarithm of the odds, $\pi_{j,t}$ is the average probability of simultaneous trading of stock j and SPY in year t, and $X_{j,t}$ is a vector of explanatory variables: logarithms of dollar volume, turnover, market capitalization and book-to-market ratio. St.err are clustered at the firm level, t-statistics in parenthesis.

Panel A: Time invariant specification

	(1)	(2)	(3)
Bid-Ask Spread	0.021		
	(1.09)		
Price Impact		-0.014	
		(-0.58)	
Amihud Ratio			0.012
			(0.25)
Dollar Volume	0.418	0.405	0.425
	(14.41)	(13.75)	(7.96)
Turnover	0.119	0.128	0.121
	(3.03)	(3.23)	(2.60)
Mcap	-0.061	-0.063	-0.062
	(-2.33)	(-2.41)	(-2.43)
Book to market ratio	0.023	0.025	0.024
	(1.97)	(2.17)	(2.08)
Year sector effects	Yes	Yes	Yes
Time varying γ_t, δ_t	No	No	No
R-sq within	0.81	0.81	0.81
R-sq between	0.76	0.76	0.75
N obs	2,876	2,877	2,877
N groups	444	444	444

Panel B: Time varying effects of liquidity on probability.

	(1)		(2)		(3)
Bid-Ask Spread -		Price Impact-		Amihud Ratio-	
2010	-0.057 (-1.59)		-0.087 (-2.14)		-0.166 (-2.92)
2011	-0.012 (-0.56)		-0.073 (-2.60)		-0.059 (-1.54)
2012	0.032 (1.66)		0.053 (2.17)		0.052 (1.98)
2013	0.031 (1.48)		0.039 (1.37)		0.030 (0.9)
2014	0.011 (0.57)		0.025 (0.96)		0.061 (1.85)
2015	-0.015 (-0.86)		-0.063 (-2.03)		0.066 (1.47)
2016	0.004 (0.22)		-0.053 (-1.65)		0.161 (3.67)
Controls	Yes		Yes		Yes
Year sector effects	Yes		Yes		Yes
Time varying γ_t, δ_t	Yes		Yes		Yes

Table 4: Correlations of our measure of the intensity of arbitrage with the four measures of the intensity of algorithmic trading from Weller (2018).
Average (over time) cross-sectional correlation of 4 measures of AT with our measure of the intensity of arbitrage

	Trade-to-Order Volume Ratio	Cancel-to-Trade Ratio	Odd Lot Volume Ratio	Average Trade Size
Probability of Simultaneous Trading	0.49	-0.47	-0.54	0.49

Median (across all S&P 500 firms) time-series correlation of 4 measures of AT with our measure of the intensity of arbitrage for each year separately

	Trade-to-Order Volume Ratio	Cancel-to-Trade Ratio	Odd Lot Volume Ratio	Average Trade Size
Probability of Simultaneous Trading-2012	0.44	-0.40	-0.39	0.21
2013	0.45	-0.38	-0.37	0.22
2014	0.26	-0.18	-0.32	0.13
2015	0.18	-0.17	-0.22	0.10
2016	0.22	-0.23	-0.37	0.25

by total volume traded and the average trade size. Higher number of cancellations and smaller average trade size indicate more intensive algorithmic trading. It can immediately be seen that first two measures tend to capture market making activity, while the last two measures can also pick up variation in order execution activities (see also Hendershott et al. (2011)). Weller (2018) finds that all four measures are highly correlated within each group (-0.83 and -0.94), but less than perfectly correlated between groups (in absolute value the correlations range from 0.34 to 0.56), which confirms different informational content of the measures.

We use all four measures to investigate the relationship between arbitrage intensity and algorithmic trading. Table 4 displays various time-series and cross-sectional correlations. The results show strong and stable correlations both in the cross-section and across time within firms (we only report the median), moreover, the correlations with our measure are comparable in magnitude with the correlations reported in Weller (2018). We observe a negative correlation of our measure with the cancel-to-trade ratio and the odd lot volume ratio and a positive correlation with the trade-to-order ratio and the average trade size. Hence, our results consistently indicate that the intensification of activity that is picked up by our measure is associated with *less intensive algorithmic trading*. The correlations with the measures of the intensity of market making activity are especially important, as they indicate that our measure indeed picks up variation in arbitrage intensity.

Test 1: Additions to and deletions from the S&P 500 index

To further investigate the informational content of our measure, we study additions to and deletions from the S&P 500 index[24]. When a stock is added to the index, it becomes exposed to index trading and thus we can expect an abrupt increase in arbitrage activity. Similarly, we should see a drop in index trading activity when a stock is deleted from the index. Our measure should reflect these changes.

Data and methodology We study S&P 500 index inclusions and deletions between January 1, 2005 and December 31, 2016. In total we have 148 inclusion events and 143 deletion events. As usual, inclusion events are excluded if the new firm is a spin-off of a firm, and deletion events are excluded if the old firm engaged in a merger or acquisition. The remaining sample contains 126 inclusion events and 72 deletion events.

For each event, we estimate the probability of simultaneous trading with SPY for the 90 working day periods before and after the event[25]. To account for possible time variation, we also perform estimation separately for each calendar year.

The changes in the index composition are preannounced. However, the channel that we are interested in is arbitrage activity between the ETFs tracking the S&P 500 and the basket. There is no reason for the new firm to be used as a part of the arbitraging basket before it is included in the index. Thus, we only consider the effective date of the inclusion or deletion event.

Results Table 5 displays the change in the probability of simultaneous trading (with SPY) for the stocks added to and deleted from the S&P 500 index. To gauge the magnitude of the effect, the third column shows the average probability of simultaneous trading before the inclusion or deletion event.

Panel A in Table 5 confirms that a stock added to the index becomes more likely to be traded simultaneously with SPY. The average effect across all events is an increase in probability of 13%. The effect is especially pronounced in the earlier subsample. In 2007, the probability of trading increases by almost 70%. In 2015, the increase is about 20%, and in 2016 about 5% on average. We focus on relative changes in probabilities rather than absolute changes, because the firms that are added to the index are relatively small firms which receive relatively *small weights* in the index. Such firms are likely to be omitted from some arbitrage portfolios as traders economize on trading costs. As a result, one cannot expect to detect large absolute changes in probabilities. Statistically, we find that an increase in probability is significant in 85 of 126 index addition events

[24]We follow large empirical literature that uses this identification scheme, see Wurgler and Zhuravskaya (2002) and Barberis et al. (2005). Similarly, Ben-David et al. (2018) use the annual reconstitution of the Russell indexes.

[25]For robustness, we repeated the estimation with longer periods. The results are similar and can be sent upon request.

Table 5: Changes in the probability of joint trading with SPY of stocks added to and deleted from the S&P 500 index.

Panel A: Stocks added to the S&P 500 index.

	Mean Δp	Median Δp	Mean p_{before}	Number of Events
2005	0.0092	0.0092	0.0255	1
2006				0
2007	0.0475	0.0417	0.0697	6
2008				0
2009	-0.0048	-0.0057	0.0931	4
2010	0.0006	0.0036	0.1076	9
2011	-0.0012	-0.0001	0.0992	13
2012	0.0079	0.0091	0.0821	14
2013	0.0186	0.0174	0.0855	16
2014	0.0112	0.0096	0.0677	12
2015	0.0196	0.0129	0.0882	24
2016	0.0039	0.0012	0.0719	27
all years	0.0110	0.0078	0.0830	126

Panel B: Stocks deleted from the S&P 500 index.

	Mean Δp	Median Δp	Mean p_{before}	Number of Events
2005				0
2006				0
2007				0
2008	0.0617	0.0617	0.1864	1
2009	-0.0273	-0.0301	0.1243	3
2010	-0.0111	-0.0077	0.1130	6
2011	-0.0146	-0.0182	0.1096	9
2012	0.0020	-0.0030	0.0964	9
2013	-0.0121	-0.0096	0.1225	12
2014	-0.0014	0.0018	0.1032	9
2015	0.0165	0.0178	0.1250	12
2016	-0.0002	0.0018	0.1211	11
all years	-0.0023	-0.0043	0.1156	72

(under the assumption that each second is an independent observation). Overall, the additions results provide strong support for our measure of index trading.

The deletion results are presented in Panel B. We see mostly negative numbers, confirming that a firm deleted from the index experiences a decrease in the probability of joint trading with SPY; on average the probability decreases by 2% when we consider all deletion events in our sample. The magnitude of the effect in relative terms can be as large as -22% in 2009, -13% in 2011, and -10% in 2010 and 2013. We find that a decrease in probability is significant in 39 of 72 index deletion events (under the same assumption as before). The number of deletions that can be used for our exercise is much smaller, as frequently a firm is deleted due to a merger or an acquisition, which we exclude from our analysis.

Overall, our results confirm that the probability of simultaneous trading with SPY increases after an addition to the index and decreases after a deletion. We consider this result as evidence in favor of using the probability of simultaneous trading as a measure of index trading intensity.

One may be concerned that our results reflect the overall intensification of trading. However, in this case, the increase in the trading volume is itself driven by index trading, and thus represents our effect.

Test 2: SPY creations and redemptions

Another way to test our measure is to analyze creations and redemptions of ETF units. Each Authorized Participant can exchange ETF shares for a basket of underlying securities[26]. Imagine that the price of ETF shares happens to be larger than the weighted price of the underlying securities. An AP purchases the underlying securities in accordance with their weights in the index and simultaneously sells ETF shares. If the accumulated position is sufficiently large by the end of the day, the AP brings the underlying securities to the fund sponsor, receives ETF shares in exchange, and uses them to cover his short position. As a result, the number of outstanding ETF shares increases and correspondingly increases the ETF's position in each of its underlying securities. Hence, intensive arbitrage transactions during the day are reflected in changes in ETF holdings. Thus, creations and redemptions can be used to test our measure of the intensity of arbitrage transactions[27].

Data and methodology We use the ETF Global database for information on daily holdings of SPY. The sample period starts in 2012. For each individual component, we

[26] Although only in large amounts and at a fee.

[27] Another reason for changes in holdings is rebalancing. However, most ETFs passively track an index and thus rarely have to rebalance their portfolio. Moreover, rebalancing activity due to changes in the index composition or weightening, typically happens at the end of the day, whereas our measure of probability is calculated as an average over the entire day.

calculate daily changes in SPY holdings and take the logarithm of the absolute change, $\log|\Delta_{j,t}|$.

For estimation purposes, we transform the probability of simultaneous trading and calculate the logarithm of the odds as $y_{j,t} = \ln\frac{\pi_{j,t}}{1-\pi_{j,t}}$. To test our hypothesis, we estimate a panel regression of the transformed daily probabilities of simultaneous trading with SPY on changes in SPY holdings. To distinguish redemptions and creations, we also estimate a specification that separates positive and negative changes in SPY holdings. We include the logarithm of the dollar volume and the full set of year-month dummies as controls. Our database covers the period from 2012 to 2016.

Results Table 6 shows our main results. The coefficients of $\log|\Delta|$ are all positive and significant. Our results confirm that changes in SPY holdings, and thus the intensity of arbitrage transactions, are associated with a higher probability of simultaneous trading.

In our main exercise, we do not distinguish redemptions and creations. The last column in Table 6 breaks down changes in SPY position into positive and negative ones. We find no difference in the estimated coefficients, which is also consistent with the interpretation of measures as reflecting the intensity of arbitrage transactions.

Test 3: Mutual fund fire sales

The unique structure of mutual funds gives us another way to test our measure. When a mutual fund experiences a large outflow of assets, its asset position has to be reduced. If the selling pressure is large enough, it can temporarily drive down the prices of individual securities. If the prices of some of the individual S&P 500 components fall, the overall value of the index decreases and thus deviates from the price of SPY shares. This opens up an arbitrage possibility. Hence, we should observe intensification of index arbitrage transactions, which should be captured by our probability of simultaneous trading with SPY.

Data and methodology We follow the literature[28] and use Thomson Returns data on mutual fund holdings and the CRSP mutual funds database to identify exposed mutual funds and calculate the overall selling pressure on each stock in each quarter. Ideally, the fire sales of mutual funds should represent a series of sell orders unrelated to any news about stock value. Thus, we follow Edmans et al. (2012) and calculate the (logarithm) of the total mutual fund hypothetical sales of stock j in quarter t, denoted as $\log|MFHS_{j,t}|$, see appendix A.8 for details. We estimate a panel regression of our transformed probability of simultaneous trading on mutual fund fire sales. We include the logarithm of the

[28] See for example Edmans et al. (2012), Dessaint et al. (2018), Phillips and Zhdanov (2013). See also the recent discussion of this identification scheme in Wardlaw (2018) and Berger (2019).

Table 6: Probability of simultaneous trading and SPY redemptions and creations.

Panel A: $y_{j,t} = c_j + \alpha_\tau + \gamma \log|\Delta_{j,t}| + \delta X_{j,t} + \varepsilon_{j,t}$, where $y_{j,t} = \ln \frac{\pi_{j,t}}{1-\pi_{j,t}}$ is the logarithm of the odds, $\pi_{j,t}$ is the average probability of simultaneous trading of stock j and SPY in year t, $\log|\Delta_{j.t}|$ is a change in SPY holdings of firm j, and controls $X_{j,t}$ contain the logarithm of dollar volume,α_τ are year-month dummies. The sample covers the 2012-2016 period. Standard errors are clustered at the firm level, t-statistics are in parenthesis.

	(I)	(II)	(III)
$\log\|\Delta\|$	0.007	0.010	
	(4.53)	(6.40)	
$\log\|\Delta\|^+$			0.010
			(6.22)
$\log\|\Delta\|^-$			0.010
			(6.58)
DollarVolume	Yes	Yes	Yes
Year-month dummies		Yes	Yes
R-sq within	0.30	0.47	0.47
R-sq between	0.72	0.70	0.70
N	368,528	368,528	368,528
N groups	449	449	449

Table 7: Probability of simultaneous trading and mutual fund fire sales.

Panel A: $y_{j,t} = c_j + \alpha_t + \gamma \log|MFHS_{j,t}| + \delta X_{j,t} + \varepsilon_{j,t}$, where $y_{j,t} = \ln\frac{\pi_{j,t}}{1-\pi_{j,t}}$ is the logarithm of the odds, $\pi_{j,t}$ is the average probability of simultaneous trading of stock j and SPY in year t, $\log|MFHS_{j,t}|$ denotes mutual fund fire sales of stock j in quarter t, and controls $X_{j,t}$ contain the logarithm of dollar volume and the indicator on zero MFHS, α_t are year-quarter dummies. We only consider domestic equity funds, and exclude sectoral funds. The flow threshold is set at 5%. Standard errors are clustered at the firm level, t-statistics are in parenthesis.

	(I)	(II)
log|MFHS|	0.074	0.032
	(7.62)	(4.34)
ZeroMHFS	Yes	Yes
DollarVolume	Yes	Yes
Year-quarter dummies	No	Yes
R-sq within	0.41	0.83
R-sq between	0.87	0.85
N	18,873	18,873
N groups	430	430

dollar volume and an indicator on zero sales as controls, and also include year-quarter dummies.

Results Table 7 displays our results. The coefficient of $\log|MFHS_{j,t}|$ is positive and significant in all specifications. Active mutual fund sales are associated with an increased probability of simultaneous trading with SPY, which is consistent with our interpretation of increased arbitrage activity.

Comparison to other measures

In what follows, we briefly identify and discuss issues of commonly used arbitrage measures.

A standard way to measure the intensity of ETF arbitrage at the stock level is to use changes in ETF holdings. Ben-David et al. (2018) construct an ETF ownership measure by multiplying assets under management by the index weight of the security. A similar approach is taken in Brown et al. (2018) and Israeli et al. (2017). Currently, ETF redemptions and creations data are available directly at daily frequency. As far as we are aware, only Da and Shive (2018) take a different approach and use ETF turnover to measure arbitrage intensity, but at the ETF level only.

To start, measures based on ETF ownership cannot account for passive investment. Growth in the assets under management does not necessarily translate into growth in the intensity of arbitrage transactions. Even for SPY, we have seen substantial fluctuations in the dollar volume and the overall intensity of trading, even though its assets have been steadily growing. Moreover, as documented using index weights and AUM may be imprecise for a sizable proportion of funds that replicate their target index by investing in a representative basket of securities rather than the entire index basket. Brogaard et al. (2019) document that at least 22% of ETFs replicate their target index, including six of the ten largest ETFs as of 2018.

Another concern is that measures based on ETF creations and redemptions do not account for the netting of positions over the day. One can only observe the end of the day holdings. However, during the day arbitrageurs can trade in both directions, arbitraging away positive and negative deviations between the ETF price and the price of the underlying basket. If the frequency and magnitude of positive and negative deviations are roughly comparable, the accumulated end-of-the-day position can be relatively small. Thus, we may have days with quite intensive arbitrage transactions, but close to zero net changes in the holdings. Hence, by using changes in ETF holdings, one may significantly underestimate the intensity of arbitrage transactions. Moreover, Evans et al. (2019) document that authorized participants have incentives to delay creation of ETF shares. ETF creation can only be done in bulk (typically 50,000 ETF shares) and at a fee, thus APs may prefer to wait until their position builds up to a size of the creation unit, which might take several days. In addition, if the underlying market is not perfectly liquid and if the ETF's order flow if expected to mean revert, APs may prefer to take the unhedged position by only selling overvalued ETF shares and waiting for the mispricing to be reversed (Evans et al. (2019) call such selling of ETF shares that have not yet been created as 'operational shorting'). Evans et al. (2019) argue that the time period between the original mispricing/arbitrage event and the corresponding creation event may be as large as six days. Hence, daily creations and redemptions represent highly imprecise and noisy measure of the intensity of arbitrage, which can explain the absence of any effects of creations and redemptions on the underlying markets documented by Box et al. (2019).

More importantly, measures that use fixed index weights are specifically constructed to reflect ETF creations and redemptions, and therefore do not account for statistical arbitrage. As we have seen, firms with larger index weight are traded more often simultaneously with SPY (controlling for the intensity of overall trading), which suggests the presence of statistical arbitrage. Not accounting for statistical arbitrage distorts the estimates of the intensity of arbitrage transactions: overestimates for the small firms, and underestimates for the large firms.

In contrast, our approach allows us to measure arbitrage intensity over any interval of the day and accounts for statistical arbitrage. Thus, it represents a more precise measure of the intensity of arbitrage transactions than existing approaches in the literature.

Our direct approach has yet another advantage. It allows us to separately measure the effect of a stock characteristic on the intensity of arbitrage transactions from its effect on the magnitude of distortion that these arbitrage transactions cause. For example, without a direct measure of arbitrage intensity, it is impossible to separately measure these two opposite effects of liquidity (see the discussion in Section 2). Shim (2018) estimates price impact to account for differential stock price sensitivity to mechanical arbitrage, but he abstracts from statistical arbitrage and thus does not consider the differential effect of liquidity on the intensity of arbitrage transactions across firms. In contrast, because we directly measure the intensity of arbitrage transactions, we can assess both effects independently. In particular, our results show a strong effect of liquidity on the market distortion, but a negative effect of liquidity on arbitrage intensity, thus further confirming our mechanism.

More generally, our measure can be viewed as complementary to existing measures of algotrading intensity. It is worth emphasizing, that our measure only requires access to a widely used commercial dataset such as TAQ, while existing measures of algotrading require access to regulatory data. Limited access to regulatory data is one of the reasons why measures of statistical arbitrage are scarce. For example, NASDAQ only provides a limited sample of randomly selected stocks (Brogaard et al. (2014) only use the data on 120 stocks) which makes it impossible to study the cross-market arbitrage. We are aware of only one study, Buyuksahin and Robe (2014) that uses a unique regulatory dataset to directly study cross arbitrage in the futures markets, by comparing the positions of the same traders at the equity futures and the commodity futures markets[29]. In contrast, our measure can be easily calculated and applied to an arbitrary combination of assets that one wishes to study over an arbitrary period of time. It should also be noted that recently created MIDAS data cannot be directly used to measure statistical arbitrage. For example, the correlations between the four measure of algotrading calculated for the SPY and for the individual constituents show very low correlations and thus fail to pick up variation in arbitrage activity.

7 Additional evidence: ETF-level analysis

So far, we have focused on one particular ETF (SPY) and explored the intensity of arbitrage transactions at the individual firm level. To further support our results, we perform an additional empirical exercise, this time exploring the reaction of ETFs to oil

[29]It should also be noted, that the task of assigning trades to the firms is not straightforward. Large trading firms such as Goldman tend to utilize multiple IDs making it impossible to comprehensively describe their trading patterns . To solve this issue, Brogaard et al. (2014) only work with 26 small independent proprietary firms.

shocks as a function of the intensity of arbitrage transactions.

Intensity of arbitrage transactions at the ETF level

To predict the intensity of arbitrage transactions at the ETF level, we propose to use a measure of liquidity at the ETF market. Indeed, arbitrage activity requires transactions with ETF shares. If the ETF market is not liquid enough, certain price deviations can stay unarbitraged for some time, as it may not be profitable to eliminate them by purchasing or selling ETF shares in a thin market. Therefore, more liquid ETFs can be more involved in intensive arbitrage.

One of the standard ways to measure liquidity is to use the bid-ask spread. The spread is basically zero for SPY, but can be extremely large for smaller ETFs trading less liquid stocks.

Methodology

We follow the same approach as described above to estimate oil betas for each ETF. Similarly to our main exercise, we then estimate the following fixed effect panel regression:

$$\beta_{j,t} = c_j + \alpha_t + \gamma_t s_{j,t} + \delta_t' X_{j,t} + \varepsilon_{j,t} \tag{3}$$

where $\beta_{j,t}$ is the estimated oil beta of ETF j in year t, $s_{j,t}$ is (the logarithm of) the average bid-ask spread on the ETF j's market over the same year, and $X_{j,t}$ is a vector of control variables. Controls include the logarithm of assets under management. The coefficient c_j captures the unobserved time-invariant individual effect, which can be interpreted as the true sensitivity to oil shocks of a particular set of firms that an ETF tracks.

Data and descriptive statistics

We consider U.S. non-energy equity ETFs. We download a list of equity ETFs from etf.com and keep those with more than \$250 mln in assets under management as of November 2018. In total that leaves us with 540 funds. As we only keep U.S. equity funds, and exclude leveraged and inverse ETFs, the remaining sample contains 302 funds. We do not have high frequency data on 18 ETFs, and thus the final sample has 284 funds. We find only a very limited number of energy-related ETFs: 5 funds invest in a broad set of energy related firms, one ETF specializes in equipment and services, and two on exploration and production. We exclude these funds from our analysis. As before, we use the Trade and Quote Database (TAQ) for the high-frequency data, and our sample covers the period from 2005 to 2016.

The main source of data on ETFs is the ETF Global database. We obtain the data on NAV and shares outstanding from 2006 to calculate assets under management for the entire period. However, other variables, including expenses, are available only from 2012. As most funds do not change expenses over time, we do not incorporate it into our analysis.

We use two main sources to calculate the bid-ask spread: CRSP and WRDS Intraday Database that reports the best bids and offers as of 1 pm. We use these quotes to calculate the bid-ask spread percentage. Unfortunately, the WRDS intraday indicators are calculated up to 2014 only, when TAQ switched to milliseconds. ETF Global also provides a measure of the bid-ask spread, but is available starting from 2012 only.

Descriptive statistics

We provide detailed information on the composition of our sample in Table 12 in the appendix. The vast majority of our ETFs are broad-based funds investing in a broad set of stocks.

Table 12 also displays estimated oil betas. We can see a familiar shift from negative to positive betas in 2008-2009. By comparing Table 12 with Table 11, we notice that oil betas are smaller for ETFs compared to the S&P 500 firms. Industry ETFs invest in a broader set of firms, including smaller firms outside major indices. Hence, we find that larger firms on average tend to display a stronger reaction to oil shocks. This result is yet another piece of evidence in favor of the index trading hypothesis.

Table 8 provides some basic statistics for our measures of the bid-ask spread. CRSP and WRDS show similar first and second sample moments, and the correlation is also large, 0.94. The ETF Global measure has a larger mean and variance, and it is less correlated with the other two measures.In our main exercise, we will use CRSP and WRDS bid-ask measures.

Results

Table 9 reports our main results. The spread coefficients are negative and significant for both measures of ETF liquidity and for most years. The negative sign means that more liquid ETFs that have lower spreads display a stronger reaction to oil shocks. As a lower spread is likely to indicate more active arbitrage transactions, this confirms our previous firm-level results that intensive index trading is associated with a stronger response to oil shocks. To see the magnitude of the effect, consider a decrease in the CRSP (log) spread from the third to the forth quantile, which equals -1.15. Such a decrease in the spread would imply an increase in oil beta by 0.022 - 0.026 in the earlier period of our sample (we use 1 minute returns). Thus, the effect is similar in magnitude to the effect of index arbitrage on oil betas for the stocks.

We find a lack of significance in the later years. One potential reason is that ETFs have grown substantially and have reached sufficient liquidity. Liquidity of ETF markets

Table 8: Different measures of bid-ask spreads.

ETF Global offers average intraday bid and asks, but only after 2012. CRSP reports the best bid and offer at market closing. WRDS Intraday Database reports the best bids and offers as of 1 pm; the data are available up to 2014. The table reports the descriptive statistics and cross-sectional correlations of different measures of spreads (in logs) over the overlapping period 2012-2014.

Panel A: Descriptive statistics

	Mean	Std.Dev	Min	Max
ETF Global	-6.26	1.39	-9.82	-1.32
CRSP	-7.25	0.92	-9.84	-3.15
WRDS, 1pm	-7.16	0.88	-9.81	-3.53

Panel B: Cross-sectional correlations

	ETF Global	CRSP	WRDS, 1pm
ETF Global	1.00		
CRSP	0.45	1	
WRDS, 1pm	0.46	0.94	1

Table 9: Impact of the probability of joint trading on oil betas in the panel of U.S. non-energy equity ETFs.

We estimate the following panel regression of ETF betas on bid-ask spread: $\beta_{j,t} = c_j + \alpha_t + \gamma_t s_{j,t} + \delta_t' X_{j,t} + \varepsilon_{j,t}$, where $\beta_{j,t}$ is the estimated oil beta of ETF j in year t, $s_{j,t}$ is (the logarithm of) the average bid-ask spread on the ETF j's market over the same year, and $X_{j,t}$ is a vector of control variables, c_j corresponds to unobserved ETF-level individual effects. Controls include the logarithm of assets under management. St.err are clustered at the ETFs level, second column in displays t-statistics.

The sample starts in 2010 and ends in 2016 when we use CRSP bid-ask data, or in 2014 if we use WRDS 1 pm data.

	CRSP		WRDS 1pm	
	1 min returns	5 min returns	1 min returns	5 min returns
Spread -				
2010	-0.023 (-2.68)	-0.011 (-1.07)	-0.031 (-4.31)	-0.038 (-3.33)
2011	-0.023 (-3.26)	-0.010 (-1.00)	-0.021 (-3.38)	-0.026 (-2.21)
2012	-0.019 (-3.16)	-0.024 (-2.30)	-0.015 (-1.76)	-0.028 (-2.15)
2013	-0.013 (-2.48)	-0.020 (-2.17)	-0.017 (-2.64)	-0.040 (-3.31)
2014	-0.001 (-0.23)	0.004 (0.52)	0.003 (0.44)	-0.008 (-0.68)
2015	-0.004 (-0.74)	0.003 (0.27)		
2016	0.005 (0.45)	0.011 (0.66)		
AUM	Yes	Yes	Yes	Yes
Time varying				
- α_t	Yes	Yes	Yes	Yes
- γ_t, δ_t	Yes	Yes	Yes	Yes
R-sq within	0.31	0.58	0.24	0.58
R-sq between	0.33	0.23	0.39	0.35
N	1,385	1,385	1,010	1,010
N groups	247	247	247	247

is no longer an issue, and the intensity of arbitrage transactions is driven only by the liquidity of the underlying markets. In this case, the bid-ask spread on the ETF market becomes a poor predictor of the intensity of arbitrage transactions.

Overall, our ETF-level analysis confirms our previous results showing that index trading intensifies the market reaction to fundamental oil shocks.

8 Conclusion

Exchange traded funds have transformed the investment industry. By providing access to the markets to a broader set of market participants, ETFs have enhanced trading opportunities and have improved risk sharing. However, the side effects of such a drastic transformation can be substantial. As arbitrage is an essential feature of ETF trading, we investigate the effect of arbitrage on price discovery. We introduce a new measure of the intensity of arbitrage transactions and use it to show that ETF arbitrage affects price efficiency by distorting the market reaction to fundamental shocks.

However, our results are likely to underestimate the effect of the ETF presence on the market reaction to oil shocks. So far, we have only considered the mechanical arbitrage channel that describes how the *propagation* of fundamental shocks can be distorted, whereas, for example, the informational channel can explain why certain shocks become more likely to hit the market, as informed traders reoptimize their trading strategies when ETFs are traded.

In a recent theoretical paper, Bhattacharya and O'Hara (2018) study the informational effects of ETFs in a factor model of prices. A newly introduced ETF attracts new clientele with additional information about the fundamental values of the underlying securities, both the systemic factor and individual shocks. The market makers on the underlying markets can now extract information from the ETF price. However, if informed traders use short-term trading strategies, the introduction of an ETF facilitates a herding equilibrium (in the spirit of Froot et al. (1992)). In a herding equilibrium, informed traders rationally choose to trade based on information on the systematic factor only and completely disregard all idiosyncratic information. Each informed trader on the market for each individual security, can accurately anticipate the information flow into their own market once market makers learn from observing other markets and the ETF price, and thus can foresee the price bump and profit from it. As a result, in equilibrium asset prices exhibit lower informational efficiency. In other words, the market reaction to news is distorted, as all individual information is disregarded and never becomes reflected in prices. It should also be noted that endogenous information acquisition and processing further reinforce the herding. Over time, any incentives to acquire and process idiosyncratic information cease if informed traders use short-term strategies and trade based on a systematic piece of information only.

Some of our results are suggestive of herding. We document a dramatic shift in average oil sensitivities at the end of 2008. In our analysis, we restrict the sample to start in 2010 to avoid this shift, and we include sector by year indicators to pick up any fluctuations in average industry sensitivities over time. We maintain the assumption that any changes in average industry sensitivities to oil shocks are not related to ETF arbitrage and potentially reflect fundamental changes in the U.S. market (perhaps due to the development of unconventional oil). However, instead, these changes may be yet another consequence of the presence of ETFs, perhaps due to a switch to a herding equilibrium. More research is needed to test the informational channel, we plan to pursue this in future research by extending our ETF-level analysis.

References

Agarwal, V., P. Hanouna, R. Moussawi, and C. W. Stahel (2018). Do ETFs increase the commonality in liquidity of underlying stocks? Working Paper.

Anatolyev, S., S. Seleznev, and V. Selezneva (2018). The role of inventories in shaping the behavior of oil prices. Working paper.

Anatolyev, S., S. Seleznev, and V. Selezneva (2019). Does index arbitrage distort market reaction to oil shocks? Working paper.

Barberis, N., A. Shleifer, and J. Wurgler (2005). Comovement. *Journal of Financial Economics 75*(2), 283–317.

Ben-David, I., F. Franzoni, and R. Moussawi (2018). Do ETFs increase volatility? *Journal of Finance 73*(6), 2471–2535.

Berger, E. (2019). Selection bias in mutual fund flow-induced fire sales: Causes and consequences. Working paper.

Bhattacharya, A. and M. O'Hara (2018). Can ETFs increase market fragility? Effect of information linkages in ETF markets. Working paper.

Bollerslev, T., S. Z. Li, and V. Todorov (2016). Roughing up beta: Continuous versus discontinuous betas and the cross section of expected stock returns. *Journal of Financial Economics*.

Bollerslev, T., V. Todorov, and S. Z. Li (2013). Jump tails, extreme dependencies, and the distribution of stock returns. *Journal of Econometrics 172*(2), 307–324.

Box, T., R. Davis, R. Evans, and A. Lynch (2019). Intraday arbitrage between ETFs and their underlying portfolios. Working Paper.

Brogaard, J., D. Heath, and D. Huang (2019). The rise of ETF trading and the bifurcation of liquidity. Working Paper.

Brogaard, J., T. Hendershott, and R. Riordan (2014). High-frequency trading and price discovery. *The Review of Financial Studies 27*(8), 2267–2306.

Brown, D. C., S. Davies, and M. Ringgenberg (2018). ETF flows, non-fundamental demand, and return predictability. Working paper.

Buyuksahin, B. and M. A. Robe (2014). Speculators, commodities and cross-market linkages. *Journal of International Money and Finance 42*, 38–70.

Chaboud, A. P., B. Chiquoine, E. Hjalmarsson, and C. Vega (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance 69*(5), 2045–2084.

Coval, J. and E. Stafford (2007). Asset fire sales (and purchases) in equity markets. *Journal of Financial Economics 86*(2), 479–512.

Da, Z. and S. Shive (2018). Exchange traded funds and asset return correlations. *European Financial Management 24*(1), 136–168.

Dessaint, O., T. Foucault, L. Fresard, and A. Matray (2018). Noisy stock prices and corporate investment. Rotman School of Management Working Paper 2707999.

Edmans, A., I. Goldstein, and W. Jiang (2012). The real effects of financial markets: The impact of prices on takeovers. *The Journal of Finance 67*(3), 933–971.

Evans, R., R. Moussawi, M. Pagano, and J. Sedunov (2019). ETF short interest and failures-to-deliver: Naked short-selling or operational shorting? Working Paper.

Foucault, T., R. Kozhan, and W. W. Tham (2016). Toxic arbitrage. *The Review of Financial Studies 30*(4), 1053–94.

Froot, K. A., D. S. Scharfstein, and J. C. Stein (1992). Herd on the street: Informational inefficiencies in a market with short-term speculation. *The Journal of Finance 47*(4), 1461–1484.

Glosten, L. R., S. Nallareddy, and Y. Zou (2016). ETF activity and informational efficiency of underlying securities. Working paper.

Hasbrouck, J. (2009). Trading costs and returns for us equities: Estimating effective costs from daily data. *The Journal of Finance 64*(3), 1445–1477.

Hendershott, T., C. M. Jones, and A. J. Menkveld (2011). Does algorithmic trading improve liquidity? *The Journal of Finance 66*(1), 1–33.

Hong, H., J. D. Kubik, and T. Fishman (2012). Do arbitrageurs amplify economic shocks? *Journal of Financial Economics 103*(3), 454–470.

Israeli, D., C. M. Lee, and S. A. Sridharan (2017). Is there a dark side to exchange traded funds? an information perspective. *Review of Accounting Studies 22*(3), 1048–1083.

Kilian, L. and C. Park (2009). The impact of oil price shocks on the us stock market. *International Economic Review 50*(4), 1267–1287.

Lee, C. and E. M. Watts (2018). Tick size tolls: Can a trading slowdown improve price discovery? Working Paper.

O'Hara, M., C. Yao, and M. Ye (2014). What's not there: Odd lots and market data. *The Journal of Finance 69*(5), 2199–2236.

Phillips, G. M. and A. Zhdanov (2013). R&d and the incentives from merger and acquisition activity. *The Review of Financial Studies 26*(1), 34–78.

Saglam, M., T. Tuzun, and R. Wermers (2019). Do ETFs increase liquidity?. Working Paper.

Shim, J. J. (2018). Arbitrage comovement. Working paper.

Staer, A. and P. Sottile (2018). Equivalent volume and comovement. *Quarterly Review of Economics and Finance 68*, 143–157.

Todorov, V. and T. Bollerslev (2010). Jumps and betas: A new framework for disentangling and estimating systematic risks. *Journal of Econometrics 157*(2), 220–235.

Wardlaw, M. (2018). Measuring mutual fund flow pressure as shock to stock returns. Working paper.

Weller, B. M. (2018). Does algorithmic trading reduce information acquisition? *The Review of Financial Studies 31*(6), 2184–2226.

Wurgler, J. and E. Zhuravskaya (2002, October). Does arbitrage flatten demand curves for stocks? *Journal of Business 75*(4), 583–608.

Abstrakt

Ukazujeme, že ETF arbitráž deformuje reakci trhu na fundamentální šoky. Potvrzujeme tuto hypotézu vytvořením nové míry intenzity arbitrážních transakcí na úrovni jednotlivých akcií a s použitím případové analýzy odhadujeme tržní reakci na ekonomické šoky. Naše míra intenzity arbitráže představuje pravděpodobnost simultánních obchodů ETF podílů a podkladových akcií a je odhadnuta s použitím vysokofrekvenčních dat. Náš přístup je přímý a bere v potaz statistickou arbitráž, pasivní investiční strategie a vyrovnání arbitrážních pozic během dne, čehož existující míry nejsou schopny. Provádíme řadu empirických testů, zahrnujících použití quasi-přirozeného experimentu, k potvrzení, že naše míra zachycuje fluktuace v intenzitě arbitrážních transakcí. Zaměřujeme se na ropné šoky, jelikož obsahují výrazné idiosynkratické komponenty, které usnadňují identifikaci našeho mechanismu a interpretaci výsledků. Ropné šoky jsou identifikovány s použitím týdenních oznámeních o zásobách ropy.

Working Paper Series
ISSN 1211-3298
Registration No. (Ministry of Culture): E 19443

Individual researchers, as well as the on-line and printed versions of the CERGE-EI Working Papers (including their dissemination) were supported from institutional support RVO 67985998 from Economics Institute of the CAS, v. v. i.

Specific research support and/or other grants the researchers/publications benefited from are acknowledged at the beginning of the Paper.

Published by
Charles University, Center for Economic Research and Graduate Education (CERGE)
and
Economics Institute of the CAS, v. v. i. (EI)
CERGE-EI, Politických vězňů 7, 111 21 Prague 1, tel.: +420 224 005 153, Czech Republic.
Printed by CERGE-EI, Prague
Subscription: CERGE-EI homepage: http://www.cerge-ei.cz

Phone: + 420 224 005 153
Email: office@cerge-ei.cz
Web: http://www.cerge-ei.cz

Editor: Byeongju Jeong

The paper is available online at http://www.cerge-ei.cz/publications/working_papers/.

ISBN 978-80-7343-458-8 (Univerzita Karlova, Centrum pro ekonomický výzkum
a doktorské studium)
ISBN 978-80-7344-515-7 (Národohospodářský ústav AV ČR, v. v. i.)