Stability Effects of Arbitrage in Exchange Traded Funds: An Agent-Based Model

Megan Shearer
University of Michigan–Ann Arbor
Ann Arbor, Michigan, USA
shearerj@umich.edu

David Byrd
Bowdoin College
Brunswick, Maine, USA
d.byrd@bowdoin.edu

Tucker Hybinette Balch
J.P. Morgan AI Research
New York, New York, USA
tucker.balch@jpmchase.com

Michael P. Wellman
University of Michigan–Ann Arbor
Ann Arbor, Michigan, USA
wellman@umich.edu

ABSTRACT

An index-based exchange traded fund (ETF) with underlying securities that trade on the same market creates potential opportunities for arbitrage between price deviations in the ETF and the corresponding index. We examine whether ETF arbitrage transmits small volatility events, termed mini flash crashes, from one of its underlying symbols to another. We address this question in an agent-based, simulated market where agents can trade an ETF and its two underlying symbols. We explore multiple market configurations with active and inactive ETF arbitrageurs. Through empirical game-theoretic analysis, we find that when arbitrageurs actively trade, background traders' surplus increases because of the increased liquidity. Arbitrage helps the ETF more accurately track the index. We also observe that when one symbol experiences a mini flash crash, arbitrage transmits a price change in the opposite direction to the other symbol. The size of the mini flash crash depends more on the market configuration than the arbitrageurs, but the recovery of the mini flash crash is faster when arbitrageurs are present.

KEYWORDS

ETF arbitrage; mini flash crash; empirical game-theoretic analysis

ACM Reference Format:
Megan Shearer, David Byrd, Tucker Hybinette Balch, and Michael P. Wellman. 2021. Stability Effects of Arbitrage in Exchange Traded Funds: An Agent-Based Model. In *2nd ACM International Conference on AI in Finance (ICAIF'21), November 3–5, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3490354.3494414

1 INTRODUCTION

An *exchange traded fund* (ETF) is a portfolio of securities that trades on the stock market. The *underlying securities* in an ETF's portfolio can be any traded entity, such as stocks, bonds, or commodities. ETFs have become a popular investment vehicle, as they provide convenient access to portfolio trades, offering average investors apparent liquidity [2]. Investing in an ETF, rather than each security in its portfolio, requires fewer trades, and therefore offers the opportunity for diversification with lower trading costs [18].

We focus on *index-based ETFs*, which are ETFs whose underlying securities are based on a *market index*. Some examples of market indexes are the Standard and Poor's (S&P) 500 and Dow Jones Industrial Average (DJIA) which track market performance, or the Volatility Index (VIX) which tracks market volatility. Though designed to track the index, the actual price of the index-based ETF is determined through trading it as a symbol on the stock market. When the market is open, participants can simultaneously observe the trading price of the ETF and calculate the index. Any disparity composes a potential arbitrage opportunity between the ETF and its underlying securities.

Arbitrage trading can help an ETF's price track its corresponding market index [2]. It may also introduce or reinforce other dependencies. For example, arbitrage may tether the price volatility of the ETF's underlying symbols [4]. The rise in market-wide price volatility in recent years has raised concerns about economic growth and investor trust [11, 21]. Price co-movements in ETF portfolios have led some to question the role ETFs play in amplifying volatility [4, 28]. We study this issue, focusing specifically on *mini flash crashes*, which are short volatility events where the price rapidly changes, then quickly reverts [31]. ETF arbitrage may channel mini flash crashes from one underlying symbol to others. ETF arbitrage may also have other effects, for example on the distribution of surplus among traders.

To explore these questions, we develop an agent-based model (ABM), populated by trading agents implementing algorithms commonly employed in agent-based finance literature. Our market consists of multiple stocks which compose the portfolio of an ETF. We study our simulated market with active and inactive ETF arbitrageurs. We then induce a mini flash crash to one of the underlying securities, by injecting a trader who submits a series of large marketable limit orders (a common pattern of behavior triggering such mini flash crashes [6, 17, 31]). We then analyze the effect of arbitrage activity on the market response.

Using agent-based simulation allows us to examine otherwise identical market environments with and without the presence of

ICAIF'21, November 3–5, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-9148-1/21/11…$15.00
https://doi.org/10.1145/3490354.3494414

ETF arbitrage. We explore the impact of mini flash crashes by comparing the two scenarios, with strategic equilibration of trading agents under each setting. We implement two strategies that trade on arbitrage opportunities between the ETF and its underlying symbols. This allows us to observe whether trading strategies dependent on an ETF and its portfolio directly impact their underlying symbols' price volatility.

We employ a simulated market model using a standard limit order book and message system similar to the US stock exchange NASDAQ [10]. Our market contains two symbols which compose an ETF's portfolio. The ETF also trades on the market. Our model is populated with 27 background trading agents and four arbitrageurs. A conservative and an aggressive arbitrageur trade only on the ETF, and a conservative and an aggressive arbitrageur trade on both ETF and underlying symbols. We also use one impact agent to submit a series of large, marketable sell orders to create a mini flash crash.

Through empirical game-theoretic analysis, we determine the optimal trading strategies for all agents when arbitrageurs are active and inactive, in order to examine the impact of ETF arbitrage on market welfare and volatility. Arbitrageurs are highly profitable and background agents' average surplus is impacted significantly by the arbitrage activity. Background traders' final payoffs increase with active arbitrage because arbitrageurs increase liquidity by submitting marketable orders. Active arbitrage also increases the volatility of its underlying symbols around events like mini flash crashes. When one underlying symbol experiences a mini flash crash, the other symbol experiences a price change in the opposite direction when arbitrageurs are active. With active arbitrage there is a faster price reversion of the mini flash crash in the symbol which originally experienced the event. The competitiveness of the background traders, rather than the arbitrageurs, influences the magnitude of the mini flash crash. The price of the index, ETF, and symbol that experiences the mini flash crash are lower than the price preceding the event. When arbitrageurs are active, the average price of the other underlying symbol is higher than the price preceding the event. Overall, the study demonstrates the effect of ETF arbitrage on welfare and how ETFs can spread volatility events through their portfolios and contribute to market-wide volatility.

This paper is organized as follows. We discuss prior work in Section 2. In Section 3 we provide an overview of an ETF's market structure and the arbitrage opportunities this structure creates. We then describe the market mechanism for our simulated ETF environment in Section 4. Section 5 presents our findings on the impacts of mini flash crashes in a market with active and inactive ETF arbitrageurs. We conclude in Section 6.

2 RELATED WORK

Most prior work uses historical data and quantitative models to analyze the impact of ETFs on the volatility of their underlying symbols. Using high frequency data, Anatolyev et al. [1] found that ETF arbitrage distorts reactions to market shocks in the US stock market. Ben-David et al. [4] constructed a quantitative model and used historical data to conclude that ETFs increase the volatility of their underlying symbols. Da and Shive [13] also analyzed historical data and found that ETFs contribute to price co-movement in the symbols in their portfolios. In contrast, Madhavan and Morillo [23] found that underlying symbol price co-movements correlate with macro-market movements, rather than the presence of the ETF. Lastly, Lynch et al. [22] found that an implementable trading strategy could generate profits using ETF arbitrage given realistic portfolio price co-movements.

There are also previous studies using historical data to examine what leads to ETF arbitrage opportunities. Box et al. [5] found that a price shock or order imbalance in an ETF's underlying portfolio typically precedes ETF arbitrage opportunities. Marshall et al. [24] found that spreads increase before ETF arbitrage opportunities.

Our agent-based market simulation builds on that of Byrd et al. [10]. Simulated market models allow other traders to strategically react given the current state of the market, which changes each iteration. It also enables us to control factors of influence, and study the effect of single factors (e.g., presence of ETF arbitrage).

Numerous other ABM studies have examined order and trade activity around volatility events. Several have replicated aspects of the Flash Crash on May, 6 2010 to determine which trading practices contributed to the price drop and recovery [27, 34, 35]. Some other ABM studies [3, 19, 20] have explored abstract mini flash crashes and price bubbles to provide insight into what might exacerbate these events. Paulin et al. [29] incorporated a dependency network to study the micro and macro impacts of flash crashes across multiple securities.

Two other prior ABM studies analyze ETF arbitrage. Mizuta [25] examined the impact of trading costs and ETF arbitrage, finding that lower trading costs and higher price volatility lead to more arbitrage opportunities, more trading volume, and stronger correlation between the index and ETF price. Torii et al. [32] model an index fund and two underlying symbols, and examine the impacts of ETF arbitrage on response to a downward volatility event in one of the symbols.

The model of Torii et al. [32] is in fact quite similar to ours, but with a few key differences. First their background traders employ technical trading strategies (i.e., consider price trends), whereas our agents are purely fundamental traders. Second, they generate volatility events by directly manipulating the fundamental value. We initiate a different type of volatility event, a mini flash crash, through an agent who quickly submits multiple large orders. In a mini flash crash, the price momentarily drifts from the asset's true value because of trading activity. Such events commonly occur in the real stock market when a trader who wants to leave a large position and starts to sell off all of their holdings [6, 17, 31]. These two methods of creating volatility events exhibit entirely different order-book dynamics, generating qualitatively different arbitrage opportunities and outputs, and differences in response after the shock. Third, the prior work does not model the primary market of an ETF, which is also an essential element in defining the performance of arbitrageurs. Despite these differences, both the study of Torii et al. [32] and our own find that when one underlying symbol experiences a downward price shock, the other underlying symbol experiences a price shock in the opposite direction.

3 ETF MARKET STRUCTURE

A special feature of ETFs in our model is that the ETF security actually trades on two markets. The *primary market* allows purchase and sale of the ETF based on the prices of its underlying securities. On the *secondary market*, the ETF trades directly like any other stock symbol. The potential divergence between prices for the ETF on these two markets is what gives rise to the arbitrage opportunity.

3.1 Primary Market

Access to the primary market is limited to select parties, commonly referred to as authorized participants (APs) [26]. In our model, APs are allowed to trade daily at the *Net Asset Value* (NAV), which is calculated based on underlying security prices at the close of the stock market. Once the primary market price (i.e., NAV) is determined, APs may submit orders (termed *basket orders*) to exchange between the ETF and underlying securities [15].

3.2 Secondary Market

The secondary market is the stock market. On the stock market, an ETF trades like any other symbol [14], with each share representing a fraction of the ETF's portfolio. When a trader submits an order, it either transacts with an order on the opposite side, or rests in limit order book until either transacted or canceled. For the ETF and any other security on the stock market, we define the *bid price* and *ask price* as the highest offers to buy and sell, respectively, on the security's limit order book.

An ETF's secondary market lowers barriers to entry compared to other portfolio funds [30], which is a principal reason for the popularity of ETFs [18]. Participants in the primary market can also trade in the secondary market.

3.3 Arbitrage between the Primary and Secondary Markets

An ETF trades at the NAV and trading price on its primary and secondary markets, respectively. Even though the NAV is calculated only once per day, it can be estimated throughout the secondary market's trading day by calculating the market index. When the bid price rises above the market index, a trader can profit by buying the underlying symbols and selling the ETF. Similarly, when the ask price falls below the market index, it is profitable to buy the ETF and sell the underlying symbols. Such arbitrage is expected to help the trading price closely track the index [2].

Whereas any secondary market participant can trade on price deviations between the ETF and market index, participants with access to both the primary and secondary market can implement arbitrage more directly [30]. Such a participant can arbitrage in the secondary market, then submit basket orders in the asset it is long in for shares of the asset it shorted.

4 MARKET MECHANISM

4.1 ABIDES

We employ a discrete event market simulation constructed on the ABIDES platform [10].[1] The platform provides a continuous double auction market with securities priced in cents, a set of typical background agents, and a kernel which drives the simulation with nanosecond resolution while permitting sparse activity patterns to be efficiently computed.

For the ETF secondary market, we use the provided ABIDES exchange agent, which operates in a manner similar to NASDAQ. The market is open from 09:30 to 16:00, lists any number of securities for trade, and provides a distinct order book mechanism for each security. The exchange accepts limit orders of any share volume, and cancellation of same, and transacts (including partial execution) those orders against a security's limit order book with a typical price-then-FIFO matching algorithm. The exchange responds to requests for market hours, last trade prices, and market depth quote requests, with depth one representing the current bid and ask.

4.2 Market Index

We define this model's *bid market index*, $\iota_{t,b}$, at time t as the sum of the bid prices of a bundle of n stocks:

$$\iota_{t,b} = \sum_{i=1}^{n} \max\left(b_t^{(i)}\right). \tag{1}$$

Where $\max\left(b_t^{(i)}\right)$ is the bid price of underlying symbol i at time t. We also define this model's *ask market index*, $\iota_{t,a}$, at time t as the sum of the ask prices of a bundle of n stocks:

$$\iota_{t,a} = \sum_{i=1}^{n} \min\left(a_t^{(i)}\right). \tag{2}$$

Where $\min\left(a_t^{(i)}\right)$ is the ask price of underlying symbol i at time t. Last, we define this model's *mid market index*, $\iota_{t,m}$, at time t as the sum of the mid prices:

$$\iota_{t,m} = \sum_{i=1}^{n} m_t^{(i)}. \tag{3}$$

Where $m_t^{(i)}$ is the midpoint between the best bid and ask prices, more formally:

$$m_t^{(i)} = \frac{1}{2}\left(\max\left(b_t^{(i)}\right) + \min\left(a_t^{(i)}\right)\right).$$

4.3 Primary Market in ABIDES

After the exchange agent stops accepting orders, the primary market receives the *close price*, $p^{(i)}$, of each underlying symbol i in the ETF portfolio, defined as the price of i's last trade on the secondary market. The primary market uses these closing prices to calculate the value of the index and uses this value as the NAV:

$$NAV = \sum_{i=1}^{n} p^{(i)}. \tag{4}$$

Then the primary market opens for basket orders. Every basket order it receives is executed at the NAV, and the agent is notified of its transaction.

4.4 Exogenous Price Time Series

We assume the existence of a mean-reverting exogenous *fundamental value series* for each underlying security, which represents an immeasurable consensus of the "true" value at any point in time. We model the fundamental as following a sparse discrete Ornstein-Uhlenbeck (OU) process [33] augmented with periodic "megashocks" of higher variance. Given a value at time 0 for the

[1]The ABIDES source code is available at `https://github.com/abides-sim/abides`.

fundamental, a value can be obtained for any subsequent time t by sampling from a normal distribution [12], with:

$$\begin{aligned} \mathbf{E}\left[Q_t^{(i)}\right] &= \mu + \left(Q_0^{(i)} - \mu\right) e^{-\gamma t} \\ \mathbf{Var}\left[Q_t^{(i)}\right] &= \frac{\sigma^2}{2\gamma}\left(1 - e^{-2\gamma t}\right), \end{aligned} \tag{5}$$

where μ is the long-term mean fundamental value, $Q_0^{(i)}$ is the time 0 fundamental value of symbol i, γ a mean reversion rate, and σ a volatility value. This fundamental process allows us to obtain $Q_t^{(i)}$ directly from $Q_0^{(i)}$ without the requirement to compute all intermediate $Q_1^{(i)} \ldots Q_{t-1}^{(i)}$. That is to say, it permits the simulation to "skip time" whenever no agents will arrive at the market.

The OU process produces a time series with a single scale of variance, which we configure to provide appropriate noise during mean reversion. To obtain a price time series more similar to a stock price over a longer window, we augment the OU process with the application of stochastic "megashocks" [9] that arrive according to a Poisson process and perturb the fundamental by a higher variance, σ_s^2, bimodal normal distribution with mean zero, and positive and negative modes substantially away from zero. This is intended to represent exogenous events which can substantially alter value perception and periodically provides much larger price moves from which the OU process can revert.

An ETF's fundamental is derived from the fundamental value of its underlying symbols:

$$Q_t^{(\mathrm{ETF})} = \sum_{i=1}^{n} Q_t^{(i)}.$$

4.5 Background Agents

We employ a population of background traders for our experiments modeled after those by Brinkman [7]. Each agent is assigned a single security of interest. They arrive according to a Poisson process with rate λ_a, cancel any outstanding orders, and place an order with equal probability to buy or sell. When arriving at the market, an agent receives a noisy observation of the exogenous price time series described above and uses this observation to update an individual Bayesian belief concerning the current fundamental value of the security. The agent then projects this belief forward to the end of the market trading period to obtain an estimated final valuation for the security.

When other agents trade, they supply relevant information about the value of the asset. Our background agents use price information from past transactions as if they are noisy observations of the fundamental at the time of the transaction. When a transaction occurs, the background agent will update its estimate of the final fundamental, as if the current transaction price is $E[Q_t^{(i)}]$, and was drawn from a Gaussian distribution around the fundamental.

An agent uses this final valuation and a private valuation to select a limit price, employing an extended form of the Zero Intelligence (ZI) strategy [16], with parameters $0 \le R_{\min} \le R_{\max}$ and $\eta \in [0, 1]$. The ZI strategy draws $R \sim U[R_{\min}, R_{\max}]$, and sets its limit price such that it would achieve surplus R if transacted at that price based on the current valuation estimate. If a market order would guarantee surplus at least ηR, the agent places that order instead of its computed limit order.

4.6 Impact Agent

We introduce the *impact agent* to induce a mini flash crash. Beginning at time τ, this agent submits a rapid series of n marketable sell orders of size q, δ seconds apart. These orders consume the buy side of the order book, causing a precipitous price drop. As subsequent traders arrive and submit orders, the price typically recovers.

4.7 ETF Arbitrage Agent

We develop four strategies for arbitrage between the ETF and its underlying symbols. Like the background agents, arbitrage agents enter the market at Poisson arrival rate λ_a. On arrival at time t, a conservative agent calculates the difference between the ETF bid price, $b_t^{(\mathrm{ETF})}$ and the ask market index $\iota_{t,a}$:

$$\Delta_{t,1} = \max\left(b_t^{(\mathrm{ETF})}\right) - \iota_{t,a}.$$

A conservative arbitrage agent must also find the difference between the bid market index, $\iota_{t,b}$, and the ETF ask price, $a_t^{(\mathrm{ETF})}$:

$$\Delta_{t,3} = \iota_{t,b} - \min\left(a_t^{(\mathrm{ETF})}\right).$$

A more aggressive agent finds the difference between the ETF mid price, $m_t^{(\mathrm{ETF})}$ and the mid market index $\iota_{t,m}$:

$$\Delta_{t,2} = m_t^{(\mathrm{ETF})} - \iota_{t,m}.$$

An aggressive agent must also find the difference between the bid market index, $\iota_{t,b}$, and the ETF ask price, $a_t^{(\mathrm{ETF})}$:

$$\begin{aligned} \Delta_{t,3} &= \iota_{t,b} - \min\left(a_t^{(\mathrm{ETF})}\right). \\ \Delta_{t,4} &= \iota_{t,m} - m_t^{(\mathrm{ETF})}. \end{aligned}$$

These agents submit only marketable orders, so their sales are at the bid and their buys at the ask.

ETF Single Asset Arbitrageur. This single asset arbitrageur trades exclusively on the ETF security. It decides when to trade based on a threshold $\varepsilon \ge 0$:

$$\begin{cases} \Delta_{t,1} > \varepsilon & \text{Sell ETF,} \\ \Delta_{t,3} > \varepsilon & \text{Buy ETF.} \end{cases}$$

An alternative, less conservative version of this arbitrageur makes decisions based on the midpoint prices ($\Delta_{t,2}$ and $\Delta_{t,4}$) instead.

ETF Multiple Asset Arbitrageur. The multiple asset arbitrageur trades both the ETF and its underlying symbols. Its trades are also triggered by a threshold $\varepsilon \ge 0$:

$$\begin{cases} \Delta_{t,1} > \varepsilon & \text{Sell ETF and buy underlying symbols,} \\ \Delta_{t,3} > \varepsilon & \text{Buy ETF and sell underlying symbols.} \end{cases}$$

There is also a variant version using midpoints.

These agents also trade on the primary market. When the primary market opens, they receive the NAV from Equation 4, then decide if they should submit basket orders by comparing the NAV and close price of the ETF, $p^{(\mathrm{ETF})}$. More formally:

$$\begin{cases} NAV - p^{(\mathrm{ETF})} > 0 & \text{ETF shares} \rightarrow \text{underlying symbol shares,} \\ p^{(\mathrm{ETF})} - NAV > 0 & \text{Underlying symbol shares} \rightarrow \text{ETF shares.} \end{cases}$$

If they submit basket orders, then they can hopefully end the day net zero. The market makers need to be the fastest agents in the system in order to be profitable when trading on so many symbols.

5 EMPIRICAL GAME-THEORETIC ANALYSIS

We introduce a set of heuristic strategies for both ETF arbitrageurs and the background agents. Using agent-based simulations and empirical game-theoretic analysis, we find the combination of strategies that agents utilize in equilibria. We determine the impact of ETF arbitrage on market welfare volatility by analyzing the market under equilibrium settings when arbitrageurs are active and inactive, as well as when arbitrageurs are active but the background agents have not recalibrated their trading strategies from when the arbitrageurs were inactive.

5.1 Market Environment Settings

We examine a variety of market environments to analyze the robustness of our results. Each environment contains one exchange agent, one ETF primary market agent, two symbols, and an ETF whose portfolio is composed of the two other symbols. The exchange agent accepts orders between 12:30 and 13:30. The ETF primary market accepts orders between 17:00 and 17:01. Following Wang et al. [37], we consider three market environments vary by market shock σ_s^2 and observation noise σ_n^2. The first consists of low shock and high observation noise (**LSHN**) with $\sigma_s^2 = 5 \times 10^4$ and $\sigma_n^2 = 10^7$. The second consists of medium shock and medium observation noise (**MSMN**) with $\sigma_s^2 = 5 \times 10^5$ and $\sigma_n^2 = 5 \times 10^6$. Lastly, the third holds high shock and low observation noise (**HSLN**) with $\sigma_s^2 = 5 \times 10^6$ and $\sigma_n^2 = 10^6$.

For each non-ETF symbol, we generate a fundamental (5), with mean $\mu = 10^5$, reversion $\gamma = 1.67 \times 10^{-13}$, and market shock $\sigma = 0$. The sparse fundamental experiences a series of megashocks throughout the trading period, and these arrive according to a Poisson distribution with $\lambda = 2.78 \times 10^{-13}$. We draw the size of these megashocks from a binomial normal distribution with means $\mu_{s,1} = 0$ and $\mu_{s,2} = 10^3$, and varying values of σ_2^2.

This market is populated with 27 background agents, where each agent is randomly assigned one symbol with equal probability to trade for the duration of each market run. Table 1 specifies the strategies of the background agents. The background agents arrive to the market according to a Poisson distribution with $\lambda_a = 10^{-11}$. These agents submit orders of size $q = 100$, but can hold a maximum number of units at any time $q_{\max} = 10^3$. When the background agents consider past transactions to update their estimate of the fundamental, they use a variance of $\sigma_p^2 = 10^3$. Lastly, the private value variance is $\sigma_{PV}^2 = 5 \times 10^6$.

We create a mini flash crash in one underlying symbol with a single impact agent. This impact agent is assigned an underlying symbol to trade on with equal probability at the beginning of each market run. It then submits $n = 5$ trades beginning at $\tau = 13{:}00$ with size $q = 100$ and $\delta = 6$ seconds between each trade.

We implement four ETF arbitrage strategies, with one conservative ETF one asset (SA-C) agent, one aggressive ETF single asset (SA-A) agent, one conservative ETF multiple asset (MA-C) agent, and one aggressive ETF multiple asset (MA-A) agent. When arbitrageurs actively trades, all arbitrageurs use a strategy with $\varepsilon = 10^3$. However, when the arbitrageurs are inactive, they only exercise a strategy where $\varepsilon = 10^{12}$. Both arbitrageurs submit orders of size $q = 100$. All ETF arbitrageurs arrive to the market according to a Poisson process with $\lambda_a = 5 \times 10^{-3}$.

5.2 EGTA Process

Empirical game-theoretic analysis (EGTA) is a method to find equilibria in games by a heuristic strategy space and simulated payoff data [39]. We utilize EGTA to find equilibria in varying market settings. EGTA iteratively searches for potential equilibria in subgames, and incrementally adds strategies to confirm or refute these potential equilibria by examining potential deviations. Previous studies utilize EGTA to examine complex environments where applying a standard analytic method is hard [8, 36–38].

We model our market as a role-symmetric game, where players are divided into five roles: background traders, ETF-SAs, and ETF-MAs. We utilize *deviation-preserving reduction* (DPR), which approximates a many-player game by aggregating a game with fewer players [40], because a game grows exponentially in players and strategies. DPR has shown to generate good approximations of the full game in multiple settings.

Our game consists of 27 background traders, one ETF-SA, and one ETF-MA, and reduces to three background traders and one of each arbitrageur when using DPR. These quantities of traders ensure that the required aggregations from DPR come out as integers. In this setting one background agent deviates to a new strategy while the other 26 background agents are further reduced to two. For a specific strategy profile, we sample between 200 and 10,000 simulation runs to reduce sampling error from stochastic market features.[2]

5.3 Impact on Market Welfare

We use EGTA to analyze the impact of ETF arbitrage on market welfare across varying market environments. Figure 1 depicts the surplus of each role when ETF arbitrageurs are active, when arbitrageurs are active but background agents have not adjusted to their presence, and when arbitrageurs are inactive in each market setting. The background agents are unadjusted to active arbitrageurs when

[2]Details of strategies employed in the Nash Equilibria identified will we presented in an online appendix (link to be provided in final version).

Table 1: Strategies employed by the background traders, where each agent chooses their desired surplus between $R_{\min}$ and $R_{\max}$, and have hyperparameter η.

Strategy	ZI_1	ZI_2	ZI_3	ZI_4	ZI_5	ZI_6
$R_{\min}$	2000	2000	2500	3000	3000	3500
$R_{\max}$	2500	3000	3000	3500	4000	4000
η	1.0	0.8	1.0	1.0	0.8	1.0

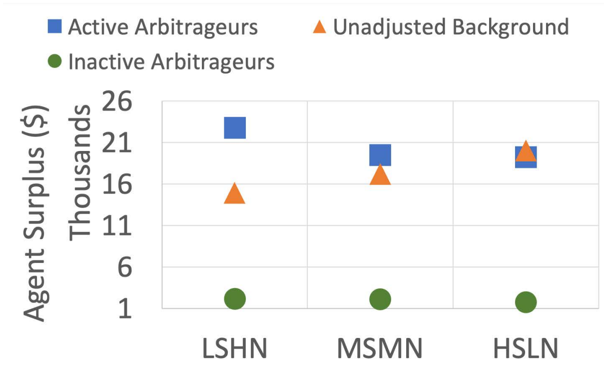

Figure 1: The average payoff over 200 simulation runs for the background agents when arbitrageurs are active and inactive in the three market environments.

they utilize the strategies optimal with inactive arbitrageurs in a setting when those agents are actually active.

Background agents are better off when arbitrageurs are active. When arbitrageurs are active they provide more marketable liquidity, particularly during a mini flash crash where there are many arbitrage opportunities. Thus, background traders have more opportunities to trade when arbitrageurs are active. These traders submit less competitive orders when arbitrageurs are active. A less competitive ZI strategy is when the agent selects a higher desired surplus range. When arbitrageurs are active, they are willing to trade with these high-margin orders, so the background traders' profit increases. The background traders realize similar payoffs when they do not adjust to arbitrage.

When the arbitrageurs are inactive, the surplus of the SA-Cs, SA-As, MA-Cs, and MA-As is always zero. ETF-MAs have a higher surplus than ETF-SAs because they have the opportunity to trade on all symbols, and the ETF primary market where ETF-MAs receive the difference in value between the ETF and index.

5.4 Impact on Symbol Price

We analyze ETF arbitrageurs' impact on market volatility, when one underlying symbol experiences a mini flash crash. To assess this impact we examine and compare the price of the underlying symbols in environments with active arbitrageurs, active arbitrageurs and unadjusted background agents, and inactive arbitrageurs. An environment with active arbitrageurs and unadjusted background agents happens when background agents do not recalibrate their strategies to consider arbitrage. In each figure we represent the price as the midpoint price over 200 simulation runs.

The price of the ETF should track the market index. Figure 2 shows the midpoint price of the ETF and Figure 3 depicts the market index of each environment. This allows us to see how arbitrage impacts the ETF trading price and market index, and what trading opportunities arbitrageurs have around the mini flash crash in an underlying symbol. The ETF trading price sees a decrease in price when arbitrageurs are active. The market index crashes a large amount in every environment at the time of the mini flash crash. In environments active and inactive arbitrage, the index drops because as the midpoint price of an underlying symbol drops, the market index or the sum of prices s_0 and s_1 also drops. When ETF arbitrageurs are present, as the market index falls below the ETF price, the arbitrageurs sell the ETF and buy the underlying symbols. As the arbitrageurs sell the ETF, the bid side of the order book reduces and the price drops. When the arbitrageurs buy the underlying symbols they absorb the ask side of the order book, which leads to the market index recovering quicker and back to a higher price when an ETF is present. Thus, arbitrage helps the ETF track the index.

In all market environments, the impact agent causes a distinct mini flash crash in the underlying symbol. This symbol experiences similar trends to the market index, where more competitive background agents creates a smaller spread and smaller price drop. Figure 4 shows each environment in a reduced time frame around the mini flash crash, and active arbitrageurs, inactive arbitrageurs, and unadjusted background agents are depicted together. The mini flash crash recovers faster with active arbitrage because the arbitrageurs submit marketable buy orders to the underlying symbol and marketable sell orders to ETF, causing the price of the index and underlying symbols to rise faster. In each environment, the price of this symbol does not recover to the price before the mini flash crash, though the price does revert to a higher level when arbitrageurs are active.

Figure 5 depicts the average price of the underlying symbol where the impact agent does not trade. When background traders are more competitive, there is a small but distinct upward price movement at the time of the mini flash crash in the other symbol. This happens because the arbitrageurs buy the ETF's underlying symbols when the index is lower than the ETF trading price, causing the price increase. In environments with arbitrage, the average price is higher than the price preceding the mini flash crash. The increase in price This implies that the presence of ETF arbitrageurs can impact the trading price of a symbol because of trading activity independent of the symbol itself.

6 CONCLUSION

We analyze an agent-based, simulated market model with a stock market and ETF primary market. We explore varying market environments that contain an ETF and two symbols which compose the ETF's portfolio. The market is also populated with numerous background agents, an impact agent which creates a mini flash crash, and ETF arbitrage agents. To determine the impact of ETF arbitrage, we examine equilibria in each environment under settings when arbitrageurs are active and inactive. The arbitrageurs are extremely profitable when they actively trade, and the surplus of the background agents increases when they are conservative and decreases when they are more competitive. We also find that this type of arbitrage may transmit a volatility event, like a mini flash crash, throughout its underlying symbols. The other underlying symbol experiences a price change in the opposite direction of the other underlying symbol at the time of the mini flash crash. The magnitude of the mini flash crash in the original symbol is impacted more by the competitiveness of the background traders than the arbitrageurs.

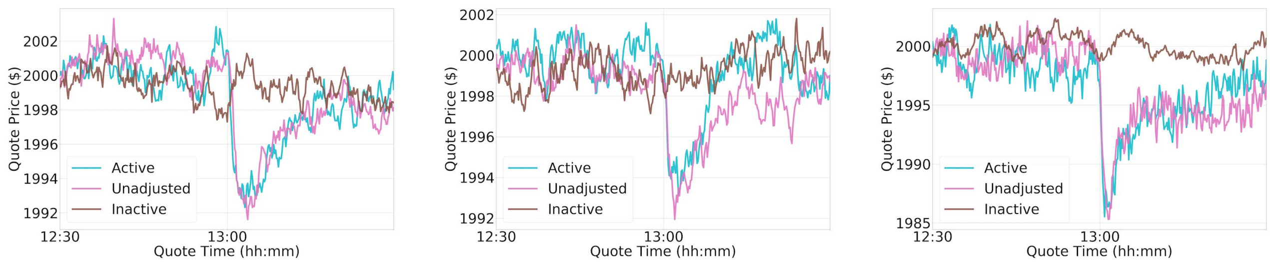

(a) LSHN, Active Arbitrage: ETF

(b) MSMN, Active Arbitrage: ETF

(c) HSLN, Active Arbitrage: ETF

Figure 2: Average price time series over 200 simulation runs of the ETF. The ETF in a market is meant to track the market index. The ETF is only present in three of the market environments.

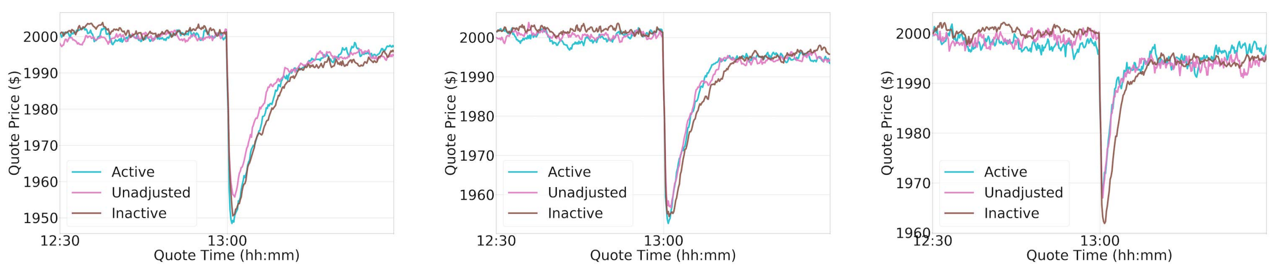

(a) LSHN, Active Arbitrage: Index

(b) MSMN, Active Arbitrage: Index

(c) HSLN, Active Arbitrage: Index

Figure 3: Average time series over 200 simulation runs of the market index. This index is composed of two symbols. An underlying symbol experiences a mini flash crash at 13:00.

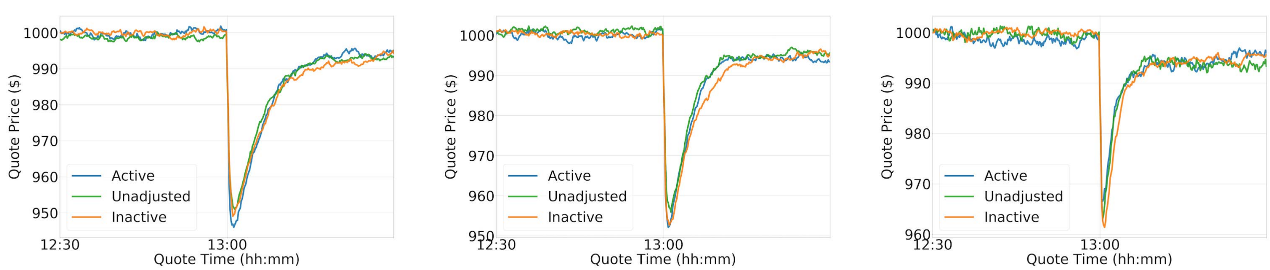

(a) LSHN: s_0

(b) MSMN: s_0

(c) HSLN: s_0

Figure 4: Average price time series over 200 simulation runs of underlying symbol, s_0, that is randomly selected to experience a mini flash crash at 13:00 through an impact agent submitting a series of large trades. This shows active arbitrageurs, inactive arbitrageurs, and unadjusted background traders on the same plot.

Most real-world ETFs represent portfolios with many symbols. Our study models an ETF with only two symbols, in order to focus the effect on a single relationship. We believe this captures qualitative properties of realistic ETFs as well, but it could be that further insights would be revealed by extending the model to cover more symbols. This could be readily incorporated in our ABM, with linear impact on computational cost of simulation, though perhaps imposing somewhat more complexity on the arbitrage strategy and the analysis of results. An ETF with a larger portfolio could cause arbitrage to have a larger impact on the ETF than the underlying symbols, so arbitrage may help the ETF track the index, but introduce less volatility to the underlying symbols.

Another potential limitation is in the space of background trading strategies. We focus exclusively on background agents that

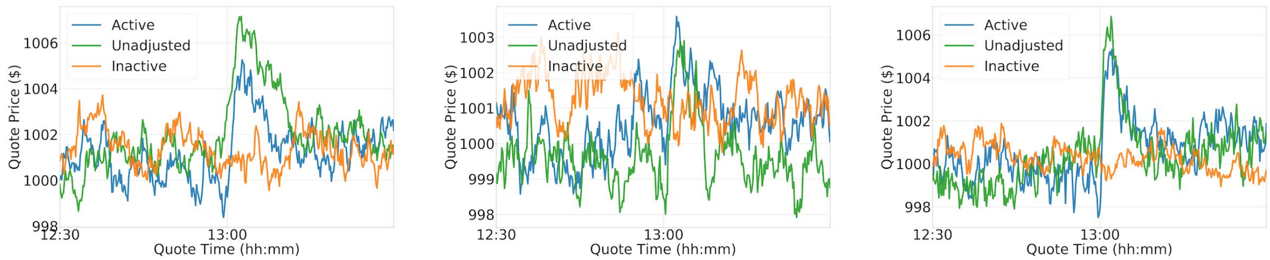

(a) LSHN, Active Arbitrage: s_1 **(b) MSMN, Active Arbitrage:** s_1 **(c) HSLN, Active Arbitrage:** s_1

Figure 5: Average price time series over 200 simulation runs of underlying symbol without a mini flash crash. An impact agent does not trade on this symbol.

use a ZI strategy which consider previous transactions in their estimate of an asset's fundamental value. It could be beneficial to utilize more trading strategies dependent on the order book and price movement, such as market makers or trend followers, because these agents might exacerbate mini flash crashes.

This paper provides insight into the impact of ETFs on market welfare, market volatility, and stock valuation. Previous studies have used historical data to examine associations between ETF activity and price volatility in its underlying symbols. With agent-based simulation we are able to examine causality through a direct A/B test in market environments with and without active ETF arbitrage. We find other agents are better off with arbitrage if they are more conservative, but arbitrage reduces their profits if they trade more competitively. An implication is that inclusion in an index ETF may impact the pricing of a stock without any actual change in the stock's fundamental value, and solely due to trading activity independent of the stock itself.

ACKNOWLEDGEMENTS

This material is based on research supported in part by the National Science Foundation under Grant no. 1741026, and by two JPMorgan AI Research Fellowships.

This paper was prepared for informational purposes in part by the Artificial Intelligence Research group of JPMorgan Chase & Coȧnd its affiliates ("JP Morgan"), and is not a product of the Research Department of JP Morgan. JP Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

REFERENCES

[1] Stanislav Anatolyev, Sergei Seleznev, and Veronika Selezneva. 2020. Does Index Arbitrage Distort the Market Reaction to Shocks? *SSRN Electronic Journal* (2020), 1–50.
[2] Rochelle Antoniewicz and Jane Heinrichs. 2014. Understanding Exchange-Traded Funds: How ETFs Work. *SSRN Electronic Journal* (2014), 1–40.
[3] Erhan Bayraktar and Alexander Munk. 2017. Mini-flash crashes, model risk, and optimal execution. *SSRN Electronic Journal* (2017), 1–48.
[4] Itzhak Ben-David, Francesco Franzoni, and Rabih Moussawi. 2015. Do ETFs Increase Volatility? Securities and Exchange Commission Comments.
[5] Tavis Box, Ryan Davis, Richard B. Evans, and Andrew A. Lynch. 2019. Intraday Arbitrage Between ETFs and their Underlying Portfolios. *SSRN Electronic Journal* (2019), 1–73.
[6] Tobias Braun, Jonas A. Fiegen, Daniel C. Wagner, Sebastian M. Krause, and Thomas Guhr. 2018. Impact and recovery process of mini flash crashes: An empirical study. *PLOS One* 13, 5 (2018), e0196920.
[7] Erik Brinkman. 2018. *Understanding financial market behavior through empirical game-theoretic analysis.* Ph.D. Dissertation. University of Michigan.
[8] Erik Brinkman and Michael P. Wellman. 2017. Empirical mechanism design for optimizing clearing interval in frequent call markets. In *18th ACM Conference on Economics and Computation.* 205–221.
[9] David Byrd. 2019. Explaining Agent-Based Financial Market Simulation. *eprint arXiv:1909.11650* (2019), 1–10.
[10] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. 2020. ABIDES: Towards High-Fidelity Multi-Agent Market Simulation. In *ACM SIGSIM Conference on Principles of Advanced Discrete Simulation* (Miami) *(SIGSIM-PADS '20).* 11—-22. https://doi.org/10.1145/3384441.3395986
[11] Ben Carlson. 2019. Spike, Swoon, Repeat: What to Make of a Volatile Stock Market—That Has Basically Gone Nowhere for a Year and a Half. Fortune.
[12] Tanmoy Chakraborty and Michael Kearns. 2011. Market making and mean reversion. In *12th ACM Conference on Electronic Commerce.* 307–314.
[13] Zhi Da and Sophie Shive. 2017. Exchange traded funds and asset return correlations. *European Financial Management* 24, 1 (2017), 136–168.
[14] Robert Engle and Debojyoti Sarkar. 2006. Premiums-Discounts and Exchange Traded Funds. *Journal of Derivatives* 13, 4 (2006), 27–45.
[15] Gary L. Gastineau. 2004. The Benchmark Index ETF Performance Problem. *Journal of Portfolio Management* 30, 2 (2004), 96–103.
[16] Dhananjay K. Gode and Shyam Sunder. 1993. Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy* 101, 1 (1993), 119–137.
[17] Anton Golub, John Keane, and Ser-Huang Poon. 2017. High frequency trading and mini flash crashes. *SSRN Electronic Journal* (2017), 1–22.
[18] Ben Golub, Barbara Novick, Ananth Madhavan, Ira Shapiro, Kristen Walters, and Maurizio Ferconi. 2013. Exchange Traded Products: Overview, benefits, and myths. Black Rock.
[19] Sandrine Jacob Leal and Mauro Napoletano. 2019. Market stability vs. market resilience: Regulatory policies experiments in an agent-based model with low- and high-frequency trading. *Journal of Economic Behavior and Organization* 157 (2019), 15–41.
[20] Blake LeBaron, Brian Arthur, and Richard Palmer. 1999. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 9–10 (1999), 1487–1516.
[21] Yun Li. 2019. Greenspan says the stock market will determine whether there's a recession. CNBC News.
[22] Hailey Lynch, Sébastien Page, Robert A. Panariello, James A. Tzitzouris Jr., and David Giroux. 2019. The Revenge of the Stock Pickers. *Financial Analysts Journal* 75, 2 (2019), 34–43.
[23] Ananth Madhavan and Daniel Morillo. 2018. The Impact of Flows into Exchange-Traded Funds: Volumes and Correlations. *Journal of Portfolio Management* 44, 7 (2018), 96–107.

[24] Ben Marshall, Nhut H. Nguyen, and Nuttawat Visaltanachoti. 2013. ETF arbitrage: Intraday evidence. *Journal of Banking and Finance* 37, 9 (2013), 3486–3498.

[25] Takanobu Mizuta. 2019. Agent-Based Model of Liquidity and Arbitrage Cost Between ETF and Stocks. In *2019 8th International Congress on Advanced Applied Informatics (IIAI-AAI)*. 685–688.

[26] Barbara Novick, Ananth Madhavan, Samara Cohen, Sal Samandar, Sander Van Nugteren, and Alexis Rosenblum. 2017. A Primer on ETF Primary Trading and the Role of Authorized Participants. BlackRock.

[27] Mark Paddrik, Roy Hayes, Andrew Todd, Steve Yang, Peter Beling, and William Scherer. 2012. An agent based model of the E-Mini S P 500 applied to flash crash analysis. In *2012 IEEE Conference on Computational Intelligence for Financial Engineering Economics (CIFEr)*. 1–8.

[28] Marco Pagano, Antonio Sánchez Serrano, and Josef Zechner. 2019. Can ETFs contribute to systemic risk? *Reports of the Advisory Scientific Committee* 9 (2019), 1–40.

[29] James Paulin, Anisoara Calinescu, and Michael Wooldridge. 2018. Agent-Based Modeling for Complex Financial Systems. *IEEE Intelligent Systems* 33, 2 (2018), 74–82.

[30] James Poterba and John Shoven. 2002. Exchange Traded Funds: A new investment option for taxable investors. *American Economic Review* 92, 2 (2002), 422–427.

[31] Megan Shearer. 2020. The Phases and Catalysts of Mini Flash Crashes. *SSRN Electronic Journal* (2020), 1–44.

[32] Takuma Torii, Kiyoshi Izumi, and Kenta Yamada. 2016. Shock transfer by arbitrage trading: Analysis using multi-asset artificial market. *Evolutionary and Institutional Economics Review* 12, 2 (2016), 395–412.

[33] George E. Uhlenbeck and Leonard S. Ornstein. 1930. On the theory of the Brownian motion. *Physical Review* 36, 5 (1930), 823.

[34] Iryna Veryzhenko and Nathalie Oriol. 2012. Post Flash Crash Recovery : An Agent-based Analysis. In *Proceedings for the 8th International Conference on Agents and Artificial*, Vol. 17. 190–197.

[35] Tommi A. Vuorenmaa and Liang Wang. 2014. An Agent-Based Model of the Flash Crash of May 6, 2010, with Policy Implications. *SSRN Electronic Journal* (2014), 1–39.

[36] Elaine Wah, Sébastien Lahaie, and David M. Pennock. 2016. An empirical game-theoretic analysis of price discovery in prediction markets. In *26th International Joint Conference on Artificial Intelligence*. 510–516.

[37] Xintong Wang, Yevgeniy Vorobeychik, and Michael P. Wellman. 2018. A Cloaking Mechanism to Mitigate Market Manipulation. In *27th International Joint Conference on Artificial Intelligence*. 541–547.

[38] Xintong Wang and Michael P Wellman. 2017. Spoofing the limit order book: An agent-based model. In *16th Conference on Autonomous Agents and MultiAgent Systems*. 651–659.

[39] Michael P. Wellman. 2016. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems* 30 (2016), 1175–1189.

[40] Bryce Wiedenbeck and Michael P. Wellman. 2012. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*. 931–938.