PERUVIAN ECONOMIC ASSOCIATION

The Impact of ETFs on Asset Markets: Experimental Evidence

John Duffy

Jean Paul Rabanal

Olga A. Rud

Working Paper No. 154, December 2019

The Impact of ETFs on Asset Markets: Experimental Evidence

John Duffy * Jean Paul Rabanal† Olga A. Rud‡

December 6, 2019

Abstract

We examine how exchange traded funds (ETFs) affect asset pricing, volatility and trade volume in a laboratory asset market. We consider markets with zero or negative correlations in asset returns and the presence or absence of composite ETF assets. We find that when the returns on assets are negatively correlated, the presence of an ETF asset reduces mispricing and price volatility without decreasing trading volume. In the case where returns have zero correlation, the ETF asset has no impact. Thus, our findings suggest that ETFs do not harm, and may in fact improve, price discovery and liquidity in asset markets.

JEL Codes: G11, G12, G14, C92.
Keywords: ETF, asset pricing, volatility, volume, experimental finance.

*Department of Economics, University of California, Irvine (Corresponding Author).
†Department of Banking and Finance, Monash University.
‡Department of Economics, RMIT University.

1 Introduction

Exchange traded funds, or ETFs, currently represent 35 percent (Ben-David et al., 2018) of all equity trades in the United States. Their meteoric rise in popularity as an investment vehicle has democratized investing, providing retail investors with access to products which were once available only to institutional investors 20 years ago (Hill, 2016). ETFs are investment products which aim to track a particular index, and may be one of the most important financial innovations in recent history (Lettau and Madhavan, 2018). The advantages of these investment products are easy to appreciate: (i) they help diversify market risk by allowing investors to hold a bundle of assets (the index or ETF), (ii) they have lower associated management fees, and (iii) they are traded continuously on an exchange, making them more liquid than mutual funds. In addition, ETFs are appealing to institutional investors who are looking to turn a profit by engaging in arbitrage.

However, the appeal and ubiquity of ETFs might have a destabilizing role for markets if they also attract speculators who add noise to the price discovery process, or who generate excessive volatility in asset prices. According to Bogle (2016), "Most of today's 1,800 ETFs are less diversified, carry greater risk, and are used largely for rapid-fire trading —speculation, pure and simple." In this paper, we provide evidence relevant to the debate about the impact of ETFs in asset markets by conducting a laboratory experiment where subjects trade assets either in the presence or in the absence of an ETF asset, so that we can understand the impact of ETFs on asset pricing, volatility and trading volume. In the experimental literature, there are a some studies that examine trading in multiple assets, but there are no studies we are aware of that address the role of composite, tradeable assets.

Our laboratory market builds on the seminal design of Smith et al. (1988) (hereafter SSW) and extends it to two assets, A and B. The asset returns are either (i) perfectly negatively correlated, as in our $2N$ treatment, or have (ii) zero correlation, as in our $2Z$ treatment. The dividend process for asset A has an expected value of zero in every period, while the dividend process for asset B has a structural break: for the first t periods, it follows the same dividend process as asset A so that its expected value is zero, and beginning in period $t+1$, the expected value of the dividend process jumps to one until the terminal period T. These two dividend processes generate either flat or declining paths for the fundamental values of the assets, which correspond to the two most commonly studied fundamental paths in the experimental asset pricing literature.

Our second design innovation concerns the *number* of assets in the market, which may be two or three. The two asset markets involve trading of assets A and B. In the three asset market, a composite ETF asset, referred to as asset C, also exists and can be traded. This ETF asset is a claim to one unit of asset A and one unit of asset B and so its fundamental value is the equal-weighted fundamental values of assets A and B. In addition to determining and seeing the price of ETF asset C, market participants also learn its net asset value (NAV), the sum of the market price of one unit of asset A and one unit of asset B in each period to facilitate arbitrage.

To preview our results, we find that in the negative correlation treatment, price volatility and mispricing are significantly reduced with the introduction of the ETF asset. Thus, the ETF asset provides an important benchmark to help traders properly price the underlying assets. Indeed, we find that ETF prices are close to the NAV and get even closer with experience in the negative correlation treatment. By contrast, in the zero correlation environment, we do not find significant differences in price volatility and mispricing between markets with and without the ETF asset, though the ETF asset continues to closely follow the NAV. In our design, the ETF asset represents 50 percent of total assets,[1] and yet we do not find any effect from the introduction of the ETF asset on trading volume in the underlying assets in either correlation case. In fact, traders actively participate in markets for *all* assets across all four of our experimental treatments.

Our motivation for introducing assets with perfectly negatively correlated returns is to highlight the insurance value of the ETF asset to subjects, since holding the ETF asset in the negative correlation case provides perfect insurance against aggregate risk. One possible interpretation of this perfectly negative correlated case is that investors hold a portfolio which consists of two assets with perfectly correlated returns, where the investor takes a long position in one asset and a short position in the other asset.

In addition to the three advantages of ETFs mentioned earlier, ETFs also have some important institutional characteristics that should be considered. For example, the majority of ETFs are composed of equities which seek to track large cap indices, sector indices or other indices.[2] These investment products are then traded in two separate markets: (i) the primary market, where Authorized Participants (APs), or

[1] The supply of assets in our laboratory markets is fixed.

[2] According to the *Wall Street Journal*, bond ETFs, have passed 1 trillion in assets in July 2019, a market that did not exist 20 years ago (https://www.wsj.com/articles/bond-exchange-traded-funds-pass-1-trillion-in-assets-11561986396).

large financial institutions, issue and redeem ETF shares, and (ii) the secondary market, where ETF shares are traded by the public. APs help ensure that the ETF price closely tracks the basket of underlying securities, as measured by the Net Asset Value (NAV), by taking advantage of any arising arbitrage opportunities. Our experimental design simplifies a number of these institutional characteristics since our goal is to isolate the effect of ETFs on asset prices, volatility and volume. Consequently, in this paper we focus on the secondary market, which represents about 90 percent of daily ETF activity (ICI, 2019),[3] and we provide asset market participants with information on the NAV of the ETF in every period.

While there is an empirical literature exploring the impact of ETFs on financial markets that we discuss in the next section, we resort to a laboratory experiment for several reasons. First, the laboratory provides us with control over the fundamental value of the assets under study so that we can accurately assess the extent to which agents are able to correctly price individual assets as well as composite assets such as ETFs. Second, we consider laboratory environments with and without ETF assets in order to clearly identify the impact of ETF assets on asset prices, volatility and trading volume. In the field, ETFs are now ubiquitous in all markets and so it would be more difficult to identify their impact. Finally, we can change other variables, such as the correlation in asset returns that might matter for the impact of ETFs on financial markets.

2 Related literature

The existing literature on the effects of ETFs on price discovery, volatility and liquidity of the underlying assets is mixed.[4] There is some evidence that ETFs can improve intraday price discovery of securities (Hasbrouck, 2003, Yu, 2005; Chen and Strother, 2008; Fang and Sanger, 2011; Ivanov et al., 2013), particularly if the individual securities are less liquid than the ETF. The improvement in price discovery comes from faster response time to new information on earnings (especially the macro-related component) and the subsequent trading of the lower cost ETFs. The fluctuations in ETF prices can help guide the prices of the underlying securities to integrate new information.

[3]For a detailed overview of ETFs, see, e.g., Lettau and Madhavan (2018).

[4]At the macro level, Converse et al. (2018) find that total cross-border equity flows and prices are significantly more sensitive to global financial conditions in countries where ETFs hold a larger share of financial assets.

Hasbrouck (2003) provides some empirical evidence for this phenomenon using index futures. Similar results are also found by Glosten et al. (2016) who document how ETFs positively affect informational efficiency at the individual stock level, in particular with respect to information on earnings. Huang et al. (2018) show similar positive effects for industry level ETFs, while Bhojraj et al. (2018) find a positive effect on efficiency for sector level funds, and a negative a effect for non-sector level funds. Agapova and Volkov (n.d.) determine that when corporate bonds are included in ETFs, the returns are less volatile than for bonds which are not included. Lastly, there is also some evidence that market liquidity of the underlying assets improves with the introduction of ETFs (Hegde and McDermott, 2004; Nam, 2017). On the other hand, Hamm (2014) finds that when lower quality individual stocks are included in ETFs, the market can become less liquid as uninformed investors move away from investing in these stocks in favor of the ETF, where asymmetric information problems are mitigated. Since evidence suggests that ETFs affect the prices of underlying assets, ETFs may also lead to price volatility and affect market efficiency. For example, a positive change in an asset's fundamental value, perhaps due to favorable news, should lead to an upward price adjustment. However, if the movement is instead driven by noisy ETF traders, then one would expect a price reversal in the near future, thus increasing the volatility of the underlying assets.

Arbitrage can also transmit pressure to the underlying assets as mispricing of the ETFs is passed through to the basket of individual securities. This can occur due to (i) trades by uninformed investors, and/or (ii) traders who participate in long-short strategies involving other mispriced securities. Ben-David et al. (2018) find that ETF arbitrage activity increases non-fundamental volatility of underlying stocks due to noisy traders. Madhavan and Sobczyk (2016) decompose the price of the ETF relative to its NAV (ETF premium) into two components, one corresponding to price discovery and the other to transitory liquidity. They find that an ETF-led price discovery following a change in fundamentals can lead to excess volatility when the composite assets are illiquid. Baltussen et al. (2019) also provide evidence of price reversals and noisy shocks to index products.

ETF assets also have some features in common with derivative assets such as futures which track an index. However, unlike futures, ETFs do not have a maturity date, which can erode performance for investors with broader horizons, and ETFs are not derivative assets since they can be directly traded. Noussair and Tucker (2006) studied the impact of futures in an SSW laboratory environment with a single asset and found that a complete set of futures markets, where one matures every period, can correct spot market price bubbles. They also observe widespread mispricing in the futures markets. In a follow-up study, Noussair et al. (2016) constrain the number of futures contracts to one and find that a longer maturity can help reduce mispricing, despite an increase in price volatility observed in some sessions. Porter and Smith (1995) find a very modest mispricing correction in the spot market when the single contract matures half-way through the life of the asset.

The dividend process we use in our experimental design is based on the previous experimental asset pricing literature. Similar dividend processes have been studied for a single asset in Kirchler et al. (2012) and Breaban and Noussair (2015). Kirchler et al. (2012) find that a constant fundamental value, (i.e., the case where the expected dividend value is 0 and there is some final, positive terminal value), facilitates price discovery, while Breaban and Noussair (2015) show that a constant fundamental value followed by a decreasing trend can also reduce mispricing relative to an environment with a constant fundamental value followed by an increasing trend.[5] A two-asset SSW market also appears in a recent study by Charness and Neugebauer (2019). They find that the law of one price holds when asset returns have a perfectly positive correlation, and fails to hold when the correlation is zero. The structure of dividends in Charness and Neugebauer (2019) follows the classical SSW environment with decreasing fundamentals, which tends to generate larger price deviations relative to the fundamental values. As pointed out by Kirchler et al. (2012), if the structure of the fundamental process is rather flat, then one should expect prices closer to fundamental values, and convergence to the law of one price.

3 The environment

Our experimental design builds upon the seminal work of SSW where market participants trade an asset with a common dividend process and a decreasing fundamental value for a finite number of periods. We extend the SSW environment to two and three asset markets, where assets are subject to different dividend processes, utilizing a 2×2 experimental design. The first treatment variable pertains to the number of

[5]The mispricing in bearish markets disappears in Marquardt et al. (2019) where the earnings are subject to a trend shock, and there are no interim dividends.

assets traded in the market. A market can have either (i) two assets A and B, or (ii) three assets, which includes an ETF asset C, a composite asset of A and B using equal weights. The second treatment variable is the correlation in the dividends earned by the two assets, A and B, which can be either perfectly negative N or zero Z.

The fundamental value of an asset, assuming no discounting, is the expected dividend payments remaining over the life of the asset in periods $T-t+1$, and the asset's terminal value TV, such that $FV_{j,t} = \sum_{s=t}^{T} E_s[D_{j,s}] + TV_j$, where j refers to asset type. We assume that $T = 15$ and specify the fundamental value of each asset as

$$\begin{aligned} FV_{A,t} &= 10 \\ FV_{B,t} &= \begin{cases} 18 & \text{for } t \leq 8 \\ (T-t+1)+10 & \text{for } t > 8 \end{cases} \\ FV_{C,t} &= \begin{cases} 28 & \text{for } t \leq 8 \\ (T-t+1)+20 & \text{for } t > 8. \end{cases} \end{aligned} \tag{1}$$

All market participants are endowed with a bundle of cash and a portfolio of assets, such that the distribution of wealth across players is equal. We specify the initial allocation of assets $\Omega = \{A, B, C\}$, and cash for all players in Table 1. In each session, players can participate in up to three separate call markets, each consisting of $T = 15$ trading periods. In every period $t = \{1, \ldots, T\}$ asset A pays a dividend $D_A \in \{-1, 1\}$, which is decided by a fair coin flip such that the expected dividend $E[D_A] = 0$. Following period T, asset A pays a terminal value $TV_A = 10$. Thus, the fundamental value FV_A is constant and equal to 10.[6] In periods $t = \{1, \ldots, t^*\}$, asset B follows the same dividend structure as asset A, such that $D_B \in \{-1, 1\}$ and $E[D_B] = 0$. For periods $t = \{t^* + 1, \ldots, T\}$, there is a structural break so that $D_B \in \{0, 2\}$, which is decided by a fair coin flip such that the expected dividend $E[D_B] = 1$. Following period T, asset B pays a terminal value $TV_B = 10$. Hence, asset B has a constant fundamental value until t^*, and a decreasing trend thereafter.

In the zero correlation environment the realizations of D_A and D_B are drawn independently of each other, and the independence of these realizations is known. By

[6] There is a small probability (equal to 0.059) that $FV_A < 0$ if the realized dividends for asset A are negative for at least 11 of the total 15 periods, given that the terminal value for A is equal to 10. However, in expectation, the value of the dividends is zero and therefore this should not be an issue for forward-looking agents.

contrast, in the perfectly negative correlation environment the realizations of D_B are exactly opposite to the realizations of D_A. That is, in a negative correlation environment in periods $t = \{1, \dots, t^*\}$, when $D_A = -1$, then $D_B = 1$, and when $D_A = 1$, then $D_B = -1$. In periods $t = \{t^* + 1, \dots, T\}$, when $D_A = -1$, then $D_B = 2$, and when $D_A = 1$, then $D_B = 0$. The exact timing of the structural break and the perfect negative correlation in dividends is known. In the three asset environment, we introduce an ETF asset to the two asset market, which we call asset C. This ETF asset is a composite asset composed of one unit of asset A and one unit of asset B.

Table 1: Endowment bundles across subjects in all treatments

		3 assets (3N, 3Z)			2 assets (2N, 2Z)	
Subjects	Cash	A	B	C	A	B
1-3	444	8	2	0	8	2
4-6	396	2	8	0	2	8
7-9	280	0	0	10	10	10

Note: we assume an initial total cash-asset ratio of three, and an equal distribution of wealth across participants. The initial fundamental value for assets {A,B,C} is {10,18,28}.

Dividend earnings from all assets held by a player in each period are stored in a separate account and are converted into cash earnings at the end of the terminal period T. The number of shares available for trade at any given time is fixed such that $s^A, s^B, s^C = \{30, 30, 30\}$.[7] Since we assume that one share of the ETF asset C, is composed of one share of asset A and one share of asset B, the net asset value (NAV) of the ETF asset C is

$$NAV_C := \frac{s_A^C \times p_A + s_B^C \times p_B}{s_C} = p_A + p_B. \tag{2}$$

Note that the dividends received for holding one share of the ETF asset C follow

$$D_C := \frac{s_A^C \times D_A + s_B^C \times D_B}{s_C} = D_A + D_B. \tag{3}$$

Therefore, for $t = \{1, \dots, t^*\}$ the ETF asset C pays $\{-2,0,2\}$ with probability $\{1/4, 1/2, 1/4\}$ when the correlation between the underlying assets is zero, or zero when the correla-

[7] Our ETF environment assumes a fixed supply of assets and is therefore different from an open-ended fund, where shares are created and redeemed in response to market forces. However, given that this is the first paper to study the impact of ETFs on market behavior, we assume a simple environment, focusing on the secondary market.

tion between the underlying assets is negative, with certainty. For $t = \{t^* + 1, \ldots, T\}$, ETF asset C pays either $\{-1,1,3\}$ with probability $\{1/4, 1/2, 1/4\}$ when the correlation between the underlying assets is zero and one when the correlation between the underlying assets is negative, with certainty.

We assume an initial total cash-asset ratio endowment of three in each experimental session. Moreover, each subject enters the market with the same wealth, though we vary the composition of wealth across subjects. Specifically, subjects 1-3 are endowed with eight shares of asset A and two shares of asset B, subjects 4-6 are endowed with two shares of asset A and eight shares of asset B, and subjects 7-9 are endowed with ten shares of the ETF asset C. In a two-asset market, where only assets A and B are traded, the distribution of shares is the same for the subjects 1-6 and subjects 7-9 now receive ten shares each of assets A and B, which is equivalent to ten shares of the composite ETF asset C. It is important to highlight that the only difference across treatments is the presence of the ETF asset. The distribution of assets and cash is similar across treatments.

3.1 Market format

As a market clearing mechanism, we employ a call market which produces a single uniform price for each asset traded in time t.[8] In a three asset environment, subjects can trade in up to three separate call markets simultaneously, with one market assigned to each asset. Similarly, in the two asset environment, subjects can trade in up to two separate call markets, with one market assigned to each asset.

In any given market, after all bids and asks are submitted, our computer program sorted bids in a descending order and asks in an ascending order to derive market demand and market supply curves to clear each asset market at a uniform price. Each period t, market participants can submit one buy order and/or sell order for each of the three asset markets. They can also choose not to participate in one, two or all three markets. A complete buy order specifies a single bid price and a number of units desired at that price. Similarly, a complete sell order includes a single ask price and a number of units for sale at that price. In our environment, participants can only

[8] We employ a call market for three reasons: (i) it requires less effort on the part of subjects tasked with following the prices of multiple assets, (ii) subjects are required to interact in each market, even if they wish to place an order for zero units, and (iii) there is a large tradition of using call markets as market institutions. The convergence of asset prices to fundamental values in call markets has been demonstrated by Smith et al. (1982), Cason and Friedman (1997) and Plott and Pogorelskiy (2017).

sell assets currently in their portfolio– that is, short-selling is not allowed. Moreover, there is no borrowing, as each participant can only place bid orders which satisfy their budget constraint (current cash available).

At the end of each period, subjects learn the market prices of either two or three assets (p_A, p_B and p_C), depending on the assigned treatment and the net asset value of the ETF asset C or $NAV_C = p_A + p_B$, if appropriate. By design, the NAV_C can depart from p_C, providing a potential for arbitrage opportunities.

3.2 Hypotheses

Assuming no limits to arbitrage, ETF prices should not deviate from their fundamental net asset values (NAVs), which depend on the market prices of the assets underlying the ETF. It follows that there should be no difference in various market measures, such as mispricing, price volatility and trade volume between our three asset market with ETFs and the comparable two asset market without ETFs. We will examine these measures in detail in section 5. We further explore how different risk profiles arising from the two different correlations between asset returns might affect market prices, assuming risk-averse subjects.

Hypothesis 1: *Price volatility in a market with ETF assets is the same as in a market without ETF assets.*

According to Ben-David et al. (2018) ETF arbitrage activity can increase non-fundamental volatility of the underlying stocks due to noisy traders. However, arbitrage also indicates an active market, which should help with price discovery and therefore reduce the price volatility of the underlying assets.

Hypothesis 2: *A market with ETF assets will exhibit the same level of mispricing as a market without ETF assets.*

The evidence is mixed with regards to how ETFs affect market prices and price discovery. As we noted earlier, there are some studies suggesting an improvement in efficiency, though this is sensitive to scope (sector, industry, etc.), and asset types (greater efficiency gains for bond ETFs, where underlying assets are less liquid). Lastly, experimental evidence suggests that depending on asset return correlations, we should see some evidence of correct relative pricing.

Hypothesis 3: *The volume of trade in markets with ETF assets is the same as in a markets without ETF assets.*

We expect this outcome for two reasons: (i) arbitrage activity will ensure an active market, and (ii) although ETFs may be thought of as passive instruments, they are actually actively traded.

Hypothesis 4: *The price of the ETF follows the NAV.*

If the price of the ETF does not follow NAV, then arbitrage between underlying securities and the composite assets will exist.

Hypothesis 5: *Asset prices are not affected by the correlation in the returns of the underlying assets.*

According to Charness and Neugebauer (2019), the correlation in asset returns matters for the pricing of assets. Therefore, we expect that when the correlation in asset returns is perfectly negative, asset prices could differ from those observed under zero correlation in returns. However, to stay consistent with our no limits to arbitrage assumption, we test the null hypothesis of no effect of correlation in returns on asset pricing.

4 Laboratory Procedures

The experiment was conducted at the Experimental Social Science Laboratory (ESSL) of the University of California, Irvine. Participants included undergraduate students from all fields who were recruited online using the SONA systems software. Subjects were assigned to participate in one of the four treatments: $\{\mathit{2Z}, \mathit{3Z}, \mathit{2N}, \mathit{3N}\}$. In each session, subjects were asked to complete a quiz after reading the instructions. Upon completion, the experimenter checked the answers and if a subject made any incorrect responses, the correct answers were given and explained privately to the individual. Before the market opened for trading, the subjects also completed a risk-elicitation task

(Crosetto and Filippin, 2013, implemented in oTree by Holzmeister and Pfurtscheller, 2016).

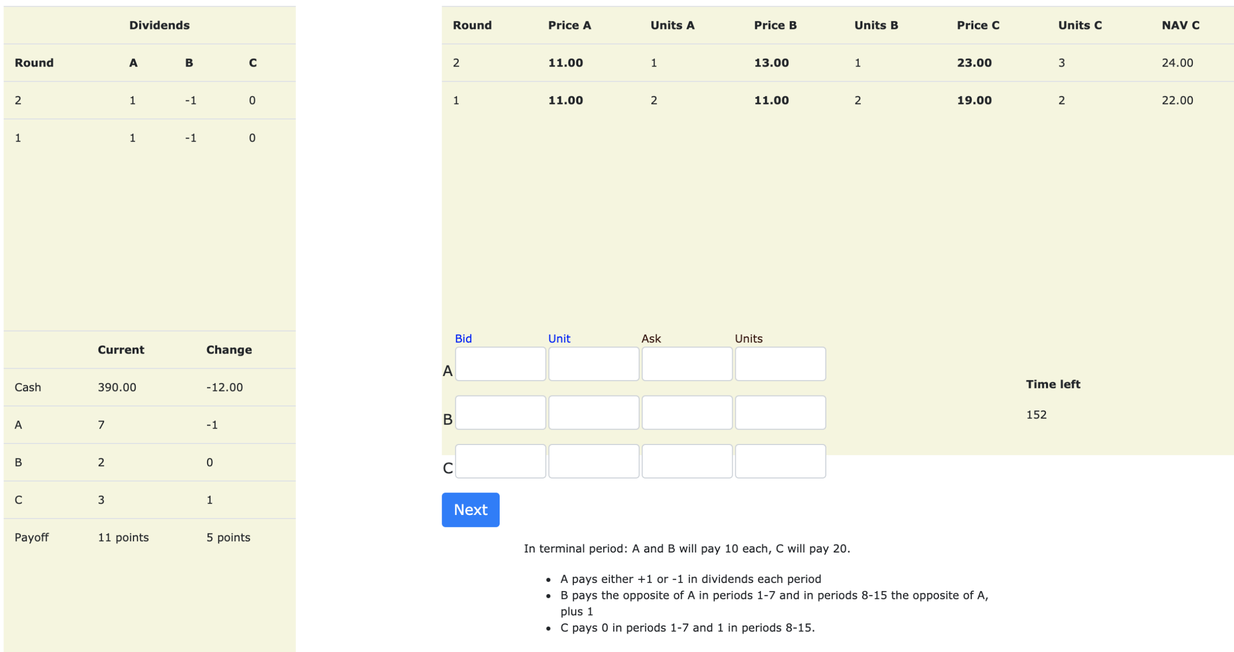

Figure 1: User interface in the $3N$ treatment.

Each session consisted of two 15 period markets. In each period, each subject had the option to input buy and/or sell orders, subject to the constraints that their buy orders did not exceed their endowment and their sell orders were for assets currently in their possession (i.e., no borrowing or short selling was allowed). Once a subject had decided on their order, they had to confirm their order by clicking on a button. For an example of the user-interface, designed in oTree (Chen et al., 2016), please refer to Figure 1.[9] Among the information provided to the subjects each period was market data (prices and units transacted) from all previous periods, and the NAV of

[9]On the decision screen shown in Figure 1, subjects had to fill in all twelve boxes (0 was always an option). The input boxes in the first column were for the per unit bids for assets A, B and C; the boxes in the second column were the number of units to buy of assets A, B, and C; the boxes in the third column were the per unit asks for assets A, B and C, and the boxes in the fourth column were the number of units to sell of assets A, B, and C. While there was a timer counting down 180 seconds, there was no binding time constraint; market prices were not determined until all 9 subjects had confirmed their orders (completed all 12 boxes and submitted the order).

the ETF asset C (in treatments where there was such an asset). In the case where current price information for either asset A or asset B was not available (because a market clearing price could not be determined), the NAV was computed using the most recently available market price.

Any dividends accrued by subjects over the course of the market were put into a separate account which was converted into cash at the end of the session.[10] The dividend yield was displayed to subjects if and only if she held the relevant asset. Current asset and cash holdings were provided, as well as the previous change in holdings, which appeared as a positive (negative) value for assets that the subject bought (sold), and a negative (positive) value for the cash paid (received). The interface for the two-asset markets, $2N$ and $2Z$, is similar to the one presented in Figure 1, except that all information related to asset C is omitted, including prices, NAV, dividends and buy/sell orders. Subject endowments of cash and assets for all sessions are presented in Table 1.

Table 2: Overview of sessions

Treatment	Sessions	Subjects per session	Payoff (USD, without show-up fee)
$2Z$	5	9	24.64
$3Z$	5	9	24.80
$2N$	5	9	24.64
$3N$	5	9	25.03
Total	20	180	24.78

In total, we conducted 20 sessions, with five sessions per treatment, and nine subjects per session. We present an overview of all sessions from our experiment in Table 2. All subjects participated in two, 15-period markets. At the end of the experiment, one of these two markets was randomly selected and subjects' total point earnings from the market were converted into US dollars at the known exchange rate of \$0.04 per point. Subjects market earnings were equal to the sum of their dividends received and/or storage costs paid over all 15 rounds from assets held plus their remaining cash balance and the value of their asset position at the end of the 15th round. The latter value was determined by summing together 1) the number of units of asset A held multiplied by 10, 2) the number of units of asset B held multiplied by 10, and 3) (in the three asset markets), the number of units of the ETF asset C held multiplied by 20. On average, each session lasted two hours and the average earnings were \$24.78.

[10] This procedure, following Kirchler et al. (2012), ensures that there is no increase in the cash-to-asset ratio over time.

In addition, subjects received a show-up fee of \$7, bringing the average total to \$31.78.

5 Results

We begin our analysis by presenting the median market prices, and the fundamental values, of assets A, B and C in Figure 2 across all markets and treatments. The x-axis provides the period number for each of the two markets, with each lasting 15 periods. In the three asset treatments, $3N$ and $3Z$, we also include the NAV which is computed as the sum of the last available prices of assets A and B. In Appendix A we also include the same figures using a subset of data limited to the second market, when subjects have more experience. In all of the statistical analysis that follows, we use data only from the second market to account for subject learning.[11] Overall, the price of asset A closely follows its flat fundamental value across all treatments except $2N$, where we observe a higher level of mispricing. Such behavior is consistent with trading of assets with flat fundamental values and a constant cash to asset ratio as shown in Kirchler et al. (2012). The mispricing of asset B increases at the structural break where the fundamental value changes trend, and then decreases by the terminal period. We also observe that the price of the ETF asset C is closer to the NAV than to the fundamental value in the negative correlation treatment.

In treatment $3N$ the ETF asset C provides perfect insurance over market outcomes while in treatment $3Z$, the insurance is incomplete. Therefore, if investors are risk-averse, then we need to account for the difference in insurance provided by the different market environments.[12]

Table 3 presents a summary of our results using averages for each of the four treatments: $2Z$, $3Z$, $2N$, and $3N$, respectively.[13] To study market behavior, we employ the following measures: price amplitude, volatility, relative absolute deviation (RAD), relative deviation (RD), relative pricing, asset turnover, relative absolute deviation from the NAV (RAD-NAV), and relative deviation from the NAV (RD-NAV). Our first measure, the price amplitude of an asset relative to its fundamental value (PA-FV), was introduced by Porter and Smith (1995). We define this measure as the difference be-

[11] Our results are qualitatively similar if we use the data from both markets.

[12] Using a call market, Biais et al. (2017) find that the price of a risky asset is lower in an environment with aggregate risk than in an environment without such risk. In multiple asset markets, a risk-premium also appears under aggregate risk in Bossaerts and Plott (2004) and Bossaerts et al. (2007).

[13] Appendix B includes the measures presented in Table 3 for each second market across all sessions.

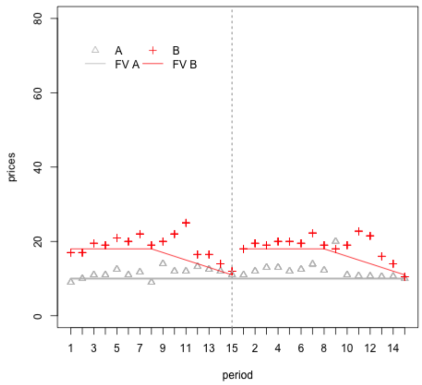

(a) treatment 2Z

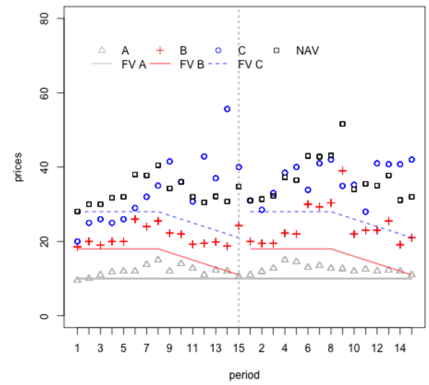

(b) treatment 3Z

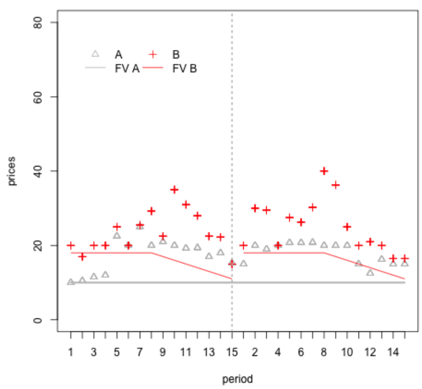

(c) treatment 2N

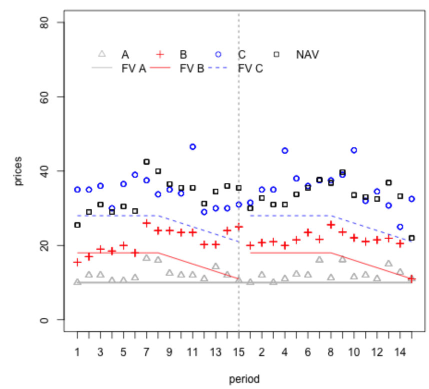

(d) treatment 3N

Figure 2: Median asset price per period for each treatment (5 sessions) and fundamental values. Each session contains two markets, each lasting 15 periods.

tween the maximum deviation from the fundamental value and the minimum deviation from the fundamental value of an asset, normalized by the initial fundamental value, such that

$$\mathrm{PA}_j - \mathrm{FV}_j := [\max(P_{j,t} - FV_{j,t}) - \min(P_{j,t} - FV_{j,t})]/FV_{j,1}, \tag{4}$$

where $t = 1, \dots, 15$ and $j \in \{A, B, C\}$. We also compute the volatility of the prices of each asset using the standard deviation of the prices in a market. To measure the extent of mispricing, or the price deviation from the fundamental value, we follow Charness and Neugebauer (2019), and define the relative absolute deviation of each

Table 3: Summary of results

	First market				Second market			
	2Z	3Z	2N	3N	2Z	3Z	2N	3N
Asset A								
Price amplitude	0.52	1.04	2.92	0.89	0.62	1.25	2.93	1.11
Volatility	1.81	3.12	8.88	2.61	2.03	4.13	8.70	3.49
RAD	0.26	0.32	1.01	0.51	0.39	0.70	1.27	0.40
RD	0.22	0.27	0.99	0.45	0.39	0.69	1.21	0.39
Turnover	1.22	1.51	1.45	1.35	1.48	1.48	1.25	1.14
Asset B								
Price amplitude	0.65	0.76	1.79	1.06	0.45	1.07	1.15	0.48
Volatility	3.54	3.31	8.82	5.90	3.65	5.71	7.21	2.25
RAD	0.31	0.47	0.70	0.52	0.34	0.88	0.60	0.36
RD	0.25	0.46	0.62	0.44	0.32	0.87	0.55	0.35
Turnover	1.18	1.33	1.30	1.54	1.19	1.06	0.99	1.11
RAD B/A	0.22	0.33	0.33	0.55	0.17	0.41	0.31	0.15
Asset C								
Price amplitude	–	0.99	–	0.43	–	0.35	–	0.52
Volatility	–	7.28	–	3.13	–	3.66	–	5.40
RAD	–	0.55	–	0.45	–	0.59	–	0.47
RD	–	0.43	–	0.48	–	0.52	–	0.47
RAD-NAV	–	0.34	–	0.17	–	0.29	–	0.19
RD-NAV	–	0.04	–	0.08	–	0.09	–	0.10
Turnover	–	1.15	–	0.99	–	0.83	–	0.90

asset from fundamental value as:

$$\text{RAD-FV}_j := \frac{1}{T}\sum_{t=1}^{T} |P_{j,t}/FV_{j,t} - 1|, \tag{5}$$

where T represents the total number of periods in which assets are transacted.[14] Our results do not significantly change when using a geometric average, as suggested by Powell (2016). We also consider a relative deviation measure, (6) which is same as RAD, but does not take the absolute value of the deviation. RD is useful in tests where the null hypothesis assumes zero differences across variables, and by design the RAD measure is always bounded away from zero.

$$\text{RD-FV}_j := \frac{1}{T}\sum_{t=1}^{T} P_{j,t}/FV_{j,t} - 1, \tag{6}$$

[14] We drop the periods in which there are no observations for a given asset from our calculations.

To account for multiple assets, we extend the RAD formula in equation (4) and assign weights according to the number of assets transacted relative to overall number of market transactions. Thus, the total relative absolute deviation is

$$\text{RAD-TOTAL} := \frac{1}{Q}\sum_{j=A}^{C}\sum_{t=1}^{T}|P_{j,t}/FV_{j,t} - 1| \times q_{j,t}, \tag{7}$$

where Q is the total number of market transactions across all assets, and $q_{j,t}$ is the quantity of asset j transacted in time t. If we divide Q by the total supply of assets, which differs across two and three asset markets, then we obtain our measure of asset turnover.

We also present the relative mispricing of assets A and B using a relative RAD B/A, which we define as

$$\text{RAD-B/A} := \frac{1}{T}\sum_{t=1}^{T}|\frac{P_t^B/FV_t^B}{P_t^A/FV_t^A} - 1|, \tag{8}$$

where T is the total number of periods with joint transactions of A and B. In the case of the ETF asset C, we measure how far it is priced from the NAV. We follow the definition above, and measure the price deviation with respect to NAV as

$$\text{RAD-NAV} := \frac{1}{T}\sum_{t=1}^{T}|P_t^C/NAV_t - 1|. \tag{9}$$

Similarly, we measure how for the ETF asset C is priced from NAV using RD measure, which does not take the absolute value of the deviation. Table 3 shows that in the negative correlation treatments, there is less price volatility and prices are closer to the fundamental value when there is an ETF asset ($3N$) compared to the case without the ETF asset ($2N$). For example, the volatility and RAD of asset B dramatically improve (from 7.21 in $2N$ to 2.25 in $3N$ and from 0.60 in $2N$ to 0.36 in $3N$ respectively, in the second market). On the other hand, in the zero correlation treatment, the introduction of the ETF asset is less helpful in terms of pricing assets A, and B, and reducing volatility.

Such outcome is intuitive, because when the correlation in asset returns is negative there is a greater demand for both assets in order to eliminate market risk. The increase in demand for assets moves prices away from their fundamental values. Once an ETF asset is introduced, the demand for individual assets decreases. In the zero correlation treatment the role of the ETF asset is less evident for subjects, which explains why its effect is generally not significant in subsequent analysis. The measure of asset turnover, which tells us how active the market is similar across all treatments.

Notice that in Table 3 there is little change in most measures for asset A between markets 1 and 2. However, most measures of asset B —price amplitude volatility and RAD— are considerably lower in the second market relative to the first. Given that asset B is harder to price due to a decreasing trend in the fundamental value following period eight, we need to allow for subject learning. For this reason, we focus most of our subsequent analysis on the second market.

Lastly, the results for asset C suggest that when the correlation between asset returns is negative, the introduction of the ETF asset improves most of the measures reported in Table 3. The price deviations for asset C with respect to NAV are smaller than with respect to the fundamental value. This indicates that subjects are pricing assets correctly in the relative sense. In the subsequent paragraphs, with the help of parametric as well as non-parametric tests, we formalize our results, using data from the second market to control for possible learning effects.

Result 1: *In the case where asset returns are negatively correlated, assets A and B exhibit lower volatility when an ETF asset is traded than in a market without an ETF asset.*

In our non-parametric analysis we use each second market from our data set as a unique observation to perform a two-sided Wilcoxon test to evaluate differences, if any, across treatments (with five observations per treatment). Table 4 shows the direction of these differences and the relevant p-values for at least a 10 percent significance level. The volatility of assets A and B in the $3N$ treatment is lower than in the $2N$ treatment (with p-values of 0.095 and 0.032 for assets A and B, respectively). The higher p-value for asset A relative to B may be due to the differences in the path of the fundamental values of the two assets. There is a greater chance of observing a deviation from the fundamental value of asset B because it is changing across time, unlike asset A which has a constant fundamental value. Thus, there is more room for improvement in the case of asset B.

The ETF asset does not appear to play a significant role in the environment with zero correlation in asset returns. We do not observe any differences between the $2Z$ and

Table 4: Wilcoxon Test: volatility (p-values)

		A			B		C
Market	2Z	3Z	2N	2Z	3Z	2N	3Z
2Z							
3Z	=			=			
2N	> (0.056)	> (0.095)		> (0.056)	=		
3N	=	=	< (0.095)	=	=	< (0.032)	=

$3Z$ treatments. In the zero correlation treatment, aggregate risk still exists, however the ETF asset cannot offer the same level of insurance as in the negative correlation case. Using another measure for the dispersion of prices in Table 5, price amplitude, we find similar outcomes. The ETF asset decreases price amplitude for assets A and B when the correlation is negative, but not when the correlation is zero.

Table 5: Wilcoxon Test: price amplitude (p-values)

		A			B		C
Market	2Z	3Z	2N	2Z	3Z	2N	3Z
2Z							
3Z	=			=			
2N	> (0.048)	> (0.087)		> (0.032)	=		
3N	=	=	< (0.087)	=	=	< (0.079)	=

Result 2: *When asset returns are negatively correlated, the price of asset A, and relative price of asset B with respect to A is closer to the fundamental value in a market with an ETF asset than in a market without an ETF asset.*

We follow equation (5) to compute the RAD for each asset, and equation (8) to study the relative price dispersion with respect to fundamentals. The introduction of the ETF asset appears to affect only the asset with a constant fundamental value (A), and only in the case where the ETF asset provides a perfect diversification of market risk (p-value of 0.056). Similarly, the relative price of asset B with respect to A is closer to the fundamental value (p-value of 0.032) in the environment with negative correlation. There is also improvement in relative pricing B/A in the three asset environment (with an ETF) when the correlation between individual assets is negative, $3N$ versus $3Z$ (p-value of 0.095).

Table 6: Wilcoxon Test: RAD-FV (p-values)

	A			B			B/A			C
Market	2Z	3Z	2N	2Z	3Z	2N	2Z	3Z	2N	3Z
2Z										
3Z	=			=			=			
2N	> (0.056)	=		=	=		=	=		
3N	=	=	< (0.056)	=	=	=	=	< (0.095)	< (0.032)	=

Complementary to the non-parametric analysis, we perform OLS regressions using our measures of RAD from equation (5) at the period level, for each asset as the dependent variable. That is, we do not aggregate the absolute deviation across periods in this analysis. Table 7 presents the results. As independent variables we include a set of dummies, Z, $Three$, and $Three \times Z$, which capture marginal effects of the different treatments. The intercept measures the effect of the baseline treatment $2N$.

Table 7: OLS regressions: RAD

	RAD-FV A		RAD-FV B		RAD-B/A		RAD-FV C	
	(I)	(II)	(III)	(IV)	(V)	(VI)	(VII)	(VIII)
Intercept	1.26^{***}	1.29^{***}	0.64^{***}	0.37^{*}	0.37^{***}	0.27^{**}	0.44^{***}	0.27
	(0.31)	(0.42)	(0.21)	(0.21)	(0.12)	(0.12)	(0.16)	(0.18)
Z	-0.84^{**}	-0.84^{**}	-0.30	-0.30	-0.21	-0.21	0.36	0.36
	(0.36)	(0.37)	(0.24)	(0.24)	(0.13)	(0.14)	(0.51)	(0.51)
$Three$	-0.82^{**}	-0.82^{**}	-0.25	-0.22	-0.21^{*}	-0.20^{*}	–	–
	(0.37)	(0.37)	(0.25)	(0.24)	(0.12)	(0.12)		
$Three \times Z$	1.12^{*}	1.12^{*}	0.73	0.69	0.41^{**}	0.38^{**}	–	–
	(0.64)	(0.64)	(0.46)	(0.53)	(0.17)	(0.17)		
Period	–	-0.00	–	0.04^{***}	–	0.01^{**}	–	0.02^{*}
		(0.03)		(0.01)		(0.01)		(0.01)
R^2	0.11	0.11	0.09	0.15	0.11	0.15	0.06	0.08
N	240	240	242	242	205	205	103	103

Note: Standard errors are clustered at the session level and computed via bootstrapping.
$^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

We define our dummy variables as follows: (i) Z takes the value of one when the correlation of asset returns is zero, and zero otherwise, and (ii) $Three$ takes the value of one when an ETF asset is traded in the market, and zero otherwise. Hence the coefficients Z and $Three$ capture the $2Z$ and $3N$ treatments, respectively. We also include an interaction term $Three \times Z$ which captures the effect of the $3Z$ treatment. To control for group effects, we cluster the standard errors by session, and compute via bootstrapping. The trend variable, *Period* controls for time and learning. For every asset, we present the specification with and without the time trend. In the last two specifications of Table 7, the interpretation of the intercept is different because the ETF asset C is traded only in the $3N$ and $3Z$ asset markets. Hence, for specifications (VII) and (VIII) the constant captures the baseline $3N$ treatment.

The negative sign on the coefficient of *Three* in specifications (I)-(VI) indicates that RAD is smaller in $3N$ compared to $2N$. The difference is significant at the five percent level for asset A and at the ten percent level for the relative price B/A. Specifically, the RAD for asset A is smaller by 82 percentage points in the $3N$ relative to $2N$. We test the effect of the ETF asset on RAD within a zero correlation environment using a Wald test, where the null hypothesis states that the sum of coefficients for the $Three$ and $Three \times Z$ is equal to zero. In all specifications, we cannot reject that RAD of individual assets in $3Z$ is equal to the RAD in $2Z$. For the ETF asset C, specifications (VII-VIII) do not show any important differences across treatments. The positive time trend coefficient ($Period$) captures the effect of the decreasing trend in the fundamental value of assets B and C, which complicates pricing behavior.

Table 8: Wilcoxon Test: RAD-TOTAL (p-values)

	2Z	3Z	2N
2Z			
3Z	$=$		
2N	$>$ (0.095)	$=$	
3N	$=$	$=$	$<$ (0.056)

To study the overall mispricing relative to fundamentals we refer to equation (7). Consistent with our earlier results, we find that mispricing is smaller in $3N$ than in $2N$, with a p-value of 0.056 as indicated in Table 8. Thus, when the correlation between assets is negative ETFs help reduce the overall level of mispricing.

Result 3: *Introducing the ETF asset does not affect the volume of trade.*

To determine whether the ETF asset affects market activity, or the volume of trade, we compare our two-asset markets (without ETFs) and our three-asset markets (with ETFs) and use asset turnover as an indicator of market activity. The results are reported in Table 18, Appendix B. We define asset turnover as total assets transacted in the market divided by the total supply. Note that in our environment, the ETF asset represents a significant portion of market supply, or 50 percent of assets. We do not find any difference in the trading volume across the two and three asset markets. This result is robust to extending the analysis to the number of bids and asks in the market.

Result 4a: *The price of the ETF asset C closely follows its net asset value.*

We test whether the price of the ETF asset is equal to the NAV using the relative deviation (RD) measure.[15] For each of the treatments, $3N$ and $3Z$, we cannot reject the null ($p > 0.10$) that the price of the ETF asset is equal to the NAV. Further, we do not find any statistical difference in the market price of the ETF asset across the $3Z$ and $3N$ treatments.

Table 9: OLS Regressions: P_C-NAV

	P_C-NAV (I)	P_C-NAV (II)
Intercept	3.824	5.178*
	(2.608)	(3.083)
Z	6.836	4.270
	(9.185)	(10.147)
Period	−0.231	−0.418*
	(0.213)	(0.221)
Period × Z	–	0.352
		(0.427)
R^2	0.04	0.07
N	103	103

Notes:
a. Standard errors are clustered at the session level and computed via bootstrapping.
$^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

We also analyze the effect of time and correlation in asset returns on arbitrage. We look for evidence that price differences from NAV are arbitraged over time by regressing the difference between the price of the ETF asset and its NAV per period on

[15] In this case, RD is a more appropriate measure because RAD is constrained by the lower bound of zero, and thus a test using RAD would fail on the account of the definition of the measure.

time as well as the correlation between asset returns in Table 9. The results suggest that in specification (I), there is no difference between the price of the ETF asset and its NAV in both the zero and negative correlation treatments. The *Period* variable, while not significant, has a negative sign which would suggest learning. Specification (II) allows further analysis with the addition of the interaction term $Period \times Z$. In this specification, the time trend coefficient, -0.418 is significantly negative (with a p-value of 0.059). This suggests that most of the arbitrage over time occurs in the negative correlation treatment. Thus, we can conclude that the ETF asset serves as an important a benchmark for price discovery.

Result 4b: *The bids across all assets, A, B, and C are more consistent with the no arbitrage prediction in treatment $3N$ than in treatment $3Z$.*

We further consider whether there is any inconsistency between the (i) individual bids for the underlying assets, A and B, and the bids for the composite asset C, and (ii) individual asks for the underlying assets, A and B, and the asks for the composite asset C. We use individual bid data to measure the absolute difference between the sum of the bids (b) for assets A and B and the bids for asset C, as defined by:

$$\text{AD-BIDS} := \frac{1}{T \cdot N} \sum_{i=1}^{N} \sum_{t=1}^{T} |b_{it}^{A} + b_{it}^{B} - b_{it}^{C}|, \tag{10}$$

where T is the number of periods, and N the number of instances with joint bids for all three assets A, B and C. The same approach is used to measure the absolute difference between the sum of the asks for assets A and B and the asks for asset C (AD-ASKS).

Table 10: Bid and ask consistency in ETF markets

	AD: BIDS		Count (subjects)		AD: ASKS		Count (subjects)		AD: BIDS + ASKS		Count	
Market	3Z	3N	3Z	3N	3Z	3N	3Z	3N	3Z	3N	3Z	3N
I	5.43	4.76	38 (6)	6 (2)	7.13	6.43	12 (3)	7 (4)	5.84	5.66	50	13
II	9.46	3.97	21 (4)	19 (4)	31.5	4.76	5 (4)	27 (5)	13.70	4.43	26	46
III	7.46	3.06	24 (6)	16 (3)	7.56	–	4 (3)	0 (0)	7.47	3.06	28	16
IV	18.35	1.82	13 (6)	47 (6)	23.69	2.77	8 (3)	33 (4)	20.38	2.21	21	80
V	27.67	6.12	6 (3)	39 (3)	18.67	9.01	3 (1)	25 (2)	24.67	7.25	9	64
Mean	13.67	3.95	21 (5.0)	19 (3.6)	17.01	5.74	6.4 (2.8)	18.4 (3)	14.41	4.52	26.8	43.8

Table 10 reports on AD-BIDS, AD-ASKS and AD-BIDS + AD-ASKS along with counts of the number of instances of joint bids or asks for all three assets and the number of subjects making these joint asks or bids (in parentheses). Using a Wilcoxon test, we find that the absolute deviation of bids (left columns) asks (middle columns) and bids and asks (right columns) presented in Table 10, is smaller in the $3N$ treatment relative to the $3Z$ treatment, with a p-values of 0.016, 0.063, 0.016, respectively. Thus, consistent with the results reported in Table 9, arbitrage appears to be stronger in $3N$ than in $3Z$.

Result 5: *Risk averse investors have a more balanced portfolio in the negative correlation treatment.*

We analyze how risk attitudes affect portfolio allocations between assets A and B by constructing a measure of portfolio imbalance as the absolute difference between an individual's holdings of A and B divided by their total holdings of assets A and B. In both environments there is an equal supply of assets A and B (though the total supply of these assets depends on whether there is an ETF or not), so the market portfolio would consist of holding equal numbers of both assets. Our portfolio imbalance measure helps quantify how far an individual portfolio is from the market portfolio. Since holding a composite asset C is equivalent to holding one A and one B, when a subject adds asset C to her portfolio it increases the denominator in our imbalance measure by a count of two. If a subject does not hold any assets, which accounts for 13 percent of participants, we drop that subject from the database.[16] When the portfolio imbalance measure is zero, a subject holds the market portfolio, or equal units of assets A and B, while a value of one means that a subject's portfolio only consists of one asset, indicating extreme imbalance.

In order to estimate the effect of risk aversion on portfolio imbalance, we regress the individual portfolio imbalance measure on risk attitudes, as measured by the bomb elicitation task Crosetto and Filippin (2013). The mean number of the boxes collected is 35.83, which is below the risk-neutral benchmark of 50. The standard deviation of boxes collected is 16. We measure the *Risk* as the maximum number of boxes (100) minus the number of boxes collected. Thus risk aversion range is from 100, maximal risk aversion, to 0, maximal risk tolerance. Table 11 summarizes the regression results of our analysis. All regressions use holdings at the end of the session, and compute the standard errors (clustered at the session level) using bootstrapping. Specification

[16]There is no significant difference in the number of boxes collected by the dropped players relative to other players in our bomb elicitation task.

(I) shows that on average a subject who collects all boxes ($Risk = 0$) will have an imbalance of 0.603, and that risk aversion has a small but positive effect on the portfolio imbalance, at the 0.05 significance level.

Table 11: OLS Regressions: Portfolio Imbalance

	(I)	(II)
Intercept	0.603***	0.781***
	(0.110)	(0.160)
Risk (100−Boxes)	−0.003**	−0.006***
	(0.002)	(0.002)
Z	–	−0.342
		(0.211)
Risk × Z	–	0.006**
		(0.003)
R^2	0.02	0.04
N	156	156

Notes:
a. Portfolio imbalance is measured as $|A - B|/(A + B + 2C)$.
b. Standard errors are clustered at the session level and computed via bootstrapping.
$^{***}p \leq .01$, $^{**}p \leq .05$, $^{*}p \leq .1$

To control for different correlations between asset returns, we include a dummy variable Z, which takes the value of one when the asset returns have zero correlation, and the value of zero otherwise. Lastly, we also include an interaction term $Z \times Risk$. According to specification (II) in Table 11, where the intercept is now 0.781, higher risk aversion should decrease the imbalance in the negative correlation treatment. On the other hand, risk does not seem to play a role in the zero correlation treatment. The interaction between the Z treatment and the risk measure results in a positive but small increase in the imbalance measure.

Result 6: *There is heterogeneity in trading strategies.*

We classify trader behavior according to three types, (i) fundamental value trader, (ii) momentum trader, and (iii) rational speculator closely following the framework of Haruvy and Noussair (2006), Haruvy et al. (2013), and Breaban and Noussair (2015). An individual's behavior is defined as consistent with the fundamental value trader type at period t if $s_{i,t} < s_{i,t-1}$ when $p_t > f_t$, where p_t is the price, f is the fundamental value and $s_{i,t}$ is the number of assets that individual i holds in period t. This means that when asset prices rise above the fundamental value, trader i is a net seller. Similarly, if asset prices go below fundamental value, then trader i will be a net buyer. Trader behavior is consistent with being a momentum trader if $s_{i,t} < s_{i,t-1}$ when $p_{t-1} < p_{t-2}$. Similarly, if $p_{t-1} > p_{t-2}$ then $s_{i,t} > s_{i,t-1}$. A momentum trader is a net buyer in period t when there is an increasing price trend over the last two periods, and a net seller when there is a decreasing price trend. Lastly, trader behavior is defined as consistent with the rational speculator trader type if $s_{i,t} < s_{i,t-1}$ when $p_{t+1} < p_t$. Likewise, if $p_{t+1} > p_t$ then $s_{i,t} > s_{i,t-1}$. A rational speculator is assumed to anticipate next period's price in an unbiased manner, and makes positive net purchases when the price is about to increase and positive net sales when the price is about to decrease.

To classify a subject as a type, we first count the number of periods for each asset when subject behavior is consistent with a particular type. We then sum across these counts to determine which trading behavior the subject follows most often per asset. Finally, in order to classify subject behavior according to one of the three aforementioned types, we require that a subject behaves as a particular type for at least two of the three assets.[17] We present these results in Table 12. In the case where there is no dominant behavior, the subject is classified as a combination of two or three types, which we display in the last four rows of Table 12.

Table 12: Proportion of trader types

Type	2Z	3Z	2N	3N
Fundamental	0.09	0.22	0.18	0.29
Momentum	0.13	0.13	0.09	0.18
Rational Speculator	0.18	0.38	0.20	0.20
Fundamental & Momentum	0.22	0.00	0.11	0.02
Fundamental & Rational Speculator	0.13	0.07	0.24	0.02
Momentum & Rational Speculator	0.20	0.13	0.11	0.20
All 3 types	0.04	0.07	0.07	0.09

Overall, we find that all three types reported in the earlier literature, involving only single assets, are also present in our multi-asset markets. The frequency of these trading strategies also varies, potentially in a response to the environment. For example, the proportion of traders who follow the fundamental strategy for buying and selling is 0.22 in the $3Z$ treatment and 0.29 in the $3N$ treatment.

[17] This construction allows comparisons across treatments with two and three asset markets.

6 Conclusion

Exchange traded funds now comprise 35 percent of all equity trades in the United States and have been exponentially growing in popularity. Understanding whether and how such assets affect market measures such as prices and volatility is important for policymakers concerned with financial stability. In this paper, we have taken a novel first step toward addressing this question with the first laboratory market experiment to use composite, exchange traded assets.

We find that for the most part, ETFs do not foster mispricing, do not contribute to price volatility nor reduce market activity. When the incentive for holding an ETF asset is especially salient, as in our negative correlation treatment where it provides perfect insurance, these composite assets actually reduce mispricing, by providing an important benchmark for price discovery. We also find that risk-averse traders prefer more balanced portfolios in our negative correlation treatment, and therefore providing an asset which facilitates holding the market portfolio may result in faster convergence to the Sharpe ratio.

Our markets with and without ETF assets make use of a simple call market environment, where the composite asset completely covers the market. In future research, we propose to study continuous time trading, which would improve opportunities for arbitrage and to consider ETFs which do not cover the market or which are not representative of market capitalizations. We further propose to study ETFs in a CAPM environment (Bossaerts et al., 2007). These projects, currently in progress, will improve our understanding of how ETFs affect diversification strategies and contribute to a lower market risk premium. Lastly, it would be interesting to study, the role of APs in providing liquidity to markets, or how multiple ETFs may affect the efficiency of markets. We leave these questions for future work.

Acknowledgements

We gratefully acknowledge research support for this project from the Small Grants Program of the International Foundation for Research in Experimental Economics (IFREE). The experiment was approved by the IRB at UC Irvine. We are grateful for comments and feedback received from Stephen Brown, Charles Noussair, Tibor Neugebauer, Peter Bossaerts, Dan Friedman, Chris Veld, Neal Galpin, Sascha Füllbrunn, as well as from participants in numerous seminars and workshops.

Appendix A: Price plots for each session

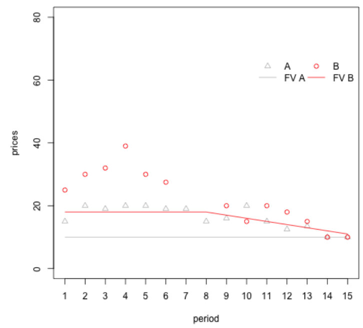

(a) Session 1

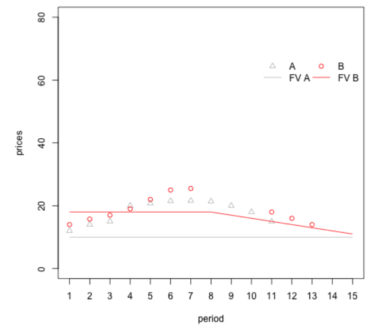

(b) Session 2

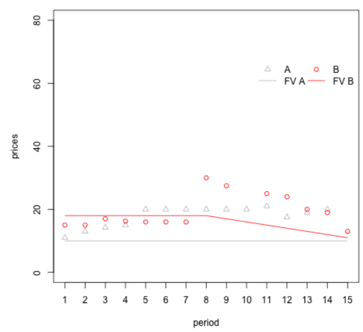

(c) Session 3

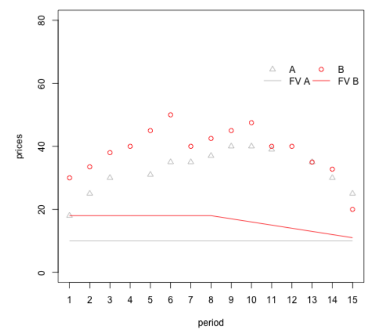

(d) Session 4

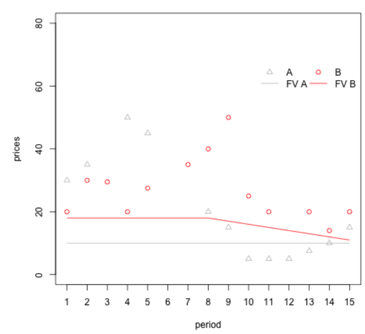

(e) Session 5

Figure 3: Asset prices and fundamental values per period for treatment 2N in the second market.

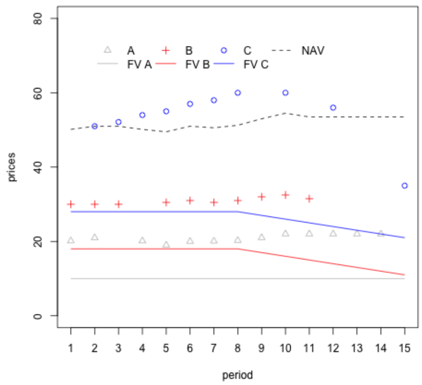

(a) Session 1

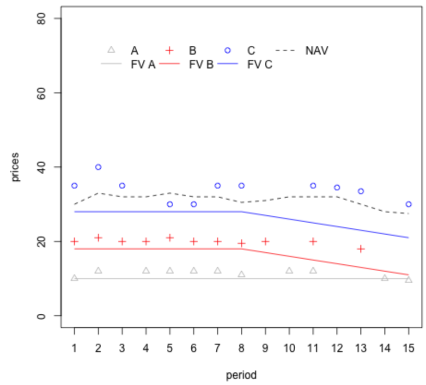

(b) Session 2

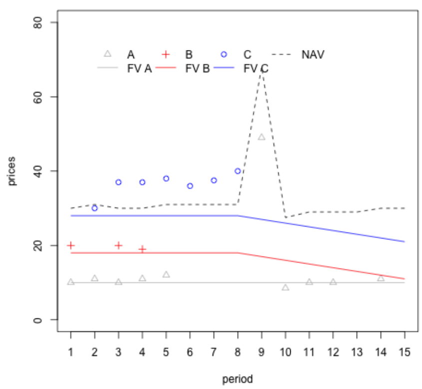

(c) Session 3

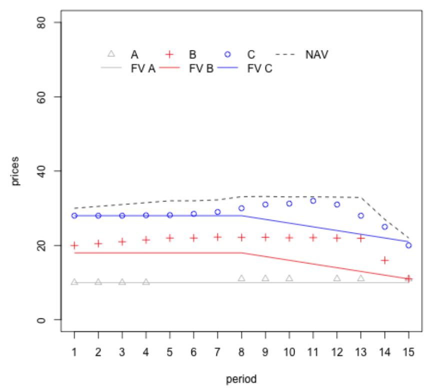

(d) Session 4

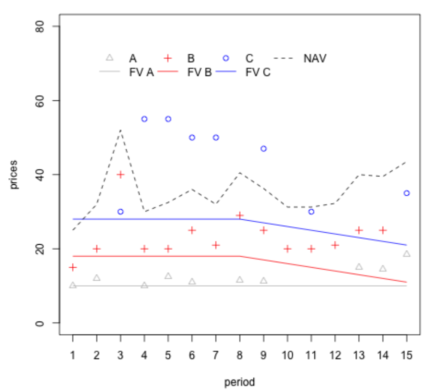

(e) Session 5

Figure 4: Asset prices and fundamental values per period for treatment 3N in the second market.

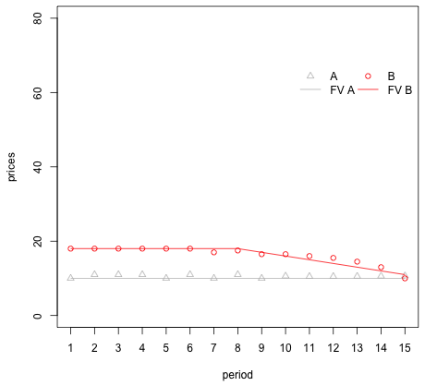

(a) Session 1

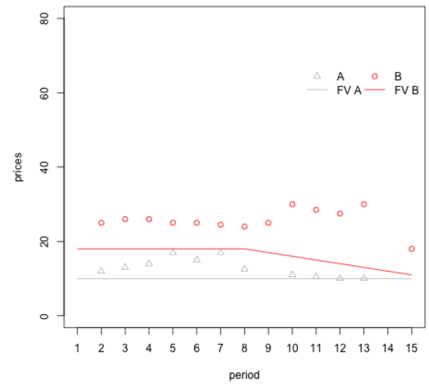

(b) Session 2

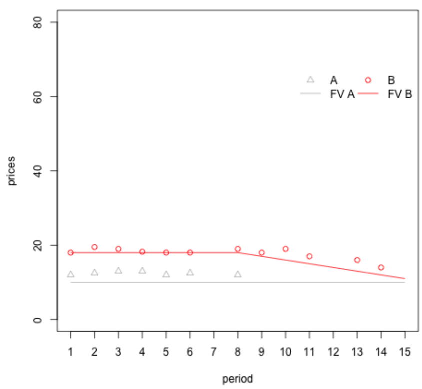

(c) Session 3

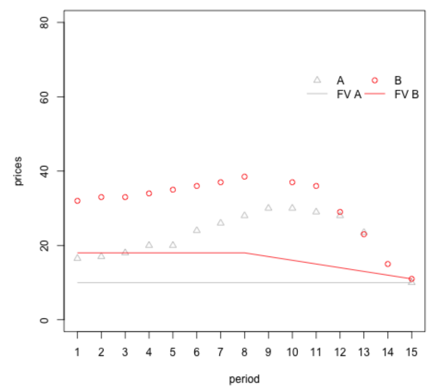

(d) Session 4

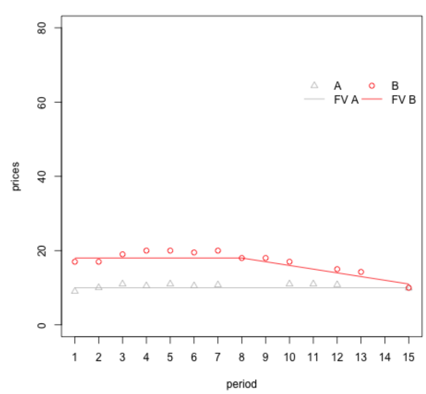

(e) Session 5

Figure 5: Asset prices and fundamental values per period for treatment 2Z in the second market.

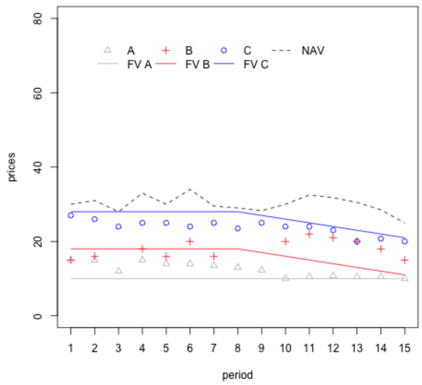

(a) Session 1

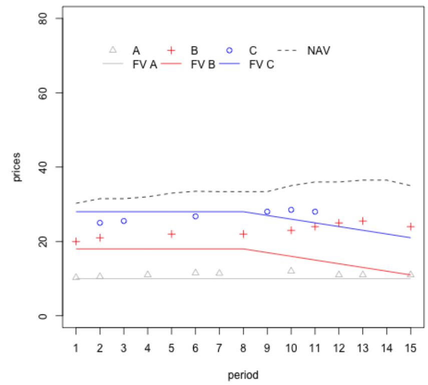

(b) Session 2

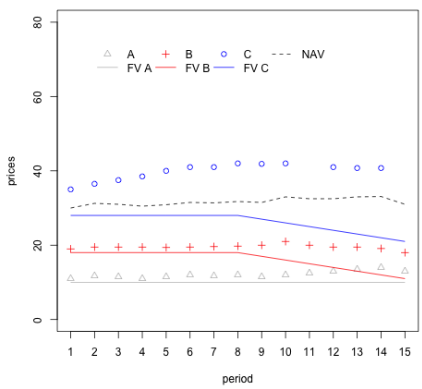

(c) Session 3

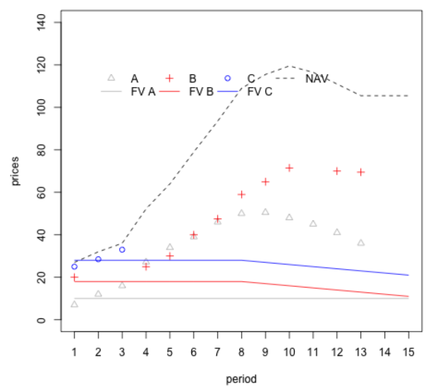

(d) Session 4

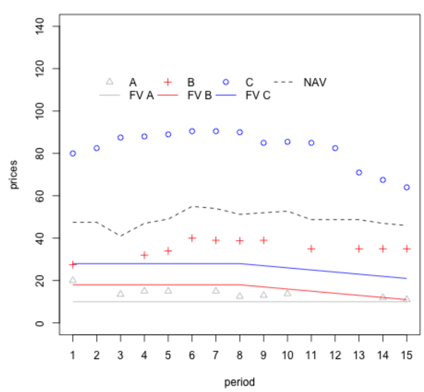

(e) Session 5

Figure 6: Asset prices and fundamental values per period for treatment 3Z in the second market.

Appendix B: Market level data

Table 13: Price Amplitude A, B and C assets

	A				B				C	
Market	2Z	3Z	2N	3N	2Z	3Z	2N	3N	3Z	3N
I	0.10	0.50	1.00	0.30	0.14	0.56	1.28	0.25	0.13	0.71
II	0.70	0.17	0.96	0.25	0.61	0.61	0.64	0.19	0.21	0.36
III	0.10	0.30	1.00	4.05	0.17	0.34	0.83	0.06	0.42	0.36
IV	2.00	4.35	2.20	0.10	1.17	3.03	1.28	0.49	0.29	0.29
V	0.20	0.90	9.50	0.85	0.17	0.81	1.72	1.39	0.70	0.89
Mean	0.62	1.25	2.93	1.11	0.45	1.07	1.15	0.48	0.35	0.52

Note: Price Amplitude is measured as PA-FV $:= [\max(P_i - FV_i) - \min(P_i - FV_i)]/FV_1$

Table 14: Volatility A, B and C assets

	A				B				C	
Market	2Z	3Z	2N	3N	2Z	3Z	2N	3N	3Z	3N
I	0.40	1.93	3.62	1.03	2.30	2.47	8.96	0.88	2.07	7.27
II	2.58	0.52	3.49	1.01	3.08	1.85	4.22	0.79	1.45	2.99
III	0.45	0.89	3.23	12.25	1.53	0.63	5.27	0.58	2.26	3.12
IV	6.10	14.84	6.65	0.52	8.48	19.94	7.60	3.09	4.01	2.93
V	0.63	2.48	26.94	2.66	2.87	3.67	9.98	5.93	8.49	10.66
Mean	2.03	4.13	8.79	3.49	3.65	5.71	7.21	2.25	3.66	5.40

Table 15: RAD A, B and C assets

	A				B				B/A				C	
Market	2Z	3Z	2N	3N	2Z	3Z	2N	3N	2Z	3Z	2N	3N	3Z	3N
I	0.06	0.24	0.63	1.09	0.04	0.27	0.42	0.79	0.06	0.37	0.16	0.13	0.09	1.01
II	0.29	0.11	0.81	0.14	0.61	0.52	0.19	0.17	0.42	0.46	0.36	0.06	0.08	0.30
III	0.24	0.21	0.79	0.46	0.08	0.24	0.34	0.09	0.17	0.09	0.31	0.09	0.51	0.30
IV	1.29	2.52	2.23	.06	0.88	2.09	1.41	0.28	0.14	0.38	0.25	0.22	0.10	0.10
V	0.07	0.41	1.90	0.26	0.07	1.29	0.66	0.46	0.04	0.74	-0.81	0.24	2.16	0.65
Mean	0.39	0.70	1.27	0.40	0.34	0.88	0.60	0.36	0.17	0.41	0.30	0.15	0.59	0.47

Note: RAD-FV $:= \frac{1}{T}\sum_{i=1}^{T}|P_i/FV_i - 1|$ and RAD-B/A $:= \frac{1}{T}\sum_{i=1}^{T}|\frac{P_i^B/FV_i^B}{P_i^A/FV_i^A} - 1|$

Table 16: RD A, B and C assets

	A				B				C	
Market	2Z	3Z	2N	3N	2Z	3Z	2N	3N	3Z	3N
I	0.06	0.24	0.63	1.09	0.01	0.19	0.37	0.78	-0.09	1.01
II	0.29	0.11	0.81	0.13	0.61	0.52	0.11	0.17	0.00	0.30
III	0.24	0.21	0.79	0.43	0.08	0.24	0.22	0.09	0.51	0.30
IV	1.29	2.47	2.23	0.06	0.88	2.09	1.41	0.28	0.03	0.09
V	0.05	0.41	1.63	0.26	0.04	1.29	0.66	0.44	2.16	0.65
Mean	0.39	0.69	1.21	0.39	0.32	0.87	0.55	0.35	0.52	0.47

Note: RD-FV $:= \frac{1}{T}\sum_{i=1}^{T} P_i/FV_i - 1$

Table 17: RAD Total Assets

Market	2Z	3Z	2N	3N
I	0.05	0.19	0.62	0.92
II	0.42	0.19	0.53	0.20
III	0.11	0.28	0.55	0.51
IV	1.06	1.62	1.84	0.15
V	0.07	1.42	1.01	0.48
Mean	0.34	0.74	0.91	0.45

Table 18: Turnover A, B and C assets

	A				B				C	
Market	2Z	3Z	2N	3N	2Z	3Z	2N	3N	3Z	3N
I	4.15	3.43	1.97	1.47	2.12	1.77	0.10	1.47	1.63	0.60
II	0.93	0.70	0.97	1.13	1.00	0.40	0.83	0.73	0.47	1.00
III	0.30	1.60	1.63	1.27	0.77	1.70	1.57	0.30	1.07	0.83
IV	1.27	1.23	0.97	137	1.12	0.80	0.77	1.67	0.33	1.47
V	0.73	0.40	0.72	0.47	0.95	0.63	0.75	1.37	0.63	0.60
Mean	1.48	1.48	1.25	1.14	1.19	1.06	0.99	1.11	0.83	0.90

Note: Turnover is measured as total assets transacted divided by the total supply.

Table 19: C - NAV

	Amplitude		RAD		RD	
Market	3Z	3N	3Z	3N	3Z	3N
I	0.03	1.10	0.21	0.11	-0.20	0.04
II	0.10	0.40	0.20	0.11	-0.19	0.09
III	0.22	0.40	0.26	0.20	0.26	0.19
IV	0.61	0.07	0.09	0.09	-0.09	-0.09
V	1.14	1.88	0.69	0.43	0.69	0.26
Mean	0.36	0.79	0.29	0.19	0.09	0.10

Appendix C: Risk elicitation responses

We elicit risk attitudes following the protocol of Crosetto and Filippin (2013) implemented in oTree (Holzmeister and Pfurtscheller, 2016). The median boxes collected in the treatments $2Z$, $3Z$, $2N$, and $3N$ are 35, 40, 35 and 36, respectively. Using the Wilcoxon test, we cannot reject that the distribution of boxes is equal across treatments. Figure 7 shows the frequency of boxes collected in the two treatments.

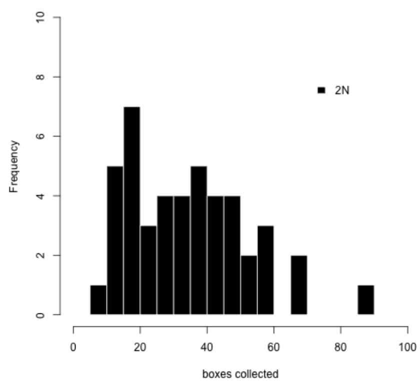

(a) treatment 2N

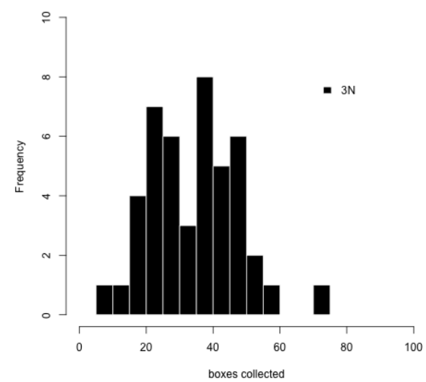

(b) treatment 3N

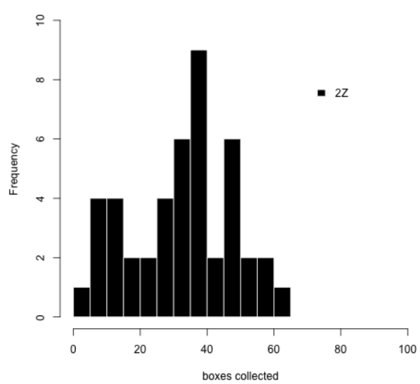

(c) treatment 2Z

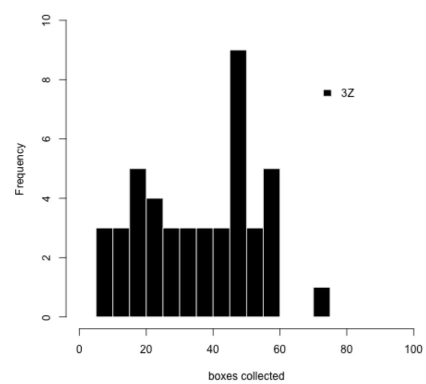

(d) treatment 3Z

Figure 7: Histogram of the risk elicitation task across treatments.

References

Agapova, Anna and Nikanor Volkov, "ETFs and Price Volatility of Underlying Bonds."

Baltussen, Guido, Sjoerd van Bekkum, and Zhi Da, "Indexing and stock market serial dependence around the world," *Journal of Financial Economics*, 2019, *132* (1), 26–48.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, "Do ETFs increase volatility?," *The Journal of Finance*, 2018, *73* (6), 2471–2535.

Bhojraj, Sanjeev, Partha S Mohanram, and Suning Zhang, "ETFs and information transfer across firms," *Rotman School of Management Working Paper*, 2018, (3175382).

Biais, Bruno et al., "Asset pricing and risk sharing in a complete market: An experimental investigation," Technical Report 2017.

Bogle, John C, "The index mutual fund: 40 years of growth, change, and challenge," *Financial Analysts Journal*, 2016, *72* (1), 9–13.

Bossaerts, Peter and Charles Plott, "Basic principles of asset pricing theory: Evidence from large-scale experimental financial markets," *Review of Finance*, 2004, *8* (2), 135–169.

_ , _ , and William R Zame, "Prices and portfolio choices in financial markets: Theory, econometrics, experiments," *Econometrica*, 2007, *75* (4), 993–1038.

Breaban, Adriana and Charles N Noussair, "Trader characteristics and fundamental value trajectories in an asset market experiment," *Journal of Behavioral and Experimental Finance*, 2015, *8*, 1–17.

Cason, Timothy N and Daniel Friedman, "Price formation in single call markets," *Econometrica: Journal of the Econometric Society*, 1997, pp. 311–345.

Charness, Gary and Tibor Neugebauer, "A Test of the Modigliani-Miller Invariance Theorem and Arbitrage in Experimental Asset Markets," *The Journal of Finance*, 2019, *74* (1), 493–529.

Chen, Daniel L, Martin Schonger, and Chris Wickens, "oTreeAn open-source platform for laboratory, online, and field experiments," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 88–97.

Chen, Guang and T Shawn Strother, "On the contribution of index exchange traded funds to price discovery in the presence of price limits without short selling," *Available at SSRN 1094485*, 2008.

Converse, Nathan, Eduardo Levy-Yeyati, Tomas Williams et al., "How ETFs amplify the global financial cycle in emerging markets," *Institute for International Economic Policy Working Paper*, 2018, *1.*

Crosetto, Paolo and Antonio Filippin, "The bomb risk elicitation task," *Journal of Risk and Uncertainty*, 2013, *47* (1), 31–65.

Fang, Yanhao and Gary C Sanger, "Index price discovery in the cash market," in "Midwest Finance Association 2012 Annual Meetings Paper" 2011.

Glosten, Lawrence R, Suresh Nallareddy, and Yuan Zou, "ETF activity and informational efficiency of underlying securities," *Columbia Business School Research Paper*, 2016, (16-71).

Hamm, Sophia JW, "The effect of ETFs on stock liquidity," *Available at SSRN 1687914*, 2014.

Haruvy, Ernan and Charles N Noussair, "The effect of short selling on bubbles and crashes in experimental spot asset markets," *The Journal of Finance*, 2006, *61* (3), 1119–1157.

_ , _ , and Owen Powell, "The impact of asset repurchases and issues in an experimental market," *Review of Finance*, 2013, *18* (2), 681–713.

Hasbrouck, Joel, "Intraday price formation in US equity index markets," *The Journal of Finance*, 2003, *58* (6), 2375–2400.

Hegde, Shantaram P and John B McDermott, "The market liquidity of DIAMONDS, Q's, and their underlying stocks," *Journal of Banking & Finance*, 2004, *28* (5), 1043–1067.

Hill, Joanne M, "The evolution and success of index strategies in ETFs," *Financial Analysts Journal*, 2016, *72* (5), 8–13.

Holzmeister, Felix and Armin Pfurtscheller, "oTree: The bomb risk elicitation task," *Journal of Behavioral and Experimental Finance*, 2016, *10*, 105–108.

Huang, Shiyang, Maureen O'Hara, and Zhuo Zhong, "Innovation and informed trading: Evidence from industry ETFs," *Available at SSRN 3126970*, 2018.

ICI, *Fact Book*, Investment Company Institute, 2019.

Ivanov, Stoyu I, Frank J Jones, and Janis K Zaima, "Analysis of DJIA, S&P 500, S&P 400, NASDAQ 100 and Russell 2000 ETFs and their influence on price discovery," *Global Finance Journal*, 2013, *24* (3), 171–187.

Kirchler, Michael, Jürgen Huber, and Thomas Stöckl, "Thar she bursts: Reducing confusion reduces bubbles," *American Economic Review*, 2012, *102* (2), 865–83.

Lettau, Martin and Ananth Madhavan, "Exchange-traded funds 101 for economists," *Journal of Economic Perspectives*, 2018, *32* (1), 135–54.

Madhavan, Ananth and Aleksander Sobczyk, "Price dynamics and liquidity of exchange-traded funds," *Journal of Investment Management*, 2016, *14* (2), 1–17.

Marquardt, Philipp, Charles N Noussair, and Martin Weber, "Rational expectations in an experimental asset market with shocks to market trends," *European Economic Review*, 2019, *114*, 116–140.

Nam, Jayoung, "Market accessibility, corporate bond ETFs, and liquidity," *Kelley School of Business Research Paper*, 2017, (18-1).

Noussair, Charles and Steven Tucker, "Futures markets and bubble formation in experimental asset markets," *Pacific Economic Review*, 2006, *11* (2), 167–184.

Noussair, Charles N, Steven Tucker, and Yilong Xu, "Futures markets, cognitive ability, and mispricing in experimental asset markets," *Journal of Economic Behavior & Organization*, 2016, *130*, 166–179.

Plott, Charles R and Kirill Pogorelskiy, "Call market experiments: Efficiency and price discovery through multiple calls and emergent Newton adjustments," *American Economic Journal: Microeconomics*, 2017, *9* (4), 1–41.

Porter, David P and Vernon L Smith, "Futures contracting and dividend uncertainty in experimental asset markets," *Journal of Business*, 1995, pp. 509–541.

Powell, Owen, "Numeraire independence and the measurement of mispricing in experimental asset markets," *Journal of Behavioral and Experimental Finance*, 2016, *9*, 56–62.

Smith, Vernon L, Arlington W Williams, W Kenneth Bratton, and Michael G Vannoni, "Competitive market institutions: Double auctions vs. sealed bid-offer auctions," *The American Economic Review*, 1982, *72* (1), 58–77.

_, Gerry L Suchanek, and Arlington W Williams, "Bubbles, crashes, and endogenous expectations in experimental spot asset markets," *Econometrica: Journal of the Econometric Society*, 1988, pp. 1119–1151.

Yu, Lei, "Basket securities, price formation, and informational efficiency," *Price Formation, and Informational Efficiency (March 25, 2005)*, 2005.