ETF indexing strategies and asset prices: Experimental evidence

Peter Bossaerts John Duffy Jean Paul Rabanal

Olga A. Rud Nitin Yadav*

September 9, 2024

Abstract

We examine how the indexing strategy used by exchange traded funds (ETFs) affects the prices of the underlying constituent assets. We study this issue in both the primary market (ETF creations and redemptions using bots as authorized participants) and the secondary market. The experiment includes three environments: (i) no ETF, (ii) an equal weighted ETF, and (iii) a market cap weighted ETF. We find that ETFs significantly affect the value of the constituent assets, in particular the value of assets that are in shortest supply.

*Peter Bossaerts (plb32@cam.ac.uk) is at the University of Cambridge, John Duffy (duffy@uci.edu) is at University of California, Irvine and ISER Osaka University; Jean Paul Rabanal (jean-paulrab@gmail.com) and Olga A. Rud (olga.rud@gmail.com) are at the University of Stavanger, and Nitiv Yadav (nitin.yadav@unimelb.edu.au) is at the University of Melbourne. We are grateful for the comments received from Stephen Brown, Dan Friedman, Debrah Meloso, Bernt-Arne Ødegaard as well as from participants in several seminars and workshops. We also gratefully acknowledge research support for this project from the Australian Research Council (ARC) DP180102284; University of Melbourne R@MAP Chair, and Leverhulme Trust (Leverhulme International Professorship). The pre-registration material can be found at https://osf.io/u2wa9, and the experiment has been approved by the Office of Research Ethics and Integrity at the University of Melbourne (Ref. number 2022-24368-30511-3).

1 Introduction

Index investing using exchange traded funds (ETFs) has become an increasingly popular, low-cost method by which investors can gain exposure to a wide range of assets, while mirroring the performance of a specific market index. The rapid growth in the demand for ETF index assets has increasingly led to experimentation with the indexing strategies that underlie the ETF index assets. While the traditional, market capitalization-based weighting scheme for ETF index funds remains the dominant approach, there has recently been very rapid growth in thematic, leveraged and "smart beta" indexing strategies including the targeted weighting of certain sectors or industries, the use of leverage and the targeting of leverage ratios, as well as the targeting of the inverse of an index's returns (Merlo, 2022). The introduction of such exotic ETFs can be viewed as a creative reaction by fund managers to the evolving investment landscape which has seen a shift away from actively managed funds and toward indexed products. These new ETF index products are *not* market-capitalization-based and are marketed as a kind of middle ground between active management-like strategies and the cost-effective and tax-efficient structure of an ETF. These ETF products have also become quite popular. According to Jackson (2023), investors put \$161.2 billion into such funds in 2022, for an annual growth rate of 9.8%.

The simplest example of a non-market-cap-based, "smart beta" ETF index product is one that assigns equal weight to all assets in some index (Fernand, 2021). For instance, in a market index such as the S&P 500 under an equal weighting scheme, one assigns a weight of 0.20% per stock. The largest equal weight S&P 500 ETF which was started 20 years ago, the \$45 billion Invesco S&P 500 Equal Weight ETF, was among the most popular exchange-traded funds in 2023 with a net inflow that year of \$10.1 billion (Wang, 2023).

Equal weighting is thought to have several advantages, including better diversification across all assets in the index as there is less concentration in assets with high market caps.[1] Indeed, a concern with market cap weighting is that some indexes can become concentrated in just a few stocks, for instance, the S&P 500 is currently dominated by 6 large technology firms (Apple, Microsoft, Nvidia, Amazon, Meta Platforms and Alphabet) that together account for 27.5% of the index value. This limits diversification and can distort the overall performance of the index, exposing investors to wild price swings (see Treynor, 2005 for a discussion).

From a theoretical point of view, market-cap weighting is the only weighting scheme that is consistent with optimal demand in a capital asset pricing model (CAPM) equilibrium; by contrast, equally weighted indices provide the holder with an inferior risk/return trade-off. In this paper we concentrate on these two indexing strategies, market cap and equal weighting for the ETF, in a setting with just two constituent assets. Our approach to equal weighting echoes the $1/N$ rule used in markets because the expected payoffs of the two constituent assets are the same. Thus, our equal weighting scheme is close to emulating the equal weighting scheme used in the field, without unnecessary complications.

We ask whether the presence of ETF assets, along with their particular indexing strategies, matters for (i) the pricing of the constituent assets that make up the ETF (ii) the liquidity/trade volume of those assets, and (iii) the pricing of the ETF index product itself. For instance, while equal weighting and market cap indexing may refer to assets in the same index class (e.g., the S&P 500), the different indexing strategies imply that the prices of these two ETF index assets will generally be quite different. One of our main questions of interest is whether traders appreciate these differences, or whether they just blindly buy ETF index assets for diversification purposes.

[1] For work related to the $1/N$ heuristic used by investors, please see Yuan and Zhou (2023), Magnani et al. (2022), Huberman and Jiang (2006), and Benartzi and Thaler (2001).

In principle, ETFs are derivative assets, and hence, their introduction should not affect the prices of constituent assets because any payoff distribution that can be attained with ETFs can also be attained without them. Such prediction is strengthened in the presence of smart Authorized Participants (APs) who —often using bots, as we shall do in our experiment— can maintain arbitrage-free prices, such that the price of the ETF and its Net Asset Value (NAV) are equal while maintaining a zero net balance. However, this argument ignores trading costs (pecuniary or cognitive). An agent may prefer the ease of trading in ETFs (one trade leads to the acquisition of more than one asset) even if the resulting payoff distribution is not optimal. If this leads to a strong demand for an ETF asset, then there will be an increase in demand for the constituent assets in proportion to their presence in the ETF. This can lead to price pressure, and thus, a particular ETF indexing strategy may distort the prices of constituent assets.

Rather than relying on archival data, we study these research questions using an experimental approach that affords control in a way which cannot be obtained in the field. When a new type of ETF asset is introduced in financial markets, the resulting effect on prices may depend on investor preferences for the particular ETF strategy, and/or the introduction of the ETF *by itself*. In the laboratory, we are able to design a true counterfactual scenario with and without ETFs to make *causal inferences* about the impact of ETFs on asset prices. If market participants only care about final payoffs, then the ETF should have no impact on prices: participants can attain their preferred payoff distribution — given their homegrown preferences — equally well under our alternative scenarios.

From past experimental research (e.g. see Bossaerts and Plott, 2004 and Duffy et al., 2022), we know that CAPM captures the preferences that participants bring to the lab. Therefore, CAPM is the theoretical framework with which we design the experiment and analyze the results. Still, the validity of CAPM is not really crucial, since our goal is to study whether there are pricing effects after introducing an ETF, and if so, what their nature is.

To preview our results, we find that there is a strong demand for ETF products as reflected by the net positive issuance of ETFs under both equal weighting and market cap weighting schemes for the ETF. The strong demand for ETF products creates a positive order imbalance in the difference between the number of bids versus asks for the constituent assets. These effects are larger for the asset in lower supply, as the price of this asset increases relative to the case with no ETFs. This price impact affects the bid-ask spread of this underlying asset most significantly when the indexing strategy uses equal weights. Finally, we find that there are no arbitrage opportunities in the creation and redemption of ETF shares, since the ETF shares are priced following the value of the constituent assets. Overall, our results suggest that ETF products distort the prices and spreads of the constituent assets, and that the extent of these distortions depends on the indexing strategy used.

2 Related Literature

Our paper contributes to the literature studying the impact of ETF assets on stock prices. Looking at the mechanical reconstitution of market indexes, Greenwood (2005) and Chang et al. (2015) have documented the existence of an index premium. Duffy et al. (2024) in a laboratory experiment with included and excluded assets with perfect correlation have found that this premium is likely driven by a preference for diversification. In an environment where fund managers are evaluated against a benchmark index, Pavlova and Sikorskaya (2023) find that an increase in the demand for assets in these indices makes it harder for arbitrageurs to absorb the effect of demand shocks, which leads to more permanent index premiums, and expected lower returns.

In an overlapping generation model with experts and non-experts (who prefer index products), Jiang et al. (2024) show that the rise of passive investment can disproportionately raise stock prices of firms with a larger systematic component. They test their predictions using calibrations and archival data. Similarly, using heterogeneous investors, Davies (2024) predicts that the price impact from trading the index product is not equal across stocks. In our environment, the demand for ETF shares has a price impact that varies across stocks. We observe a larger effect on the stock in short supply, which has a lower risk given our parameter values.

Recent archival studies find that investors often struggle with investments in specialized, non-market cap weighted ETF index products, leading to costly investment decisions. Ben-David et al. (2023) illustrate that specialized ETFs (those that invest in a specific sector or theme) are overvalued, while Bhattacharya et al. (2017) report that ETF investors in Germany have poor timing and selection, and Brown et al. (2021) show that investors pick dominated ETF products.[2]

The experimental literature to date studying ETFs (Duffy et al., 2021, 2024) has focused on a single weighting scheme reflecting the proportion of assets in the market, and therefore, makes use of a market cap weighting scheme. The contribution of this paper is that we can compare the effect of different ETF index strategies against a benchmark scenario without an ETF asset in order to identify the casual effect of different index strategies for prices of the constituent assets, trading volume and the price of the ETF asset itself.[3]

[2]Huang et al. (2023) provides evidence on the performance of another strategy, smart beta indexes. They show that the good performance of such products only exists in backtests; the performance deteriorates after ETF listings.

[3]While Duffy et al. (2021) find that ETFs may offer a good benchmark, they focus only on trade in the secondary market, and do not consider the impact of the primary market on asset prices.

3 Environment

There are three types of assets in our experimental markets: (i) a risk-free asset or *Note*, (ii) two risky assets, $\theta = \{A, B\}$, and in most of our treatments, (iii) an *ETF* index asset that assigns either equal (treatment E) or unequal weights (treatment U) to the two constituent risky assets, θ. As a baseline, we consider a market environment where there are no ETF assets (treatment N).

Traders are of two types: I and II, and are initially endowed with experimental cash and units of assets A and B as shown in Table 1. The supplies of assets A and B are fixed for the duration of a market. If there is an equal number of traders (where an increase of two traders in the market implies an increase of each type equally), then it follows that the supply of asset B is always double that of asset A.

Table 1: Initial endowments

Type	Cash	A	B	$Note$	ETF
I	10	4	10	0	0
II	60	2	2	0	0

Notes: Initial endowment for market participants includes cash and constituent assets according to trader type (I or II). Notes and ETF shares are in zero supply initially. Assets A, B and Notes can be sold short (up to 5 units). ETFs can only be bought in the marketplace (no short-selling) once they are created by an Authorized Participant (AP) who then covers its position by buying the constituent assets.

Our environment makes use of the CAPM experimental market design of Bossaerts and Plott (2004). In that setting, the payoff to an asset share is state-contingent where $s \in \{\text{X, Y, Z}\}$ and is drawn with an equal probability at the end of a market trading round. The state contingent asset payoffs for the constituent assets θ and the *Note* are presented in Table 2. Traders are perfectly informed about the probabilities of the three states and the state contingent payoffs.

In the unequal weighing case U, the *ETF* asset weights reflect the relative supply

Table 2: Asset payoffs

Asset/State	X	Y	Z
A	10	0	5
B	0	5	10
ETF (treatment E)	10	5	15
ETF (treatment U)	10	10	25
Note	5	5	5

Notes: E refers to equal weight treatment where $ETF = 1A + 1B$, while U refers to unequal weight treatment where $ETF = 1A + 2B$.

of the constituent assets in the market with one unit of asset A and two units of asset B. We regard this scheme as the market capitalization weighing since, as Tables 1-2 reveal, the fundamental market value of asset B is always twice that of asset A. By contrast, we approximate the equal weighing strategy E for the ETF asset using one unit of asset A and one unit of asset B.

To determine the equilibrium price of the constituent assets θ, we assume mean-variance utility $U = \mu - \frac{b}{2}\sigma^2$. According to the CAPM model (see, e.g., Hirshleifer and Riley, 1992), the predicted prices of assets A and B are given by

$$p^*_\theta = \mu - b\Delta\bar{Q}, \tag{1}$$

where $\mu = [5, 5]$ is a vector with the expected value of the assets θ,

$$\Delta = \begin{bmatrix} 5^2 \cdot 2/3 & -5^2/3 \\ -5^2/3 & 5^2 \cdot 2/3 \end{bmatrix}$$

is the covariance matrix of risky assets, and $\bar{Q} = [3, 6]$ is the total supply of assets per capita. Assuming a moderate degree of risk-aversion, $b = 0.01$, equation (1) yields $p^*_A = 5$ and $p^*_B = 4.25$.

Note that the equilibrium price of A will always be 5 since the covariance with the market portfolio is zero, and therefore it should be not affected by risk aversion. For our purposes, this is convenient experimentally, as we do not really know the participants' risk aversion. Note that from the moment that the supply of the constituent assets deviates from the target ratio of 2 B for each unit of A, the price of A will no longer be 5. When the ratio decreases, A will become cheaper. Conversely, when the ratio increases, the equilibrium price of A increases above the expected payoff. This last result derives from the fact that A and B have a negative covariance; so when more (less) units of B are available in the market, the price of A increases (decreases).

Assets A, B and the risk-free $Note$ can be short sold by market participants to obtain more cash, which is another risk-free asset. A short-seller receives a price for selling the asset in the marketplace, keeps the sales price, but commits to paying from final period earnings the payout on the short asset position, which is state dependent. The buyer pays the purchase price out of cash holdings, and receives the payout on the share of the asset from the short-seller after trading stops. Each participant can short sell up to 5 shares.[4]

When the ETF asset is introduced and sold in the market, the constituent assets, which depend on the indexing strategy, are removed from the market and held in trust by an AP in order to make state contingent payouts. Thus, both the $Note$ and the ETF assets are offered in zero net supply, which means that the CAPM asset pricing predictions are not affected across treatments E and U relative to the baseline environment without ETF assets, N. The creation of ETF shares is automated and performed by the AP; in section 3.1 we provide greater detail on how the AP introduces

[4]The restriction on short-selling is introduced to ensure that the short-sellers can deliver the payout. In very rare circumstances, default is possible, but it requires the short-seller to use cash acquired to purchase risky securities in an unbalanced way and end up in a state (say X) where dividends on held assets (for state X, asset B) are minimal. In that case, the experimenter compensates the buyer. Bossaerts and Plott (2004) introduced a penalty scheme to further limit the chance of default, but the scheme requires total experiment earnings to depend on final holdings of all trading rounds, not just one.

new shares of ETF assets to the market.

Market participants trade using a continuous double-sided open-book market,[5] with strict time and price priority. There are no hidden limit orders and trading is anonymous. Short-selling is permitted, with a constraint of up to 5 units, for $Notes$ and assets θ but not ETF shares. Market participants initially trade only assets θ, and $Notes$, until an Authorized Participant (AP) posts new ETF shares for sale (details in the next section). Once an ETF share is sold to a market participant, the AP buys the constituent assets in the underlying markets, and the ETF share can be traded in the secondary market. Each round of trade consists of 10 minutes, during which traders can submit limit orders. When a limit order crosses the other side of the market (the bid is greater than or equal to the best standing ask, or the ask is less than or equal to the best standing bid), a trade takes place, at the price of the standing order. Otherwise the limit order is posted in the electronic book of the corresponding market. After markets close, the state of the world, $s \in \{X,\ Y\ Z\}$, is drawn. Traders receive asset payoffs depending on this state and their final asset positions. Subjects only learn of the state realization at the very end of the session, after completing 6 rounds of trading, and only for the round for which they will be compensated.[6]

3.1 Authorized participants

In our environment, APs are bots (automated trading algorithms) who create and redeem ETF index asset shares, offsetting any creation/redemption activity by buy-

[5]The continuous double-sided open-book system works like the continuous double auction (DA) mechanism, but limit orders that do not execute immediately are kept in a "book" that all traders can monitor, until they are executed or canceled, whichever comes first. Orders in the book are referred to as *standing* orders. Orders that execute immediately are referred to as *market* orders since their prices are irrelevant: they trade at the price of the standing order they are paired with; their price only determined that they will be executed.

[6]This choice was made because Bossaerts and Plott (2004) found that subjects did not believe in the independence of draws of the state variable. In our design, investors are forced to consider that all three states are equally likely in every market.

ing/selling the component assets in the underlying markets. The creation of ETFs by the AP precedes any redemption. The process proceeds as follows: (i) first, the AP scans the order books of the constituent assets for the best bids and asks available, (ii) then using the best available outstanding asks for the constituent assets, the AP places an ask for the ETF asset (that is, creates a share).

If the ETF offer is accepted and therefore one share of ETF is issued, the bot places bid limit orders for the constituent assets at the ask prices used in the calculation of its own sell offer in the ETF market. This process is automated and is extremely fast – the bot's reaction times are of the order of 10 milliseconds, including communication with the market server through the internet.[7]

Since the bot is a two-sided market maker in the ETF market, this means that once ETFs assets have been created, the AP also posts bids for the ETF assets at prices that reflect the bids in the markets of the constituent assets. Thus, when there is a redemption of an ETF share (a participant sells the ETF to the AP), the bot places sell orders in the underlying markets at the best bids. Again, the process is fast, to ensure that the bot's actions have minimal impact on market risk and to safeguard the bot from losing money as market maker.

For example, suppose that in treatment E, the best asks for assets A and B are 4 and 3, respectively. The AP-bot will then post a sell order for an equal weighted ETF index asset equal to 7 in the ETF market. If a subject submits a bid order of 7 or higher for that ETF asset, then the bot sells the ETF at a price of 7, and simultaneously places bid limit orders for the constituent assets A and B at 4 and 3,

[7]Internet delay times are the most constraining: the bot is not collocated with the Flex-E-Markets market server. The bot is continuously informed about book updates in the underlying markets, so it can continuously change its own offer in the ETF market. If its purchase orders in the underlying markets are not executed because it submitted orders at stale prices, the orders become standing orders, and if after 15s the orders are still not executed, the bot will cancel the existing orders and submit new orders at the current best standing asks. This protocol ensures that the bot's inventory is almost never unbalanced, and hence, the effective supply of risks in the marketplace remains constant. While the bot is not making money, the protocol also limits losses as much as possible.

respectively. Thus, the AP-bot buys a unit of each asset A and B, and maintains a net zero asset position. The bot's profit is equal to zero since the price received for the ETF is equal to the price paid for the constituent assets.

For treatment U, using the same values of 4 and 3 for the best asks for assets A and B, respectively, the AP-bot will post an ask order equal to 10 $(4 + 3 \times 2)$ for a market-cap weighted ETF index share. When the ETF is purchased, then the AP-bot will submit bid limit orders for 1 unit of A and 2 units of B. The bot's profit can be negative if the current best ask in the order book is higher than 3 for B, and higher than 4 for A at the time trades are executed. In other words, if the orders for the constituent assets are no longer available at the original terms, which can occur if the orders are taken by someone else, or cancelled, then to maintain a zero net supply, the AP would need to cover its position on less favorable terms. However as noted already, due to the speed with which the bot operates, and its ability to monitor the books in real time (modulo internet communication delays), losses are minimal.[8]

3.2 Testable Hypotheses

In this section, we present testable hypotheses that we formulate based on previous empirical and theoretical work.

Our first hypothesis pertains to the demand for ETF shares. In our market, traders can replicate ETF shares individually by trading the constituent assets in the desired proportions. Nevertheless, we expect a significant demand for ETF shares since these assets offer a convenient way to acquire a more diversified portfolio. Such behavior is consistent with the net issuance of $ETFs$ as observed by ICI (2023) over the past 10 years, as well as with experimental work by Duffy et al. (2024). Purchasing ETF

[8] In Appendix D, we find that bot profit is zero in E, and negative, though very small, in treatment U.

shares instead of submitting individual orders for constituent assets not only reduces cognitive effort, it diminishes execution risk, ensuring transactions at desired prices, thereby reducing price uncertainty. Thus, we arrive at our first hypothesis:

Hypothesis 1. *There is a net issuance of ETF index assets in both treatments E and U.*

The second hypothesis concerns the prices of constituent assets across treatments. Since the ETFs are offset in ways that do not affect market risk, we expect the prices of the constituent assets to remain the same when ETFs are introduced in treatments E, U.

Hypothesis 2. *The prices of the constituent assets A and B are the same across treatments N, E, and U.*

Nevertheless, the convenience of holding $ETFs$ may potentially generate demand pressure for the constituent assets. When an ETF is sold, constituent assets are removed from the market (decreasing their supply), and this may put upward pressure on the prices of those assets relative to a market without ETF assets. However, the supply of constituent assets is not fully binding: in our design, individual traders can short-sell up to 5 units each of assets A, B or the $Note$. The ability to short sell assets should work to reduce price pressures from changes in the available supply of the constituent assets due to ETF purchases.

Beyond asset prices, we can also study other market measures such as *order imbalance* (i.e., the difference between the number of buy and sell orders) and *bid-ask spreads*. Specifically, we adopt a measure of order imbalance, where $d \in (0, 1]$ is a weight parameter that discounts orders more heavily the further they are away from the midpoint $m(t)$ between the current best bid and ask prices, and $Q(p, t) \in Z$ is the number of outstanding + sell (− buy) orders at price p at time t. The order imbalance is defined as

$$z(t,d) := -\sum_p Q(p,t)d^{|p-m(t)|}. \tag{2}$$

In our analysis we use $d = 0.99$ to give almost full weight to all serious orders and much lower weight only to extreme outliers.[9] We measure time t up to 1/10th of a second, in order to capture sudden jumps in the spreads.

In line with our hypotheses above, that the introduction of ETFs does not affect asset prices, we expect that:

Hypothesis 3. *The order imbalance and spreads for assets A and B are the same across treatments E and U, and equal to the baseline treatment N.*

A reasonable alternative hypothesis is that trade in $ETFs$ generates demand pressure on the constituent assets, causing more positive order imbalances and larger spreads in the underlying markets.

4 Laboratory Procedures

The experiment was conducted using Flex-E-Markets[10] and employed a within-subject design. We recruited subjects from all fields of study at the University of Melbourne.[11] In all, 83 subjects participated in the experiment. Each subject took part in a single session, which consisted of one practice market, and six 10 minute markets interchangeably referred to as "rounds" or "periods." The order of the 3 treatments in a session N, E and U varied according to Table 3, though we always started with the baseline

[9] Since the tick size is one cent, if the midpoint is 5.00, a buy order at 6.00 would be weighted at only $0.99^{100} = 0.366$.

[10] Flex-E-Markets is a SaaS platform where users can flexibly design marketplaces with multiple simultaneous public and private markets using the continuous open-book mechanism, and invite participants to trade directly through an online manual user interface or indirectly through robots that exploit the API. Developed by Quantahm, in Salt Lake City, UT (USA); see quantahm.com.

[11] Ethics approval was obtained from the university where the experiment took place, Project ID number 24368.

treatment of no $ETFs$ (N).[12] In all of our sessions, except one (IV) we have a balanced number of subjects that are endowed with either ($10, 4, 10) or ($60, 2, 2) units of cash, asset A and asset B, to incentivize trade.

In each session, subjects were given written instructions (see Appendix A) which were read aloud by the experimenter. A long practice session followed, to familiarize the participants with the trading interface and the marketplace setting. After a short break, 6 rounds were run with the order of treatments shown in Table 3 for each session. At the end of the experiment, one round and a state of the world were chosen

Table 3: Summary of sessions

Session	Order of treatments	Participants			
		Type 1		Type 2	
		Initial**	Count***	Initial	Count
I	N, U, E, U, U, E*	($10,4,10)	7	($60,2,2)	7
II	N, E, E, U, U, E	($10,4,10)	9	($60,2,2)	9
III	N, E, U, E, U, E	($10,4,10)	8	($60,2,2)	8
IV	N, E, E, U, E, U	($10,4,10)	9	($60,2,2)	8
V	N, U, U, U, E, E	($10,4,10)	9	($60,2,2)	9

Notes: **This round had an issue with the internet communication between the AP and the market server, and is therefore omitted from the analysis. ** Holdings: (cash, asset A, asset B); holdings of Note and ETF were always zero. Short-selling of up to 5 units of A, B and Note were allowed. *** Number of participants of the indicated type.*

randomly to assess participants' take-home earnings; payoffs were computed based on individual final holdings for the round, and payments were made in cash. Each session lasted approximately 150 minutes. Subjects' point totals from the chosen round were converted into Australian dollars (AUD) at the rate of AUD 5 per 10 points. Average earnings were AUD 48, which included a show-up reward of AUD 10.[13]

[12]This permutation of possible orderings is not exhaustive. The advantage of our within subject design is that the same participants are exposed to all three treatments, which means that any differences in response to the treatments can be more confidently attributed to the treatments themselves rather than to individual differences between participants.

[13]Payments were constrained to lie between AUD 40 and 60, per ethics committee request. Theoretical average expected earnings equaled AUD 50. Actual average earnings depended on states drawn.

4.1 Trading interface

Figure 1 shows the market user interface (UI). The left panel of the UI is divided into 2 parts. The top part of the left panel shows current asset holdings including cash. Unlike in field financial markets, settlement is immediate, meaning that once a trade takes place, holdings are updated. The holdings based on settlements are listed in the column "Settled." The column "Available" lists cash and asset units available after taking into account units committed in standing orders and potential shortsales. The bottom part of the left panel allows subjects to submit orders by selecting the number of units and price for their bid (ask) limit orders. The order form applies to the market that the user highlights in the right panel of the UI.

The right panel of the UI displays the 4 asset markets (A, B, ETF and $Note$). In each asset market, a subject can view the order book with standing bids (blue) and asks (red), as well as all past traded prices (trade history). Bids in the order book are sorted from highest to lowest and asks are sorted from lowest to highest.

Subjects submit limit orders by clicking on the button labeled "Place Buy order" ("Place Sell order"). A subject's submitted order is immediately posted to the order book for the relevant asset and is identified on their screen by a cross next to the order. Subjects can cancel their own order if it is not yet executed by clicking on the cross. Limit order can also be submitted by clicking on any standing order on the other side of the book. This will immediately cause a trade unless the standing order disappears in the meanwhile.

Transactions occur when a limit bid (ask) order comes in above (below) a standing limit ask (bid) order. Trade occurs at the price of the best standing ask (bid). All transactions are recorded and presented in the Trade History column. The transactions are color-coded: buyer-initiated trades are in blue, and seller-initiated trades are in

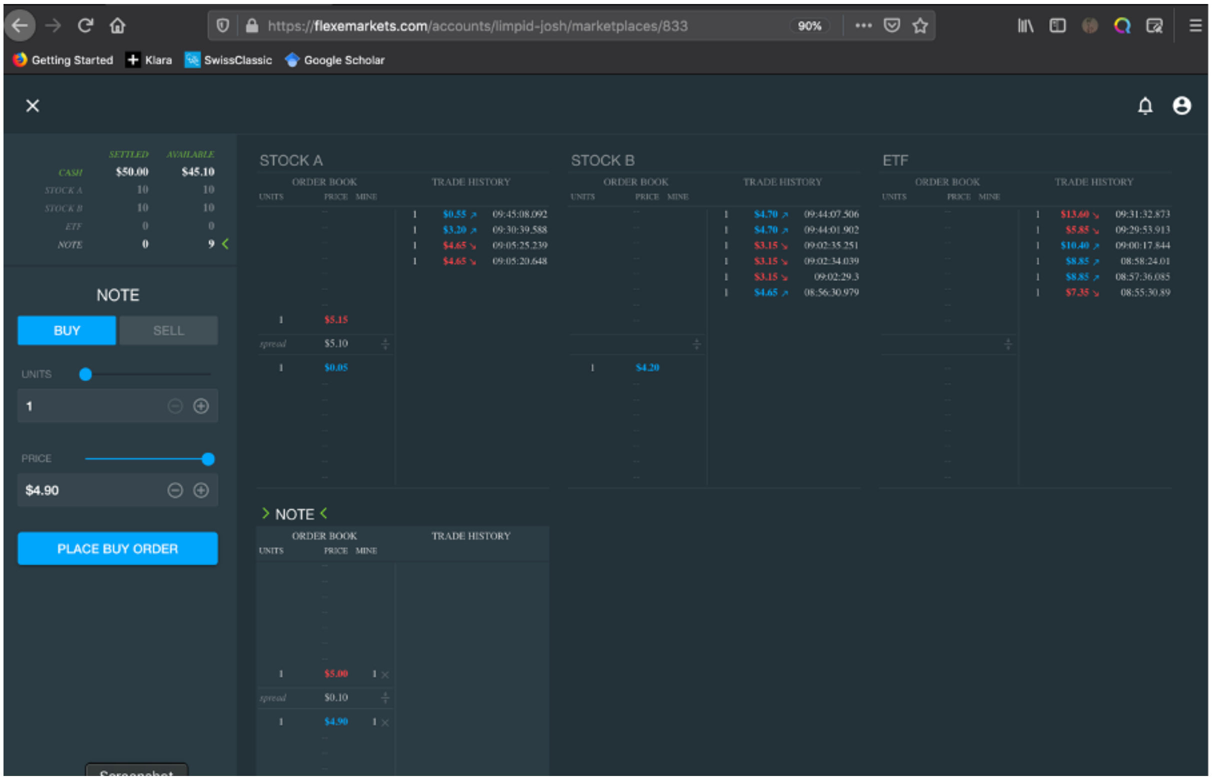

Figure 1: User interface (UI): *Buy (sell) orders appear in blue (red). There are 4 simultaneous markets for assets A, B, ETF and the Note. The AP posts asks and bids for the ETF using the current best bids and asks of the constituent assets. A subject submits limit orders by clicking on the market; selecting the number of units and the price, and clicking on the "Place order" button. Limit orders can also be submitted by clicking on the price level of a standing order on the other side of the book. The UI also presents information on the asset holdings including cash settled and available. See text and instructions (Appendix A) for further explanation.*

red.[14]

In all treatments, bids are rejected when they exceed the subject's cash balance. Subjects can short-sell up to 5 units of assets A, B, and the *Note*. If the final position in any of those assets is negative, then the subject's period earnings were adjusted downward by the value of the asset in the realized state after markets close. The experimenter covers any shortfall when a subject ends with negative final earnings due

[14] A buyer-initiated trade is one where a newly arriving bid causes a trade; conversely, a seller-initiated trade is one where the newly arriving order is a sell.

to shortsales. Parameters are chosen such that default is rare.

The ETF asset is a composite asset that is formed as a treatment-specific combination of the constituent assets. Prior to the start of each period of markets where ETFs can be traded subjects learn whether the ETF will have equal or unequal weights. When the market opens none of the subjects hold any $ETFs$ (and $ETFs$ cannot be shorted), however, the AP will offer units of $ETFs$ for sale as soon as other subjects enter ask orders for the required constituent assets (depending on the indexing strategy). Subjects can submit bids for the ETF at any point in the market, but they can submit asks for the ETF asset if only if they have previously bought a ETF unit.

5 Results

In this section, we report on the results from our experiment. We focus our analysis around the evaluation of our four main hypotheses.

Result 1. *There is a net issuance of ETF index assets in both treatments.*

The demand for ETF products is reflected in the large net issuance of $ETFs$; that is, the number of ETF asset shares created is larger than the number of shares redeemed. In each experimental session, following an initial round of trade without $ETFs$ (baseline), the market opens with a zero supply of ETF shares. In order for the first ETF to appear in the market, an ask needs to be submitted by the AP. Recall, that the AP will post an ask for an ETF share as long as there are existing asks for the constituent assets A and B. Thus, from the moment that there are offers in markets A and B, the AP can quote prices in the ETF market and ETF shares can be created and sold.

Table 4 presents the volume of transactions in the ETF market (value-based), per treatment, and whether the trade involved the AP bot or was strictly between

Table 4: *ETF* transaction volume across primary and secondary markets (value)

Treatment	Primary Market				Secondary Market
	Total	AP Bot-Buy	AP Bot-Sell	Net Issuance	Between-Participant
E	389,805	23,620	366,185	342,565	77,105
U	349,110	6,210	342,900	336,690	38,495
Total	738,915	29,830	709,085	679,255	115,600

"Humans." We refer to the former type of trade as the "primary market," and the latter as the "secondary market."[15] We find that 94% of transactions (366,185/389,805) correspond to ETF creations in the primary market for treatment E, and the number increases to 98% (342,900/349,110) in treatment U.

We also observe significant turnover of $ETFs$ in the secondary market: 77,105 (38,495) in treatment E (U). The difference in turnover rates across the 2 treatments can be explained by the difference in the expected values of the ETF shares across treatments: in treatment E, the expected price of the ETF asset is 1000, while in treatment U the expected value is 1500. The higher expected value in the latter requires higher liquidity if prices follow expected value. Since buy orders require cash availability in the account of the buyer, this constrains the ability to trade as frequently.

Thus, we can conclude that ETF shares are traded quite frequently, and that the majority of market participants employ a buy-and-hold strategy, confirming the prediction in hypothesis 1. In the next result, we examine hypothesis 2 and ask whether (and how) the prices of underlying assets change when $ETFs$ are introduced to the market under different indexing strategies.

Result 2. *The price of asset A in treatments E and U is greater than in the baseline*

[15] It should be noted that a buyer who purchases an ETF does not necessarily know whether the counterparty is an AP or another participant: the purchase could also be from another participant who bought a unit earlier in the round.

treatment N.

This result contradicts the prediction of Hypothesis 2, which states that there should be no difference in prices across treatments *N*, *E*, and *U*.

In Table 5 we present the mean prices observed across all sessions. When an *ETF* is introduced, in treatments *E* and *U*, prices of constituent assets appear to be significantly higher. Surprisingly, we find that prices of the constituent assets, as well as the prices of the *ETF* shares, are similar regardless of how whether we include all transactions in our analysis, or focus only on secondary markets. To test whether the presence of *ETF* assets affects the prices of the constituent assets, we run Fixed Effects (FE) regressions on the mean asset prices in a given round using dummies for: asset type (*B*, or *ETF*), treatment (*E*, or *U*), session (2 through 5), and all relevant interaction terms (please see Table 10 in Appendix *C*). We then conduct an appropriate *F*-test to determine significance. A FE regression using mean asset prices controls for the time-dependence within a session. Adding fixed effects at the session level allows us to control for the different number of participants across sessions.

Table 5: Asset prices

Treatment	*N*	*E*	*U*
All transactions			
A	509	581	572
B	481	515	534
ETF	–	1 103	1 653
Secondary market			
A	509	580	570
B	481	512	532
ETF	–	1 105	1 654

We begin by first comparing the price of asset *A*, which is the asset in short supply, when *ETFs* are not available in the market (treatment *N*) relative to the price in treatment *E*, which introduces an equal weighted *ETF* to market. Using an F-test, we find that the price of asset A is significantly higher in treatment E than in treatment N (p-value of 0.008, based on Table 10). In particular, the price of asset A is higher by about a magnitude of 70 when an ETF is available in treatment E. A similar analysis of prices in treatments N and U also finds that the price of asset A is significantly higher in U than in N by about 60, with a p-value of 0.003). In the case of asset B, we do not find a statistically significant difference in prices between treatments E and N (p-value of 0.402) or between treatments U and N (p-value of 0.451). The higher price of asset A relative to B is consistent with the CAPM theory, as discussed earlier, which suggests that asset B requires a higher risk premium, and should therefore sell at a lower price than asset A. We also do not observe any differences in the prices of the risky assets A and B between treatments E and U (p-value of 0.72 for asset A, and 0.79 for asset B).

In addition to studying the prices of the constituent assets, we also analyze whether the ETF prices follow the NAV. That is, for treatment E we check whether the ETF price is equal to the sum of the prices of assets A and B, while for treatment U, we check whether the ETF price is equal to the sum of the prices of 1 unit of A plus 2 units of B. We find that, indeed there are no arbitrage opportunities available, as the ETF prices are equal to the NAV across both treatments (p-value of 0.583 for treatment E and 0.663 for treatment U).

We next address Hypothesis 3, concerning the absence of spread and order imbalances for assets A and B across different treatments. Contrary to Hypothesis 3, we have the following result:

Result 3. *The bid-ask spread for the asset in short supply (A) is higher in treatment E compared to the baseline treatment N.*

Table 6 presents the average order imbalance and bid-ask spread observed for the assets in our experiment. We compute the order imbalance, or the difference between the number of bid and ask orders in the book, and spread of the orders, every 0.1 seconds, and then present the average for the last 8 minutes of trading. We observe a negative order imbalance for the constituent assets A and B (more outstanding asks), and a positive order imbalance for the ETF and risk-free asset N (more outstanding bids). The larger negative order imbalance for asset B might indicate a selling pressure on asset B, due to an insufficient risk-premium, or a market that has not yet settled.

Similar to the analysis of asset prices, we also compare the order imbalance and spread for each asset across treatments by using FE regressions and computing the relevant F-test (see Table 11 in Appendix C). We find that the order imbalance for the constituent assets is not significantly different in the E and U treatments relative to treatment N (for asset A, p-value of 0.161 in treatment E and 0.612 in treatment U; for asset B, p-value of 0.509 in E and 0.466 in U).[16] Note that the lower p-value for asset A in treatment E can be explained by the larger demand pressure generated by ETF creations, and the relative supply of asset A to asset B.

Table 6: Order imbalance and bid-ask spread

Treatment	N	E	U
Order Imbalance			
A	-3.47	-0.53	-2.70
B	-9.08	-7.96	-9.74
ETF	–	0.22	0.54
N	2.59	0.22	0.22
Bid-Ask Spread			
A	23.10	70.78	44.67
B	41.44	37.07	29.91
ETF	–	62.67	66.58
N	18.17	19.57	14.41

In terms of the bid-ask spread, we encounter a large increase in the spread for asset A in treatment E as compared to treatment N (p-value of 0.011), and no significant

[16]These results are consistent if we use the last 5 or 6 minutes of trading, instead of 8 minutes as presented above.

change in treatment U when compared with treatment N (p-value of 0.134). In treatment E, which requires equal shares of both assets A and B, there is more pressure on the price of asset A since it is in short supply relative to asset B. This leads to higher bids for asset A than in treatment U, where the indexing strategy resembles the relative market supply of assets. For asset B, we do not find significant differences in bid-ask spreads between treatments when compared with the baseline treatment N (p-value of 0.771 for treatment E and 0.616 for treatment U).[17]

6 Discussion and Conclusion

Our paper makes an empirical contribution toward understanding how ETF indexing strategy affects the prices of individual securities in the aggregate economy. Due to the convenience offered by ETF assets, there is high demand for these products and a positive net issuance of ETFs in our experiment. This is consistent with what we have seen in the field since the first ETF asset, the SPDR S&P 500 ETF, was launched in January 1993. As a result of this high demand, we observe that the price of asset A, which is in short supply, trades at a premium compared to an environment without ETFs. This price increase can be explained by standard asset pricing theory, which suggests that asset B should require a higher risk premium, and therefore sell at a relatively lower price (which it did). Surprisingly, we do not find significant differences of prices of A and B across the two types of ETFs. We would have expected greater price distortions in treatment E, since the optimal portfolio is value-weighted rather than equally weighted. The lack of significance may reflect a power issue. In our sessions, risk premia were relatively small, even in the base treatment N. The size of

[17] The large spread for asset A in treatment E compared to N is consistent when we use the last 6 minutes of trading (p-value of 0.003), but the result fades when using the last 5 minutes (p-value of 0.178). This difference in significance suggests that the buy-and-hold ETF strategy has a larger impact on the initial periods of trading compared to the latter.

the risk premium is cohort-dependent, and we happened to have selected relatively risk tolerant cohorts. If, however, the nature of ETFs is irrelevant for price distortions, the emergence of exotic weighting schemes in the field is nothing to be worried about per se.

Our results can motivate future theoretical work that incorporates the convenience offered by ETF assets, not only due to transaction costs but also due to cognitive costs since the ETF offers a simpler way to diversify a portfolio. Other experimental studies have also documented that the behavior of certain types of investors can have an indirect effect on prices because they change the per capita amount of risk that is shared among the marginal investors (see, e.g., Bossaerts et al., 2010 for the case of ambiguity-averse investors).

Future work can also explore alternative AP-bot algorithms, such as ones that ensure a positive bot profit. However, such algorithms have the disadvantage that they may require that the bot's inventory stay unbalanced for significantly longer periods. We experimented with profit-making bots and observed that a sale by a bot in the ETF market may take a minute or more to be completely offset in the two underlying markets. In the meantime, market risk is changed, since the bot will have created, through the ETF, a unit of asset A and up to 2 units of asset B, either or both of which are not offset by sales in the markets for A and B. This temporarily upsets the equilibrium pricing. Granted, such profit-making algorithms presumably may better reflect the actions of APs in the field. However, the temporary changes to equilibrium pricing stand in the way of a clean test of the presence of ETFs.

References

Ben-David, Itzhak, Francesco Franzoni, Byungwook Kim, and Rabih Moussawi, "Competition for Attention in the ETF Space," *Review of Financial Studies*, 2023, *36* (3), 987–1042.

Benartzi, Shlomo and Richard H Thaler, "Naive diversification strategies in defined contribution saving plans," *American Economic Review*, 2001, *91* (1), 79–98.

Bhattacharya, Utpal, Benjamin Loos, Steffen Meyer, and Andreas Hackethal, "Abusing ETFs," *Review of Finance*, 2017, *21* (3), 1217–1250.

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.

_ , Paolo Ghirardato, Serena Guarnaschelli, and William R Zame, "Ambiguity in asset markets: Theory and experiment," *Review of Financial Studies*, 2010, *23* (4), 1325–1359.

Brown, David C, Scott Cederburg, and Mitch Towner, "Dominated ETFs," *Available at SSRN 3694592*, 2021.

Chang, Yen-Cheng, Harrison Hong, and Inessa Liskovich, "Regression discontinuity and the price effects of stock market indexing," *Review of Financial Studies*, 2015, *28* (1), 212–246.

Davies, Shaun, "Index-linked trading and stock returns," in "Proceedings of Paris December 2021 Finance Meeting EUROFIDAI-ESSEC" 2024.

Duffy, John, Daniel Friedman, Jean Paul Rabanal, and Olga Rud, "The Impact of ETF Index Inclusion on Stock Prices," *Management Science*, 2024.

_ , Jean Paul Rabanal, and Olga A Rud, "The impact of ETFs in secondary asset markets: Experimental evidence," *Journal of Economic Behavior & Organization*, 2021, *188*, 674–696.

_ , _ , and _ , "Market experiments with multiple assets: A survey," in Füllbrunn Sascha and Ernan Haruvy, eds., *Handbook of Experimental Finance*, Edward Elgar Publishing, 2022, pp. 213–224.

Fernand, Larissa, "Equal weighted vs. Market cap weighted," *Morningstar India*, 2021.

Greenwood, Robin, "Short-and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage," *Journal of Financial Economics*, 2005, *75* (3), 607–649.

Hirshleifer, Jack and John G Riley, *The analytics of uncertainty and information*, Cambridge University Press, 1992.

Huang, Shiyang, Yang Song, and Hong Xiang, "The smart beta mirage," *Journal of Financial and Quantitative Analysis*, 2023, pp. 1–69.

Huberman, Gur and Wei Jiang, "Offering versus choice in 401 (k) plans: Equity exposure and number of funds," *Journal of Finance*, 2006, *61* (2), 763–801.

ICI, *Fact Book*, Investment Company Institute, 2023.

Jackson, Ryan, "The Strategic-Beta ETF Market Has Matured, but That Doesn't Mean the Party's Over," *Morningstar*, 2023.

Jiang, Hao, Dimitri Vayanos, and Lu Zheng, "Passive Investing and the Rise of Mega-Firms," *Available at SSRN 4851266*, 2024.

Magnani, Jacopo, Jean Paul Rabanal, Olga A Rud, and Yabin Wang, "Efficiency of dynamic portfolio choices: An experiment," *Review of Financial Studies*, 2022, *35* (3), 1279–1309.

Merlo, Lorenzo FP, "On the Impact of Thematic ETFs on Index Investing," *Available at SSRN 4048988*, 2022.

Pavlova, Anna and Taisiya Sikorskaya, "Benchmarking Intensity," *Review of Financial Studies*, 2023, *36*, 859–903.

Treynor, Jack, "Why market-valuation-indifferent indexing works," *Financial Analysts Journal*, 2005, *61* (5), 65–69.

Wang, Isabel, "This S&P 500 equal-weight ETF saw record inflows in 2023. Why the 'anti-momentum' fund is back in favor with investors," *Market Watch*, 2023.

Yuan, Ming and Guofu Zhou, "Why Naive Diversification Is Not So Naive, and How to Beat It?," *Journal of Financial and Quantitative Analysis*, 2023, pp. 1–32.

Online Appendix: Not for Publication

A. Instructions

Summary

Your task today is to trade up to four securities which will generate part of your earnings. You buy securities with cash and you get cash for securities. You will start out with some random endowment of the securities. Two of these securities generate payoffs that change with a randomly drawn state. The third security earns an amount that equals a fixed, pre-announced combination of the payouts on the first two securities. The last security is risk-free, meaning that it always pays the same, no matter which state is drawn.

You will repeat this task in a sequence of periods. Your take-home earnings will depend on your cash allocation at the end of one randomly selected period.

Setting

The entire session consists of **6** trading periods. Periods will last at most **10** minutes each; they can be shortened if trading subsides. There is a ~60-second break between periods.

We start the experiment by going over the instructions, familiarizing you with the trading interface, and running a trial period. After a break, the actual experiment takes place, where your actions impact your take-home pay. We spend quite a bit of time on instructions and practice, because we do not want you to be confused about anything.

There will always be a minimum of three securities. These securities expire at the end of a period, and pay a liquidating dividend based on the state, as noted in Table 1 below. The liquidating dividends on the securities you hold at the end of a period, plus your final cash, contribute towards your experiment earnings. You start the period with an endowment of cash and securities. This endowment may change from period to period.

Securities "**Stock A**" and "**Stock B**," pay a random dividend, determined by the random drawing of a state, "X," "Y," or "Z." The states are equally likely to occur. The third security is risk free (always pays the same amount) and is referred to as "Note." The table below lists the dividends for Stock A and B as a function of the state.

If State Is...	X	Y	Z
Stock A Pays	10	0	5
Stock B Pays	0	5	10
Note Pays	5	5	5

[Table 1]

Certain periods, there will be a fourth security, called "**ETF**." This security also pays a random dividend, but the dividend is simply a fixed combination of the dividends of Stocks A and B. The exact combination may change, but this will always be made clear before the start of the period.

We only consider two dividend combinations, referred to with the following terminology:

- Equal-Weighted ETF: 1 dividend of both A and B
- Unequal-Weighted ETF: 1 dividend of A, 2 of B

The Equal-Weighted ETF pays 1 dividend of A and 1 dividend of B. For instance, if the state drawn is Z, this then implies that the ETF would pay 1*5 + 1*10 = $15. Since A and B both pay $5 in expectation, the *ETF would be expected to pay 1*5 + 1*5 = $10.* With security ETF added, the payoff table would look as follows.

If State Is...	X	Y	Z
Stock A Pays	10	0	5
Stock B Pays	0	5	10
Equal-Weighted ETF Pays	10	5	15
Note Pays	5	5	5

[Table 2]

The Unequal-Weighted ETF pays 1 dividend of A and 2 dividends of B. For instance, if the state drawn is Y, this then implies that the ETF would pay 1*0 + 2*5 = $10. Since A and B both pay $5 in expectation, the *ETF would be expected to pay 1*5 + 2*5 = $15.* With security ETF added, the payoff table would look as follows.

If State Is...	X	Y	Z
Stock A Pays	10	0	5
Stock B Pays	0	5	10
Unequal-Weighted ETF Pays	10	10	25
Note Pays	5	5	5

[Table 3]

You will be allowed to *shortsell* the Note, Stock A and Stock B, but not the ETF. This means that you can take a negative position: you can sell a unit even if you do not own any and keep the purchase price. However, when the period ends, the dividend on the security that you sold short will be *subtracted* from your period earnings, times the number of security units you sold. E.g., if you shortsell one Note, since the final payoff on the Note is $5, this amount will be subtracted from your period earnings. Likewise, if you sold short 2 units of B and the state drawn is Z, then 2*10 = 20 will be subtracted from our period earnings. You will be allowed to shortsell up to 5 units.

You will face this setting for 6 periods. At the end of the experiment, we will randomly draw a period, and a state that determines the liquidating dividends for that period. You will be paid these dividends in accordance with your final holdings in that period and the state drawn, and your final cash holdings. This then constitutes your take-home pay, plus a sign-up reward.

All accounting is done in Experimental Dollars, which will be converted at the end to Australian Dollars at the rate of 1:2 (1 Australian Dollar per 2 Experimental Dollars).

The sign-up reward is always 10 Australian Dollars. In expectation, you earn 50 Australian Dollars. Since risk is involved (you do not know which state will be drawn in a period), we will make sure that your take-home pay is at least 40, but not more than 60 (Australian Dollars).

Trading Interface

The trading platform, Flex-E-Markets, works like a simple version of a real-world electronic stock market such as the Australian Stock Exchange (ASX).

You submit limit orders: orders to buy a chosen number of securities at a chosen price (or lower), or to sell a chosen number of securities at a chosen price (or higher). Transactions take place from the moment a buy order with a higher price crosses a sell order with a lower price. All trade occurs at the price specified by the best standing order. In other words, if a trade occurs, the price of the earlier best order determines the price. Orders at a better price execute first. Given a price, orders arriving earlier execute first.

Orders remain valid until cancelled or the marketplace closes. You cancel an order by clicking on the "x" next to the order in the book. The number next to the "x" indicates how many orders you have for a given limit price; you cancel one order at a time. See Figure below.

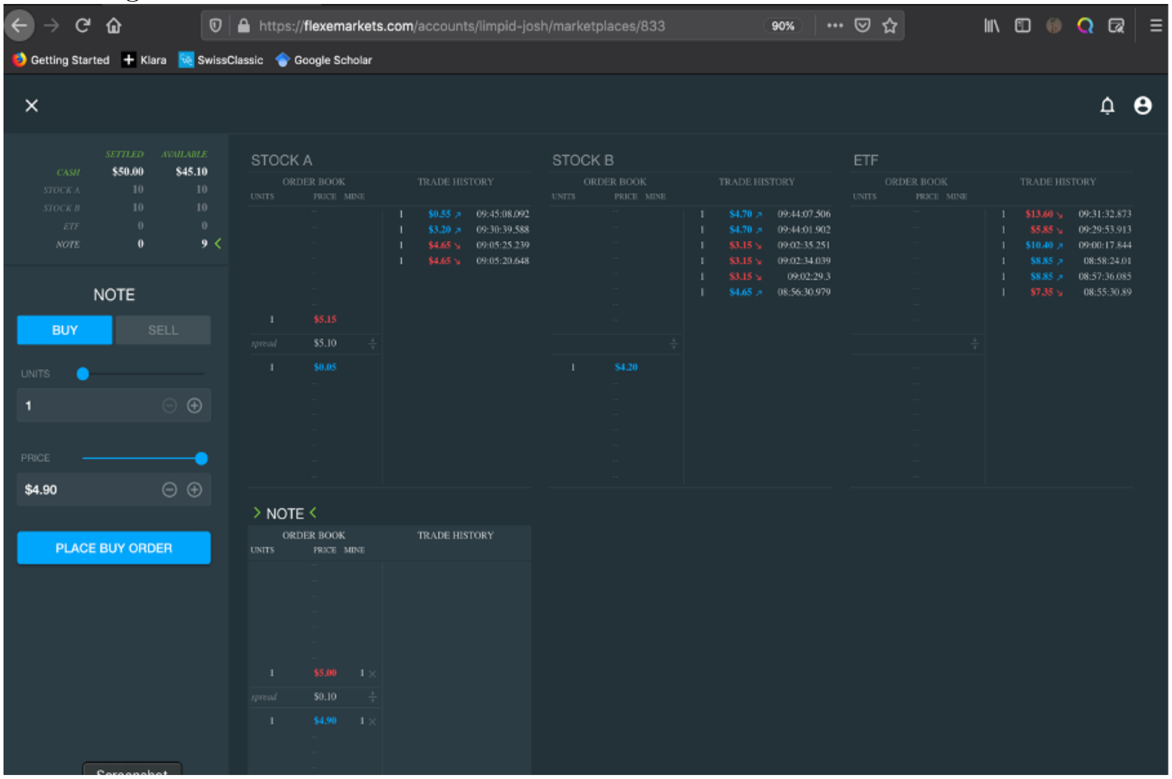

[Figure 1]

For each market, Flex-E-Markets shows you the book of all standing orders. Buy orders are in blue; sell orders in red; the "spread" is the difference between the best buy and sell orders. All orders are shown anonymously: you do not know who submitted them; nor will others know what your orders are.

The order form to the left allows you to submit orders, to buy or sell. When you click on the book of a market, say the "Note," the order form will apply to that market, and the name of the market is displayed in the order form. You adjust prices and quantities using a slider, before submitting your order.

There is one major difference between Flex-E-Markets and real-world stock markets. When you submit an order, the platform will check beforehand whether you have enough cash (for a buy order) or securities (for a sell order). The platform takes into account all your standing orders; it treats them as commitments (of cash or securities) even if they have not been executed yet.

Because of this, in the interface, you see two versions of your holdings: "settled'' and "available." "Settled" refers to quantities actually held. "Available" refers to cash/securities still available for further order submission. E.g., if your cash holdings equal \$50, and you submitted a buy order for \$4.90 (in the Notes market), the cash "settled" will equal \$50.00, while the cash "available" will be \$45.10 (= 50 – 4.90). If you submitted an order to (short)sell a Note, your Notes "settled" may show 0 (you own none), but your Notes "available" will show 4 (you can sell 5 units short, but already have one order to sell outstanding, for a total of 5 – 1 = 4).

More on ETF

Nobody starts with any units of security ETF. We employ a trading robot to facilitate trade, and who will introduce units of ETF to the market. If the robot sells you one unit of ETF, it will immediately cover the position by buying 1 unit of both Stock A and B (if the ETF is Equal-Weighted) or 1 unit of Stock A and 2 units of Stock B (if the ETF is Unequal-Weighted).

When others have acquired ETFs, you may be able to buy them directly from them, without robot interference. When you sell ETF, either the robot will buy back from you (and execute offsetting sales in the underlying markets), or others will buy from you. Since quotes and trades are anonymous in our trading system (Flex-E-Markets), you will not be able to tell whether you interact with a robot or with others in your cohort.

We use a sophisticated robot that is instructed to cover all positions it takes in the ETF market. For example, if a robot sells 6 units in the ETF market when the ETF is Unequal-Weighted, the robot is short 6 ETFs; to cover this position, the robot will try to buy 6 units of A and 12 units of B.

You do not have to worry about how the robot does this. Suffice it to say that the robot will be active in the markets for A and B from the moment it trades in the ETF market.

Please do ask questions when something is not clear! We are there to help you perform as best as you can.

Good Luck!

B. Session graphs

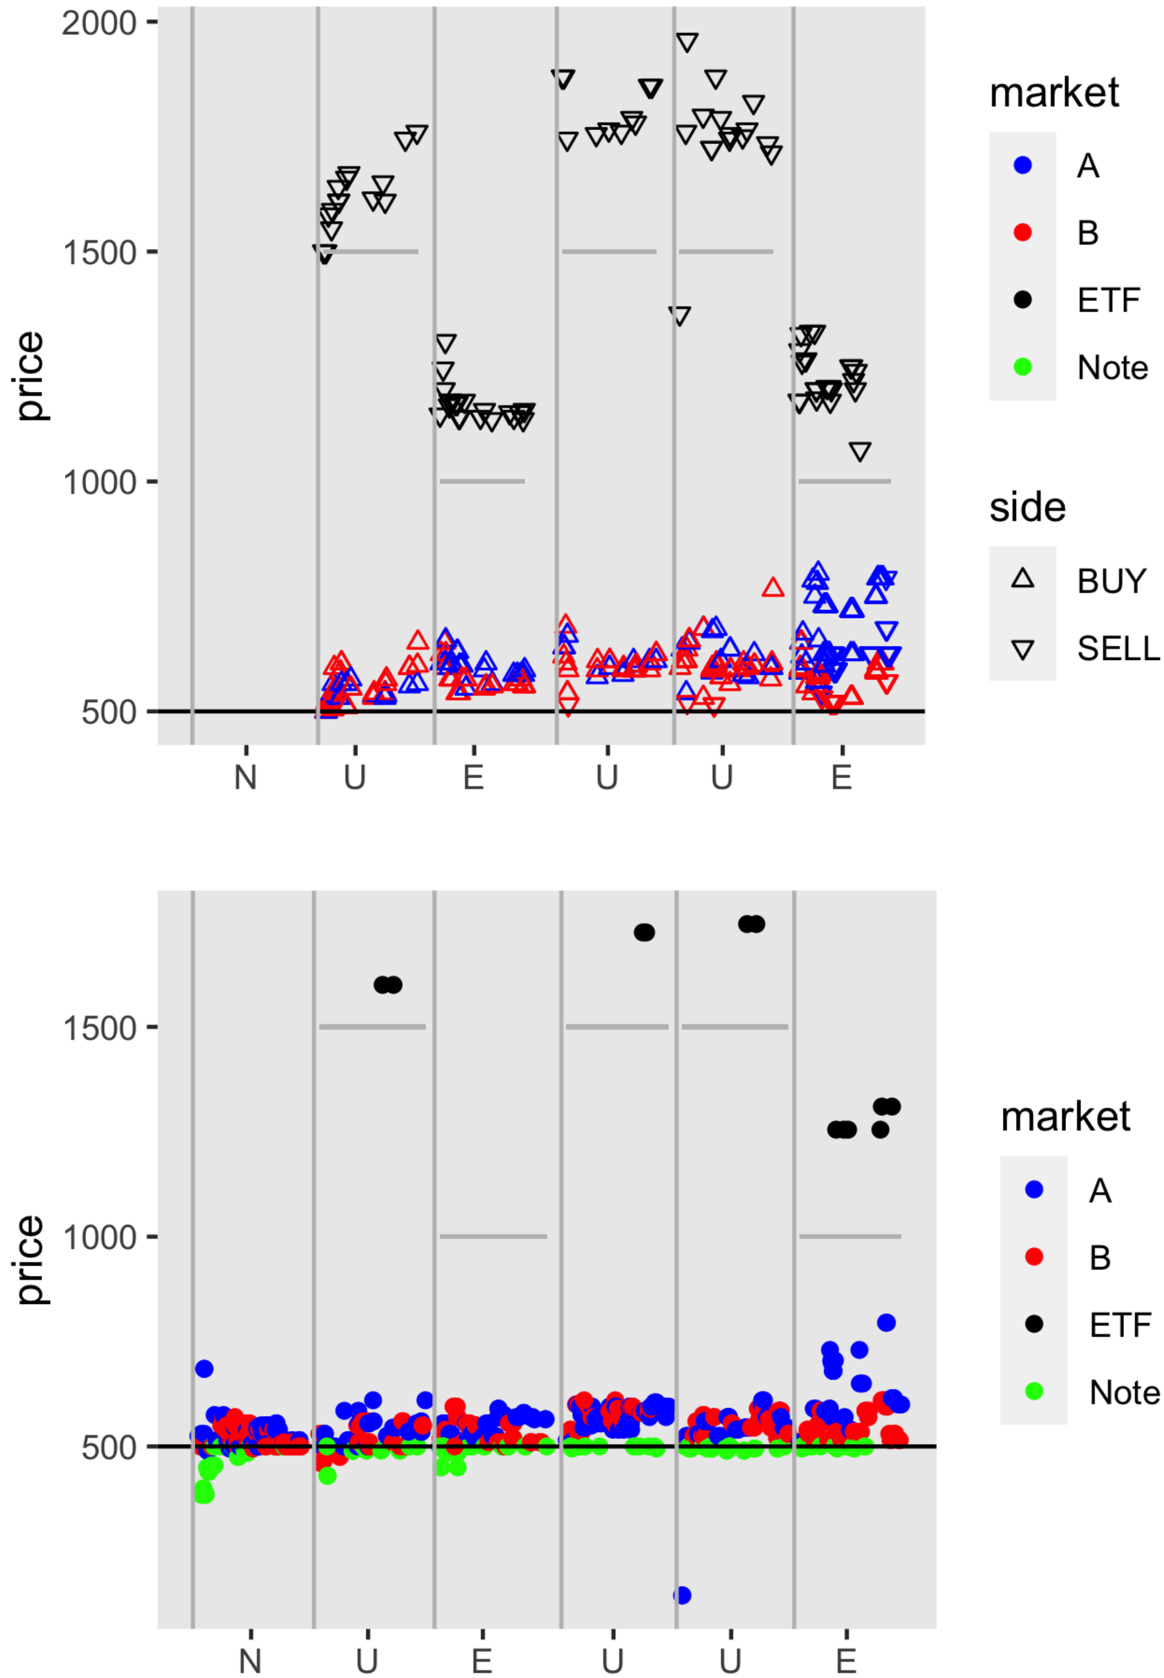

Figure 2: Traded price by bots (top) and humans (bottom) - Exp. 1

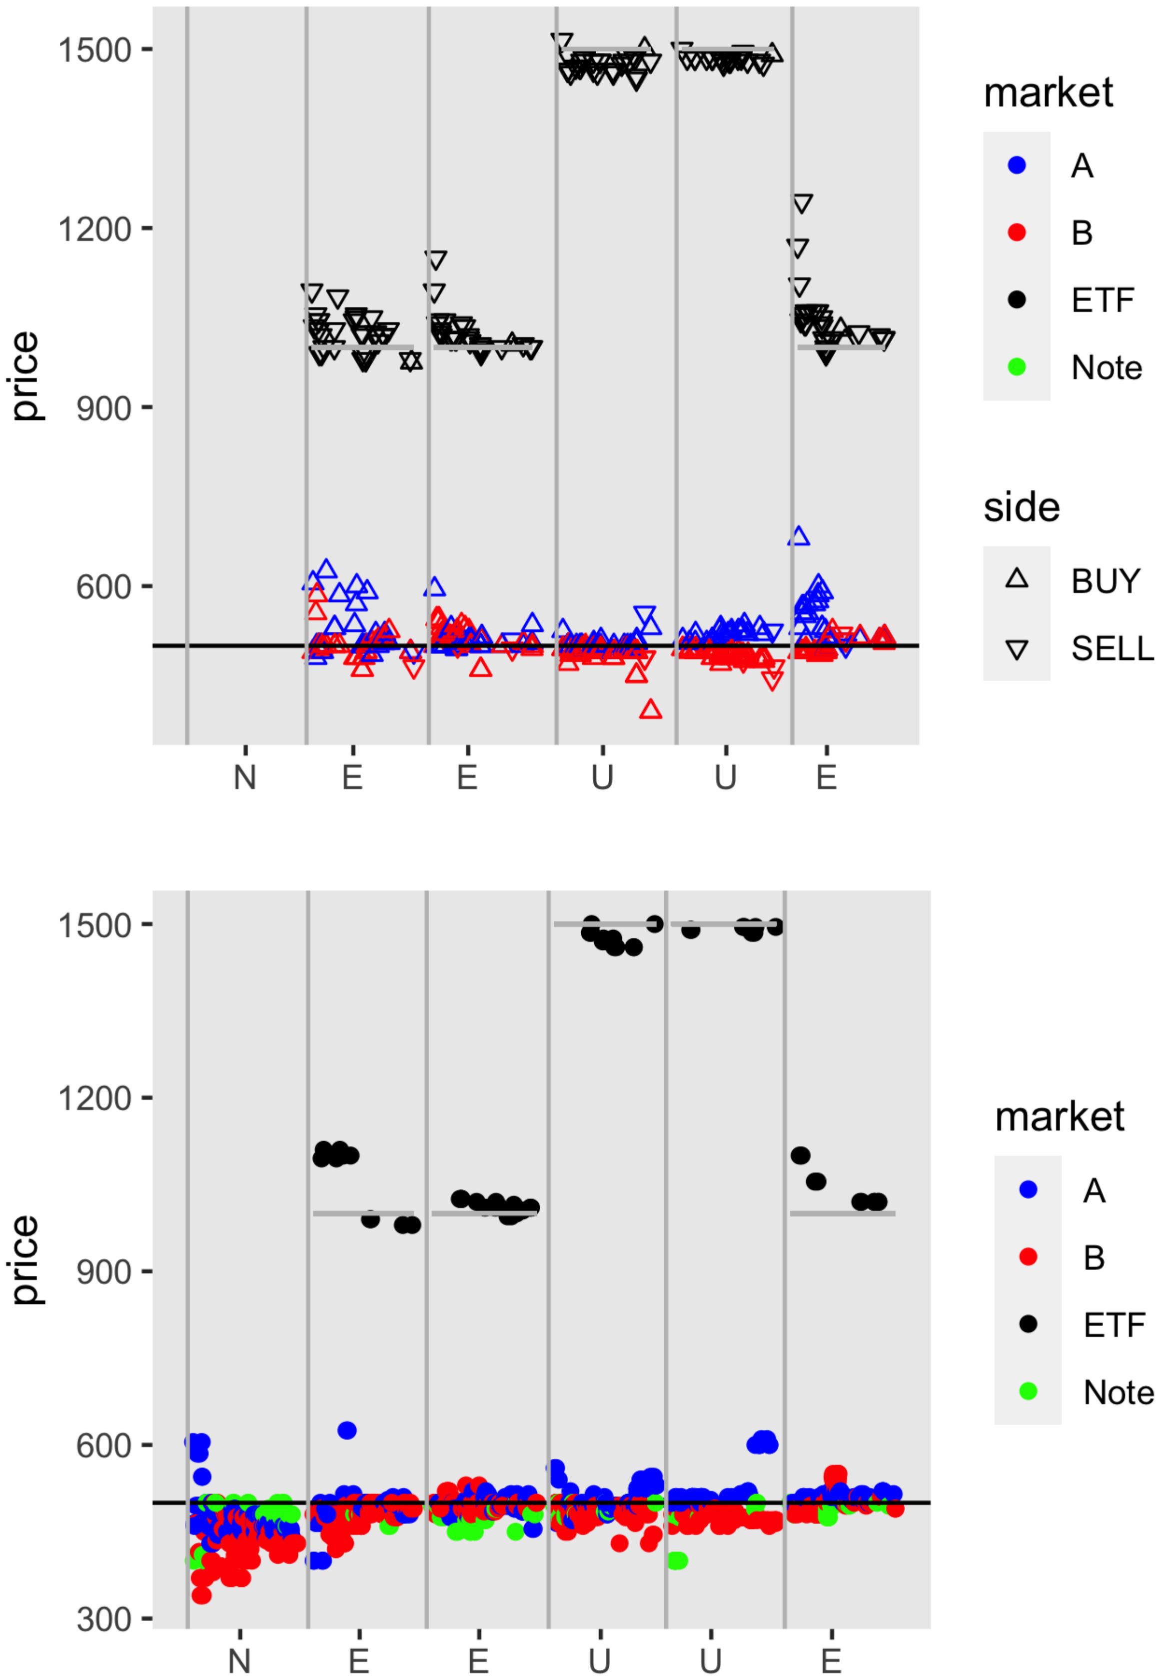

Figure 3: Traded price by bots (top) and humans (bottom) - Exp. 2

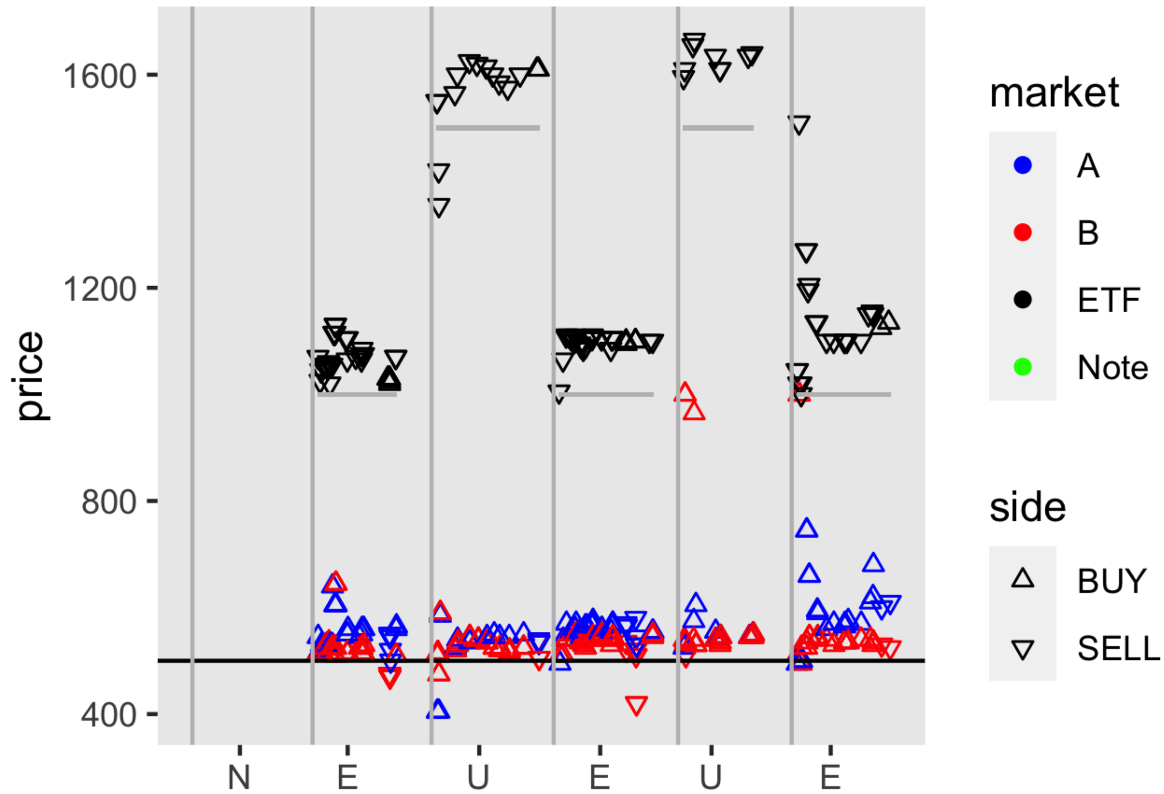

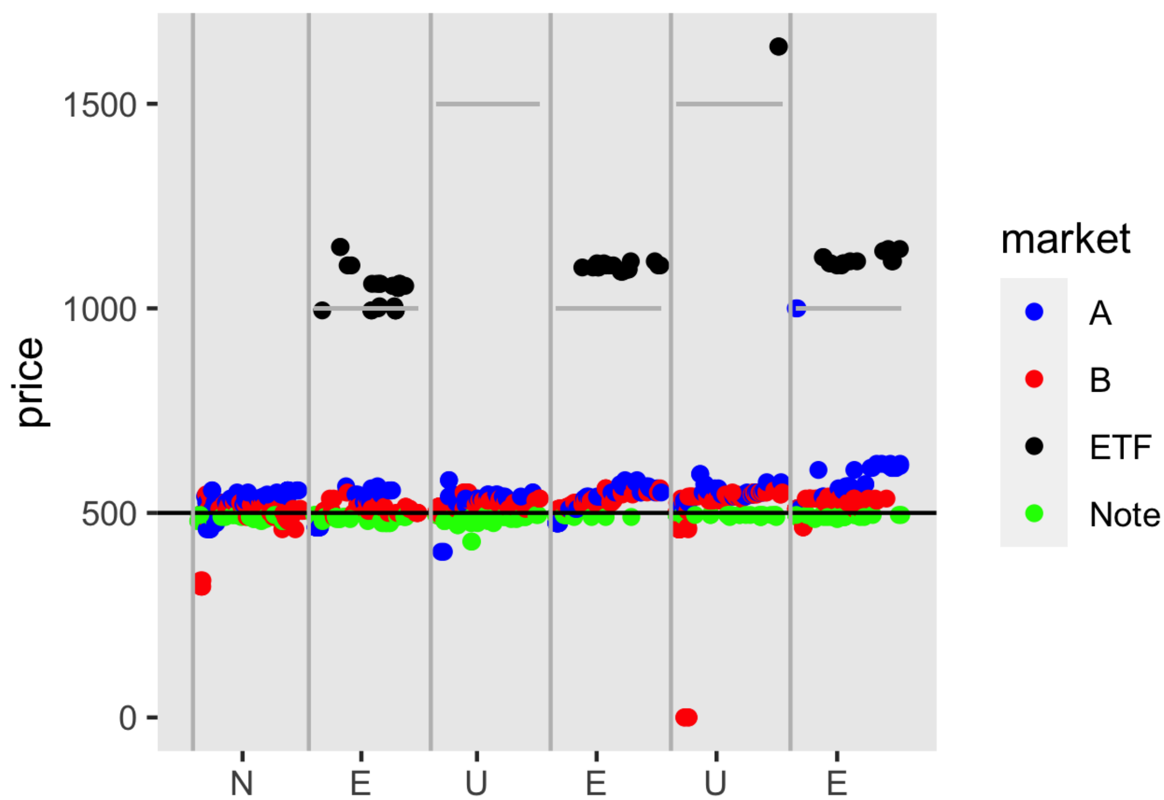

Figure 4: Traded price by bots (top) and humans (bottom) - Exp. 3

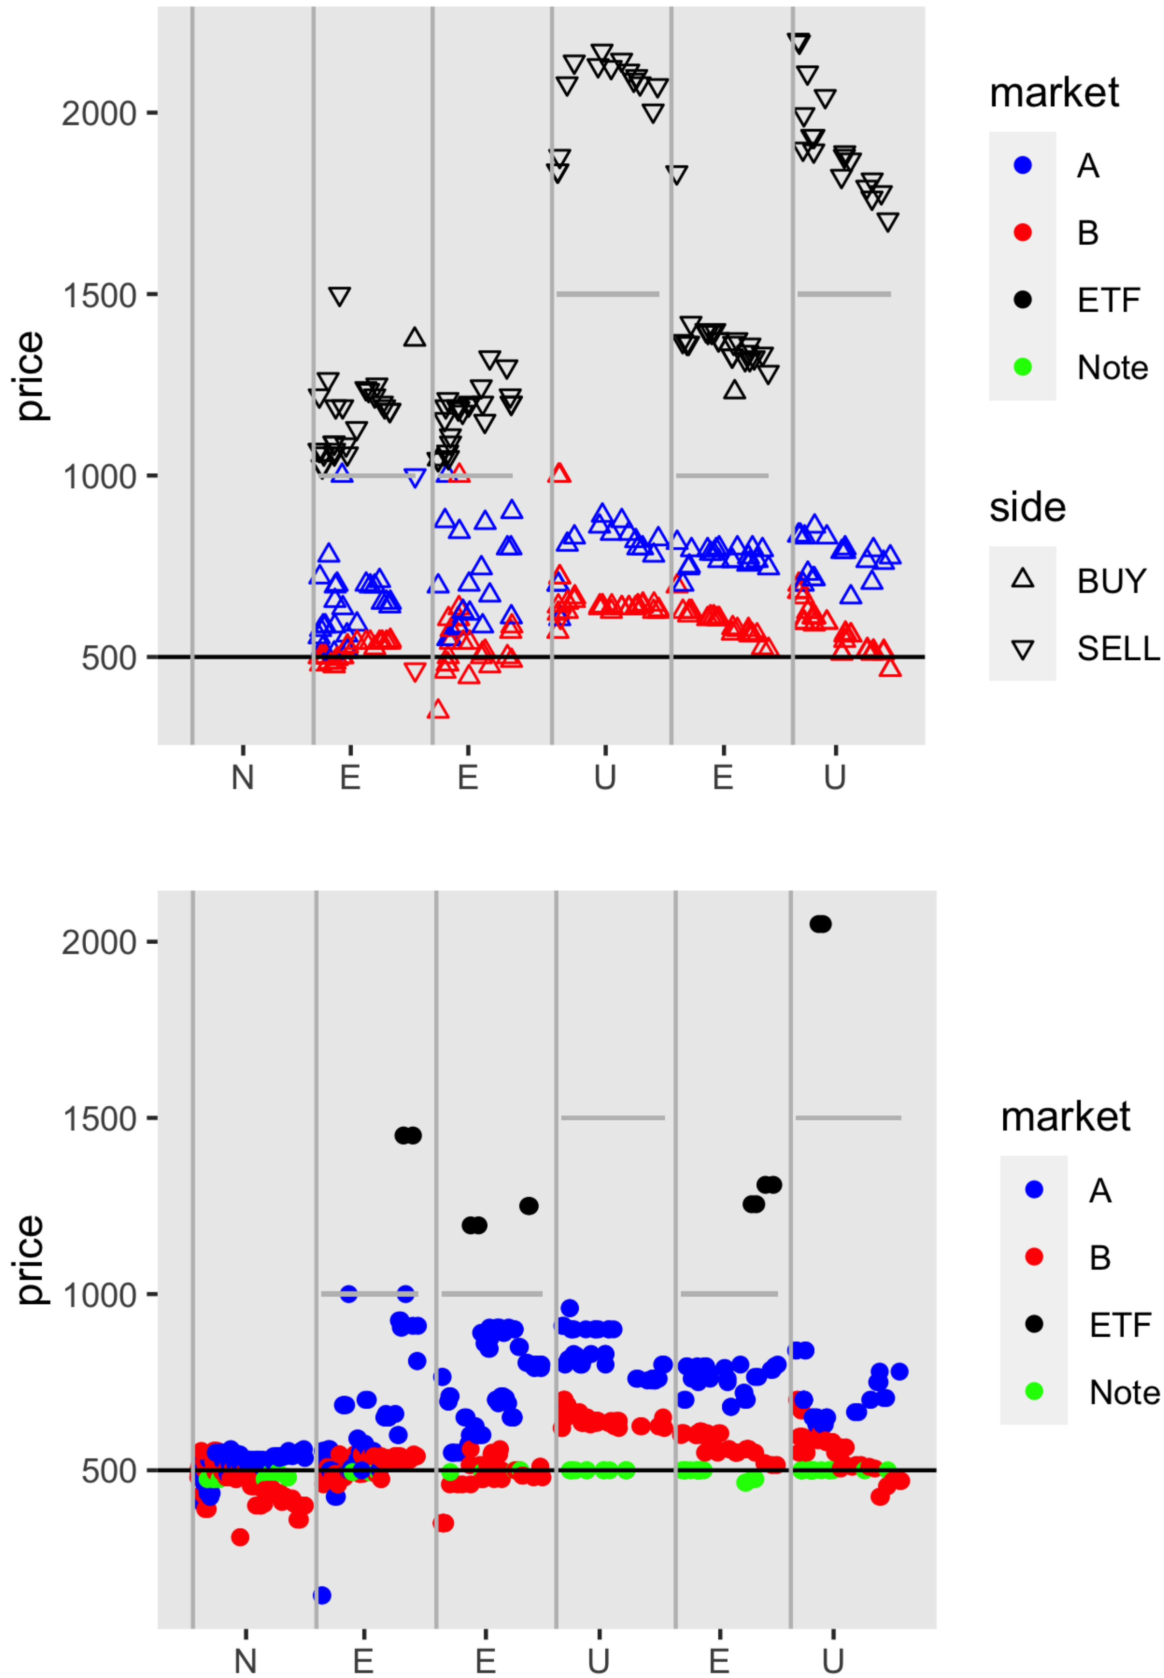

Figure 5: Traded price by bots (top) and humans (bottom) - Exp. 4

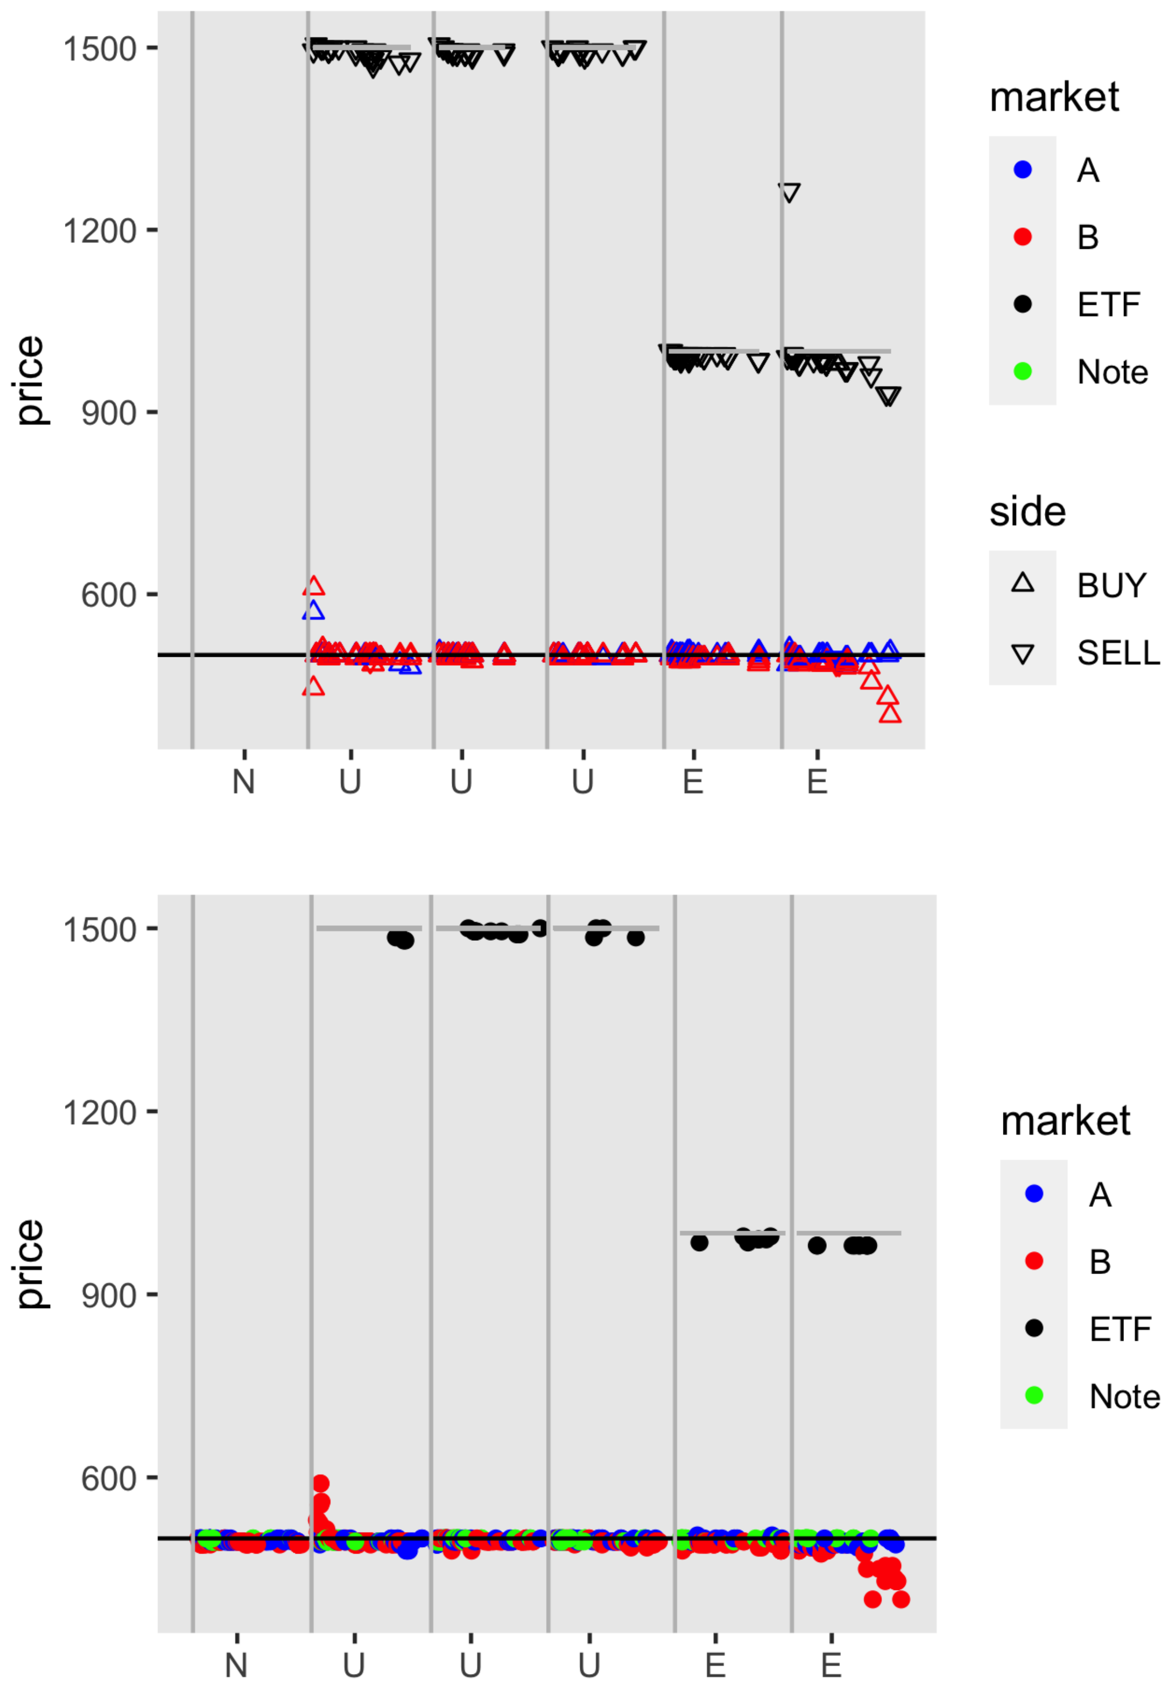

Figure 6: Traded price by bots (top) and humans (bottom) - Exp. 5

C. Regressions

Table 10: FE regression (Price)

	(1) price
Asset B	2.056 (0.972)
Treatment E	47.46 (0.426)
Asset B × Treatment E	-31.79 (0.705)
Asset ETF × Treatment E	581.9∗∗∗ (0.000)
Treatment U	50.94 (0.297)
Asset B × Treatment U	3.275 (0.962)
Asset ETF × Treatment U	1165.8∗∗∗ (0.000)
Asset B × Session 4	-83.67 (0.322)
Asset B × Session 3	-30.51 (0.717)
Asset B × Session 5	-4.212 (0.960)
Asset B × Session 2	-34.05 (0.686)
Treatment E × Session 4	148.7 (0.058)
Treatment E × Session 3	-17.73 (0.817)
Treatment E × Session 5	-46.69 (0.553)
Treatment E × Session 2	-0.811 (0.992)
Asset B × Treatment E × Session 4	-72.22 (0.507)
Asset B × Treatment E × Session 3	17.25 (0.874)
Asset B × Treatment E × Session 5	17.12 (0.878)
Asset B × Treatment E × Session 2	47.95 (0.659)
Asset B × Treatment U × Session 4	-112.0 (0.266)
Asset B × Treatment U × Session 3	14.28 (0.887)
Asset B × Treatment U × Session 5	-2.205 (0.982)
Asset B × Treatment U × Session 2	-5.484 (0.956)
Asset ETF × Treatment E × Session 4	-44.12 (0.521)
Asset ETF × Treatment E × Session 3	-56.16 (0.415)
Asset ETF × Treatment E × Session 5	-96.01 (0.192)
Asset ETF × Treatment E × Session 2	-72.87 (0.291)
Treatment U × Session 4	192.3∗∗ (0.009)
Treatment U × Session 3	-45.89 (0.518)
Treatment U × Session 5	-49.95 (0.468)
Treatment U × Session 2	1.498 (0.983)
Asset ETF × Treatment U × Session 4	43.32 (0.426)
Asset ETF × Treatment E × Session 3	-96.61 (0.080)
Asset ETF × Treatment E × Session 5	-172.0∗∗∗ (0.001)
Asset ETF × Treatment E × Session 2	-200.5∗∗∗ (0.001)
Session 4	11.30 (0.849)
Session 3	21.66 (0.716)
Session 5	-22.48 (0.705)
Session 2	-55.87 (0.350)
Constant	518.5∗∗∗ (0.000)
N	82
R^2	0.995

Notes: p-values in parenthesis.

Table 11: FE Regression (Order imbalance and spreads)

	(1) z	(2) spr
Asset B	-4.909	-2.353
	(0.349)	(0.953)
Treatment E	-4.649	2.077
	(0.375)	(0.959)
Asset B × Treatment E	5.341	11.67
	(0.470)	(0.836)
Asset ETF × Treatment E	0.342	10.32
	(0.948)	(0.796)
Treatment U	-2.928	18.02
	(0.493)	(0.581)
Asset B × Treatment U	3.799	8.975
	(0.529)	(0.846)
Asset ETF × Treatment U	-1.233	65.05**
	(0.682)	(0.007)
Asset B × Session 4	-0.268	23.73
	(0.971)	(0.675)
Asset B × Session 3	3.784	46.81
	(0.608)	(0.409)
Asset B × Session 5	-12.53	3.091
	(0.095)	(0.956)
Asset B × Session 2	5.429	29.83
	(0.463)	(0.598)
Treatment E × Session 4	9.651	139.9**
	(0.156)	(0.009)
Treatment E × Session 3	8.875	36.54
	(0.192)	(0.480)
Treatment E × Session 5	12.78	3.916
	(0.069)	(0.941)
Treatment E × Session 2	6.251	13.42
	(0.355)	(0.795)
Asset B × Treatment E × Session 4	-4.418	-136.3
	(0.643)	(0.067)
Asset B × Treatment E × Session 3	-9.867	-80.92
	(0.303)	(0.270)
Asset B × Treatment E × Session 5	-15.01	-4.423
	(0.129)	(0.953)
Asset B × Treatment E × Session 2	-7.824	-59.25
	(0.413)	(0.418)
Asset B × Treatment U × Session 4	-3.376	-119.6
	(0.701)	(0.081)
Asset B × Treatment U × Session 3	-4.894	-52.04
	(0.578)	(0.440)
Asset B × Treatment U × Session 5	1.117	-8.188
	(0.896)	(0.900)
Asset B × Treatment U × Session 2	-18.83*	-45.39
	(0.037)	(0.501)
Asset ETF × Treatment E × Session 4	0.421	-28.14
	(0.944)	(0.542)
Asset ETF × Treatment E × Session 3	-1.119	-24.57
	(0.853)	(0.595)
Asset ETF × Treatment E × Session 5	4.239	-3.950
	(0.508)	(0.936)
Asset ETF × Treatment E × Session 2	-0.482	-21.00
	(0.936)	(0.649)
Treatment U × Session 4	5.582	73.67
	(0.370)	(0.126)
Treatment U × Session 3	4.625	-13.73
	(0.457)	(0.773)
Treatment U × Session 5	4.443	-14.62
	(0.462)	(0.751)
Treatment U × Session 2	4.996	-20.49
	(0.422)	(0.667)
Asset ETF × Treatment U × Session 4	3.863	-84.44*
	(0.419)	(0.025)
Asset ETF × Treatment U × Session 3	5.296	-34.21
	(0.269)	(0.350)
Asset ETF × Treatment U × Session 5	11.46**	-58.42
	(0.010)	(0.079)
Asset ETF × Treatment U × Session 2	0.676	-48.11
	(0.887)	(0.191)
Session 4	-9.993	-17.87
	(0.060)	(0.655)
Session 3	-9.508	-26.49
	(0.074)	(0.508)
Session 5	-16.84**	-35.64
	(0.002)	(0.374)
Session 2	-6.411	-9.214
	(0.223)	(0.818)
_cons	5.077	40.90
	(0.173)	(0.152)
N	82	82
R^2	0.891	0.806

Notes: p-values in parenthesis.

D. Other results

Result 4. *The bot profit is zero (negative) in treatment E (U)*

Table 12: AP-bot trader profit

Treatment	Median Profit	Freq. of positive or zero profit
E	0	0.55
U	-185	0.14

Table 12 shows that the relative frequency with which the AP-bot makes a zero or positive profit is 0.55 (0.14) for treatment E (U). We test whether the profit is zero in treatment E using a Wilcoxon test, and find that we cannot reject that it is equal from zero (p-value of 0.36 for or treatment E and 0.06 for U). Note that the median profit in treatment U is negative, though the value is small (185).

Lower bot profits in treatment U can be explained by the mechanism used to post ETF ask prices, and the higher prices paid for the constituent assets. The bot may purchase a second unit of asset B at a price that is higher than the first unit, if the same price is not available (recall, the ETF ask price is based on best asks for the constituent assets, and therefore uses the best ask for B for both units), creating small losses.[18]

[18] To explore whether this mechanism was exploited by subjects, we study the value of asks submitted by investors across treatments U and E. Using a Wald-test, we fail to reject that asks submitted are the same between the two treatments (p-value of 0.42). Therefore, the mechanism which determines the ETF price seems to be the most appropriate explanation for the small losses observed in treatment U.