



Munich Personal RePEc Archive

## **Arbitrage bots in experimental asset markets**

Angerer, Martin and Neugebauer, Tibor and Shachat, Jason

29 June 2019

Online at <https://mpra.ub.uni-muenchen.de/96224/>  
MPRA Paper No. 96224, posted 28 Sep 2019 09:35 UTC

# Arbitrage bots in experimental asset markets\*

Martin Angerer<sup>a</sup>, Tibor Neugebauer<sup>b</sup> & Jason Shachat<sup>c</sup>

- a. University of Liechtenstein
- b. University of Luxembourg
- c. Durham University and Wuhan University

*Abstract:* While algorithmic trading robots are a proliferating presence in asset markets, there is no consensus whether their presence improves market quality or benefits individual investors. We examine the impact of robots seeking arbitrage in experimental laboratory markets. We find that the presence of algorithmic arbitrageurs generally enhances market quality. However, the wealth of human traders suffers from the presence of algorithmic traders. These social costs can be mitigated as we find high latency algorithms harm investors less than low latency algorithms; while the improvements in market quality are indistinguishable between algorithm latency levels and whether they provide liquidity or not.

**Keywords:** asset market experiment, arbitrage, algorithmic trading

**JEL Codes:** C92, G12

\*We acknowledge funding from Economic and Social Research Council (ESRC) and Luxembourg National Research Fund (FNR).

## 1. Introduction

There has been proliferation of algorithmic trading systems (subsequently abbreviated with AT) on electronic exchanges.<sup>1</sup> Kirilenko and Lo (2013) report that after 2009 the trading volume in the Dow Jones Industrial Index doubled in less than three years, accelerating well beyond the previous rate of doubling every seven years. Most of this increase is due to the use of AT. Carrion (2013) examined a sample of NASDAQ data in which 68 percent of the total trading volume could be assigned to AT activity. Hendershott, Jones and Menkveld (2011) reported high frequency trading, a subset of AT, accounted for 73 percent of US trading activity in 2009. According to estimates of Glantz and Kissell (2013), AT-activity is responsible for 85 percent of equity trading volume in 2012, up from 15 percent in 2003. New market structures and anomalies<sup>2</sup> arising amid the predominance of AT presents have opened a host of research questions.

One such question is whether AT activity improves or deteriorates market efficiency. Researchers hold generally two diverging views.

(i) The first view suggests that ATs can be benevolent market makers like designated dealers and specialists in traditional markets. Hendershot, Jones and Menkveld (2011) use the rate of “electronic message traffic” (including submissions, cancellations and transactions) on the NYSE as a proxy measure for AT activity. They suggest that, for large stock in particular, ATs increase market liquidity and enhance the pricing efficiency. Hasbrouck and Saar (2013) study two NASDAQ order level

---

<sup>1</sup> Kirilenko and Lo (2013) define algorithmic trading as the automation process of buying and selling orders of securities via sophisticated mathematical models, high performance computers and telecommunication networks.

<sup>2</sup> "The flash crash" on 6 May 2010, when the Dow Jones Industrial Average dropped for a short time by 998.5 points only to recover half an hour later almost to its initial value, brought the phenomenon of automated trading systems, in particular high frequency trading, into the public's eye. Easley, Lopez de Prado and O'Hara (2011, 2012a, b), O'Hara (2014, 2015), Kirilenko, Kyle, Samadi and Tuzun (2017), Menkveld and Yueshen (forthcoming), Andersen, Bondarenko, Kyle and Obizhaeva (2016) discuss market microstructure and other causes of the flash crash.

data samples to report positive impacts of low latency algorithm traders on bid-ask spreads, order book depth, short-term volatility as well as the overall price impact of trades. Brogaard, Hendershott and Riordan (2014) examine the trading behavior of 26 “independent proprietary” high frequency trading firms within a NASDAQ data set. They find that high frequency ATs provide the best bid and offer quotes for a significant portion of the trading day. They suggest that high frequency ATs reduce volatility, aid the price discovery process and generally enhance market quality.

(ii) The second view about AT activity in electronic exchanges suggests that ATs, and in particular high frequency ATs, are predatory, rent-seeking on public information or by manipulating open books. Foucault, Hombert and Rosu (2016) point out that ATs operate with short-term time horizons and thus do not contribute information to long-term price discovery. Jarrow and Protter (2012) find that high frequency ATs can create mispricing that is disadvantageous to the slower traders. Biais, Foucault, Moinas (2015) model fast trading and find two equilibria; the first one excludes high frequency trading, the second equilibrium implies the coexistence of traders of different speeds in the market. Finally some papers discuss legal and illegal predatory strategies including stuffing, smoking and front-running of buy-side traders (Biais and Woolfen 2011; Easley et al. 2012a; O’Hara 2014; Clark-Joseph 2014; Jarrow and Protter 2015; Baron et al. 2018).<sup>3</sup>

Public knowledge about ATs’ strategies is generally limited as the strategies of (‘black box’) ATs are proprietary.<sup>4</sup> Our focus is the impact of algorithmic arbitrageurs, as we will sometimes call arbitrage bots/ATs, on market quality.

---

<sup>3</sup> Aldrich and Friedman (2018) suggest to peg orders to protect them from exploitation by high frequency traders.

<sup>4</sup> Hagströmer and Nordén (2013) point out that a suitable categorisation of AT’s is difficult as respective firms do not generally use one strategy and will dynamically modify strategies in response to market conditions. Some studies (e.g., Glantz and Kissell 2013; Brogaard, Hendershott and Riordan 2014) provide broad AT classifications.

Research has identified the presence of arbitrage ATs in financial markets. Chaboud, Chiquoine, Hjalmarsson and Vega (2014) present data on a triangular arbitrage opportunity in currency exchange. ATs have an advantage over humans when the euro-yen exchange rates are out of line with dollar-yen and euro-dollar exchange rates. Chaboud et al. report a reduction in arbitrage opportunities associated with liquidity taking arbitrage bots. They suggest this improvement in market quality imposes an adverse selection cost on slower traders. Menkfeld and Yueshen (forthcoming) report on arbitrage in fragmented markets where cross-asset arbitrage effectively connects buyers and sellers. The profitability of arbitrage bots depends on their search and transmission speed, i.e., its latency relative to the latency of the other traders (e.g., Carrion, 2013; Hasbrouck and Saar, 2013; Brogaard et al., 2014; Biais, Foucault, and Mfoinas, 2015; Budish, Cramton and Shim, 2015; O’Hara, 2015; Wah, 2016; Baron, Brogaard, Hagstromer and Kirilenko, 2018; Brogaard and Garriott, forthcoming).<sup>5</sup>

In the paper, we experimentally explore the market impacts of alternative ATs that seek riskless arbitrage in a fragmented market. The ATs differ by either providing liquidity by submitting limit orders or by absorbing liquidity by exclusively submitting market orders. Through experiments with concurrent human and AT participation, we address the following questions:

- (i) Can an arbitrage AT induce the law-of-one-price and generally eliminate any mispricing in the experimental asset market?<sup>6</sup>

---

<sup>5</sup> Latency arbitrage opportunities are plentiful according to Wah (2016). Wah measured in 2014 across 11 US exchanges 69 arbitrage opportunities per security per day in 495 securities of the S&P 500, which existed for 0.87 seconds; a \$3 billion market for high frequency traders, according to her estimates.

<sup>6</sup> Mispricing has two dimensions in this study; (i) price discrepancies across twin shares and (ii) deviations from fundamental dividend value. The former point reflects the fact that arbitrage opportunities happen in real time, and the elimination of arbitrage opportunities does not guarantee one

- (ii) What is the impact of arbitrage bots on volatility?
- (iii) How does the latency of arbitrage bots impact pricing efficiency?
- (iv) What is the differential impact for liquidity providing and liquidity absorbing on total market liquidity?
- (v) What are the arbitrage AT's share of transactions, and what gains do they reap from human subjects?

We approach these questions adopting an experimental asset market design for twin-shares trading with perfectly correlated cash flows (Charness and Neugebauer 2019). Price discrepancies between the two twin-shares offer an arbitrage opportunity, but as Charness and Neugebauer report, human subjects fail to generally extinguish arbitrage opportunities. In this study, we introduce algorithmic arbitrageurs that detect and exploit arbitrage opportunities. In total we consider five treatments: (1) The “Baseline” treatment replicates Charness and Neugebauer with solely human participants; (2) The “NoBot” treatment adds a statement that an algorithmic trader may participate in the market to the instructions of Baseline treatment; (3) The “FastBot” treatment incorporates a low latency liquidity taking AT; (4) The “SlowBot” treatment incorporates a high latency liquidity taking AT; and (5) The “LiqBot” treatment incorporates a liquidity providing AT. Note that our FastBot and SlowBot treatments investigate differences in nominal latency, and not in relative latency (as suggested in Biais et al. 2015, and Baron et al. 2018). Our Baseline and

---

average price. The latter point refers to the traditional view that market equilibrium requires that the no-arbitrage condition holds (e.g., Harrison and Kreps 1979). Shleifer (2000) states the traditional view in the context of efficient markets as follows (p. 4); “the process of arbitrage brings security prices in line with their fundamental values even when some investors are not fully rational and their demands are correlated, as long as securities have close substitutes.” In our setting such an impact of arbitraging is not straight forward.

NoBot treatments allow us to check if simply announcing the potential presence of ATs impacts market outcomes (Farjam and Kirchkamp 2017; Leal and Hanaki 2018).

Our results suggest that arbitrage bots enhance market quality but do not establish price efficiency. Arbitrage AT participation generates a regression towards the law-of-one-price. This effect is greater when the arbitrage bot provides liquidity rather than only absorbing liquidity. The effect on mispricing, however, is not great. Prices are not systematically closer to fundamental dividend value under AT participation than without. Given that we observe no reduction of volatility through AT participation suggests that pricing efficiency is unaffected beyond cross-asset price effects. We have mixed results vis-à-vis liquidity improvement. In the case of the liquidity providing arbitrageur, we observe more limit order submissions and more frequent bid-offer spread than in the other treatments. Surprisingly we find no immediate effect on liquidity by the presence of liquidity absorbing arbitrage bots, but speed matters for transaction volume. The number of AT involvements in transactions positively impact AT's gains from trade. When the AT responds slowly, the market quality enhancement is still present in the data, but the adverse trading costs to subjects are relatively low due to the reduced number of transactions.

Our study also contributes to two modest streams of the experimental asset market literature on (i) arbitrage and (ii) AT participation.

(i) The literature on arbitrage in laboratory experiments includes O'Brien and Srivastava (1991) and Abbink and Rockenbach (2006) who look at subjects' skills to choose replicating portfolios of options and stocks. Rietz (2005) reports on a prediction-market experiment with state-contingent claims, where arbitrage opportunities are easily spotted. Charness and Neugebauer (2019) conduct an asset market experiment with perfectly correlated twin-shares. Both studies suggest that

subjects fail to exploit arbitrage opportunities in the laboratory.<sup>7</sup> Bossaerts, Shachat and Xie (2018) study the drivers of arbitrage opportunities in a one-asset setting. Bossaerts et al. report that more competition and the existence of higher endowed traders are factors that reduce arbitrage opportunities whereas relaxing margin requirements or restrictions on short sales do not.

(ii) We provide a partial review of the nascent literature using experimental methods to study ATs in financial markets. Rietz (2005) is the only experimental study prior to ours where an algorithmic arbitrageur was used to automatically eliminate each price discrepancy. Rietz reported that the automatic arbitrageur was involved in most trades in the experiment, but he did not report on the impact of this activity on market quality. Farjam and Kirchkamp (2017) and Leal and Hanaki (2018) assess mispricing in experimental asset markets when the potential presence of an algorithm is announced, but when actually no algorithm participates in the market. The studies report mixed results on the announcement effect. Farjam and Kirchkamp find that mispricing of human subjects is reduced when they expect a trading algorithm in the market. However, Leal and Hanaki fail to reproduce this effect in a very similar setting. In Aldrich and López Vargas (forthcoming) subjects choose and tune predefined market maker or sniper ATs, and then decide on costly improvements in algorithm latency. They compare two alternative market mechanisms, the frequent batch auction (FBA) and the continuous double auction (CDA). Aldrich and López Vargas find the FBA is less prone to predatory trading behavior than the CBA, it implies lower wasteful investments in low-latency communication and lower transaction costs, and it has lower volatility in spreads and liquidity. Asparouhova, Bossaerts, Rotaru, Wang, Yadav and Yang (2019) study single-asset market behavior

---

<sup>7</sup> Oliven and Rietz (2004) report a considerable share of arbitrage opportunities, which market participants fail to exploit, also outside of the lab in a three months long prediction market (i.e., the IOWA election market).



when subjects choose algorithmic strategies or manual trading. Their results suggest that mispricing and volatility are not enhanced by AT usage.

The paper is organized as follows. In section 2, we present the details of the experimental design, and in section 3 we discuss testable hypothesis. Section 4 presents the results of our study, and section 5 concludes.

## **2. Experimental design**

We present our experimental design by discussing the following in sequence: asset structures and cash, trader endowments, the continuous double auction, the different trading algorithms, and the experimental procedures.

### **2.1 Economic environment**

Following Charness and Neugebauer (2019), we implement arbitrage markets in the Smith et al (1988) formulation. This setting has three goods: cash (numerated in experimental currency units or ECU's), asset A and asset B. All three goods live for ten periods. ECU's neither pay dividends or accrue interest. Units of each asset pay a dividend in each of the ten periods and then expire without a redemption value. Each period the units of asset A generates a common ECU dividend that is a realization from random variable with four equally-likely outcomes,  $\{0, 8, 28, 60\}$ . The ten dividends are independent random draws. Therefore the expected value of single dividend is 24 ECU, and if there are  $t$  remaining periods the expected value of the future stream of dividends is  $24t$ . Asset B also pays a dividend each period. The dividend value of a unit of asset B is always equal to that of asset A plus 24 ECU. Thus the dividends of assets A and B are identical modulo a shift, thus perfectly correlated. Further the expected value of a single dividend of asset B equals 48, and if

there are  $t$  remaining periods the expected value of the future stream of dividends from holding a unit of asset B is  $48t$ . The law-of-one-price asserts, in period  $t$ , the difference between prices of assets A and B is  $24(11 - t)$ .

In our setting there are nine human traders. Prior to period one, each receives an endowment: 1300 ECU's, and four units each of assets A and B. We impose the following leverage purchase and short sale constraints. At any point in time a trader's holding of cash must exceed -2600, and the holdings of each asset A and B must exceed -8. Note when a trader holds a negative amount of an asset and a dividend is generated, the trader pays, rather than receives, the dividends on those units. We assume traders derive utility solely from their final holdings of ECUs at the end of period ten, after all dividends have been paid. Each trader has full information regarding these preceding elements. If traders are risk neutral and seek to maximize the expected value of their terminal ECU holdings, then the respective expected fundamental dividend values of assets A and B, and corresponding prices in a rational expectation equilibrium, are respectively  $24t$  and  $48t$ , for  $t = 1, \dots, 10$ .

## **2.2 Continuous double auction**

Each period, prior to the determination and payment of dividends, traders can buy and sell units of the assets in markets for the two assets. There is a *continuous double auction market* (CDA) for each of the assets. A CDA is open for a fixed length of time in which traders may generate publicly observable quotes which can lead to bilateral trades. Traders can take four types of actions. The first two are limit orders. A *limit bid*,  $B^j$ , is an amount of ECUs at which the trader is willing to purchase a unit of asset  $j$ . A *limit ask*,  $A^j$ , is an amount of ECUs a trader is willing to accept to provide a unit of asset  $j$ . These limit bids and asks are publicly displayed in the *order book*.

Limit bids are listed from highest to lowest, and limit asks are listed from the lowest to highest. The lowest limit ask and the highest limit bid define the bid-ask spread. A trader may submit multiple limit bids (asks) for an asset as long as their holdings of the asset (cash) does not fall below -8 units (-2600 ECUs).

There are two other actions a trader may take: *market buys* and *market sells*. A trader submits a market buy when she wishes to purchase a unit at the lowest limit ask in the order book. This generates a transaction in which the trader submitting the market buy and trader who submitted the lowest current ask trade at that ask price. Similarly, a trader submits a market sell when she wishes to sell a unit of the asset at the highest limit bid in the order book. This generates a transaction in which the trader submitting the market sell and the trader who submitted the current highest bid trade at the bid price. Note that whenever a transaction occurs the involved limits order(s) are deleted from the order book. Further any remaining limit asks for the seller – and limit bids of the buyer - are also deleted from the order book of the asset. We forbid traders from submitting market orders that would trade with one of their own limit orders. We clear the order books for both assets when the trading period expires. Screen shots and instructions are appended to the paper.

### **2.3 Algorithmic trading algorithms**

We consider two types of algorithmic trader robots. Each seeks out riskless arbitrage utilizing only the same information available to the traders. One type of algorithmic trader robot only makes market orders, absorbing market liquidity. We call these *Market Order Robots*. The other type generates arbitrage by creating limit orders in one market that, if accepted, can be paired with a market order in the other market for a certain profit. This type of robot provides market liquidity. We call this type *Limit*

*Order Robot* or *LiqBot*, short for *Liquidity Order Robot*. In any pair of markets for assets A and B, there is at most one robot trader.

## 2.4 Market Order Robots

Market Order Robots exploit the bid-ask spreads that violate the law-of-one-price in view of the perfect correlation of dividends. In trading period  $t$  the law-of-one-price dictates that price of asset B is equal to the price of asset A less twenty-four times the number of remaining dividends; i.e.,  $P_B - P_A = 24(11 - t)$ . There are two cases in which market fragmentation allows for a riskless arbitrage using simultaneous market orders. First, when the highest limit bid for asset B exceeds the lowest limit ask for asset A by more than  $24(11 - t)$  then market order pair sell B and buy A yields a certain increase in cash flow equal to  $B^B - A^A - 24(11 - t)$ . Second, when the difference between the lowest limit ask for asset B and the highest limit bid for asset A is less than  $24(11 - t)$ , then the market order pair Sell A and Buy B results in a certain increase in cash flow of  $24(11 - t) - (B^B - A^A)$ . We implement two variants of the Market Order Robot which differ in their speed market order execution.

A *FastBot* executes the required market order pair immediately when an arbitrage opportunity arises. Effectively the Fastbot reaction time is quicker than any human trader to the extent human traders do not get to observe these arbitrage opportunities. A *SlowBot* executes the required pair of market orders, conditional upon the arbitrage opportunity still existing, only after a total of four human generated market actions are taken upon an arbitrage opportunity arises. We implement this SlowBot to allow human traders the opportunity to exploit arbitrage opportunities.

## 2.5 Limit Order Robots

A LiqBot monitors the order books for both assets. When a human generated limit order arrives in one market it formulates a corresponding limit order (including an arbitrage premium) in the other market. For example, if a *human* trader submits a limit bid for A, say  $B_h^A$ , then the LiqBot formulates a corresponding limit bid for B,  $B_{Liq}^B$  that would generate a certain profit if a human traders accepts it with a market sell and at the same time the LiqBot accepts  $B_h^A$  with a market sell. In order to ensure that this is an arbitrage gain,  $B_{Liq}^B$  is determined as

$$B_{Liq}^B = B_h^A + 24(11-t) - \varepsilon$$

where  $\varepsilon$  is both the minimal acceptable size of the arbitrage gain and a random variable distributed uniformly over the interval  $[0, \frac{24(11-t)}{2}]$ . Thus the LiqBot seeks an arbitrage no more than one-half the difference in the fundamental dividend values which can be as small as zero. The LiqBot only submits  $B_{Liq}^B$  if it is an enhancing bid, i.e., if  $B_{Liq}^B$  exceeds all outstanding bids for asset B, with a delay of one human generated market action. The LiqBot immediately cancels its bid,  $B_{Liq}^B$ , as soon as its corresponding human bid,  $B_h^A$ , is either accepted with a market sell or cancelled. If a human trader accepts the LiqBot's limit order,  $B_{Liq}^B$ , with a market sell, the LiqBot immediately accepts by market sell the best outstanding bid in the market for asset A, thus, generates at least an arbitrage gain of  $\varepsilon$ .

When a human trader submits a limit ask for asset A,  $A_h^A$ , then the LiqBot formulates a limit offer for asset B,  $A_{Liq}^B$ , according to

$$A_{Liq}^B = A_h^A + 24(11-t) - \varepsilon$$

A similar logic as before prevails for the submission and cancellation of  $A_{Liq}^B$ . Note, that a new  $\varepsilon$  is drawn independently for each human trader limit order. The LiqBot follows the symmetric process for human traders' limit orders for asset B.

## 2.6 Experimental procedures

We conducted all of our experiments at the experimental economics laboratory in the Newcastle University Business School. We recruited subjects through e-mail invitations from random selection from a pool of economics and science students at Newcastle and Northumbria universities via ORSEE (Greiner 2004). A subject could participate in only one experimental session. A session consisted of the following timeline which lasted approximately two and one-half hours.

1. Informed consent,
2. An investment task to elicit individual risk attitudes,
3. A Cognitive Reflection Task (CRT) to elicit individual propensities for Level 2 thinking,
4. A public reading of the market instructions including a quiz to ascertain understanding of the asset and dividend structures, and two separate three-minute practice rounds with the CDA,
5. A sequence of three iterations of market of ten trading periods – new endowments each iteration. Trading periods lasted 180 seconds in the first iteration, and 120 seconds in the last two iterations<sup>8</sup>,
6. A random die roll by one of the subjects determined which of the three market iterations contributed to the subjects' earnings,
7. Subjects completed a debriefing questionnaire, and
8. Subjects were privately paid a £3 show-up fee + earnings from the investment task + earnings from the CRT task + earnings from the randomly selected market iteration. In case of a negative final cash balance, not including the show-up fee, a subject's payoff would be zero.

---

<sup>8</sup> We allowed more time in the first market for people to get accustomed to the decisions. There is no evidence (including questionnaire reports) that subjects were short of time in the shorter intervals.

The investment task was introduced by Charness and Gneezy (2010) to provide a simple and intuitive assessment of an individual's degree of risk aversion. A subject chose an amount  $0 \leq X \leq \text{£}10$  to allocate to a risky asset that paid with equal probability 0 or  $\text{£}2.5X$ , and to a safe asset  $\text{£}10 - X \geq 0$  to be paid out with certainty.<sup>9</sup> We randomly selected one of the nine participants in the asset market to receive the payoff from their investment decision.

The second task was the CRT (Frederick 2005), which consisted of three questions asked in a random order.<sup>10</sup> Subjects had 90 seconds to answer the questions and were rewarded with  $\text{£}1$  per correct answer. These questions are designed to separate whether the responder adopts Level 1 thinking (quick response without reflection) or Level 2 thinking. We developed these tasks and our CDA implementation using the software ztree (Fischbacher 2007).

## 2.7 Experimental treatments

Our experiment design varies the presence and type of algorithmic trading robot, as well as the instructions presented to subjects about the potential presence of a trading robot. Also we adopt a between-subject design in which a group of nine traders experience the same robot trader presence/type in all three market iterations. In total we have five treatments: Baseline, NoBot, SlowBot, FastBot and LiqBot (see Table A1). In all treatments except Baseline we included the following *announcement* in the instructions, of potential AT presence,

---

<sup>9</sup> Note that the payoffs in the investment and CRT tasks were expressed in British Pounds. We introduced the ECU in the instructions for the market trading tasks, which we distributed only after the Investment and CRT tasks.

<sup>10</sup> (1) A hat and a suit cost \$110. The suit costs \$100 more than the hat. How much does the hat cost? If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets? (3) In a lake, there is a patch of lily pads. Every day the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of it?

*“Besides the participants in the room, a computerized trading algorithm may participate in the market. The computerized algorithm may take the same actions as you. It can buy and sell in the market. The details of the strategy followed by the algorithm are not revealed to you, and you will not be informed if the computerized trading algorithm actually acts in the market or not.”*

The Baseline treatment allows us to account for any “announcement” effect invoked by mentioning the potential participation of algorithmic trading robots in the market iterations.

### **3. Measures and hypotheses**

To assess the impact of the algorithmic arbitrageur, we measure market performance in terms of equilibration, liquidity and volatility, and arbitrageur's gains from trade.

#### **3.1 Equilibration**

A benevolent view on arbitrageurs in markets is that they establish the law-of-one-price and further enhance the equilibration process. In fact, arbitrage is made in real time by selling high and buying low. Obviously, arbitrage does not immediately imply the same average prices. Charness and Neugebauer (2019) report that despite the fact that prices were the same on average under perfect correlation, significant differences in average prices persisted throughout the experiment. We address the question of whether the presence of the arbitrageur implies a decrease in the average price discrepancy between assets. Following Charness and Neugebauer, we measure the cross-asset *price discrepancy* as follows;

$$(1) \quad PD = T^{-1} \sum_{t=1}^T \left| \frac{P_t^B}{P_t^A + (F_t^B - F_t^A)} - 1 \right|$$



where  $P_t^i$  and  $F_t^i$  denote the average price and the fundamental dividend value of securities A and B in period  $t$ . This measure will inform us about the impact of arbitrage on the law-of-one-price.

A more subtle question is how arbitrageurs impact mispricing in the market. Mispricing is frequently measured by the Relative Absolute Deviation from fundamental dividend value (Stoeckl et al. 2010).

$$(2) \quad RAD^i = (\bar{F}^i T)^{-1} \sum_{t=1}^T |P_t^i - F_t^i|$$

Here,  $\bar{F}^i$  is the average fundamental dividend value of asset  $i = \{A, B\}$  over the  $T = 10$  periods of the asset market. We expect that the presence of the AT and the speed of the AT positively impact equilibration on fundamentals and support the law-of-one-price.

**Hypothesis 1.** *PD* and *RAD* will be lower in presence of algorithmic trading.

### 3.2 (II-)liquidity measures and volatility

A benevolent view is that ATs enhance market liquidity. We consider both liquidity providing and taking arbitrageur ATs. Takers might negatively impact liquidity rather than enhance it. To check the impact, we measure liquidity at the end of each period via spread, transaction volume, and order flow.

*Spread, transaction volume, and order flow*

A common liquidity measure in financial economics is the spread between the offering and the bidding price, shortly we refer to the (percentage) *spread*,  $S_t$ , at the

end of the period. The spread is the difference between the best outstanding offer and the best bid relative to the bid-ask midpoint; see equation (4) in the Table 1. Here  $p_{At}$  is the best asking price and  $p_{Bt}$  is the best bidding price outstanding in period  $t$ , the denominator accounts for the price level. The spread is a measure of illiquidity; the larger the spread, the lower the liquidity. It measures the loss incurred by simultaneously buying and selling a share.

A second measure of liquidity is the transaction volume, which we denote by  $V_t$ . It measures the number of shares transacted during a period, i.e., it informs us about the speed at which the equity capital of a company is turned over, and the number of periods we need to sell a share or all the shares in the company. A third measure of liquidity is the number of order submission in a period, sometimes referred to as order flow, which we denote by  $\zeta_t$ . The higher the transaction volume and the higher the order flow, the more liquid is the market.

Volatility is liquidity correlated. In our setting we must take account of the declining fundamental value when we measure volatility. We measure volatility on logarithmic changes in the price-value ratio (see the following table). The logarithmic price-value change is as follows;

$$\ln \Delta \frac{P_t}{F_t} = \ln \frac{P_t}{F_t} - \ln \frac{P_{t-1}}{F_{t-1}},$$

where  $F_t$  denotes the fundamental dividend value of the asset, and  $P_t$  is the average price in period  $t$ . In case of a missing price-value ratio in any period, we reduce the number of measure-periods correspondingly.

**Table 1: Liquidity and volatility measures**

---

(3) (Volume) number of transactions,

$$V_t, \zeta_t$$

number of orders

(4) (bid-ask) spread

$$S_t = \frac{P_{At} - P_{Bt}}{(P_{At} + P_{Bt})/2}$$

(5) Volatility

$$\sigma = \sqrt{\text{VAR}(\ln \Delta \frac{P_t}{F_t})}$$

---

Our *testable hypothesis* is that liquidity will be impacted in our FastBot and SlowBot treatments relative to the NoBot treatment, because the AT absorbs liquidity and contributes to the transactions in the market. A priori we would expect no impact on order flow, but conjecture a positive impact on transaction volume. Given the liquidity absorption of our AT, a spread increase could be the outcome, but that is not clear. Hence, we suggest the following.

**Hypothesis 2.** Transaction volume, volatility and liquidity are impacted by AT.

### 3.3 Arbitrageur gains and exposure

In efficient markets, arbitrageurs make no gains, but the question is relevant. How much can an arbitrageur skim the returns of market participants? In our treatments, we can measure potential and actual arbitrage gains. We can also measure the arbitrageur's exposure, i.e., the credit capacity including short and long positions the arbitrageur requires to implement the strategy. We measure the arbitrageur's exposure by the maximum short position in cash and shares. Assuming more transaction

activity of the AT in the FastBot than in the SlowBot treatment, we expect the following.

**Hypothesis 3.** Gains and exposure of the arbitrageur are impacted by speed.

#### 4. Results

We first evaluate whether the announcement that an AT is participating in the market, but actually does not, impacts market behavior by comparing our treatments Baseline and NoBot. Second, we examine pricing efficiency across treatments by examining the *RAD* and the *PD* measures. Third, we investigate if the alternative ATs affect our liquidity measures. We conclude by examining the portfolio positions of the ATs.

An individual cohort is our unit of analysis. In Table 2, we present the number of cohorts for each treatment. In addition, we present the total number of participants per treatment (8 or 9 each in a cohort), the proportion of female participants, the average number of correct CRT responses, and the average proportion of tokens invested in the risky option in the investment task. In total, there are 40 cohorts (8 per treatment) involving 270 participant. Participants earned £22 on average. We recruited participants by sending invitation letters to randomly selected members of a database consisting of community members from Newcastle and Northumbria Universities.

**Table 2: Treatment and participant information**

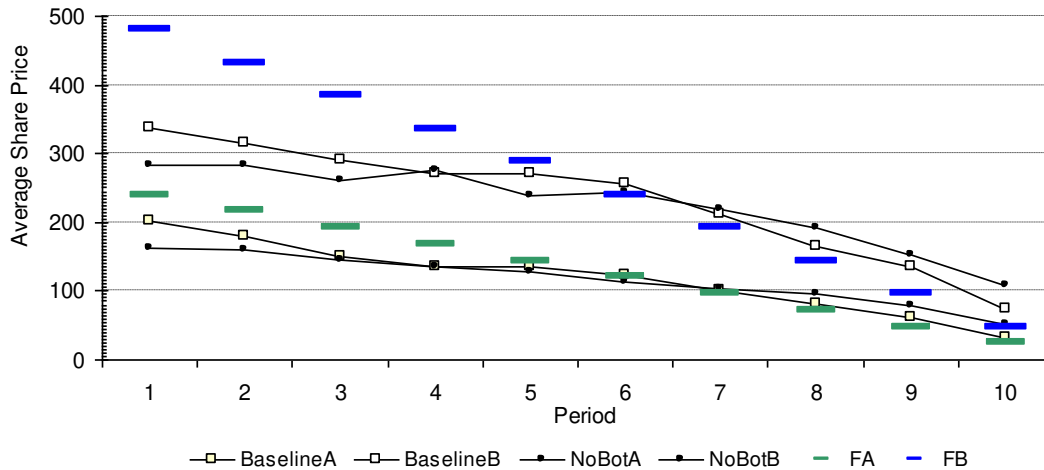
Pooling	Treatment	# participants	# cohorts	Average CRT-score	Female proportion	Investment task risky share
NoAT:	Baseline	70	8	0.978	0.653	0.376
	NoBot	67	8	0.556	0.502	0.329

AT:	SlowBot	68	8	0.500	0.705	0.385
	FastBot	68	8	0.750	0.720	0.293
	LiqBot	71	8	0.777	0.541	0.373
Total		344	40	0.690	0.619	0.349

#### 4.1 Evaluating announcement effects

**Observation 1:** We find no measurable market performance effect following the announcement of potential AT market participation. More specifically, the levels of *PD* and *RAD* are not statistically different between treatments Baseline and NoBot.

**Support:** Figure 1 exhibits the average price trajectories over all markets and cohorts for the NoBot and Baseline treatments. The average paths are quite similar in both treatments. Table 3 shows the average price discrepancy, *PD*, and relative absolute deviation, *RAD*, for each treatment. The test results of the two tailed Mann-Whitney test are reported in the bottom rows. The respective *PD* measures for the Baseline treatment and the NoBot treatment are .369 and .273, and the respective *RAD* measures are .428 and .319. As the test results indicate, the differences are not significantly different at the ten percent level; the p-values in both comparisons are 0.600 and 0.208. Our lack of announcement effect is consistent with Leal and Hanaki (2018) and counter to Farjam and Kirchkamp (2018).



**Figure 1.** Average asset prices in treatments without AT participation, NoBot vs. Baseline

#### 4.2 Algorithmic trading and differential price efficiencies

Our data suggest AT participation increases compliance with the law-of-one-price in the market, but does not lead prices closer to fundamental values.

**Observation 2:** Pricing discrepancy across assets, as measured by  $PD$ , is reduced by AT participation. The ATs effectively reduce discrepancies from the law-of-one-price. The LiqBot treatment generates the lowest cross-asset price discrepancy. The treatments SlowBot and FastBot give nearly identical cross-asset price discrepancy measures, which are lower than the NoBot treatment.

**Support:** We report the  $PD$  and  $RAD$  measures for each treatment in Table 3. Notice the  $PD$  measure for the NoBot treatment, 0.341, is much larger than of any AT treatment. Mann-Whitney non-parametric two-sample tests (reported in the lower half of Table 3) confirm the significant differences of  $PD$  measures for the NoBot – more generally, the NoAT treatments– from each of the three AT treatments. Further, we reject the null hypothesis that the  $PD$  is the same across treatments LiqBot and

FastBot or SlowBot. We cannot reject the null hypothesis of an equal *PD* in the treatments FastBot and SlowBot.

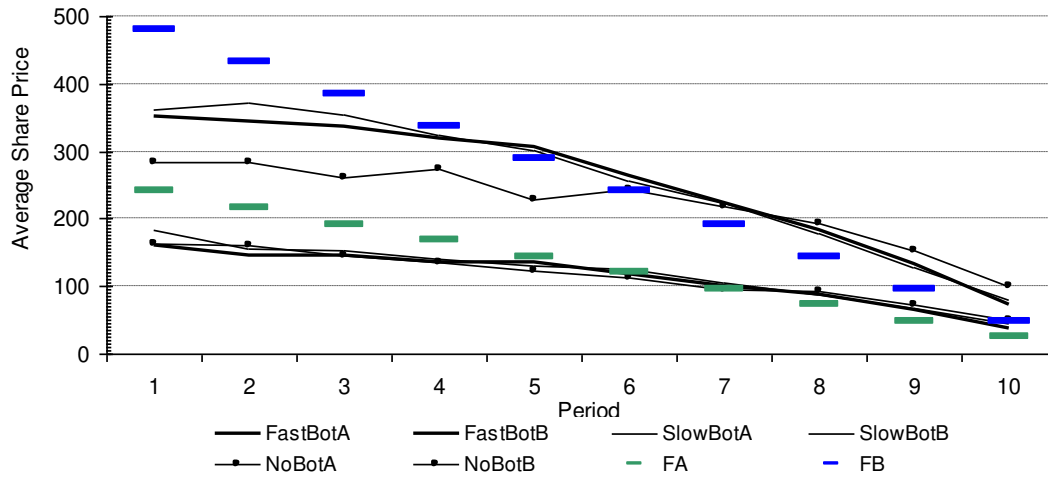
**Table 3. Price discrepancy and relative absolute deviation**

Treatment	PD	RAD
NoAT: Baseline	0.299	0.327
NoBot	0.341	0.434
AT: SlowBot	0.182	0.281
FastBot	0.184	0.337
LiqBot	0.095	0.279
<u>Two tailed Mann-Whitney test results (p-value):</u>		
AT vs NoAT	0.001***	0.122
NoBot vs Baseline	0.753	0.208
SlowBot vs NoBot	0.036**	0.115
FastBot vs NoBot	0.036**	0.345
SlowBot vs FastBot	0.674	0.462
Slow+FastBots vs LiqBot	0.004**	0.713
LiqBot vs NoBot	0.002***	0.059*
LiqBot vs SlowBot	0.027**	0.834
LiqBot vs FastBot	0.006***	0.674

*Note:* \*\*\* indicates a *p*-value of less than 0.01, \*\* indicates a *p*-value of less than 0.05, and \* indicates a *p*-value of less than 0.1.

**Observation 3:** Mispricing, as measured by *RAD*, is unaffected by AT participation.

**Support:** Inspecting Table 3 reveals no nominal ordering of *RAD* measures among our various treatments. This finding is borne out by uniform failure across our set of Mann-Whitney hypothesis tests comparing the *RAD* measures between various pairs of treatments and groups of treatments. Further, Figure 2 shows the similarity of the trajectories of average prices for the NoBot and various AT treatments.



**Figure 2.** Average asset prices in announcement treatments NoBot, SlowBot, FastBot, LiqBot

### 4.3 Liquidity impact of AT participation

We find there are no treatment effects on volatility. Further, there is no difference in the volumes and the number of limit orders per period between the NoAT treatment and the SlowBot treatment. Both the FastBot and LiqBot treatments exhibit similar volume, which exceeds that of other treatments. However, the LiqBot algorithm generates an increase in the average number of limit orders in a trading period.

**Observation 4:** Price volatility is unaffected by treatment variation.

**Support:** We report the calculated price volatility measure for each treatment in the second column of Table 4. These measurements span the narrow range of 0.211 and 0.263, for which no pairwise difference in an exhaustive set of Mann-Whitney two-sample tests as we report in the bottom part of Table 4.

**Observation 5:** The SlowBot does not affect the per-period transaction volume or the number of limit order submissions compared to the NoAT treatments.



**Support:** Table 4 records the average number of transactions and the average number of limit order submissions in each period and asset. The Baseline, NoBot, and SlowBot all generate the same average volume and limit order counts of 9 and 30/31 respectively. Accordingly, no pairwise test indicates any significant treatment effect between these three treatments.

**Observation 6:** The FastBot treatment has a positive effect on transaction volume, while at the same time leaving the number of limit orders unaltered. The LiqBot algorithm generates the same transaction volume as the FastBot, while at the same time increasing the number of limit orders.

**Support:** We observe, as reported in Table 4, more transactions per period and asset in the FastBot and LiqBot treatments (on average 13 and 14 respectively) than in the other treatments (9 each in SlowBot and NoAT treatments).<sup>11</sup> The significance of these differences is confirmed by the corresponding Mann-Whitney two-sample test. The average submission of limit orders per period in the FastBot treatment is 31, virtually the same as the one in the SlowBot and NoAT treatments. Surprisingly the ATs which only utilize market orders do not absorb liquidity in a way that impacts our standard measures. The LiqBot induces human subjects to submit more limit orders, 35, and the algorithm submits an additional 10 limit orders per asset – on average – each period. This increase is significant as substantiated by all of the appropriate comparisons via the Mann-Whitney test.<sup>12</sup>

---

<sup>11</sup> The differences between the FastBot and the other treatments are particularly notable in the first market, amid a general decline in the number of transactions from market 1 to markets 2 and 3.

<sup>12</sup> All reported significances of increased limit order submissions result from algorithm submissions and subject submissions. The subject-impacted increase of limit orders alone is not significant.

**Table 4. Liquidity: transactions, limit orders per asset and period, and spread**

Treatment		Volatility	Transaction volume	Limit orders	Spread measured at end of period <sup>a)</sup>	
		$\sigma$	$V$	$\zeta$	Offer-Bid spread, $S_t$	Spread exists %
NoAT:	Baseline	0.245	9	30	.669	96
	NoBot	0.263	9	30	.765	96
AT:	SlowBot	0.211	9	31	.668	95
	FastBot	0.211	13	31	.757	95
	LiqBot	0.231	14	35+10 bot	.672	100
<u>Two-tailed Mann Whitney test results (p-value):</u>						
AT vs NoAT		0.720	0.067*	0.079*	0.910	0.710
NoBot vs Baseline		0.462	0.958	0.916	0.248	0.911
SlowBot vs NoBot		0.529	0.635	0.752	0.462	0.825
FastBot vs NoBot		0.345	0.267	0.752	0.916	0.619
SlowBot vs FastBot		0.916	0.025**	0.916	0.462	0.743
Slow+FastBots vs LiqBot		0.540	0.002***	0.079*	0.582	0.017**
LiqBot vs NoBot		0.834	0.039**	0.009***	0.401	0.084*
LiqBot vs SlowBot		0.674	0.001***	0.021**	1.000	0.034**
LiqBot vs FastBot		0.529	0.425	0.016**	0.345	0.030**

Note: <sup>a)</sup> Bid is set equal zero when there are no limit bids, and ask is set equal 2000 where there are no limit asks. \*\*\* indicates a p-value of less than 0.01, \*\* indicates a p-value of less than 0.05, and \* indicates a p-value of less than 0.1.

**Observation 7:** The LiqBot treatment more effectively maintains a bid-ask spread as measured at the end of trading periods. Bid-offer spreads are the same in all treatments.

**Support:** In the last two columns of Table 4 we report, by treatment, the percentage of periods that close with both active limit bids and offers as well as the average width of the closing spread. In all market periods of the LiqBot treatment but one the market closes with active limit offers. However, the average width of the closing spread is between 0.668 and 0.765 around midpoint in all treatments.

#### 4.4 AT performance and margin positions

In this section, we summarize the benefits of the three considered ATs. The LiqBot algorithm is most effective in reducing deviations from the law-of-one-price. The LiqBot also does the most to increase market liquidity in terms of the number of limit orders and maintaining a consistent bid-ask spread. Both, LiqBot and FastBot are most effective in generating transaction volume. Further, we find the Liqbot doesn't take more extreme short-sale positions than the FastBot does.

We now evaluate the cost associated with the potential benefits of algorithm trader market participation. Subjects have symmetric information about dividend structure of the assets. In its core, our experimental asset market is a zero sum game. No positive gains from trade are possible, except through differing risk attitudes or subjective beliefs. Profits earned by an AT in this setting extract surplus from the human traders.

**Observation 8:** The SlowBot algorithm has a smaller detrimental impact on subjects' wealth than FastBot and LiqBot.

**Support:** The average per period profits of the algorithm trader in FastBot and LiqBot treatments are three times as high (99 vs 33 cash units) or higher (125 vs 33) than the in the SlowBot treatment. These differences are significant; the respective  $p$ -values of the two-tailed Mann Whitney are 0.000 and 0.046 as reported in the fourth column of Table 5.

**Table 5: AT transaction gains and maximum short positions**

	# per period per asset	average gain/trade	average gain/period	average max. short position
SlowBot	1.0	36	33	6
FastBot	3.1	35	99	19
LiqBot	4.9	25	125	16
<u>Two-tailed Mann Whitney test results (p-value):</u>				
SlowBot vs FastBot	0.012**	0.753	0.000***	0.018**
SlowBot+FastBot vs LiqBot	0.002***	0.807	0.002***	0.244
LiqBot vs SlowBot	0.001***	1.000	0.046**	0.014**
LiqBot vs FastBot	0.036**	0.674	0.000***	0.636

*Note:* \*\*\* indicates a  $p$ -value of less than 0.01, \*\* indicates a  $p$ -value of less than 0.05, and \* indicates a  $p$ -value of less than 0.1.

**Observation 9:** The short positions of the AT in the SlowBot treatment are less extreme than in FastBot and LiqBot treatments. Margin loans are within borrowing constraints in all treatments.

**Support:** Subjects were able to short sell and to purchase on credit. The borrowing capacity on purchases was 2,300 cash units, and 8 units of each share. Providing the AT with the same endowments and margin loan constraints as human participants, the AT reached in no treatment the borrowing capacity of subjects.<sup>13</sup> The AT violated the maximum short position requirement of subjects in 2 of 8 sessions in the SlowBot treatment, in 7 of 8 sessions in the FastBot treatment and in 7 of 8 sessions of the LiqBot treatment. The maximum short position during the session equalled or exceeded 16 (i.e., double the permitted position) in no session of the

<sup>13</sup> The cash borrowing constraint was not violated in any moment of the experiment by any arbitrage algorithm trader. The maximum margin loan was 1,980. Overall sessions, the FastBot algorithm finished eight periods (in three sessions) with a negative cash balance, i.e., on average once in 30 periods (3.3% of all periods in that treatment). The SlowBot algorithm never finished any period with a negative cash balance, and the LiqBot algorithm finished one period with a negative cash balance (0.4% of all periods in that treatment). Due to the few occasions, the differences are not significant at 10%.

SlowBot treatment, in 5 of 8 sessions of the FastBot treatment and in 4 of 8 sessions in the LiqBot treatment. The average maximum short positions in each treatment are recorded in the right column of Table 5. The ATs in the FastBot and the LiqBot treatment require on average triple capacities than in the SlowBot treatment in short positions (19 and 16 vs 6). The maximum short position measured in any session of the treatments is 36 for FastBot, 31 for LiqBot, and 14 for SlowBot. Since the differences are significant, we conclude that the SlowBot requires significantly fewer borrowing capacities than the other ATs, and that there are no significant differences in this regard between the LiqBot and the FastBot.

## **5. Conclusion**

In summary, from our study with automated arbitrageurs in the experimental asset markets with twin-shares (Charness and Neugebauer 2019) we have learned the following; algorithms can enhance the market efficiency in terms of parity pricing of equal risks, but this enhancement comes at a high adverse cost to experimental subjects.

Pricing discrepancies across assets with perfectly correlated returns are reduced when arbitrageurs act in the market. Nonetheless, even in this no-arbitrage environment we can confirm the law-of-one-price hypothesis only relative to the setting without arbitrageurs, but pricing discrepancies between twin-shares are still notable in the data. Pricing is closer to parity when the AT provides liquidity. The reason for this enhancement is most likely the fact that the AT actively sets the price spread, whereas the liquidity absorbing bots do not make any price proposals. The traditional view on arbitrageurs' effect (e.g., as suggested but also theoretically challenged in Shleifer 2000) to push market prices towards fundamentals has not been

confirmed in our study. Similar to Leal and Hanaki (2018) we also find no announcement effect in our data; mispricing is not affected by expectations of subjects that an algorithmic trader participates in the market.

The participation of the algorithmic arbitrageur implies great adverse costs on subjects in our experimental asset markets. The adverse costs are increasing with the number of transactions of the AT. The number of transactions and the adverse costs are lower, when the algorithm is liquidity taker and the reaction time of the algorithm is high (in the SlowBot treatment). With low latency, the AT skims subjects significantly more than that. The liquidity providing algorithmic arbitrageur in our experiment trades more and obtains even higher gains than the low latency liquidity absorbing AT. The gains of the liquidity providing AT could be much smaller if we reduced the margins of the algorithm at limit order submission.

Our experiment involves only mechanic arbitrageur ATs. These have been designed in first place to enhance market efficiency rather than with the objective to skim gains from experimental subjects. In the future we plan to study also ATs that are more profit oriented.

## References

- Aldrich, E. M., & Friedman, D. (2017). Order protection through delayed messaging. <https://cloudfront.escholarship.org/dist/prd/content/qt4938f518/qt4938f518.pdf>
- Aldrich, E. M., and López Vargas, K. (Forthcoming). Experiments in High-Frequency Trading: Testing the Frequent Batch Auction, *Experimental Economics*.
- Andersen, T. G., Bondarenko, O., Kyle, A. S., & Obizhaeva, A. A. (2016). Intraday trading invariance in the E-mini S&P 500 futures market. <http://pages.nes.ru/aobizhaeva/ABKO-intradayinv.pdf>
- Asparouhova, Elena, Peter Bossaerts, Kristian Rotaru, Tingxuan Wang, Nitin Yadav and Wenhao Yang (2019), Humans in Charge of Trading Robots: The First Experiment. <http://uleef.business.utah.edu/newsite/pubs/PilotStudyAsparouhova.pdf>
- Baron, M., Brogaard, J., Hagströmer, B., & Kirilenko, A. (2018). Risk and return in high-frequency trading. *Journal of Financial and Quantitative Analysis*, 1-32.
- Biais, B., Foucault, T., & Moinas, S. (2015). Equilibrium fast trading. *Journal of Financial Economics*, 116(2), 292-313.
- Bossaerts, P., Shachat, J., Xie, K. (2018). Arbitrage Opportunities: Anatomy and Remediation, Working Paper
- Brogaard, J., Hendershott, T., Riordian, R. (2014). High-Frequency Trading and Price Discovery. *The Review of Financial Studies*. 27(8), 2267-2306
- Brogaard, J., & Garriott, C. (forthcoming). High-Frequency Trading Competition. *Journal of Financial and Quantitative Analysis*. doi:10.1017/S0022109018001175
- Budish, E., Cramton, P., & Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4), 1547-1621.
- Carrion, A. (2013). Very Fast Money: High-frequency trading on the NASDAQ. *Journal of Financial Markets*. 16(4), 680-711
- Chaboud, A., P., Chiquoine, B., Hjalmarsson, E., Vega, C. (2014). Rise of the Machines: Algorithmic Trading in the Foreign Exchange Market. *The Journal of Finance*. 69(5), 2045-2084
- Charness, G., & Neugebauer, T. (2019). A Test of the Modigliani-Miller Invariance Theorem and Arbitrage in Experimental Asset Markets. *The Journal of Finance*, 74(1), 493-529.
- Charness, Gary and Uri Gneezy, 2010. Portfolio choice and risk attitudes: An experiment. *Economic Inquiry* 48(1), 133-146.

- Clark-Joseph, A. (2014). Exploratory trading. Unpublished job market paper. Harvard University, Cambridge, MA.
- Easley, D., Lopez de Prado, M.M., O'Hara, M. (2011). The Microstructure of the Flash Crash: Flow toxicity, liquidity crashes and the Probability of Informed Trading. *Journal of Portfolio Management*. 39(1), 19-29.
- Easley, D., Lopez de Prado, M.M., O'Hara, M. (2012a). The Volume Clock: Insights into the High- Frequency Paradigm. *Journal of Portfolio Management*. 39(1), 19-29
- Easley, D., Lopez de Prado, M.M., O'Hara, M. (2012b). Flow Toxicity and Liquidity in a High- frequency World. *The Review of Financial Studies*. 25(5), 1457-1493
- Farjam, Mike, and Oliver Kirchkamp, 2017. Bubbles in hybrid markets: How expectations about algorithmic trading affect human trading, *Journal of Economic Behavior & Organization*, 146, 248-269.
- Foucault, T., Hombert, J., Rosu, I. (2016). News Trading and Speed. *Journal of Finance*. 71(1), 335- 382
- Frederick, Shane, 2005. Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*. 19(4), 25-42.
- Glantz, Morton and Robert Kissell. Multi-Asset Risk Modeling: Techniques for a Global Economy in an Electronic and Algorithmic Trading Era. Chapter Academic Press, 2013, p. 258.
- Gomber, P., Gomber, P., & Haferkorn, M. (2015). High frequency trading. In Encyclopedia of Information Science and Technology, Third Edition, IGI Global, 1-9.
- Hagströmer, B., Nordén, L. (2013). The diversity of high-frequency traders. *Journal of Financial Markets*. 16(4), 741-770
- Harrison, J. M., & Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3), 381-408.
- Hasbrouck, J. & Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*. 16(4), 646–679
- Hendershott, T. & Riordan, R. (2013). Algorithmic Trading and the Market for Liquidity. *Journal of Financial and Quantitative Analysis*. 48(4), 1001-1024
- Hendershott, T., Jones, C., M., Menkveld, A., J. (2011). Does Algorithmic Trading Improve Liquidity? *The Journal of Finance*. 66(1), 1-33



- Jarrow, R. A., & Protter, P. (2012). A dysfunctional role of high frequency trading in electronic markets. *International Journal of Theoretical and Applied Finance*. 15(03), 1250022.
- Jarrow, R., & Protter, P. (2015). Liquidity suppliers and high frequency trading. *SIAM Journal on Financial Mathematics*. 6(1), 189-200.
- Kirilenko, A., Lo, A. (2013). Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents. *The Journal of Economic Perspectives*. 27(2), 51-72
- Kirilenko, A., Kyle, A. S., Samadi, M., & Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72(3), 967-998.
- Leal, Jacob Sandrine and Nobuyuki Hanaki. Algorithmic Trading, What if it is just an Illusion? Evidence from Experimental Financial Markets. No. 2018-31. Groupe de REcherche en Droit, Economie, Gestion (GREDEG CNRS), University of Nice Sophia Antipolis, 2018.
- Menkveld, A.J., 2013. High frequency trading and the new market makers. *Journal of Financial Markets*. 712–740.
- Menkveld, A.J., Yueshen, B., (Forthcoming) The Flash Crash: A Cautionary Tale About Highly Fragmented Markets. *Management Science*.
- O'Hara, M. (2014). High-Frequency Trading and Its Impact on Markets. *Financial Analysts Journal*. 70(3), 18-27
- O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*. 116(2), 257-270
- Rietz, Thomas A., 2005. Behavioral Mis-pricing and Arbitrage in Experimental Asset Markets, *Working Paper*, University of Iowa.
- Shleifer 2000, *Inefficient Markets An Introduction to Behavioral Finance*. Clarendon Lectures in Economics. Oxford: Oxford University Press.
- Smith, Vernon L., Gerry L. Suchanek and Arlington W. Williams, 1988. Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica*. 56(5), 1119-1151.
- Stöckl, Thomas, Jürgen Huber, and Michael Kirchler, 2010, Bubble measures in experimental asset markets. *Experimental Economics*. 13(3), 284-298.
- Wah, E. (2016). How prevalent and profitable are latency arbitrage opportunities on US stock exchanges?. SSRN 2729109, <http://www.smallake.kr/wp-content/uploads/2015/07/SSRN-id2729109.pdf>

Zhang, S., & Riordan, R. (2011). Technology and market quality: the case of high frequency trading. ECIS 2011 proceedings.

## Appendix

**Table A1: Experimental treatments and characteristics**

Treatment	AT strategy	Instructions refer to potential AT
FastBot	Immediate exploitation of arbitrage opportunity by market order ; liquidity taker	YES
SlowBot	Delayed exploitation (4 ticks) of arb opportunity by market order ; liquidity taker	YES
LiqBot	Submits delayed matching limit order in the other market; immediate exploitation of subjects' market order; liquidity maker	YES
NoBot	---	YES
Baseline	---	NO

# Instructions

This is an experiment in market decision-making. You will be paid in cash for your participation at the end of the experiment. Different participants may earn different amounts. What you earn depends on your decisions and the decisions of others. The experiment will take place through computer terminals at which you are seated. If you have any questions during experiment, raise your hand and a monitor will come by to answer your question.

## I. The Situation

In this experiment, you will participate in a market of 9 participants. The identities of the other market participants will not be revealed to you.

Each market participant will be initially given **Cash** (1300 units) and **Shares** (4 units of asset "A" and 4 units of asset "B"). Shares generate **Dividends** (income) over the 10 periods of the experiment, but have no value at the end of the session. There are four possible dividends that can be paid in a period, and each is equally likely to be paid (random drawing). At the end of EACH period, EACH share will pay the owner a dividend.

When the experiment starts, you will participate in a market where the **Shares** can be bought and sold between participants. You pay out of your Cash when you buy a share, and you get Cash when you sell a share.

The experiment is divided into **10** consecutive trading **Periods**. Within each period, the market is open for trading Shares. When a period is over your Cash and Shares will carry over to the next period.

After the payment of the last dividend at the end of period **10**, all shares will be worth nothing. Your earnings will be based on the amount of cash that you have at the end of the 10 periods. You accumulate cash by buying and selling shares, and/or by holding shares and collecting dividends.

## II. Share classes and Dividends

Trading will occur in two classes of shares, one is called "A" and one is called "B". The dividend of an "A" share per period can be **0, 8, 28 or 60 Cash Units (CU)**, with equal chances. The dividend of a "B" share in each period will be identical to the one paid on "A" shares **plus a fixed 24 CU, i.e., 24, 32, 52, 84 CU**. Thus, if an "A" share pays 0, a "B" share pays 24 CU; if an "A" share pays 8 CU, a "B" share pays 32 CU, etc. The dividends will be added to your cash amount immediately.

Note that the average dividend per period per "A" share is 24 CU, since this is the average of 0, 8, 28, and 60, the equally likely dividends that can be paid. That is, over many periods, the expected average dividend per period tends to be 24 CU per "A" share. Likewise, the expected dividend per share of "B" is 48 CU, as this is the average of 24, 32, 52, and 84.

The minimum total dividend that could possibly be paid on an "A" share is 0 (if 0 is drawn in each of the 10 periods) and the maximum that could possibly be paid on a "A" share is 600 (if 60 is drawn in each of the 10 periods). Similarly, the minimum total dividend that could

possibly be paid on a “B” share is 240 (if 0 is drawn on “A” in each of the 10 periods) and the maximum that could possibly be paid on a “B” share is 840 (if 60 is drawn on “A” in each of the 10 periods).

**Understanding expected dividends:**

The following table summarizes the sum of remaining dividends per share class:

Period	Remaining dividends incl. same period	X	Range and Expected dividend per “A” share	=	range and Sum of expected dividends per “A” share	Expected dividend per “B” share	range and Sum of expected dividends per “B” share
1	10		0..24..60		0..240..600	24 + “A”	240..480..840
2	9		0..24..60		0..216..540	24 + “A”	216..432..756
3	8		0..24..60		0..192..480	24 + “A”	192..384..672
4	7		0..24..60		0..168..420	24 + “A”	168..336..588
5	6		0..24..60		0..144..360	24 + “A”	144..288..504
6	5		0..24..60		0..120..300	24 + “A”	120..240..420
7	4		0..24..60		0..96..240	24 + “A”	96..192..336
8	3		0..24..60		0..72..180	24 + “A”	72..144..252
9	2		0..24..60		0..48..120	24 + “A”	48..96..168
10	1		0..24..60		0..24..60	24 + “A”	24..48..84
End	0		-		0	-	0

*Exercise:* To check your understanding of the table, please answer now the quiz questions on your screen. Please inform the instructor if you need any help.

**III. How to Trade Shares?**

We are interested in the price you are bidding to pay and the price you are asking to sell. In order to buy shares, you need cash. Alternatively you can **borrow cash** (with no interest) up to 2600 CU. The cash you own is shown on the screen. In order to sell shares, you need shares. The number of shares you own is indicated at the top of your screen for “A” and “B” shares, respectively. If you do not own (enough) shares and wish to sell (more) shares anyway, you can **borrow to sell up to 8 class “A” shares AND up to 8 class “B” shares**. If you sell more shares than you own your share holdings will be negative. For given negative share count at the end of the period, the dividend on these negative shares will be subtracted from your cash.

During a period, you may buy or sell shares (see Figure 1 at the end of the Instructions). **Note that you can only buy or sell one share at a time.**

**1. Submit an ASK (a proposed selling price)** for one share. You can offer a share from your share holdings for sale by entering the asking price to sell one share in the space underneath the button ASK. You confirm the ask by a click on the button. The ask is then added to the list of outstanding asks. The outstanding asks are publicly recorded in increasing order, i.e. the best outstanding ask (the cheapest proposed selling price) being placed at the top of the list. All market participants can see this list.

*Note:* you can submit as many asks as you like to sell one share. Upon selling one share, all your outstanding asks (for that share class) are cancelled. **To sell another share of that share class you then must submit a new ask.**

**2. Submit a BID** (a **proposed buying price**) for one share. You can bid to purchase a share by entering your bidding price for one share in the space underneath the button **BID**. You confirm your bid by a click on the button. The bid is then added to the list of outstanding bids. The outstanding bids are publicly recorded in decreasing order, i.e., the best outstanding bid (highest proposed purchase price) being placed at the top of the list. All market participants can see this list.

*Note:* If two or more orders (bids or asks) are the same, they are listed in the order of arrival, earlier orders being given priority over later ones. Upon purchasing one share, all your outstanding bids (for that share class) are cancelled. **To buy another share for this share class you then must submit a new bid.**

**3. Immediate BUY** – accept an ask: The best outstanding ask of the other market participants is marked on your screen. You can accept the asking price (i.e., entering in a purchase agreement of a share with the seller) by clicking on the button **Immediate BUY** below the list of outstanding asks.

**4. Immediate SELL** – accept a bid: The best outstanding bid of the other market participants is marked. You can accept the bid (i.e., entering in a sale agreement of a share with the buyer) by clicking on the button **Immediate SELL** below the list of outstanding asks.

*Note:* Your own orders are displayed in blue, while the other orders are visible to you in black. You cannot accept your own orders. You cannot purchase shares if the ask exceeds your cash plus credit line. If your holding of “A” shares is -8, you cannot sell any further “A” shares. If your holding of “B” shares is -8, you cannot sell any further “B” shares.

#### **IV Transaction and price announcement**

Upon acceptance of a bid or ask, via **Immediate BUY** or **Immediate SELL**, a transaction is completed. The accepted order is the transaction price. The transaction price is recorded on your screen in between the lists of bids and asks. Next to the price you are informed if you participated as buyer or seller in the transaction. The more recent prices are listed first. The most recent prices are also recorded for each share class in the middle of the screen below the cash amount.

Upon transacting the price is debited from the buyer’s cash balance and credited to the seller’s cash balance. The purchased share is added to the buyer’s share holdings and subtracted from the seller’s share holding.

*Note:* Next you are going to participate in a **Practice Session** of trading. You trade for 3 minutes on your screen with the other participants. There are **NO payoff consequences** linked to trading in the **Practice Session**. During the Practice Session please practice submissions of bids and asks, immediate selling and buying. You may want to practice selling more shares than you own to end up with a negative share count. You may also want to practice buying

more shares than you can pay with your own money to end up with a negative cash balance. During the **Practice Session** none of your actions will have any payoff consequences.

## **V. Information**

You will receive real-time updates on bids, asks and prices for both share classes. Information regarding the two share classes “A” and “B” are given on the screen on the left-hand and on the right-hand side, respectively. You will receive summary information about the prices at opening of the period, the high, the low and the average price during the period.

In each period you will be reminded on screen about the expected future dividends, and the sum of expected dividends for the remaining periods. Finally, the realized past dividends are shown. The **latest** paid out **dividend** of the prior period is highlighted.

The **past prices** are shown in a table on the bottom of the screen, including the prices at opening, closing, the high, low and **average** of each past period. **Alternatively to the past prices**, you receive past information on your share and cash holdings at the end of the period, buys and sells during a period, and the past period dividends. You can alternate the **past information** with the **past prices** by clicking on the button. (see Figure 1)

## **VI. Endowment and earnings**

The experiment involves **3 rounds** of **10 periods** of trading. Each trading period in the first round will last 180 seconds, and 120 seconds in the later rounds. Before the first round of each 10 periods starts, you will be given another **Practice Session** during which you can practice making offers and transactions. The Practice Session will allow you again to trade three minutes without any payoff consequences.

At the beginning of each of the 3 rounds, you will be endowed with shares and cash. You will receive 4 “A” and 4 “B” shares and 1300 cash units. If you run out of cash, you will be able to borrow up to 2600 cash units to purchase shares. If you run out of shares, you will be able to sell 8 borrowed shares of “A” and 8 borrowed shares of “B”. Once again, note that if you have borrowed shares, your share count is negative. For each negative share at the end of the period, the dividend to be paid on one share will be subtracted from your cash.

At the end of the experiment, cash units CU will be converted to Pound Sterling, at an exchange rate of £1 = 200 CU. Your final payment will be equal to the cash you had at the end of the decisive round. The final payment will be made to you in private; you will receive an envelope delivered to your seat in exchange for your signed receipt. (see Figure 2)

## **VII. Trading Algorithm**

Besides the participants in the room, a computerized trading algorithm may participate in the market. The computerized algorithm may take the same actions as you. It can buy and sell in the market. The details of the strategy followed by the algorithm are not revealed to you, and you will not be informed if the computerized trading algorithm actually acts in the market or not.

## VIII. Summary

1. You will be given an initial amount of Cash and Shares at the very beginning.
2. Each “A” share pays the owner a dividend of either 0, 8, 28 or 60 CU at the end of EACH of the 10 trading periods. Each of these amounts is equally likely to be drawn at the end of the period. The average dividend per period per “A” share is 24 CU. The dividend of the “B” share is 24 CU higher than the dividend of the “A” share, thus the average dividend of the “B” share is 48.
3. You can submit offers to BUY shares and offers to SELL shares. You can make immediate trades by buying at the lowest ask (offer to sell) or selling at the highest bid (offer to buy).
4. You will participate in 3 rounds of 10 periods. At the end of the experiment one round of ten periods is selected for payment. A participant will roll the die. The first outcome of the die roll (smaller or equal 3) will determine the payment-decisive round for every participant in the experiment.
5. Note that if you borrow cash or shares you may end a round with a negative cash balance. If a round is chosen for payment in which you incur losses, you will earn nothing.
6. A computerized trading algorithm may participate in the market. However, you will never be told if the algorithm acts in the market and what it is programmed to do.
7. The instructions are over. If you have any question, raise your hand and consult the monitor. Otherwise, please wait for the following Practice Session of three minutes.



Period 1 out of 10 Remaining time [sec]: 104

class "A" shares

#shares holding 4

Number of A shares

Remaining expected dividends of "A" will be displayed here

Sum of expected Dividends 240

dividends: 60, 28, 8, 0

Submit an ask

Submit a bid

Outstanding asks of "A"

Outstanding bids of "A"

Immediate buy / sell

Past information "A" (Past prices "A")

#buys	#sales	#shares	Dividend	#shares x Divi	Cash	Period
Record of your returns Latest dividend of "A" will be displayed here						

CASH 1300

Price of "A" shares

Price of "B" shares

Current transaction price

Average price, Low price, high price and opening price, will be displayed here

class "B" shares

#shares holding 4

Number of B shares

Remaining expected dividends of "B" will be displayed here

Sum of expected Dividends 480

dividend: divi "A" + 24

Submit an ask

Submit a bid

Outstanding asks of "B"

Outstanding bids of "B"

Immediate buy / sell

Past information "B" (Past prices "B")

#buys	#sales	#shares	Dividend	#shares x Divi	Cash	Period
Latest dividend of "B" will be displayed here						

Figure 1: Screen shot of trading screen

Period 10 out of 10 Remaining time [sec]: 78

A shares

market participation PAYOFF history						Price History
Period	#buys	#sales	#shares	Dividend	#shares x Divi	Cash
Your changes in holdings of "A" shares and changes in cash including dividend payments are recorded here						

Your eligible payoffs including payoff from market participation and from forecasting will be recorded for rounds 1-3

Your Payoff in Round 1	XXX
Your Payoff in Round 2	XXX
Your Payoff in Round 3	XXX

The End

B shares

market participation PAYOFF history						Price History
#buys	#sales	#shares	Dividend	#shares x Divi	Cash	Period
Your changes in holdings of "B" shares and changes in cash including dividend payments are recorded here						

Figure 2: Screen shot of payoff screen