

Arbitrage bots in experimental asset markets*

Martin Angerer^a, Tibor Neugebauer^b & Jason Shachat^c

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

Abstract: Trading algorithms are an integral component of modern asset markets. In twin experimental markets for long-lived correlated assets we examine the impact of alternative types of arbitrage-seeking algorithms. These arbitrage robot traders vary in their latency and whether they make or take market liquidity. All arbitrage robot traders we examine generate greater conformity to the law-of-one-price across the twin markets. However, only the liquidity providing arbitrage robot trader moves prices into closer alignment with fundamental values. The reduced mispricing comes with varying social costs; arbitrage robot traders' gains reduce the earnings of human traders. We identify factors which drive differences in human trader performance and find that the presence of an arbitrage robot trader has no disproportionate effect with respect to these factors on subjects' earnings.

Keywords: asset market experiment, arbitrage, algorithmic trading

JEL Codes: C92, G12

*We acknowledge funding from Economic and Social Research Council (ESRC), Luxembourg National Research Fund (FNR, Grants: INTER/RCUK/16/11555355 and INTER/MOBILITY/2022/16991801) and Forschungsförderungsfond Liechtenstein.

Corresponding author's e-mail address: Tibor.Neugebauer@uni.lu

1. Introduction

We study algorithmic arbitrageurs (which we call Arbitrage Robot Traders or ARTs) in the laboratory and assess their type-specific benefit for the capital market and the social costs of wealth extraction. ARTs are significant actors in financial equity markets.¹ The profitability of ARTs 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 Moinas, 2015; Budish, Cramton and Shim, 2015; O'Hara, 2015; Wah, 2016; Baron, Brogaard, Hagstromer and Kirilenko, 2018; Brogaard and Garriott, 2019).²

We experimentally explore the market impacts of alternative ARTs that seek riskless arbitrage in fragmented markets. In a controlled laboratory setting we can measure the effects relative to baseline treatments that are not available in the real world as, e.g., the unambiguously defined expected dividend value of assets. We consider ARTs that differ by either providing liquidity through the submission of limit orders or by taking liquidity through the exclusive submission of market orders. In the latter group we also analyse latency differences.

Through experiments with concurrent human and ART participation, we address the following questions with an emphasis on social costs and benefits:

- (i) In terms of pricing efficiency across assets, do ARTs help to induce the law-of-one-price?
- (ii) In terms of fundamental value, do ARTs help to increase pricing efficiency?³

¹ Menkveld and Yueshen (2019) report on arbitrage in fragmented markets where cross-asset arbitrage effectively connects buyers and sellers.

² 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.

³ 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 average price. The latter point refers to the traditional view that market equilibrium requires that the no-

- (iii) How does an ART's latency impact its performance?
- (iv) To what extent do ARTs transact in markets, and what are the social costs in terms of gains that ARTs reap from human subjects?

We approach these questions by implementing 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 generally fail to exploit arbitrage opportunities. In this study, we introduce algorithmic arbitrageurs that detect, or attempt to generate, and then exploit arbitrage opportunities. In total we consider four treatments: (1) The “NoART” treatment replicates Charness and Neugebauer with solely human participants; (2) The “FastART” treatment incorporates a low latency liquidity taking ART; (3) The “SlowART” treatment incorporates a high latency liquidity taking ART. SlowART and FastART are liquidity taking in that they only place market orders. Finally, (4) the “LiquidART” treatment incorporates a liquidity providing ART that posts both market and limit orders.

Our results suggest that the three ARTs push markets towards the law-of-one-price. This effect is greater when the ART provides liquidity rather than only absorbing liquidity. Also, liquidity provision seems crucial for improving pricing efficiency vis-à-vis fundamentals as only the LiquidART improves price efficiency. The reduced mispricing comes at a social cost incurred through the extraction of aggregate wealth from human traders in the market.⁴ Even in our quite thin markets this cost never amounts to more than 30% of the expected value of a single trader's endowment, which

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.

⁴ The asset structures and induced preferences in our experiment do not permit welfare generation via gains from exchange. Thus, the ARTs can only capture wealth from the human participants.

is about one-half the standard deviation (57%) of a subjects' market earnings. Further, the ART's wealth extraction does not discriminate against any subject with respect to the subject's trading strategy and individual characteristics such as cognition, gender, or risk attitude.

Our study contributes to two modest streams of the experimental asset market literature on (i) arbitrage and (ii) ART 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.⁵ 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 endowments of traders are factors that reduce arbitrage opportunities whereas relaxing margin requirements or restrictions on short sales do not.

(ii) We contribute to the nascent literature using experimental methods to study ARTs in financial markets. A recent survey of the literature on algorithmic trading in experimental asset markets suggests that there are a handful of related studies on ARTs in such markets (Bao, Nekrasova, Neugebauer and Riyanto 2021). Harrison (1992) suggests that an ART can help informational efficiency and decrease mispricing between spot and futures markets. Rietz (2005) and Grossklags and Schmidt (2006)

⁵ 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).

study arbitrage in a prediction market and report mixed evidence. Rietz reports that the arbitrageur was involved in most trades in the experiment, and Grossklags and Schmidt find a lower rate of ART's involvement and no significant effect on market efficiency in terms of the law-of-one-price. Berger, DeSantis and Porter (2020) study latency arbitrage within one market in which a sniper algorithm exploits market orders to find increased transaction activity. Finally, Neugebauer, Shachat and Szymczak (2020) test the Modigliani and Miller theorem of dividend policy irrelevance with and without ART to find, similarly to what we report here, that the law-of-one-price is more supported with than without ART.

We do not investigate competition among algorithmic arbitrageurs in this study. Related research studies algorithmic market makers in a single market setting report, depending on the research question, market quality enhancements or no enhancement with algorithmic trading (Aldrich and López Vargas 2019; Asparouhova, Bossaerts, Rotaru, Wang, Yadav and Yang 2019). Differently from the above cited ART experimental studies that investigate the effect of centralized arbitrage, subjects choose algorithms for trading in the latter studies. These studies thus make progress on the important issue of endogenous algorithmic trader adoption.

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

2. Experimental design

We present the following aspects of our experimental design: the economic environment including asset structure and trader endowments; the constraints on short sales and leveraged purchases; the continuous double auction market institution; our ARTs' strategies; and the experimental procedures.

2.1 Economic environment

Following Charness and Neugebauer (2019), we implement arbitrage markets in the Smith et al (1988) formulation of multi-period lived assets. This setting has two financial securities, asset A and asset B, that generate cash dividends for ten periods and can be traded for cash in the continuous double auction market. 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 generate a common cash 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 the single dividend is 24 cash units. If there are t remaining periods the expected sum of future 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 cash units. Therefore, the dividends of assets A and B are identical modulo a shift, and thus perfectly correlated. Asset B can be thought of as mixed financial fund that combines (at a 50:50 ratio) the risky asset A with a fixed income security that pays 24 cash units per period with certainty. Further the expected value of a single dividend of asset B equals 48 cash units, and if there are t remaining periods the expected sum of future dividends for holding a unit of asset B is $48t$ cash units. The law-of-one-price

asserts the difference between prices of assets A and B must be $24t$ cash units, where t indicates the remaining number of dividend payments.

In our setting, which replicates the correlation treatment of Charness and Neugebauer (2019), there are nine human traders. Prior to period one, every subject receives an endowment: 1300 cash, and four units each of assets A and B. Following Charness and Neugebauer (2019), and to remain consistent with common arbitrage practice, we allow limited leveraged purchases and short sales. 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, in case of a short position, 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 cash holdings 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 cash 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 $48t$, given the remaining number of dividend payments $t = 1, \dots, T, T = 10$.

2.2 Continuous double auction

Each period, prior to the determination and payment of dividends, traders can simultaneously buy and sell units of the assets in separate 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*, bid^j , is an amount of cash units at which the trader is willing

to purchase a unit of asset j . A *limit ask*, ask^j , is an amount of cash units a trader is willing to accept to dispose of 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 difference between the lowest limit ask and the highest limit bid defines 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 cash units).

There are two other actions a trader may take: *market buys* and *market sells*. A trader submits a market buy by accepting to purchase a unit at the lowest outstanding limit ask in the order book. This generates a transaction in which the trader submitting the market buy and the trader who submitted the lowest outstanding ask trade at that ask price. Similarly, a trader submits a market sell by accepting to sell a unit of the asset at the highest outstanding limit bid in the order book. This generates a transaction in which the trader submitting the market sell and the trader who submitted the highest outstanding bid trade at the bid price. Note that whenever a transaction occurs the involved limit 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.⁶ 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.

⁶ Charness and Neugebauer (2019) used this implementation, which makes the submission process for subjects easier. Clearing a trader's position prevents subjects from inadvertently taking on larger short positions than intended, and subjects can react fast to changes in the order book as they do not need to cancel outstanding bids before submitting enhancing ones.

2.3 Arbitrage robot trading

We consider two types of ARTs. Each type seeks out riskless arbitrage utilizing only the same information available to the traders. One ART type only makes market orders, exclusively *absorbing* market liquidity. We call these *Market Order ARTs*. 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 ART provides market liquidity. We call this type *Limit Order ART* or *LiquidART*, short for *Liquidity providing ART*. In any pair of markets for assets A and B, there is at most one ART and up to 9 human trader subjects. The following is a summary of the treatments and the corresponding ART strategies:

1. FastART: Immediate exploitation of arbitrage opportunity by market orders - liquidity taker;
2. SlowART: Delayed exploitation of an arbitrage opportunity by market orders - liquidity taker. The ART waits until four subsequent human market actions have been taken before executing the arbitrage opportunity if it stills exists;
3. LiquidART: Submits a delayed matching limit order in the other market; immediate exploitation of subjects' market order - liquidity maker and taker. The ART waits until (at least) two subsequent human market actions have occurred before submitting the matching limit order. The ART's delayed limit order appears on the order book if it is the best in the market and whilst the matched human order is outstanding. Arbitrage is immediate upon a subject's acceptance of the ART's order, the ART simultaneously accepts the matched human order; and,
4. NoART: No ART involvement; only human traders.

2.3.1 Market Order ART

Market Order ARTs exploit the bid-ask arbitrage opportunities across markets that violate the law-of-one-price in view of the perfect correlation of dividends. In each trading period the law-of-one-price dictates the price of asset B is equal to the price of asset A less twenty-four (i.e., the expected dividend) times the number of remaining

dividend draws, i.e., $P_B - P_A = 24t$. 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 $24t$ cash units. Then the ART simultaneously sells B and buys A in a pair trade by market order. This pair trade yields a certain, riskless gain equal to $bid^B - ask^A - 24t$. Second, when the difference between the lowest limit ask for asset B and the highest limit bid for asset A is less than $24t$ cash units, then the market order pair (Sell A, Buy B) results in a certain gain of $24t - (bid^B - ask^A)$. We implement two variants of the Market Order ART which differ in their latency of market order execution.

The *FastART* executes the required market order pair immediately when an arbitrage opportunity arises. Effectively, this renders the FastART's reaction time shorter than that of any human trader as human traders do not even get to observe an arbitrage opportunity in the order book. The *SlowART* executes the required pair of market orders, conditional upon the arbitrage opportunity still existing, only after a total of four human generated market actions have been taken once the arbitrage opportunity arises. We implement the SlowART to allow human traders the opportunity to exploit arbitrage opportunities. Despite the fact that the SlowART has probably very limited relevance for real-world trading, we use this treatment in the lab to control for the speed effect vis-à-vis the FastART treatment.

2.3.2 Limit Order ART

The LiquidART 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) for the other market. For example, if a *human* trader submits a limit bid for asset A, say bid_h^A , then the LiquidART formulates a

corresponding limit bid for asset B, bid_{Liquid}^B . This limit bid is chosen such that if a human were to accept it and the LiquidART were to accept bid_h^A simultaneously, then the LiquidART receives a certain profit. To ensure this is arbitrage, bid_{Liquid}^B is determined as

$$bid_{Liquid}^B = bid_h^A + 24t - \varepsilon,$$

where ε is the profit margin of the arbitrage determined by a random variable distributed uniformly over the interval $[0, \frac{24t}{2}]$. This interval sets the profit margin of arbitrage gain to no more than one-half the difference in the assets' dividend values. Note, a new ε is drawn independently for each limit order submitted by a human trader.

The LiquidART only submits bid_{Liquid}^B if it exceeds all outstanding bids for asset B, with a delay of two human generated market actions.⁷ The LiquidART immediately cancels its bid if the corresponding human bid, bid_h^A , is either accepted with a market sell or cancelled. If a human trader accepts the LiquidART's limit order, bid_{Liquid}^B , the LiquidART immediately accepts the best outstanding bid for asset A, generating an arbitrage gain of at least ε .

When a human trader submits a limit ask for asset A, ask_h^A , then the LiquidART formulates a limit offer for asset B, ask_{Liquid}^B , according to

$$ask_{Liquid}^B = ask_h^A + 24t + \varepsilon$$

⁷ We implemented this response delay in an attempt to obscure the market participation of the LiquidART from the subjects.

A similar logic as before prevails for the submission and cancellation of ask_{Liquid}^B , as well the execution of arbitrages. Finally, we note that the LiquidART follows symmetric processes in market A for human traders' limit orders for asset B.

2.4 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 without ART participation,
5. A sequence of three iterations of markets 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^{8,9},
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 (risk elicitation) task + earnings from the CRT (Cognitive Reflection 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 market earnings would be zero.

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

⁹ For the analysis we have pooled the data from the three markets. In Appendix A we provide figures that show the outcomes of the single markets.

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.¹⁰ 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.¹¹ 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).

2.5 Experimental treatments

Our experiment design varies the presence and type of ART.¹² Also, we adopt a between-subject design in which a group of nine traders experience the same ART presence/type in all three market iterations. In total we analyze four treatments: NoART¹³, SlowART, FastART and LiquidART (see Table A1).

¹⁰ 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.

¹¹ (1) A hat and a suit cost \$110. The suit costs \$100 more than the hat. How much does the hat cost? (2) 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?

¹² When there is an ART participant, there is exactly one. It can be viewed as dispatched by an entity outside the set of traders and its earnings are retained outside the group of traders. The instructions indicate that there may be a computerized participant. We didn't explicitly state its earnings are retained by the experimenter. Due to the competitive nature of the market experiment, we don't model the possibility this invokes other-regarding preferences.

¹³ We conducted one-half of the NoArt sessions with an announcement in the instructions there was the possibility of an algorithmic trader participating. We refer to this as the NoBot treatment, and the other half of those sessions, which we call the Baseline treatment, we gave no such announcement. We

3. Measures and hypotheses

3.1 Law-of-one-price

One benevolent view on arbitrageurs in markets is that they establish the law-of-one-price. To assess the extent the law-of-one price is satisfied in our setting we adopt the *cross-asset price discrepancy* measure introduced by Charness and Neugebauer;

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

where P_t^i and F_t^i denote respectively the average price and the fundamental value of securities A and B when t dividend draws remain. We expect that the ART's presence positively impact conformity to the law-of-one-price.

Hypothesis 1. PD will be lower in the SlowART, FastART and LiquidART treatments than in the NoART treatment.

3.2 Market mispricing

A subtler 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).

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

introduced this variation to test for a potential announcement effect as documented by Farjam and Kirchkamp (2017). We find no significant announcement effect, which is consistent with the findings of Leal and Hanaki (2018). We provide the details of this analysis in Appendix B.

Here, \bar{F}^i is the average fundamental dividend value of asset $i = \{A, B\}$ over the $T = 10$ periods of the asset market. Below we refer to the two respective measures as RAD^A and RAD^B .

The arbitrage strategies of our ARTs have no direct trading relation to fundamentals. We note that the potential of an ART to reduce RAD measurements, i.e. correct mispricing, in a market might depend upon the location of bids and asks relative to the fundamental values when the ART transacts and on the type of ART. Since we have no behavioral theory, we cautiously offer the following hypothesis on the ability of ARTs to improve price efficiency.

Hypothesis 2. RAD^i is unaffected by ART participation in markets.

3.3 Arbitrageur gains and exposure

In an attempt to inform social policy and assess ARTs' impact on wealth distribution we ask, how much does an arbitrageur skim the returns of market participants? In our experiment, we can measure unrealized arbitrage opportunities and realized arbitrage gains. By construction, the FastART transacts more than the SlowART. Accordingly we expect the following.

Hypothesis 3. The FastART experiences greater gains than the SlowART.

The hypothesis is not trivial, during the preceding time SlowART acts upon a pricing discrepancy, the discrepancy may disappear or widen. If the latter case is more frequent, the hypothesis is flawed. In the experimental literature of algorithmic trading in experimental asset markets, algorithmic traders that act at the pace of human traders seem to perform better than very fast algorithmic traders (see Gjerstad 2007, Peng, Shachat, Wei, Zhang 2020, and the survey of Bao et al. 2021), but evidence on the impact of arbitrageur latency is missing.

3.4 Individual Trader Characteristics

Subjects in economic market experiments are typically paid according to their investment performance. Trader characteristics that have a positive impact on the individual earnings of market participants include cognitive abilities (see Charness and Neugebauer 2019, and the survey of Bosch-Rosa and Corgnet 2021). On the other hand, excessive trading has a negative impact on performance (e.g., Odean 1999; Carbone, Hey and Neugebauer 2021). We speculate that ART participation may reinforce the shortcomings of poor trading behavior and potentially increase income inequality in the market.

Hypothesis 4. The earnings of subjects with high cognitive abilities are less impacted by ART participation than subjects with low cognitive abilities. Also subjects who engage in less frequent trading are less proportionately impacted by ART participation than subjects who engage in more frequent trading.

4. Results

In this section, we first examine pricing efficiency across treatments by assessing the *PD* and *RAD* measures. Thereafter, we report on the ARTs' profits, inform on the how individual characteristics and trading strategies drive human subjects' market earnings, and assess whether our ARTs impact the effects of these factors.¹⁴

Our unit of analysis is an individual session. In Table 1, we present the number of sessions for each treatment. In addition, we present the total number of participants per treatment (8 or 9 each in a session), the proportion of female participants, the average number of correct CRT responses, and the average proportion of tokens

¹⁴ In Appendix C, we discuss market quality measures including spread and volatility.

invested in the risky option in the investment task. In total, there are 40 sessions (8 per treatment) involving 344 participants, earning £22 on average.

Table 1: Treatment and participant information

Pooling	Treatment	# participants	# sessions	Average CRT	Female proportion	Investment task risky share
	NoART	137	16	0.767	0.578	0.353
ART	SlowART	68	8	0.500	0.705	0.385
	FastART	68	8	0.750	0.720	0.293
	LiquidART	71	8	0.777	0.541	0.373
	Total	344	40	0.690	0.619	0.349
KW test: p-value				0.247	0.067*	0.223

Note: * indicates a p -value of less than 0.1 according to two-tailed multiple sample *Kruskal-Wallis test*. According to the test, we find no significant differences across NoArt and ART treatments, besides the different proportion of females.

4.1 Algorithmic trading and differential price efficiencies

Our data suggest ART participation increases compliance with the law-of-one-price. However, but only the presence of the LiquidART algorithm generates prices closer to fundamental values.

Observation 1: Pricing efficiency across assets, as measured by PD , is increased by ART participation. The ARTs effectively reduce deviations from the law-of-one-price. The LiquidART treatment generates the lowest PD . The treatments SlowART and FastART give nearly identical PD measures, which indicate higher pricing efficiency than the NoART treatment. **This observation confirms Hypothesis 1 and that the presence on an ART brings stronger compliance with the law-of-one-price.**

Support: We report the PD and RAD measures for each treatment in Table 2. Notice the PD measure for the NoART treatment, 0.341, is larger than any other ART 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 NoART treatments for each of the three ART treatments. Further, we reject the null hypothesis that the PD is the same across treatments LiquidART and FastART or SlowART. But we cannot reject the null hypothesis of an equal PD in the FastART and SlowART treatments.

Observation 2: Mispricing vis-à-vis fundamentals, as measured by RAD , is minimally affected by the liquidity taking SlowART and FastART. The presence of the LiquidART reduces mispricing vis-à-vis fundamentals. **This observation presents a partial refutation of Hypothesis 2. Despite fundamental valuation having no direct motivation behind LiquidART its presence does reduce mispricing.**

Support: Inspecting Table 2 suggests no significant differences of RAD^A and RAD^B measures for NoART and market order ARTs. For the independent units of observation, we take the session average of the RAD measures of the three market iterations. However, there are differences in mispricing across treatments for asset B. The LiquidART treatment has a lower RAD^A and RAD^B than the NoART, FastART and SlowART treatments. Figure 1 (see also the figures in Appendix A and Appendix B) shows the differences of the trajectories of average prices for the NoART and various ART treatments and the fundamental dividend values.¹⁵

¹⁵ Figure 1 reveals a negative average deviation from fundamentals in each treatment. While this pattern is perfectly in line with theoretical assumptions of a positive risk premium, the status quo would be to observe positive deviations from fundamentals in the single asset market à la Smith et al. (1988). One potential explanation of this deviation from the status quo could be that we allow for short sales which have shown to imply lower asset prices (as, e.g., in Haruvy and Noussair 2006, or Füllbrunn and

Table 2. Price discrepancy and relative absolute deviation

Treatment		PD	RAD ^A	RAD ^B
	NoART	0.320	0.379	0.382
ART	SlowART	0.182	0.311	0.251
	FastART	0.184	0.386	0.289
	LiquidART	0.095	0.341	0.217

Two tailed Mann-Whitney test results (z-stat [p-value]):

NoART - ART	3.313 ^{***} [.001]	1.049 [.294]	1.215 [.294]
NoART – SlowART	1.837 [*] [.066]	0.122 [.903]	0.000 [1.00]
NoART – FastART	2.021 ^{**} [.043]	0.306 [.760]	-0.245 [.807]
NoART – LiquidART	3.491 ^{***} [.001]	1.898 [*] [.058]	2.939 ^{***} [.003]
SlowART – FastART	-0.420 [.647]	0.105 [.916]	0.000 [1.00]
SlowART – LiquidART	2.205 ^{**} [.027]	1.155 [0.248]	1.785 [*] [.074]
FastART – LiquidART	2.731 ^{***} [.006]	1.680 [*] [0.093]	2.941 [*] [.004]

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 according to the two tailed Mann-Whitney test. [*p*-values in brackets]

Comment: The reason for the different success of the ARTs in reducing mispricing vis-à-vis fundamentals can be seen in the different regions where arbitrage happens. ART activity around the fundamentals likely assists price discovery of the

Neugebauer *in press*). To test this and other alternative explanations, we regressed the relative deviation, which is frequently used as bubble measure, on the number of short sales. We did do so for every repetition and every asset and overall. We also ran separate regressions on female share, cognitive abilities, and risk aversion. There were no significant results. Another potential explanation could be the fact that we have a market for multiple assets rather than single asset trading and results for multiple assets are mixed (see the recent survey of Duffy et al. 2022). As said, however, we use the design of Charness and Neugebauer (2019) including the same instructions. Their price data are very close to fundamentals, but rather above than below. In conclusion, we can only speculate why we find the negative deviation from fundamentals. It could be a subject-pool effect, England versus California. Anyway, there is no problem as we are interested in relative pricing and the behavior is alike in all treatments.

fundamental value. Across all ART treatments, however, only a small fraction of the ART transactions lay in that area. In the other areas, we would not expect any impact for the market order ART. The limit order ART, unexpectedly for us, contributes to the equilibration process because it helps to increase the buying pressure when the market trades below fundamentals and increases the selling pressure when above fundamentals.

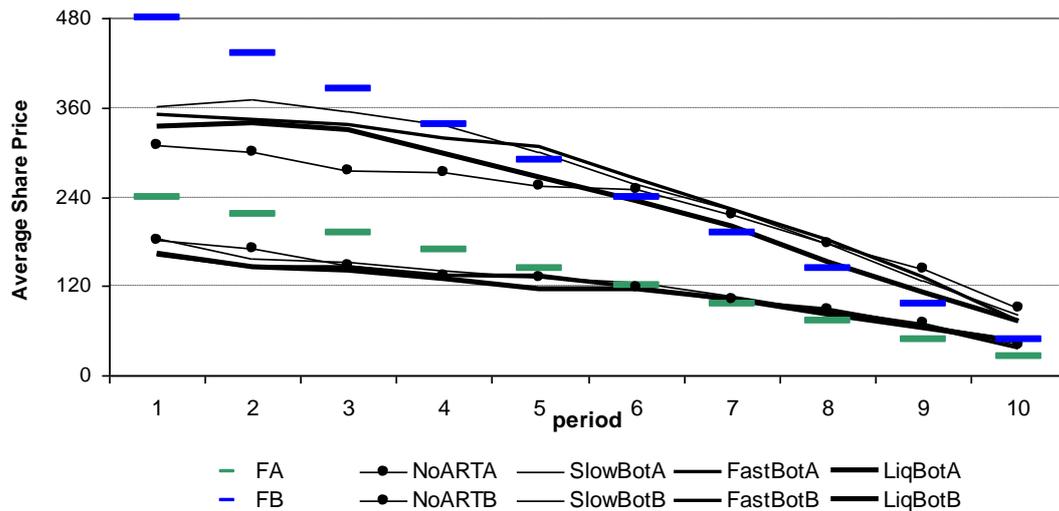


Figure 1. Average asset prices in treatments NoART, SlowART, FastART, LiquidART

4.2 The earnings of ARTs and human traders

We turn our attention from the relative benefits of our alternative ARTs to their respective social costs. Given our asset structure, absent of possible wealth gains from trade, any ARTs' earnings can only come at the expense of a loss to human traders. Accordingly, we use an ART's earnings as measurement of the social cost of the reduced mispricing it generates. We analyse the absolute and relative levels of these costs across ART types.

Table 3: ART transaction gains

	#arbitrage transactions per period	average gain/trade	average gain/period	average market earnings
SlowART	1.0	36	33	329
FastART	3.1	35	99	993
LiquidART	4.9	25	122	1220

Two tailed Mann-Whitney test results (z-stat [p-value]):

SlowART – FastART	-2.836*** [.005]	-0.315 [.753]	-2.521** [.012]
SlowART – LiquidART	-3.363*** [.001]	0.000 [1.00]	-3.361*** [.001]
FastART – LiquidART	-1.997** [.046]	0.420 [.674]	-2.100** [.036]

*** 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 3a: The average earnings of each ART type is less than 30% of the expected value of trader’s initial endowment. The LiquidART exceeds the market order ARTs in gains and in transaction frequency. The SlowART has smaller average earnings than either the FastART or LiquidART. **This observation confirms Hypothesis 3 that asserts the FastART experiences greater gains than the SlowART.**

Support: The expected value of each trader’s initial portfolio is 4180. The average market earnings of SlowART, FastART and LiquidART are 330, 990 and 1220 respectively. The economic size of these costs seems small. With respect to the relative earnings of the ARTs, the average per period profits of the algorithm trader in FastART and LiquidART treatments are three times as high (99 vs 33 cash units) or higher (122 vs 33) than the in the SlowART treatment. These differences are statistically significant; the respective *p*-values of the two-tailed Mann Whitney are 0.000 and 0.012, as reported

in the third column of Table 3. The LiquidART implies more transactions and more gains than either market order ART.

Observation 3b: ART transactions happen both early and late in a period. The share of ART transactions remains constant throughout the experiment.

Support: We compared the ART transaction share in the first and second half of the period transactions, and we compared the share of ART transactions across the three market rounds to find no significant differences overall. However, in the LiquidART treatment the transaction share of ART is at 0.286 lower in the third repetition than in the first repetition at 0.366, whereas in the FastART and the SlowART treatments the transaction share of ART increases to 0.266 and 0.146 from 0.211 and 0.084 respectively.

The overall economic cost of the ART's presence is surprisingly low; however, we should have concerns that ARTs prey on human traders with certain personal characteristics or those who adopt certain trading strategies. We evaluate the validity of this concern through a regression analysis on subjects' average market earnings.

We consider two types of primary drivers of the variation in market earnings across human subjects. The first type include the personal characteristics risk preference (as measured by responses in the Investment task), cognitive reflection (as measured by the number of correct answers in the CRT task) and gender. The second type are trading strategy properties. We consider the tendency to execute trades with market orders versus limit orders with a factor we call Take ratio (see also Stöckl 2014), the proportion of a trader's transactions made with market orders. Next, we measure the amount of excessive trading by intra-period speculation with a factor we call the Retrade ratio (see

also Carbone et al 2021). Finally, we use the number of limit orders a trader submits, #limitorders, to measure the willingness to provide liquidity.

4.3 Determinants of human trader performance

Observation 4: Human traders' earnings are positively correlated with their CRT scores, and with their willingness to use limit orders. Human traders' earnings are negatively correlated with excessive trading and their proportion of transactions made via market orders. The presence of an ART does not impact subjects' earnings differentially upon the cognitive abilities or trading strategies we consider. **This observation refutes the first component of Hypothesis 4 which asserts that traders with lower cognitive abilities are disproportionately harmed by ARTs. It also refutes the second component of the Hypothesis 4 which asserts those who trade more frequently are disproportionately harmed by ARTs.**

Support: Table 4 reports on the results of ordinary least squares regressions, with robust standard errors clustered at the experimental session level, of a human trader's market earnings on individual subject characteristics and measures of their trading strategies. Then we consider models that include the interaction of these factors with a dummy variable for the presence of an ART.

Table 4: Determinants of human traders' market earnings

Dependent variable: average human trader's earning in a market iteration						
Intercept	7183*** (17.2)	6824*** (14.7)	6853*** (13.32)	7465*** (10.2)	7262*** (8.52)	7508*** (7.49)
<i>CRT</i>		328*** (4.26)	304*** (3.48)		288** (2.00)	286 (1.66)
<i>Female</i>			-433* (-1.83)			-495 (-1.31)
<i>Safe Invest</i>			3.48 (.92)			.25 (.04)
<i>Retrade ratio</i>	-4450*** (-8.31)	-4371*** (-7.89)	-3983*** (-6.39)	-4462*** (-5.32)	-4336*** (-4.21)	-3991*** (-3.43)
<i>Take ratio</i>	-4495*** (-8.45)	-4150*** (-7.81)	-4146*** (-8.08)	-5046*** (-5.80)	-4703*** (-5.60)	-4772*** (-5.98)
<i>#limitorders</i>	3.28** (4.70)	3.03*** (4.93)	2.52** (3.87)	3.20* (1.80)	2.27* (1.73)	2.00 (1.27)
<i>ART</i>	-172 (-1.31)	-168 (-1.16)	-117 (-.69)	-621 (-.71)	-830 (-.85)	-1086 (-.96)
<i>ART × CRT</i>					73 (.43)	39 (.20)
<i>ART × Female</i>						65 (.13)
<i>ART × Safe Invest</i>						4.74 (.63)
<i>ART × Retrade ratio</i>				-41 (-.04)	-37 (-.03)	15 (.01)
<i>ART × Take ratio</i>				901 (0.83)	854 (.24)	941 (.92)
<i>ART × #limitorders</i>				0.18 (.10)	1.04 (.69)	0.77 (.44)
R-squared	0.473	0.523	0.518	0.475	0.525	0.521
#observations	332	307	279	332	307	279
#clusters	40	37	34	40	37	34

*** indicates a p-value of less than 0.01, ** indicates a p-value of less than 0.05, * indicates a p-value of less than 0.1. (t-stats in parenthesis)

CRT indicates the individual's number of correct answers to the CRT questions {0;3}. Data are missing for three NoART sessions.

Take ratio is the individual ratio of #market orders/#transactions in the session, [0;1].

Retrade ratio is the average individual ratio of min(#sells/#buys; #buys/#sells) per asset and period in the session, [0;1].

#limitorders is the individual's number of submitted limit orders in the course of the session, {0;1,185}.

ART is a binary dummy variable taking value 1 in treatments LiquidART, FastART, SlowART and 0 otherwise.

safeinvest is the amount of tokens not invested in the risky gamble of the investment game, {0;100}.

female is a binary dummy variable taking value 1 for female, and 0 for male. Data are missing for two FastBot sessions.

With respect to individual characteristics, only the CRT score has a strongly significant effect, which is positive, on earnings. In terms of trading strategies, the tendency to execute transactions with market orders, the factor Take ratio has a strong negative impact on earnings.¹⁶ Engaging in intra-period excessive trading yields expected lower earnings, as indicated by the highly significant and negative estimated coefficients of Retrade ratio. While exhibiting a willingness to provide liquidity, reflected in the significant and positive estimated coefficient for #limitorders, leads to higher market earnings. While these factors are all significant determinants of human trader earnings, the interaction of these factors with the presence of an ART is insignificant in all cases. In other words, while we find strong variance in subjects' performances in asset market trading, the presence of an ART does not shift the impact of the key factors driving this variance.¹⁷

6. Conclusion

We report data on experimental asset markets with twin-shares and arbitrage robot traders and show the following: arbitrage algorithms increase compliance with the law-of-one-price, but only arbitrage algorithms that provide liquidity reduce mispricing vis-à-vis fundamentals.

¹⁶ The Take ratio has also been analyzed in Stöckl (2014) and Stöckl and Kirchler (2014), who studied the behavior of traders with inside information. Insiders seem to prefer market orders over limit orders in their design, because market orders possibly imply a lesser amount of information leakage to the liquidity traders. Stöckl and Kirchler (2014) also explored submission and execution rates, which we do not report here because we detected autocorrelation with the Take ratio and at the same time, we found no increase in the explained proportion of the variance (R-squared) for our data. Accordingly, when we include these two suggested variables, we observe the standard errors on their coefficients to blow-up in value. This classic sign of multi-collinearity is confirmed when we calculate the variance inflation factor and find its value is approximately 10.

¹⁷ Unreported analysis of Gini coefficients across treatments find that ART participation leaves income inequality unaffected.

Our results contribute to the recent regulatory debate on make rebates/take fees (see, e.g., Foucault et al. 2013; Malinova and Park 2015). The proliferation of algorithmic trading profoundly impacted the liquidity of order books. In response, real-world exchanges offered monetary rewards for liquidity provision and charged fees for liquidity absorption. In light of our data this kind of discriminatory regulation is sensible. In our experiment, the benefits of reduced mispricing are larger with limit order arbitrageurs than with market order arbitrageurs.¹⁸

We note our results are based on monopolistic arbitrageurs. This approach is relevant from the perspective of a central regulator who carefully contemplates the trade-off between social benefit and costs of dispatching an ART. Our data suggest that the participation of an algorithmic arbitrageur imposes moderate costs on the trader subjects. The social costs are increasing with the number of transactions of the ART. When the algorithmic arbitrageur is liquidity taker and the reaction time of the algorithm is high, in the SlowART treatment, the number of transactions and adverse costs are lower. With low latency the market order ART imposes higher social costs in terms of wealth extraction than with high latency, whereas the benefits are the same. In fact, the liquidity providing algorithmic arbitrageur in our experiment trades more and obtains even higher gains than the low latency liquidity taking ART. An open question is how much the margin of our liquidity providing ART contribute to this result.

There are two obvious vectors for expanding the setting of our experiments beyond a monopolistic structure. One is the exogenous introduction of competition of multiple ARTs. We conjecture the introduction of competition, particularly between liquidity providing ARTs, would increase pricing benefits and at the same time reduce

¹⁸ In the appendix we report that also the market quality measures including spread and volatility for the liquidity maker arbitrageur are always as good or better than with the liquidity taker or without arbitrageur.

social costs. A second vector is to have traders endogenously choose to adopt algorithm on their behalf and to further modify the algorithms. Interesting initial studies that consider decentralized competition of market makers and sniper algorithmic traders (Aldrich and Lopez Vargas 2019) or "reactionary bots" (Asparouhova et al. 2019) led to welfare losses and market instabilities in continuous double auction experiments. It would also be an interesting effort to evaluate the social costs and benefits of human trader subjects employing ARTs that include the option to fine tune the trading strategy.

References

- Aldrich, Eric M., and Kristian López Vargas. (2019). Experiments in high-frequency trading: comparing two market institutions. *Experimental Economics* 1-31.
- Asparouhova, Elena, Peter Bossaerts, Kristian Rotaru, Tingxuan Wang, Nitin Yadav and Wenhao Yang (2019), Humans in Charge of Trading RoARTs: The First Experiment. <http://uleef.business.utah.edu/newsite/pubs/PilotStudyAsparouhova.pdf>
- Bao, Te, Elizaveta Nekrasova, Tibor Neugebauer, and Yohanes E. Riyanto, 2022. Algorithmic Trading in Experimental Markets with Human Traders: A Literature Survey, In Sascha Füllbrunn and Ernan Haruvy (Ed.) Handbook of Experimental Finance.
- Baron, M., Brogaard, J., Hagströmer, B., and Kirilenko, A. (2018). Risk and return in high-frequency trading. *Journal of Financial and Quantitative Analysis*, 1-32.
- Berger, N., DeSantis, M., & Porter, D. (2020). The Impact of High-Frequency Trading in Experimental Markets. *The Journal of Investing*, 29(4), 7-18.
- Biais, B., Foucault, T., and Moinas, S. (2015). Equilibrium fast trading. *Journal of Financial Economics*, 116(2), 292-313.
- Bosch-Rosa, Ciril and Brice Corgnet. Cognitive Finance, In Sascha Füllbrunn and Ernan Haruvy (Ed.) Handbook of Experimental Finance.
- Bossaerts, P., Shachat, J., and Xie, K. (2018). Arbitrage Opportunities: Anatomy and Remediation, Working Paper
- Brogaard, J., Hendershott, T., Riordan, R. (2014). High-Frequency Trading and Price Discovery. *The Review of Financial Studies*. 27(8), 2267-2306

- Brogaard, J., and Garriott, C. (2019). High-Frequency Trading Competition. *Journal of Financial and Quantitative Analysis*. 54(4), 1469-1497.
- Budish, E., Cramton, P., and 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.
- Carbone, E., Hey, J., & Neugebauer, T. (2021). An Experimental Comparison of Two Exchange Economies: Long-Lived Asset vs. Short-Lived Asset. *Management Science* 67(11):6946-6962.
- Carrion, A. (2013). Very Fast Money: High-frequency trading on the NASDAQ. *Journal of Financial Markets*. 16(4), 680-711
- Charness, G., and 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.
- Duffy, John, Jean-Paul Rabanal, and Olga Rud, 2022. Market Experiments with Multiple Assets: A Survey. In Sascha Füllbrunn and Ernan Haruvy (Ed.) *Handbook of Experimental Finance*.
- 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, Thierry, Kadan O., and Kandel, E. (2013). Liquidity Cycles, and Make/Take Fees in Electronic Market. *Journal of Finance* 68, 299-341.
- Frederick, Shane, (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*. 19(4), 25-42.
- Füllbrunn, Sascha, and Tibor Neugebauer (in press). Testing market regulations in experimental asset markets—the case of margin purchases. *Journal of Economic Behavior & Organization*.
- Gjerstad, Steven, (2007). The competitive market paradox. *Journal of Economic Dynamics and Control* 31 (5), 1753-1780.
- Grossklags, J., & Schmidt, C. (2006). Software agents and market (in) efficiency: a human trader experiment. *IEEE transactions on systems, man, and cybernetics, part C (applications and reviews)*, 36(1), 56-67.
- Harrison, G. W. (1992). Market dynamics, programmed traders and futures markets: beginning the laboratory search for a smoking gun. *Economic Record*, 68, 46-62.

- Harrison, J. M., and Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3), 381-408.
- Haruvy, Ernan, and Charles N. Noussair, 2006. The effect of short selling on bubbles and crashes in experimental spot asset markets. *The Journal of Finance*, 61(3), pp.1119-1157.
- Hendershott, T. and Riordan, R. (2013). Algorithmic Trading and the Market for Liquidity. *Journal of Financial and Quantitative Analysis*. 48(4), 1001-1024
- Leal, Jacob Sandrine and Nobuyuki Hanaki. (2018) 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.
- Malinova, K., and Park, A. (2015). Subsidizing liquidity: The impact of make/take fees on market quality. *The Journal of Finance*, 70(2), 509-536.
- Menkveld, A. J., and Yueshen, B. Z. (2019). The Flash Crash: A cautionary tale about highly fragmented markets. *Management Science*, 65(10), 4470-4488.
- Neugebauer, T., Shachat, J., & Szymczak, W. (2020). A test of the Modigliani-Miller theorem, dividend policy and algorithmic arbitrage in experimental asset markets.
- Odean, T. (1999). Do investors trade too much?. *American economic review*, 89(5), 1279-1298.
- O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*. 116(2), 257-270
- Peng, Y., Shachat, J., Wei, L., & Zhang, S. S. (2020). Speed Traps: Algorithmic Trader Performance Under Alternative Market Structures.
- Rietz, Thomas A., (2005). Behavioral Mis-pricing and Arbitrage in Experimental Asset Markets, *Working Paper*, University of Iowa.
- Shleifer, Andrei (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, (2014). Price efficiency and trading behavior in limit order markets with competing insiders. *Experimental Economics*, 17(2), 314-334.
- Stöckl, Thomas, Jürgen Huber, and Michael Kirchler, (2010). Bubble measures in experimental asset markets. *Experimental Economics*. 13(3), 284-298.

Stöckl, Thomas, and Kirchler, Michael, (2014). Trading behavior and profits in experimental asset markets with asymmetric information. *Journal of Behavioral and Experimental Finance*, 2, 18-30.

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>

Appendix A Average prices by treatment disaggregated by market iteration

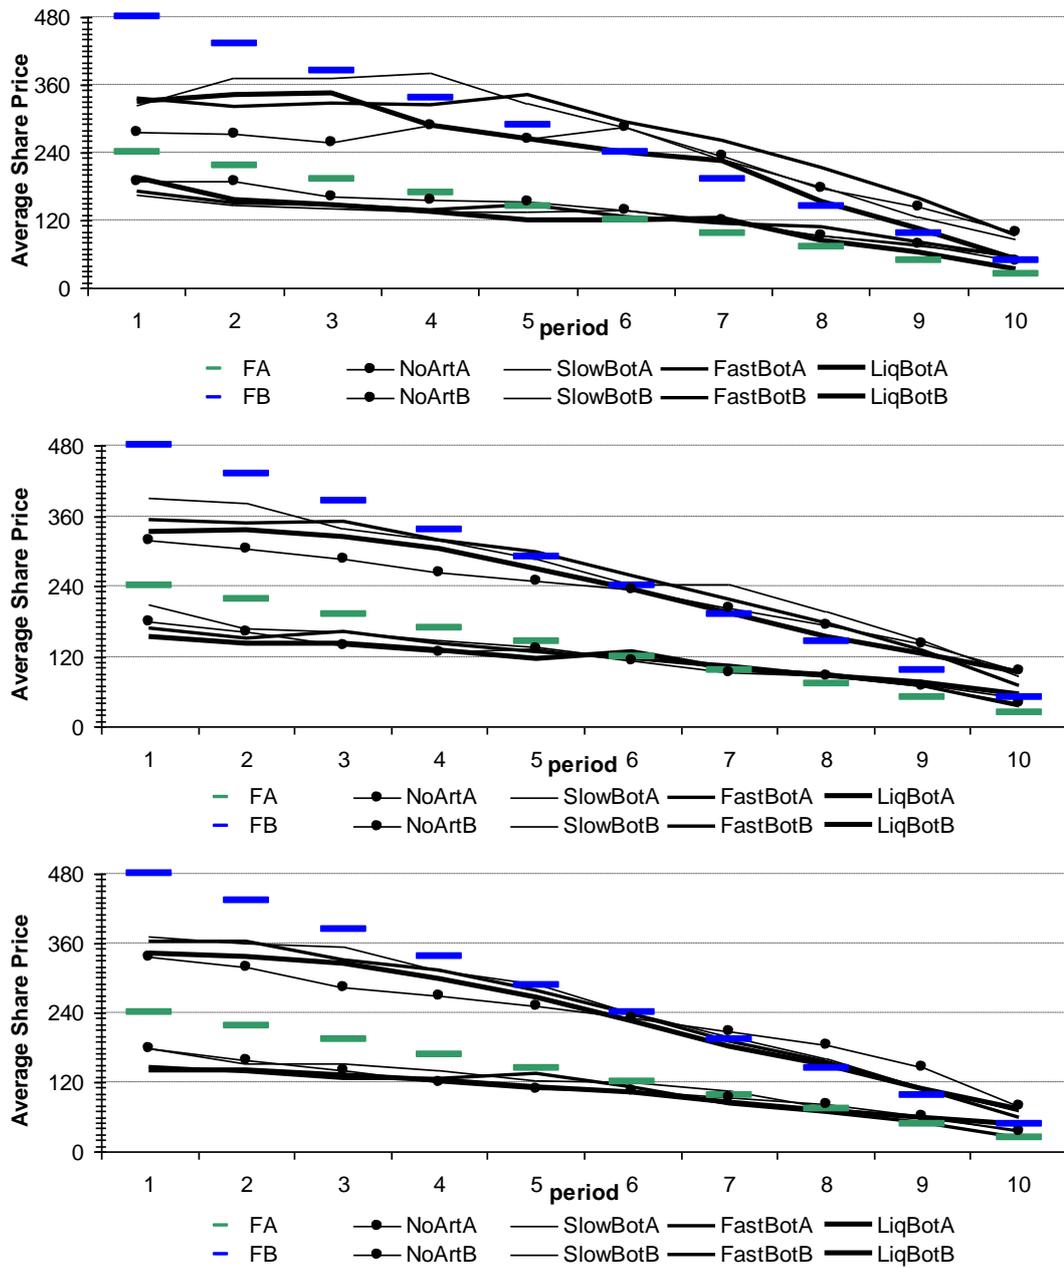


Figure A1. Average asset prices in treatments NoArt, SlowART, FastART, LiquidART (Top: first market round; middle: second market round; ARTtom: third market round)

Appendix B Evaluating announcement effects

The NoART sessions involved two treatments which we call NoBot and Baseline. We kept the instructions the same across all treatments with one exception. In all treatments except Baseline we included the following *announcement* in the instructions, of potential ART presence,

“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 in the market.

Figure A2 exhibits the average price trajectories over all markets and sessions 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 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).

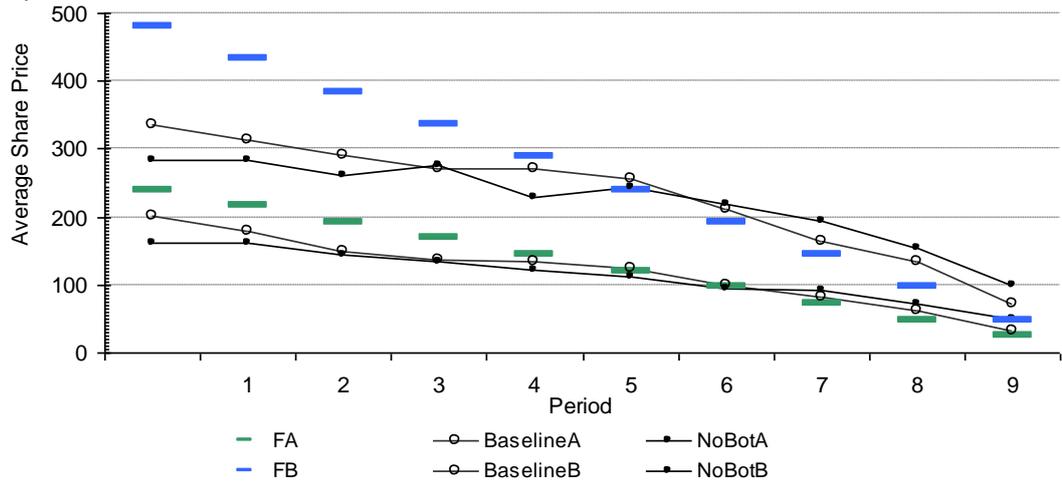


Figure A2. Average asset prices in treatments without ART participation, NoBot vs. Baseline

Observation Announcement: We find no measurable market performance impact following the announcement of potential ART market participation. Specifically, the levels of *PD* and *RAD* are not statistically different between the Baseline and NoBot treatments.

Appendix C – Market liquidity and Volatility

In this part, we report market quality measures for our data including liquidity and volatility.

(II-)liquidity measures

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. Abusing notation, 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 via market orders.

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 ψ_t . As fourth measure of liquidity, ζ_t , we count the number of outstanding orders per share at the end of the period, which is a measure of limit order intensity. This measure shows the number of potential transactions remaining in the market. The higher the transaction volume, the higher the order flow and the number of outstanding orders, the more liquid is the market.

Volatility measures

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 Table A2). The logarithmic price-value change is as follows;

$$(3) \quad \ln \Delta \frac{P_t}{F_t} = \ln \frac{P_t}{F_t} - \ln \frac{P_{t-1}}{F_{t-1}},$$

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. This volatility measure is on the average price across periods.

Table A1: Liquidity and volatility measures

(4) (Volume) number of transactions, number of limit orders, outstanding orders	V_t, ψ_t, ζ_t
(5) (bid-ask) spread	$S_t = \frac{P_{At} - P_{Bt}}{(P_{At} + P_{Bt})/2}$
(6) Volatility	$\sigma_t = \sqrt{\text{VAR}(\ln \Delta \frac{P_t}{F_t})}$
(7) Transitory volatility	$\sigma_\tau = \sqrt{\text{VAR}(\ln \Delta p_\tau)}$

We also examine transitory (or short-term) volatility through which we attempt to measure price impacts of the transactions within a period. Fast traders can obviously

have an impact on transitory volatility if they are involved in transactions. We measure the impact on the price changes as suggested in Table A1. We compute the price growth, $\ln \Delta p_\tau$,¹⁹ between each two prices of the same asset within the period (but not between periods) and take the standard deviation of all price changes in the session.

We believe that liquidity will be impacted in our FastART and SlowART treatments relative to the NoART treatment, because the ART 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 and a negative impact on outstanding orders. Given the liquidity demand of our market order ART, a spread increase could be the outcome, but that is not clear. On the other hand, we expect that the LiquidART helps to increase the number of outstanding orders and positively impacts order flow and transaction volume simply because the number should increase if there is an additional (automatic) submitter of limit orders in the market.

Observation Liquidity 1: The SlowART treatment does not affect the per-period transaction volume or the number of limit order submissions compared to the NoART treatments.

Support: Table A2 records the average number of transactions and the average number of limit order submissions in each period and asset. The Baseline, NoBot, and SlowART 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.

¹⁹ $\ln \Delta p_\tau = \ln p_\tau - \ln p_{\tau-1}$, $\tau > 1$, where the index indicates the τ -th transaction within a period and p_τ the transaction price.

Table A2. Liquidity: transactions, limit orders per asset and period, and spread

Treatment	Volatility σ_t	Transitory volatility σ_τ	Transaction volume V	Submitted Limit orders ψ	Outstanding limit orders at end of period ζ	Spread measured at end of period ^{a)}	
						Ask-Bid spread, S_t	Spread exists %
NoART:	0.254	.194	9	30	11	.717	96
Baseline	0.245	.194	9	30	11	.669	96
NoART	0.263	.194	9	30	11	.765	96
ART:							
SlowART	0.211	.195	9	31	11	.668	95
FastART	0.211	.182	13	31	9	.757	95
LiquidART	0.231	.175	14	36+15 ART	15-3 ART	.672	100
Two tailed Mann-Whitney test results (z-stat [p-value]):							
NoART – ART	0.359 [.720]	1.877*** [.061]	-1.830* [.067]	-1.051 [.2939]	0.442 [.659]	0.221 [.825]	-0.372 [.710]
NoBot – Baseline	-0.735 [.462]	0.000 [1.00]	0.053 [.958]	-0.105 [.916]	-0.105 [.916]	-1.155 [.248]	0.112 [.911]
NoART – SlowART	0.490 [.624]	-0.122 [.903]	0.277 [.782]	-0.031 [.976]	0.000 [1.00]	0.490 [.624]	0.420 [.675]
NoART – FastART	0.490 [.624]	1.592 [.111]	-1.664* [.091]	-0.460 [.645]	0.980 [.327]	-0.306 [.756]	0.708 [.480]
NoART – LiquidART	-0.184 [.854]	2.694*** [.007]	-2.686*** [.007]	-1.843* [.065]	0.000 [1.00]	0.306 [.756]	-2.071** [.038]
SlowART – FastART	-0.105 [.916]	1.890* [.059]	-2.240*** [.025]	-0.105 [.916]	1.260 [.208]	-0.735 [.462]	0.327 [.743]
SlowART – LiquidART	-0.420 [.674]	3.361*** [.001]	3.275*** [.001]	-1.163 [.245]	0.420 [.674]	0.000 [1.00]	-2.122** [.034]
FastART – LiquidART	-0.630 [.529]	0.840 [.401]	-0.797 [.425]	-0.686 [.493]	-1.680* [.093]	0.945 [.345]	-2.176** [.030]

Note: a) 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 Liquidity 2: The FastART treatment has a positive effect on transaction volume, while at the same time leaving the number of limit orders unaltered vis-à-vis the NoART treatment. The LiquidART algorithm generates the same transaction volume as the FastART, while at the same time increasing the number of limit orders.

Support: We observe, as reported in Table A2, more transactions per period and asset in the FastART and LiquidART treatments (on average 13 and 14 respectively) than in the other treatments (9 each in SlowART and NoART

treatments).²⁰ 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 FastART treatment is 31, virtually the same as the one in the SlowART and NoART treatments. Surprisingly the ARTs which only utilize market orders do not absorb liquidity in a way that impacts our standard measures. The LiquidART 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.²¹ It suggests that algorithm trading enhances the order book resiliency, i.e. the replenishment of limit orders after transactions.

Observation Liquidity 3: The LiquidART treatment more effectively maintains a bid-ask spread as measured at the end of trading periods. Bid-ask spreads are the same in all treatments.

Support: In the last two columns of Table A2 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 LiquidART 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.

Observation Liquidity 4: The number of subjects' outstanding limit orders (at the end of the period) is not larger in the LiquidART and FastART treatments than in

²⁰ The differences between the FastART 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.

²¹ 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.

the other treatments, but in the LiquidART treatment the number of outstanding limit orders is larger than in the FastART treatment.

Support: In the LiquidART treatment more limit orders are outstanding on average at the end of the period than in the other treatments (result of a one-tailed Mann-Whitney test as reported in Table A2). Subtracting the number of outstanding algorithmic limit orders, however, the result remains significant with respect to the FastART treatment only (indicated by the parenthesis around the asterisk in the table.)

Comment: The above reported increase in the number of human limit order submissions (observation liquidity 2) could be an indication of competitive bidding. From observation liquidity 4 we see, however, that the increased number of limit orders is probably an artefact of the experimental design. The increased number of limit orders apparently follows from the automatic post-transaction cancellations and subjects' replacement submissions. There are no indications of a higher number of competitive spread-splitting limit orders that respond to the actions of the algorithmic traders. The unaffected bid-ask spread across treatments (observation 8) underlines this finding.

Observation Volatility: Price volatility across periods is unaffected by treatment variation, but algorithmic trading impacts transitory volatility.

Support: We report the calculated price volatility measure for each treatment in the second column of Table A2. 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 A2. The data reported in the third column shows that the treatments FastART and especially LiquidART have smaller levels of transitory volatility. The magnitude of price changes within a period are on average smaller than in the other treatments.

Appendix D

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 ARTh 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 ARTom 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 $\text{£}1 = 200 \text{ CU}$. 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