23. Algorithmic trading in experimental markets with human traders: A literature survey[1]

Te Bao, Elizaveta Nekrasova, Tibor Neugebauer, Yohanes E. Riyanto

1. INTRODUCTION

The financial world has witnessed a skyrocketing volume of algorithmic (algo or robot) trading since the beginning of the 2000s. Today most transactions in financial markets are executed by automated trading systems. According to Treleaven et al. (2013), algorithmic trading already accounted for more than 70 percent of the US stock trading volume in 2011. Flash crashes, which until 6 May 2010 were unprecedented phenomena of extreme short-term volatility triggered by high-frequency algorithmic trading (Kirilenko et al. 2017), prominently demonstrate that algorithms have radically changed the financial market environment. It seems fair to say that, today, without understanding the impact of algorithmic trading, a thorough understanding of market behavior would become almost impossible.

To understand the market behavior, we need answers to many questions. For example, what are the impacts of the usage of trading algorithms on market quality, what would be the return to such algorithms and how will they affect human traders' performance and behavior? How do humans, especially individual investors, respond to trading in markets dominated by algorithms; does interaction with algorithmic traders generate emotional or psychological responses in human traders that impact market sentiment and thus price fluctuations? What type or features of algorithms are either helpful or rather harmful to investors; what kind of regulation is sensible? We contribute to this research by reviewing the answers given to some of these questions by the experimental literature focusing on the interaction between algorithms and humans in laboratory markets.

Related literature reviews discuss algorithmic trading in the real-world financial marketplace. Kirilenko and Lo (2013) provide an initial review of algorithmic trading in the real-world financial marketplace. The authors acknowledge types of automated trading, including passive strategies such as market making, arbitrage trading and more aggressive high-frequency trading. They recount important historical events in the age of machine trading, including the 2010 flash crash and cases of high-frequency trading manipulation such as spoofing. Finally, they reflect on potential regulatory measures, particularly in view of the presence of high-frequency trading algorithms, including speed bumps and Tobin taxes. Goldstein et al. (2014) provide a literature survey on algorithmic trading, including theory and studies based on real-world data. Miller and Shorter (2016) survey recent developments in high-frequency trading strategies, focusing on recent efforts in regulatory measures. Beckhardt et al. (2016) provide a broad survey on high-frequency trading strategies, including simulation analyses of profitability. We refer the interested reader to the literature reviews mentioned above as these issues go beyond the scope of the current review, which is limited to controlled laboratory studies.

A related area of interest is agent-based modeling, that is, computer simulations where algorithms interact with algorithms. Duffy (2006) provides an excellent survey. He summarizes the literature on zero-intelligence agents, learning and evolutionary algorithm models of agent behavior. Duffy also reviews the literature that compares human laboratory results with simulation results. Brewer (2008) and De Luca et al. (2011) also review zero-intelligence agents and their extensions. In the following section, we review some relevant algorithms for the interaction with humans discussed in that literature.

Closer to us in terms of coverage, March (2019) provides a broad literature review on the interaction of computer players with human subjects, including experiments on strategic reasoning, social dilemmas, markets, auctions, and bargaining and negotiation, among other topics. Naturally, this literature review has some overlaps with March's review, notably with his section on market experiments. Nevertheless, our focused approach allows a more detailed report of the studies in question, and we also include unpublished work.

This chapter provides an overview of the experimental literature on algorithmic trading in experimental financial markets, focusing on human-robot interaction.[2] The reported research is interdisciplinary. The interaction of man and machine is of general interest to the behavioral sciences and the computer sciences. The findings of this research can have implications for regulation. That said, the laboratory research that we report here is nascent. Based on the literature survey, we propose, without the ambition of being conclusive, some interesting questions for future research in this area and possible policy implications.

The remainder of the chapter is organized as follows. Section 2 reviews the literature on experimenter-induced algorithms that concentrate on the profitability of algorithms versus human traders, studies that look at market quality, the behavioral effects of algorithm response speed, manipulation, arbitraging activity and the question of subjects' aversion to interacting with algorithms. In Section 3, we review studies in which the experimenter leaves in human subjects' hands the decision to employ algorithms for trading. In Section 4, we finally conclude and discuss future directions.

2.   ALGORITHMS IN THE HANDS OF THE RESEARCHER

In this section, we review the literature on market experiments involving subjects playing the role of traders competing against algorithms programmed by researchers. In Section 3, we turn to studies in which subjects can choose to use an algorithm for trading. As a guide to the reader, Table 23.1 provides a non-exhaustive overview of the referenced algorithmic traders and their meaning.

Loosely following a historical perspective, we begin by briefly reviewing the literature in which the efficiency of markets populated by algorithms is compared to experimental markets with only human subjects before we look at the performance in hybrid markets, that is, markets in which human subjects interact with algorithms. We dedicate a sub-chapter to discuss the effects of the algorithm's reaction speed, another to discuss arbitrage in multiple markets, and yet another to cover the manipulation with algorithms. Finally, we turn to how the announcement of possible market participation of an algorithm can impact human subject behavior.

Table 23.1   Overview of algorithmic traders and main features

Algorithmic trader	Main features
Zero intelligence (ZI)	Algorithm randomly submits orders subject to minimal budget constraint without profit maximization (Gode and Sunder 1993).
Kaplan's Sniping Agent	The algorithmic sniper seller (buyer) sends a limit order to sell (buy) at the market best bid (offer) if at least one of three conditions is met; the best bid (offer) is at least as good or better as the high (low) transaction price of the previous period; the bid-offer spread is small ($\leq$ const$_s$) while the expected profit is more than a minimum profit factor ($\geq$ const$_p$); few instances of time left until the closing of the market period ($\leq$ const$_t$) (Rust et al. 1994).
Zero-intelligence plus (ZIP)	ZI algorithms with adaptive profit margin which is defined as the difference between the target price and the submitted price. The profit margin is increased/decreased by the algorithm after successful/failed transactions (Cliff and Bruten 1997).
Gjerstad/Dickhaut − GD agent	Algorithm submits orders to the market that maximize its expected surplus based on an updated belief distribution. The GD agent forms a belief about an offer or bid being accepted at price p based on the recent market history of accepted and unaccepted (including inframarginal and extramarginal) orders at that price (Gjerstad and Dickhaut 1998).
Arbitrage trader	Algorithm risklessly exploits mispricing across markets of perfectly correlated assets (for example, Angerer et al. 2019).
Spoofing	Algorithm aims at manipulating market sentiment by placing a huge limit order (which is not attempted for transaction and which is quickly cancelled upon attempt) while simultaneously trading in the opposite direction (for example, Leal and Hanaki 2018).
Manipulator	Algorithm, aimed at manipulating demand and supply, repeatedly buys shares by overbidding the best outstanding bid and subsequently sells the shares by undercutting the best offer (Veiga and Vorsatz 2009, 2010)
Market maker	Algorithm submits bid-offer spreads around the target price (for example, Aldrich and López Vargas 2020)
Reactionary robot	Algorithm reacts to a (range of) specified limit orders by submitting a market order as soon as a specified order enters the market (Asparouhova et al. 2020)

2.1   Comparison of Algorithms in Simulations with Human Traders in Experimental Markets

The early experimental studies involved no interaction between algorithms and human subjects. These studies compare the outcomes of interactive experiments with human subjects to the ones of interacting algorithms in the continuous double auction (hereafter CDA) markets. The focus of these studies is typically on market efficiency or trader performance.

The documented academic research on algorithms in asset markets seemingly started in the late 1980s. Sunder (2003, p. 10f) recounts his approach: the press blamed the stock market crash of 1987 on algorithmic trading. Skeptical of this claim, Sunder designed and taught a course at Carnegie Mellon University on algorithmic trading to learn about the structure of trading strategies and the behavior of the CDA market. Being challenged by the students in the course, he and Dan Gode programmed a random algorithm − later labeled 'zero intelligence or ZI traders' (Gode and Sunder 1993), which adheres to a budget constraint. The chosen CDA market environment followed the one in Smith (1962), with induced values and costs. ZI traders provide liquidity to the market by repeatedly submitting orders; ZI buyers submit random bids between 0 and the experimenter induced value; ZI sellers submit random offers between the upper bound of the cost distribution and the induced cost level. Since traders have zero intelligence, they do not profit maximize, remember or learn. In Gode and Sunder (1993), a ZI transaction occurs whenever the best bid exceeds the best offer (each for one unit), the transaction price being equal to the earlier submitted one of the two.

The result was that ZI agents achieved an allocative efficiency of 99 percent across different sessions, comparable to the one found in data from experiments with human subjects. Gode and Sunder (1993, p. 134) conclude that 'the high allocative efficiency of double auctions is [caused by] market discipline imposed on traders' and not by profit maximization, learning and intelligence. This study, in particular, and more generally the literature on zero-intelligence traders (summarized in Duffy 2006; De Luca et al. 2011) provides a very important step towards a micro-foundation of general equilibrium theory by showing that market efficiency does not rely on perfect individual rationality and utility maximization behavior. Nonetheless, the trajectories of equilibrium market prices with human subjects are relatively flat, whereas those with ZI agents produce continued volatility around the equilibrium price.

Cliff and Bruten (1997) show that with ZIP agents, which are ZI agents with an adaptive profit margin (see Section 2.2.), the market price convergence ~~trajectories are~~ similar to human behavior.

In a more complex CDA market design with a multi-period lived asset, which frequently generates bubbles and crashes in laboratory studies (that is, the design of Smith et al. 1988),[3] Duffy and Ünver (2006) also report price and volume paths similarities of their enhanced 'near' ZI-trader markets with the human-trader experimental results of Smith et al. (1988). In particular, the authors are able to reproduce bubble and crash patterns in the simulated 'near ZI' market. Different from the ZI of Gode and Sunder (1993), the 'near ZI' trader's bids and offers are not purely random but are biased towards the past period's average price.

Arifovic (1996) finds in an experimental macroeconomic setting that the market price behavior of human experimental subjects shares similarities to that of a genetic algorithm.[4] The genetic algorithm selects a decision rule which is updated from one generation to the next using the three genetic operations in offspring generation; reproduction, crossover and mutation.[5]

Rust et al. (1994) report on the Santa Fe Institute double auction tournament − SFDAT − in 1990/91. For the SFDAT, 30 colleagues submitted the strategies of profit-maximizing algorithms, including quite complex ones, to trade with another in the Smith (1962) CDA market. To their surprise, the tournament winner involved relatively simple liquidity absorbing (profit-making) strategy − later labeled Kaplan's Sniping Agent. The Sniper seller (buyer) sends a limit order to sell (buy) at the market best bid (offer) if at least one of three conditions is met. They are, *firstly*, the best bid (offer) is at least as good or better than the high (low) transac-

tion price of the previous period; *secondly*, the bid-offer *spread* is small ($\leq const_s$) while the expected *profit* is more than a minimum profit factor ($\geq const_P$); and *thirdly*, few instances of time left until the closing of the market period ($\leq const_t$). Later simulation studies highlighted that this Kaplan's sniper could only be profitable if few agents apply it, as this trading strategy does not play a best response against itself.[6]

   In conclusion, markets inhabited with algorithmic traders, even with the simplest algorithms, show convergence to the equilibrium similar to the one in the markets of human traders only. This conclusion validates the robustness of the continuous double auction market institution in terms of price discovery and equilibration. However, in the real world, algorithms interact not only with other algorithmic traders, but also with human traders, and the interaction of humans and algorithms in hybrid markets can lead to frictions. The remainder of the chapter focuses on hybrid experimental markets in which algorithms and human subjects transact assets and cash.

2.2   Performance of Algorithmic and Human Traders in Hybrid Experimental Markets

Das et al. (2001) study how agent-human interaction influences human traders' market outcome and trading performance in a hybrid experimental asset market setting within a CDA environment with induced values (Smith 1962). In each of their experimental markets, there are 6 human traders and 6 algorithmic traders (3 buyers and 3 sellers each).[7] Depending on the treatment, the algorithmic trader may adopt one of two types of adaptive trading strategies: (1) the 'zero-intelligence plus (ZIP)' algorithm (Cliff and Bruten 1997) provides liquidity to the market, similarly to ZI. Still, its orders involve a private profit margin updated over time if a limit order either fails to transact or transacts immediately. When a trade takes place, all agents adjust their bids towards the transaction price. If no trade occurs in t seconds,[8] all agents adjust their bids to improve the best existing bid. (2) The GD algorithm (Gjerstad and Dickhaut 1998) submits orders to the market that maximize its expected surplus based on an updated belief distribution. The GD agent forms a belief about an offer or bid being accepted at price p based on the recent market history of accepted and unaccepted (including inframarginal and extramarginal) orders at that price. Das et al. (2001) find that compared to CDA markets with all-human design or all-algorithm design, the market price shows slower convergence to the equilibrium price in their experimental hybrid market. Meanwhile, human traders underperform algorithms by about 20 percent in trading surplus.

   Gjerstad (2007) studies how different market structures and paces of submitting bids and offers influence the trading performance of humans and the GD algorithm in a CDA market with induced values (Smith 1962). There are 6 buyers and 6 sellers in the experimental market. In the hybrid markets involving interaction between human and GD agents. The experiment involves both unbalanced markets (with 6 algorithmic traders on one side and 6 human traders on the other side) and balanced markets (where 3 human buyers/sellers and 3 automated buyers/sellers are on each side of the market). Interested in submission pace, the author differentiates treatments between 'patient' and 'impatient' algorithmic traders regarding waiting time before submitting a new order. Patient traders submit bids and offers at a slower pace than impatient ones. The result of the paper shows that all markets achieve a very high level of efficiency (usually more than 99.5 percent). Meanwhile, impatient algorithmic traders' profitability seems to be lower compared to the patient ones. If algorithmic buyers/sellers are too actively submitting new limit orders, the price moves up/down quickly thus adversely

impacting their profits. In general, the profit of patient algorithmic traders is highest, followed by the impatient ones, and human traders' profit is lowest. The latter result is impacted by the low performing human sellers; human buyers and automated buyers have a similar performance. Unbalanced markets result in greater differences between the performance of humans and agents than balanced markets.

Peng et al. (2020) study balanced, unbalanced and uncompetitive markets with 'fast' and 'slow' ZIP agents. The authors confirm the result of Das et al. (2001) that algorithms outperform human traders in the 'fast' treatment in a balanced market. The result seems to be due to 'fast' ZIP buyers benefitting from low-priced offers faster than human buyers. In all other conditions, contrary to the result of Das et al. (2001), Peng et al. report that human traders outperform 'fast' algorithmic traders. Again, these results are obtained for the simplistic CDA market with induced values (Smith 1962).[9]

In a more complex environment, in which earnings depend on the share price at period end, Feldman and Friedman (2010) study human-algorithm interaction in a hybrid experimental CDA market. The studied algorithmic traders are adaptive optimizers adjusting their portfolio composed by a riskless asset and a risky asset in accordance with their changing payoff expectation. Their experimental treatments vary in the composition and the size of markets. Human traders interact with algorithmic traders in large markets (1 human and 29 robots or 5 human and 25 robots) and small markets (5 human traders and 5 robots). The key findings of their study include: (1) the average trading gain of human traders is generally smaller than of algorithmic traders, but human traders may outperform algorithmic traders in market crashes; (2) human traders tend to destabilize small markets and neither stabilize nor destabilize large markets; (3) ~~human traders respond to the payoff gradient similarly~~ as the algorithmic trader. In their study, it is interesting to note that human traders earned higher profits during crashes (that is, lose less with extreme market volatility) and tend to sell faster after experiencing a loss, although generally exhibiting similar trading behavior as the algorithms.

Tai et al. (2018) let one human subject interact in CDA markets populated with ZI traders or with adaptive algorithmic traders of SFDAT, including Kaplan's Sniper, GD and ZIP. Surprisingly, subjects' earnings are higher in the treatment with adaptive algorithmic traders than with ZI traders. The authors conjecture that subjects' cognitive working memory capacity impacts their trading acuity and test this hypothesis in asymmetric and symmetric CDA markets of Smith (1962) type.[10] The result confirms the hypothesis; subjects with high elicited working memory capacity earn higher profits than those with low elicited working memory capacity; the difference is pronounced in the more complex environments, that is, asymmetric markets and adaptive agents.

Akiyama et al. (2017) implement an algorithm that trades on fundamentals in a Smith et al. (1988) call-auction asset-market design involving belief elicitation on future prices. The authors propose two treatments to study the question of strategic uncertainty as a cause for bubbles and crashes: treatment with 6 human subjects and treatment with 1 human trader and 5 algorithmic traders committing transactions at fundamental value. In the second scenario, strategic uncertainty is eliminated while participants have perfect information about the algorithm's presence and its performed strategy. The results of Akiyama et al. (2017) suggest that strategic uncertainty might partly explain observed mispricing in this market. Using the same experimental setting, Hanaki et al. (2018) show that traders' performance is negatively correlated with their confidence in their short-term price forecast. In a related study, Ahrens et al. (2020) also use this experimental design with the fundamentalist algorithm to investigate subjects' overconfidence

in their price forecast to find that the level of overprecision (that is, the narrowness of the predicted confidence interval) may be endogenously determined or influenced by the observed market price dynamics. It tends to go up (down) when the asset price goes up (down).[11]

To conclude, experiments in hybrid markets shed light on the limitations of both algorithmic and human traders. On average, algorithmic traders gain higher profits than humans particularly in early trading phases,[12] but human traders may learn and adapt more quickly to extreme volatility and more complex market environments than simple algorithms. Most of the reported results were achieved in a convergent environment (Smith 1962). It is an open question for further research, how 'near ZI' traders (Duffy and Ünver 2006), GD traders or snipers would perform in a more complex hybrid CDA market design such as in Smith et al. (1988), which is known for mispricing. It would be interesting to see if such algorithmic traders would rather have a stabilizing or destabilizing effect on the market.

2.3   Algorithm Speed

Faster than human response speed to profit from trading has been one of the main reasons for the adoption of algorithmic trading in asset markets, and therefore, it has been an innate research question how much algorithmic traders' profit from low latency, that is, the minimal response delay.

In the previous section, we have already suggested that fast algorithmic traders do not necessarily outperform slower algorithmic traders. Gjerstad (2007) found that patient GD agents perform better than impatient GD agents, and both perform better than human traders in the hybrid market.

Das et al. (2001) vary the ZIP algorithm's response speed, introduced by a 'sleep-wake cycle', to examine the interaction between humans and algorithmic traders. The 'fast' algorithm would be idle for $t = 1$ second and become active when a new quote or trade is made. The 'slow' algorithm would be idle for $t = 5$ seconds and only become active when a trade is made. When active, the algorithm would update its orders by submitting a new order or updating the existing order. Das et al. (2001) observe more transactions in the 'fast' than in the 'slow' ZIP treatment. In both treatments, algorithmic traders tend to trade among themselves first before trading with human traders and price trajectories converge to the competitive equilibrium. Das et al. (2001) suggest that algorithmic ZIP traders may outperform their human counterparts in a balanced human-algorithm market, but there are too few observations to draw any conclusion on the performance effect of speed.

Contributing data to this point, Cartlidge and Cliff (2013) study the effect of algorithm response speed in the hybrid market involving a further enhanced ZIP agent (called 'aggressive adaptive strategy'). They consider four different sleep-wake cycle treatments including the two time cycles of Das et al. (2001) and two extreme ones; $t = \{0.1, 1, 5, 10\}$. The data seem to suggest that the average algorithmic trader earns a lower positive margin over humans in the 'fast' treatment ($t = 1$) than in the 'slow' treatment ($t = 5$). However, the positive marginal gain of algorithmic traders over humans is not a monotonous function of sleep time t, as superhuman speed ($t = 0.1$) earns a high margin and the very slow speed ($t = 10$) a low one.

Peng et al. (2020) follow up on Das et al. (2001) to investigate the role of ZIP response speed in different market structures involving symmetric and asymmetric demand-supply schedules and balanced and unbalanced markets. Human traders outperform 'fast' ZIP traders in all conditions other than the balanced market, and humans outperform 'slow' ZIP traders in

all balanced and unbalanced competitive markets, that is, where the number of buyers equals the number of sellers. 'Slow' ZIP traders outperform human traders and 'fast' ZIP traders in situations of market power. Peng et al. report interesting convergence patterns to the competitive equilibrium; the trajectories seem biased in favor of the trader side that has market power or patience or both. For instance, if there is only one human seller and 6 'slow' ZIP buyers, price trajectories approach the equilibrium from above; or if there are 6 'fast' ZIP sellers and 6 human buyers, price trajectories approach the equilibrium from below.

Cartlidge et al. (2012) conduct a series of laboratory experiments assessing the role of algorithms' super-human speed in market efficiency and their performance in an environment where human traders and algorithms interact in the market. They find that the purely simulated market inhabited by slower algorithms, whose trading speed resembles the speed of human traders, seems to converge closer to a competitive equilibrium with enhanced market efficiency. Also, in a hybrid experimental market where human traders interact with algorithms, Cartlidge and Cliff (2013, 2018) investigate the impact of the millisecond-by-millisecond speed of the stock price movement. They argue that there is a price-movement-speed threshold above which human traders can still engage in market transactions and trade with human and algorithmic traders. Below the threshold, latency is too low for human traders to react, so humans can no longer participate in the market. Essentially, the threshold is a tipping point that creates a phase transition from a mixed human-algorithm phase to a algorithm-algorithm phase so that algorithmic traders interact with other algorithms instead of algorithmic traders interacting with human traders. Cartlidge and Cliff (2013, 2018) coined this as a robot-phase transition or market disintegration. They also find that very fast algorithmic traders can cause lower market efficiency besides market disintegration.

In summary, the effects of the algorithm's response speed are mixed depending on the timing of the sleep-wake cycle, on the strategy of the algorithmic traders and the market structure. Algorithmic traders that trade at the response time of humans may perform better than algorithms that respond faster or slower. The robot-phase transition described by Cartlidge et al. (2012) and Cartlidge and Cliff (2013, 2018) poses interesting questions for further research. The lower market efficiency observed by the authors in the market disintegration case contrasts with the results obtained in the previous studies without interactions between algorithmic and human traders. Thus, it seems that the inclusion of fast algorithms into the experiment might help to explain flash crashes happening in real-world markets.[13] Another interesting direction marked by the existing research is the adaptability of human and algorithmic traders to the market conditions: increased volatility, unbalanced or uncompetitive markets or non-converging markets, and, of course, to multiple simultaneous markets.

2.4   Arbitrage Algorithms

In real-world exchanges, financial assets are traded in fragmented markets because the regulatory authorities seek to enforce competition among exchanges to avoid monopoly fees for transactions. Market fragmentation can lead to situations in which an identical asset is demanded or offered at different prices at different venues, thus creating an arbitrage opportunity (for example, see Figure 23.1a). Algorithms can also provide arbitrage price discrepancies between an exchange-traded index fund and the assets that compose the index. Similar price discrepancies can arise with two or several different exchange-traded funds based on the same index or between a derivative financial contract and the underlying asset. Automation

is usually much faster at exploiting arbitrage opportunities than manual transmissions and, therefore, arbitrage algorithms have been among the most frequently applied algorithmic traders in financial markets (Kirilenko and Lo 2013).

Harrison (1992) studies an 8-period lived asset with imperfect payoff information. Including two one-period-ahead futures markets, for period 4 and period 8, he implements an algorithm that arbitrages between spot and futures CDA market (in treatment 4). Harrison (1992) concludes that arbitrageurs could be crucial for ensuring the spot market's informational efficiency and help to constrain the length of any mispricing in spot prices in the study.

Angerer et al. (2019) study algorithmic arbitrage in the setting of Charness and Neugebauer (2019), which allows for trading in twin markets of the Smith et al. (1988) type. The dividends in the two markets A and B are perfectly correlated modulo a shift, that is, the B-share pays in each period the same dividend as the A-share plus a fixed payment of 24 cash units. The authors investigate two liquidity absorbing algorithmic arbitrage traders called FastBot (see Figure 23.1a) and SlowBot, the liquidity providing algorithmic arbitrage trader LiquidBot (see Figure 23.1b), and two control treatments, that is, NoBot (in which the potential participation


Figure 23.1   Liquidity (a) taking and (b) making algorithmic arbitragers in Angerer et al. (2019) P, B, O,   F, ΔF and ξ denote price, bid, offer, fundamental value, the difference in fundamental value and a random variable with support on the interval [0, F/2], respectively. The algorithmic trader exploits an arbitrage opportunity by selling high and buying low an identical claim of cash flows transacting at prices P_A and P_B. The sequence of events is numbered 1–4; 5 indicates the size of the arbitrage gain

of an algorithm is announced, but no algorithm participates) and Control (with no announcement and no algorithm). The FastBot arbitrage trader immediately exploits arbitrage opportunities in real time when they arise, while the SlowBot arbitrage trader trades with a delay. The study suggests that algorithmic arbitrage traders moderate the extent of mispricing. The algorithmic arbitrage traders help to approximate the law of one price and marginally amend the discovery of the fundamental value. The market quality is generally enhanced. Volatility is lower, transaction volume higher, and, particularly in the LiquidBot treatment, liquidity is enhanced relative to the NoBot treatment. The arbitrage traders reap some earnings from human subjects upon transaction by design.

Nonetheless, subjects' earnings are not significantly lower compared to the treatments without algorithms. Interestingly, the SlowBot algorithm amends market efficiency similar to the other two algorithms, although it earns only a fraction of what the other algorithms earn. Finally, the authors find no announcement effect (see the following subsection) comparing the treatments Control and NoBot.

Neugebauer et al. (2020) test the Modigliani-Miller theorem of dividend policy irrelevance[14] involving a FastBot algorithmic arbitrager (as in Angerer et al. 2019) and the trading of two 4-period lived assets in a complete asset market. Each asset pays a dividend at the end of the period, which is drawn without replacement from a set of four dividends. After the four regular dividends, shareholders receive a liquidating dividend which is high or low with equal probability. Owed to the fact that the remaining regular dividends are known, the difference in the fundamental value of the two assets is known in each period. Hence, if order in one market crosses the spread in the other market, an arbitrage opportunity arises (step 2 in Figure 23.1). In the treatment with the algorithmic arbitrager, such arbitrage opportunities are immediately exploited. The result of the study is that the law of one price (and thus dividend policy irrelevance) holds with and without arbitrager if dividend streams of both assets are identical. If dividends are not identical, the Modigliani-Miller theorem of dividend policy irrelevance can only be supported in the presence of (and must be rejected without) the algorithmic arbitrager. Hence, the result of the study adds laboratory evidence that an algorithmic arbitrage traders may amend mispricing.

Rietz (2005) studies index arbitrage in a 15-times repeated one-period CDA setting. At the beginning of each period, subjects are endowed with green and blue assets in a prediction market. One of the assets generates a dividend of \$0.50 and the other a dividend of \$0.00. The dividend-paying asset is determined by drawing from a bag with 14 green and 6 blue balls at the end of the trading period. Hence, the fundamental dividend value for the green asset is \$0.35 and \$0.15 for the blue asset, and predicted relative prices are \$0.15/\$0.35. During the period, subjects trade green and blue assets for cash, and subjects can buy a bundle containing one green and one blue asset from the experimenter or sell the bundle to the experimenter for the bundle's dividend value of \$0.50 in cash. Arbitrage opportunities arise whenever the sum of bids (offers) for the two assets totals more (less) than the bundle value. If such an opportunity arises, the arbitrage trader exchanges the bundle for the two assets. In the treatment with the arbitrage trader, subjects are informed about its functioning in the instructions. The results of the study are as follows: the arbitrage trader is involved in most of the transactions, and transaction volume and volatility increase significantly; prices drop relative to the treatment without arbitrage trader, where prices usually are above fundamentals; but relative prices are driven away from their predicted value, with prices of blue assets above and green assets below their fundamentals; and individuals holding less diversified portfolios. Hence, this evidence

suggests that the arbitrager supports the law of one price but not always aids market quality and the price discovery of single asset fundamentals.

Grossklags and Schmidt (2006) also study arbitrage in a prediction CDA market. Differing from Rietz (2005), where the bundle involves two securities, Grossklags and Schmidt's bundle involves five securities and increased complexity. The algorithmic arbitrage trader is involved in about every fifth transaction. Surprisingly, mispricing in terms of the law of one price is not enhanced with algorithmic arbitrage. Even more surprising, in one treatment, the algorithmic arbitrage trader's presence is not announced, and in that treatment, mispricing is significantly worse than without the participation of the arbitrage trader.

Berger et al. (2020) study latency arbitrage in a repeated CDA market for a one-period lived asset with induced values, hence similar to Smith (1962) but with challenges to price discovery. In this setting, an algorithmic 'HFT' trader, that is not announced, basically front-runs incoming orders to book an immediate gain. The first one, labeled directional trading algorithm, realizes an immediate gain via the submission of two orders within the queue when a subject submits a market order; for example, if two offers to sell are outstanding at 100 and 101 and a market buy order is submitted, the algorithm buys at the best offer of 100 front running the market buy order to which it sells at 100.9 just a point below the second-best offer price. The second one, labeled arbitrage algorithm, front runs any incoming limit order that crosses the spread realizing an immediate gain (such as $P_B - P_A$ in Figure 23.1a, but within one market) through two transactions. For example, if one offer to sell is outstanding at 100 and a bid at 101 is submitted to the market, the algorithm buys at the best offer of 100 and sells to the incoming bid at its limit of 101, thus realizing an immediate gain of 1 cash unit. Berger et al. (2020) report market quality enhancements, including an increase in transaction volume and bid depth of the order book, in the human-algorithm hybrid experimental market relative to the baseline market with human subject only.

In conclusion, most existing studies on algorithmic arbitrage in experimental markets suggest that algorithmic arbitrage traders enhance mispricing.[15] However the question of the social cost of arbitrage activity requires further research; particularly, if arbitrage is rather detrimental or beneficial for the market participants and the society. Exploitation of temporary arbitrage opportunities is the business of institutional investment including algorithmic trading companies; including but not limited to stock, index, or foreign exchange arbitrage.[16] The industry evidently extracts value from investors. Therefore, a careful evaluation of this activity is needed to see if the industry should be further regulated.

A relevant application of algorithmic arbitrage is theory testing, as financial economic theories are frequently built on the assumption that no arbitrage opportunity should exist in the market. Since experimental subjects seem not aware of arbitrage opportunities without being shown, algorithmic arbitrage can help to test theories built on no arbitrage assumptions in the laboratory. The observed unawareness of arbitrage opportunities is a matter of financial illiteracy and thus another question for further research.

2.5   Manipulation

Market manipulations have always been a concern of market participants. Putniņš (2012) surveys manipulative practices in real-world exchanges, the theoretical and empirical literature. A great advantage of laboratory experiments on market manipulations compared to real-world discovery is that the experimenter can unequivocally identify manipulation in real time and its effects of price distortion relative to fundamentals.

Leal and Hanaki (2018) address the HFT practice of market making and the manipulative practice of spoofing in a CDA market with long-lived assets (Smith et al. 1988). 'Spoofing' involves intentionally manipulating prices by placing an order to buy or sell a security and then canceling it shortly thereafter, at which point the spoofer consummates a trade in the opposite direction of the canceled order' (Kirilenko and Lo 2013, p. 66). Leal and Hanaki (2018) do not concentrate their analysis on the direct effects of spoofing and market making but report the effects on subjects' beliefs of the potential presence of such an algorithm.[17] We report their experimental design and results in the following subsection.

Veiga and Vorsatz (2009, 2010) investigate the impact on price distortions from manipulation (similar to a 'pump-and-dump' scheme, that is, an attempt to boost the price of the stock to sell it high) performed through an algorithm in an experimental hybrid market. Veiga and Vorsatz (2009) set up an experimental CDA market for an asset that generates a high or low payoff with equal probabilities to study information dissemination similar to Plott and Sunder (1982). The authors consider two treatments. In the control treatment and manipulation treatment, one third of the 12 market participants are privately informed with certainty about the asset value, whereas the others are uninformed, that is, have only prior information. In the manipulation treatment, subjects know about the presence of the algorithm but not its strategy. The manipulator algorithm is programmed to buy 10 shares out of 24 after the market is open for 25 seconds and subsequently to sell them back to the market until 50 seconds before market closing. To buy, the algorithmic manipulator overbids the best outstanding bid by a random number drawn from a small positive interval, thus pushing up the price. To sell, similarly, the algorithmic manipulator undercuts the best outstanding offer. After each transaction the algorithm removes all its outstanding orders and after a delay of a few seconds continues with the placing of a new order. The authors find that the algorithmic manipulator distorts price discovery when the asset's actual value is low. When the actual asset value is high, the algorithmic manipulator does not seem to distort price discovery.[18] Also, the algorithmic manipulator usually loses money as it is not programed for profitability.

In the follow-up study, Veiga and Vorsatz (2010) investigate manipulation in a CDA market with partially informed traders as in Plott and Sunder (1988) to study the impact of their algorithmic manipulator on information aggregation. In this set-up, the participants trade one asset, taking three possible values x, y or z with equal probabilities. In the first treatment, all participants are partially informed about the asset's value. There is no aggregate uncertainty as 6 subjects are informed that it is not x and the other 6 subjects are told that it is not y when the actual asset value is z. In the second treatment, 6 subjects are partially informed (no aggregate uncertainty), while the other 6 are uninformed. Finally, in the third treatment, similarly to Veiga and Vorsatz (2009), 1/3 of the participants are perfectly informed, while the others are uninformed. The authors report that manipulation has a lasting effect on price discovery only with perfectly informed insiders when the asset's actual value is low. Also, unlike in Veiga and Vorsatz (2009) the algorithmic manipulator earns at least a positive average payoff across human traders.[19] Veiga and Vorsatz's (2009, 2010) two laboratory experiments provide an argument in favor of the regulation obliging market insiders to disclose their transactions. Other experimental studies on manipulation do not use algorithmic traders but offer incentives to subjects to distort market prices (Hanson et al. 2006; Comerton-Forde and Putniņš 2011).[20]

In summary, manipulation in financial markets is an important issue. Since manipulation of market prices is illegal, real-world data on the effects of manipulation are difficult to obtain. There are only few experimental market studies addressing the effects of market manipulation

and the mixed evidence of these effects raises questions. Why is the price discovery of low values affected in Veiga and Vorsatz (2009, 2010), but not the price discovery of intermediate or high values after a run up in the price? Does the application of a pump-and-dump manipulation rather lead to inflated prices, or do inflated prices follow from the experimental design as short sales are banned or because subjects do not know about the strategy of the algorithm? Would human traders be able to counteract and to return the prices to fundamentals if they were familiar with the manipulator's strategy? Under what conditions can an algorithmic manipulator achieve profitability? How would a perhaps more sophisticated, profit-oriented algorithm implement a pump-and-dump manipulation, and what effect would it have on the market? These questions deserve further exploration.

Generally, it is surprising that experimental research on market manipulation is scarce. We think that this area of experimental research is very relevant and should attract more interest of experimental finance researchers. The laboratory seems to be a good place to study algorithmic manipulation.

2.6   Announcement Effect

Today, a person committing transactions in the financial market should reasonably expect an algorithmic trader as his or her counterparty. At the same time, the impact of an algorithm's presence or the possibility of its presence on humans' actions and expectations might be non-trivial. Thus, an important question regarding investor psychology is whether the possibility of interacting with an algorithm has a measurable influence on human behavior and the market. The evidence is mixed.

As pointed out above, Grossklags and Schmidt (2006) study a prediction market with an algorithmic arbitrager. The paper suggests that the announcement of the presence of algorithms facilitates price discovery raising the rate of price convergence to the equilibrium relative to the setting where the algorithm is present, but this presence is not announced. Within the experiment, three treatments are investigated: no algorithm and no announcement (baseline); algorithm and no announcement; algorithm and announcement. Overall, announcement leads to the increased informational efficiency, defined as deviations between prices and fundamental values, but at the same time, the algorithm's presence without announcement results in a decrease in the convergence rate in comparison to the baseline treatment. The authors explain that arbitrage algorithms tend to decrease the trading opportunities for humans, which results in a lower number of trades and distortion of the information aggregation process. However, when the presence of the algorithms is announced, subjects adapt their behavior by switching to more conservative trading strategies bidding closer to the fundamentals.

Farjam and Kirchkamp (2018) also suggest a positive announcement effect. Their subjects seem to behave more rationally following the announcement, bringing transaction prices closer to the fundamental value than without the announcement. The experimental design involves a six-subject CDA market with one multi-period lived asset (Smith et al. 1988). The study compares price deviations from the fundamental value across the two treatments: either subject is told that the algorithm may be present in their market or that the algorithm is not present. Meanwhile, no algorithm participates in the experiment. The authors align subjects' expectations by asking early participants to describe the algorithm and then sharing the prepared wordle[21] with the other subjects claiming that the algorithm is programmed based on this description.[22]

Leal and Hanaki (2018) suggest no announcement effect on prices but find an effect on the elicited first-period beliefs. The experiment involves three treatments that differ in the instructions only. The treatment human only (HO) makes no reference to algorithmic traders. In the instructions to the treatments spoofing (SP) and market making (MM), subjects receive the information that they may interact with an algorithmic trader in the market, and the general strategy of the algorithms MM and SP are explained. SP is supposed to be taking advantage of human traders, while MM is supposed to provide more liquidity to the market. Surprisingly, the result of the experiment shows little difference between the two types of market. The results suggest that in MM and SP, relative to HO, initial average price forecasts are higher and more volatile. Initial orders are submitted later. Besides these effects, the market price in MM and SP deviates more from the fundamental value than in HO.

Finally, as pointed out above, Angerer et al. (2019) find no announcement effect and no pricing difference relative to fundamentals in the CDA market study with two perfectly correlated assets. The authors compare their control treatment without the announcement of potential algorithm participation with their NoBot treatment in which the potential participation is announced, but no algorithm participates. Different from the studies above, no information is disclosed on the strategy of the algorithm.

To sum up, there is mixed evidence. A psychological phenomenon as the announcement effect seems to be dependent on the design of the experiment, including the format and the framing of the announcement, and possibly on the subject pool. For example, simplified visual forms, as for example wordle, possibly have a different impact on the subjects than verbal descriptions of the algorithm's strategy. Subjects' background or experience with algorithms may lead to a different response. For further research, it may be interesting to investigate the impact of various formats and framing of the announcement on the presence of algorithmic traders to market participants, and the development of the announcement effect over time in the laboratory.

3.   ALGORITHMS IN THE HANDS OF THE SUBJECT

While many experiments sought to treat human traders independently from algorithmic traders, Aldrich and López Vargas (2020) and Asparouhova et al. (2020) allowed their subjects to choose to employ algorithmic traders in market experiments.

Aldrich and López Vargas (2020) asked subjects to choose a predefined market maker or sniper algorithm or to exit the market to trade a single asset. In the former two cases, subjects also decide on costly improvement in roundtrip-messaging latency to and from the exchange. The market maker algorithm submits the subject's chosen, symmetric bid-offer spread around the fundamental value to transact with algorithmic noise traders that place market orders at random times to buy or sell one asset. The sniper is a predator algorithm that takes advantage of a market maker's delayed response to changes in fundamentals. At random times the fundamental value jumps and market maker and sniper algorithms adapt to random jumps in fundamental value at established latency. With a large jump, the outstanding spread of a market maker is shortly mispriced. The sniper algorithm attempts to take advantage of the mispricing transacting before and after the spread is readjusted. Aldrich and López Vargas (2020) consider two market environments: the CDA market and the frequent batch auction (FBA) format. The FBA refers to a clearing house mechanism where traders submit orders in continuous time

and a call auction clears at uniform price all crossing orders in discrete time intervals. The purpose of the study is to see if FBA, compared to CDA, leads to less wasteful investment in low-latency technology and less predatory behavior, measured by the number of snipers in the market. The authors conclude that FBA induces higher liquidity, price efficiency and less volatility than the CDA. Furthermore, FBA shows fewer snipers present in the market and causes a lower investment in low-latency technology than CDA.[23] In the CDA market, the algorithms produce permanent mispricing, and the authors report flash crashes in the first period.

Asparouhova et al. (2020) allow subjects to trade manually or deploy algorithms, and they are assumed to be aware of the potential presence of traders employing algorithms. The trading environment is a CDA market with the declining fundamental value of the underlying asset used in Smith et al. (1988). The algorithms either act as a *market maker* or a *reactionary* robot. The market-maker robot provides liquidity by submitting a buy order (market-maker buyer) for one unit of an asset at 5 cents below or a sell order (market-maker seller) for one unit of an asset at 5 cents above the subject's target price, or both. The reactionary robot is a sniper that takes liquidity; it submits a buy order for one unit at the subject's target price when a sell order arrives at 5 cents below the subject's target price and submits a sell order at the subject's target price when there is a buy order submission at 5 cents above the subject's target price. Asparouhova et al. (2020) report that subjects utilize algorithms frequently, and roughly between 67 percent and 80 percent of trades employed algorithms. They are interested in evaluating whether putting algorithms in the hands of subjects reduces the extent of asset mispricing but find no evidence to that effect. Price bubbles occur as frequently as without algorithms in the market. Further, they show that subjects who use algorithms do not earn higher earnings than manually trading subjects, and the use of algorithms causes a higher frequency of price surges in the first rounds of trading.

In summary, subjects received an opportunity to employ algorithms for trading in the experimental market. Existing experimental evidence shows that human traders utilize the algorithms when they are available to them. Depending on the social cost of algorithmic trading, the regulator can consider the implementation of speed bumps or batch auctions as alternative trading mechanisms to the CDA.[24] Further research shall investigate the impact of other types of algorithmic traders beyond market-making and reactionary algorithms within the framework used in Asparahouva et al. (2020). Generally, the interactions between algorithmic traders and human traders deserve further investigation by considering also other trading strategies and market designs. Designing and teaching a course on algorithmic trading could be a good starting point.

4. CONCLUSION

The experimental research on the interaction between human and algorithmic traders in experimental markets can be organized into studies in which algorithms are in the researcher's hands and those in which the researcher leaves the use of algorithms in the hands of the human subject.

In the first category, the algorithm in the researcher's hand, the reported studies have addressed research questions concerning the performance of algorithms and humans, the impact on market quality and investor psychology. The answers to the questions are not unambiguous. The results suggest that algorithms (particularly the 'fast' ones) frequently

outperform humans in simple market settings. However, in more complex market situations, algorithms (particularly the 'fast' ones) may do worse. Similarly, market quality would usually be enhanced in human-algorithm markets relative to all-human markets, particularly with passive algorithms like arbitrager traders, but may be worsened with manipulators. Investors' behavior and market prices may be attracted closer to fundamentals when the experimenter announces a possible interaction with an algorithmic trader, or no difference may be visible in the data. It seems to depend on the experimental design, and more data are needed to conclude.

In the second category, the algorithm in the subjects' hand, it was reported that real-world phenomena like flash crashes could be reproduced in the laboratory when strategies of inexperienced subjects align. According to the available studies, the efficiency of the CDA market may be unaffected if subjects take algorithms in their hand or if they trade by submitting orders. Again, it would be good to have more data, possibly involving other algorithms than market makers and snipers.

Although the literature is still at an early development stage, the reported results of the studies can have several valuable implications for policy design and discussion:

(1) The results of the experimental literature show that 'fast' algorithms may destabilize the market (Cartlidge et al. 2012; Cartlidge and Cliff 2013). Institutional investors' expenses in latency improvements can moreover be extremely wasteful. Experimental evidence shows that policymaker can regulate wasteful expenses not only by prohibition, but also through an amendment of the continuous double auction market institution. Batch auctions (Aldrich and López Vargas 2020) or speed bumps (Khapko and Zoican 2020) can help to make latency investments obsolete.
(2) The experimental studies suggest that arbitrage algorithms can enhance market quality (Harrison 1992; Angerer et al. 2019; Berger et al. 2020). However, the extraction of wealth from investors by arbitraging activity can be substantial and the policymaker must evaluate if the benefit outweighs the cost for the society or if a redistributive policy for arbitrage gains is desirable. Front running, as discussed in Berger et al. (2020), was already prohibited in order to protect investors.
(3) Algorithms built for the purpose to manipulate the market price and profit from the mispricing (like 'pump-and-dump') must be banned. However, even non-profitable algorithms can systematically impact mispricing (Veiga and Vorsatz 2009) whether algorithmic traders have the intention to manipulate the price or not. Therefore, the code of the algorithmic trader including a verbal description should be registered with the market regulator. The market regulator can then decide, based on the results of experiments or simulations, if the algorithmic trader can be approved or must be rejected.

For future studies in hybrid experimental markets, we think it may be interesting for researchers to explore or continue exploring the following areas:

(1) Many contributions have used the induced value paradigm of Smith (1962). This environment is very well behaved as trading seems to converge always to the competitive equilibrium. To learn about the detrimental or beneficial impact of algorithmic trading in the hybrid market we need more data of market settings that are less well behaved (for instance, Smith et al. 1988). More complex settings with multiple assets or funds of assets would be interesting.

(2) Many studies have focused on ZI agents. The research has not been exhaustive yet. For instance, the discussion on the optimal response speed in the hybrid market is not conclusive. However, other algorithmic traders, for example, as the genetic algorithm, for which we have little or no evidence in hybrid markets, would be very interesting for the study of finance in the laboratory. Experiments in which subjects employ algorithms for trading could ~~be interesting with~~ a richer choice set than just market makers and snipers. For instance, it would also be interesting to see subjects employ algorithmic arbitrage traders to benefit from pricing discrepancies in a multiple market setting.

(3) More studies with manipulator algorithms are required to learn the response of the human-trader market to manipulation, including manipulators built to profit and manipulators to attract attention of investors. The laboratory is an ideal place to study manipulation with algorithms in hybrid markets.

To conclude, experimental hybrid market studies can be informative to researchers, traders and regulators, whereas evidence from real-world observation is sometimes guesswork or sometimes impossible to obtain. We hope that this literature review contributes towards a growing interest in the topic and that an update will be required soon.

NOTES

1. We would like to thank the editor, Sascha Fullbrunn, Steve Gjerstad, one anonymous referee, Stefan Altmann and as well as the participants of the ZIMI2 conference 2021 for helpful comments. We acknowledge fundings from National Research Fund of Luxembourg in collaboration with ESRC (INTER/RCUK/16/11555355) awarded to Tibor Neugebauer, Ministry of Education (MOE) Academic Research Fund Tier 1 (RG163/18) and National Research Foundation (NRF) – CREATE Future Resilient System (FRS) II awarded to Yohanes E. Riyanto and Ministry of Education (MOE) Academic Research Fund Tier 1 (RG69/19) and NTU-WeBank JRC (NWJ-2020-003) awarded to Te Bao.
2. Concretely, we focus on the studies in which algorithms traders interact with human traders in experimental asset markets. Another remotely related strand of literature investigates the role of robot advisors or chatbots in marketing (for example, Luo et al. 2019; Alemanni et al. 2020; Back et al. 2021; Filiz et al. 2021; Hodge et al. 2021). The scope of this research is very different from ours in terms of focus and experimental design and is not included in this review.
3. The design of Smith et al. (1988) is described as follows: nine subjects, initially endowed with cash and assets, can buy or sell assets between each other during $T = 15$ periods in a CDA market. Note that differently from the Smith (1962) market there no induced values and costs, but common values, and each trader can buy and sell assets. No margin purchases and no short sales are permitted. Assets and cash carry over between periods. At the end of the period, a dividend is paid to the asset holders, which takes one of four values in cash units, {0, 8, 28, 60}, and is independently and identically drawn in each period. At the end of the last period, the assets are redeemed at 0 cash units. Hence, the fundamental dividend value is constantly declining across periods.
4. Arifovic studies exchange rate behavior in an overlapping generations model with fiat money. Endowed with units of the consumption good in two periods, the decision makers decide on their consumption when young and their savings in two currencies, which both allow the purchase of the consumption good when old. Intertemporal consumption is valued with a utility function, which translates to a fitness value in the genetic algorithm.
5. Selection for reproduction involves the random mating of two parental decision rules to produce a new generation of decision rules. A decision rule in Arifovic (1996), encoding the savings and portfolio decision of a new agent, is a binary string involving 30 digits, each digit being either 0 or 1. The probability of selection of each parental decision rule depends on its fitness value, which is the ex-post value of the utility function. Reproduction implies an identical copy of the binary strings of each parent to begin with. Crossover is the exchange of parts of the initial strings, for example, the first 20 digits are copied from one parent and the last 10 digits from the other. Mutation is a random change from 0 to 1 or 1 to 0 of a position within a string. The initial two generations in the genetic algorithm are randomly determined decision rules. The following generations are offspring of the young generations. Arifovic (1996) additionally uses the intergeneration enhancement rule *election* as a fourth operator in offspring creation. Kirman (1993) suggests a related model of mutation of heterogeneous opinions, for example, chartists and fundamentalists (see also Brock and Hommes (1998) and Lux and Marchesi (1999); Hommes (2006) surveys the literature on heterogeneous agents).

6. Varying the share of snipers and ZI agents, Brewer and Ratan (2019) find (in an all-algorithm setting) that market efficiency and snipers' profits are strongly impacted when 20 percent or more of traders are snipers.
7. There are three demand and supply schedules, so that for each algorithmic trader there is one human trader with the same induced value or cost structures in balanced markets.
8. As reported in the following section, Das et al. (2001) implement both fast ($t = 1$) and slow ($t = 5$) ZIP agents.
9. Algorithms seem to obtain better returns than human traders also in more complex environments (Sato et al. 2002), but the relative superiority of the algorithms can depend on the market conditions. That is a result of Sato et al. (2002) who report a hybrid human-algorithm market in which students interact with algorithmic traders programmed by teams of researchers at a conference.
10. Symmetric and asymmetric markets refer to the relative shape and intersection of the demand and supply curves. In symmetric markets, the demand curve decreases at any price as much as the supply curve increases, so that the equilibrium price is the axis of symmetry. Asymmetric markets refer to markets where the supply and demand curves are not symmetric, but where typically the demand curve is steeper than the supply curve or vice versa. Equilibrium price discovery usually takes longer in experiments with asymmetric markets than with symmetric markets.
11. Besides the experimental studies investigating the role of algorithmic traders in financial markets, some studies employ algorithmic traders and do not choose the impact of algorithmic traders as the primary research question. For example, Cason and Friedman (1997) used algorithms trading at fundamentals to train subjects. Plat (1997) and Schnitzlein (2002) include liquidity-providing algorithms to create incentives for trading to their experimental subjects. In the learning to forecast experiment by Hommes et al. (2005), the authors also include a fundamental algorithmic trader in the market who constantly predicts and trades based on the asset's fundamental value. The purpose of including these robot traders is to mimic the mean-reverting forces in financial markets.
12. As pointed out above, there is some mixed evidence; Peng et al. (2020) suggest that ZIP traders outperform human traders only under certain conditions.
13. Of course, it may be challenging to embed a fast-speed algorithm in a highly stylized laboratory experimental asset-market setting due to the limited number of subjects, the short and highly compressed duration of a trading period and the relatively low number of trading activities.
14. According to the MM theorem, in the absence of taxes, the market value of a corporation is independent of its dividend policy.
15. There is also some mixed evidence, that is, the study of Grossklags and Schmidt (2006). However, this result might be at least partially explained by the complexity of their experimental design.
16. At the time of writing, further experimental studies are prepared for index arbitrage by Duffy, Friedman, Rabanal and Rud, and for foreign exchange arbitrage by Mitzkewitz and Neugebauer.
17. The authors report data on the hybrid markets in the appendix of the paper. The transaction volume is increased relative to the markets without algorithms impacting mispricing and slowing convergence on fundamentals.
18. The algorithm in Veiga and Vorsatz (2009, 2010) aims at pushing up the price. Thus, when the actual value of the dividend is high, the algorithm could even facilitate price discovery.
19. In Veiga and Vorsatz (2010) the algorithmic features minor changes to the one in Veiga and Vorsatz (2009). The authors thus adapt it to the different market design. The algorithm features increased price increments and decrements, involving identical distributions for buy and sell orders. Most importantly, the algorithm is endowed with 4 shares as human traders, and in the selling mode it can sell up to 14 shares (10 initially bought plus endowment). However, Veiga and Vorsatz (2010) do not clarify whether these changes enhance the algorithm's profitability.
20. Hanson et al. (2006) study price manipulation in a prediction market in which manipulator subjects receive a bonus payment based on price distortions. The authors find that the attempts to distort the price are short lived due to successful counteractions of other market participants. However, it is important to emphasize that in this setting everyone was informed about the manipulators' presence, their objective function and the direction of the manipulation. Therefore, the authors call for further research with the relaxed assumptions before claiming that (prediction) markets cannot be manipulated. Comerton-Forde and Putniņš (2011) investigate the impact of closing price manipulation. One of the main findings is that market manipulations through aggressive buying or selling activities just before the market closing can effectively distort the price. They investigate the possibility of punishment of manipulators by the other market participants but find that others are not always able to identify manipulation. In fact, despite closing price manipulation practices being illegal as they create an illusion of market interest and hinder the price discovery process, it seems complicated to prove and prosecute manipulations in financial markets.
21. Wordle or tag cloud is a visual data representation where the size and the color of a word shows its relative importance.
22. The authors also conducted a treatment in which an algorithm participates trading on fundamental value but report no data of that treatment.
23. In a related study, Kahpko and Zoican (2020) investigate whether a speed-bump policy in a continuous auction environment could have a similar effect. The experimental treatments involve the submission of an order, the first arriving order wins or the winning order is chosen randomly if several orders arrive first. Orders generally

arrive delayed. Subjects can make latency investments to decrease the arrival time. The authors find that subjects do invest in low-latency trading technology in the control treatment without speed bumps. In the experimental treatment, in which speed bumps artificially delay arrival times, investment in low-latency technology is not reduced if the speed bumps are identical to everyone. Only if speed bumps are heterogenous, investments in low latency technology drop by 20 percent relative to the control treatment. This result seems robust whether the time delay involved with the speed bumps is certain or uncertain.

24. Alternative market structures deserve to be tested in the hybrid human-robot experimental market. An interesting paper in this respect is the simulation study of Brewer et al. (2013). The authors simulate limit orders at random arrival times in the CDA market to study market erosion amid a large order that introduces a flash crash. They find that the severity of the erosion depends on the market structure. Their study includes the following alternative structures; minimum resting times, trading halts, and switching to the call auction mechanism amid a flash crash. Their results suggest that the temporal switch to a call auction from the continuous double auction may be most effective in rebuilding the order book liquidity and recovery of the price level amid a flash crash. It would be interesting to confirm these effects in experimental hybrid markets.

REFERENCES

Ahrens, S., Bosch-Rosa, C., & Roulund, R. P. (2020). *Asset price dynamics and endogenous trader over-confidence.* Working Paper. Available at: https://www.macroeconomics.tu-berlin.de/fileadmin/fg124/Ahrens/Ahrens_BoschRosa_Roulund.pdf

Akiyama, E., Hanaki, N., & Ishikawa, R. (2017). It is not just confusion! Strategic uncertainty in an experimental asset market. *The Economic Journal, 127*, 563–580.

Aldrich, E. M., & López Vargas, K. (2020). Experiments in high-frequency trading: Comparing two market institutions. *Experimental Economics, 23*(2), 322–352.

Alemanni, B., Angelovski, A., di Cagno, D. T., Galliera, A., Linciano, N., Marazzi, F., & Soccorso, P. (2020). Do investors rely on robots? Evidence from an experimental study. *CONSOB Fintech Series, 7*, 1–61.

Angerer, M., Neugebauer, T., & Shachat, J. (2019). *Arbitrage bots in experimental asset markets*. Working Paper. Available at: SSRN 4010055

Arifovic, J. (1996). The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political economy*, *104*(3), 510–541.

Asparouhova, E. N., Bossaerts, P., Rotaru, K., Wang, T., Yadav, N., & Yang, W. (2020). *Humans in charge of trading robots: The first experiment*. Working Paper. Available at: SSRN 3569435.

Back, C., Morana, S., & Spann, M. (2021). *Do Robo-advisors make us better investors?* Working Paper. Available at: SSRN 3777387.

Beckhardt, B., Frankl, D. E., Lu, C., & Wang, M. I. (2016). *A survey of high-frequency trading strategies*. Working Paper. Available at: https://web.stanford.edu/class/msande448/2016/final/group5.pdf

Berger, N., DeSantis, M., & Porter, D. (2020). The impact of high-frequency trading in experimental markets. *The Journal of Investing*, *29*(4), 7–18.

Brewer, P. J. (2008). Zero-intelligence robots and the double auction market: A graphical tour. *Handbook of Experimental Economics Results* (vol 1, pp. 31–45). North-Holland, Amsterdam.

Brewer, P., Cvitanic, J., & Plott, C. R. (2013). Market microstructure design and flash crashes: A simulation approach. *Journal of Applied Economics*, *16*(2), 223–250.

Brewer, P., & Ratan, A. (2019). Profitability, efficiency, and inequality in double auction markets with snipers. *Journal of Economic Behavior & Organization*, *164*, 486–499.

Brock, W. A., & Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, 22(8–9), 1235–1274.

Cartlidge, J., & Cliff, D. (2013, February). Evidencing the 'robot phase transition' in human-agent experimental financial markets. *Proceedings of 5th International Conference on Agents and Artificial Intelligence ICAART (1)*, 345–352.

Cartlidge, J., & Cliff, D. (2018). Modelling complex financial markets using real-time human-agent trading experiments. In S-H. Chen, Y-F. Kao, R. Venkatachalam, & Y-R. Du (Eds.), *Complex Systems Modeling and Simulation in Economics and Finance* (pp. 35–69). (Springer Proceedings in Complexity). Springer International Publishing AG. https://doi.org/10.1007/978-3-319-99624-0_3

Cartlidge, J., Szostek, C., De Luca, M., & Cliff, D. (2012, February). Too fast too furious: Faster financial-market trading agents can give less efficient markets. *Proceedings of 4th International Conference on Agents and Artificial Intelligence ICAART* (2), 126–135.

Cason, T., & Friedman, D. (1997). Price formation in single call markets. *Econometrica: Journal of the Econometric Society*, 65(2), 311–345.

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.

Cliff, D., & Bruten, J. (1997). Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report HP-97-91, Hewlett–Packard Research Labs, Bristol, England.

Comerton-Forde, C., & Putniņš, T. J. (2011). Pricing accuracy, liquidity and trader behavior with closing price manipulation. *Experimental Economics*, 14(1), 110–131.

Das, R., Hanson, J. E., Kephart, J. O., & Tesauro, G. (2001, August). Agent-human interactions in the continuous double auction. *Proceedings of 17th International joint conference on artificial intelligence (IJCAI-01)*. Lawrence Erlbaum Associates Ltd, *1*, 1169–1178.

De Luca, M., Szostek, C. S., Cartlidge, J., & Cliff, D. (2011, September). *Studies of interaction between human traders and algorithmic trading systems*. Foresight Report – The Future of Computer Trading in Financial Markets, DR13.

Duffy, J. (2006). Agent-based models and human subject experiments. *Handbook of Computational Economics* (vol. 2, pp. 949–1011). North-Holland, Amsterdam.

Duffy, J., & Ünver, M. U. (2006). Asset price bubbles and crashes with near-zero-intelligence traders. *Economic Theory*, *27*(3), 537–563.

Farjam, M., & Kirchkamp, O. (2018). Bubbles in hybrid markets: How expectations about algorithmic trading affect human trading. *Journal of Economic Behavior & Organization*, 146, 248–269.

Feldman, T., & Friedman, D. (2010). Human and artificial agents in a crash-prone financial market. *Computational Economics*, *36*(3), 201–229.

Filiz, I., Judek, J. R., Lorenz, M., & Spiwoks, M. (2021). Reducing algorithm aversion through experience. *Journal of Behavioral and Experimental Finance*, 31, 100524.

Gjerstad, S. (2007). The competitive market paradox. *Journal of Economic Dynamics and Control*, *31*(5), 1753–1780.

Gjerstad, S., & Dickhaut, J. (1998). Price formation in double auctions. *Games and Economic Behavior*, *22*(1), 1–29.

Gode, D. K., & Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1), 119–137.

Goldstein, M. A., Kumar, P., & Graves, F. C. (2014). Computerized and high-frequency trading. *Financial Review*, *49*(2), 177–202.

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.

Hanaki, N., Akiyama, E., & Ishikawa, R. (2018). Behavioral uncertainty and the dynamics of traders' confidence in their price forecasts. *Journal of Economic Dynamics and Control*, 88, 121–136.

Hanson, R., Oprea, R., & Porter, D. (2006). Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior & Organization*, *60*(4), 449–459.

Harrison, G. W. (1992). Market dynamics, programmed traders and futures markets: Beginning the laboratory search for a smoking gun. *Economic Record*, 68, 46–62.

Hodge, F. D., Mendoza, K. I., & Sinha, R. K. (2021). The effect of humanizing robo-advisors on investor judgments. *Contemporary Accounting Research*, 38(1), 770–792.

Hommes, C. H. (2006). Heterogeneous agent models in economics and finance. *Handbook of Computational Economics* (vol. 2, pp. 1109–1186). North-Holland, Amsterdam.

Hommes, C., Sonnemans, J., Tuinstra, J., & Van de Velden, H. (2005). Coordination of expectations in asset pricing experiments. *The Review of Financial Studies*, 18(3), 955–980.

Khapko, M., & Zoican, M. (2020). Do speed bumps curb low-latency investment? Evidence from a laboratory market. *Journal of Financial Markets*, 55, 100601.

Kirilenko, A. A., & Lo, A. W. (2013). Moore's law versus murphy's law: Algorithmic trading and its discontents. *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.

Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108(1), 137–156.

Leal, J. S., & Hanaki, N. (2018). *Algorithmic trading, what if it is just an illusion? Evidence from experimental financial markets*. Technical report, Groupe de Recherche en Droit, Economie, Gestion (GREDEG CNRS), Université Côte d'Azur, France.

Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. humans: The impact of artificial intelligence chatbot disclosure on customer purchases. *Marketing Science*, *38*(6), 937–947.

Lux, T., & Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719), 498–500.

March, C. (2019). *The behavioral economics of artificial intelligence: Lessons from experiments with computer players*. Working Paper. Available at: SSRN 3485475.

Miller, R. S., & Shorter, G. (2016). *High frequency trading: Overview of recent developments*. Congressional Research Service Report No. 4, Washington, DC.

Neugebauer, T., Shachat, J., & Szymczak, W. (2020). *A test of the Modigliani-Miller theorem, dividend policy and algorithmic arbitrage in experimental asset markets*. ESI Working Paper 20-14. Available at: https://digitalcom-mons.chapman.edu/cgi/viewcontent.cgi?article=1309&context=esi_working_papers

Peng, Y., Shachat, J., Wei, L., & Zhang, S. S. (2020). *Speed traps: Algorithmic trader performance under alternative market structures*. ESI Working Paper 20-39. Available at: https://digitalcommons.chapman.edu/esi_working_papers/334

Plat, C. G. (1997). *A double auction market with signals of varying precision*. Working Paper. Available at: https://econwpa.ub.uni-muenchen.de/econ-wp/exp/papers/0508/0508004.pdf

Plott, C. R., & Sunder, S. (1982). Efficiency of experimental security markets with insider information: An application to rational-expectations models. *Journal of Political Economy* 90, 663–698.

Plott, C. R., & Sunder, S. (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica: Journal of the Econometric Society*, 56(5), 1085–1118.

Putniņš, T. J. (2012). Market manipulation: A survey. *Journal of Economic Surveys*, *26*(5), 952–967.

Rietz, T. A. (2005). *Behavioral mis-pricing and arbitrage in experimental asset markets*. University of Iowa Working paper.

Rust, J., Miller, J. H., & Palmer, R. (1994). Characterizing effective trading strategies: Insights from a computerized double auction tournament. *Journal of Economic Dynamics and Control*, 18(1), 61–96.

Sato, H., Matsui, H., Ono, I., Kita, H., Terano, T., Deguchi, H., & Shiozawa, Y. (2002). Case report on U-Mart experimental system: Competition of software agent and gaming simulation with human agents. In A. Namatame, T. Terano, & K. Kurumatani (Eds.), *Agent-Based Approaches in Economic and Social Complex Systems, Frontiers in Artificial Intelligence and Applications* (pp. 167–178). Tokyo, Japan: IOS Press & Ohmsha Ltd.

Schnitzlein, C. R. (2002). Price formation and market quality when the number and presence of insiders is unknown. *The Review of Financial Studies*, 15(4), 1077-1109.

Smith, V. L. (1962). An experimental study of competitive market behavior. *Journal of Political Economy*, *70*(2), 111–137.

Smith, V. L., Suchanek, G. L., & Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society*, 56(5), 1119–1151.

Sunder, S. (2003). *Markets as artifacts: Aggregate efficiency from zero-intelligence traders*. Yale ICF Working Paper 02-16. Available at: SSRN 309750.

Tai, C. C., Chen, S. H., & Yang, L. X. (2018). Cognitive ability and earnings performance: Evidence from double auction market experiments. *Journal of Economic Dynamics and Control*, 91, 409–440.

Treleaven, P., Galas, M., & Lalchand, V. (2013). Algorithmic trading review. *Communications of the ACM*, 56(11), 76–85.

Veiga, H., & Vorsatz, M. (2009). Price manipulation in an experimental asset market. *European Economic Review*, 53(3), 327–342.

Veiga, H., & Vorsatz, M. (2010). Information aggregation in experimental asset markets in the presence of a manipulator. *Experimental Economics*, *13*(4), 379–398.