For Better or For Worse: Algorithmic Choice in Experimental Markets

Elena Asparouhova University of Utah Debrah Meloso TBS Business School Wenhao Yang

UC Berkeley

Christine A. Parlour

Xiaoqin Cai

UIUC

Chinese University of Hong Kong

Jan Nielsen

Ad Hoc Markets LLC

December 22, 2022

Abstract

Participants in an experimental market choose to enter private value trades manually and/or algorithmically. Each algorithm or trading robot makes or takes liquidity based on the trader's current marginal valuation modulo a spread chosen by the trader. We evaluate experimental outcomes against both competitive equilibrium and equilibrium of the strategic game if all participants choose robots. Data from six laboratory experimental sessions support many of the theoretical findings. Most traders deploy an algorithm whenever available (the average trader deploys a robot in 82% of the rounds, and only 4% of subjects never deploy a robot), and learn to use them with experience. Compared to rounds with only manual trading, algorithms improve allocative efficiency. Realized gains from trade increase from 55% to 84% . While the allocative efficiency increases across the board, those who benefit most are the traders who perform poorly in manual trading. Our results highlight how algorithm choice can affect relative outcomes and market observables.

We acknowledge financial support from the Fondation Banque de France. We thank Noah Alex, Austin Brooksby, Dylan Finlayson, Matthew Froberg, and Gus Stevenson, who provided invaluable research assistance. We also thank seminar attendants at the University of Alicante, the University of Geneva, the University of Montpellier, the University of Murcia, University of Stavanger, the University of Toronto, and TBS Education. We are also grateful for comments received from conference audiences at the Durham Algorithm Trading Workshop, the Fintech: Innovation, Inclusion, and Risks conference, the First Zero/Minimal Intelligence conference, the French Association for Experimental Finance conference, the Galapagos Experimental Economics conference, the Society for Experimental Finance's annual conference, the Southern European Experimental Team's annual meeting, and the Third Behavioral Finance and Capital Markets conference.

I. INTRODUCTION

Trades in today's financial markets are overwhelmingly consummated by trading algorithms. The *Economist* reports that as of the end of 2019 "90% of equity-futures trades and 80% of cash-equity trades are executed by algorithms without any human input." In spite of the widespread adoption of these trading tools, open questions remain: do algorithms improve aggregate outcomes in financial markets or simply redistribute gains from trade to their users? How does the combination of manual and algorithmic trading affect allocations and prices?

To investigate these questions, we run a series of experiments in which participants can design and implement trading robots. We use a trading environment similar to a simple financial market in which agents compete to consummate gains from trade. In our experiments, agents have linear induced demand in an asset and trade in a continuous double auction also known as a limit order book. Gains from trade arise as all the agents are trying to diversify their portfolios – there is no private information.¹ To assist them in their trades, agents can engage robots. There are two types of robots, those that supply or make liquidity, i.e., post prices, and those that demand or take liquidity, i.e, lift posted prices. Once deployed, these robots calculate their agent's instantaneous valuation for the asset and submit orders that are centered around the valuation plus or minus a spread selected by the agent. By design, only one robot can be deployed, and once deployed a robot cannot be terminated.² In sum, the traders in our setup can choose the role (maker or taker) and the spread of each deployed algorithm. Importantly, the traders can also trade manually even if they have chosen to deploy a robot.

There are three notable features of our experiments. First, all trades are private value trades and thus transaction prices should reflect these private motives for trade. Second, by design, our robots can only improve a user's welfare. In particular, agents using robots can never "make mistakes" that lead to losses, nor do the agents have to engage in complicated updating. Finally, starting with Smith (1962), the literature has found that the continuous double auction (CDA)— in particular, the shared information feature of the open book— leads to equilibrium in human-only experimental markets.

¹For a theoretical treatment of Algorithmic Market-Makers using Q-learning algorithm in an environment of adverse selection, see Colliard et al. (2022)

²Multiple algorithm deployment and pausing, stopping, and redeploying new algorithms is technically possible. The choice of simplicity is for tractability, both theoretical and statistical.

We have two natural benchmarks to evaluate outcomes: social efficiency and strategic equilibrium with only deployed robots. These benchmarks allow us to evaluate the effect of additional manual trading. In the socially efficient outcome, all gains from trade are exhausted. Given the symmetry we impose on payoffs, this means that all agents should end up with the same portfolio. We can measure how close agents are to this benchmark individually and collectively. Of course, as our agents are strategic, each is trying to appropriate more of the gains from trade. This is done through robot and spread choice, which deters some of the gains from trade. Conditional on robot choices, we can calculate these strategic outcomes.

In a framework with an equal number of buyers and sellers and no manual trading we characterize the equilibrium choice of robot and spread, efficiency losses, and price deviations from a model of rational manual trading. There is an equilibrium in which (*i*) one side of the market (buyers or sellers) chooses the role of takers and sets the spread to the minimal available spread; (*ii*) the other side of the market chooses maker robots and sets the spread as high as possible; (*iii*) the gains from trade depend on the spreads, and the distribution of gains is highly dependent on robot parameter choices.

We report the results from six experimental sessions where we find that when not provided access to robots, humans realize 55% of the theoretical gains from trade available in the market. When given access to robots, 96% of traders use them at least once, with the fraction of participants deploying robots varying between 2/3 to 1 in any given round (with an average of 82%). In spite of robot use, manual trading remains popular: The fraction of participants who trade manually alongside a robot ranges from 0% to 87.5%, and manual volume as a percentage of trading volume ranges from 17.4% to 60.2%.

Based on survey responses, we document that participants who understood the robots were more likely to deploy them. Conditional on deploying a robot, the participants do so mostly within 20 seconds of the start of round. As predicted by the theory, the chosen robots are equally split between makers and takers but the participants do not appear to be able to solve the coordination problem that equilibrium choice of robots presents. The participant's choice of robot type is sticky, for both types of robots, across rounds. As predicted by theory, maker algorithms are set with higher spreads than taker algorithms. Most importantly, the realized gains from trade increase from 55% in manual to 84% in rounds where robots were available. Those who benefit most from the robots are the traders that perform the worst in manual trading. We identify traders with different trading skill by ranking them based on manual rounds. In our experiment, the worst tercile of manual traders increases their relative performance from 5% to 74% of maximal gains when using robots and as a result, those who do best in manual trading decrease their relative performance from 108% to 98%. This follows from the fact that the robots provide utility improvements, they are easy to understand, and relieve the user from the cognitive burden associated with monitoring the market and trading. (Whether humans understand, trust, and deploy robots is part of a large literature on human-robot interaction (HRI).)

More surprising are the cross-sectional implications of robot use. Good manual traders deploy robots more effectively and so continue their superior performance in the rounds where robots are available. However, there is an offsetting effect: The overall gains from trade, or the "size of the pie," increases and so even though the skilled manual traders improve performance with robot usage, their relative performance decreases. Part of the reason these good traders are profitable is that they are more likely to sell at higher prices and buy at lower ones. Thus, transaction prices do not necessarily reflect traders' ordered marginal valuations but also their strategic skill.

It is natural to examine the overall and cross-sectional effectiveness of algorithms in an experimental setting. This topic is difficult to address theoretically with Bayesian expected utility maximizers, which are effectively modelled as perfect algorithms. Gode and Sunder (1993), an early study using algorithms, showed that (naive) centralized price competition and continuous trading, both features of the CDA, converge to competitive equilibrium. In the context of multiple assets, Asparouhova et al. (2020a) identifies boundedly-rational individual bidding rules that have been theoretically shown and experimentally verified to lead to equilibrium in CDA markets. We use both Gode and Sunder (1993) and Asparouhova et al. (2020a) as the motivation for our study. The algorithms that we create are in the spirit of Gode and Sunder (1993), called zero-intelligence (ZI) traders.³ They are augmented in that traders choose a spread. The spread parameter creates individual bid and ask reservation prices, where the bid (ask) price is equal to the marginal valuation of the asset minus (plus) half the spread. After simplifying the bidding rule in Asparouhova et al. (2020a), the resulting algorithms—our *Mean-Variance Optimizer*, or MVO robots—constitute a natural extension of the ZI robots of Gode and Sunder (1993) to the CAPM-driven design.

 $^{^{3}}$ The original ZI traders were deployed in a standard demand-supply setup, as in Smith (1962), where each trader is only a buyer or a seller. Each buyer (seller) robot sends a limit offer to the market to buy (sell) at the marginal price minus/plus a random amount.

The study of algorithm participation in financial markets has focused on the effect of algorithms (often *high frequency traders* (HFT)) on welfare, fairness (do some parties benefit disproportionately at the expense of others?), and price stability.⁴ In contrast to this work, our algorithms work on behalf of the humans who deploys them.

A subset of the literature considers hybrid experiments in which humans participate alongside (experimenter-controlled) algorithms to see which "win," how human behavior is affected by the presence of algorithms or their specific type, and how efficiency changes as a consequence. By 2006, engineers at IBM had found a robot that led them to proclaim that "The Trader Is Dead, Long Live The Trader" (Das et al., 2001; Dence et al., 2006). Little after, the human trader was resuscitated by studies pointing out that the success of rigid algorithm rules and speed is highly dependent on market parameters (Gjerstad, 2007; Shachat et al., 2020) and that earnings to traders deploying trading algorithms have not been abnormal on average (Kissell and Malamut, 2005). In the context of financial markets, Grossklags and Schmidt (2006) and Farjam and Kirchkamp (2015) show that when human traders know or think that robots are present, allocations are more efficient than when humans are alone. However, when humans are not aware of the presence of robots, the hybrid market is less efficient (Farjam and Kirchkamp, 2015).

In real markets, the complexity of the markets populated by algorithms "has created a new class of finance professionals known as "power users," who are highly trained experts with domain-specific technical knowledge of algorithmic trading." (Kirilenko and Lo, 2013, p. 67) Our work pertains to this category, where humans *control* robots. To the best of our knowledge, Asparouhova et al. (2020b) and Aldrich and López Vargas (2020) are the only other studies examining this form of human-robot interaction. Bao et al. (2021) provide a recent and comprehensive review of the literature.

Asparouhova et al. (2020b) studies the results from the engagement of algorithms (trading robots) in a financial markets experiment where participants have access to a set of robots imple-

⁴Das et al. (2001) find that human traders underperform algorithms by about 20% in trading surplus. Gjerstad (2007) studies how different market structures and paces of submitting bids and offers influence the trading performance of humans and algorithms in a CDA market with induced values (Smith (1962)). In this study, impatient algorithmic traders' profitability seems to be lower compared to the patient ones. Shachat et al. (2020) study balanced, unbalanced and noncompetitive markets with "fast" and "slow" ZIP robots. The authors confirm the result of Das et al. (2001) that algorithms outperform human traders in the "fast" treatment in a balanced market. In all other conditions, contrary to the result of Das et al. (2001), the authors report that human traders outperform "fast" algorithmic traders. For the various results on the effect of robots on convergence to equilibrium, among others, see Shachat et al. (2020), Furse et al. (2012), and De Luca et al. (2011)

menting user-specified bidding strategies—participants specify the type of orders (market or limit), the side (buy or sell), and the price their robots should submit to markets. As the robots in this work trade in the seminal Smith et al. (1988) framework, individual and collective benefits from trade are not well defined (barring knowledge of participant risk aversion). This stands in opposition to our setup, where individual and collective welfare are clearly defined. This enables us to provide our participants with goal oriented robots while still allowing for the emergence of inefficiencies. As expected, our results differ from those in Asparouhova et al. (2020b), as we find a positive effect of robot usage on efficiency, improving the performance of all participants, not only those who trade manually.

Aldrich and López Vargas (2020) allows humans to set a role and a costly speed for their robots in the context of two differently-organized markets—a CDA and a *frequent batch auction* (FBA) market. The roles of robots in their setting are that of market maker or *sniper*, a form of opportunistic liquidity-taker who preys on speed differences and the rules of the CDA. They find that humans choose different robots in the CDA than in the FBA, optimally reacting to the incentives provided both by the market type and the consequent characteristics of other robotic traders. While we differ from this work in our research question—theirs is a question of market design—we share with this study the prisoner's dilemma-like robot choice game in the CDA. Instead of costly speed, we use robot spread as a strategy that increases individual performance and decreases efficiency. Unlike in Aldrich and López Vargas (2020), we find significantly less inefficiency than expected, as participant choices cut short of the predicted "spread" arms race.⁵

The remainder of the article is organized as follows. In Section II. we present the basic structure of the market, preferences, and algorithms that are used in our experiment. We analyze the experimental results in Section III. A brief conclusion follows in Section IV. Details of the theoretical framework, survey questions, and additional results are presented in the Appendix.

II. FRAMEWORK

In our environment mean variance robots seek gains from trade within a fixed time period.

⁵A recent and important literature on algorithmic advice has focused on human trust in algorithms. Our experiment is not designed to address the question of robotic advisor trust and adoption. For recent experimental contributions to this literature, see Ben-David and Sade (2021) and Greiner et al. (2022).

Utility gains from trade: All gains from trade come from portfolio diversification. There are two risky assets, *A* and *B*, each characterized by an expected return and standard deviation, denoted (μ_i, σ_i) , i = A, B with covariance σ_{AB} . For notational compactness, we sometimes denote μ as the vector of expected returns, while Σ is the variance covariance matrix. Traders have mean variance preferences with common risk aversion parameter γ . All agents are endowed with the same amount of *B*, (e_B) , and cash, (c), but differ in their initial endowments of asset A. Trade is in asset A.

Let $\mathbf{x}(\mathbf{t})^j = (x_A^j(t), x_B^j(t))^T$ and $c^j(t)$ denote agent j's position in the risky assets and cash at time t. Then their expected utility at time t is

$$U_t^j(\mathbf{x}(t)^j, c^j(t)) = c^j(t) + \mu \mathbf{x}^j(t) - \frac{\gamma}{2} (\mathbf{x}^j(t))^T \Sigma \mathbf{x}^j(t)$$
(1)

Agents maximize the difference between the expected utility of their final (after trading) endowment and their initial endowment or $U_T^j - U_0^j$.

Throughout our analysis we consider two benchmarks. First, social efficiency and second, the equilibrium of a strategic trading game.

Social Efficiency: The agents are *ex ante* identical, and for 2*N* agents, their aggregate expected utility is

$$\Omega = \sum_{j=1}^{2N} U_T^j - U_0^j$$
 (2)

The maximal or socially efficient outcome is generated by a set of feasible allocations of asset *A*, $x^{1}(T), \ldots x^{2N}(T)$ that maximizes Ω subject to the resource constraint $\sum_{j=1}^{2N} e_{A}^{j}$.

Given that the only source of heterogeneity is the initial endowment of asset *A*, it is immediate that in the socially efficient outcome, each participant receives an identical allocation of $x_A^* = \sum_{j=1}^{2N} \frac{e_A^j}{2N}$. This allocation implies a collective valuation of asset *A* or, in other words, a price. To see this, suppose that agent *j* at time *t* has acquired Δ units of asset A. Given their initial endowment of e_A^j , the current valuation for an additional unit of asset *A* is

$$\rho^{j}(e_{A}^{j}+\Delta) = \mu - \gamma \Big[\sigma_{A}^{2}\Big(e_{A}^{j}+\Delta\Big) + \sigma_{AB}e_{B}\Big].$$
(3)

This is the amount that any agent would be willing to pay for an additional unit of *A*. We formalize this discussion in the following lemma.

Lemma 1 In a two asset economy with 2N mean variance agents with a common risk aversion parameter and endowments of cash and asset B,

- i. in the socially efficient allocation each agent holds $x_A^* = \sum_{j=1}^{2N} \frac{e_A^j}{2N}$.
- ii. at the socially efficient allocation the price is $\rho(x_A^*)$.

The socially efficient allocation and attendant social welfare represent the maximal welfare that can be obtained in this environment and are an important benchmark which is not necessarily reached through trade.

A. TRADING GAME

The market for stock A is open for a finite length of time, *T*. There is no short selling or borrowing. Trade is organized as an open electronic limit order book. A limit order to buy (bid) or sell (ask), specifies a quantity and a price that the trader is willing to buy/sell. A non-marketable limit order is one that does not match any of orders in the book and, thus, sits in the book until matched or cancelled. By contrast, marketable limit orders specify a price that matches a preexisting order and, thus, immediately executes. If the quantity exceeds that of the preexisting order, a large marketable order is partially executed, and the remainder accumulates in the book.

Orders are submitted either manually or by deployed robots. A manual trader chooses whether to submit a marketable or non-marketable limit order, the quantity, and the price. Similarly, the decision of agent *j* deploying a robot is a pair (r, δ), where $r \in \{M, T\}$ is the robot role and $\delta \ge 0$ is the spread.

Definition 1 A robot

- i. is a maker or makes liquidity if it posts liquidity that rests on the limit order book,
- ii. is a taker or takes liquidity if it executes against previously posted liquidity.

The price at which a robot makes or takes is determined by the spread chosen by agent j, δ^{j} . The spread is fixed throughout each realization of the market (i.e. *a round* experimentally). The spread is centered around agent *j*'s current valuation for the asset (as specified in Equation 3). Specifically, the algorithm calculates an agent's expected utility and marginal valuation of the asset.

A maker robot posts personalized, single quantity, offers and bids:

$$a^{j}(x) = \rho^{j}(x) + \frac{\delta^{j}}{2}$$
(4)

$$b^{j}(x) = \rho^{j}(x) - \frac{\delta^{j}}{2}.$$
(5)

Such bids and offers are adjusted (i.e., cancelled and resubmitted) after any change to an agent's portfolio. By contrast, a taker robot only submits a single quantity order if it can execute immediately. Specifically, suppose there is a posted price *p*, then for a taker robot

if
$$p \begin{cases} \geq \rho^{j}(x) + \frac{\delta^{j}}{2} & \text{sells} \\ \leq \rho^{j}(x) - \frac{\delta^{j}}{2} & \text{buys.} \end{cases}$$
 (6)

If the taker fails to consummate the trade it submits a cancel order immediately. The robots submit single quantity orders every two seconds when conditions are satisfied.

As the only motive for trade in this economy is private hedging, any trade consummated by a robot increases the agent's expected utility.

B. MARKET OUTCOMES WITH ROBOTS

To build intuition for the role that robots play, we present market outcomes for two traders and discuss the 2*N* person version in the appendix. First notice, if there are two traders with different endowments of asset *A*, then if $e_A^i > e_A^j$, agent *i* is the natural seller of *A* in the economy. Going forward, we refer to this trader as the supplier of asset *A*, denoted S, while the other agent is a demander, denoted D.

As outlined above, the decision of each agent is a pair (r, δ) , where $r \in \{M, T\}$ is the robot type and $\delta \ge 0$ is the spread. First, suppose that spreads for both are constrained to be zero and thus robots post or are willing to accept posted prices at their valuations. Given the endowments, the traders' valuations of the asset are ranked so that $\rho^S < \rho^D$. Thus, the prices at which trade occurs depends on the robot choice. Specifically, if the supplier has chosen to be a maker, they will post prices at their valuation, which slowly decreases after successive trade rounds. If the demander has chosen to be a taker, trade will occur at the supplier's valuation. Conversely, if the buyer is the maker and the seller is a taker, trade will occur at the buyer's valuation. If both choose to be takers no trade will occur, and if both choose to be makers there is an equal chance of trade at either trader's valuation.⁶ Clearly, given the expected gains to trade, agents have an incentive to choose different roles.

Lemma 2 Suppose there are two traders with different endowments facing a fixed spread of $\delta \ge 0$, then in all pure strategy equilibria, one trader chooses to be a maker while the other chooses to be a taker.

Strategic play ensures that traders will choose trade enhancing roles. However, even if one trader supplies liquidity and one demands liquidity (i.e, one is maker and one is a taker) not all gains from trade may be consummated. This is because spread choice affects how the gains from trade are split. *Ceteris paribus*, the larger the spread choice that an agent makes, the more of the gains from trade that they appropriate. However, larger spread choices may deter trade.

Figure 1 illustrates valuations as a function of quantity traded and how the valuations are distorted because of the spreads. Maximum individually rational trade occurs under the socially efficient outcome which exhausts all gains from trade. In this outcome, both traders hold the same final portfolio, and so the total quantity traded is $\frac{1}{2}(e_A^S + e_A^D)$. By contrast, if both agents submit a spread, then there is less trade, and the last traded unit is $x^*(\delta^S, \delta^D) = \frac{e^S - e^D}{2} - \frac{\delta^S + \delta^D}{4\gamma \sigma_A^2}$.

We measure social welfare loss relative to the socially efficient outcome.

Lemma 3 Efficiency loss Γ in the two-player MVO agent-choice game for exogenous spreads $\{\delta^S, \delta^D\}$ is

$$\Gamma = \frac{1}{16} \frac{\left(\delta^S + \delta^D\right)^2}{\gamma \sigma_A^2}.$$

⁶This is an assumption of our theoretical framework, but we also verify it with simulations. With only two traders, the frequencies of trade at either trader's valuation, are very close to 1/2. With 2N > 2, there is a higher probability that the side of the market setting the first price, also set it later in the round. In the appendix, where we treat this case, we state where the assumption is necessary.

Figure 1: Two-person Spreads and Efficiency Loss. The thick lines represent marginal values as a function of units bought by the buyer (pink) and sold by the seller (blue). The distance between the pink and blue line as they intersect the vertical axis, is the gains from trade. The thin lines are cum spread demand and supply. The intersection of thin lines determines quantity and last trading price, and the gray shaded triangle represents the dead weight loss with respect to the zero-spread benchmark (socially efficient).



Proposition 1 (Two traders, role and spread choice.) When traders choose both their role and their spread, there are two pure strategy equilibria with the following characteristics

- 1. The taker sets a spread equal to zero.
- 2. The maker sets a spread equal to $\Delta = \gamma \sigma_A^2 (e^S e^D)$, the difference between the buyer and the seller's initial marginal values for Stock A.
- 3. The traded quantity is half the zero-spread (efficient) quantity.
- 4. The last trading price is biased away from the zero-spread price in the maker's advantage.

- 5. The efficiency loss in equilibrium equals $\Gamma^* = \frac{1}{16} \gamma \sigma_A^2 (e_A^S e_A^D)^2$.
- 6. The maker obtains twice as large a payoff as the taker.

Proposition 1 tells us that the game that arises when traders can set both the role and the spread of their MVO agents is akin to a prisoner's dilemma. As in the prisoner's dilemma, driven by individual rationality, the dead-weight loss in equilibrium is large. Agents' competition to obtain a larger share of the trade surplus in this game leads to a reduction of the available surplus.

C. EXPERIMENTAL SESSIONS

Each session lasted approximately three hours; half of which was devoted to instruction and practice, and half to for-pay rounds and the exit survey. The timing of a session is displayed in Figure 2 below.

Figure 2: Timeline of an Experimental Session



Subjects participated online with cameras on for the entire duration. Instructions, trading aids, and the trade interface were available over the web, and were shared through interactive video.

Trading Rounds. Ten trading rounds are relevant for participant payoff. In two of these rounds (R1 and R2), participants only trade manually, without the help of robots. We call these *manual rounds*. In the eight remaining rounds (R3 to R10), participants can deploy algorithms to trade on their behalf, while still being able to trade manually. We call these *algo-possible* rounds. Initial asset and cash holdings as well as the payoff obtained in each round are independent of choices made in other rounds.

Endowments in Trading Rounds. In each round, all participants are endowed with equal units of Stock B and cash. Participants differ in their initial holdings of Stock A, thus ensuring that no participant is fully diversified and mutual gains from trade exist. Two endowment types differing only in the number of units of Stock A, coexist in each round. We call these endowments *type I* and

type II and denote them e_I and e_{II} , respectively. Approximately half the participants are allocated each endowment type.⁷ Cash endowments always equal 150 ECU (Experimental Currency Units), but risky asset endowments e_I and e_{II} vary across rounds. Table 1 presents the five endowment setups used in one session.

Setup Number	1	2	3	4	5
Rounds	All practice rounds; R1, R2 (manual)	R3, R4	R5, R6	R7, R8	R9, R10
Stock B (both types)	160	80	122	122	52
Stock A type I Stock A type II	0 160	225 5	240 4	20 224	20 224

Table 1: Endowment setups across experimental rounds

Between-Rounds Feedback. At the end of each round, participants are told their payoff for the round (after conversion to USD). Between algo-possible rounds participants also receive feedback on the algorithms deployed the previous round. Specifically, participants are told the number of maker and taker algorithms deployed by "odd-numbered" (type I) and "even-numbered" (type II) participants.

Instructions, Practice, and Aids. Instructions are provided as a slide show (https://shorturl. at/ijB69) containing several videos to enhance replicability. The experiment has two practice rounds: one for manual and one for algo-possible rounds. Each practice round consists of a first, guided part (explained in instructions) and a second, free trading part. Practice rounds do not count for the final payoff. At all times, participants can use two round-specific tools called *Valuation Tool* and *Performance Tool* (Figure 3) to aid with trade choices. The Valuation Tool displays the marginal value of Stock A (given the mean-variance utility used for payoff) as a function of holdings. The Performance Tool displays (gross) payoff as a function of units of Stock A, under the assumption that cash holdings remain constant at 150 ECU. The tools are interactive and display the coordinate values as the cursor moves over the plotted curve.

⁷In recruiting participants for the experiment, we aimed to have equal numbers of type I and II participants, for which an even number of volunteers is needed. When an odd number of participants shows up to the session or remains after instructions and practice, types differ in their number. In the sessions reported here, they differ by at most 1.





(a) Performance Tool: By positioning the cursor on the (b) Valuation Tool: By positioning the cursor on the graph at the holdings of Stock A, participants are provided the gross payoff (GP) of their position. To compute the (net) payoff, the change in cash needs to be added to the GP.



graph at the current holdings of Stock A, participants are provided with the marginal valuation of the next unit of the stock.

Final Participant Payoff. Participant payoffs for the session consist of three parts: the show-up fee of USD 5, a lump-sum payoff of USD 15 for participation in the practice rounds, and a variable part that is normalized to an average across participants of USD 30. This means that average participant payoff is USD 50 for three hours of participation. The variable part of payoff equals participant payoff for two randomly-chosen rounds. The payoff in a given round equals the difference between the mean-variance utility of the final and the initial holdings of risky assets and cash (as in Equation 1). Risk aversion is set to $\gamma = 0.006$. The payoff in ECU is next converted to USD with an exchange rate that ensures the average per-round payoff across participants is USD 15. Within the variable payoff, while the mean was USD 30, payoffs were bell-shaped distributed around the mean, with the lowest equal to USD 0, and the highest equal to USD 120.

Exit Survey. At the end of an experimental session, participants answer a questionnaire (see Appendix), choose a payment method, and receive payment.⁸

⁸Payment methods were (1) electronic visa card, emailed immediately to the participant, or (2) cash, paid in person in the laboratory after the session.

D. EXPERIMENTAL ROUND

Securities. Participants are endowed with two risky assets and cash. One asset, *Stock A*, can be traded, while the other, *Stock B*, cannot. Expected payoffs (μ) and variance-covariance matrix (Σ) for the risky assets Stock A and Stock B are:

$$\boldsymbol{\mu} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \ \boldsymbol{\Sigma} = \begin{bmatrix} 2/3 & -1/3 \\ -1/3 & 2/3 \end{bmatrix}.$$

Interface for Trade and Algorithm Deployment. The trading mechanism is part of online software called Flex-E-Markets, which provides a software as a service (SaaS) platform to organize and manage multiple, simultaneous online marketplaces (see https://adhocmarkets.com). The left panel of Figure 4 shows the interface for submitting manual orders alongside the open book displaying all currently available limit orders and past trades of Stock A. Orders and trades are displayed anonymously and without reference to their manual or algorithmic origin. The right panel of Figure 4 closes up on the algorithm deployment tool, which is integrated in the lower left part of the trading interface whenever algorithm deployment is allowed.





Explanations: In Flex-E-Markets, participants observe the open book of buy (blue) and sell (red) orders, submit/cancel orders, and trade with each other. When available, algorithms can be deployed via a menu placed underneath the "Place Buy Order" button, enlarged for clarity on the right. Before starting an algorithm, it is possible to set its role (maker/taker) with a drop-down menu, and its spread with a slider.

Information on submitted orders and trades, algorithm connection to the trading platform (algorithm deployment time), and algorithm parameters, is collected directly from the trading platform, Flex-E-Markets. Each submitted order is identified (to us, the experimenters) as coming from a human (identified by their trader ID) or a specific algorithm type (identified by the parameters of the algorithm deployed and also by the deploying trader's ID).

E. Algorithms

In the algo-possible rounds, participants can deploy Mean-Variance Optimizing (MVO) algorithmic robots. As outlined in Subsection A., as the MVO robot trades, the robot recomputes the marginal (mean-variance) utility of Stock A given current holdings of both stocks and uses this as a reservation price for trades. The parameters we used were $\mu_A = 1$ as the expected dividend of Stock A and of Stock B, $\sigma_A^2 = 2/3$ as the variance of Stock A dividends, and $\sigma_{AB} = -1/3$ as the covariance of Stock A and Stock B dividends. With this, the valuation is

$$\rho_t^J = \left(1 + \frac{0.006}{3}e_B\right) - \left(\frac{0.012}{3}\right)x_A^j(t),$$

a linear function of the holdings at *t* of Stock A.

Participants in an algo-possible round choose whether to deploy an algorithm or not, and if they do deploy one, they can choose the **spread**, the **role**, and the time of deployment. Following Equations 4 and 6, spreads were chosen from $\delta \in [0.02, 0.5]$.

It is important to stress that the algorithms in our experiment are not independent economic agents, since they do not have endowments of their own. Instead, they participate in markets only if a human participant deploys them, and all their actions affect the payoff of this human participant. Once deployed, an algorithm cannot be stopped or changed, but manual trades can be submitted alongside it.

III. EXPERIMENTAL RESULTS

We ran six online experimental sessions. Participants were recruited from two pools: University of Utah students registered on the Utah Laboratory of Experimental Economics and Finance (ULEEF)

sign-up system (SONA), and members of the Discord group Beermoney India, subsequently registered on the ULEEF system.⁹ Number of participants of each endowment type, market portfolios and equilibrium prices for each endowment setup in each experimental session, are all reported in Table 2.

			Ses	sion	
		1	2	3	4, 5, 6
Participants with endowm	ent type I	4	6	5	7
Participants with endowment type II		5	6	4	7
Total Participants		9	12	9	14
	Cotup 1	(88.9, 160)	(80,160)	(71.1,160)	(80,160)
	Setup 1	0.96	1	1.04	1
	Cotup 2	(102.8,80)	(115,80)	(127.2,80)	(115,80)
	Setup 2	0.75	0.7	0.65	0.7
Mkt. portfolio (\bar{e}_A, \bar{e}_B);	Sotup 2	(108.9, 122)	(122, 122)	(135.1,122)	(122, 122)
Equilibrium price p*	Setup 5	0.81	0.76	0.7	0.76
	Sotup 4	(133.3, 122)	(122, 122)	(110.7, 122)	(122, 122)
	Setup 4	0.71	0.76	0.8	0.76
	Sotup E	(133.3, 52)	(122,52)	(110.7, 52)	(122, 52)
	Secup 5	0.57	0.62	0.66	0.62

Table 2: Participants, Market Portfolio, and Equilibrium Price per Session.

To implement our benchmarks, we use the parameters indicated in Table 2. Specifically, to evaluate aggregate outcomes and the performance of an individual trader, we compare volume, prices, efficiency, and individual realized payoffs to predictions under the socially efficient outcome.

The experiment was not preregistered. Instead, all sessions were run via a script that automatically started and stopped rounds, effectively preventing deviations from a preset design.¹⁰

⁹The subject pool used to be restricted to enrolled University of Utah students but was opened up during the Covid-19 pandemic. We only discovered on the spot—thanks to the exit survey—that participants of our fourth session were not University of Utah students. As a reaction, we asked for additional information from these participants and created a formalized connection to the Discord group that allowed us to conduct another session with participants from Discord, resided in India at the time of the experiment.

¹⁰The script can be viewed at the link https://shorturl.at/AHOXY



Figure 5: Round final prices vs. competitive equilibrium prices.

Explanations: Average price of the last 20 trades of each round vis-à-vis socially-optimal price. Black lines show Nash equilibrium prices (algorithms only), where the natural buyers (lower line) or the natural sellers (upper line) deploy maker algorithms.

A. Aggregate Outcomes: Pricing, Volume, and Efficiency

We compare aggregate outcomes in manual and algo-possible rounds to each other and to the socially-optimal outcome. We start by looking at prices, which we measure as the average price of the last twenty trades in a round.¹¹

Figure 5 summarizes the behavior of final prices. The gray, 45-degree line, indicates equality between socially-optimal prices and the price observed in the experiment. The black lines indicate the window of Nash equilibrium predicted prices (depending on what side adopts the maker role).

Figure 6 presents the average difference between the maximal and minimal price within each decile of trading times within a round. This difference proxies for bid-ask spread and the ten bins

¹¹By construction, MVO algorithms start a round trading at prices far from the socially-optimal price. Past experiments suggest this is the case also with human traders. Therefore, to compare to the social optimum, we focus on the end of each round.

proxy for its evolution over time. It is evident that spreads are larger in the robot-able rounds and they remain larger through time within rounds. This is consistent with the theoretical prediction of the robot-only game. Should everyone deploy a robot and not trade manually, prices follow marginal valuations of market makers modulo a spread. While in equilibrium, there is perfect coordination and all makers are on one side of the market (buyers or sellers), traders in the experiment do not perfectly coordinate and so prices bounce between the marginal valuations (modulo a spread) of buyers and sellers. Such following of marginal valuations bounce is not expected in the manual rounds and the prediction would be that the bid-ask spread in such rounds is lower than what would be predicted in the model of algorithms only. The data and the corresponding analysis confirm the predictions.

Both in manual and algo-possible rounds, the distance from socially-optimal prices is small but significant. As reported in Table 3, this distance is significantly larger in manual than in algo-possible rounds, despite bid-ask spreads being larger in the latter.



Figure 6: Price Spread Over Time Intervals

Explanations: This figure presents the bar plot of the average proxies for bid-ask spread using binned data in each trading round. For each bin, we group all transaction prices by deciles in the order in which they are executed and compute the *Bid-Ask Proxy* as the maximum transaction price minus the mininum transaction price within the corresponding decile. Then we plot the average *Bid-Ask Proxy* over time-intervals across session rounds for manual rounds and robot rounds separately.

We measure trade volume as the sum of all trades by participants in a round. We use the ratio of this volume to the trade volume needed for all participants to arrive at socially optimal allocations, to compare to this benchmark. We cannot reject the null hypothesis that trade volume in a round equals socially-optimal volume, neither for manual nor algo-possible rounds (Wilcoxon paired test with p-value 0.22 on the full sample, 0.15 for manual, and 0.19 for algo-possible rounds). We can

	Manual	Robot	Diff	t-stat
Price abs. difference (in cents)	8.923	4.845	4.078	1.927
Volume ratio	95.570%	94.547%	1.023%	0.080
Efficiency	55.649%	84.053%	-28.404%	-6.109
No. obs	12	48		

 Table 3: Aggregate outcomes. Difference between manual and algo-possible rounds.

Explanations: Aggregate indicators price, volume, and trade surplus, with respect to the social optimum. For price we use the absolute difference, for volume and trade surplus we use a ratio.

also not reject the null hypothesis that trade volume is equal in manual and algo-possible rounds (Table 3).¹²

The possibility of measuring efficiency of outcomes is one of the advantages of our experimental setup. For each set of parameters used and number of participants of each endowment type, we know what the trade surplus is at the social optimum given the induced utility. We can compare the sum of utility gains of all participants in a round to the total surplus using a ratio that is only equal to 100% if all gains from trade are reaped. Whereas full efficiency is never reached, the maximal ratio in a round is of 98.24%. As reported in Table 3, the difference in efficiency between manual and algo-possible rounds is large and significant.

Efficiency depends crucially on all participants' (equal) final holdings of Stock A. There is thus a close link between trade volume and efficiency, apparently broken in the results above: no difference in volume, but a large difference in efficiency between manual and algo-possible rounds. To reconcile these results, we turn briefly to the type of trading taking place in either type of rounds. We define a participant in a round to be a *speculator* if they both buy and sell Stock A. The fraction of speculators is significantly larger in manual than in algo-possible rounds (t-test statistic of 3.78). Thus, while trade volume is large because of speculation, final holdings of Stock A in manual rounds remain dispersed, significantly more so than in algo-possible rounds (Brown-Forsythe test with p-value less than 0.001).

Result 1: Prices, Volume, and Efficiency. Final prices and trade volume in all experimental rounds are close to the social optimum. Between manual and algo-possible rounds, final price is

¹²If, instead of round-wide volume, we use the volume traded by each participant and its ratio to per capita volume to achieve the social optimum, we find that volume is significantly smaller than predicted. There is, however, no significant difference between manual and algo-possible rounds.

closer to the social optimal in the latter, and volume is indistinguishable. Algo-possible rounds come close to full efficiency (84%), significantly more than manual rounds.

We now turn to individual trader behavior and performance in manual and algo-possible rounds.

B. Algorithm Deployment

We first examine participant behavior in algo-possible rounds and ask to what extent they delegated trading to algorithms, and possibly why. Figure 7 shows the distribution of the number of rounds in which a participant deploys an algorithm. Most participants deploy algorithms whenever available (eight rounds), and in all rounds of all sessions, a majority of participants of both endowment types deploys algorithms: the fraction of participants deploying algorithms in any given round ranges from 2/3 to 1.

Result 2: Algorithm Deployment, Trade Volume, and Delay Across Rounds. A majority of participants deploy algorithms in all rounds, and most deploy them within twenty seconds after the round's start. The number of algorithms deployed does not significantly change across rounds. The percentage of trade performed by algorithms decreases in later rounds.



Figure 7: Histogram of Algorithm Deployment Frequency.

Explanations: Number of participants who deployed robots in *N* of the 8 algo-possible rounds, N = 0, 1, ..., 8.

The percentage of participants deploying a robot who also perform manual trades in a given round ranges from 0% to 87.5%, and manual volume as a percentage of total traded volume in

	Robot Volu	me (Percentage)	Deployment Delay		
	(1)	(2)	(3)	(4)	
tech	0.012	0.038*	5.816	6.514*	
	(0.550)	(2.060)	(1.824)	(2.048)	
understand	0.061**	0.043*	-7.432*		
	(2.881)	(2.252)	(-2.227)		
understand robot				-5.329***	
				(-3.358)	
algopref	0.195***	0.064***	-13.455***	-12.077***	
	(8.751)	(3.157)	(-3.788)	(-3.403)	
makerpref	0.059***	0.065***	0.981	0.290	
	(3.615)	(4.975)	(0.432)	(0.128)	
spreadpref	0.020	0.034***	2.508	3.392*	
	(1.640)	(3.431)	(1.429)	(1.918)	
Constant	-0.292	-0.066*	62.864**	56.237**	
	(-1.945)	(-0.526)	(2.866)	(2.901)	
Adjusted R ²	0.170	0.143	0.059	0.071	
Sample	A11	Robot	Robot	Robot	
Sample	All	deploying	deploying	deploying	
Ν	569	498	498	498	

 Table 4: Algorithmic Trade Percentage and Deployment Delay Regressions.

a round ranges from 17.4% to 60.2%. On average, both the number of participants deploying robots who also perform manual trades and the percentage of total volume that is manually traded, increase in later rounds (a linear Round effect is significant at the 1% level).

Between participant differences in reliance on algorithms for trading, significantly correlate with exit survey responses. We group the survey responses in five indicators of technological affinity (*tech*), understanding of the determinants of each participant's own payoff (*understand*), the belief that algorithm deployment increases payoff (*algopref*), the belief that deploying a maker increases payoff (*makerpref*), and the belief that a high spread increases payoff (*spreadpref*)—see the Appendix for a detailed explanation. As can be seen in Table 4, we find that *understanding*, *algopref* and *makerpref* positively correlate with algorithmic trading volume in general. When considering all participants the variables positively correlate with the combined choice of deploying an algorithm and of limiting manual trades. When we limit the analysis to only participants who deploy a robot,

Explanations: Explanatory variables are indicators constructed with exit survey responses. Only algo-possible rounds are included in all regressions, and the *robot deploying* sample includes only participants who deploy an algorithm. All regressions include session fixed effects.

the variables correlate with the choice to limit manual trades, i.e., to delegate most of the the trades to the robot. In this latter subset, *tech* also correlates positively with algorithmic volume.

Result 3: Between-Participant Differences in Algorithm Use. Participants who perform a higher percentage of trades via their algorithmic robot and who deploy algorithms earlier in a round, report better understanding of algorithm deployment and a belief that algorithms increase their performance. They do not score higher than other participants on technological affinity questions.

C. CHOICE OF ALGORITHM PARAMETERS

We want to know what roles and spreads participants assign to their algorithms and what factors, if any, affect these choices. We want to test whether equilibrium predictions such as high maker spread and the sorting into roles by endowment type, arise in later experimental rounds.

C.1. Choice of Role

Result 4: Algorithm Role Choice. On average, an equal number of participants chooses the role of maker or taker. Participant role choice is reactive to past experience as role choice is sticky and high spreads increase the likelihood of being a maker in the future.

Figure 8 shows the percentage of all participants of each endowment type that deploys an algorithm (solid lines) with the role of maker (bars) in each round. Type I endowment is displayed on the bottom, and type II endowment is displayed on the top of each figure. As previously mentioned, a majority of participants deploys an algorithm in all algo-possible rounds. There is, on average, no preference for role: the median fraction of makers across rounds is 54.2% and the mean maker fraction is 55.2% (significantly larger than taker fraction at the 1% level). While the figures suggest the choice of role becomes polarized across endowment types in later rounds, statistical tests do not corroborate this suggestion.

In later rounds—starting with round 4— 60% of participants who choose the role of *taker* chose the same role in the preceding round, and only 40% were either makers or did not deploy an algorithm. Analogously, twice as many makers (2/3) were a maker in the preceding round than a taker or no algorithm. Significance and robustness of these results to the inclusion of other variables



Figure 8: Maker Algorithm Deployment Per Round.

Explanations: The figures display percentages of all participants of each endowment type. Endowment type I is on the bottom part, and endowment type II is on the top part of each figure. Lines display percentages deploying an algorithm, and bars percentages (of all participants) deploying a maker algorithm.

are displayed in Table 11 in the Appendix. The regression reported in this table also shows that the algorithm environment (algorithms deployed by other participants in previous rounds) have only a mild effect on participants' odds of switching role.

C.2. Choice of Spread

Figure 9, panel (a), shows a histogram of spreads of maker (pink) and taker (blue) algorithms set by each participant in each round. Across all sessions and rounds, the median and mean spread set by makers (mean of 0.174 ECU) is significantly larger than the median and mean spread of takers (mean of 0.154 ECU). Panel (b) shows how the average effect of role on spread is driven by experience, where experience is defined as the number of rounds during which a participant has deployed an algorithm with the same role. Inexperienced participants deploying a maker or a taker pick the same spread. With experience, maker spread increases and taker spread decreases, creating significant wedge between the two.



Figure 9: Spreads of Maker and Taker Algorithms.

Table 12 in the Appendix corroborates that role has a significant effect on spread, while the effect of experience is not significant. It also shows a correlation between performance and spread, as participants chosing high spreads tend to outperform in any given round.

Result 5: Spread Choice. Participants deploying maker algorithms set higher spreads than participants deploying takers. The difference increases with participant experience in the role.

D. INDIVIDUAL PERFORMANCE, DISTRIBUTION, AND EFFICIENCY

We now turn to understanding performance and the effect of algorithm deployment and parameter choices on individual and collective performance. As we do so, we analyze data from both manual and algo-possible rounds.

D.1. Algorithm deployment, performance, efficiency, and equality

We measure individual performance in each round as each individual's payoff as a proportion of the theoretically possible gain and call this benchmark-*performance ratio*. A performance ratio of 0 indicates that none of the theoretically possible gains from trade corresponding to an individual are realized. A ratio of 1 indicates that 100% of the theoretical gains from trade are realized. Performance ratios may be larger than 1 or smaller than 0, since theoretically possible gains from trade are computed at equilibrium prices and the experimental markets take a while to equilibrate,





Explanations: In both figures, participants are ranked according to their average performance ratio in manual rounds. Bars are color coded by the frequency of algorithm usage in algo-possible rounds.

resulting in a sizeable portion of the trades being conducted at off-equilibrium prices. Off equilibrium prices, some participants may exploit others, thus generating performance ratios outside the [0,1] range. For reference we present the performance measures in Table 5 below.¹³

The upper panel of Figure 10 illustrates the payoffs of all participants in the manual trading round in decreasing order. The lower panel displays the payoffs of the same participants in the robotic rounds. While the best performing participants in manual rounds also do very well in algopossible rounds, the worse-performing participants in manual rounds are not the worse-performing in algo-possible rounds. Importantly, fewer individuals are exploited on average (across all rounds) in algo-possible than in manual rounds, as fewer participants have negative average performance

¹³The round payoff was normalized to USD 15 per the Institutional Review Board requirements for one of the cites where we performed our experiments. Thus, if one would like to compare the average amount actually paid to participants for each round, this amount does not change and is always equal to USD 15. However, we claim it is a meaningful exercise to compute the pre-normalization average payoff and compare this amount across the manual and robotic trading rounds. The participants cannot know the overall gains from trade that are being realized in each round, thus the only objective they have is to do as well as they can for themselves. We think we can confidently conjecture that the results would be unchanged if a normalization did not take place (but then the study would be in violation of the local university's rules of engagement).

ratios. Below we formalize the following conjectures stemming from the figure. (i) those who perform worse in the manual rounds are those who benefit the most from the robots; and (ii) the inequality between participants decreases in robotic rounds.

			Robot -	Manual	
	Manual	Robot	Diff	t-stat	Nobs
Low	0.05	0.74	0.69	8.12	25
Medium	0.54	0.81	0.28	3.75	25
High	1.08	0.98	-0.10	-1.29	22
Overall (individual paired)	0.55	0.84	0.30	4.92	72

Table 5: Performance Ratio Results

Explanations: The performance ratio (realized utility gain divided by the equilibrium utility gain) across three groups (low, medium, and high) based on subjects' manual trading performance. The t-stat is from a paired t-test.

Overall, as Table 5 shows, the performance increases by 0.3. Thus, when trading manually, participants utilize slightly more than half of the possible gains from trade on average, while when they utilize algorithms, this percentage increases to 84%, for a relative increase of 53% in algorithmic vs. manual rounds. The above measures are also those for the efficiency of the market, an average of 100% would indicate that the markets have exhausted all possible gains from trade. Who benefits the most from the increase of individual performance that is enjoyed in the robotic rounds? To investigate this, we split participants in three groups, Low, Medium and High, based on their performance ratio in the manual trading rounds. Participants classified in the group Low, fall in the bottom third tercile, when ranked based on their performance ratio in the two manual rounds. Groups High and Medium correspond to the first and the second terciles. As the table shows, the disparity between groups in terms of performance ratio is large (and statistically significant). Keeping the classification from manual rounds, we compute the corresponding average performance ratios for the three groups in the rounds that allow algo trading. The largest improvement is of the Low group, followed by the Medium and then the High. Both the Low and the Medium groups enjoy economically and statistically significant improvements, while the slight decline in performance of the High group is not significant.

Notice, the High group has a performance ratio in excess of 100%, i.e, these traders realize more than their equilibrium gains from trade. This is possible due to some of them assuming

market-making roles and in addition to the gains in utility, realizing gains from the spread in market-making. Also, note that in algorithmic rounds, such possibilities disappear, but still, the highly skilled traders make all of the possible gains from trade.

Table 6 presents the formal test of the notion that the total (bigger) pie in the robotic sessions is one that is more evenly split between the three groups of participants. For each session one group, we collect the share of payoffs distributed to each group (note here that normalizations play no role, or minimal role at best). The difference of the proportion of total payoffs between manual and robotic rounds is computed for each session and group. This difference is the largest for the Low group, followed by the Medium, and lastly by the High.¹⁴. This comes to show, one more time, that the usage of the robots significantly decreases the inequality of payoffs among participants based on manual trading skills.

	Share of Pie			
	Manual	Robot	Diff (Robot - Manual)	
Low	0.0241	0.2940	0.270	
Medium	0.3115	0.3203	0.009	
High	0.6645	0.3857	-0.279	
	Wilcoxon Signed Rank Test, H ₀		p-value	
	Diff(low) = Diff (Medium)		0.03	
	Diff(Medi	um) = Diff (High)	0.03	

Table 6: Proportion of Performance Gai
--

Explanations: This table describes how the utility gains are split among three groups defined based on manual performance. Each number is computed as $\pi_j / \sum_j \pi_j$, where π stands for the performance ratio and j = low, medium, high.

Result 6: Private and Social Gains from Algorithm Availability. The availability of algorithms (algo-possible rounds) increases individual performance across the board. It also increases efficiency and decreases inequality. Increased volume of trade by algorithms improves individual performance of algorithm-deploying participants, and decreases inequality, but has no effect on efficiency.

¹⁴We use Wilcoxon signed rank test to establish the statistical significance

D.2. Algorithm parameters, trading behavior, and performance

Table 7 presents the spread choices of the deployed maker and taker robots across the three manual trading skill level groups. The theoretical model predicts that makers post maximal spread, while takers should have minimal spreads. The null that the spreads are equal is rejected only in the high-skilled group. The skilled traders have the lowest spreads among both maker and taker robots. Off-equilibrium best response in the game of choosing spreads would dictate slight undercutting of spreads, and we hypothesize that these spread choices of the high-skilled traders contribute to their better performance.

			Taker –	Maker
	Maker	Taker	Diff	t-stat
Low	0.1709	0.1706	0.000	0.04
Medium	0.1862	0.1650	-0.021	-1.14
High	0.1623	0.1193	-0.043	-3.18
Overall	0.1745	0.1484	-0.026	-2.56

Table 7: Spread Choice across Manual Performance Groups

Explanations: This table describes how makers and takers make spread choices across three groups defined based on manual performance (Low, Medium, High) and the whole sample (Overall).

Table 8 shows how algorithm parameter choices affect the performance ratio. As can be seen in Table 8, the performance ratio is significantly larger for participants who deploy taker algorithms, execute a high volume of trades with their algorithms, and set large spreads. Theory shows that maker algorithms perform at least as well as takers because *in equilibrium* such algorithms are deployed with very large spreads. Our results reflect the prior observation that maker spreads are significantly, but not largely different from taker spreads. A such, we are far from the equilibrium-predicted spread gap (and its consequent negative effect on efficiency).

	(1)	(2)	(3) Performance	(4)	(5)
Taltar	0.050***	0.005***	Terrormanee	0.000***	0.041
laker	0.258	0.335		0.280	0.041
	(4.65)	(3.54)		(5.20)	(0.27)
taker×spread		-0.639			0.864**
		(-1.41)			(2.34)
Spread		1.059^{***}		0.866***	
		(3.41)		(3.89)	
Robot Volume			0.003***	0.004***	0.006***
			(3.10)	(4.52)	(4.53)
hvbrid				-0.019	0.050
5				(-0.30)	(0.88)
Human Volume				0.002**	0.002**
Truman vorume				(2,33)	(2.16)
makaryanroad				(2.55)	1 501***
maker×spreau					1.521
,					(4.83)
maker					-0.415**
					(-2.40)
Constant	0.743***	0.600***	0.665***	0.258***	0.342***
	(21.20)	(9.03)	(9.41)	(2.61)	(3.10)
Adj. R ²	-0.008	0.016	-0.020	0.069	0.077
Ν	573	573	573	573	573
Round FE	Yes	Yes	Yes	Yes	Yes

Table 8: Performance Regression.

Explanations: *hybrid* indicates the participants used both algorithms and manual trades. *Human Volume* measure the trading volume executed by manual trades. *Robot Volume* measures the trading volume executed by algorithms. *taker* is an indicator for deploying taker-type algorithm.

We cannot rule out the possibility that our top performers rely on manual trades that complement algorithmic trades to enhance their performance. As can be seen in table 9, there is no significant difference in the total volume and the fraction of this trade that is performed by algorithms, between the three performance groups. There is, however, a significant difference in the fraction of all robotic trade being performed by a taker robot, with the High performance type using takers significantly more. Thanks to this usage and accurate manual trading, this type generally trades at more convenient prices. Figure 11 displays the empirical cumulative distribution functions of buy and sell prices across the three types. It is immediately evident that the High type's sale prices first order stochastically dominate those of the other groups. It is less obvious that their buy prices are stochastically dominated by the other types. A regression, reported in Table 13 in the Appendix, shows that this type's buy prices are indeed smaller than for other types, while sell prices are higher.

	Total*	Robot	Taker	Nobs			
Low	0.933	0.690	0.290	25			
Medium	0.993	0.666	0.312	25			
High	0.953	0.728	0.487	22			
T-statistic of difference between groups							
Low v. Medium	-0.578	0.251	-0.287				
Low v. High	-0.210	-0.435	-2.472				
Medium v. High	0.399	-0.209	-2.006				

 Table 9: Algo-possible Rounds' Volume by Performance Type.

Explanations: Total volume is given as a fraction of the socially-optimal trading volume. Robot volume is given as a fraction of total volume, and Taker volume is given as a fraction of all robotic volume.



Figure 11: Trading Prices by Performance Type.

Explanations: Performance types are created on the basis of performance in manual rounds only. Displayed buy (upper figure) and sell (lower figure) prices are for algo-possible rounds only.

IV. CONCLUSION

The effect of automation on individual financial market performance is both important and difficult to observe. In this work we use experimental market to investigate how agents program and use simple robots. On the one hand, robots relieve agents from the burden of computing their optimal trading strategy throughout the trading session which benefits poor traders. However, those that understand robots can deploy them or also engage in manual trading to better effect.

Overall, our work highlights that while in theory there is no difference between a Bayesian decision maker and a perfect algorithm, in the experimental markets we have analyzed there are differences. The choice of robots affected not only the size of the gains from trade and how they are split among the traders, but also the prices generated by the market. Indeed, even though allocations may be close to efficient, prices can deviate significantly because of how the gains from trade are split through robot and spread choice. In real markets we do not observe valuations, but rather infer them from price. It is useful to know the extent to which prices and trading volume differ because of algorithms both relative to efficient benchmarks and manual trading.

Finally, we note that all our agents had the same choice of robots, but this still generated heterogeneity in chosen spreads and aggressiveness (role choice). In real markets, participants have access to many different algorithms which could exacerbate the cross-sectional effects of robot-choice.

REFERENCES

- Aldrich, E. M., and K. López Vargas. 2020. Experiments in high-frequency trading: comparing two market institutions. *Experimental Economics* 23:322–352.
- Asparouhova, E., P. Bossaerts, and J. Ledyard. 2020a. Price Formation in Multiple Simultaneous Continuous Double Auctions, With Implications For Asset Pricing. Available from author's webpage and on https://ssrn.com/abstract=2774301 or http://dx.doi.org/10.2139/ssrn.2774301.
- Asparouhova, E., P. Bossaerts, K. Rotaru, T. Wang, N. Yadav, and W. Yang. 2020b. Humans in Charge of Trading Robots: The First Experiment. Available from author's webpage.
- Bao, T., E. Nekrasova, T. Neugebauer, and E. Y. Ryanto. 2021. Algorithmic Trading in Experimental Markets with Human Traders: A Literature Survey. Available from author's webpage and on https://ssrn.com/abstract=3908065.
- Ben-David, D., and O. Sade. 2021. Robo-Advisor Adoption, Willingness to Pay, and Trust–Before and at the Outbreak of the COVID-19 Pandemic. Available on SSRN.
- Colliard, J.-E., T. Foucault, and S. Lovo. 2022. Algorithmic Pricing in Securities Markets HEC Paris Working Paper.
- Das, R., J. E. Hanson, J. O. Kephart, and G. Tesauro. 2001. Agent-human interactions in the continuous double auction. In *International Joint Conference on Artificial Intelligence*, vol. 17, pp. 1169–1178.
- De Luca, M., C. Szostek, J. Cartlidge, and D. Cliff. 2011. Studies on interactions between human traders and Algorithmic Trading Systems. *Foresight, The Future of Computer Trading in Financial Markets, Driver Review DR13, Crown Copyright.* .
- Dence, S., D. Latimore, and J. White. 2006. The trader is dead, long live the trader!-A financial markets renaissance. *IBM Institute for Business Value*.
- Farjam, M., and O. Kirchkamp. 2015. *Bubbles in hybrid markets-How expectations about algorithmic trading affect human trading*. CESifo Working Paper Series.

- Furse, D., P. Bond, D. Cliff, C. Goodhart, K. Houstoun, O. Lington, and J. Zigrand. 2012. The Future of Computer Trading in Financial Markets: An International Perspective. *Foresight, Government Office for Science, London, UK*.
- Gjerstad, S. 2007. The competitive market paradox. *Journal of Economic Dynamics and Control* 31:1753–1780.
- Gode, D. K., and S. Sunder. 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of political economy* pp. 119–137.
- Greiner, B., P. Grünwald, T. Lindner, G. Lintner, and M. Wiernsperger. 2022. Incentives, Framing, and Trust in Algorithmic Advice: An Experimental Study Working Paper.
- Grossklags, J., and C. Schmidt. 2006. Software agents and market (in) efficiency: a human trader experiment. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 36:56–67.
- Kirilenko, A. A., and A. W. Lo. 2013. Moore's Law versus Murphy's Law: Algorithmic Trading and Its Discontents. *Journal of Economic Perspectives* 27:51–72.
- Kissell, R., and R. Malamut. 2005. Understanding the profit and loss distribution of trading algorithms. *Algorithmic Trading* pp. 41–49.
- Shachat, J., L. Weib, P. Yanc, and S. S. Zhangd. 2020. Speed traps: Algorithmic trader performance under alternative market structures. Tech. rep., Chapman University ESI Working Papers Series.
- Smith, V. L. 1962. An experimental study of competitive market behavior. *The Journal of Political Economy* pp. 111–137.
- Smith, V. L., G. L. Suchanek, and A. W. Williams. 1988. Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica* 56:1119–1151.
- Venkatesh, V., J. Y. Thong, and X. Xu. 2012. Consumer Acceptance and Use of Information Technology: Extending the Unified Theory of Acceptance and Use of Technology. *MIS Quarterly* 36:157–178.

EXIT SURVEY QUESTIONS, ANSWERS, AND CONSTRUCTED VARIABLES

After participating in all experimental rounds, and before payment, participants participated in an exit survey administered on using Qualtrics, and containing the following questions:

- 1. What was your trader ID? (M#)
- 2. Please indicate your age.
- 3. What is your gender?
- 4. If you are a student, please tell us your major.
- 5. To what extent do you agree/disagree with the following statements?
 - (a) I find technology useful in my daily life.
 - (b) Using technology helps me accomplish things more quickly.
 - (c) Using technology increases my productivity.
 - (d) Learning how to use technology is easy for me.

Five-level Likert scale ranging from strongly disagree (1) to neither agree nor disagree

- (3), to strongly agree.
- 6. To what extent were each of the following aspects of the experiment clear to you?
 - (a) How to use the interface to place orders and trade manually
 - (b) How to use the interface to deploy an algorithm
 - (c) How your payoff was calculated from your holdings in a given round
 - (d) How your overall payoff for the experiment was determined
 - (e) How your valuation for Stock A changed as you traded
 - (f) How algorithms determined the trades they tried to make on your behalf

Five-level Likert scale ranging from *I* don't know (0), and not at all clear (1) to to very clear (5).

7. To what extent do you believe the following choices were good or bad for your payoff in a given round?

- (a) Launching an algorithm
- (b) Trading manually alongside an algorithm
- (c) Trading manually only
- (d) Conditional on launching an algorithm, choosing it to be a maker
- (e) Conditional on launching an algorithm, choosing it to be a taker
- (f) Setting a high spread when you launched a maker algorithm
- (g) Setting a high spread when you launched a taker algorithm

Six-level Likert scale ranging from *I* don't know (0) and very bad (1) to very good (includes *irrelevant* (3)).

- 8. Do you believe the choice of algorithm type (maker or taker) by other participants affected your payoff? **Three levels of answer:** *Yes, No, and I don't know*
- 9. Do you believe the choice of spread set on the algorithms deployed by other participants affected your payoff? **Three levels of answer**: *Yes*, *No*, and *I don't know*
- 10. What feedback, if any, would you like to provide to the experimental team?

Answers to survey questions, tabulated in Table 10, were used to create several indicators, as follows:

- *Tech*: Sum of answers to survey questions 5(a) to 5(d). Questions 5(a) to 5(c) correspond to the *Performance Expectancy* questions on technology adoption taken from Ben-David and Sade (2021), in turn based on Venkatesh et al. (2012). Question 5(d) corresponds to the first of their *Effort Expectancy* questions on technology adoption.
- *Understand Basic*: Sum of answers to question 6(a) and 6(b). An indicator of basic understanding of the trading interface.
- *Understand*: Sum of answers to questions 6(a) to 6(f). An indicator of understanding of all aspects useful for trading.
- *Understand robot*: Sum of answers to questions 6(b) and 6(f). An indicator of understanding of the usage and functioning of algorithms.

	Question 5						
Possible Responses	(a)	(b)	(c)	(d)			
Strongly disagree	1	1	1	1			
Somewhat disagree	2	3	3	1			
Neither agree nor disagree	2	1	5	5			
Somewhat agree	11	13	26	30			
Strongly agree	56	54	37	35			
Total	72	72	72	72			
			Ç	Question 6			
Possible Responses	(a)	(b)	(c)	(d)	(e)	(f)	
Not at all clear	0	1	6	0	0	2	
A little clear	8	11	20	11	14	16	
Clear	33	27	21	29	31	31	
Very clear	30	30	23	29	24	21	
I don't know	1	2	2	3	3	2	
Total	72	71	72	72	72	72	
			Ç	Question 7			
Possible Responses	(a)	(b)	(c)	(d)	(e)	(f)	(g)
Very bad	1	3	6	0	2	1	2
Bad	10	6	14	7	8	14	14
Irrelevant	6	5	5	8	7	9	9
Good	34	23	26	25	32	24	24
Very good	21	34	18	28	18	19	16
I don't know	0	1	3	4	5	5	7
Total	72	72	72	72	72	72	72
Possible Responses	Ques	tion 8	Ques	stion 9			
I Don't Know		2		6			
No		5		9			
Yes	6	65		57			
Total	5	72		72			

 Table 10: Frequency of Survey Responses.

- *Algopref*: The answer to question 7(a) plus 5 minus the answer to question 7(b), plus 5 minus the answer to question 7(c). Answers "I don't know" to any of the questions have a contribution of 0 to the indicator. An indicator of participant belief that launching an algorithm is more beneficial for performance than manual trading.
- *Makerpref*: The answer to question 7(d) plus 5 minus the answer to question 7(e). Answers "I don't know" to either question have a contribution of 0 to the indicator. An indicator of

participant belief that launching a maker algorithm is more beneficial for performance than a taker algorithm.

• *Spreadpref*: Sum of answers to questions 7(f) and 7(g). An indicator of participant belief that setting a high spread is beneficial for performance.

PROOFS

Proof of Lemma 1 Follows immediately from arguments in the text.

Proof of Lemma 2

The triangle whose area determines the efficiency loss has a base equal to $(\delta^S + \delta^D)/2$ and a height equal to the difference between the zero-spread trade quantity, $(e_A^S - e_A^D)/2$, and the actual trade quantity, x*. The result follows.

Proof of Lemma 2

Follows from arguments in the text.

Proof of Proposition 1

From the proof of Proposition 2 when spread choice is endogenous, it is optimal to choose to be a taker when the other player is a maker.

1. Assume the buyer is a maker and the seller, a taker, wants to optimally set their spread, δ_s . They do so buy maximizing their payoff, given below.

$$\Pi_{S}(\delta_{S},\delta_{D},T,M) = \left(\frac{\rho_{DA}(0) - \rho_{SA}(0)}{2} - \frac{\delta_{S} + \delta_{D}}{4}\right)x^{*} + \frac{\delta_{S}}{2}x^{*}$$
$$= \left(\frac{\gamma\sigma_{A}(e_{S} - e_{D})}{2} + \frac{\delta_{S} - \delta_{D}}{4}\right)x^{*}$$
$$= \left(\frac{\gamma\sigma_{A}(e_{S} - e_{D})}{2} + \frac{\delta_{S} - \delta_{D}}{4}\right)\left(\frac{e_{S} - e_{D}}{2} - \frac{\delta_{S} + \delta_{D}}{4\gamma\sigma_{A}}\right)$$

The above concave function of δ_S has a zero derivative if and only if $\delta_S = 0$.

2. We find the optimal spread of the maker (buyer, in our case) given that the seller sets a spread of zero. When facing a taker, the maker can only reap the spread from each transacted unit, leading to a payoff

$$\Pi_D(0,\delta_D,T,M) = \frac{\delta_D}{2} x^* = \frac{\delta_D}{2} \left(\frac{e_{SA} - e_{DA}}{2} - \frac{\delta_D}{4\gamma \sigma_A^2} \right).$$

With this, the first order condition for optimality, yields

$$\frac{\partial \Pi_D(\delta_D; \cdot)}{\partial \delta_D} = 0 \iff \frac{\delta_D}{2} x^* = \frac{\delta_D}{2} \left(\frac{e_{SA} - e_{DA}}{2} - \frac{\delta_D}{4\gamma \sigma_A^2} \right) \iff \delta_D(\delta_S) = \gamma \sigma_A^2(e_{SA} - e_{DA}) = \rho_{DA}(0) - \rho_{SA}(0).$$

Notably, this spread size means that at the onset (first traded unit), the maker absorbs half the trade surplus, and this fraction becomes even larger at ensuing units.

- 3. The traded quantity is found by replacing the maker's spread in the expression for x^* derived before. This immediately gives $\frac{(e_{SA}-e_{DA})}{4}$, which is half the quantity when spreads of both traders are zero (efficient outcome).
- 4. We replace the traded quantity in the expression for the buyer's bid function or the seller's ask function to obtain the price of the last traded unit

$$p^* = \mu_A - \gamma \left[\sigma_A^2 \frac{e_{SA} + e_{DA}}{2} + \sigma_{AB} e_B \right] - \gamma \sigma_A^2 e_{DA} = p^*_{0-spread} - \gamma \sigma_A^2 e_{DA}$$

In our proofs we have taken the buyer to be the maker and we have found the last trade price to be *lower* than the last trade price with zero spreads. Thus, the bias is in favor of the maker.

5. The efficiency loss in equilibrium is found by replacing the equilibrium spreads (0 and δ_D^*) in the expression of Lemma 1. This yields

$$EL^* = \frac{1}{16} \gamma \sigma_A^2 (e_{SA} - e_{DA})^2 = \frac{1}{16} \frac{(\rho_{DA} - \rho_{SA})^2}{\gamma \sigma_A}$$

The last part of point 5 requires that we find optimal spreads if both players are makers. This is obtained from payoff maximization with each player having equal odds of being the price setter. This is

$$\begin{split} \Pi_D(\delta_S, \delta_D, M, M) &= \frac{1}{2} \bigg(\frac{\rho_{DA}(0) - \rho_{SA}(0)}{2} - \frac{\delta_S + \delta_D}{4} \bigg) x^* + \frac{\delta_D}{2} x^* \\ &= \frac{1}{2} \bigg(\frac{\gamma \sigma_A^2(e_{SA} - e_{DA})}{2} + \frac{3\delta_D - \delta_S}{4} \bigg) \bigg(\frac{e_S - e_D}{2} - \frac{\delta_S + \delta_D}{4\gamma \sigma_A^2} \bigg), \end{split}$$

and the first-order condition gives:

$$\begin{split} \frac{\partial \Pi_D(\delta_D; \cdot)}{\partial \delta_D} &= 0 \Longleftrightarrow \frac{1}{2} \left(\frac{e_{SA} - e_{DA}}{2} - \frac{\delta_S}{4\gamma \sigma_A^2} \right) = \frac{3\delta_D}{8\gamma \sigma_A^2} \Longleftrightarrow \\ \delta_D(\delta_S) &= \frac{2\gamma \sigma_A^2 (e_{SA} - e_{DA})}{3} - \frac{\delta_S}{3}. \end{split}$$

By symmetry, $\delta_S(\delta_D) = \frac{2\gamma\sigma_A(e_S - e_D)}{3} - \frac{\delta_D}{3}$ and, thus, optimal buyer and seller spreads are equal and equal to $\delta_D = \delta_S = \frac{1}{2}(\rho_{DA} - \rho_{SA})$. This is half as large a spread as in the maker-taker case. However, as this spread is set by both traders while in the maker-taker case the taker sets a zero spread, the efficiency loss is the same as in equilibrium.

6. To obtain the maker's payoff we replace the equilibrium spread and traded units in the maker buyer's payoff function:

$$\Pi_D\left(0,\delta_D^*,T,M\right)=\frac{\delta_D^*}{2}x^*=\frac{\gamma\sigma_A^2}{8}\left(e_{SA}-e_{DA}\right)^2,$$

To obtain the seller's payoff, we notice it equals the size of the triangle with base equal to half the initial surplus (at x = 0) and height the total traded units in equilibrium, $x^* = \frac{e_{SA} - e_{DA}}{4}$. This yields

$$\Pi_{S}(0,\delta_{D}^{*},T,M) = \frac{1}{2} \left[\frac{e_{SA} - e_{DA}}{4} \times \frac{\gamma \sigma_{A}^{2}(e_{SA} - e_{DA})}{2} \right] = \frac{\gamma \sigma_{A}^{2}(e_{SA} - e_{DA})^{2}}{16}.$$

Thus, in equilibrium, the maker obtains twice as large a payoff than the taker.

2N-PERSON ALGORITHM SELECTION GAME

Consider a replica economy, with *N* players of each of the two endowment types introduced in the 2-person game above. Thus, there are *N* players with e_{SA} initial units of Stock A and an equal number of players with e_{DA} initial units of A. The intuition and, to some extent, formal results of the 2-player game carry over to this more complex setting, as shown in the propositions below.

Proposition 2 (2N traders, role and spread choice.) When traders choose both their role and their spread, the following actions of players constitute two equilibria of the game:

- 1. One side of the market (buyers or sellers) is all taker robots and all takers set a spread equal to zero.
- 2. The other side of the market is all maker robots where the spreads are decreasing as follows

$$\delta_{N-r} = \frac{2N-r+\frac{1}{2}}{2N-r+1}\delta_{N-r+1},$$

for r = 1, ..., N - 1 and δ_i denotes the spread of trader *i*, i = 1, ..., N - 1. For ease of notation, traders have been numbered in increasing order of their equilibrium spreads. Also, the largest spread, δ_N , equals

$$\delta_N = \Delta = \gamma \sigma_A^2 (e_{SA} - e_{DA})$$

3. As $N \to \infty$, gains from undercutting other makers' spreads to increase the number of traded units become infinitesimal, effectively equating the game to the 2-person game. As a result, like in the 2-person game, all makers set $\delta_i = \Delta$, the traded quantity is half the zero-spread (efficient) quantity, and the makers obtain twice as large a payoff as the takers.

The asymmetry present in the equilibrium described in Proposition 3, stems from a mixture of robot and market mechanism (CDA) characteristics, that is relevant also for commercial markets and algorithms therein. However, it is unappealing as a model to test in the reduced time of an experimental session: participants not only must coordinate in their choice of role (across endowment types), but also in their choice of ranked spreads when makers. The game can be simplified by restricting participant choices of spread, which is what we do in the experiment. Recalling that $\Delta = \gamma \sigma_A^2 (e_{SA} - e_{DA})$ is the gain from trade when no trade has yet occurred, let

the individual choice of spread be constrained above by $\bar{\delta} \in [\frac{1}{2}\Delta, \frac{5}{8}\Delta]$. The equilibria of this constrained game are characterized in the following proposition.

Proposition 3 (2N traders, role and constrained spread choice.) When traders choose both their role and their spread but spread is constrained to be $\delta_i < \overline{\delta}$, with $\overline{\delta} \in [\frac{1}{2}\Delta, \frac{5}{8}\Delta]$, two equilibria exist where the traders adopt different roles. The equilibria have the following characteristics:

- 1. One side of the market is all taker robots and all takers set a spread equal to zero.
- 2. The other side of the market is all maker robots where they all choose

$$\delta_i = \bar{\delta}$$

ADDITIONAL RESULTS

	(1)	(2)	(3)	
	taker	maker	switch	
Perf Ratio _{$t-1$}	0.033	0.075	-0.896***	
	(0.19)	(0.43)	(-4.65)	
taker _{t-1}	1.511***	-0.902***		
	(7.17)	(-4.35)		
Robot Volume $_{t-1}$	0.007**	0.006**	-0.001	
	(2.32)	(2.13)	(-0.48)	
Human Volume $_{t-1}$	-0.000	-0.005**	-0.004**	
	(-0.20)	(-2.27)	(-1.96)	
Spread_{t-1}	-1.391	2.500***	1.904**	
	(-1.58)	(2.95)	(2.26)	
Type II	0.163	-0.399**	-0.324	
••	(0.79)	(-1.97)	(-1.59)	
samemakers $_{t-1}$	0.130	-0.068	-0.134*	
t I	(1.57)	(-0.84)	(-1.67)	
othermakers $_{t-1}$	0.104	-0.038	0.031	
	(1.32)	(-0.51)	(0.41)	
Constant	-2.049***	0.032	0.661*	
	(-4.85)	(0.08)	(1.78)	
Pseudo. Rsq.	0.103	0.081	0.059	
N	499	499	499	

Algorithm Role and Spread Choice

Table 11: Logistic Model of Choice of Algorithm Role.

Explanations: Dependent variables are dummies equal to 1 when a participant chooses role "taker" for their deployed algorithm (column 1), or the role "maker" (column 2), or switches roles between the previous and the current round (column 3).

Table 12 investigates if other factors besides algorithm role affect spread choice. First, while makers set higher spreads, the added effect of experience illustrated in Figure 9 is not significant (coefficient of *serial maker*). Further, the only environmental element that significantly affects chosen spreads is past average spread, showing a form of inertia or arms race in spread setting. Importantly, participants setting a high spread tend to outperform the median participant (column (1), variable Outperform).

	(1)	(2)
	Spread	Spread
serial maker _t	0.007	0.013
	(0.29)	(0.76)
maker _t	0.059***	
	(4.82)	
Robot Volume _t	-0.000	
	(-0.44)	
Outperform _t	0.028***	
	(2.70)	
Average Spread $_{t-1}$		0.581^{***}
		(9.49)
$Outperform_{t-1}$		0.005
		(0.54)
$maker_{t-1}$		0.003
		(0.27)
Constant	0.105***	0.056***
	(9.19)	(6.43)
Adj. Rsq.	0.109	0.378
Ν	573	499
Round FE	Yes	Yes

Table 12: Choice of Spread (independent of the type of robot chosen).

Explanations: Dependent variable in both columns is the bid-ask spread participants choose for their algorithms. serial maker indicates when a participant has deployed a maker in five or more preceding rounds. Outperform indicates whether the participant's performance is above the median level in a given round. Average spread is the average of all participants' spreads in a given round.

TRADE PRICES BY PERFORMANCE TYPE

Table 13 displays regression results showing that participants of the High performance type sell at significantly higher prices and buy at significantly lower prices than participants of either the Medium or Low type.

	(1) Buy Price	(2) Sell Price	(3) Buy Price	(4) Sell Price
Constant	73.89	70.34	74.81	70.76
	(210.84)	(197.86)	(213.41)	(202.19)
Performance Type Low			-0.92	-0.42
			(-3.39)	(-1.60)
Performance Type Medium	0.92	0.42		
	(3.39)	(1.60)		
Performance Type High	-3.72	6.51	-4.64	6.09
	(-13.48)	(23.48)	(-17.23)	(22.31)
Adj. R ²	0.112	0.123	0.112	0.123
Ν	28011	28011	28011	28011
Round FE	Yes	Yes	Yes	Yes

 Table 13: Buy and Sell Prices (in US\$ cents) by Performance Type.

Explanations: Regression of Buy price (Columns (1) and (3)) and Sell price (Columns (2) and (4)) as a function of performance type. We control for Round fixed effects, as socially-optimal prices vary from round to round. Columns (1) and (2) have as reference the prices of the Low type. Columns (3) and (4) have as reference the Medium type. That way, the significantly more convenient prices of the High type are revealed with respect to either other type.