

Marner-Hausen, Mark

Working Paper

Developing a framework for real-time trading in a laboratory financial market

ECONtribute Discussion Paper, No. 172

Provided in Cooperation with:

Reinhard Selten Institute (RSI), University of Bonn and University of Cologne

Suggested Citation: Marner-Hausen, Mark (2022) : Developing a framework for real-time trading in a laboratory financial market, ECONtribute Discussion Paper, No. 172, University of Bonn and University of Cologne, Reinhard Selten Institute (RSI), Bonn and Cologne

This Version is available at:

<https://hdl.handle.net/10419/262334>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

ECONtribute
Discussion Paper No. 172

**Developing a Framework for Real-Time
Trading in a Laboratory Financial Market**

Mark Marner-Hausen

June 2022

www.econtribute.de



Developing a Framework for Real-Time Trading in a Laboratory Financial Market*

Mark Marner-Hausen[†]

This version: January 31, 2022

Abstract

One of the challenges that economic experiments that use artificial financial markets to explore high-frequency trading face, is the development of a sufficiently sophisticated software. Moreover, it is not trivial to adequately communicate the complex financial market rules to non-experts. The present paper is part of an ongoing project with Peter Cramton, Daniel Friedman, Kristian Lopez Vargas, and Axel Ockenfels in which a novel framework enabling algorithmic real-time trading at millisecond speeds is being developed. This novel framework provides a more accurate laboratory replication of the financial market mechanisms relevant to high-frequency trading than has been achieved up to this point. This will provide a basis for comparing the current financial market design with new, exciting market design approaches, both under normal and stressful market conditions. The ongoing project includes the development of the theoretical foundations, as well as the experimental design and the analysis of the corresponding data. The contribution of the present study consists of deriving parameters for the replication of short-lived financial market crashes that can be adopted by the new framework for real-time trading; to provide means for an adequate communication of complex financial market rules to non-experts; and to provide solutions to technical and conceptual difficulties encountered in the preparation for the realization of the experiment. In addition, a thorough review of the literature most relevant to the new framework for real-time trading is provided.

*Marner-Hausen gratefully acknowledges funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 741409) and funding from by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866. Special thanks is due to all collaborators of the joint project, in Santa Cruz, this includes Daniel Friedman and Kristian Lopez Vargas, in Cologne, Peter Cramton and Axel Ockenfels, for developing the theoretical background, the experimental design, and the derivation of the research questions along with the methodology to answer them. In addition, I would like to express my gratitude to the LEEPS lab at UCSC, led by Kristian Lopez Vargas, who developed the software for the new real-time trading framework, and to Max R.P. Gorßmann, representing the CLER lab at UoC, for providing crucial support in running the experiment.

[†]Mark Marner-Hausen: University of Cologne, marner-hausen@wiso.uni-koeln.de

JEL Codes: XX

Keywords: Financial Market Design, High-Frequency Trading, Frequent Batch Auction, Experimental Economics, Artificial Financial Markets

1 Introduction

High-Frequency Traders (HFTs) are market participants that can typically be identified by the following distinct characteristics. They are best known for algorithmic trading almost at the speed of light and unlike traditional investors, HFTs have no intrinsic motivation to own or owe a particular asset. Instead, the objective of HFTs is to profit from trading through strategies based on turning over assets quickly. The inventory positions they built up during the respective short holding periods, is subject to the risk of adverse price movements. This risk is evidently relevant to HFTs, as they only maintain an average inventory of less than 3.95% of their daily trading volume and tend to end the trading day with an inventory of zero or close to zero (Anand and Venkataraman, 2013; Korajczyk and Murphy, 2019). HFTs can further be identified by the very specific investments they undertake, which form the basis of their competitive speed advantages. They invest in new links between exchanges to ensure the shortest possible transmission time of information (Laughlin et al., 2014), they buy sites near or in the same building as the exchanges on which they operate and subscribe to proprietary data feeds (Budish et al., 2019; Baker and Gruley, 2019; Rogow, 2012), and they spend significant sums on human capital to develop the fastest possible algorithms (Lockwood et al., 2012; Aquilina et al., 2020). Lastly, HFTs differ from traditional market participants in that they frequently compete in a race for near risk-free profits. These races are fought among HFTs and usually start as soon as a virtual arbitrage opportunity presents itself. This can essentially happen whenever new price-affecting information reach the market. The price-time priority in the continuous time of modern financial markets facilitates such races in the first place, as even symmetric and publicly available information generates near-arbitrage opportunities. Only the first to react can profit from such opportunities, which leads to a constant competition in speed between HFTs.

That raises both academic and public-policy concerns. Fundamental economic theory would typically assume that arbitrage opportunities are competed away over time, yet, the pioneering work by Budish et al. (2015) reveals that the annual size of virtually risk-free profit opportunities in modern financial markets remains constant. In a subsequent analysis, Aquilina et al. (2020) estimate that the annual global arbitrage price pool in equity markets amounts to about \$5 billion. Competition is limited to technological speed

improvements which caused a reduction in information transmission time between the Chicago Mercantile Exchange (CME) and the New York Stock Exchange (NYSE) from roughly 16 milliseconds in 2010 to nearly the speed of light today (light requires about 4 milliseconds to travel the associated distance). The question is whether these incremental increases in communication speed, which accelerate price discovery but do not result in new or improved information, are desirable or rather socially wasteful. If the latter is true, it may indicate a flawed market design, which allows HFTs to generate substantial, almost risk-free profits at the expense of ordinary investors due to their speed advantages and complex trading algorithms. However, HFTs regularly dismiss this notion as more of a fabrication. Largely due to data limitations, it has been difficult for the literature to quantify the impact of high-frequency trading and thus confirm either of the two opposing perspectives. The estimation of Aquilina et al. (2020) suggests that the reality lies somewhere between, with the liquidity costs associated to arbitrage opportunity races being almost negligible for retail investors, but not for institutional traders (more on this in Section 2.1). However, this raises the overarching question of whether HFTs improve market quality overall, and if not, whether market interventions exist that incentivize HFTs to exhibit market quality-improving behavior.¹ The following section discusses several potential approaches, while the remainder of this section outlines additional public concerns that should be accounted for in either approach.

Automated trading at high frequency has arguably been around since the introduction of fully electronic trading on NASDAQ in 1983. Decades later, in June 2009, after automated trading had captured more and more market share, the *Wall Street Journal* was among the first to bring the issue to public attention. The related article Rogow (2009) discusses the growing popularity of algorithmic trading and claims that trading volume generated by HFTs had more than doubled from nearly a quarter of daily trading volume in 2004 to nearly two-thirds in 2009. Although retail trading has risen over the past decade, due to FinTech innovations² such as "Robinhood" and "smartphone trading" (Welch, 2020; Kalda et al., 2021), and although this trend has been exacerbated by the COVID-19 lockdowns (Ozik et al., 2020), today's literature agrees that HFTs are still

¹Market quality is defined such that an improvement leads to generally desirable market conditions for traditional investors.

²FinTech is a new financial industry that seeks to develop technologies that make financial markets more accessible to the general public, see https://en.wikipedia.org/wiki/Financial_technology.

responsible for about 50% of daily trading volume.

Aside from making headlines due to their market presence, HFTs have been ubiquitous in the media after two short-lived financial market crashes in 2010 and 2012. On May 6, 2010, the well-known "Flash Crash" transpired, during which a large sell order allegedly triggered a chain reaction of algorithms causing a downward spiral of prices. Within 20 minutes the prices of more than 300 financial instruments fell by more than 60%, just to recover quickly towards the end of the very same trading day. This raised the question of whether algorithms sacrifice risk assessment and sophisticated economic decision making in exchange for speed advantages. While the role of HFTs during the Flash Crash is not obvious, the August 2012 incident clearly suggests a lack of risk assessment and sophistication. According to an article by Popper (2012) published by *The New York Times*, the high frequency trading company *Knight Capital Group* (KCG) made a loss in the amount of approximately \$440 million within the first 45 minutes of trading on August 1, 2012. A software error that caused the company's algorithm to buy overpriced assets caused this loss. The incident eventually resulted in a 70% drop of KCG's share price, which ultimately led to the company's acquisition by the high-frequency trading firm *Getco LLC* later the same year. What happened had stock market-wide repercussions, leading, e. g., to *Wizzard Software Corporation's* stock price rising by more than 400%, from \$3.50 to \$14.76, for no underlying reason in the fundamental value (Valetkevitch and Mikolajczak, 2012). While one would assume that safety measures are in place to prevent errors that potentially threaten a company's survival, such measures would slow down the algorithms' operation time and thus might be knowingly sacrificed by HFTs. The public attention generated by these events, as well as the subsequent legal reappraisal, led to the associated market outcomes being particularly well documented, making both events very interesting for the parameter tuning of the new framework, as discussed in Section 4.

It is a widely raised concern that HFTs act as liquidity³ taker in stressful market conditions, such as during these two crashes. The liquidity provided by HFTs is therefore often referred to as "phantom liquidity".⁴ An article by Zuckerman et al. (2018) published by *The Wall Street Journal* shares this assessment and further argues, that automated

³"Liquidity refers to the efficiency or ease with which an asset or security can be converted into ready cash without affecting its market price", according to <https://www.investopedia.com/terms/l/liquidity.asp>.

⁴The empirical facts on liquidity withdrawal by HFTs under stressful market conditions are discussed in Section 2 and 4.

trading generally may not trigger a decline in stock prices, but potentially amplifies a decline through liquidity withdrawal.

Most recently, HFTs have even been the lead story in *The Economist*, which covered the topic under the title "March of the machines" in its October 5, 2019 issue. Thereby, *The Economist* takes on opposing perspectives with its editorials eco (2019a) and eco (2019b). The former article points out that the last decade's technological advances in artificial intelligence and machine learning could exacerbate the competitive advantages of HFTs, raising concerns about wealth distribution. The latter article, eco (2019b), in contrast, argues that today's greater data source diversity, time horizons, and algorithmic strategies potentially helps curb bearish market moves. Not only because of an improved price discovery, but also because machines do not panic.

The latter contradicts what has been stated so far, which highlights a point also made by Aquilina et al. (2016), namely, that the question of the impact of high-frequency trading on market quality has not yet been conclusively resolved. Part of the reason is that data is often not made available. Moreover, specific observations are not possible in reality, e.g., it is not possible to suspend HFTs from all exchanges in the world in order to identify their impact on market quality. Similarly, it is not feasible to identify appropriate interventions by running market wide trials. Thus, although, as discussed in Section 2, individual exchanges can be observed that have adopted new market designs with the goal of improving market quality, it is ambiguous what impact a market-wide adoption of these designs would have. Consequently, regulators face the difficult challenge of identifying appropriate market interventions without having the relevant data to guide them. This is where the joint project with Peter Cramton, Daniel Friedman, Kristian Lopez Vargas, and Axel Ockenfels as well as the present complementary work step in. The new framework for real-time trading, developed by the LEEPs lab, builds an artificial financial market capable of testing various interventions and novel market designs in a controlled laboratory environment. In the laboratory, it is possible to investigate the consequences if the financial market were to switch market designs overnight or if a global intervention were to be introduced. This allows the pursuit of the following concrete research questions: Could a change in market design (i) stop the arms race for speed; (ii) create incentives for HFTs to improve market quality, both under normal and stressful conditions; (iii) be beneficial for ordinary investors? It is important to note that the

potential guidance for policymakers stems not only from the experiment, but from the combination with simulations, parameters based on empirical observations, and the insights of professional financial market participants. Indeed, the results of the experiment alone are relevant to reality only to the extent that the new framework is capable of capturing the financial market mechanisms relevant to high-frequency trading.

In pursuit of these research questions, the remainder of the paper is organized as follows. Section 2 outlines previous work in terms of the empirical facts, the theoretical foundation, and previous laboratory experiments. Section 3 introduces the new framework for real-time trading as developed by the LEEPS lab at the USCS and discusses its novelties. In the course of that section, the instructions, which were developed in close cooperation of all collaborators with the objective to make the complex rules of financial markets comprehensible for non-experts, will also be presented. The main contribution of this work starts with the elaboration of different alternatives to induce stressful market conditions in Section 4. Section 5 contributes with an evaluation of the first pilots and a description of the improvements undertaken, followed by the conclusion in Section 6.

2 Previous Work

In a broader context, the relevant market design literature includes Roth and Xing (1994, 1997), who study the timing of transactions, introducing the discussion of serial versus batch processing; Roth and Ockenfels (2002) as well as Ariely et al. (2005), who introduce the idea of bid sniping, based on observable behavior in online auctions; and Du and Zhu (2017), who study optimal frequencies in double auctions. Beyond that, the experimental financial market literature is naturally most relevant, in particular studies dealing with double auctions. This includes early studies in which subjects could still place orders orally or by hand signals (Friedman et al., 1984; Plott and Sunder, 1988), but also later studies, featuring computer-based auctions (Smith et al., 1988; Friedman, 1993). In this respect, the experimental literature developed in a similar fashion as the market itself, which experienced the introduction of fully electronic trading by NASDAQ at the same time in 1983, as previously described. Subsequently, however, the developments drifted apart, such that the highly visible experimental financial market literature still tends to focus on trading carried out by humans (Brunnermeier and Morgan, 2010; Moinas and

Pouget, 2013), while the market has experienced the outlined advances of algorithmic trading. The literature that investigates algorithmic trading, which is most closely related to the the new framework, includes Aldrich and Vargas (2019) and Aldrich et al. (2020), both studies are discussed in more detail in Section 2.3.

The predominant market design of the financial market is briefly outlined first, before discussing the empirical facts, the theoretical foundations, and the experimental literature in a more narrow context.⁵ The predominant market design features a *limited order book* and matches incoming orders based on a *double auction in continuous time*.⁶ These features function as follows. Limit orders are messages directed to the market composed of four objects: (i) *direction*: expresses the desire to buy (called bid) or sell (called ask); (ii) *quantity*: specifies the quantity of shares to be traded; (iii) *limit price*: defines the highest acceptable price for a bid or the lowest acceptable price for an ask; (iv) *expiration date*: defines the time when the order is to be deleted from the *order book*. The latter, the order book, is a list of all outstanding limit orders, sorted by price and time of arrival.

Consider the exemplary order book of a fictitious asset in Figure 1: For simplicity, assume without loss of generality, that the quantity of each order is normalized to one and that the expiration date is set to be in the future (e.g., at the end of the trading day). Orders are sorted first by price, which is illustrated by a smaller distance to the opposing market side. Secondly, in the case of a price tie, the orders are sorted by the time of arrival. From the latter, it follows that the ask in *State 1* line (2) must have arrived at the market before the ask in line (1). The ask marked in blue in line (3) of *State 1* specifies the lowest price for which the asset can be bought at that moment, this order is called **best ask (BA)**. Similarly, the order that defines the highest price for which the asset can be sold at that moment is called **best bid (BB)**, marked in orange. The market design, in this example CDA, defines how incoming messages are processed: Suppose two new bids at \$105 and \$98 arrive at the order book of *State 1*. The first bid immediately triggers a transaction with the **BA** at \$102, causing the ask in line (2), at \$103, to become the new **BA**. The second bid does not trigger a transaction, nor does it change the **BB**, it enters the order book after the outstanding bids. The resulting order book is shown in *State 2*, with the

⁵Over the past decade, several exchanges have implemented new modifications of this market designs based on so-called "speed bumps," which will be discussed in Section 2.2.

⁶For the remainder of this paper, the market design will be referred to as a *continuous double auction* (CDA).

Figure 1: Exemplary Order Book CDA

	Price	Direction		Price	Direction		Price	Direction
1.	\$103	Ask	1.	\$103	Ask	1.	\$103	Ask
2.	\$103	Ask	2.	\$103	Ask	2.	\$103	Ask
3.	\$102	Ask	3.	\$98	Bid	3.	\$101	Bid
4.	\$98	Bid	4.	\$98	Bid	4.	\$98	Bid
5.	\$97	Bid	5.	\$97	Bid	5.	\$98	Bid
						6.	\$97	Bid

State 1 State 2 State 3

second bid at \$98 to be found in line (4). Suppose that yet another bid arrives in *State 2*, this time at \$101. The bid is below the **BA**, such that there is no immediate transaction, but it is above the **BB**, making it the new **BB**, as depicted in *State 3*. The CDA rules apply analogously to incoming asks. It follows that the **BA** is always higher than the **BB** by design, the difference being called *market spread*. The market spread is a measure for the cost of liquidity. To illustrate, note that the midpoint between **BB** and **BA** represents the fundamental value of the asset, thus in *State 3* the fundamental value would be \$102. Assuming an investor wants to own the asset, he or she may place a limit order at or above the **BA**, which will execute immediately at \$103. Note that the ask in line (1) of *State 3* becomes the new **BA**, also at \$103, keeping the fundamental value constant. In sum, the investor buys the asset at \$103 even though it is worth only \$102, the desire to own the asset consequently induces liquidity cost of \$1, or, equal to one-half of the market spread. Note that an investor wishing to own or owe the asset could also place a so-called *market order*, which is designed to instantly transact with the **BB** or **BA**.

2.1 Empirical Foundation

The highly visible work of Budish et al. (2015) identifies persistent mechanical arbitrage opportunities in micro-structure trading data, sparking immense academic interest in high-frequency generated trading and HFTs. More concretely, the authors analyze the trading data of the SPDR S&P500 Trust (SPY), traded on the *New York Stock Exchange* (NYSE), and the S&P500 E-mini Futures (ES), traded on the *Chicago Mercantile Exchange* (CME), between 2005 and 2011. Both are financial instruments tracking the Standard & Poor's 500 (S&P500)⁷ stock market index, such that a near perfect correlation follows by design.

⁷The S&P500 is a stock market index composed of the 500 largest companies listed on U.S. stock exchanges. Here is a list of the associated companies

For a time horizon of one minute, the correlation is indeed nearly perfect, as shown in the left panel of Figure 2.⁸ Panel B on the right zooms in, depicting the price movements over a time horizon of mere 250 milliseconds. Over this time horizon, the near-perfect correlation appears to break down, which the authors indeed prove empirically. Combined with the price-time prior in continuous time under CDA, this creates mechanical arbitrage opportunities, where the price movement of one instrument can be understood as a near-perfect signal of the immediate upcoming price changes of the other instrument.

Figure 2: Micro-Structure Price Movements SPY and ES

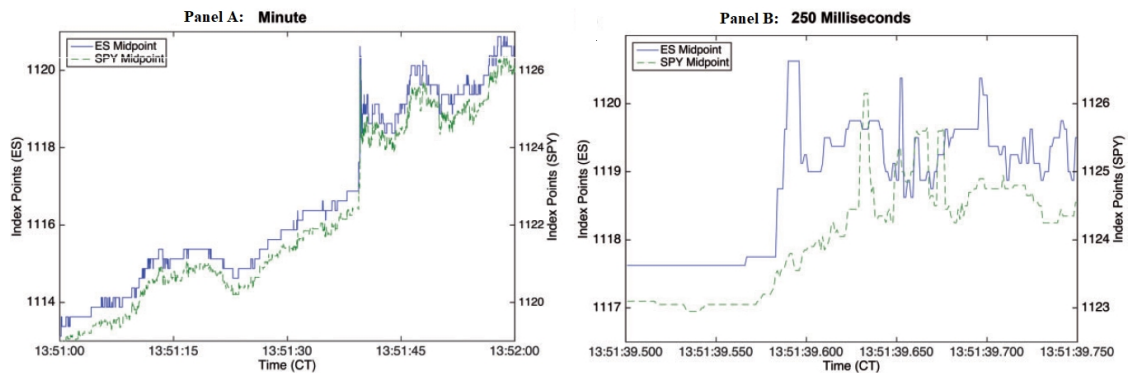


Figure (2) is taken from Budish et al. (2015). It shows the price movement of SPY and ES for different time horizons on August 19, 2011. The blue lines depict the ES midpoint, while the green depict the SPY midpoint. Panel A on the left hand side shows the data over the time horizon of one Minute, Panel B shows the data over the course of 250 milliseconds.

Consider the following example to illustrate the mechanical arbitrage opportunities. In Panel B just before the 13:51:39.600 time stamp, the ES midpoint jumps up. The SPY midpoint lags behind by about 30 milliseconds reaching a corresponding height at about 13:51:39:630. A HFT located at the NYSE observing the ES jump wishes to buy SPY at the stale price⁹, only to sell it again after 13:51:39:630. HFTs who actively attempt to pick off outstanding orders before their instigators are able to update them to the latest price information, are referred to as *Snipers*. Note that the number of stale quotas at the BA price is finite, if enough volume is generated on the buy side, it will lead to a new BA at a higher price. This subsequently increases the midpoint, causing a reduction or elimination of the arbitrage opportunity. Consequently, as previously noted, only the

https://en.wikipedia.org/wiki/List_of_S%26P_500_companies.

⁸Note that the different levels of SPY and ES are due to differences between NYSE and CME, which are not relevant to the present discussion.

⁹A price is called stale, if it does not reflect the latest information, according to, e.g., <https://www.nasdaq.com/glossary/s/stale-price>.

first to react is able to profit from the arbitrage opportunities. Alongside Snipers, the market participants whose orders Sniper target, also frequently participate in the races for arbitrage opportunities. These market participants are called Liquidity Providers (LP) and represent the second role in which HFTs may enter the market. LPs provide liquidity in form of limit orders, trying to adjust their outstanding orders according to new price information before Snipers can pick them off. If a Sniper is successful, this results in a loss for the LP equal to the Sniper's profit, which is explained in more detail in the section 2.2.1. Analyzing the trading data over the course of multiple years, between 2005 and 2011, Budish et al. (2015) find that while the time span of arbitrage opportunities shrinks from a median of 97 milliseconds in 2005 to a median of 7 milliseconds in 2011, the amount of arbitrage opportunities remains constant. Competition thus appears to be limited to speed increases, while leaving the arbitrage opportunities amounting to approximately \$75 million per year, untouched. The authors conclude, therefore, that the predominant market design is flawed in that it permits constant arbitrage opportunities while setting the stage for a potentially socially wasteful arms race for speed.

Admittedly, \$75 million seems almost negligible in the context of financial markets, but recall that these are only the arbitrage opportunities between ES and SPY, i.e., only between two financial instruments. Globally, there are manifold more financial instruments with similar correlations and beyond that, publicly observable information like annual reports or even social media posts might even create arbitrage opportunities. Think of politicians or celebrities sending stock prices up or down by posting social media post expressing a positive or negative sentiment regarding future market performance. Among the most bizarre recent examples might be a tweet by Elon R. Musk on Jan 26, 2021, which potentially sent Game Stop's (GME) stock price skyrocketing even after an already intense frenzy. While such tweets represent symmetric information, they make outstanding orders stale in the instant they go live, incentivizing HFTs to anticipate the sentiment as quickly as possible and react accordingly. The countless potential causes of mechanical arbitrage opportunities make it difficult for the literature to provide an estimate of the global annual arbitrage prize pool. Existing approaches to quantification the prize pool include Aquilina et al. (2016) who estimate the prize pool across UK dark pools to be GBP4.2 million per year, as well as Ding et al. (2014) and Wah (2016) who estimate latency arbitrage opportunities based on frequencies and price differences,

leading to an estimate of \$3 up to \$4 billion annually across all assets listed on the S&P500. A recent study by Aquilina et al. (2020) attempts to estimate the global arbitrage opportunities based on a novel data set called "message data". The data includes records of trade and cancellation attempts in addition to ordinary transaction order book data, allowing the authors to identify races directly by their entrants for the first time. The data is retrieved from the London Stock Exchange (LSE) and includes all stocks listed in the FTSE 350 index¹⁰ over the course of two months in 2015. The main results imply that the annual global arbitrage price pool on equity markets amounts to approximately \$5 billion. Thereby, an average race provides profits in the amount of just GBP2, but races happen frequently, on average once per minute per stock listed in the *FTSE 100*, accounting for 22% of daily volume. Additional results show that winners beat losers by mere 5-10 microseconds, and in 70% of the cases compete for strictly positive returns. This is consistent with previous literature concerned with HFTs and their profits, showing that marginally faster HFTs have significantly higher profits (Brogaard et al., 2014; Baron et al., 2019). Additionally, Aquilina et al. (2020) derive that the top 6 HFTs win 82% of races, but also loose 87% of them, and that 90% of the races are won by a snipe rather than a cancellation. These results are in close agreement with the theoretical prediction of Budish et al. (2015), as will be discussed in the next section. Not entirely in agreement with the theory is the empirical finding that while HFTs have a tendency to act as either Sniper or LP, they do not confine themselves exclusively to either role: 2/6 HFTs are responsible for only 28% of wins, but account for 61% of successful cancellations, while displaying a ratio of liquidity-taking to liquidity-providing behavior in races of 2:3. The remaining 4/6 HFTs are responsible for 54% of wins, only 21% of cancellations, and display a ratio of 5:1. Lastly, the authors estimate that a market design capable of eliminating arbitrage opportunities could lead to a 17% reduction in liquidity costs, which is of considerable significance, at least for institutional investors.

In light of the pending questions raised in the previous section, the adverse effects of high-frequency trading on liquidity costs certainly argues that HFTs have a negative impact on market quality. However, there exist many more effects to consider, with almost an entire field of research devoted to resolving the question regarding the overall

¹⁰Similar to the S&P500, the FTSE 350 Index is a stock market index consisting of the 350 largest companies listed on the LSE. It is a combination of the [FTSE 100 Index](#) and the [FTSE 250 Index](#).

impact of HFTs. This literature distinguishes between liquidity-providing and liquidity-taking activities of HFTs. The liquidity providing behavior of HFTs is mostly associated with an improvement of market quality (Hagströmer and Nordén, 2013; Jovanovic and Menkveld, 2016; Menkveld and Zoican, 2017). The results for liquidity-taking behavior are more ambiguous, as although these activities are associated with short-term informed price impact and thus improved price discovery, they are also associated with increased negative prices for slow traders and increased short-term volatility (Brogaard et al., 2014; Zhang and Riordan, 2011). On the resulting question of the net effect of HFTs, most of the literature finds an overall positive effect on market quality (Hasbrouck and Saar, 2013; Breckenfelder, 2019; Brogaard and Garriott, 2019). In particular liquidity is often found to be improved (Hendershott et al., 2011; Menkveld, 2013). At this point, however, it should be clearly stressed that these results do not imply that latency arbitrage and the associated arms race for speed are necessary for HFTs to exhibit market quality-improving behavior. In this respect, a highly original empirical study by Shkilko and Sokolov (2016) shows that a disruption of the microwave network between the CME and the NYSE through rain, causes HFTs to slow down, resulting in improved liquidity. Breckenfelder (2019) further contributes to this debate by studying not only the differences between liquidity-provision and liquidity-withdrawal, but also disentangling the effect of increasing high-frequency trading volume from increasing competitive intensity between HFTs. The main results suggest that increased competitive intensity raises the ratio of liquidity-taking to liquidity-providing behavior, which leads to a decrease in market quality. In contrast, an increase in volume generated by HFTs in isolation is found to have either a positive or neutral impact on market quality. Consequently, volume generated by HFTs may be desirable for market quality, while competition between HFTs may not.

Finally, it should be emphasized that although HFTs may virtually see the future, have supremely sophisticated algorithms, and trade near the speed of light, they are not entirely devoid of risk. During their pursuit of profiting from trading, HFTs build up an average inventory equal to 25%-35% of their daily trading volume, which exposes them to adverse price movements.¹¹ In this context, HFTs not only actively seek to keep their inventory as small as possible (Anand and Venkataraman, 2013; Korajczyk and Murphy, 2019), but also close the trading day with an inventory of zero (or close to

¹¹For a theoretical discussion, see Herrmann et al. (2020) or Guilbaud and Pham (2013).

zero) in 57% of the cases, to avoid the risk of adverse overnight price movements. Beyond that, both studies are relevant to the present paper in that they compare the behavior of HFTs with market participants obligated to provide liquidity, also under stressful market conditions. Stress conditions might expose HFTs to greater risks, e.g., through dry spells, unbalanced order flows, or high volatility. Under all stressful market conditions, except high volatility, HFTs tend to act as net liquidity takers, whereas the same is not true for obligated liquidity providers. The authors find that in general, obligated liquidity providers tend to earn less on days when HFTs are not providing liquidity, suggesting that HFTs provide liquidity only when it is profitable. In sum, this may support the public concern expressed in Section 1, namely, that HFTs provide phantom liquidity, leaving obligated liquidity providers as the only reliable source of liquidity in stress situations. The absence of an effect by increased volatility could be due to the notion that while volatility increases inventory risk, it compensates HFTs by increasing the likelihood of arbitrage opportunities. The latter is empirically supported by (Aquilina et al., 2020), who find that arbitrage opportunities are strongly correlated with volume and volatility.

The following section presents the experiments' theoretical foundation. Simplifying assumptions are made with respect to the empirical facts presented here, some of which will be relaxed in later sections for the sake of real-world proximity.

2.2 Theoretical Foundation

Before discussing the prevailing financial market design, as well as alternative designs, it is worthwhile to mention a simplifying assumption that carries throughout this section: Trading is concentrated on a single exchange. This assumption is largely based on the Securities Exchange Commission's (SEC) rules on *Unlisted Trading Privileges* and the *Regulation of the National Market System*.¹² The former essentially ensures that an asset can be purchased on one exchange and sold on another, if it is listed on both exchanges. The second ensures that investors trade at the statewide best bid or best ask, while search and execution on the various exchanges proceeds friction-less. For U.S. investors, therefore, the limit order book is effectively the sum of all limit order books on U.S. exchanges, such that the simplifying assumption poses little trouble. For

¹²See <https://www.sec.gov/rules/final/34-43217.htm> and <https://www.sec.gov/rules/final/34-51808.pdf>, respectively.

HFTs, on the other hand, the assumption would imply that all HFTs are on the same exchange. Although this is a stronger abstraction, it still holds that a price jump on any exchange, on which HFTs are not resident, could create a signal, effectively similar to the previous example. In addition, it holds that a price jump is likely to happen earlier on one exchange than the other (Hagströmer and Menkveld, 2019), such that signals could still create arbitrage opportunities.

An additional theoretical consideration with respect to exchanges, that holds for the prevailing market design and underlines the importance of the present study, stems from Budish et al. (2019). The authors model the incentive system of exchanges under the CDA market design and derive that the exchanges will not have the desire to adopt a new market design, i.e., the market will not fix itself in the light of the CDA flaws. The main reason for this is that exchanges profit from the very flaws that alternative market designs seek to eliminate. Namely, the high trading volume, the willingness to buy proprietary data feeds, and the willingness to settle nearby of HFTs are all highly profitable for exchanges. But to attract and retain HFTs, exchanges need to generate volume from traditional, slow investors, who in turn feel disadvantaged by the presence of HFTs and are crowded out of the exchanges. The associated dilemma of introducing a new market design poses a repeated prisoner's dilemma (O'Hara, 2015; Budish et al., 2019). If only one exchange "defects", i.e., adopts a new market design that eliminates arbitrage opportunities, it experiences a rise in volume that leads to substantial profits. If all other exchanges defect subsequently, the volume-increasing effect vanishes, leaving everyone with less profit than under CDA. Since the game can be thought of as infinitely repeatable, there exists a cooperation equilibrium in which the financial market currently finds itself. It also follows that regulators might not necessarily have to mandate a particular market design, but instead could incentivize a first mover sufficiently to bring about a defection equilibrium. That said, the overarching question remains, what market design should regulators be encouraging in the first place?

The remainder of this section provides theoretical foundations of potentially suitable market designs, beginning with the *BCS* model introduced by Budish et al. (2015).

2.2.1 Continuous Double Auction

Budish et al. (2015) build the BCS model upon the following simplifying assumptions beyond those mentioned above. A single artificial asset, a , is traded and the asset is perfectly correlated with a publicly observable signal, s . The signal can be understood as any information causing latency arbitrage opportunities, with the peculiarity that it directly reflects the change in a 's fundamental value. Moreover, it is assumed that a can be liquidated costlessly at its fundamental value at any time. Given the inventory risk outlined previously, this assumption poses a strong abstraction; relaxing it will be one among the most important distinctive features of the new framework for real-time trading. It follows that the BCS model creates a best-case scenario for liquidity provision and price discovery, such that the following theoretical predictions should be interpreted as boundary solutions.

The BCS model distinguishes between two distinct, risk-neutral market entrants, *HFTs* and *Investors*. The latter are market participants who withdraw liquidity based on an intrinsic motivation to either own or owe the asset.¹³ Investors arrive at the market according to a Poisson process with arrival rate λ_{invest} , the arrival at either market side is equally likely. Investors are assumed to be impatient and observe changes in s strictly later than HFTs. In equilibrium, Investors wish to trade immediately and hence place market orders rather than limit orders. In that, Investors do not behave strategically. HFTs, on the other hand, use limit orders to profit from trading, they do not derive utility from owning or owing the asset. HFTs either provide liquidity in the role of LP or withdraw liquidity in the role of Sniper. LPs determine the market spread, ms , at any point in time by constantly holding a limit bid at $s - \frac{ms}{2}$ and a limit ask at $s + \frac{ms}{2}$. If either order is executed, the LP immediately replaces it. If it is executed against an Investor's order, the LP immediately books a profit equal to one-half of the market spread. Note that this coincides precisely with the liquidity cost of the last investor from the example of Figure 1. If, however, the order of an LP executes against a Sniper's order, the LP might book a loss. To illustrate: Suppose the fundamental value changes to $s_{t'} > s_t$ at time $t' > t$, if the LP fails to cancel its outstanding ask in time, the stale quote will result in a payment of: $s_t + \frac{ms}{2} - s_{t'}$. This payment results in a loss for the LP and a profit for Sniper in the same amount, if the price

¹³These Investors are similar to the *liquidity takers* originally introduced by Glosten and Milgrom (1985) or Kyle (1985), with the distinction that the BCS model abstracts from asymmetric information.

change is sufficiently large: $s_{t'} - s_t > \frac{ms}{2}$. It follows that Snipers engage in racing only if the price jump is sufficiently large. In the case of a price drop, the mechanism works analogously. If, instead, the LP succeeds in canceling its outstanding orders, no trade happens and hence the LP books no loss. The BCS model assumes price jumps in s to follow a compound Poisson jump process, with price changes occurring at random times according to the arrival rate λ_{jump} , while their magnitude is drawn from the distribution F_{jump} . The distribution is assumed to have finite support, to be symmetric, and to have mean zero.

The model assumes that HFTs can either be slow or fast, with no difference in latency within these categories but between them. HFTs decide before the trading period whether to take an investment in speed in the amount of c_{speed} or not. The investment can be understood as the sum of all costs associated with the competitive speed advantages outlined in Section 1. Fast HFTs observe the change in s strictly before slow ones. Assuming that all fast HFTs have the same latency leads to all their messages arriving at the exchange at the exact same time. The resulting tie in arrival time is settled by random selection. Although this may seem arbitrary, it can be understood as marginal noise in message delivery that determines the winner of the race, which is consistent with the empirical results in Section 2.1: The six fastest HFTs win but also lose most of the races.

The equilibrium of the BCS model is based on two conditions. First, HFTs need to be indifferent between both roles, i.e., the expected profit of a Sniper and LP must be equal; and second, HFTs need to generate zero-profits in expectation, investing all their potential profits in speed. With these conditions and given λ_{invest} , λ_{jump} , and F_{jump} , the equilibrium market spread ms^* is uniquely determined, independent of both the number of HFTs, N , in the market and c_{speed} . Note, that while the market spread may not depend on N and c_{speed} , it increases with increasing arbitrage opportunities. According to the previous discussion this may suggest that if λ_{jump} and F_{jump} are such that the market is particularly volatile, the cost of liquidity is particularly high. The equilibrium spread, together with c_{speed} , determines the equilibrium entry N^* . HFTs sort themselves endogenously into a single LP and $N^* - 1$ Sniper. Empirically, this is supported by the observation that the ratio of sniping to liquidity providing behavior increases in competition and that 70% of competing HFTs act on the same market side (Breckenfelder, 2019). Moreover, it

follows from random selection in case of a tie in arrival times, and with strictly more than two HFTs in the market, that races are more often won by liquidity taking than liquidity providing activities. Indeed, with six HFTs in the market, Aquilina et al. (2020) find that 90% of races are won by liquidity taking. Notably, a slow LP would always be sniped, while a slow Snipers would never be successful, both are unprofitable strategies, such that only fast HFTs enter the market in equilibrium, leading to an aggregate speed investment of $N^* \cdot c_{speed}$. In equilibrium, the LP provides a single bid at BB and a single ask at BA. The economic intuitions behind the results are as follows. First, competition, which increases the cost of speed, has no positive impact on market quality, it neither increases liquidity nor reduces the cost of liquidity. Second, the total return of HFTs, which is equal to the total investment in speed, is borne entirely by Investors. Budish et al. (2015) emphasize that it is not HFTs per se, but sniping activities that have a negative impact on market quality. The underlying idea is that the LP compensates for a higher risk of being sniped by charging Investors more for liquidity, and that the LP's risk increases proportionally with the liquidity provided. As shown below, an arguably small change in the market design might resolve these deficiencies in such a way, that Investors are better off due to an improved market quality.


2.2.2 Frequent Batch Auction

The *Frequent Batch Auction* (FBA) features some similarities to to the CDA market design, the key difference, however, is that the trading day under FBA is divided into discrete time intervals of equal length τ , called *batches*. At the end of each batch, a *uniform price auction* is held. The auction takes into account all market and limit orders placed during the preceding batch, as well as all orders from the order book. The order book carries over all outstanding limit orders from previous batches into the next auction. Limit orders work exactly as before, while market orders need to be slightly adjusted. Namely, market orders do not transact immediately, but lead to a transaction in the next batch auction upon their arrival. Therefore, market orders are still a viable tool for impatient investors, although they now include a positive delay of at most τ . This delay, as well as the associated costs it imposes, is one of the main drawbacks of the FBA market design as modeled by Budish et al. (2015). The uniform market clearing price, if it exists, is determined by the intersection of the aggregate demand and supply functions

at the end of each batch. If the market clearing quantity is not unique and the time priority is not sufficient to resolve the tie, a random selection is made. If the market clearing price is not unique, the midpoint of the interval of market clearing prices is selected. If supply and demand do not cross, no trade occurs and all orders are carried over to the next batch. The market clearing price and quantity are announced after each auction, whereas the activities during a batch are not publicly observable.

The FBA design reduces the effectiveness of speed advantages quite mechanically by substantially reducing the number of arbitrage opportunities. Under FDA, a change in s creates an arbitrage opportunity only if it takes place near the end of a batch, such that a fast HFT is still able to react before the batch auction, while a slow HFT is unable to. This is essentially a derivative of the market design feature that any message that arrives at the exchange before the batch auction is indeed processed. It holds that if a Sniper is able to transmit an attempt to snipe a stale quota to the exchange in time, a LP is able to transmit a cancellation. As a result, only slow LPs are vulnerable to sniping. For illustration, see Figure 3, which shows a sample batch between t_0 and t_1 of length τ . Suppose a fast HFT has latency δ_{fast} , while a slow HFT operates at δ_{slow} , with $\delta_{slow} > \delta_{fast}$. The change in s must occur between $t_1 - \delta_{slow}$ and $t_1 - \delta_{fast}$ to create an arbitrage opportunity. The associated time period is $(\tau_1 - \delta_1) - (\tau_1 - \delta_{slow}) = \delta_{slow} - \delta_{fast}$, which corresponds to the speed difference between a fast and a slow HFT. This time period is likely very small. It follows that under FBA, speed advantages are significant only for the fraction $\frac{\delta}{\tau}$ of the trading day, which, according to a rough estimate by the authors, is in the range $[0.01\%, 0.1\%]$, while, recall, that speed under CDA is meaningful for 100% of the trading day.

Figure 3: Time-Span of Latency Arbitrage under FBA



This illustration follows Budish et al. (2015) most closely.

Not only does FBA mechanically reduce arbitrage opportunities, it also creates new incentives with desirable consequences in this direction. While an opportunity to snipe

outdated quotas of slow LPs may still exist during the time interval of Figure 3, note that the attempts of all fast Snipers arrive in the same batch and are thus all processed. Since only the Sniper with the best price executes with the stale quota, a Bertrand price competition unfolds. This ultimately leads to the stale quota being executed at the updated fundamental value, which means that the slow LP does not book a loss, even if it is sniped after a sufficiently large price jump.

In equilibrium, all HFTs endogenously choose to enter the market as LP, none enter as Sniper. If τ is sufficiently large relative to δ , there is no investment in speed. If multiple LPs enter the market, they compete in price, leading to $ms^* = 0$. It follows that HFTs make zero profits just as before. Investors are impatient and have no liquidity costs; they trade in the batch in which they arrive, at the fundamental value. HFTs provide at least enough liquidity to meet the demand in each batch. Since they do not observe demand before the auction, it follows that the liquidity has to be improved compared to CDA. Besides the improved liquidity, the economic interpretation of the equilibrium is that cost for liquidity equal zero and that investors have a welfare gain of $c_{speed} \cdot N^*$ compared to CDA. A market spread in the amount of $ms^* = 0$ has a further implication, namely, that every trade is executed at the fundamental value, potentially allowing incoming investors to execute against the orders of other investors from the opposing market side. The market design discussed in the following section has a similar feature by design, which will require an adjustment of the experiment software compared to FBA and CDA.

2.2.3 Investors' Exchange

The third market design to be considered in the new framework for real-time trading is the so-called *IEX* market design. The name is derived from the *Investors' Exchange LLC*, which was the first exchange to introduce the respective market rules.¹⁴ The SEC approved this market design on June 17, 2016, allowing the Investors' Exchange to operate as a public exchange for the first time. The name already gives away the objective of IEX, which is to improve the market quality for ordinary investors. The innovation of IEX stems from so-called *speed bumps*, which are supposed to delay incoming orders by 350 microseconds, with the purpose of protecting slow investors from Snipers. The

¹⁴See <https://iextrading.com>.

delay might, e.g., be achieved by routing each incoming message through a 38-mile long cable, as is done by the Investors' Exchange. Originally, the IEX was intended to route messages from all market participants through this cable (symmetric). However, since the approval by the SEC, there are various versions that only delay messages from specific market participants (asymmetric). Apart from the Investors' Exchange, the *Aequitas NEO Exchange* and the *TSX Alpha Exchange* in Canada; the *Thomson Reuters* and the *NYSE* in the U.S.; and the *Eurex Exchange* in Germany have implemented speed bumps in some form.¹⁵ The delays vary from 350 microseconds to 9 milliseconds and from symmetric to asymmetric message delays.¹⁶ Although not true for some of these variants, the original idea of the IEX is consistent with the objective of the FBA, i.e., to improve the welfare of ordinary investors. This, combined with the fact that IEX has already been approved by the SEC, renders the market design a natural candidate for the comparison in the new real-time trading framework.

Aldrich and Friedman (2022) introduce a theoretical model of the original IEX market design, inspired by the BCS model. In this context, IEX is very closely related to the CDA market design and, unlike the FBA market design, does not mechanically reduce the number of arbitrage opportunities, but instead protects stale quotas from being sniped. Symmetric speed bumps alone are not sufficient to achieve this goal, they only work in combination with so-called *pegged orders*. These orders are updated according to new information by the exchange itself. Although there are three different types of pegged orders, the authors show that investors almost exclusively use *midpoint pegs*, which is why these are the focus of the following analysis. Pegged orders in general are a common order type on virtually all exchanges. The novelty of the IEX design is that the speed bump gives the exchange sufficient time to update pegged orders before sniping messages reach the market. An additional feature of pegged orders is that they are hidden, i.e., they are not visible to all market participants in the order book, unlike the "lit" limit orders discussed so far. Limit orders are beneficial to the market because they facilitate price discovery, which is not the case for hidden orders, leading to exchanges charging higher fees for hidden orders and giving them lower time priority. Accordingly, hidden orders have the implicit disadvantage of always being queued behind lit orders, resulting in non-

¹⁵Baldauf and Mollner (2020) provides a thorough overview of all exchanges that have implemented some form of speed bumps so far.

¹⁶Section 2.3 discusses an experimental study examining the efficiency of different types of delays.

trivial delay costs due to fixed tick sizes.

Midpoint pegs are limit orders pegged at the midpoint of BB and BA, to illustrate their functioning, consider the exemplary order book in Figure 4. The order book works essentially as in the previous example from Figure 1, except that it incorporates the hidden midpoint pegs marked in gray, as shown, e.g., in lines (3) and (4) of *State 1*. In the order book, midpoint pegs can originate only from one market side by design, as incoming midpoint pegs immediately transact against outstanding midpoint pegs from the opposing market side, if available. Suppose that in *State 1*, two investors arrive at the market. The first investor is more impatient and places a market order, while the second is more patient, placing a midpoint peg. The market order, which would normally be executed at the BA of line (2), is executed against the hidden ask at \$100, which is favorable for the incoming investor. The hidden peg of the second investor enters the order book since there are no additional outstanding orders at the fundamental value. This resulting order book is depicted in *State 2*, which notably would have been the same if both investors had placed a midpoint buy peg. Suppose that in *State 2* an investor arrives who places a lit ask at \$101, which becomes the new BA. The resulting order book is shown in *State 3*. Note in particular that the price of the outstanding midpoint peg has been updated to \$99.5. Without any outstanding hidden pegged orders, the order book would function exactly as under CDA.

Figure 4: Exemplary Order book IEX

	Price	Direction
1.	\$102	Ask
2.	\$102	Ask
3.	\$100	Ask
4.	\$98	Bid
5.	\$98	Bid

State 1

	Price	Direction
1.	\$102	Ask
2.	\$102	Ask
3.	\$100	Bid
4.	\$98	Bid
5.	\$98	Bid

State 2

	Price	Direction
1.	\$102	Ask
2.	\$102	Ask
3.	\$101	Ask
4.	\$99.5	Bid
5.	\$98	Bid
6.	\$98	Bid

State 3

In the model of Aldrich and Friedman (2022), investors decide endogenously whether to enter the market with pegged or lit orders. The investor flow and price jumps are exogenous events, as before in the BCS model. It is also assumed that market participants who provide liquidity do not purchase speed, that Snipers can immediately reverse their transaction, and that pegged orders are effectively protected from sniping. Under these assumptions and with the authors' baseline parameters, the model predicts that in the

equilibrium a majority of investors will choose pegged orders. However, the delay costs and higher fees ensure that at least some investors choose lit orders in equilibrium. It should be pointed out that the aforementioned feature of pegged orders, which allows investors to trade among themselves, reduces the profit opportunities of LPs, which needs to be incorporated into the experimental software. Compared to CDA the theoretical analysis predicts (i) a lower ratio of Snipers to LPs, (ii) transaction prices that deviate less from the fundamental value, and (iii) lower liquidity costs. Compared to the FBA, a lower liquidity at fundamental value is predicted. The economic interpretation is that IEX does indeed improve market quality, although not to the same extent as FBA. Empirically, there is suggestive evidence of a decrease in high frequency trading of about 20% under IEX compared to CDA. However, it is unclear whether this decline is due to a reduction in sniping activity, which would be beneficial for market quality, or a reduction in liquidity-providing activity. Indeed, other empirical studies find that speed bumps have a negative impact on liquidity (Chen et al., 2017) and no detectable impact on the market spread (Anderson et al., 2021). The IEX model further predicts that more volatility leads to fewer LPs, but not fewer Snipers, such that the ratio of sniping to liquidity-providing behavior could be particularly adverse under stressful market conditions. Moreover, higher speed costs are found to reduce the number of active HFTs in the market, but not to eliminate the arms race, similar to CDA. Brolley and Cimon (2020) agree and further argue that the IEX market design is an insufficient intervention to resolve the arms race. However, there is a lack of conclusive empirical evidence to support this, as IEX has not yet been adopted market-wide.

These three market designs are the first candidates to be compared in the new real-time trading framework. However, as the literature advances, new and exciting approaches for improving market quality should not be overlooked. These approaches include market designs such as the fully continuous exchange presented by Kyle and Lee (2017) or *Flow Trading* presented by Budish et al. (2020). Flow Trading is essentially an extension of the FBA market design that allows investors to express their desire to own or owe an asset in a sophisticated manner, rendering it a highly aspirational candidate for the comparison. Investors place their orders in the form of piece-wise linear, downward-sloping demand curves that are continuous in price and quantity.¹⁷ In addition, investors specify a trading

¹⁷Supply curves can be understood as demand curves with negative quantity.

rate and a maximum quantity to be achieved. The order is executed at a rate of zero if the price is above the upper bound set by the piece-wise linear demand function; it is executed at the specified maximum rate if the price is below the lower bound; and it is executed at a linearly interpolated rate, if the price is between these bounds. As under FBA, the market is cleared after each batch, with the intersection of the aggregate downward-sloping demand curves determining the uniform price. It is worth pointing out that Flow Trading not only counteracts the arms race for speed, but also addresses potential weaknesses of CDA already pointed out by Tong (2014) and Kyle (1985). The latter study suggests that the best strategy of institutional traders under CDA might be to split large orders into as small a pieces as possible in order to escape the negative price impact and not allow HFTs to recognize one's intention. Indeed, Van Kervel and Menkveld (2019) show that institutional investors are exposed to higher trading costs when confronted with HFTs. Through the execution rate, Flow Trading has a built-in feature that removes these obstacles.

As with FBA and IEX, the problem with Flow Trading is the absence of real-world observations, and for now, there can be none. The laboratory comparison could provide initial insights into what the impact of a market-wide adoption of either of these market designs might be. The next section discusses prior experimental financial market literature that analyzes the behavior of HFTs.

2.3 Experiments in High Frequency Trading

The complex coding involved might be one of the reasons behind the relatively small size of the financial market literature featuring experiments in which subjects take on the role of HFTs. Khapko and Zoican (2020) circumvent this issue by considering a rather abstract artificial financial market, in which trade happens in discrete time, with each round presenting exactly one arbitrage opportunity. Subjects compete against each other in the role Sniper and against an artificial LP. Before each round, subjects decide whether to invest in speed, which increases the probability of winning the race, or not. The study focuses on the efficiency of different speed bump variations in reducing speed investments. The following properties of speed bumps create the associated treatments: (i) deterministic & symmetric; (ii) deterministic & asymmetric; and (iii) probabilistic

& asymmetric.¹⁸ The authors find that the size of the speed bump is important, with a one standard deviation increase in delay, reducing speed investments by 8.33%. In contrast to previous studies, e.g., Aoyagi (2019), no difference between deterministic and probabilistic speed bumps could be identified. The main results are that treatment (ii) significantly reduces speed investment, while the same is not true for treatment (i). This is consistent with the IEX model by Aldrich and Friedman (2022), in that speed bumps alone are not sufficient for reducing investment in speed.

Aldrich et al. (2020) provide the foundation for the experimental software that Aldrich and Vargas (2019) employ and on which the new framework for real-time trading is built. The study designs an architecture for artificial financial markets that enables real-time bidding in continuous time at a speed of 10-20 milliseconds. This represents a significant reduction in the degree of abstraction compared to the previous financial market literature. One of the reasons are technical challenges, as widely used experimental software such as zTree (Fischbacher, 2007), or standard application of oTree (Chen et al., 2016), do not evaluate subjects' decisions until they have actively been submitted. Similarly, for the distribution of information regarding the decisions of other subjects. Consequently, previous trading on artificial financial markets mostly took place in discrete time only. The extension of oTree by the so-called *web-socket* links enables continuous communication between all subjects (Fette and Melnikov, 2011), since both sides of a link can initiate a message at any time. All subjects are linked to the experiment server, which additionally is linked to a so-called *exchange server*. The exchange server derives its name from the fact that, like a real financial market exchange, it maintains an order book and processes incoming orders according to the respective market design. For the sake of realism and to ensure that subjects receive updated information as fast as possible, the communication between the experimental server and the exchange server is carried out via the so-called OUCH protocol, which is also used by NASDAQ.¹⁹ The complete route of any order is as follows: Subjects submit their order to the experimental server via a sophisticated user interface. The experimental server translates the message to OUCH and forwards it to the exchange. The exchange server processes the message depending on the market

¹⁸Note, asymmetric in this section means in favor of LPs, as for example implemented by *TSX Alpha Exchange* located in Canada. Probabilistic means that not all messages are delayed, only randomly selected ones.

¹⁹Detailed information regarding OUCH is available at Nasdaq Trader (2021).

design, e.g., serially in case of CDA, and either makes an adjustment in the order book or executes the order, if applicable. The exchange sends the updated information back to the experiment server via OUCH in either case. The experiment server encodes the message before broadcasting it to all subjects simultaneously via the web-socket links. There is an alternative to this route, however, as subjects are not required to perform a conscious decision to place an order. The architecture includes a feature that allows subjects to switch on a trading algorithm, which performs the trading on their behalf. Each subject has its own trading algorithm, located on the experiment server, which, if turned on, responds to changes in the exogenous events of the BCS model, the fundamental value and the investor flow. Note that the exogenous events are also processed by the experimental server. As a result, an algorithm may technically have already made a trade on behalf of a subject prior to the subject observing the respective consequences on its user interface. In the laboratory, this delay is mostly below 100 milliseconds. Considering that a spontaneous blink of the eye is in the range of 200 milliseconds (Królak and Strumiłło, 2012), it follows that the information is available to subjects virtually instantaneous. This discussion may seem overly technical, but the consideration will become important in Section 5.

Aldrich and Vargas (2019) are the first to utilize this architecture in order to compare CDA and FBA in a laboratory setting. It is, therefore, the most closely related work to the joint project, and this paper. Technical differences include that Aldrich and Vargas (2019) are not using oTree as an experimental server. The assumption of immediate reversibility of HFTs' transactions at the fundamental value is implemented by subjects immediately generating payments as soon as one of their orders is executed, according to the profit scheme outlined earlier in this section. The strategy space of subjects includes the decision to enter the market as LP or Sniper, or to stay out of the market. Subjects can revise their decision at any time. If they enter the market, the decision space further includes the possibility to subscribe to a faster communication technology, which creates costs per second used. This faster communication technology reduces the time it takes for a message send by the trading algorithm to be processed by the exchange server. In addition, subjects acting as LP need to manually set the desired distance of their orders from the fundamental value. Note that the market spread is determined by the subject placing its bid and ask closest around the fundamental value. In the event of a price

jump, the LPs' algorithms seek to keep this distance to the fundamental value constant by deleting the outstanding orders and replacing them with updated ones according to the new fundamental value. Snipers' algorithms attempt to intervene in this process and exploit the stale quotas.

Compared to the set of baseline parameters, Aldrich and Vargas (2019) consider the following two sets of market conditions as treatments. First, a market with increased volatility, lower investor flow, and constant speed costs; Second, a market with increased volatility, higher investor flow, and double the speed costs. Intuitively, higher volatility exposes LPs more to sniping, while a thicker investor flow makes up for it. The main results include that compared to CDA, the authors observe under FBA: (i) lower Sniper to LP ratio; (ii) lower subscription to speed technology; (iii) lower liquidity costs; (iv) lower volatility in market spread and liquidity; and (v) improved liquidity. The results are generally highly statistically significant and consistent with the equilibrium predictions of Section 2.2, although the magnitudes are moderated. The treatments provide first insights regarding the public concerns raised in Section 1. The liquidity-improving effect of FBA is found to be highest for the most volatile market with the thinnest investor flow, suggesting that FBA might be a tool to resolve the phantom liquidity provided by HFTs. Moreover, the authors elaborate on the concern, that while HFTs may not cause crashes, they potentially amplify them under CDA. The authors derive that under CDA a shock causing an increase in the ratio of the number of price changes to the number of investors arrivals, leads to a short-term, significant increase in the fraction of Snipers to LPs and a weakly significant increase in the minimum market spread. The same does not hold for FBA, suggesting that FBA potentially attenuates the amplification of crashes by HFTs.

A result of the laboratory study, that was not accounted for by the BCS model, is LP-to-LP sniping. With baseline parameters, Aldrich and Vargas (2019) find that fast LPs account for roughly 60% of all snipes.²⁰ Suppose, for illustration, that strictly more than one LPs are active in the market, which would only be possible outside of the CDA equilibrium behavior, then the underlying mechanics behind LP-to-LP sniping are as follows. Assume that there exists a fast and a slow LP in the market, both placing orders at the $BB = 101$ and $BA = 103$ of the order book in *State 3* of Figure 1. Suppose the fundamental value jumps from $s = 102$ to $s' = 106$, which causes both LPs to update

²⁰For the other treatments, this effect is much less pronounced.

their outstanding orders accordingly, trying to keep the distance to the fundamental value constant. The fast LP manages to do so in advance of the slow LP, causing the new bid of the fast LP at $\overline{BB} = 105$ to transact with the stale ask of the slow LP at $BA = 103$. In other words, even if not intended, the fast LP snipes the slow one. This trade, however, is quite advantageous for the fast LP as it yields a profit in the amount of: $s' - s - \frac{ms}{2} = 3$. Notably, this coincides with the profit a successful Sniper would obtain. The adoption of LP-to-LP sniping along with the relaxation of the assumption of instant reversibility of transactions, as well as the elimination of the conscious decision to enter the market as either LP or Sniper, constitute the defining features of the new framework for real-time trading, as discussed in the next section.

Finally, note that this section provides an initial overview of the relevant literature with a special emphasis on the most closely related studies and that it is by no means complete. The financial markets literature is extensive, for a review of the literature on high-frequency trading, consider Jones (2013), Biais et al. (2014) or Menkveld (2016).

3 A new Framework for Real-Time Trading

The relaxation of simplifying assumptions in the new framework of real-time trading creates a more realistic environment. This comes at the cost of losing analytic predictions regarding the equilibrium behavior and properties of the various market designs. The study is thus of exploratory nature. Simulations designed and run by the LEEPS lab are required to determine the market parameters that lead to equilibrium behavior under CDA consistent with the real world observations. Once these parameters are determined, laboratory tests are performed to subsequently confirm whether human subjects indeed converge to this equilibrium, if at all, as well as whether the behavior under FBA or IEX indeed deviates. The experiment, thereby, captures both the behavior and the associated impact of HFTs on market quality under normal, as well as stressful market conditions.

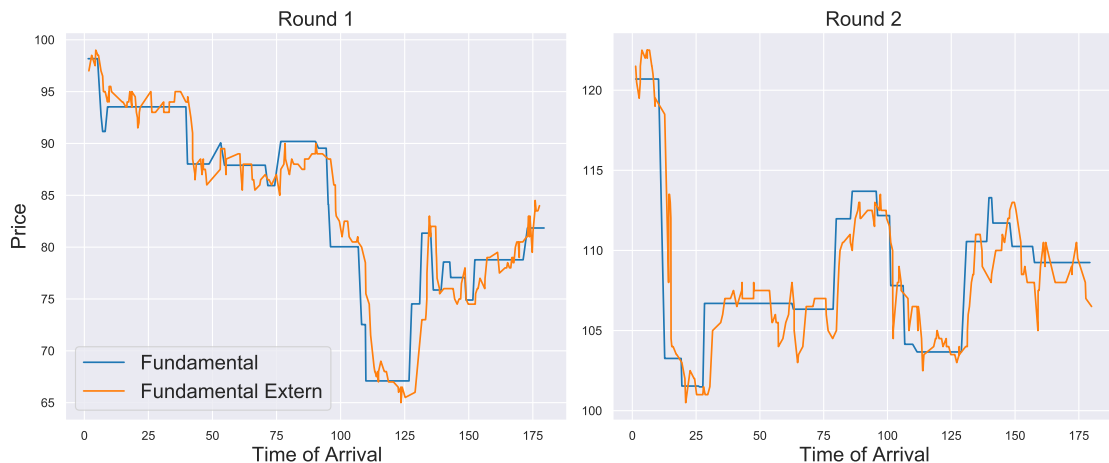
The new framework, which is forthcoming in the joint project, follows Budish et al. (2015) in that a single artificial asset is traded and the same two exogenous events are considered, i.e., price jumps and investor flows. The baseline parameters, which are not subject to simulations, are as follows. Each laboratory session is supposed to last roughly

90 minutes, consisting of 12 trading periods of 3-4 minutes each. The remaining time is allocated for reading the instructions and for a subsequent questionnaire. Subjects begin each trading period with an initial endowment of ECU 100 and the first two trading periods are trial periods that allow participants to get familiar with the experiment without needing to be concerned about their payouts. As suggested by Aquilina et al. (2020), six subjects compete over the profit opportunities generated by the exogenous events in the role of HFTs. That is, subjects do not generate any payments by owning or owing the artificial asset, but rather by buying the asset at a lower price than selling it. The strategy space initially includes a *out*, *manual* and *automated* mode. Subjects selecting the *out* option choose to stay out of the market altogether, which means they will neither book losses nor profits. The two other options result in market entry, with the *manual* option granting subjects full control over their trading at the cost of limiting the speed of trading to their human reaction time, since subjects are required to place their orders by hand. Alternatively, the selection of the *automated* mode turns on the algorithm that trades on subjects' behalf, as outlined in Section 2.1. The automated mode opens up an additional strategy space, which allows subjects to fine-tune the functioning of their trading algorithm, composed of a *speed* option and two *sensitivity sliders*. The *speed* option reduces the communication time between the algorithm and the exchange server, similar as before, from 250 ms to 50 ms, at c_{speed} per second used. The faster service can be turned on and off at any time during the trading period. This represents an abstraction from the real world, as most of the investments of HFTs described in Section 1 are of a more permanent nature, yet, it assures a much improved learning process. The *sensitivity sliders* are slightly more complex and discussed in more detail towards the end of this section.

The conscious decision to operate in the market as either Sniper or LP is no longer part of the strategy space, which is one of the key distinctions from previous literature and which has significant implications for the underlying dynamics. Subjects, although still participating as HFTs, no longer book profits the moment any of their orders are executed. Instead, subjects will now maintain an inventory that increases by one with each executed bid, while each executed ask decreases it by one. Negative inventories are the result of short selling. Subjects realize their profit at the end of the trading period, depending on the revenues generated by their trading activities, their inventory and the stock price. The

stock price matters, since the units in a subject's inventory at the end of the trading period are priced based on a reference price P . The reference price is a weighted average of the most recent transaction prices, with more recent prices having a greater weight. As outlined in Section 2.1, the introduction of an inventory exposes subjects (HFTs) to the risk of adverse price movements.

Figure 5: Fundamentals of Trading Rounds 1 and 2



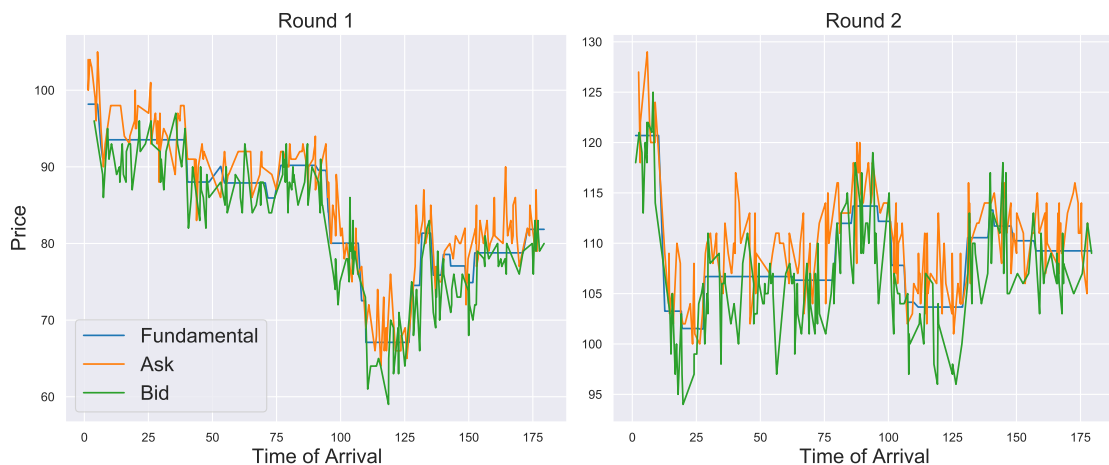
To illustrate, consider the blue graph in Round 1 on the left side of Figure 5, which plots an approximation of the fundamental value of the artificial asset over the course of one trading period. Suppose a subject buys the artificial asset for ECU 98 at the beginning of the trading period.²¹ Shortly after, the subject will book a loss on this transaction independent of the timing it decides to resell the asset. The reason is that the fundamental value will not reach the level of ECU 98 again for the remainder of that trading period. As an alternative to selling the asset, the subject could hold the share until the end of the trading period. However, the reference price at the end of the trading period with which the asset is valued will also be lower than the purchase price, so that the subject will also book a loss in this scenario. To illustrate that subjects using the automated mode are similarly exposed to adverse price movements and that the selection into the role of LP or Sniper takes place endogenously, continuing with Round 1 on the left side of Figure 5: Suppose a trading algorithm buys in second 75 because it anticipated the upcoming price increase due to the latest market information and signals. At this moment, the algorithm, or rather its subject, operates as a Sniper, trying to exploit stale quotas. After the price increase in second 75, the algorithm intends to resell the purchased unit at the

²¹ECU means *Experiment Currency Unit* and represents the currency used during the experiment.

higher price in order to generate a profit. Consequently, it places an ask, becoming an LP. Assume further that in the time interval from second 76 - 90, before the price eventually drops again, the algorithm is not successful in reselling the unit. Similar to the previous example, the algorithm will book a loss on that trade in any case.

Two out of the non-trial rounds will feature stressful market conditions with varying causes of stress, discussed in more detail in Section 4. The causes of stress will be alternated between groups within a treatment. In general, a price jump occurs on average every 10 seconds, with the size drawn randomly from a uniform distribution with support $[-20, 20]$. The investor flow is such, that on average an investor arrives every half a second on either market side. Figure 6 illustrates an exemplary implementation of the two processes. It shows the imprecise price knowledge of artificial investors, which often causes their orders to be executed as soon as they arrive at the BB or BA. In this respect, investors behave naively, generating profit opportunities for subjects.

Figure 6: Investors' Orders in Rounds 1 and 2



The first sensitivity slider to fine-tune the functioning of the trading algorithm, the *inventory sensitivity* slider, is a tool to not only control the risk of adverse price movements, but also the risk of tax exposure. Subjects are charged a tax at the end of the trading period for each unit they own or owe. The purpose of this tax is to provide subjects with an incentive similar to that provided by adverse overnight price movements to HFTs in the real-world, namely, to end the trading period with an inventory of zero or close to zero. Increasing the slider causes the trading algorithm to bring the subject's inventory balance close to zero: Suppose a subject with a positive inventory increases the inventory sensitivity, the algorithm will tend to be less aggressive on the buy-side by lowering its

bids and more aggressive on the sell-side by also lowering its asks. Consequently, buy transactions become less likely, while sell transactions become more likely. In case of a negative inventory, the algorithm works vice versa. The inventory sensitivity slider can be adjusted on a scale from 0 to 1 in 0.1 increments. Turning the slider all the way up to 1 results in a special case, wherein the algorithm will continue to place a bid (ask) equal to the BA (BB) until the negative (positive) inventory returns to zero. This strategy is potentially very costly, as the algorithm no longer seeks to maximize profit, but rather aims primarily at reducing inventory. Nevertheless, it might be useful under certain scenarios, e.g., towards the end of a trading period, in order to avoid high tax payments; or to reduce the exposure to adverse price movements. Since the exposure to adverse price movements grows proportionally with each unit in the inventory, the latter could be particularly relevant under stressful conditions, where the previous assumption of immediate reversibility of transactions might be most ambiguous.

The *external feed sensitivity* slider determines the sensitivity of the algorithms to price movements in an external market, which is another distinguishing feature of the new framework. There exists an external market listing the same artificial asset and receiving essentially the same investor flow. The external market has its own best bid and best ask, called BBext and BAext, and thus its own fundamental value at the midpoint of these two values. A corresponding implementation of its fundamental value is illustrated by the orange line in Figure 14. Price movements in one of the two markets can be understood as perfect signals for the other market, as outlined in the BCS model. Consistent with Hagströmer and Menkveld (2019), the blue line sometimes runs ahead of the orange line and vice versa.²² If a subject increases the sensitivity of its algorithm to the fundamental value on the external market, by means of the external feed sensitivity slider, the algorithm places orders closer to the external market's signals. This is a valuable strategy precisely when the orange line is rushing ahead of the blue line. In these instances, an algorithm that is more sensitive to the external feed will react faster to changing prices than a less sensitive algorithm. To illustrate, consider the fundamental values in Round 2 on the right

²²A short note: The blue line, while certainly plausible, is only an imperfect representation of fundamental value. The reason for this is that the fundamental value is only captured by the software whenever an artificial investor places an order or whenever a subject makes a decision. In this respect, the blue line displayed may be smoother than it actually is and there could be more or fewer cases where the orange line leads the blue line. Following this realization, the experimental software has been extended to include a measure that records the fundamental value at each point in time.


hand side of Figure 5. Within the first 10 seconds of the trading period, the fundamental value on the external market drops ahead of the fundamental value on the subjects' market. An algorithm more sensitive to the external market observes this signal and adjusts its outstanding orders accordingly, i.e., it clears its own bid and tries to snipe as many stale bids as possible. As previously, the slider can be adjusted on a scale from 0 to 1 in 0.1 increments, with setting the slider all the way up to 1 creating a special case. In this case, the algorithm places and adjusts its orders solely based on changes in BBext and BAext. Note that this strategy is always disadvantageous whenever the blue line leads the orange line.

3.1 Trading Controls and Information

This section introduces the *user interface* available to all subjects during each of the trading periods, which includes the strategy space mentioned above and provides crucial information to the subjects. Previous studies often had to truncate the trading periods as the first few seconds would usually be used to enter one's strategy. The experimental software of the joint project avoids this by introducing an *Initial-Strategy-Page*, which allows subjects to specify the strategy they intend to enter the market with. This page is a simplified version of the user interface depicted in Figure 7, containing only the trading control buttons that correspond to the strategy space described earlier. The user interface of Figure 7, to which subjects are redirected after the Initial-Strategy-Page, is available throughout the trading period and consists of three panels: The top panel is a graphical representation of the order book as a horizontal price line. Here, blue dots represent asks and orange dots represent bids. Larger dots mean more liquidity at the respective price. The lighter dots represent the subject's own orders, while the black borders indicate the BB and BA prices. The BBext and BAext are denoted above the price line by the smaller orange and blue dots. The middle box contains the trading control buttons of the Initial-Strategy-Page. In addition, it displays the BB and BA, as well as the subject's bid and ask in written instead of graphical format, and provides a visual cue whenever one of the subject's orders is executed. The remaining items *PBD* (Payoff Before Deduction), *Inventory*, and *Cash* are performance indicators that are explained in more detail below. The bottom panel is a graphical representation of the subject's decision regarding speed,

depicted by the shaded area in the "Sensitivities" graph between 01:15 and 01:30; the subject's decision regarding inventory sensitivity, depicted by the blue line in the same graph; and the subject's decision regarding external feed, depicted by the purple line. Finally, the progression graph in the lower part of the bottom panel shows the evolution of the subject's PBD value, which provides the main measure of performance and is discussed next.

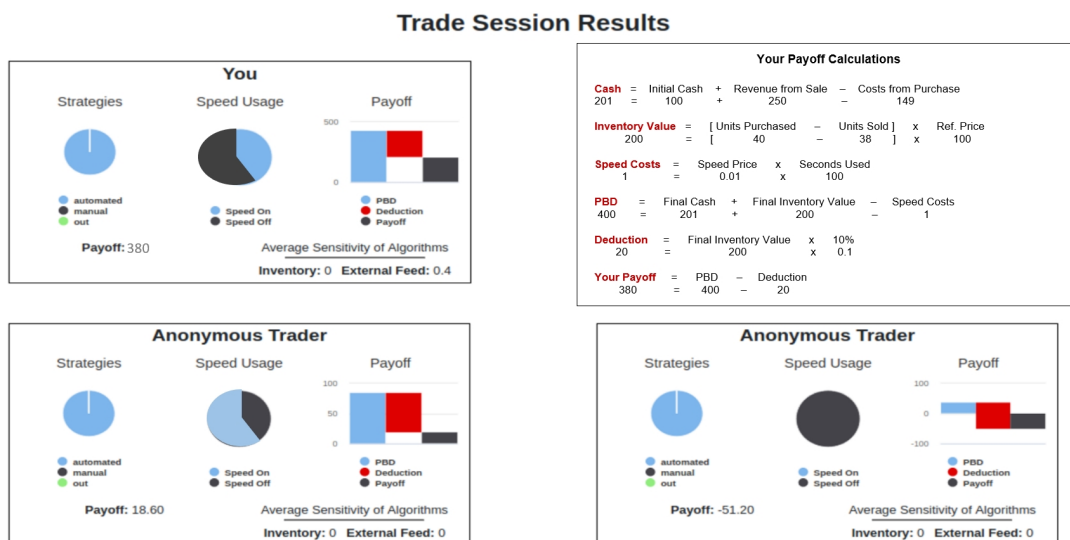
Figure 7: Subjects' User-Interface



The PBD is a reference value for the profit a subject would realize if the trading period ended at that moment. It depends on the other two measures in the middle panel, the inventory and the cash position. The inventory position works as outlined above, i.e., it increases by one with each executed bid, while each executed ask decreases it by one. Similarly, an executed bid decreases the cash balance by the transaction price, while an executed ask increases it by the respective price. Subjects start the trading period with an inventory of zero and a cash balance of ECU 100, i.e., they hold their initial endowment in cash. The PBD is displayed at each instant of the trading period, calculated by multiplying the units in the subject's inventory at that instant by the current reference price, P , adding the subject's current cash balance, and subtracting the cost of speed in real time. Consequently, the PBD displayed in the very last second of a trading period coincides with the subject's profit in that period, before taxes. This is where the name of this position originates from. Although it is essentially a tax, the term "tax" has been replaced by "deduction" to avoid confronting subjects with the sentiment associated with taxes.

The results page of Figure 8 is displayed to subjects after each trading period and contains the following input in addition to the information available during the trading period. The "Your Payoff Calculations" table on the top right shows a detailed calculation of the subject's profit. In the example in Figure 8, the subject ended the trading period with a cash balance of ECU 201, an inventory of 2 units at a reference price of $P = 100$ ECU, and a tax rate of 10%. The upper left box contains a graphical summary of the respective subject's profit and chosen strategies. The other boxes, sorted from highest to lowest payout, show similar graphical summaries for the other human subjects in that market. These information allow subjects to reconsider their strategies. The pilots in Section 5 confirm to some degree that they thereby pursue a process similar to that implemented in the simulations, which are discussed next.

Figure 8: Post Session Results Page



3.2 Simulations

The simulations are designed and performed by our colleagues in Santa Cruz, their work will be briefly outlined in this section, with the clarification that the simulations are not part of my efforts. For computational feasibility, the simulations seek to identify symmetric equilibria exclusively. Parameters that are subject to the simulations include, e.g., the batch length under FBA, the delay for IEX, or the speed cost per second. The simulations undergo the following iterative steps. Six simulated agents trade in the same market, initially all entering the market with no speed, no inventory sensitivity, and no

sensitivity to the external feed. In the first step, five agents maintain this behavior while one agent moves through all possible combinations in the strategy space of the automated mode. In this process, the simulations are compressed by more frequent investor flow and price jumps to one-twentieth of the usual duration of a trading period. This allows multiple simulated markets to be considered for each strategy combination, reducing the risk of results by chance. The iteration step ends with the non-distinguishable agents adopting the most profitable strategy of the distinguishable agent. Throughout the next iteration step, all of the non-distinguishable agents execute this previously most profitable strategy, while the distinguishable agent cycles through all combinations of a narrower strategy space, closer to the non-distinguishable agents' strategy. The simulation proceeds with five additional iterations before concluding with the non-distinguishable agents all playing the latest most profitable strategy, while the distinguishable agent cycles through the entire strategy space again. The objective of the last step is to verify that if all non-distinguishable agents play the latest most profitable strategy, the distinguishable agent has no incentive to choose a different strategy, i.e., that the strategy is a candidate for symmetric equilibrium. In addition to the compression, the simulations have another practical difference from the actual trading periods. An inventory sensitivity of 0.1, for instance, does not imply that the inventory sensitivity is 0.1 throughout the simulated trading period, but rather that it is below this threshold at the beginning of the period and above it towards the end.

For each new set of parameters, simulations are first executed under CDA with the purpose of achieving an equilibrium that is consistent with real-world behavior. The next step is to run simulations with the same parameters for FBA (or IEX) to verify whether a different equilibrium behavior is reached. Although one would expect this based on the results of Section 2.2, recall that the theoretical predictions may no longer hold true due to the more realistic environment. Moreover, it could be that the parameters are chosen such that the incentives for subjects are unreasonably strong, i.e., such that the simulations lead to the same equilibrium for each market design. The latter could arise, e.g., as a result of c_{speed} being set too high, causing the equilibrium behavior under each market design to display no speed investment at all. Choosing the right c_{speed} is not trivial, as it is a proxy for a variety of investments and therefore difficult to quantify empirically. The latest set of parameters led to the following desirable differences between CDA and FBA: (i) higher

sensitivity to external feed under CDA; (ii) profits are more sensitive to strategy choice regarding external feed under CDA; (iii) turning on speed leads to a slightly lower payoff on average under FBA, while the opposite is true for CDA. These results are essentially consistent with the theoretical prediction, (i) and (ii) can be interpreted as suggesting that observing the external market signals is more valuable under CDA, which is to be expected. The interpretation of (iii) is straightforward and implies that investing in speed is of lower value under FBA than under CDA. Looking at the fundamental values in Figure 14, it is striking that the external feed seems to lag behind most of the time. The latest simulations capture this by determining a sensitivity of the external feed under CDA equal to $[0.2, 0.25]$ to be optimal, implying that subjects should indeed not place too much emphasis on the external feed. If this set of parameters is ultimately found to be suitable, then experiments will be necessary to confirm whether human subjects do indeed behave as simulated. The equilibrium behavior identified by the simulations could depend on the very specific path taken by the simulations. The pilots in Section 5 provide suggestive evidence, however, that subjects at least to some extent behave as simulated, in that they copy the strategy of the most successful player from the previous round, similar to the non-distinguishable agents in the simulations.

3.3 Empirical Strategy

A fundamental reason for economists to conduct experiments in the laboratory, which is particularly relevant in the context of financial markets, is the data generating process in a highly controlled environment. In the real world, neither the fundamental value, the aggregate arbitrage opportunities, nor all activities of HFTs are generally observable. There are reasonable measures for the size of the price pool based on the original ideas of Glosten (1987) or Stoll (1989), but a direct measure is lacking. The same is fortunately not true in the laboratory. The generated data includes all profits and actions of HFTs, such that the size of the arbitrage price pool can be measured directly. The same is true for liquidity, since the data includes a position that captures the volume at BB and BA at each instance of the trading period. In addition, the data includes standard limit order book data such as a record of BB, BA, and the transaction prices. This enables a calculation of the market spread at any point during the trading period. Additionally, it allows the

calculation of the mean square deviation between transaction prices and the fundamental value, which serves as a measure of price efficiency. Beyond the measures mentioned so far, aggregate speed investments and aggregate role choices represent the main dependent variables of interest, denoted by y in the following regression equation:

$$y_{gt} = \sum_{i=1}^5 [\alpha_i \cdot MC_{it} + \beta_i MC_{it} \times FBA_g + \gamma_i MC_{it} \times IEX_g] + \varepsilon_{gt} \quad (1)$$

The dependent variable is indexed at the group level, denoted by g , and the round, denoted by t . The independent variables of the regression equation are dummy variables. The term MC_{igt} identifies the market condition that is considered in round t of group g , with $i \in \{1, \dots, 5\}$. It holds that $MC_{1gt} = 1$ if round t of group g features the stress condition S_1 , which is introduced in the next section, and $MC_{1gt} = 0$ otherwise. Note that $MC_{5gt} = 1$ denotes a special case, namely, it marks that in round t of group g one of the normal market conditions is considered, without specifying more precisely which one. Moreover, $FBA_g = 1$ holds if the group g features the market design FBA, otherwise the term is equal to 0; analogously for IEX_g . Consequently, this analysis will generate 15 non-zero coefficients. Thereby, α_i captures the mean of y under market condition i in CDA; β_i captures the mean difference in y between CDA and FBA under market condition i ; while δ_i captures an analogous measure between CDA and IEX. As a side note, role choices are not straightforward candidates for a standard regression, since they are discrete choices that take on three different values, required to be weighted by the time used. Note that the dependent variables *average market spread* and *average liquidity* are also required to be weighted by time. Fortunately, the experiment captures a time stamp, such that weighting by time is straightforward.

Since the final parameter set has not yet been found, a precise power analysis has not yet been performed. A rough first estimation indicates that at least 24 different groups with 6 subjects each should be considered. This would result in 8 groups per market design, with two alternating causes of stressful market conditions per group. The four different potential causes of stress are introduced next.

4 Stress Test

This section introduces the stress conditions that will shape the market in two out of the ten rounds. The stress conditions extend the analysis beyond just the comparison of the three market designs on an average trading day. This not only increases the robustness of the results, but also enables addressing several of the public concerns outlined in Section 1, including phantom liquidity and the amplification of stressful market conditions by HFTs. As with the normal exogenous fundamental values, the implemented stressful price movements are generated according to a random process, partly following the empirics of the Flash Crash. This particular crash, although basically any market crash or drastic price drop may cause stressful market conditions, indeed even extreme booms may as well, is of special interest as the crash phase is concentrated on just a few minutes. This allows a replication of related price movements and liquidity drain outs within a single laboratory trading period. Moreover, the Flash Crash has the previously outlined advantage of being very well documented, especially with regard to the behavior of HFTs, as discussed next.

4.1 Flash Crash and Related Consequences

In the early afternoon of May 6, 2010, the Flash Crash unfolded, during which more than 300 securities were executed at prices more than 60 percent away from their respective opening prices (Kirilenko et al., 2017). Individual securities such as *Apple Inc.* (AAPL) and *Accenture Plc.* (ACN) traded at absurd prices, with the former selling for nearly \$100,000 per share, with an opening price of \$36.26, while the latter fell from \$39.98 at 2:46 p.m. to 1 cent at 2:49 p.m., only to rebound to \$39.51 a minute later (Fox et al., 2015). The Figures 19 and 20 in the Appendices depict the price movements and market depth of AAPL and ACN, respectively. But as spectacular as such price jumps are, they were mostly the result of orders executed against "stub quotes"²³ after one side of the market had completely dried out. These orders were subsequently reversed by the exchanges since they were executed at "clearly unrealistic prices under difficult market conditions". This raises the question of which price jumps would still be permissible causing actual profits or losses. The **Limit Up-Limit Down Rules** (LULD), approved by the SEC

²³Stub quotes are orders that market participants set intentionally far away from the fundamental value in order to meet liquidity obligations without intending for those orders to ever be executed.

on May 31, 2012, as a response to the Flash Crash, set clear guidelines in that regard (Hughes et al., 2017). The LULD stops trading of a security whenever the price jumps too drastically in either direction over a "prolonged" period of time. The idea behind the trading pause is that it provides sufficient time for price discovery, avoiding amplifying panic situations. A jump is classified as drastic, based on the following criteria:

$$\text{Limits} = \text{Reference Price} \pm \text{Reference Price} \times \text{Percentage Parameter}$$

with the *Reference Price* being the mean transaction prices over the preceding five minutes and the *Percentage Parameter* depending on the particular security's price, as depicted in the following Table 1.

Table 1: Limit Up-Limit Down Percentage Parameters

Tier 1	
Previous Closing Price	Percentage Parameter
Greater than \$3.00	5%
\$0.75 up to and including \$3.00	20%
Less than \$0.75	75%

Tier 2	
Previous Closing Price	Percentage Parameter
Greater than \$3.00	10%
\$0.75 up to and including \$3.00	20%
Less than \$0.75	75%

The market is divided into Tier 1 and Tier 2 securities, the former are securities listed in the S&P500 index, the *Russell 1000* index or listed on the *LULD Tier 1 ETP List*, the latter list contains all other securities.

If a security experiences a price jump outside of these limits lasting for more than 15 consecutive seconds, which can be understood as "prolonged" in the context of financial markets, trading of that security is halted for 5 minutes. The SEC later expanded the LULD by issuing a rule that during the first and last 15 minutes of a trading day, limits are calculated using twice the percentage parameters, which essentially means, that larger price jumps are permissible at the beginning and end of a trading day.

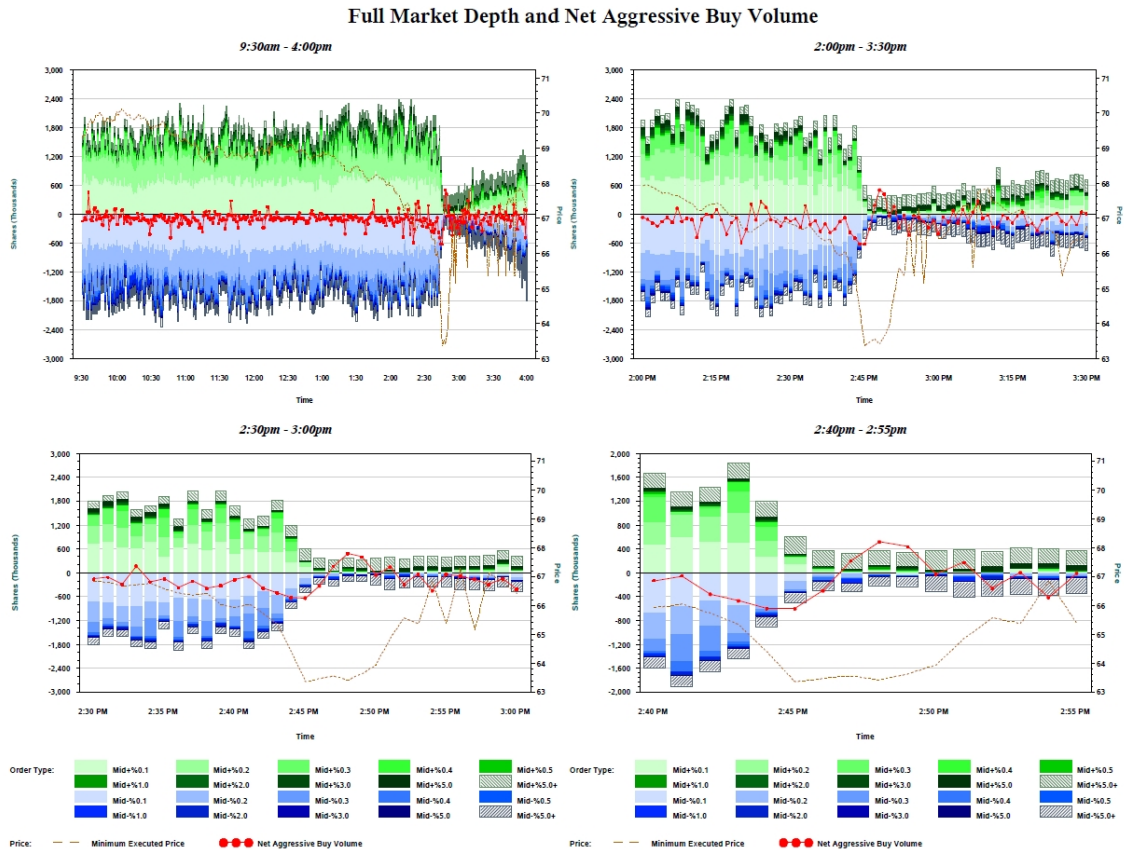
In general it appears, that trading pauses tend to happen more frequently under stressful market conditions. More recent examples include March 9, 2020, and March 16, 2020,

during the aftermath of the COVID 19 pandemic, and January 25, 2021, during the GameStop frenzy, which triggered trading halts nine times on NYSE that day. This implies that it is not uncommon for trading halts to be triggered during stressful market conditions. So it could be a fair strategy to generate fundamental values for the laboratory replication, that are close to triggering a trading halt, but do not actually do so, as otherwise the laboratory trading period would essentially end immediately.

In the experiment, stress can not only be induced by such price movements, but also by the second exogenous event, the investor flow. When the inventory sensitivity slider was introduced in the beginning of Section 3, it was established that the automated mode exposes subjects to the risk of adverse price movements, especially, if the algorithm is not able to balance the inventory at the desired time. The drying out of investor flow increases the risk of an algorithm being unable to balance the inventory timely. The Flash Crash was supposedly triggered by a large sell order of 75,000 e-mini futures worth about \$4.1 billion. This order was said to have caused a liquidity imbalance that spread throughout the market. Figures 19 and 20 in the appendices show that in the case of AAPL and ACN, the sell side indeed strongly prevails over the buy side, at least deep in the order book, while Figure 9 shows that the same is not true for the highly liquid *iShares Russell 2000 Index*. Moreover, these figures do not reveal whether an imbalance in liquidity existed at the fundamental level, which would be most relevant for HFTs. More recently, Aldrich et al. (2017) derive that a liquidity imbalance exclusively existed deep down in the order book and that an imbalance is therefore unlikely to have triggered the Flash Crash. This conclusion is consistent with the stylized facts established by Bouchaud et al. (2009) and Eisler et al. (2012), in that a price move in one direction generates order flow at the fundamental value on the opposing market side. In short, based on these insights, a significant price decline should not be implemented in the laboratory along an imbalance in the order books at the fundamental value, but certainly along a drying out of liquidity on both sides of the market, see Figure 9. The substantial drop in liquidity that accompanied the drop in prices during the Flash Crash may partly explain the behavior of HFTs discussed next.

Traders classified as HFTs by the SEC in the U.S. Securities and Exchange Commission (2010) report, traded 140,000 E-mini contracts between 2:41 p.m. and 2:44 p.m., during the period when the E-mini experienced the sharpest price decline. The SEC reports

Figure 9: Market Depth and Prices of iShares Russell 2000 Index on May 6, 2010



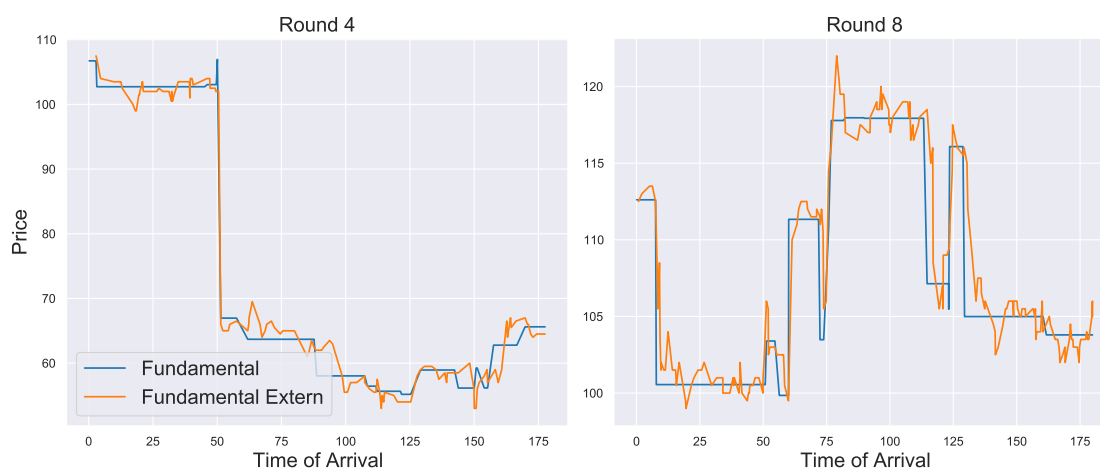
This figure is taken from the U.S. Securities and Exchange Commission (2010) report. Market depth is calculated by aggregating outstanding limit orders at the minute level. The blue bars show the market depth on the buy side, while the green bars show the market depth on the sell side. The brightness of colors indicates the distance from the midpoint. Four different time periods are considered, with the peak of the Flash Crash lying between 14:40 and 14:55. The dashed yellow lines show the minimum transaction price per minute. The dotted red line represents the net aggressive buy volume. Aggressive orders are market orders or limit orders at or below (above) the BB (BA).

that most HFTs held positive inventory averaging 3,300 contracts before the crash, which could be a result of relatively high prior selling pressure. They sold a net average of 2,000 units during the three-minute period, which corresponded to two-thirds of their holdings, implying that they acted as liquidity takers during that time. The 140,000 contracts, which accounted for 33% of the total volume in this period, are thereby below the average of 50% under normal conditions reported in Section 1. A potential reason could be that a significant number of HFTs withdrew from the market altogether or relied on manual trading despite the speed disadvantage. As the market normalized toward the end of the day, the percentage of contracts traded by HFTs increased in proportion to the increasing trading volume, which is highly consistent with the empirical facts in Section 2.1. However, not all HFTs continued trading after the peak of the Flash Crash, some abstained from trading completely for the rest of that trading day. Since the strategy space of the experiment includes the out and manual mode, subjects can potentially behave exactly as HFTs did under stress in reality. This would not only reinforce the external validity, but also allow for a comparison of market designs. More specifically, based on the theoretical results, one would expect the magnitude of crashes to be more moderated under FBA (and possibly under IEX) with an improved market quality. Possibly as a result of HFTs providing a positive net liquidity and competing in price rather than speed. The implementation of potential stress conditions is discussed next.

4.2 Stressful Market Parameters

After consultation with all collaborators, the experiment is intended to consider four stress conditions, S_1, \dots, S_4 . The first three are intended to represent stress conditions related to those experienced during the Flash Crash. Condition S_3 essentially creates a market similar to the most intense minutes of the Flash Crash. Whereas S_1 and S_2 aim to disentangle the effects of the price decline from the effects of the drying out of liquidity. In S_4 , a so-far unmentioned stress condition is induced, namely, higher volatility, i.e., price jumps are amplified. While this may well be called a stressful market condition, the question is, whether it actually causes stress for HFTs or whether it is instead rather beneficial for them. Below are the random processes that generate the implemented stress conditions:

Figure 10: Fundamental values of S3 and S4



- S1 Temporary Crash:** ECU 40 drop at a random time in $[45, 90]$, λ_{invest} , λ_{jump} and F_{jump} remain unchanged
- S2 Dry Out:** λ_{invest} is decreased by 50% at a random time in $[45, 90]$, λ_{jump} and F_{jump} remains unchanged
- S3 Temporary Crash + Dry Out:** At a random time in $[45, 90]$, combination of S1 and S2
- S4 High Volatility:** At a random time in $[45, 90]$, the support of the uniform price jump distribution is increased by a random selection in $[100\%, 200\%]$

The fundamental values of the implemented S3 and S4 for Pilot 2 and Pilot 3 are shown in Figure 10. The stress conditions were scheduled to run in Rounds 4 and 8 for Pilot 2, respectively, hence their names in the graph. It was randomly chosen that the price drop, along with the thinning of investor flow in S3 happens in second 51. The ECU 40 price drop, corresponding to a roughly 40% decline, which would not trigger a trading halt if the artificial asset, regardless of tier classification, were to trade in the real world within the first or last 15 minutes of the trading day between $[\$0.75, \$3]$ or regardless of the timing but below $\$0.75$. The timing of the doubling of volatility under S4 was randomly chosen to start at second 65.

At this point, as well as again in Section 5.2, I would like to draw the readers' attention to possible adjustments to these stress conditions, which I am convinced of, would improve real world proximity. Although prices fell extraordinarily fast during the Flash Crash, it

did not happen from one second to the next. Looking at Figures 19 and 20, even the prices of the individual assets with most extraordinary price movements declined over the course of several minutes. As of now, subjects in the laboratory experience a sudden drop in prices in *S1* and *S3* to which they are unable to respond. The damage is done from one second to the next, without subjects even having the chance to click on the out or manual button. Therefore, subjects basically have no opportunity to behave as one would expect based on the real world observations. After the price drop, when subjects have realized what has happened, it is already too late, prices have settled at the new level between ECU 55 and ECU 65, see Figure 10. At this point, there is no incentive to switch the strategy to manual mode, or to exit the market, as the trading algorithm can now continue to seek to profit from trading, just as it did before the crash. The sole difference being that orders are now executed at a lower price. So instead, the laboratory settings should consider a more or less continuous and significant drop in prices over the duration of, e.g., one minute. This would not only be more in line with reality, but would also provide subjects with enough time to realize that they are experiencing a market crash and to reconsider their chosen strategy. However, this is only the first part of the proposal to improve the stress test, the second part is discussed in the following section.

5 Insights from Pilot Sessions

In early 2020, a preliminary version of the real-time trading experiment software was completed, originally intended to be hosted at the Cologne Laboratory for Economic Research (CLER). Unfortunately, the laboratory facilities had to be closed shortly after due to COVID-19 regulations in Germany. Meanwhile, the interim time has been used to carry out a series of pilots that proved valuable for various reasons. The pilots provide insights beyond the scope of the simulations by examining the extent to which subjects understand the implications of their decisions and whether the instructions and the user interface are satisfactory in terms of handling and comprehension. The latter is particularly important as subjects of the CLER subject pool are typically students from a wide variety of academic backgrounds. Thus, from the outset it is fair to assume that creating a level playing field, in which subjects are equally familiar with the complex rules of financial markets, is challenging. The pilots are a useful tool to check to what extent the instructions

succeed in this, and where improvements are needed. In addition, the pilots are useful to detect errors in the experimental software. With well over 100,000 lines of code, the software is among the most sophisticated in our field of research, which also implies that there exist countless potential sources for errors. Minimizing errors is important, as an interruption of the laboratory session will most likely render the associated data unusable. The pilot studies presented in the remainder of this section do not seek to identify causalities, but instead attempt to detect errors, confounding, and undesirable exploitation opportunities by analyzing summary statistics and subsequent questionnaires. At this point it should be noted that the figures shown so far represent the current state of the experiment (December 2021), i.e., most of the insights from Pilot 1, Pilot 2 and Pilot 3 have already been implemented, while the execution of possible additional pilots is still pending.

5.1 Pilot 1

Pilot 1, which was carried out on July 3, 2020, operated on a very preliminary version of the experiment, featuring only a single repeated implementation of the exogenous events. This implies that Pilot 1 did not include stress conditions and in some sense became obvious after the first round. Basically, the subjects could have known the upcoming price movements at the beginning of the second round. In fact, some subjects figured this out and took advantage of their knowledge using the manual mode. The subjects consisted of six colleagues from Prof. Ockenfels' chair who were unfamiliar or only partly familiar with the project. Subjects were instructed to maximize their profits, but they were not compensated beyond their wages based on their decisions in the pilot. The experiment began with three trading periods under the CDA market design, followed by three trading periods under FBA and IEX each. The duration of each trading period was set to 4 minutes, the tax rate to 20%, and the $c_{speed} = 0.02$ ECU. In Rounds 1 and 2 of the CDA market design, subjects' right mouse button did not function, making the manual mode virtually unusable. In round 3, this bug was fixed, but no artificial investor flow arrived at the market. This and other technical errors were fixed by our colleagues in Santa Cruz prior to Pilot 2.

Subsequent analysis of the data, as well as the questionnaires, revealed the following

insights. The external feed was selected in 27% of the cases, but subjects indicated that they did not sufficiently understand the concept and functioning of this slider, which led to the current version of the instructions placing more emphasis on explaining the external feed. Subjects expressed skepticism as to whether the switched-on algorithm works as intended based on the confirmation they received from the user-interface. Indeed, the inventory position, f.ex., is only a very imperfect indicator of the functioning of the algorithm, if at all, since changes in the inventory often happen faster than the human eye can comprehend. The solution has been the introduction of a visual cue in form of red flashing dots in the user interface, which provide direct feedback on the functioning of the algorithm. The results page of Pilot 1 contained only the graphical summaries, which the subjects felt did not provide enough information about their own performance to thoroughly update their expectations and derive new strategies. Consequently, the "Your Payoff Calculations" table was added to the results page. The graphical summaries of the results page were initially randomly ordered. Subjects indicated that they would prefer sorting from best to worst, since they would focus primarily on the previously most successful player. Subjects explicitly stated that they based their own strategy very closely on that of the previous round's best player, which, as highlighted previously, supports the approach of the simulations. The current sorting from best to worst might even reinforce this behavior, which is suggestively supported by the behavior in Pilot 2, discussed below. According to subjects' accounts, the usage of the inventory slider was as intended, i.e., particularly towards the end of the trading period for clearing the inventory position. Subjects further reported that making positive profits in the experiment was non-trivial and that they viewed the out option as a profitable strategy. The out option was played about 8% of the time, but this result may be biased as subjects also indicated that they avoided the out option for the sake of the experimenter. Subsequent interviews revealed that the tax rate might have been the driving factor, since it was perceived as too high, which is why it was reduced to 10% for the following pilots. Similarly, speed was chosen only 17% of the time, prompting a reduction from ECU 0.02 to ECU 0.01 per second. All adjustments have been implemented prior to the second pilot, which is discussed next.

5.2 Pilot 2

Pilot 2, conducted on December 18, 2020, ran exclusively under the CDA market design. The pilot was scheduled for 10 rounds and, unlike its predecessor, was to include different implementations of the exogenous events in each round, in particular including $S3$ and $S4$ in rounds 5 and 6, respectively. An unforeseen error ended the experiment prematurely after the sixth round. The fundamental values of all rounds considered are shown in Figure 13 in the Appendices. This and other technical errors were resolved post-experiment in Santa Cruz. Each of the six rounds lasted 4 minutes and involved six subjects with different scientific backgrounds and degrees. The subjects included two mathematicians with at least a bachelor's degree, two economics students, including one undergraduate student, one doctoral and one post-doctoral student economics. Due to the ongoing COVID-19 lockdown in Germany at that time, this experiment was carried out online, which constituted the first time this experiment was held remote. At that time, the online version was considered to be only an imperfect transitional solution, suitable for debugging and assessing subjects' understanding of the instructions, but not feasible as a permanent solution. However, follow up analysis suggested that an online version could prove to be a valuable alternative, which is why it is the main focus of Pilot 3 in Section 5.3. Prior to the experiment, all subjects received the instructions and were directed to the *Initial-Strategy-Page*. After the end of the experiment, a brief interview session was held with all subjects via Zoom, after which the questionnaire depicted in Figure 17 was distributed.

The main findings include that the external feed continues to cause difficulties for the subjects. For example, some reported to have understood the external market to be "sometimes faster, but never slower", leading them to always set the external feed slider to 1, because it was understood to be "a weakly dominant strategy". Therefore, the instructions explaining the external market have been revised to minimize the risk of misleading. The current version of the instructions, as depicted in Figure 21 in the Appendices reads: "This sensitivity [external feed] can be useful because the external market will sometimes (not always) react faster than your own market to changes in patterns of automated investor orders, allowing your bot to act on market information earlier than other players. However, if you pay no attention to your own market, you might miss the times when your own market has reacted faster [...]". Subjects reported

Table 2: Speed Costs and Net Worth

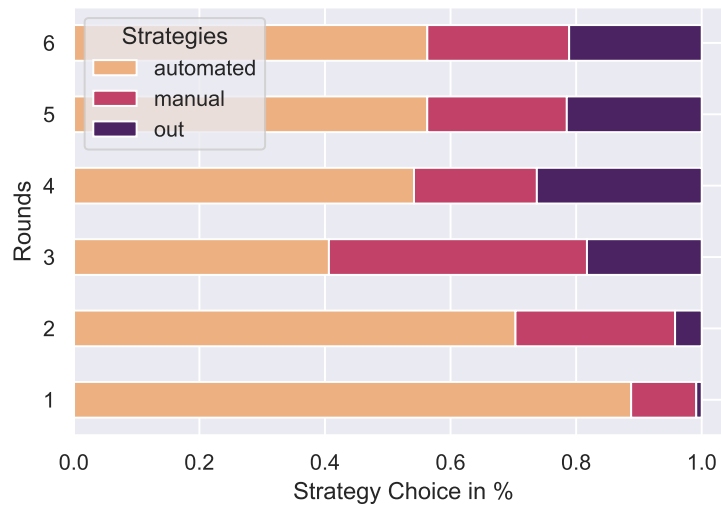
Round	Speed Investment				Profits			
	Number	min	max	mean	std	min	max	mean
1	0.0	2.1	0.7	0.8	-60.0	130.0	29.3	63.0
2	0.0	2.2	0.8	0.9	-120.6	144.6	-4.8	94.7
3	0.9	2.4	2.0	0.6	-251.0	154.0	-21.3	141.7
4	0.0	2.4	1.5	1.2	37.0	187.6	95.2	50.7
5	0.0	2.4	1.5	0.9	-125.8	323.6	130.4	161.6
6	0.0	2.4	1.9	0.9	-3.0	165.7	91.4	56.6

difficulties in distinguishing between bids, asks, best bid, and best ask in the instructions, which resulted in subjects now being trained in the instructions with the same colors that can later be seen in the user interface. The readers of this paper could experience for themselves how learning by coloring functions, based on the description of BB and BA in the exemplary order book in the Section 2. Moreover, subjects indicated that they had difficulties understanding the "PBD" position, which is why the position is now explained in more detail in the instructions, using the newly introduced "Your Payoff Calculations" table as an example. Note that the table itself was also revised after the second pilot, since subjects complained about the calculations being too "complicated and overwhelming". The current, simplified version of the table is as depicted in Figure 8. Prior to the experiment, it was determined that the experiment would begin once all subjects had read the instructions and entered their initial strategy. According to the subjects, this created a sense of time pressure, especially because reading the instructions took a perceptually long time of about 25 minutes. This prompted some subjects to start the experiment even before finishing reading the instructions.

Regarding strategies, subjects chose the faster service more often, as intended, see Table (2). After the second round, at least one subject had speed turned on for the entire period, and speed was turned on for more than half of the period on average. Here, the correction from Pilot 1 to Pilot 2, i.e., changing c_{speed} from ECU 0.02 to ECU 0.01, may actually have overshoot the goal, though it should be noted that c_{speed} is fortunately subject to simulations, i.e., it could still change with the final parameter set.

Overshooting might also apply to the tax reduction: 3/6 subjects explicitly stated that they used large stock holdings to "gamble". They used the *manual* or *automatic* mode to achieve a large negative (positive) inventory, and shortly thereafter stopped trading or selected the *out* mode in the hope that the reference price would decrease (increase)

Figure 11: Strategy Choices



subsequently. Figure (11) shows that the strategies *out* and *manual* indeed were chosen significantly often and in relative consistent frequency after the second round. Subjects that adopted this strategy indicated that they found it profitable across trading periods, even though they regularly finished the trading periods with large inventories, accepting relatively high tax payments, Table (3). Given the very limited number of observations, it is not possible to find empirical evidence for the profitability of this strategy. In addition, it is important to recall that the experiment ended prematurely, so it is quite possible that subjects, given more time to learn, would have abandoned this strategy. Gambling is far from new in the context of financial markets, as discussed for instance by Kumar (2009) or Bali et al. (2011), but based on real-world observations one would expect HFTs to rely rather on their sophisticated algorithms than on gambling. If the gambling behavior does not dominate and does not completely ignore the intention of the tax, there may be no reason to worry for the time being. Unfortunately, unlike speed costs, the tax rate is not subject to the simulations, so the existence of this behavior is re-examined in Pilot 3.

The post-experiment questionnaire asked subjects several questions concerning stress conditions. This included a question asking whether a certain trading period was particularly present to their minds. All subjects were able to recall the drying out of the market in the fifth round, but none could recall experiencing a crash during that exact same trading period. A possible explanation might be that subjects had previously experienced similar prices and therefore did not perceive the *S3* price decline as a "market crash". Table 4

Table 3: Speed Costs and Net Worth

Round	inventory			tax_paid			
	min	max	std	min	max	mean	std
1	-5	10	5.6	0.0	81.0	31.0	33.8
2	-8	6	5.0	0.0	97.6	44.7	42.0
3	-1	7	2.9	0.0	45.5	13.0	16.9
4	-1	0	0.4	0.0	10.1	1.7	4.1
5	-11	5	5.4	0.0	70.4	19.2	28.0
6	0	4	1.5	0.0	43.6	14.5	16.4

shows that subjects experienced prices as low as ECU 64 and ECU 58 in the first and third round, respectively. These prices are only one and a half or one average standard deviation away from the overall minimum of ECU 51 during the crash under *S3*. In sum, although the fifth round had the lowest price and the highest standard deviation according to Table 4, market participants may have been familiar with the associated price movements based on the preceding trading periods. In this respect, *S3* may not have "shocked" subjects, which in reality a market crash would possibly tend to do. This might have prevented panic behavior and the associated negative effects on market quality.

The second part of my proposal to improve the stress tests therefore includes a general amplification of the price decline. For instance, it is conceivable that the increase would still take place within the limits of the LULD, if one assumes that the artificial asset is such that it would trade below \$0.75 in reality. Consequently, price declines of up to 75% would be permissible, or even up to 100%, if only the first or last 15 minutes of a trading day are considered. In this context, I suspect that a price decline in the range of 75% could be sufficient to be perceived as a market crash by the subjects. The confinement to a penny-stock-like artificial financial instrument certainly reduces external validity. That said, the volatility in our artificial market is on relatively high, such that we potentially also need larger price declines to convey the sense of a market crash. Furthermore, I believe this adjustment would improve the experiment overall, in part, because shock and panic are an essential component of a stock market crash and could potentially drive behavior. The amplified price declines should take place over a longer period of time, from several seconds to minutes, as previously described.

Another implication is that the subjects also failed to recall the double volatility in the sixth round. Again, an increase in the magnitude of stress could be considered in order to capture the attention of the subjects. After consultation with all scientists involved in

Table 4: Realized Prices

Round Number	Best Bid			Best Ask			Reference Price		
	min	max	std	min	max	std	min	max	std
1	64.0	96.0	7.5	67.0	100.0	7.42	67.0	96.0	7.3
2	97.0	123.0	5.24	102.0	127.0	5.25	102.0	122.0	4.78
3	58.0	103.0	11.9	61.0	106.0	12.49	62.0	100.0	11.7
4	91.0	108.0	3.5	93.0	109.0	3.55	93.0	108.0	3.22
5	51.0	107.0	20.52	55.0	108.0	20.23	56.0	107.0	20.33
6	96.0	113.0	4.91	100.0	115.0	4.83	100.0	113.0	4.47

the joint project, we decided to adjust the stress conditions according to this proposal. However, this involves a considerable work effort, so that the new stress parameters were not yet available for Pilot 3, which is discussed in the following.

5.3 Pilot 3

Pilot 3 was carried out on September 2, 2021. It was set up to involve 18 subjects, 6 per treatment, over the course of 12 trading periods of 3 minutes each per treatment. Rounds 1 and 2 were set to be trial rounds, while rounds 8 and 9 featured $S3$ and $S4$, respectively. The instructions can be accessed via a link and are made available only during the duration of the experiment. Subjects can advance from both the instruction as well as the subsequent questionnaire page only after 10 seconds in order to minimize the likelihood of unintentional click-through. This was not only the first pilot to investigate a remote experiment as main source of data, but also the first to be carried out with subjects from the CLER subject pool. To this end, the experimental software had to be extended by standard pages such as a consent page, a remote instruction page, a subsequent remote questionnaire and a payment page. Although these are standard features of economic experiments, some technical obstacles arose leading up to the execution of the experiment, mainly due to the complexity of the software, as discussed in the following.

5.3.1 Technical Issues Beforehand

The consent page must be embedded in the experiment, as of the guidelines from the European Research Council (ERC), under which umbrella this project is being conducted. According to the ERC, subjects' consent must be anonymized and included in the same data set as the behavioral data.

Furthermore, the questionnaire, shown in Figure 18 in the Appendices, was designed to be accessed via link or QR code. The latter facilitates subjects to film their trading screen for a few seconds, using their smart-phones and upload the corresponding file, which they are asked to do in the questionnaire. This serves the purpose of getting feedback on whether the user interface works as desired at the subjects' end, f.ex., with respect to responsiveness. Using a smart-phone is preferable to standard actions such as screen recording via Zoom, since these would operate on the same device as the experiment, which potentially could have adverse affects on the latency. In addition, subjects are presented with a convenient method to take and upload screen shots in case an error occurs, which they are also prompted to do in order to help debug the software.

Challenges to paying subjects using *PayPal* include ensuring anonymity, since in this experiment the payments are frequently unique. This problem is countered by using *AnonPay*, a software programmed at CLER under the supervision of our colleague Max R.P. Grossmann. While testing the software, it became apparent that both a lower and upper bound for payouts is necessary, since exceptionally high or low total profits are possible. Consider the following example for illustration: A subjects' algorithm that sells heavily at the beginning of *S3* will make huge profits as soon as the price falls in second 51. The subject would essentially make 40 ECU on each unit in its inventory, or € 1.2 per unit. Recent tests indicate that profits of ECU 1000 or higher, which amount to € 15 or up to € 30, in a single round are possible. In contrast, a subject who would buy heavily at the beginning of *S3*, would make losses of similar magnitude. Therefore, a clause had to be included stating that subjects would never receive less than the show-up fee of € 2.5 and never more than € 50, which is still a substantial hourly wage for 90 minutes compared to standard economic experiments. However, we recognize that this is not best practice for laboratory experiments and are searching for a feasible solution.

Another substantial problem that became apparent during the preparation of Pilot 3, is that standard tools of oTree for running multiple groups simultaneously, are not feasible. Part of the reason is, that a computational core handling the experimental processes reaches its limits with 6-8 subjects, causing the experiment to break down as soon as more than one group is considered. This is basically an inheritance of serial processing, since it requires that messages are processed by a single logical core. To illustrate, physical cores typically have multiple logical cores that they might switch between to avoid waiting

times. However, this could result in a logical core already processing a message even though it arrived later in the market than another message waiting in a queue at a different location. Therefore, to guarantee serial processing, the experiment is designed such that all messages are processed on a single logical core. Unfortunately, the standard oTree tool for multiple groups does not assign each isolated market its own logical core, but performs the calculations of all groups on one core. The solution required modifying the software so that each self-contained market is assigned a distinct logical (and physical) core. The hardware used for Pilot 3 has six logical cores and three physical cores, thus allowing us to run up to three groups simultaneously.

5.3.2 Results of Pilot 3

Not all treatments run through as intended. The IEX market did not start because of one subject denying consent. The Pilot 3 version of the experiment software did neither include a feature that allows for a subject to take another subject's place if the latter decides not to consent, nor did it allow overbooking . The current experiment version (December 2021) includes both features. Due to a technical error, which has later been resolved, no automated investors arrived at the market in rounds 7, 10, 11 and 12. The corresponding data is thus not insightful and therefore excluded from the following discussion. In the CDA treatment, one subject stopped playing in the third round. The experimenter took over the ui of this subject to advance the experiment. This renders the data biased, since there is no candidate for a suitable "default" strategy. The experimenter decided to advance the missing subject using the out option. The interference of the experimenter and the absence of automated investors most certainly distorted the data, however, the data is still analyzed in the following, since the purpose of Pilot 3 was troubleshooting, the identification of the possibility to run the experiment remote and the elicitation of subjects' level of comprehension.

First, the possibility of running the experiment online is discussed. This depends heavily on latency and whether the experiment runs as intended on the subjects' hardware. As mentioned earlier, the impact of latency between the subjects and the experiment server is less significant in automated mode, as opposed to the intuitive expectation. The bot operates on the experiment server, so the operation of the bot and the speed

option is not affected. Thus, larger latencies would only mean that the subjects see their bot's decisions later on their screen. In manual mode, latency is more significant because subjects must respond to the information presented to them on the screen. The information provided is directly affected by latency, so a subject who is near the experiment server and has a fast internet connection may receive new, simultaneously broadcast information earlier than other subjects. This is clearly an advantage for subjects with low latency, but the question is, whether subjects in manual mode can leverage this advantage. With average latency across players and rounds of 69.8 and 100.1 milliseconds in the CDA and FBA treatments, respectively, the answer is most likely no.²⁴ These average latencies are well below the 200 millisecond time span that a human requires to blink spontaneously. Since subjects in manual mode must process new information in the user interface, form a strategy, select the appropriate price in the horizontal price line, and press either the right or left mouse button to perform an action, the latency seems negligible in comparison. There is still reason to be cautious for now, because the experiment in Pilot 3 only recorded the average latency for each player in each round. This means that it cannot be ruled out that there were moments during the trading period when latency indeed created an advantage. The current version of the experiment also records the maximum latency of each subject in each round, which allows to control for these moments.

The screen recordings provided by the subjects indicate that the experiment is working as intended on subjects' side and that the trading controls are responsive. A common caveat to online experiments in economics is subject attention. While there are experiments where this is a minor constraint, such as when the experiment does not continue until all subjects have given their informed decision, this does not apply to the present experiment. Real-time trading and the dynamic nature of the experiment cause data quality to decrease when subjects are not paying attention to their trading screens. Thus in the follow-up to Pilot 3 a control variable that captures the percentage of time subjects stay on the experiment screen during each trading period was introduced. Subjects are assumed to have stayed on the screen if they did not click on another area outside their trading platform. However, this is an imperfect measure because it cannot detect whether subjects leave their workspace, for example, or pay attention to a second screen. Consequently,

²⁴Note: Outliers are excluded, for example the subject who disconnected from the experiment in round 3 of the CDA treatment, as this resulted in an average latency of 29695 milliseconds.

my overall conclusion is that a remote execution of the experiment is technically feasible and can be expected to provide adequate data in light of the discussed control variables. However, if COVID-19 permits, the experiment should preferably be conducted in person at the Cologne Laboratory for Economic Research.

Second, the perceived level of comprehension is analyzed. The subsequent questionnaire displayed in Figure 18 not only prompts subjects to share their screen records and error messages, but also contains questions that provide information on how subjects perceived the instructions, the ui and experiment in general. The closed questions are as follows: (A3) "How would you rate the clarity of the instructions?"; (A5) "How would you rate the clarity of the user interface?"; and (A7) "How would you rate your understanding of the experiment?". The closed-ended questions are based on a scale from 1 (very good) to 5 (unsatisfactory), which is intuitive for German students as it corresponds to the German grading system. The 11 subjects who completed the questionnaire gave an average grade of 3 (std 1.18) for (A3), 2.45 (std 1.13) for (A5), and 3.27 (std 0.77) for (A7). This might suggest that subjects overall lacked confidence in their level of comprehension. Each closed-ended question is followed by an open-ended question asking for a more detailed description of the parts that subjects found confusing, if any. The responses to these questions, listed below, paint a similar picture:

A4 Did you think any parts of the instructions were confusing? If yes, which ones: "For non-specialists relatively difficult formulated and abstract.", "Instruction was very long and with a lot of information. I would have found it better if you had received the information beforehand.", "It was just way too much."

A6 Did you think any parts of the ui were confusing? If yes, which ones: "Yes, the part where you could set bids and asks manually."

A8 Did you think any parts of the experiment were confusing? If yes, which ones: "I found the whole experiment confusing and could not make clear decisions.", "Speed + external Feed.", "Didn't quite understand what being "out" meant."

Note that the responses of strong negative opinions might be overrepresented because answering the open-ended questions was not mandatory. Overall, the experiment appears to continue to overwhelm non-expert subjects, and subjects continue to have difficulty

understanding certain concepts. This led to one of the biggest changes following Pilot 3, the introduction of video instructions. The video instructions are based on the written instructions. Subjects will be directed to an instruction page prior to the experiment, which will provide them with a video player streaming the instruction video. The video player allows subjects to fast forward and rewind. Tests with student assistants of Prof. Ockenfels' chair have shown that especially the rewind function is valuable for non-experts to understand complex elements like the external feed and the speed option. In the experiment, participants are allowed to leave the instruction page only after at least the duration of the video, which reduces the incentives to fast-forward or skip the video. Note, this also counteracts the time pressure discussed in Pilot 2. The written instructions from Figure (21) are provided during the experiment in case subjects want to look up any segments.

In creating the video instructions, the goal was to use multimedia learning tools in a targeted manner to facilitate subjects' learning of mostly unfamiliar, complex material. Studies from experimental economics and cognitive science literature provide principles to this end. The most closely related study by Schweitzer (2017) discusses measures to overcome challenges with respect to adequately teaching non-expert subjects the complex rules of multi-unit auctions. Following a pilot experiment that failed in this regard, the author successfully overcame the obstacle by using video instructions, an intuitive way of bidding, an optimized user interface, and automatic compliance with the auction rules. As described below, the design and handling of our experiment's trading platform already meets the requirements, and automatic compliance with auction rules is also in place. But there is room for improvement regarding the instructions and Schweitzer (2017) provides a blueprint on how to use the results of the cognitive science literature on multimedia learning in this respect.

To begin with, instructions should adhere to the modality principle, meaning that showing a picture of the ui or demonstrating it in a video can make its appearance and use more accessible. In contrast, a purely verbal explanation would require subjects' minds to construct a mental image from the verbal descriptions in order to visualize the ui and its functioning. The written instructions already contain images of the user interface. However, since the present experiment involves real-time trading, subjects had to imagine what the ui would look like if prices and their own returns were changing

dynamically. By using a video format, it is possible to show subjects a live session of a trading period, making the dynamic aspect accessible. In addition, video instruction can apply the multisensory principle, where verbal delivery of information has been shown to be preferable to text for complex learning material. In addition, in accordance to the literature, visual cues are used to highlight the parts of the ui that are being discussed. Recall that subjects in Pilot 3 rated the ui better than the instructions and the experiment as a whole, which could be due to our advanced trading platform. As suggested by Schweitzer (2017), the graphical representation of the ui is already sorted and arranges information through different layouts, colors, sizes, shapes, and contexts.

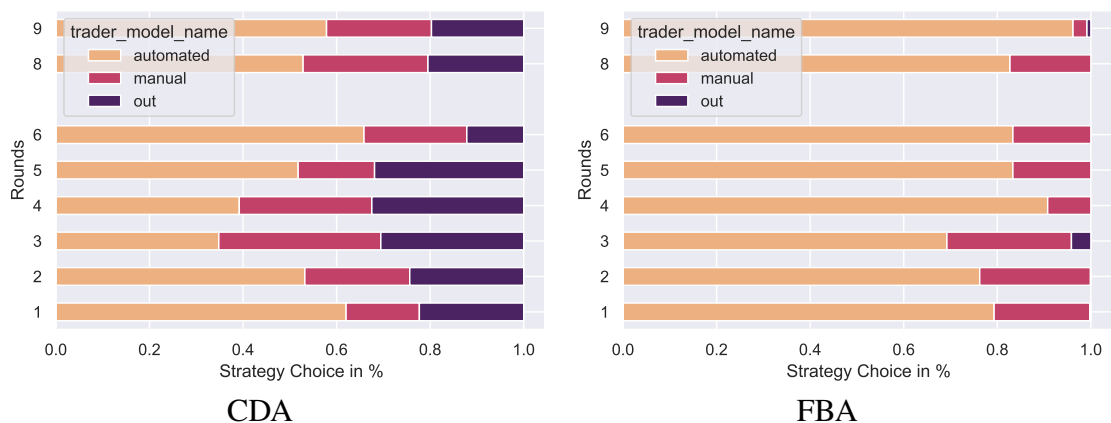
Evidence from cognitive science on multimedia learning further suggests that these principles should only be applied when appropriate, or extraneous overload could occur. Extraneous overload occurs when the essential cognitive processing required to understand the essential material in a multimedia message and the extraneous cognitive processing required to overcome extraneous material or to overcome the confusing arrangement in multimedia messages, exceeds the learner's cognitive capacity (Mayer and Fiorella, 2014). The principles we applied to prevent extraneous overload in the video instructions include: Signaling principle, by highlighting essential content in the ui and instructions; Redundancy principle, by conveying information from graphics and narratives rather than graphics, narratives, and text; Spatial principle, by grouping related information near rather than far on the screen; Temporal contiguity principles, by presenting related animations and narratives simultaneously rather than sequentially. Finally, as suggested in the literature, the ui and video instructions avoid sub-screens and pop-ups.

Incorporating insights from the cognitive science literature as well as lessons learned from Schweitzer (2017) should significantly increase subjects' level of comprehension. In order to test this hypothesis, more sophisticated measurement methods are used in the future. The subsequent questionnaire of Pilot 3 is sufficient to reveal previously undiscovered weaknesses, but provides only the perceived level of comprehension. How exactly the level of comprehension is elicited in the future is discussed in more detail in Section 5.3.4. The next section takes a brief look at the subjects' strategy choices and performances. Though recall, that the experimenter's interference and the absence of automated investors in 4 out of 12 rounds certainly biased the data, and that essentially only two independent observations are discussed, i.e., statistical power is not given.

5.3.3 Interpretation of Strategy Choices

Pilot 2 raised the question of whether the reduction of speed costs from 0.02 ECU to 0.01 ECU and the reduction of the tax rate from 20% to 10% overshoot the mark. The parameters remained unchanged between Pilot 2 and Pilot 3. In Pilot 3, the average speed cost under CDA was 1.3 ECU, which is larger than the average of 0.21 under FBA. The results are consistent with the theoretical prediction and show no unreasonable overuse of speed in either treatment. Thus, the speed reduction does not appear to have overshoot the target. With respect to tax reduction and gambling, Figure 12 shows the subjects' role choices for each round and both treatments. The strategy choices in the CDA treatment on the left are consistent with our observations from Pilot 2, i.e., there is a significant amount of manual and out use, though automatic use dominates. For the FBA treatment on the right, manual use is slightly lower and out-use is almost absent. Moreover, it can be seen from Tables (5) and (6) that the average tax payment for the payout-relevant rounds with normal market conditions is 32.7 ECU in CDA and 15.5 ECU in FBA. Subjects who use manual mode for more than 20 seconds in a given period, have a lower average tax payment of 22.3 ECU and 10.58 ECU. From this and from the fact that none of the subjects stated gambling behavior when asked about their strategies, it follows that the tax reduction did not overshoot its mark either.²⁵

Figure 12: Strategy Choices



Two features emerge from the graph and tables that are not explained by the underlying

²⁵The averages stated in this and the following segment have been adjusted to account for the intervention of the experimenter. The same is not true for Figure (12) and Tables (5) and (6), because the software of Pilot 3 did not include the participant code, discussed at the end of this section.

Table 5: CDA Performance

rounds		deduction_paid				payoff			
rounds	max	mean	min	std	max	mean	min	std	
1	60.8	25.33	0.0	24.82	100.0	-12.69	-259.8	131.35	
2	288.4	54.93	0.0	114.82	100.0	61.4	2.6	38.44	
3	73.8	17.77	0.0	28.13	107.8	77.05	-16.8	47.25	
4	108.0	27.0	0.0	40.27	111.2	66.83	-35.0	52.76	
5	96.0	24.53	0.0	36.3	183.11	48.08	-124.0	103.51	
6	66.0	37.4	6.6	23.12	173.8	36.09	-82.0	103.82	
8	52.0	15.17	0.0	18.69	153.5	85.66	2.0	65.34	
9	51.5	15.45	0.0	24.16	160.0	72.88	-19.5	66.59	

theory. First, the difference in role choice between treatments that already exists from the beginning, that is, especially before the experimenter's interference. Second, the well below initial endowment average payoffs of the payoff-relevant rounds with normal market conditions at 48.4 ECU and 37 ECU in CDA and FBA, respectively. Through the out option, each subject could secure a payout of 100 ECU. Already in Pilot 1 and Pilot 2, subjects identified the out option as a winning strategy. One might initially assume that we therefore only observe participation by risk-seeking subjects, or that participation is due to a lack of alternative occupation, as described for example by Lei et al. (2001) in experiments on financial market bubbles. But, compared to such zero-sum games, in the present experiment there exist real profit opportunities generated by the automated investors. The experiment was designed in such a way that generating profits is not trivial and requires careful tuning of the trading bot. This should lead to competition, as we observe in reality, and to only profitable subjects staying in the market. It is quite possible that in this process successful subjects even benefit from less successful subjects.

With only two observations and only 4 payout relevant periods with normal market conditions, no clear trend can be identified. Round 4 of the FBA treatment provides reason for optimism: In this round, the percentage use of the automated mode is highest and at the same time the average profits are highest, while the standard deviation is lowest, compared to all other payoff relevant rounds with normal market conditions of both treatments. The FBA subjects used automated more often and therefore had more time to learn. This is also evidenced, for example, by the correct use of the inventory slider. The FBA subjects have on average a 1.9 times higher setting of the slider at the

Table 6: FBA Performance

rounds		deduction_paid			payoff			
rounds	max	mean	min	std	max	mean	min	std
1	90.6	40.27	15.1	28.11	197.3	59.06	-128.81	126.81
2	51.5	18.88	0.0	18.9	112.6	68.62	15.2	33.96
3	24.6	16.4	0.0	10.37	119.5	74.85	44.9	25.46
4	21.6	9.0	0.0	8.13	110.5	89.33	73.7	14.76
5	18.9	12.6	0.0	9.76	105.6	21.57	-41.4	62.01
6	47.6	23.8	6.8	14.1	123.7	-37.88	-108.1	87.8
8	94.4	51.13	0.0	40.13	410.8	-174.43	-415.0	323.5
9	20.8	12.13	0.0	10.23	128.2	90.62	35.5	33.82

end of the trading period compared to the round average, while there is no difference for CDA subjects. Thus, it is conceivable to assume that with more time to learn, subjects could have reached an equilibrium that yields higher profits than ECU 100 on average. Unfortunately, the FBA players are unable to maintain the level and achieve lower profits in subsequent rounds, possibly due to increased competition. The current value of the out option can certainly slow down the learning process a lot, but at this point, there is no direct answer to whether an increase in profit opportunities will change this, or only benefit the players who are already successful. In addition, fixing the bugs, as well as introducing the new video instructions and comprehension questions could accelerate learning processes significantly. Therefore, no further changes will be made for now, but we will keep a close eye on the average payoffs in the next pilots.

It would be interesting to know if certain players have already achieved an average payoff above 100 ECU across the experiment. Unfortunately, in Pilot 3 there was not yet a way to track a subject's performance across rounds. The current software allows this through a so-called *participant code*. Finally, on External Feed Sensitivity, the slider is set higher under CDA as suspected with an average usage of 0.4 compared to only 0.14 under FBA. The following section discusses the comprehension measure that will be used in the upcoming experiment.

5.3.4 Comprehension Questions

Figure 22 shows a list of possible comprehension questions. In the experiment, each topic in this figure is presented on a separate page. As described by Schweitzer (2017), the

questions start from the top down with simple questions about basic concepts and move to more complicated logical questions with numerical examples. The comprehension questions are not only intended to test the level of understanding, but also fulfill a teaching function. Therefore, if a question is answered incorrectly twice, the experiment software displays the correct solution, which must be confirmed before the participant can proceed to the next page. In the experiment, answering each question is mandatory. As described in the literature, the number of trials, as well as the difficulty of the question, are taken into account to calculate a quantitative comprehension measure as a control variable for the analysis.

Both, the aggregate score, as well as subjects relative rank within their respective group will be recorded. The latter may be more relevant as a control variable, since subjects compete within a group for a limited amount of profit opportunities. That is, the understanding of the market rules of the other subjects within your market is likely to be crucial for your performance. The last section identifies further challenges as well as initial ideas for possible future research questions.

6 Conclusion

The new framework for real-time trading has a strong exploratory character, such that the results are only valid in the interplay of experiment, simulations, empirical observations and insights provided by professional financial market participants. The insights that the joint project of Peter Cramton, Daniel Friedman, Kristian Lopez Vargas, Axel Ockenfels and Mark Marnier-Hausen will provide are relevant to policy makers only to the extent that the new framework is capable of capturing the financial market mechanisms relevant to high-frequency trading. However, the experiment software developed by the LEEPs lab is among the most sophisticated in our research area, creating a high proximity to the real world. Most importantly, laboratory studies are currently the only way to obtain data on the functioning of alternative financial market designs, as no empirical observations are available. Even though the IEX market design has already been approved and FBA may soon be introduced on dark pools, the associated observations are insufficient to provide insights into what would happen, if the interventions were to be introduced on all markets around the world simultaneously.

There are certainly some drawbacks to this study. Most importantly, its complexity makes it challenging to human subjects in the laboratory. That said, as explained above, many design choices were made to make the platform as easy to understand as possible. Moreover, it has been shown that the behavior of laypersons and experts in experimental studies does in general not differ significantly (Fréchette, 2011), and indeed, *ex ante*, we see no reason why, for example, our laboratory subjects would react qualitatively different to extreme stress in their markets than financial market experts. Moreover, to test for robustness of our results, it would be conceivable to conduct the experiment with financial experts only. However, it remains to be seen how our subjects will handle our platform.

If the creation of a sufficiently workable version of the new framework for real-time trading is successful, it will lay the foundation for the comparison of various emerging and exciting new market design approaches. These promising approaches provide most desirable properties beyond those of FBA and IEX, in some cases they come even at lower implementation costs. A natural candidate for future research in this regard is the *Flow Trading* market design proposed by Budish et al. (2020). In addition, future research could also follow the lead off Breckenfelder (2019), who considers different levels of competition based on the number of HFTs active in the market. The author finds that as competition increases under CDA, the ratio of sniping to liquidity-providing behavior increases. The new framework could be used to examine whether the same undesirable behavior holds with increased competition under other market designs. This need not be the case; it is conceivable, for instance, that increased competition under FBA would lead to greater liquidity and thus improved market quality. Another treatment that has not yet been considered, is a "boom phase". Intuitively, a boom phase should be equally stressful for HFTs as a market crash. If HFTs hold negative inventory, i.e., have shorted the asset, and there is not enough liquidity on the opposite market side to balance their inventory, they could book similar losses during a boom phase as is possible during a crash. In reality, this type of stressful market conditions have been observed, for instance, following the recent *GameStop* (GME) frenzy in early 2021 (Umar et al., 2021). During this frenzy professional trading companies such as *Melvin Capital Management LP* indeed suffered losses amounting to roughly \$4 bn (Chung, 2021). However, if the boom phase is accompanied by higher volatility and/or higher volume, it could be very profitable for

HFTs instead.

In conclusion, a working version of the new framework for real-time trading would provide exciting opportunities, the pursuit of which may not only increase external validity, but also provide novel insights for future financial market design and policy-making.

References

- (Oct 3. 2019a). The rise of the financial machines. *The Economist*. Retrived from <https://www.economist.com/leaders/2019/10/03/the-rise-of-the-financial-machines>.
- (Oct 5. 2019b). The stockmarket is now run by computers, algorithms and passive managers. *The Economist*. Retrived from <https://www.economist.com/briefing/2019/10/05/the-stockmarket-is-now-run-by-computers-algorithms-and-passive-managers>.
- Aldrich, E. M., Demirci, H. A., and Vargas, K. L. (2020). An otree-based flexible architecture for financial market experiments. *Journal of Behavioral and Experimental Finance*, 25, page 100205.
- Aldrich, E. M. and Friedman, D. (2022). Order protection through delayed messaging. *Management Science*.
- Aldrich, E. M., Grundfest, J., and Laughlin, G. (2017). The flash crash: A new deconstruction. *Available at SSRN 2721922*.
- Aldrich, E. M. and Vargas, K. L. (2019). Experiments in high-frequency trading: Comparing two market institutions. *Experimental Economics*, pages 1–31.
- Anand, A. and Venkataraman, K. (2013). Should exchanges impose market maker obligations. *Available at SSRN 2179259*.
- Anderson, L., Andrews, E., Devani, B., Mueller, M., and Walton, A. (2021). Speed segmentation on exchanges: Competition for slow flow. *Journal of Financial Markets*, page 100632.
- Aoyagi, J. (2019). Speed choice by high-frequency traders with speed bumps.
- Aquilina, M., Budish, E. B., and O’Neill, P. (2020). Quantifying the high-frequency trading “arms race”: A simple new methodology and estimates. *Chicago Booth Research Paper*, 20-16.
- Aquilina, M., Foley, S., O’Neill, P., and Ruf, T. (2016). Asymmetries in dark pool reference prices. *FCA Occasional Paper*, 21.

- Ariely, D., Ockenfels, A., and Roth, A. E. (2005). An experimental analysis of ending rules in internet auctions. *RAND Journal of Economics*, pages 890–907.
- Baker, N. and Gruley, B. (Mar 8, 2019). The gazillion-dollar standoff over two high-frequency trading towers. *Bloomberg Businessweek*. Retrived from <https://www.bloomberg.com/news/features/2019-03-08/the-gazillion-dollar-standoff-over-two-high-frequency-trading-towers>.
- Baldauf, M. and Mollner, J. (2020). High-frequency trading and market performance. *The Journal of Finance* , 75.3, pages 1495–1526.
- Bali, T. G., Cakici, N., and Whitelaw, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of financial economics*, 99.2, pages 427–446.
- Baron, M., Brogaard, J., Hagströmer, B., and Kirilenko, A. (2019). Risk and return in high-frequency trading. *Journal of Financial and Quantitative Analysis*, 54.3, pages 993–1024.
- Biais, B., Foucault, T., et al. (2014). Hft and market quality. *Bankers, Markets & Investors*, 128, pages 5–19.
- Bouchaud, J.-P., Farmer, J. D., and Lillo, F. (2009). How markets slowly digest changes in supply and demand. In *Handbook of financial markets: Dynamics and evolution*, pages 57–160. Elsevier.
- Breckenfelder, J. (2019). Competition among high-frequency traders, and market quality.
- Brogaard, J. and Garriott, C. (2019). High-frequency trading competition. *Journal of Financial and Quantitative Analysis*, 54.4, pages 1469–1497.
- Brogaard, J., Hendershott, T., and Riordan, R. (2014). High-frequency trading and price discovery. *The Review of Financial Studies*, 27.8, pages 2267–2306.
- Brolley, M. and Cimon, D. A. (2020). Order-flow segmentation, liquidity, and price discovery: The role of latency delays. *Journal of Financial and Quantitative Analysis*, 55.8, pages 2555–2587.

- Brunnermeier, M. K. and Morgan, J. (2010). Clock games: Theory and experiments. *Games and Economic Behavior*, 68.2, pages 532–550.
- Budish, E., Cramton, P., Kyle, A. S., Lee, J., and Malec, D. (2020). Flow trading. Technical report, University of Cologne Working Paper.
- Budish, E., Cramton, P., and Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130.4, pages 1547–1621.
- Budish, E., Lee, R. S., and Shim, J. J. (2019). *Will the Market Fix the Market?: A Theory of Stock Exchange Competition and Innovation*. National Bureau of Economic Research.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, pages 88–97.
- Chen, H., Foley, S., Goldstein, M. A., and Ruf, T. (2017). The value of a millisecond: Harnessing information in fast, fragmented markets. *Fragmented Markets (November 18, 2017)*.
- Chung, J. (Jan 25, 2021). Citadel, point72 to invest 2.75 billion dollars into melvin capital management. *The Wall Street Journal*. Retrived from <https://www.wsj.com/articles/melvin-capital-lost-53-in-january-hurt-by-gamestop-and-other-bets-11612103117>.
- Ding, S., Hanna, J., and Hendershott, T. (2014). How slow is the NBBO? A comparison with direct exchange feeds. *Financial Review*, 49.2, pages 313–332.
- Du, S. and Zhu, H. (2017). What is the optimal trading frequency in financial markets? *The Review of Economic Studies*, 84.4, pages 1606–1651.
- Eisler, Z., Bouchaud, J.-P., and Kockelkoren, J. (2012). The price impact of order book events: Market orders, limit orders and cancellations. *Quantitative Finance*, 12.9, pages 1395–1419.
- Fette, I. and Melnikov, A. (2011). The websocket protocol.

- Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10.2, pages 171–178.
- Fox, M. B., Glosten, L. R., and Rauterberg, G. V. (2015). The new stock market: Sense and nonsense. *Duke LJ*, 65, page 191.
- Fréchette, G. R. (2011). Laboratory experiments: Professionals versus students. *Available at SSRN 1939219*.
- Friedman, D. (1993). How trading institutions affect financial market performance: Some laboratory evidence. *Economic inquiry*, 31.3, pages 410–435.
- Friedman, D., Harrison, G. W., and Salmon, J. W. (1984). The informational efficiency of experimental asset markets. *Journal of political Economy*, 92.3, pages 349–408.
- Glosten, L. R. (1987). Components of the bid-ask spread and the statistical properties of transaction prices. *The Journal of Finance*, 42.5, pages 1293–1307.
- Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics*, 14.1, pages 71–100.
- Guilbaud, F. and Pham, H. (2013). Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13.1, pages 79–94.
- Hagströmer, B. and Menkveld, A. J. (2019). Information revelation in decentralized markets. *The Journal of Finance*, 74.6, pages 2751–2787.
- Hagströmer, B. and Nordén, L. (2013). The diversity of high-frequency traders. *Journal of Financial Markets*, 16.4, pages 741–770.
- Hasbrouck, J. and Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*, 16.4, pages 646–679.
- Hendershott, T., Jones, C. M., and Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? *The Journal of finance*, 66.1, pages 1–33.
- Herrmann, S., Muhle-Karbe, J., Shang, D., and Yang, C. (2020). Inventory management for high-frequency trading with imperfect competition. *SIAM Journal on Financial Mathematics*, 11.1, pages 1–26.

- Hughes, P., Ritter, J., and Zhang, H. (2017). Limit up-limit down pilot plan and extraordinary transitory volatility. *US Securities and Exchange Commission*.
- Jones, C. M. (2013). What do we know about high-frequency trading? *Columbia Business School Research Paper*, pages 13–11.
- Jovanovic, B. and Menkveld, A. J. (2016). Middlemen in limit order markets. *Available at SSRN 1624329*.
- Kalda, A., Loos, B., Previtiero, A., and Hackethal, A. (2021). Smart (phone) investing? A within investor-time analysis of new technologies and trading behavior. Technical report, National Bureau of Economic Research.
- Khapko, M. and Zoican, M. (2020). Do speed bumps curb low-latency investment? Evidence from a laboratory market. *Journal of Financial Markets*, page 100601.
- Kirilenko, A., Kyle, A. S., Samadi, M., and Tuzun, T. (2017). The flash crash: High-frequency trading in an electronic market. *The Journal of Finance*, 72.3, pages 967–998.
- Korajczyk, R. A. and Murphy, D. (2019). High-frequency market making to large institutional trades. *The Review of Financial Studies*, 32.3, pages 1034–1067.
- Królak, A. and Strumiłło, P. (2012). Eye-blink detection system for human–computer interaction. *Universal Access in the Information Society*, 11.4, pages 409–419.
- Kumar, A. (2009). Who gambles in the stock market? *The Journal of Finance*, 64.4, pages 1889–1933.
- Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, pages 1315–1335.
- Kyle, A. S. and Lee, J. (2017). Toward a fully continuous exchange. *Oxford Review of Economic Policy*, 33.4, pages 650–675.
- Laughlin, G., Aguirre, A., and Grundfest, J. (2014). Information transmission between financial markets in Chicago and New York. *Financial Review*, 49.2, pages 283–312.

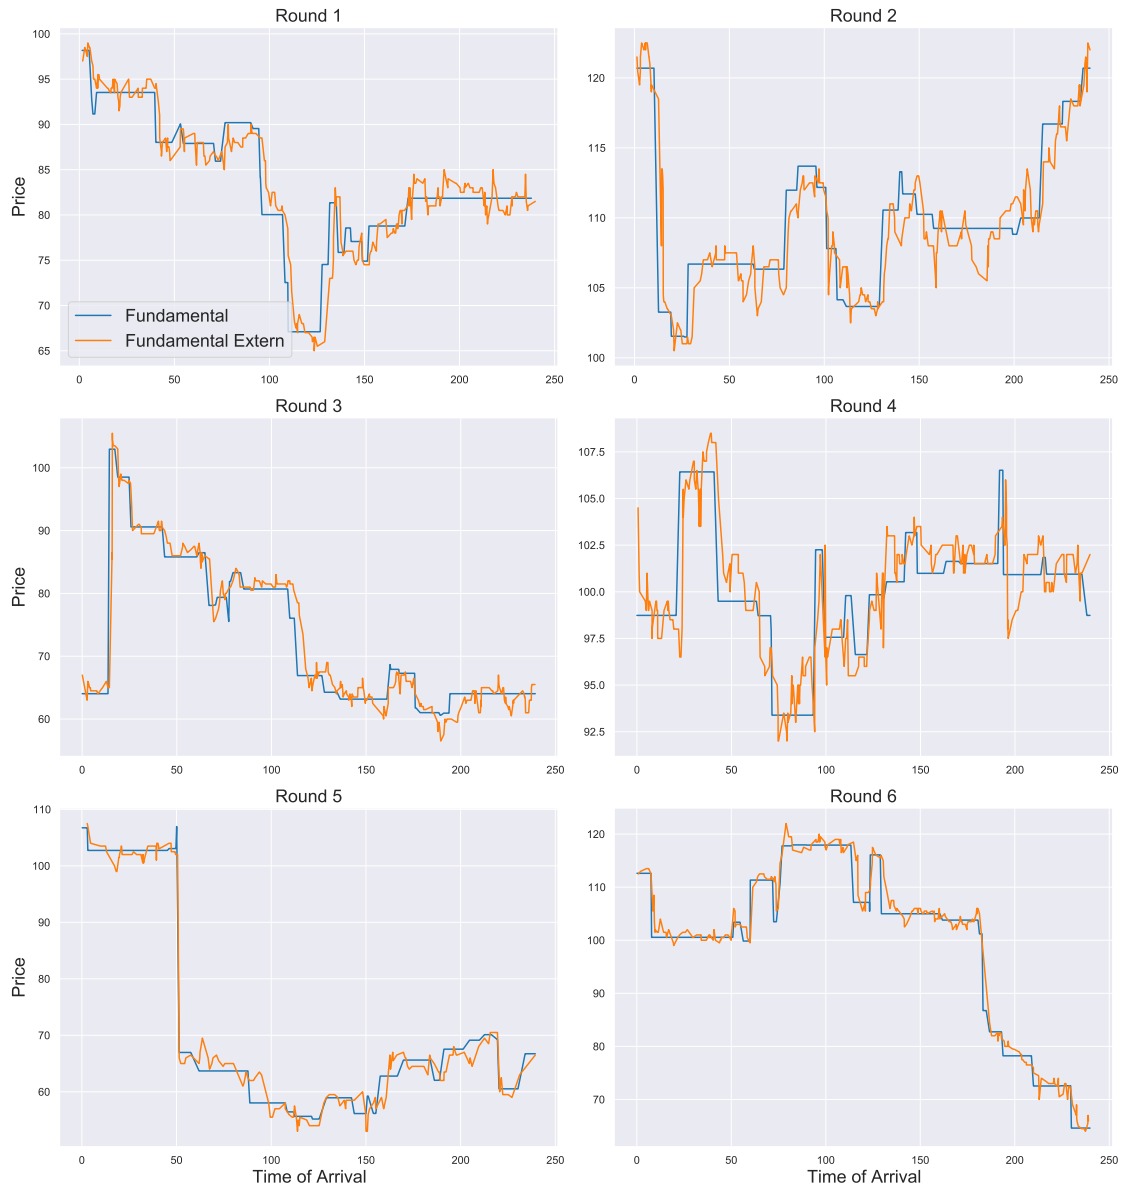
- Lei, V., Noussair, C. N., and Plott, C. R. (2001). Nonspeculative bubbles in experimental asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica*, 69(4):831–859.
- Lockwood, J. W., Gupte, A., Mehta, N., Blott, M., English, T., and Vissers, K. (2012). A low-latency library in FPGA hardware for high-frequency trading (HFT). In *2012 IEEE 20th annual symposium on high-performance interconnects*, pages 9–16. IEEE.
- Mayer, R. E. and Fiorella, L. (2014). 12 principles for reducing extraneous processing in multimedia learning: Coherence, signaling, redundancy, spatial contiguity, and temporal contiguity principles. In *The Cambridge handbook of multimedia learning*, volume 279. Cambridge University Press New York, NY.
- Menkveld, A. J. (2013). High frequency trading and the new market makers. *Journal of financial Markets*, 16.4, pages 712–740.
- Menkveld, A. J. (2016). The economics of high-frequency trading: Taking stock. *Annual Review of Financial Economics*, 8, pages 1–24.
- Menkveld, A. J. and Zoican, M. A. (2017). Need for speed? Exchange latency and liquidity. *The Review of Financial Studies*, 30.4, pages 1188–1228.
- Moinas, S. and Pouget, S. (2013). The bubble game: An experimental study of speculation. *Econometrica*, 81.4, pages 1507–1539.
- Nasdaq Trader (April 8, 2021). Ouch version 4.2. Retrived from <https://www.nasdaqtrader.com/Trader.aspx?id=OUCH>.
- O’Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics*, 116.2, pages 257–270.
- Ozik, G., Sadka, R., and Shen, S. (2020). Flattening the illiquidity curve: Retail trading during the covid-19 lockdown. *Available at SSRN 3663970*.
- Plott, C. R. and Sunder, S. (1988). Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica: Journal of the Econometric Society*, pages 1085–1118.

- Popper, N. (Aug 2, 2012). Knight capital says trading glitch cost it 440 million dollar. *The New York Times*. Retrived from <https://dealbook.nytimes.com/2012/08/02/knight-capital-says-trading-mishap-cost-it-440-million/>.
- Rogow, G. (Jun 19, 2009). Rise of the (market) machines. *The Wall Street Journal*. Retrived from <https://www.wsj.com/articles/BL-MB-8764>.
- Rogow, G. (Sep 20, 2012). Colocation: The root of all high-frequency trading evil? *The Wall Street Journal*. Retrived from <https://www.wsj.com/articles/BL-MB-44124>.
- Roth, A. E. and Ockenfels, A. (2002). Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. *American economic review*, 92.4, pages 1093–1103.
- Roth, A. E. and Xing, X. (1994). Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review*, pages 992–1044.
- Roth, A. E. and Xing, X. (1997). Turnaround time and bottlenecks in market clearing: Decentralized matching in the market for clinical psychologists. *Journal of political Economy*, 105.2, pages 284–329.
- Schweitzer, S. (2017). Cognitive-based design principles in experimental economics: How to devise more effective instructions, software, and comprehension tests. *Software, and Comprehension Tests (May 8, 2017)*.
- Shkilko, A. and Sokolov, K. (2016). Every cloud has a silver lining: Fast trading, microwave connectivity and trading costs. *Journal of Finance, Forthcoming*.
- Smith, V. L., Suchanek, G. L., and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society*, pages 1119–1151.
- Stoll, H. R. (1989). Inferring the components of the bid-ask spread: Theory and empirical tests. *the Journal of Finance*, 44.1, pages 115–134.
- Tong, L. (2014). A blessing or a curse? The impact of high frequency trading on institutional investors. In *The Impact of High Frequency Trading on Institutional Investors (October 5, 2015)*. *European Finance Association Annual Meetings*.

- Umar, Z., Yousaf, I., and Zaremba, A. (2021). Comovements between heavily shorted stocks during a market squeeze: Lessons from the GameStop trading frenzy. *Research in International Business and Finance*, page 101453.
- U.S. Securities and Exchange Commission (2010). Findings regarding the market events of May 6, 2010. Retrived from <https://www.sec.gov/oiea/investor-alerts-bulletins/investor-alerts-circuitbreakersbulletinhtm.html>.
- Valetkevitch, C. and Mikolajczak, C. (Aug 1, 2012). Error by knight capital rips through stock market. *Reuters*. Retrived from <https://www.reuters.com/article/us-usa-nyse-tradinghalts-idUSBRE8701BN20120801>.
- Van Kervel, V. and Menkveld, A. J. (2019). High-frequency trading around large institutional orders. *The Journal of Finance*, 74.3, pages 1091–1137.
- Wah, E. (2016). How prevalent and profitable are latency arbitrage opportunities on us stock exchanges? *Available at SSRN 2729109*.
- Welch, I. (2020). The wisdom of the robinhood crowd. *NBER Working Paper*, (27866).
- Zhang, S. and Riordan, R. (2011). Technology and market quality: The case of high frequency trading.
- Zuckerman, G., Levy, R., Timiraos, N., and Banerji, G. (Dec 25, 2018). Behind the market swoon: The herdlike behavior of computerized trading. *The Wall Street Journal*. Retrived from <https://www.wsj.com/articles/behind-the-market-swoon-the-herdlike-behavior-of-computerized-trading-11545785641>.

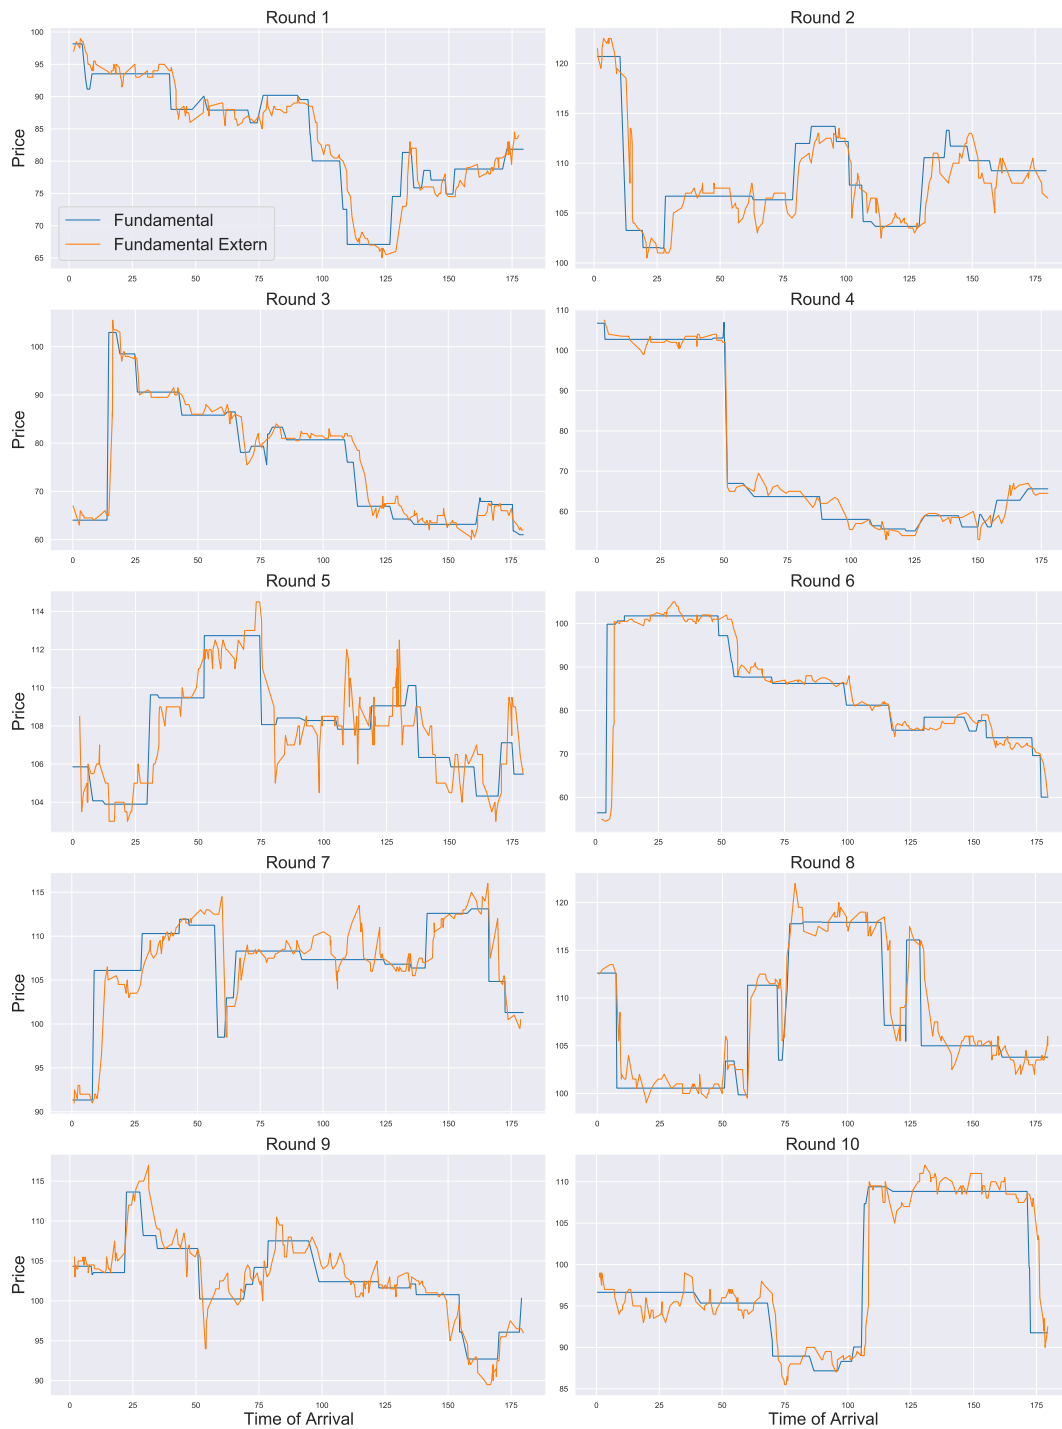
Appendices

Figure 13: Fundamental values of Pilot 2



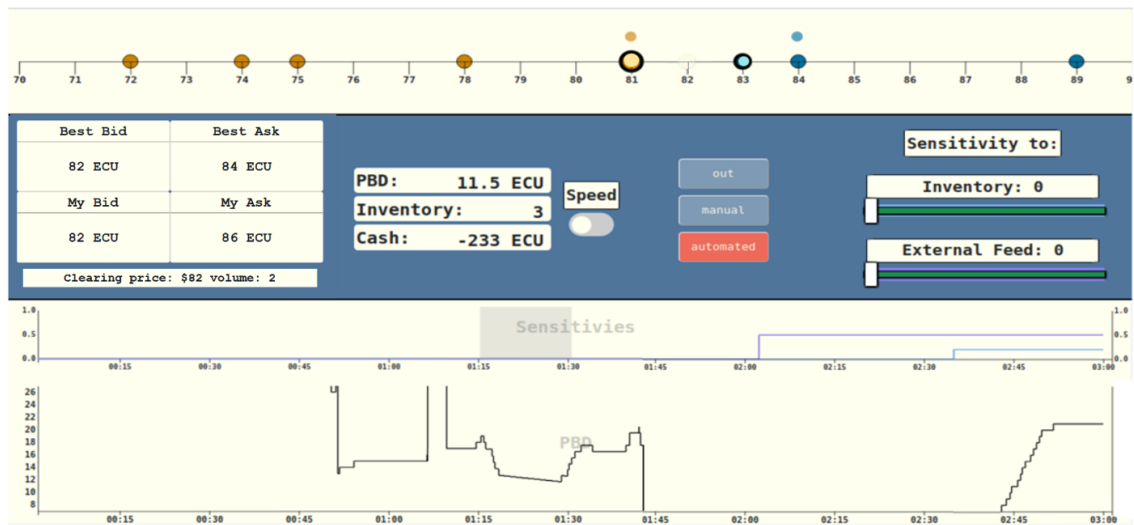
This figure captures the artificial fundamental values considered in Pilot 2. Recall that the pilot ended prematurely after six rounds, which is why this figure only contains the fundamental values that were actually considered. The blue line depicts the fundamental value on the subjects' market, whereas the orange line depicts the fundamental value on the external market. The stress conditions $S3$ and $S4$ are depicted in Round 5 and 6, respectively.

Figure 14: Fundamentals values of Pilot 3



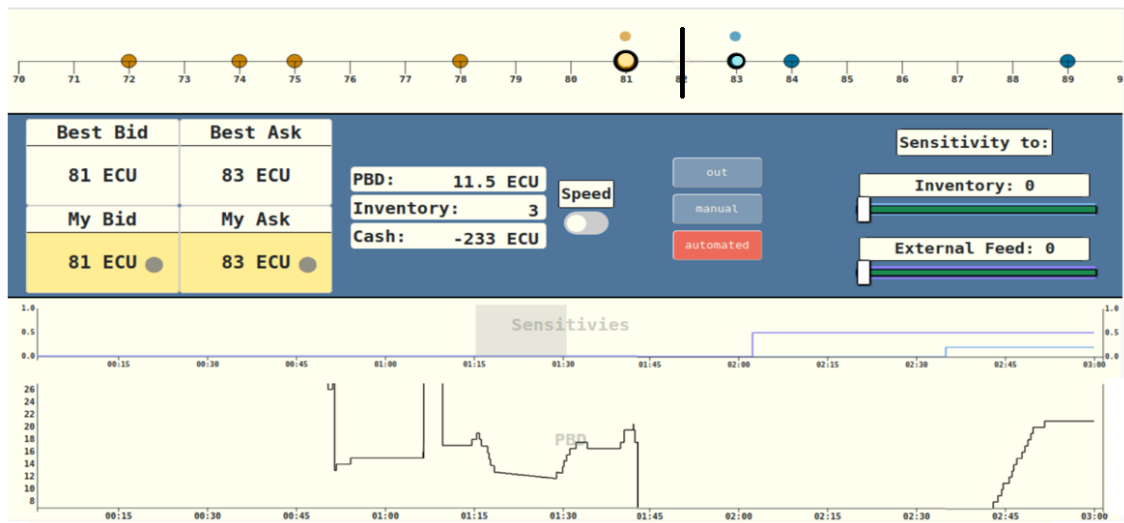
This figure captures the artificial fundamental values planned for Pilot 3. The fundamental values for the two trial rounds are not displayed, only for the rounds that are relevant to the subjects' payoffs. So far, S_3 and S_4 are depicted in Round 4 and 8, respectively. This is currently under discussion and will likely change prior to the pilot. Stress conditions could be considered in Rounds 6 and 7 to give subjects sufficient time before and after to converge to a potential equilibrium.

Figure 15: Subjects' User-Interface FBA



The FBA user interface contains an additional field on the right side of the middle panel, below the table, which displays the best bid, best ask, my bid and my ask. This field contains the *clearing price* and the *volume*. The former corresponds to the market clearing price of the most recent end-of-batch auction in which aggregate demand and supply intersected. The latter indicates the volume of this auction, i.e. the amount of assets traded. This new field flashes whenever an auction takes place to indicate to subjects the end of a batch. This works similar to the red flashing dots that indicate the transaction of the subject's own bid or ask.

Figure 16: Subjects' User-Interface IEX



The key distinguishing feature of the IEX user interface is the black vertical line in the horizontal price line in the upper panel. This vertical line is based on the following consideration: According to a random selection, the bids of investors will be lit in 50% of the time, and pegged in 50% of the time. The prices of pegged orders are updated by the experimental exchange server itself, such that they always correspond to the midpoint between BBext and BAext, graphically indicated by the black vertical line. As pegged orders are hidden, subjects are not able to observe whether bids or asks are placed at the black line.

Figure 17: Subsequent Questionnaire of Pilot 2

CLER

Date 18.12.2020

Questionnaire HFT-Experiment

How far advanced are you in your studies?

Wählen Sie ein Element aus.

Do you have a background in economics?

Wählen Sie ein Element aus.



1. Did you find parts of the instructions or user interface confusing? If yes, which ones?

2. Did you follow a strategy? If yes, which one?

4. Has your strategy changed over time? Why?

We highly appreciate any further comments that might improve our experiment and subjects' experience.

Figure 18: Subsequent Questionnaire of Pilot 3

Teil A: Fragebogen

A1. In welchem Studium befinden Sie sich?

Bachelor

Master

Ph.D.

A2. Hat Ihr Studium einen wirtschaftswissenschaftl. Hintergrund?

Ja

Nein

A3. Wie benoten Sie die Klarheit der Instruktionen?
Die Noten basieren auf der Ihnen bekannten Skala von 1 (sehr gut) bis 5 (mangelhaft).

1

2

3

4

5

A4. Fanden Sie Teile der Instruktionen verwirrend? Falls ja, welche?

A5. Wie benoten Sie die Klarheit der Benutzeroberfläche?
Die Noten basieren auf der Ihnen bekannten Skala von 1 (sehr gut) bis 5 (mangelhaft).

1

2

3

4

5

Subsequent Questionnaire of Pilot 3 p. 2



A6. Fanden Sie Teile der Benutzeroberfläche verwirrend? Falls ja, welche?

A7. Wie benoten Sie Ihr Verständnis des Experiments?

Die Noten basieren auf der Ihnen bekannten Skala von 1 (sehr gut) bis 5 (mangelhaft).

1

2

3

4

5

A8. Fanden Sie Teile des Experiments verwirrend? Falls ja, welche?

A9. Falls Sie eine Strategie verfolgt haben, wie erfolgreich war diese?

Die Noten basieren auf der Ihnen bekannten Skala von 1 (sehr gut) bis 5 (mangelhaft).

1

2

3

4

5

A10. Welche Strategie haben Sie verfolgt?

Subsequent Questionnaire of Pilot 3 p. 3



A11. Wie häufig hat sich Ihre Strategie mit der Zeit verändert?

Auf einer Skala von 1 (nie) bis 5 (sehr häufig).

- 1
- 2
- 3
- 4
- 5

A12. Falls Sie Ihre Strategie verändert haben, warum?

A13. Wir freuen uns noch über jegliche Hinweise, die unser Experiment und die Erfahrung der Versuchspersonen verbessern könnten.

A14. Sind Sie auf Fehler gestoßen, oder hatten Sie das Gefühl, dass das Experiment irgendwann nicht mehr richtig funktioniert hat? Wir würden uns sehr über eine kurze Beschreibung oder das Hochladen von Fehler-Screenshots freuen.

Beschreibung:

A15. Datei 1 (optional):

A16. Datei 2 (optional):

Subsequent Questionnaire of Pilot 3 p. 4



A17. Datei 3 (optional):

A18. Im Folgenden können Sie Ihre Bildschirmaufzeichnung des heutigen Experiments hochladen. Dafür können bis zu 2 Dateien verwendet werden.

Hinweis: Bitte warten Sie bis die Datei vollständig hochgeladen ist, bevor Sie die Umfrage beenden.

Datei 1 (optional):

A19. Datei 2 (optional):

A20. Wenn Probleme mit dem Upload bestehen, haben Sie hier auch die Möglichkeit die Datei über Sciebo, Dropbox oder Ähnliches zu teilen. Tragen Sie dafür bitte im folgenden Feld den Link ein.

Teil B: Endseite

Hiermit ist die Umfrage beendet. Durch das Betätigen der Schaltfläche "Absenden" können Sie ihre Antworten an uns übermitteln.

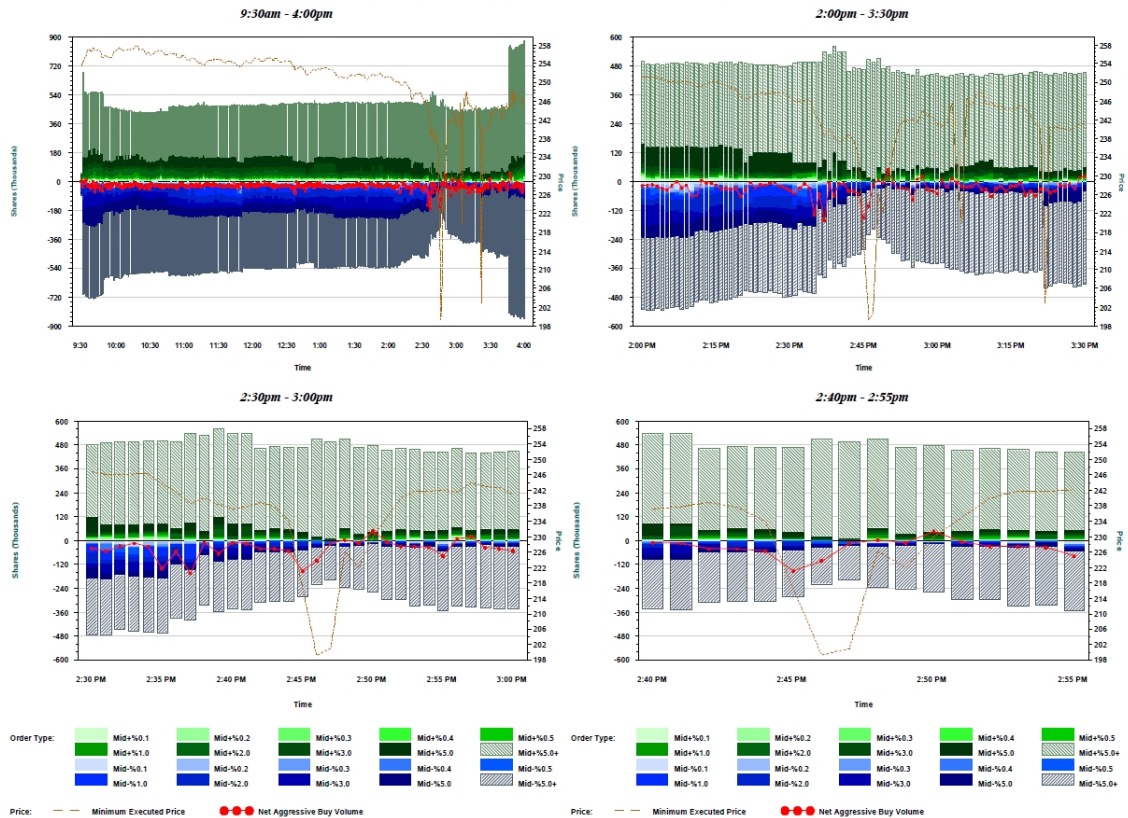
Subsequent Questionnaire of Pilot 3 p. 5



Vielen Dank für ihre Teilnahme an diesem Fragebogen. Sie können nun zu ihrem Experimentfenster zurückkehren und die „Next“-Schaltfläche betätigen.

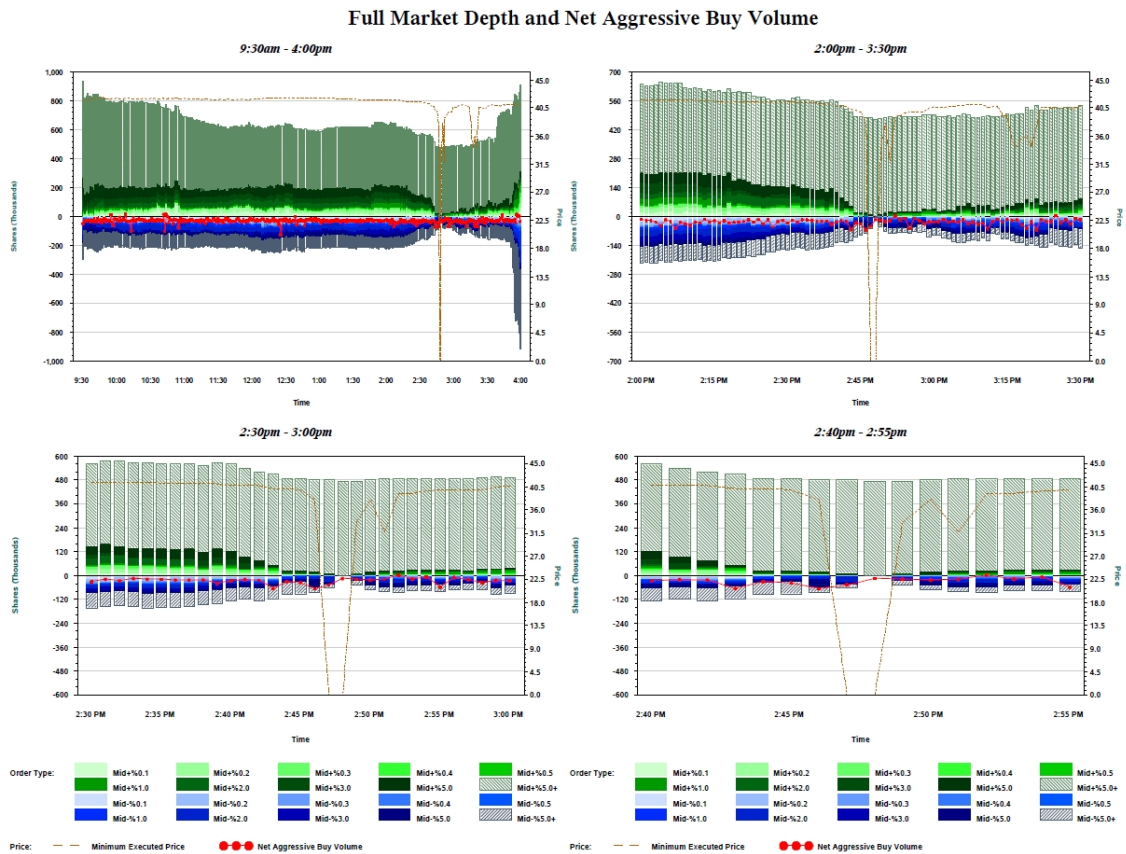
Figure 19: Market Depth and Price Movement of AAPL on May 6, 2010

Full Market Depth and Net Aggressive Buy Volume



This figure is taken from the U.S. Securities and Exchange Commission (2010) report. Market depth is calculated by aggregating outstanding limit orders at the minute level. The blue bars show the market depth on the buy side, while the green bars show the market depth on the sell side. The brightness of colors indicates the distance from the midpoint. Four different time periods are considered, with the peak of the Flash Crash lying between 14:40 and 14:55. The dashed yellow lines show the minimum transaction price per minute. The dotted red line represents the net aggressive buy volume. Aggressive orders are market orders or limit orders at or below (above) the BB (BA).

Figure 20: Market Depth and Price Movement of ACN on May 6, 2010



This figure is taken from the U.S. Securities and Exchange Commission (2010) report. Market depth is calculated by aggregating outstanding limit orders at the minute level. The blue bars show the market depth on the buy side, while the green bars show the market depth on the sell side. The brightness of colors indicates the distance from the midpoint. Four different time periods are considered, with the peak of the Flash Crash lying between 14:40 and 14:55. The dashed yellow lines show the minimum transaction price per minute. The dotted red line represents the net aggressive buy volume. Aggressive orders are market orders or limit orders at or below (above) the BB (BA).

Figure 21: Version of the instructions prior to Pilot 3. The differences to the baseline treatment CDA are marked in light blue.

CDA/HFT Sept 2019

Instructions for Participants

Welcome, and thank you for participating in this experiment. For better readability, the language forms male, female and diverse (m/f/d) are not used simultaneously. All references to persons apply equally to all genders. From now until the end of the experiment, please turn off your cell phone and do not communicate with other participants. If you have any questions, please raise your hand; an experimenter will come to answer your question. Please pay careful attention to the instructions as real money is at stake.

The Basic Idea

You and 5 other randomly-chosen participants will be traders in a virtual financial market. You will be able to buy and sell shares of a stock (also referred to as "units") that exist only in this experiment. Trading takes place over 10 periods, each lasting 3 minutes.

You make profits each trading period by buying units at prices lower than the prices at which you sell them, and you make losses if your purchase price is higher than your selling price. Some of the trading opportunities will come from the other human traders in your market, and some from fully-automated agents called *investors*. The trading activities of the automated investors take place at random times and are independent of the other market transactions. The trading patterns of automated investors might change over time.

Trading profits are denominated in the lab currency: Experimental Currency Units (ECUs). You will begin each trading period with a cash account of 100 ECUs and an inventory of 0 units. Each time you buy [or sell] a unit, your inventory increases [or decreases] by 1, while your cash account decreases [or increases] by the price of that unit. If the number of units you buy differs from the number of units you sell, you will have either a positive or negative inventory at the end of the period. A negative inventory is a result of "short selling", that is, you sell units that you do not own. Your end-of-period inventory will be converted to ECUs following a rule explained below and will become part of your payoff, along with your trading profits.

At the end of the experiment, your payoff in ECUs of each period will be added up and converted to Euro and paid to you in cash. The exchange rate is 100 ECU = XX EUR. Additionally, you will receive 4 EUR for your participation.

Trading Rules [CDA]

At any moment during a trading period, you, the other human participants, and the automated investors can make offers to buy (called **bids**) and offers to sell (called **asks**). For example, if you submit a **bid** at 99, that means you are willing to pay at most 99 ECU to buy a unit of the stock. Similarly, if you submit an **ask** for 101, that means you are willing to receive 101 or more ECU to sell a unit. To indicate the current state of the market at all times, the computer also calculates and displays the **Best Bid (BB)** and the **Best Ask (BA)**. The **Best Bid** is the **highest bid** currently made by any trader, and the **Best Ask** is the **lowest** current **ask**. For reasons explained below, the **BA** is always higher than the **BB**.

When any trader submits a new **bid** priced at or above the current **BA**, it immediately transacts at the **BA** price. If the new **bid** is priced below the **BA** but above the **BB**, then it becomes the new **BB**. If the new **bid** is at or below the **BB**, then it enters the **order book**, a list of unexecuted **bids**, and **asks** sorted by price and time, behind other **bids** at that price. For example, suppose that the **BA** (held by some other trader) is 102. If you submit a **bid** at 105, then you immediately buy a unit at price 102 (better for you than buying at 105). The next **ask** on the order book becomes the new **BA**. If instead, your **bid** was 101, then it would enter the order book, after the existing orders. It would become the new **BB** if (and only if) the current **BB** had been lower than 101.

Similarly, an **ask** submitted at or below the current **BB** is transacted at the **BB**, while **asks** submitted at higher prices enter the order book. For example, suppose that the **BB** is 89 and the **BA** is 91. Then if you submit an **ask** of 89 or less, it is immediately transacted at a price of 89 and the next **bid** in the order book will become the new **BB**. If you choose an **ask** of 90 instead, then your **ask** becomes the new **BA** since it is lower than the existing **BA** of 91. If your **ask** is 91 or above, it will be added to the order book without changing the **BA** or **BB**.

As you trade, you make profits or losses, and your inventory changes as well. As a simple example, suppose you buy a unit for 99 ECUs and sell a unit for 101 ECUs. Then, your cash position will increase as you made a profit of 2 ECUs (101-99). Also, in this case, since you are buying one unit and selling one unit too, your inventory will remain unchanged.

Trading Rules [FBA]

At any moment during a trading period, you, the other human participants, and the automated investors can make offers to buy (called **bids**) and offers to sell (called **asks**). For example, if you submit a **bid** at 99, that means you are willing to pay at most 99 ECU to buy a unit of the stock. Similarly, if you submit an **ask** for 101, that means you are willing to receive 101 or more ECU to sell a unit. To indicate the current state of the market at all times, the computer also calculates and displays the **Best Bid (BB)** and the **Best Ask (BA)**. The **Best Bid** is the **highest bid** currently made by any trader, and the **Best Ask** is the **lowest** current **ask**. For reasons explained below, the **BA** is always higher than the **BB**.

Each trading period is divided into many two-second batch intervals. At the end of each batch, an auction is conducted as follows. All bids and asks submitted in the current batch, enter the order book, a list of

unexecuted bids and asks from previous batches sorted by price and time. At the end of the batch a clearing price is calculated so that the number of units offered in the order book at that price, equals the number of units asked. The computer then accepts all those bids and those asks at the clearing price. The bids and asks that are not accepted are carried over in the order book to the next batch auction.

For example, suppose that you just entered a bid at 99 and an ask at 101 in the most recent batch and the clearing price for this auction turns out to be 98 ECUs. Your ask is above the clearing price and therefore gets carried over to the next batch auction. Your bid of 99 is accepted since it is above the clearing price of 98, so you bought 1 unit at a price of 98 (better for you than buying at 99). Your inventory, therefore, goes up by one unit and your cash account goes down by 98 ECUs. Suppose that you don't change your ask and, at a later batch auction, the clearing price is 101 ECUs. Now your ask is accepted since it is at or below the clearing price, and so you sell one unit. Consequently, your inventory goes down by one unit and your cash account goes up by 101 ECUs. Considering only these two transactions, your inventory has not changed ($1-1=0$) but your cash position has increased by $101-98 = 3$. That is, from those two transactions you have made a profit of 3 ECUs.

As you trade, you make profits or losses, and your inventory changes as well. As a simple example, suppose you buy a unit for 99 ECUs and sell a unit for 101 ECUs. Then, your cash position will increase as you made a profit of 2 ECUs ($101-99$). Also, in this case, since you are buying one unit and selling one unit too, your inventory will remain unchanged.

Trading Rules [IEX]

At any moment during a trading period, you, the other human participants, and the automated investors can make offers to buy (called **bids**) and offers to sell (called **asks**). For example, if you submit a **bid** at 99, that means you are willing to pay at most 99 ECU to buy a unit of the stock. Similarly, if you submit an **ask** for 101, that means you are willing to receive 101 or more ECU to sell a unit. As explained below, all orders by human traders and some automated investor orders will be "lit" and every player can see them on their screen. Some automated investor orders will be hidden in that their prices are not visible to anyone. To indicate the current state of the market at all times, the computer also calculates and displays the **Best Bid (BB)** and the **Best Ask (BA)**. The **Best Bid** is the **highest lit bid** currently made by any trader, and the **Best Ask** is the **lowest lit current ask**. For reasons explained below, the **BA** is always higher than the **BB**.

When any trader submits a new lit bid, four things can happen:

1. If the bid is below the **BA** and any hidden ask, and at or below the **BB**, then it enters the order book, a list of unexecuted bids and asks sorted by price and time, behind other lit bids at that price.
2. If the bid is below the **BA** and any hidden ask, and above the current **BB**, it enters the order book and becomes the new **BB**.
3. If the bid is at or above a hidden ask but below the **BA**, then it transacts immediately with the hidden ask.

4. Otherwise, if the bid is at or above the BA and there exists no hidden ask below the BA, the bid transacts immediately at the BA.

For example, suppose that the BA (held by some other trader) is 103 and there is a hidden ask at 101. If your bid is 100, then it would enter the order book, after the existing orders. It would become a BB if (and only if) the current BB were 99 or less. If instead, you submit a bid at 102, then it would immediately transact with the hidden ask and you would buy a unit at price 101 (better for you than buying at 102). Finally, suppose you submit a bid at 105, then you immediately buy a unit at price 101 (the hidden ask). But if there were no hidden ask, then you would transact at 103 (the current BA).

Similarly, when any trader submits a new lit ask, four things can happen:

1. If the ask is above the BB and any hidden bid, and at or above the BA, then it enters the order book behind other lit asks at that price.
2. If the ask is above the BB and any hidden bid, and below the current BA, it enters the order book and becomes the new BA.
3. If the ask is at or below a hidden bid but above the BB, then it transacts immediately with the hidden bid.
4. Otherwise, if the ask is at or below the BB and there exists no hidden bid above the BB, the ask transacts immediately at the BB.

For example, suppose that the BB is 89 and the BA is 91 with no hidden bids in between. Then if you submit an ask of 89 or less, your bid is immediately transacted at a price of 89 (which is better for you than selling at a lower price) and the next bid in the order book will become the new BB. If you choose an ask of 90 instead, your bid becomes the new BA, as you ask is lower than the existing BA of 91. If your ask is higher than 90, it will be added to the order book without changing the BA or BB.

As you trade, you make profits or losses, and your inventory changes as well. As a simple example, suppose you buy a unit for 99 ECUs and sell a unit for 101 ECUs. Then, your cash position will increase as you made a profit of 2 ECUs (101-99). Also, in this case, since you are buying one unit and selling one unit too, your inventory will remain unchanged.

Interface Layout

Figures 1-3 below show the user interfaces for today's experiment. Figure 1 shows the screen where you enter the initial strategy with which you want to enter the market. This interface is a simplified version of the actual trading interface shown in Figure 2. As soon as all participants in your market enter their initial strategy and select the "Next" button, they are redirected to the trading interface (Figure 2) and the trading period begins.

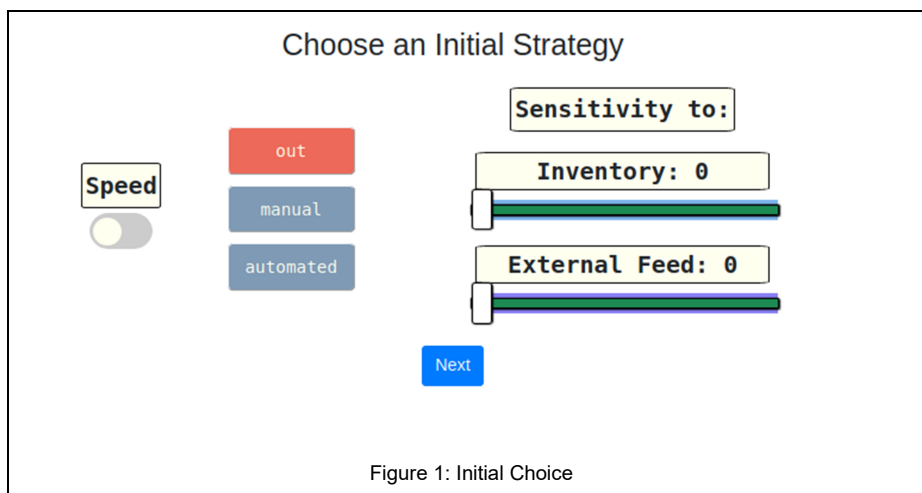


Figure 1: Initial Choice

Let us now examine the main trading interface (Figure 2) in more detail. This screen is active during the 3 minutes of trading and is divided into three boxes at the top, middle, and bottom of your screen:

1. The top box displays the current order book graphically, using a horizontal price line. The blue dots represent **asks** and the orange dots represent **bids**. Larger dots indicate more orders tied at that same price. The lighter blue [orange] dots represent your own **ask** [**bid**]. Black borders represent the **BB** and the **BA**. [IEX: When there is a transaction, the bordered dot will shrink if there are ties at that BB or BA, otherwise, the difference between BA and BB will increase. There is also a vertical black line, a midpoint between specific bids, and asks as explained below.]
2. The middle box displays information and controls. The table at the left repeats, in text, the current **BB** and **BA** and your own **bid** and **ask**. In this table, you will see a dot flashing under "my Bid" or "my Ask" when one of your orders is executed. The three lines to the right of the table show your current cash position in ECUs, your current inventory (in units), and your **payoff before deduction (PBD)**, that you would receive if your inventory were cashed out now, as explained below. The rest of the box contains the control buttons and sliders you will use to make trades, as explained below.
3. The bottom box graphically displays the **history** of your PBD and your chosen controls for trading ("Sensitivity to") for the current period so far. The sensitivities you choose on the right side of the middle box can affect your PBD as explained in the next section.

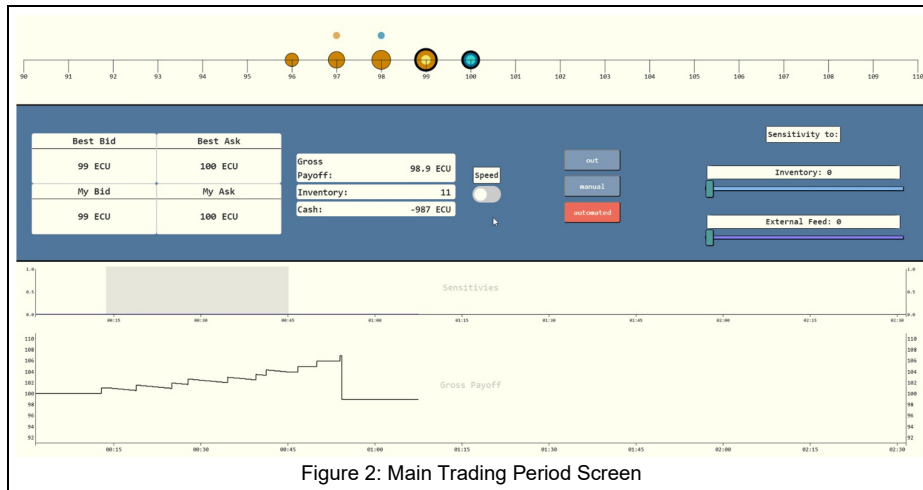


Figure 2: Main Trading Period Screen

Trading Controls

You can choose at any moment between staying **out** of the market (top rectangular button in the middle box), participating in **manual** mode (middle button), or participating in **automated** mode (bottom button). While out, you do not participate in the market and do not earn profits, nor take losses.

In manual mode, you submit **bids** and **asks** directly into the top box using your computer mouse buttons. Find the desired position on the price line and left-click to submit a **bid**, and right-click to submit an **ask**. This gives you direct control over your trades and resulting profits or losses, but it is limited by your reaction time and accuracy. If you decide to enter the trading period in manual mode, you can select and adjust your **bid** and **ask** in the price line as soon as you are redirected to the main trading interface.

In automated mode, you control a bot that submits **bids** and **asks** on your behalf. As soon as your **bid** or **ask** is accepted by another trader [**FBA: is accepted at the end of the batch**], your bot automatically submits a new **bid** or **ask** at the same price. You can fine-tune this process using the sensitivity sliders on the right side of the control box. That is, you can program your bot to react immediately to changes in market conditions and to changes in your own situation.

Here is how your bot works: It looks at the **BB** and **BA** (excluding your own current **bid** and **ask**), and just matches them if the sensitivities are set to zero. When you use the slider to increase your "**inventory sensitivity**", your bot will move your **bid** and **ask** trying to return your inventory level to zero. Suppose you have a positive inventory, then increasing inventory sensitivity will tend to lower your **bid** (making it less likely to buy) and also lower your **ask** (making it more likely to sell). If you have negative inventory, then increasing inventory sensitivity tends to raise your **bid** and raise your **ask** relative to **BB** and **BA**, respectively. That is, increasing your inventory sensitivity will help you sell off positive inventory or buy

units to eliminate negative inventory. This can be especially useful late in the period to avoid the inventory deduction explained below.

The other slider controls your bot's sensitivity to the “**external feed**”. There is an external market with no human traders that responds to essentially the same automated investor orders that reach your own market. That external market has its own **BB** and **BA**, call them **BBext** and **BAext**; these are displayed in Figure 2 in the top box on your screen as a small orange dot and blue dot above the price line. The higher you set the External Feed sensitivity slider, the more weight your bot gives to the **BBext** and **BAext** instead of the **BB** and **BA** of your own market. That is, higher sensitivity to external feed makes your **bid** and **ask** to follow more closely the best **bid** and **ask** in the external market. This sensitivity can be useful because the external market will sometimes (not always) react faster than your own market to changes in patterns of automated investor orders, allowing your bot to act on market information earlier than other players. However, if you pay no attention to your own market, you might miss the times when your own market has reacted faster.

Your **bids** and **asks** are constructed in such a way that you can never trade with yourself. For example, assume that you choose a new **bid** that is higher than your existing **ask**. In this case, your existing **ask** will be automatically deleted, otherwise, you would end up trading with your own **ask**. Similarly, if you choose a new **ask** below your existing **bid**, your existing **bid** will be automatically deleted.

There is one more control button, the “**Speed**” toggle, which is displayed to the left of the controls described so far. Even in automated mode, it takes time for the market to respond to the orders you submit. You can reduce this time by activating the speed option in automated mode (in manual mode, your mouse clicks determine the speed). If you do not activate the speed option (the default), the market will react to the new **bids** and **asks** that your bot submits on your behalf within half a second. If you switch on speed, then it will respond in one-tenth of a second, but you will pay 0.01 ECU per second for the faster service. You can activate or deactivate the speed option at any time. This will be reflected in the History box where the sensitivity graph is shaded in time intervals when speed is activated, and the PBD graph slants down. For example, in Figure 2, speed was ON for 30 seconds in the time interval 00:15 to 00:45, which amounts to a cost of this 0.3 ECUs (30×0.01).

Speed is worth the cost only if it sufficiently improves the prices you get. For example, assume that you (or your bot) expect a price increase, then you want to buy now and sell later. The speed option increases the chances that your **bid** (**ask**) will be updated according to your expectations before the other traders can react. If your (your bot's) belief regarding the price change was right, this enables you to make a higher profit.

[FBA: Speed is worth the cost only if it sufficiently improves the prices you get. For example, if you (or your bot) think that the clearing price in the next batch will be higher than in the current batch, then you want to buy in the current batch and sell in the next. If your (your bot's) belief regarding the price change is right and you have selected the speed toggle, then you will miss the current batch only if your bid is sent with less than a tenth of a second to go in the batch interval. With no speed, it will miss the current batch if it is sent with less than half a second left. Speed makes no difference for orders sent with more than half a second remaining or less than a tenth of a second remaining in the batching interval.]

Notice that in automated mode, you do not need to cancel your orders; when the market conditions change, your bot will automatically replace your previous orders with the new ones. Also, while using automated or manual mode, you can cancel all your orders by pressing the “out” button.

Other Important Details

Automated investors. You and the other human participants can all make positive profits because of the random stream of **bids** and **asks** generated by the automated investors. The automated investors do not react strategically to your decisions or those of the other human participants and behave naive in this sense. Investors' orders often create transactions at the **BB** and **BA** and otherwise will enter the order book.

PBD. At all times during the trading period, the computer generates and shows you your PBD. As this referential represents the payoff before deduction you would hypothetically make if the trading period ended at that moment, it provides you with feedback regarding your performance so far. The PBD is defined as the sum of the balance in your cash account and a referential value of your inventory. The computer calculates the latter by multiplying the positive or negative number of units in your inventory by a reference price (P). The reference price (P) is a weighted average of all previous transaction prices [[FBA: clearing prices](#)], with greater weight on transactions [[FBA: prices](#)] that are more recent.

End-of-period Inventory buyout. If you end the trading period with a positive inventory, the computer buys out your final inventory from you at the final reference price P, calculated at the end of the trading period as described above. Similarly, if you end the trading period with a negative inventory, the computer makes you buy the units you owe at the final reference price P. It follows, that the PBD you observe in the very last second of a trading period, coincides with the payoff before deduction you realize in this period. For example, consider "Your End of Period Report" at the top right of Figure 3: In this example, your cash equals 201 ECUs, the reference price P is 100 ECUs and your inventory is 2 at the end of the period. The computer will sell those 2 units in your name for 100 ECUs each. Note, the speed costs are due as soon as you select the faster service, as described above, and amount to a total of 1 ECU in this example. Thus, your PBD amounts to $201 + 2 \times 100 - 1 = 400$. If, on the other hand, you would have had a negative inventory of -2 at the end of the trading period with the same cash balance and the same P, your PBD would have been: $201 - 2 \times 100 - 1 = 0$.

Deduction. Additionally, to discourage large end-of-period inventories (positive or negative), a deduction payment equal to 10% of your inventory value at the end of the period will be imposed. In the example from Figure 3 the inventory value amounts to 200, of which you will be charged 10%, i.e. $200 \times 0.1 = 20$ ECUs. Similarly, if you had a negative inventory of -2 units, you would have to pay 10% of the absolute of $|-200|$, which also amounts to 20 ECUs.

Payoff. The payoff you realize is equal to your PBD at the end of the period minus the deduction. Thus, in the example in Figure 3, your payoff would be: $400 - 20 = 380$ ECUs.

[IEX: While human participants in the market can only submit the lit orders (bids and asks) described above, automated investors, in contrast, submit hidden pegged orders as well as lit orders. Randomly, 50% of the time, investors will submit pegged orders, and 50% of the time they will submit lit orders. Lit orders have two features: (a) their prices can only be set and updated by the traders themselves or their bots, and (b) lit orders can be seen by everyone on their screen. Pegged orders, in contrast, have these two features: (a) their prices are set by the exchange itself at the midpoint between BBext and BAext (indicated by the black vertical line in the top box on the trading screen); that is, when either the BBext and BAext changes, the exchange itself calculates the new midpoint and updates the price levels of all pegged orders. (b) Pegged orders are hidden in that they cannot be seen by anyone on their screen.

Speed bump. Every new order coming from you, the other human participants or the automated investors, that arrives at the exchange is delayed one-tenth of a second. This ensures that the exchange has enough time to update the prices of pegged orders according to changing market patterns before they can be picked off by fast traders such as yourself.]

Reporting. Figure 3 is displayed to you at the end of each trading period and includes, in addition to "Your End of Period Report", a graphical summary of your payoff and the strategies you have chosen (see top left of Figure 3). The other boxes, sorted from highest to lowest payoff, show similar graphical summaries for the other human participants in your market. Studying these reports may help you find more profitable strategies in subsequent trading periods.

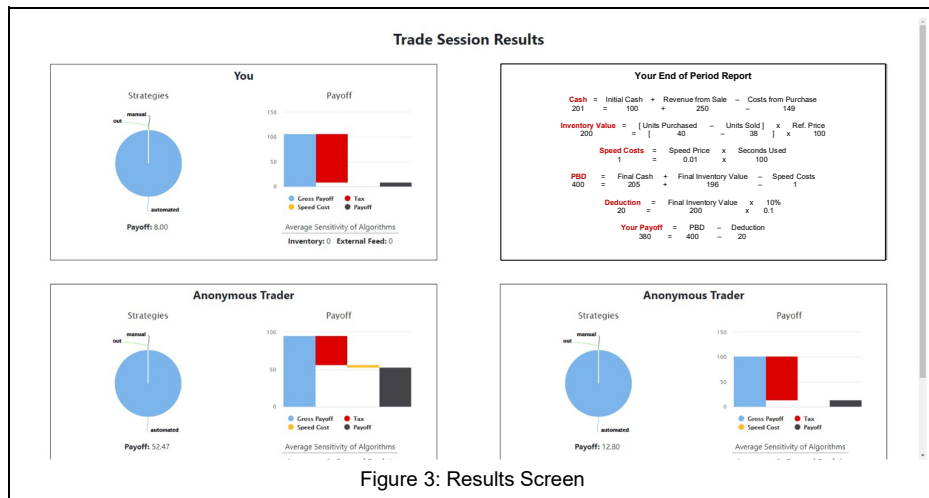


Figure 3: Results Screen

Figure 22: Comprehension Questions

General:

- What degree program are you in?

- B.Sc./M.Sc./Ph.D.

- Does your degree have a background in economics?

- Yes/No

- How would you grade the clarity of the video instructions?

- Scale from 1 = very good to 5 = deficient
- Did you find parts of the instructions confusing? If yes, which ones **[Free Text]**

- How would you grade the clarity of the user interface?

- Scale from 1 = very good to 5 = deficient
- Did you find parts of the user interface confusing? If yes, which ones? **[Free Text]**

- Do you have suggestions for improvement, encountered errors, or felt that the experiment did not work properly? We welcome any feedback that helps improve the user experience. **[Free Text]**

Inventory:

- Suppose you have 10 units of inventory at a certain point during the trading period and significantly increase the inventory sensitivity. What is likely to happen in the next few seconds?

Answer: The inventory will most likely decrease and tend to zero.

Answer: The inventory will most increase rapidly.

Answer: The inventory will stay at exactly 10 units.

Answer: The inventory will shortly after tend to minus 10 units.

Answer: Nothing will happen.

- What is a disadvantage of increasing the inventory sensitivity?

Answer: Your bot will focus less on profitable trading and more on balancing the inventory, which will most likely result in you receiving poorer prices.

Answer: Your bot will focus solely on profitable trading, which will most likely result in high positive or negative inventories and thus in a higher deduction payment.

Answer: Your bot focuses exclusively on the prices of the external market, which might cause you to miss the moments when your own market reacts faster.

Answer: Your bot focuses exclusively on the prices on your own market, which might cause you to miss the moments when the external market reacts faster.

o Answer: There are no disadvantages associated with increasing the inventory sensitivity slider.

- Suppose you end the trading period with a PBD (Payoff Before Deduction) of ECU 100 and a positive inventory of 10 units. The deduction rate is 10% and suppose the reference price is also ECU 100. What is your payoff?

o Answer: The deduction amounts to $10 \times 100 \times 0.1 = 100$ ECU. The deduction needs to be subtracted from the PBD in the amount of 100 ECU, such that the payoff is equal to 0 ECU.

o Answer: The deduction amounts to $10 \times 100 \times 0.1 = 100$ ECU. The deduction needs to be added to the PBD in the amount of 100 ECU, such that the payoff is equal to 200 ECU.

o Answer: The deduction is equal to the inventory value, which amounts to $10 \times 100 = 1000$ ECU. The inventory value needs to be added to the PBD in the amount of 100 ECU, such that the payoff is equal to 1100 ECU.

o Answer: The deduction is equal to the inventory value, which amounts to $10 \times 100 = 1000$ ECU. The inventory value needs to be subtracted from the PBD in the amount of 100 ECU, such that the payoff is equal to -900 ECU.

- Suppose you end the trading period with a PBD (Payoff Before Deduction) of ECU 100 and a negative inventory of -10 units. The deduction rate is 10% and suppose the reference price is also ECU 100. What is your payout relevant profit for this period?

o Answer: The deduction amounts to $|-10 \times 100 \times 0.1| = 100$ ECU. The deduction needs to be subtracted from the PBD in the amount of 100 ECU, such that the payoff is equal to 0 ECU.

o Answer: The deduction amounts to $|-10 \times 100 \times 0.1| = 100$ ECU. The deduction needs to be added to the PBD in the amount of 100 ECU, such that the payoff is equal to 200 ECU.

o Answer: The deduction is equal to the inventory value, which amounts to $-10 \times 100 = -1000$ ECU. The deduction needs to be added to the PBD in the amount of 100 ECU, such that the payoff is equal to -900 ECU.

o Answer: The deduction is equal to the inventory value, which amounts to $-10 \times 100 = -1000$ ECU. The deduction needs to be subtracted from the PBD in the amount of 100 ECU, such that the payoff is equal to 1100 ECU.

External Market:

- Suppose you significantly increase the external sensitivity at a certain point in the trading period. What is likely to happen if, shortly thereafter, the best bid and best ask in the external market increase significantly?

o Answer: Both, your bid and ask increase following the external market.

o Answer: Both, your bid and ask decrease opposing the external market.

o Answer: It is likely that neither your bid nor ask will change.

o Answer: It is likely that your bid increases, while your ask decreases.

o Answer: It is likely that your bid decreases, while your ask increases.

- What is a disadvantage of setting the external sensitivity slider up to 1?

o Answer: Your bot focuses exclusively on the prices of the external market, which might cause you to miss the moments when your own market reacts faster.

o Answer: Your bot focuses exclusively on the prices on your own market, which might cause you to miss the moments when the external market reacts faster.

o Answer: Your bot will focus less on profitable trading and more on balancing the inventory, which will most likely result in you receiving poorer prices.

o Answer: Your bot will focus solely on profitable trading, which will most likely result in high positive or negative inventories and thus higher deduction payment.

o Answer: There are no disadvantages associated with increasing the external sensitivity slider up to 1.

- What is a disadvantage of setting the external sensitivity slider to 0?

o Answer: Your bot focuses exclusively on the prices of the external market, which might cause you to miss the moments when your own market reacts faster.

o Answer: Your bot focuses exclusively on the prices on your own market, which might cause you to miss the moments when the external market reacts faster.

o Answer: Your bot will focus less on profitable trading and more on balancing the inventory, which will most likely result in you receiving poorer prices.

o Answer: Your bot will focus solely on profitable trading, which will most likely result in high positive or negative inventories and thus higher deduction payment.

o Answer: There are no disadvantages associated with increasing the external sensitivity slider.

Speed:

- Suppose two players (Anna and Tim) operate in the same market and have the same high sensitivity to the external market. Suppose they also have the same low inventory sensitivity. The only difference is that Anna has speed on, whereas Tim has speed off. Whose bid and ask will change faster when the prices on your own market or the external market move?

o Answer: Anna's bid and ask will change faster in response to price movements on either market.

o Answer: Tim's bid and ask will change faster in response to price movements on either market.

o Answer: Anna's bid and ask will change faster if the prices on the external market move, while Tim's bid and ask will change faster if the prices on their own market move.

o Answer: Tim's bid and ask will change faster if the prices on the external market move, while Anna's bid and ask will change faster if the prices on their own market move.

o Answer: Anna and Tim's bid and ask will change equally fast in response to price movements.

- Assuming you turn on speed for 3 Minutes, what is your cost?

o Answer: $180 * 0.01 = 1.8$ ECU

o Answer: $180 * 0.1 = 18$ ECU

o Answer: $180 * 1 = 180$ ECU

o Answer: $180 * 10 = 1800$ ECU

o Answer: The speed option does not create costs

Market Design Specific:

- Suppose there exists exactly one ask in the market at 99 ECU and that Anna places a bid at 100 ECU right at the beginning of the trading period. One second later, Tim also places a bid at 100 ECU. What would happen?

o Answer: The bids are processed serially, which means Anna's bid will be executed against the ask, while Tim's bid enters the orderbook. [CDA][IEX]

o Answer: Both bids are processed simultaneously, it is randomly determined which bid will be executed against the ask, the other bid will enter the orderbook. [FBA]

o Answer: Both bids enter the orderbook at 100 ECU.

o Answer: Anna's bid will be executed against the ask at the end of the trading period.

o Answer: Tim's bid will be executed against the ask at the end of the trading period.

- What are the two distinctive features of hidden order? [IEX only]

o Answer: The exchange automatically adjusts the price of these orders, which is indicated by the black vertical bar, and the market participants cannot observe whether there exist hidden orders at the black bar or not.

o Answer: The exchange automatically adjusts the price of these orders, which is indicated by the black vertical bar, and all market participants can observe whether hidden orders exist at the black bar or not.

o Answer: The exchange automatically adjusts the price of these orders, which is indicated by the black vertical bar, and these orders can be placed only by you and the other human market participants.

o Answer: The automated investors adjust the price of these orders according to your behavior and the behavior of the other human market participants, none of the market participants can observe whether there exist hidden orders at the black bar or not.

o Answer: Hidden orders cannot be distinguished from ordinary limit or market orders.