

Getting burned by frictionless financial markets ^{*}

Anirudh Dhawan ^a, Benjamin Loos ^b, Marco Navone ^a, and Tālis J. Putniņš ^{a,c}

^a *University of Technology Sydney*

^b *Technical University of Munich*

^c *Stockholm School of Economics in Riga*

This version: October 23, 2020

Abstract

The evolution of stock markets into highly accessible, low-cost, virtually frictionless venues has been praised by policymakers and institutional investors. But could frictionless markets actually harm individual investors by increasing impulsive trading driven by heuristics and biases? Using laboratory experiments, we examine how investor performance is impacted by various trading frictions: transaction costs, time delays in placing orders, and tasks requiring cognitive effort. High transaction costs and time delays have no effect or harm performance, while cognitive tasks benefit participants that are most prone to underperforming. Frictions can yield benefits when they help inattentive investors consider information they might otherwise neglect.

Keywords: behavioral finance, trading frictions, individual investors, cognitive effort, investor attention, experimental asset markets

JEL classification: G11, G41

* The Internet Appendix accompanying this study can be found at this link (<https://bit.ly/2AGr6bF>).

1. Introduction

Frictions that impede trade are typically considered harmful by financial economists. At the market-level, frictions harm resource allocation (Barlevy, 2003), informational efficiency (Amihud and Mendelson, 1986; Lo, Mamaysky and Wang, 2004), and arbitrage (Shleifer and Vishny, 1997). Regulators and marketplace operators go to great lengths to reduce frictions in financial markets. Large retail brokers such as Charles Schwab, Robinhood, and TD Ameritrade now allow investors to trade US stocks with zero fees on highly streamlined platforms. Investors can create trading accounts in “about 10 minutes” and trade online from anywhere using their computers, mobile phones, or by simply giving voice commands to virtual assistants such as Alexa, making markets highly accessible and virtually frictionless.¹

While reduced frictions benefit some market participants, we conjecture they may harm individual investors by increasing their impulsive, heuristic-driven trading and exacerbating their tendency to underperform the market (Barber and Odean, 2000). Studies of individual investors primarily attribute their underperformance to biases in decision-making (Barber and Odean, 2013). Psychology literature contends that individuals making decisions rely on two kinds of cognitive processing: intuitive processing that is fast-paced and automatic, and analytical processing that is slow-paced and deliberative. Intuitive cognitive processes are more prone to systematic errors, biases, and reliance on heuristics (Kahneman, 2011). Lack of deliberative thinking has been shown to lead to rash decisions in financial markets (Kocher, Lucks and Schindler, 2018). Trading frictions that increase deliberative thinking may therefore help reduce some errors caused by intuitive thinking and heuristics. We test this conjecture for a range of frictions using laboratory experiments.

In our first experiment, based on Weber and Camerer (1998), participants trade multiple assets in markets that last for multiple trading periods. Assets follow a stochastic price process, with price movements in every period; participants can trade in a trading period at the prevailing market price. In our second experiment, based on Plott and Sunder (1988), participants trade one asset in successive continuous double-auction limit-order book markets. This asset pays out one of three values at the end of a market; all traders receive a private clue about one of the incorrect payouts.

¹ Sources: Charles Schwab and TD Ameritrade (<https://bit.ly/3oeQhGn>).

We test the effect of three trading frictions, each implemented as a separate treatment, and contrast these treatments against a baseline frictionless market (*NOFRICTION* treatment). The high transaction cost treatment (*HIGHCOST* treatment) increases transaction costs by a factor of up to five. The slow markets treatment (*SLOW* treatment) adds a delay between investors' opportunities to trade. The cognitive effort treatment (*TASK* treatment) asks participants a question about their beliefs regarding the fundamental value before allowing them to trade.

Our first key finding is that participants exert more cognitive effort while making trading decisions in both the *SLOW* and *TASK* treatments. We use time between orders as a measure for cognitive effort in trading decisions. Participants spending more time between placing orders are likely to be spending more time thinking about those orders. We adjust for the mechanical effects caused by the delay (20-second waiting period) in the *SLOW* treatment by adjusting the time between orders metric for the *NOFRICTION* treatment. In this adjustment, we assume that all orders in the *NOFRICTION* treatment that would be disallowed by a 20-second waiting period occur immediately after it.² We adjust for the mechanical delays due to reading the question in the *TASK* treatment by removing all observations where participants encounter the question for the first time.³ We find that, in both experiments, participants spend more time between orders in the *SLOW* and *TASK* treatments compared to the frictionless market. Participants spend between 17–47% more time per order in the *SLOW* and *TASK* treatments, after adjusting for mechanical effects. The additional time spent per order is between 6–17% of the duration of a trading period in the multiple assets experiment and between 3–5% of the duration of an entire market in the single asset experiment. We conclude that participants exert more cognitive effort per trading decision in both the *SLOW* and *TASK* treatments.

Our second key finding is that both the *HIGHCOST* and *SLOW* treatments either have no effect or harm investor performance relative to a frictionless market, while the *TASK* treatment helps improve the performance of the participants most prone to underperformance. While high transaction costs and time delays are not beneficial, inducing cognitive effort related to trading decisions seems to benefit at least one set of participants. In the *TASK* treatment, participants in the bottom-quartile for performance underperform nearly 57% (40%) less than the *NOFRICTION* treatment in our multiple assets (single asset) experiment.

² This adjustment only applies to the single asset experiment.

³ The question remains the same in all subsequent instances. An impulsive participant can quickly input the same answer as the first instance and move to order submission with minimal delay.

Our third key finding is that the benefits of the cognitive task are primarily due to better decision-making. To reach this conclusion, we decompose the sources of participant underperformance into losses due to overtrading (overpaying transaction costs) and losses due to bad decision-making. In both experiments, we find that the performance improvement in the *TASK* treatment is almost entirely explained by better decision-making rather than less overtrading. We further decompose losses due to bad decision-making into a fundamental and non-fundamental component. In the multiple assets experiment, we find that the *TASK* treatment helps underperformers more closely match the Bayesian optimal strategy, i.e., helps them lose less money due to the fundamental component of bad decision-making. In the single asset experiment, the *TASK* treatment helps underperformers lose less money due to speculating on prices, i.e., non-fundamental component, without affecting losses due to the fundamental component. In both experiments, the *TASK* treatment targets the component that is the dominant contributor to underperformance due to bad decision-making. Interestingly, even the *SLOW* treatment helps the worst performers lose less money due to speculating on prices in the single asset experiment. However, this benefit is almost entirely negated by an increased tendency to sell the asset at a low price when it has a high value (fundamental component).

One interpretation of our results is that the cognitive task friction (but not any other friction) helps direct the attention of inattentive investors (underperformers) to important information and consider aspects of the trading decision that they would otherwise neglect or underweight.⁴ Previous research shows that investors have limited attention, and this limited attention can cause them to consistently neglect or underreact to nonsalient information and overreact to salient information (Hirshleifer and Teoh, 2003; Palomino, Renneboog and Zhang, 2009; Laudenbach, Loos, Pirschel and Wohlfart, 2020). In our experiments, prices are more salient than information about asset fundamentals. Prices are displayed prominently on a large graph that is constantly updated as the market price changes, whereas information about asset fundamentals is static and displayed less prominently on one corner of the investor's screen. Previous research finds that investors have a tendency to overweight recent price movements since they expect these price movements to extrapolate in the future (Greenwood and Shleifer, 2014; Barberis, Greenwood, Jin and Shleifer, 2015). The salience of prices in our experiments can exacerbate this tendency.

⁴ We assume a linear-log relationship between performance and attention and conjecture that the most inattentive investors are the worst performers. Gargano and Rossi (2018) document such a relationship for portfolio returns and investor attention.

The cognitive task friction assists inattentive investors in making better trading decisions by helping them approach the trading decision in the “right way” (see Enke, Gneezy, Hall, Martin, Nelidov, Offerman and Ven (2020)). To answer the question in the *TASK* treatment, participants need to use the nonsalient information about asset fundamentals and carefully consider the process by which the fundamental value is determined. In doing so, they reduce their tendency to underweight information about asset fundamentals and overweight recent price movements. However, in the *SLOW* treatment this tendency is not corrected fully.

Evidence for the attention-inducing property of the *TASK* treatment lies in the fact that participants increase deliberation before trades in both the *SLOW* and *TASK* treatments, but only in the *TASK* treatment are some participants able to earn a material performance benefit from this increased deliberation. This is likely because the increased deliberation in the *SLOW* treatment does not help investors consider important information or crucial aspects of the decision. The mechanism through which the treatments affect investor performance differently in our two experiments also supports this interpretation. In the multiple assets experiment, the tendency to overweight recent price movements and underweight asset fundamentals causes underperformers (the most inattentive investors) to deviate from the optimal strategy. The *TASK* treatment (but not the *SLOW* treatment) helps them reduce this tendency by increasing deliberation on fundamentals. In the single asset experiment, the tendency to overweight recent price movements causes inattentive investors to speculate on prices, with the expectation that prices will keep moving in the same direction. Although both the *SLOW* and *TASK* treatments help reduce this tendency, the *SLOW* treatment increases their tendency to prematurely sell when the asset has a high value. This shows that these inattentive investors are underweighting asset fundamentals in the *SLOW* treatment but not in the *TASK* treatment. Overall, the evidence suggests that frictions that induce investors to think about stock fundamentals can help them make better trading decisions.

Our findings are important since the barriers to join stock markets today are very low. Bogan (2008) documents that online trading has significantly increased household participation rates in financial markets. More recent evidence indicates that newer innovations in financial technology (Fintech), such as robo-advising, have further increased financial market participation rates (Reher and Sokolinski, 2020). In the first four months of 2020 alone, Robinhood added more

than three million users, half of whom were first-time investors.⁵ The number of individual investors using online trading to access financial markets is ever-increasing. Retail trading today is almost entirely conducted online. Out of the brokers we mentioned at the start, only Charles Schwab lets clients trade over the phone. This ease-of-access means that an increasing number of potentially inattentive individuals are taking up online trading.

It is easier for inattentive investors to neglect important information in real markets than in our experiments. This is because information in real markets is more complex, voluminous, and dispersed than in our laboratory asset markets. A stock's fundamental value can be a function of a myriad of factors. In our experiments, the process that determines an asset's fundamental value is simple and is clearly explained to participants. In our setting, there are a maximum of three information sources (market prices, the order book, and the private information in the form of a clue) among which only two sources (prices and order book) are dynamic. In real markets, investors have to sift through numerous information sources which constantly provide new information. Therefore, the effects we find in the laboratory are possibly even larger in real markets.

Our study contributes to the literature studying individual investor behavior in financial markets. Barber and Odean (2002) find that investor performance reduces when they move from trading over telephone calls to online trading. Choi, Laibson and Metrick (2002) also analyze a similar jump from telephonic trading to online trading, however, only for 401(k) accounts (retirement savings accounts). They document no difference in performance across the two environments. Both these studies compare trading in settings that implicitly have different degrees of frictions, although there are other factors that change as well such as the framing of the information about the market and the available information (Kalda, Loos, Previtro and Hackethal, 2020). Telephonic trading is costlier, slower, and more effortful. In that sense, the findings of Barber and Odean (2002) are consistent with our laboratory evidence that some frictions can help some investors reduce their underperformance. Our contribution is to separate the various frictions, test them separately, and isolate them from the effects of changes in the information environment. Importantly, we find that different frictions have different effects. Additionally, in the laboratory, we are able to overcome the self-selection effects that might influence the results of these field studies. Our renewed analysis of frictions is crucial since the

⁵ Source: Robinhood (<https://bit.ly/2FTe8dw>).

observations in Barber and Odean (2002) might not hold today as today's markets have significantly fewer frictions than even the frictionless setting in Barber and Odean (2002). For example, Charles Schwab used to charge a "reduced" \$14.95 trading commission per online trade in 2000 while it offers commission-free trading now.⁶

2. Hypotheses development

2.1. Retail investor underperformance

Behavioral finance literature finds that retail investors consistently underperform the market due to their systematically biased decision-making (Barber and Odean, 2013). Investor biases can lead them to overtrade, sell winning stocks too early and hold on to losing stocks for too long, chase trends, and under-diversify, among others. Barber, Lee, Liu and Odean (2009) provide the most comprehensive breakdown of sources of underperformance in the literature. They split the losses made by Taiwanese individual investors into four categories: trading commissions, transaction taxes, trading losses, and market timing losses. They find that the first three categories can explain the bulk of the losses these individuals make by trading.

Fundamentally, in financial markets, traders can employ two strategies, either individually or in conjunction. The first strategy involves using the information available to them to generate a belief about the asset's true value. The trader can then buy the asset if the market price is lower than their true value belief and sell the asset if the market price higher than their true value belief. Traders stand to make losses from this strategy if their belief is incorrect and they end up holding or accumulating the asset despite the true value being lower than the market price or selling the asset despite the true value being higher than the market price. The second strategy that traders can employ is to predict the future direction of market prices, and then buy low and sell high. In this strategy, traders can make losses if they are unable to exit their positions at a better price. To employ either of these strategies successfully an investor must be able to: make sound statistical inferences; correctly interpret and use private information; accurately infer and use public information contained in market prices; avoid overpaying transaction costs and/or taxes; and limit heuristics and biases. An investor falling short on one or more of these

⁶ Sources: CNNMoney (<https://cnn.it/31xJUUT>) and Charles Schwab.

parameters might underperform systematically. In our experiments, we pick two asset market designs that collectively capture all these sources of underperformance.

2.2. Psychological effects of obstacles

Psychology literature finds that people are more likely to display biases and rely on heuristics to make decisions if they use more intuitive or automatic cognitive processing and less analytical or deliberative cognitive processing (Kahneman, 2011). Additional cognitive effort has been shown to mitigate biases caused by overreliance on intuitive thinking to some extent (Enke et al., 2020). Consequently, any external stimulus that can trigger additional cognitive effort before a decision could help reduce the influence of heuristics and biases on that decision.

Further evidence from the psychology literature shows that obstacles people encounter while making decisions can have the effect of improving decision-making by inducing additional cognitive effort. When encountered with an obstacle, people tend to adopt a global processing based approach, i.e., they take a step back and see the “big picture” (Marguc, Förster and Van Kleef, 2011). The researchers find that this global processing mindset needed to tackle the obstacle carries over to unrelated tasks performed after encountering the obstacle.

We hypothesize that certain frictions that traders encounter immediately before a trade can also perform the same role of increasing global processing in the trading decision as obstacles do in Marguc et al. (2011). Such global processing could help investors make more thoroughly considered trading decisions by properly using all sources of information at hand, considering previously neglected aspects of the decision, and accounting for previously disregarded adverse contingencies. This reasoning is supported by previous research. Investors perform better when trading via phone calls than when trading online (Barber and Odean, 2002). Trading on the phone is more expensive, slower, and more effortful than online trading.⁷ Investors appear to fare better in a trading environment riddled with frictions than a relatively frictionless trading environment.

We select three frictions that can act like obstacles in Marguc et al. (2011) insofar as they can help investors take a step back and reconsider their trading decision, perhaps even helping them

⁷ Telephonic trading can also involve more cognitive effort than online trading. An investor might discuss the order with the broker before finalizing the order. This discussion could help an investor process the trading decision globally, especially if the broker presents new information or informs the investor of aspects of the decision that they previously neglected. Even a simple “are you sure?” from the broker can cause the investor to exert additional cognitive effort before a trade.

use the information at hand holistically. The first friction is high transaction costs. Making trading more expensive can prohibit investors from making ill thought out trades by ensuring that they only trade when their trades are expected to be profitable enough to justify paying the high transaction costs. Barber and Odean (2000) find that investors trade too much due to overconfidence bias and end up overpaying transaction costs. If a friction like high transaction costs does induce more cognitive effort, it could reduce investor overconfidence and cause them to reduce their trading activity. The second friction we select is time delays before orders. This friction can help investors take a step back from the fast-paced trading environment and spend more time on carefully thinking about a trading decision. Lastly, the third friction we select involves asking participants a question regarding their beliefs about the fundamental value of the asset before orders. This friction can induce participants to think about the fundamental value of the asset and help them use that information in their trading decision, in case they are not doing so already. The marginal benefit of any additional cognitive effort and reduction in biases caused by our frictions is likely to be highest for the most biased individuals. Consequently, we expect our frictions to benefit the most biased investors, e.g., the worst performing investors, more than other investors.⁸

3. Experiment design

3.2. Multiple assets experiment

Our multiple assets experiment uses a modified version of the market in Weber and Camerer (1998). In this market, participants make trading decisions (buy, sell, or do nothing) for four assets (Asset 1–4) over eight trading periods (Period 1–8). Each trading period lasts for one minute. Participants are endowed with \$1,000 (laboratory currency) in cash and four units of each of the four assets at the start of the market.

The prices of the assets are not determined endogenously through participant trading. Rather, prices of all assets start at \$180 and follow a stochastic process. The prices of all assets move by \$15 after every trading period. Each asset has a different probability of experiencing a price increase or decrease after a given period. The probability of a price increase in any given period

⁸ Corgnet, Desantis and Porter (2018) find that traders with high (low) cognitive ability, i.e., traders that are less (more) prone to biases, perform better (worse) in financial markets.

is 65% for one asset, 55% for one asset, 45% for one asset, and 35% for one asset. The probability of a price decrease is one minus the probability of a price increase. In a trading period, participants can submit orders to buy or sell at the prevailing market price (displayed on the participants' screens) for the asset. These orders execute automatically at the end of the trading period. Participants are informed about the price process and the price increase probabilities; however, they are not told the price increase probability corresponding to each asset. They need to infer this information using each asset's price path, which is displayed graphically on their screens in all trading periods. In addition to the eight trading periods, we add six non-trading periods before Period 1.⁹ These non-trading periods help in participant decision-making by providing a larger sample of price movements.

At the end of a market, a participant's portfolio value is added to their earnings for the experimental session. Their assets are valued at the prevailing market prices after the price movement in the final period. Their cash and asset balances are reset in the next market.

To be successful in this experiment, a trader must be able to: make sound statistical inferences; accurately infer and use public information contained in market prices; avoid overpaying transaction costs; and limit heuristics and biases.

3.1. Single asset experiment

Our single asset experiment uses an asset market design that is a modified version of Market 9 in Plott and Sunder (1988). Participants trade a single asset in a continuous double-auction limit-order book market in which they are free to post limit and market orders at their desired prices and volumes at any time. Each market lasts for three minutes. At the start of each market, participants are endowed with \$1,000 (laboratory currency) in cash and four units of the asset. The asset has a cash payout at the end of each market. This payout is independent of the payout in other markets. The asset does not generate any income other than this payout. Once the payout is made at the end of the market, the cash balances of all participants are recorded as earnings from the market. Their cash and asset balances are reset in the next market.

A crucial feature of this market design is the way traders are informed about the fundamental value of the security, or the end of market payout. The asset can have one of three payouts, \$50, \$240, or \$490. The probability of the \$50 payout is 35%, the probability of the

⁹ Participants can see the price movements in these non-trading periods on their price charts.

\$240 payout is 45%, and the probability of the \$490 payout is 20%.¹⁰ Before each market, all traders are given a clue about which of the three payouts is incorrect. However, all traders do not have the same clue. Half the traders are told one of the incorrect payouts and the other half are told the other incorrect payout. For example, if the asset pays out \$50 in a given market, half the traders are told that the asset payout is not \$240, and the other half is told that the asset payout is not \$490. Although each trader is partially informed about the correct payout, collectively, the market has full information.

In addition to the skills required in the multiple assets experiment, in this experiment, to be successful, a trader must also be able to correctly interpret and use private information.

3.3. Treatments

Both our experiments have four treatments: the *NOFRICTION* treatment, the *HIGHCOST* treatment, the *SLOW* treatment, and the *TASK* treatment. The first treatment is our baseline control treatment. We take the baseline *NOFRICTION* treatment and individually add one friction to create the latter three treatments. The transaction cost in the *NOFRICTION* treatment in the multiple assets experiment is \$5, while the transaction cost in the *NOFRICTION* treatment in the single asset experiment is 2% of the transaction value.

The *HIGHCOST* treatment contains the high transaction cost friction. In this treatment, we increase the transaction cost to \$20 in the multiple assets experiment and 10% of the transaction value in the single asset experiment.

The *SLOW* treatment contains the time delay friction. In this treatment, we increase the duration of a trading period from one minute to two minutes in the multiple assets experiment. In the single asset experiment, we implement a compulsory 20-second waiting period between orders. This means that after placing an order (market or limit), participants are not allowed to place another order for 20 seconds.

Lastly, the *TASK* treatment contains the cognitive effort task friction. In this treatment, we make participants answer a mandatory question about their beliefs regarding the true value of the asset(s).¹¹ In the multiple assets experiment, we ask participants about the assets that they think correspond to the most extreme price increase probabilities (see Panel A in Figure 1 for the exact

¹⁰ We follow Corgnet et al. (2018) in the modification of the payout probabilities.

¹¹ To ensure that participants take these questions seriously, we attach a small monetary reward of \$10 (in laboratory currency) for correct responses and a penalty of \$10 for incorrect responses.

question). Participants are only required to answer this question once in a trading period. Once they answer the question, they can place one or more orders in that trading period. If they want to trade in subsequent trading periods, they need to answer the question again, once per trading period. In the single asset experiment, we ask participants about their beliefs regarding the asset's true payout (see Panel B in Figure 1 for the exact question). Participants need to answer this question before every order (market or limit) they place. If they place multiple orders in a market, they need to answer the question multiple times, once before each order.

< Figure 1 here >

3.4. Additional experiment-related details

We conduct both experiments in the University of Technology Sydney Behavioral Lab. We develop the software for both experiments using z-Tree (Fischbacher, 2007). We start off with the GIMS software (Palan, 2015) as the foundation and make the necessary changes to fit our market designs. Participants in our experiments are undergraduate and postgraduate students at the University of Technology Sydney. These participants were recruited using the Online Recruitment System for Economic Experiments database. We conduct eight laboratory sessions for the multiple assets experiment and four sessions for the single asset experiment. For each experiment, we select the number of sessions to run such that we generate a large enough sample size for our statistical tests.¹² In total, 95 participants participate in the multiple assets experiment and 47 participants participate in the single asset experiment. After each session, participants are ranked based on their total earnings in the session. Participants receive a cash reward between 25 AUD and 60 AUD based on their rank.

For both experiments, we use a *within subjects* design, i.e., all our participants receive all four of our treatments in a randomized sequence. In the multiple assets experiment, we conduct four markets. Each market corresponds to one treatment. To control for learning effects, we vary the sequence of treatments such that each treatment has a similar number of subject-market observations in each position in the sequence of markets. Every participant participates in four

¹² We run more experimental sessions for the multiple assets experiment since that experiment generates fewer observations per session.

markets. We generate a total of 380 subject-market observations in the multiple assets experiment, 95 observations per treatment.

In the single asset experiment, we conduct 12 markets in each experimental session. Each market corresponds to one treatment. In total, there are three markets for each treatment in an experimental session. We generate a total of 564 subject-market observations in the single asset experiment, 141 observations per treatment.

Table 1 Panel A summarizes the structure of the asset markets and Table 1 Panel B summarizes the treatments. Participant instructions and screenshots of participant trading screens for the multiple assets experiment (single asset experiment) are in Section IA3 (Section IA4) of the Internet Appendix.

< Table 1 here >

4. Experimental results

4.1. Cognitive effort

We start by examining whether our frictions help our participants think harder before making trading decisions. To do so, we examine trade frequency and time between orders across treatments. The time in a market or trading period is limited and participants are not interrupted or asked to complete any other tasks (except the cognitive task in the *TASK* treatment) during this time, implying that their entire focus is on trading. Consequently, participants placing fewer trades and participants spending more time between orders are likely to be thinking more about each order than other participants.

4.1.1. Trade frequency

We start by examining treatment effects on trade frequency. Kocher et al. (2018) also use trading frequency as an indicator of the degree of “activeness” or “passiveness” in a trader’s decision-making. Using Baumeister, Bratslavsky, Muraven and Tice (1998), Kocher et al. (2018) highlight that in the context of a financial market, trader passiveness could result either in higher or lower trading. Based on Baumeister et al. (1998), a passive individual performs routine or expected actions without deliberation. This indicates a heavier reliance on automatic cognitive

processing than deliberative cognitive processing. In a trading context, if passive traders associate participation in financial markets with trading frequently, their automatic systems might push them to trade frequently. Their trades would be prone to biases due to the lack of deliberation. On the contrary, if traders do not associate financial market participation with frequent trading, their automatic systems would lead them to favor inaction over action, thus reducing their trading frequency.

Table 2 displays the number of trades across treatments in both experiments. This table uses mean values for subject-market observations. In the multiple assets experiment, participants make an average 14.09 trades per market in the *NOFRICTION* treatment. The number of trades reduces by 3.44 trades in the *TASK* treatment. No other treatment causes a significant reduction in number of trades.

In the single asset experiment, on average, participants trade 7.9 times per market in the *NOFRICTION* treatment. The number of trades reduces by 3.93 trades and 3.55 trades in the *SLOW* and *TASK* treatments, respectively.

The reduction in trading activity we observe can be explained by more deliberation or “activeness” in trading induced by our frictions. Based on this explanation, passive traders trade frequently since their automatic systems associate participant in financial markets with frequent trading. Our frictions (*SLOW* and *TASK*) help these traders slow down and carefully consider each trade. Further evidence of trader activeness or deliberation lies in the fact that traders also make fewer momentum trades in the *SLOW* (in the single asset experiment) and *TASK* (in both experiments) treatments (see Table 2).¹³ Passive traders making impulsive trading decisions would be inclined to follow the market and trade in the direction of price movements, i.e., buy after a price rise and sell after a price drop. The reduction of momentum trades caused by our frictions indicates that passive traders are becoming more active and thinking harder before each trade. This initial evidence indicates that our *SLOW* and *TASK* treatments appear to be inducing traders to exert additional cognitive effort and deliberative cognitive processing before trades.

4.1.2. Time between orders

Next, we examine treatment effects on time between orders. Although our trade frequency results indicate that our *SLOW* and *TASK* treatments are inducing participants to think harder

¹³ We classify a buy (sell) trade as a momentum trade if it occurs immediately after a price rise (fall).

before trades, these results could be misleading if trades are clustered together in time and/or mainly occur early in the market or trading periods. In this case, we cannot conclusively state that a reduced number of trades implies higher cognitive effort per trade. Hence, it is useful to examine the time between orders as well, to determine exactly how much time participants spending thinking about their trading decisions.

In the multiple assets experiment, we measure the average time between orders as the time of the last order in a trading period divided by the number of orders in the trading period. We generate a value for the average time between orders at the market-level by averaging across assets and trading periods. In the single asset experiment, we perform the same calculation for each market.

In the single asset experiment, the 20-second waiting period in the *SLOW* treatment can have a mechanical effect on time between orders. We control for this mechanical effect by using an adjusted time between orders metric for the *NOFRICTION* treatment in the single asset experiment. We adjust the time between orders metric to make the *NOFRICTION* treatment comparable to the *SLOW* treatment. We do so by assuming that all orders in the *NOFRICTION* that would be mechanically disallowed by a compulsory 20-second waiting period occur immediately after the waiting period. To calculate the adjusted time between orders metric, we set the time between orders for all orders occurring within 20 seconds of the previous order to 21 seconds (20 seconds waiting period and one second for order submission).¹⁴

We also recognize that answering the question in the *TASK* treatment can cause a mechanical delay in order submission. To control for this mechanical delay, for the *TASK* treatment, we exclude the first trading period in the multiple assets experiment and the first (of three) trading market corresponding to the *TASK* treatment in the single asset experiment. Excluding the first market and trading period makes sense since this is the first instance when participants encounter the question. Reading and answering the question might cause a delay in the first instance. In subsequent markets and trading periods, the question remains the same. A passive trader who wishes to ignore the question and jump straight to order submission can quickly input and submit the same answer as the first instance or any random answer, with minimal delay in order submission. They do not need to read the question again.

¹⁴ We add one second for order submission since that is the minimum time required for a participant to submit an order. In our data, the minimum time between orders for a market order is 0.22 seconds and for a limit order is 1.05 seconds.

Table 2 displays the average time between orders. In the multiple assets experiment, the average time between orders in the *NOFRICTION* treatment is 18.98 seconds. This number increases by 6.84 seconds or 36.04% ($t = 4.73$, $N = 190$) in the *SLOW* treatment. This increased time spent per order is slightly more than 10% of the total extra time (1 minute extra) participants are given per trading period in the *SLOW* treatment. Panel B in Figure 2 plots the number of orders across time in the *NOFRICTION* and *SLOW* treatments. The ordering activity in both treatments is similar initially, however, ordering activity in the *SLOW* treatment dissipates gradually. Interestingly, participants do place orders even after the normal trading period duration (1 minute). This indicates that they are utilizing the extra time given; the extra time enables them to think harder about their orders.

The time between orders in the multiple assets experiment increases from 18.98 seconds in the *NOFRICTION* treatment to 29.45 seconds in the *TASK* treatment; an increase of nearly 9.36 seconds or 46.6% ($t = 6.87$, $N = 190$). From the graph of number of orders across time in the *NOFRICTION* and *TASK* treatments (Panel C in Figure 2), it is quite evident that participants are spending more time thinking about orders in the *TASK* treatment. Ordering activity in the *TASK* treatment rises at a much slower rate than the *NOFRICTION* treatment. In the *TASK* treatment, participants are likely spending the first-half of the trading period absorbing the new information they receive and thinking about their orders (while simultaneously responding to the cognitive task) and only placing orders in the second-half of the trading period. In contrast, in the *NOFRICTION* treatment, participants are jumping straight into placing orders, with ordering activity peaking in the first-half of the trading period itself. Participants are clearly spending less time to process the new information and decide their trades in the *NOFRICTION* treatment than the *TASK* treatment.

The absolute value of the additional time between orders might not seem high in the multiple assets experiment; the increases are only between 7–10 seconds. However, even 7 seconds is valuable in this experiment since it is around 12% of the total time available to participants in a period to make a trading decision. Hence, even though the magnitudes of the increases might seem low, they are reasonably large given the context.

In the single asset experiment, the average adjusted time between orders increases from 31.42 seconds in the *NOFRICTION* treatment to 36.79 seconds in the *SLOW* treatment; a rise of 5.37 seconds or 17.09% ($t = 2.60$, $N = 282$). Panel B in Figure 3 plots number of orders in the

NOFRICTION and *SLOW* treatment in the single asset experiment across time. Like the multiple assets experiment, ordering activity in both treatments is similar initially, however, ordering activity in the *SLOW* treatment becomes lower gradually.

In the *TASK* treatment in the single asset experiment, we observe that the average time between orders increases by 32.64% ($t = 3.26$, $N = 282$) as compared to the 25.13 seconds in *NOFRICTION* treatment. This result is visualized in Panel C in Figure 3, which plots the number of orders in the *NOFRICTION* treatment and the *TASK* treatment across time. Ordering activity in the *TASK* treatment is lower than the *NOFRICTION* treatment at all times.

Like the multiple assets experiment, the increases in times between orders caused by the *SLOW* and *TASK* treatments in the single asset experiment might seem small (between 5–10 seconds). However, these times are around between 17–33% higher than the time spent on an individual order in the *NOFRICTION* treatment. These times are also almost 3–5% of the total market duration. Hence, the extra time spent thinking about an order is meaningfully large even in this experiment.

The preponderance of evidence suggests that participants are considering their trading decisions more deliberatively in both the *SLOW* and *TASK* treatments but not the *HIGHCOST* treatment. Participants spend between 17–46.6% more time between orders in the *SLOW* and *TASK* treatments. This increased deliberation is accompanied by a reduction in trading activity in the *TASK* treatment in both experiments and in the *SLOW* treatment for only the single asset experiment.

< Table 2 here >

< Figure 2 here >

< Figure 3 here >

4.2. Trading performance

Now, we examine the treatment effects on investor performance. Numerous previous studies have documented that individual investors underperform the market.¹⁵ Investors underperform in our experiments too. In the multiple assets experiment, the average investor underperforms relative to the Bayesian optimal strategy.¹⁶ The average investor underperforms the Bayesian optimal strategy by 7.86% or \$388.28 in terms of final earnings in a market. The most underperforming participants (bottom-quartile for performance) underperform the Bayesian optimal strategy by around \$938 or 21.89% more than their peers in the *NOFRICTION* treatment.

Since the single asset experiment is a zero-sum game, there is no underperformance on average. However, we can quantify the money that participants lose due to transaction costs. On average, participants lose around \$64.22 or 3.54% of their total portfolio value in a market due to transaction costs. In addition, the most underperforming participants earn an average \$432.8 or 24% less than their peers in the *NOFRICTION* treatment.

We begin by examining the overall effects of our treatments on participant performance. In the multiple assets experiment, we use participants' earnings as the measure of performance, after subtracting from it the earnings that would have been earned under the Bayesian optimal strategy.¹⁷ Earnings from a market are calculated as the total value of the cash and asset portfolio held at the end of the market.

Starting with the multiple assets experiment, we regress individual earnings in each market on a set of indicator variables for the three treatments (the frictionless treatment is the base case): $HIGHCOST_{j,k}$, $SLOW_{j,k}$, and $TASK_{j,k}$, where j is an index for laboratory sessions and k is an index for markets within a session.¹⁸ We include fixed effects for market sequence.¹⁹

Table 3 Model 1 reports the results for the performance effects of our treatments. Contrary to expectations, the *HIGHCOST* treatment reduces participant performance relative to the

¹⁵ See Barber and Odean (2013) for a comprehensive review on individual investor trading behavior.

¹⁶ Section IA2 in the Internet Appendix provides details about the Bayesian optimal strategy.

¹⁷ Performing this subtraction merely removes some of the variance in earnings that is due to different asset value realizations.

¹⁸ All our regressions (in the current and subsequent sections) are OLS regressions and all our tests use subject-market observations.

¹⁹ We vary the sequence of treatments such that each treatment has a similar number of subject-market observations in each position in the sequence of markets. Sequence fixed effects ensure that we only compare treatments in the same position.

NOFRICTION treatment. The *HIGHCOST* treatment reduces participant earnings by a magnitude of \$650.91 or 16.5% of the mean earnings level in the *NOFRICTION* treatment. Participant earnings are also lower in the *SLOW* treatment than the *NOFRICTION* treatment. The difference in earnings between both treatments is \$260.61 or 6.61% of the mean earnings level in the *NOFRICTION* treatment. Although the *TASK* has a positive coefficient, it is not statistically significant, implying that subject-market performance is not significantly different in the *TASK* as compared to the *NOFRICTION* treatment. Overall, the frictions either have no significant effect on participant performance or they reduce participant performance when pooling all participants into a single group.

We also explore possible individual-level heterogeneity in treatment effects on performance. Our frictions aim to induce additional cognitive effort; the marginal benefit of this added cognitive effort is likely to be highest for participants who exert the least cognitive effort, i.e., the worst performing participants. We divide participants into quartiles based on total performance across all treatments and test whether our treatments benefit the worst performing more than others.^{20,21} We define a new variable, UP_i , which is an indicator variable that equals one if subject i is in the bottom-quartile for performance. We perform the same regressions as the previous set of tests, only adding UP_i and interaction terms between the treatment variables and UP_i as additional regressors.

Table 3 Model 2 reports the results for the underperformance quartile. The *TASK* treatment benefits underperformers more than other participants. Underperformers earn \$453.06 more in the *TASK* treatment as compared to the *NOFRICTION* treatment. This result implies that their *NOFRICTION* treatment underperformance relative to others is mitigated by around 57% in the *TASK* treatment. The *TASK* treatment also helps reduce the average performance gap between the worst-off participants and other participants; this performance gap reduces from \$938 in the *NOFRICTION* treatment to \$254 in the *TASK* treatment.

< Table 3 here >

²⁰ For the multiple assets experiment, we generate underperformance quartiles at the experiment level, i.e., we compare participants' underperformance with all other participants in the experiment. In this experiment, we are able to compare underperformance across experimental sessions since we measure it relative to the optimal strategy.

²¹ Our results for both experiments are robust to dividing participants into quartiles based on their underperformance in just the *NOFRICTION* treatment.

We now examine the effects of our treatments on participant performance in the single asset experiment. In the single asset experiment, we use participants' earnings from a market as a measure for performance in our tests. A participant's earnings from a market are calculated as the sum of their cash balance and the total payout they receive from their asset holdings at the end of the market.

Like the multiple assets experiment, we regress performance on the set of treatment indicators and include experimental session fixed effects and market sequence fixed effects.^{22,23} We include controls for the participant's clue by including: $ClueNot50_{i,j,k}$ (equals one when participant clue is that the payout is not \$50) and $ClueNot490_{i,j,k}$ (equals one when participant clue is that the payout is not \$490). We also control for the payout in the market by including: $Payout50_{j,k}$ (equals one when the asset payout is \$50) and $Payout490_{j,k}$ (equals one when the asset payout is \$490). For conciseness, we do not report the coefficients for our control variables.²⁴

Table 4 Model 1 reports the results for the performance effects of our treatments. The only treatment that has a statistically significant effect on earnings is *HIGHCOST*. Like the multiple assets experiment, the *HIGHCOST* treatment lowers subject-market earnings as compared to the *NOFRICTION* treatment; subject-market earnings are \$134.14 lower in the *HIGHCOST* friction, or around 8% less than the mean earnings level in the *NOFRICTION* treatment. In contrast, the average subject-market earnings *SLOW* and *TASK* treatments do not differ from earnings in the *NOFRICTION* treatment in a statistically significant manner. The *HIGHCOST* treatment reduces participant earnings on average, while the *SLOW* and *TASK* treatments have no statistically significant effect on earnings when pooling all participants into a single group.

Table 4 Model 2 reports the results for the underperformance quartiles.²⁵ The *TASK* treatment benefits underperformers more than other participants as indicated by the positive and

²² We effectively conduct three sets of four markets (one for each treatment) in the single asset experiment. We control for market sequence by adding fixed effects for the set that contains the market observation.

²³ We perform an additional robustness check to control for non-linear learning effects. To do so, we add the market sequence, squared market sequence, and cubed market sequence as additional controls. The results are reported in Section IA1 in the Internet Appendix. All our results are robust to controlling for non-linear learning effects.

²⁴ Technically, the control variables are not needed due to randomization of the payoffs and clues. However, they help increase statistical power of the tests of interest by absorbing some of the otherwise unexplained variance in earnings.

²⁵ For the single asset experiment, we generate underperformance quartiles at the experimental session level, i.e., we only compare participants' underperformance with other participants in their session. We do so because participant

statistically significant coefficient for the interaction term between $TASK_{j,k}$ and UP_i . The magnitude of this incremental benefit is \$151.83, almost a 40% reduction in their *NOFRICTION* treatment underperformance relative to others. The *TASK* treatment helps reduce the average performance gap between the most underperforming participants and others; this performance gap reduces from \$432.8 in the *NOFRICTION* treatment to \$193.3 in the *TASK* treatment.

Across both experiments, we observe that our treatments either do not improve performance or reduce performance for the average participant when all participants are pooled in a single group. In the pooled tests, the *HIGHCOST* treatment reduces performance relative to the *NOFRICTION* treatment in both experiments, while the *SLOW* treatment does so only in the multiple assets experiment. However, our *TASK* treatment does help the most underperforming participants mitigate their underperformance by about 40% to 57%.

< Table 4 here >

4.3. Overtrading and bad decision-making

In this sub-section, we analyze the treatment effects on specific sources of underperformance. In particular, we aim to examine which source contributes the most to the reduction in underperformance that we observe for the worst performers in the *TASK* treatment.

We measure participant underperformance relative to a benchmark strategy in both experiments and divide this underperformance into two main sources: overtrading and bad decision-making. Overtrading captures losses due to overpaying trading costs and is calculated as the additional transaction costs a subject pays relative to the benchmark strategy. Bad decision-making captures losses due to poor stock selection in the multiple assets experiment and both poor stock selection and market timing losses in the single asset experiment. We use a simple measure for bad decision-making; participant underperformance that is not explained by overtrading.

We start by examining the treatment effects on sources of underperformance for the worst underperformers in the multiple assets experiment. Like the previous sub-section, we define the indicator variable UP_i to represent underperformers. In these tests, we use the same set of

performance in this experiment is sensitive to the asset payout distributions and price paths, which are different in each session.

independent variables as the previous sub-section, i.e., treatment dummies, underperformer dummy, treatment dummies interacted with the underperformer dummy. Our dependent variables are proxies for overtrading and bad decision-making. We measure participant underperformance relative to the Bayesian optimal strategy. Our measure for overtrading is calculated as the difference between the transaction costs paid by a participant and the transaction costs in the optimal trading strategy. Additionally, our measure for bad decision-making is calculated as the difference the optimal strategy's gross earnings, i.e., final earnings with transaction costs added back, and the participant's gross earnings. Collectively, these two variables capture all participant underperformance in the multiple assets experiment.

Table 5 reports the results for the multiple assets experiment. The incremental reduction in underperformance for underperformers in the *TASK* treatment we observe in the previous set of tests is primarily driven by better decision-making. The coefficient on the interaction term between $TASK_{j,k}$ and UP_i is large and negative when regressed on the bad decision-making proxy, indicating that losses due to bad decision-making reduce significantly for underperformers relative to other participants in the *TASK* treatment. The corresponding coefficient for the overtrading regression is small, positive, and statistically insignificant, implying that the *TASK* treatment does not help underperformers reduce overtrading significantly more than the other participants. This result indicates that the worst performers improve their trading decisions in the *TASK* treatment, and this improvement in decision-making is the primary driver of their improved performance.

< Table 5 here >

Next, we perform the same set of tests for the single asset experiment. Except here, we define our proxies for overtrading and bad decision-making differently. In this experiment, we measure participant underperformance relative to a no trade strategy, i.e., we compare participant performance to a hypothetical scenario where they did not trade at all. Consequently, our measure for overtrading is the total transaction costs paid by a subject in a market. Additionally, our measure for bad decision-making is the underperformance that is not explain by overtrading. Our bad decision-making measure is calculated as the difference between earnings from a no trade strategy and the participant's gross earnings, i.e., their final earnings with their transaction

costs added back. Collectively, these two variables capture all participant underperformance in the single asset experiment.

Table 6 reports the regression results for the single asset experiment. We confirm that, like the multiple assets experiment, the underperformance reduction that underperformers experience in the *TASK* treatment is primarily due to better trading decisions. We observe that, on average, underperformers lose \$10 due to overtrading and \$170 less due to bad decision-making in the *TASK* treatment as compared to the *NOFRICTION* treatment. Although neither value is statistically significant, they are still economically meaningful since we observe a statistically significant relation in the previous earnings tests.

We find that the *TASK* treatment helps underperformers improve their decision-making quality more than others. This improvement in decision-making, and not less overtrading, explains the reduction in underperformance experienced by underperformers in the *TASK* treatment.

< Table 6 here >

4.4. Components of bad decision-making

In this sub-section, we analyze the treatment effects on components of underperformance due to bad decision-making. Here, we wish to examine the mechanism through which our *TASK* treatment helps underperformers make better trading decisions.

We break down underperformance due to bad decision-making into two components: fundamental and non-fundamental. The fundamental component of bad decision-making captures the extent to which a participant underperforms the benchmark strategy in terms of earnings derived from an asset's fundamental value. In the multiple assets experiment, this component includes instances wherein a participant incorrectly guesses an asset's price increase probability and consequently does not trade in the same direction as the optimal strategy (e.g., the participant buys or holds the asset when the optimal strategy involves selling it). Additionally, even if the participant is able to guess the price increase probabilities correctly, they can still trade at inferior prices than the optimal strategy. In this case, underperformance due to buying (selling) a high (low) price increase probability at a higher (lower) price than the optimal strategy is also captured by the fundamental component. In the single asset experiment,

the fundamental component includes underperformance due to buying (selling) a low (high) payout asset at a higher (lower) price than the payout.

The non-fundamental component of bad decision-making includes all underperformance due to bad decision-making that is not explained by the fundamental component. In both experiments, losses due to buying high and selling low during the market are included in the non-fundamental component. The non-fundamental component is more relevant in the single asset experiment than in the multiple assets experiment. This is because, unlike the multiple assets experiment, in the single asset experiment, prices can deviate from fundamentals. In this case, a trader could profitably speculate on prices by buying a low payout asset at a low price and then selling it at a high price later during the market. However, a failure to implement this strategy successfully can cause the trader to underperform, e.g., if they buy high and sell low.

We start by examining the treatment effects on components of underperformance in the multiple assets experiment. We calculate the fundamental component of bad decision-making using the participants' portfolios at the end of a market. In the multiple assets experiment, we subtract the endowed units from the number of units of the asset in the participant's terminal portfolio. This gives us the participant's net position in the asset at the end of a market. We then multiply this net position by the difference between the terminal price of the asset and the volume-weighted average price at which they bought or sold the additional units in their terminal portfolio. This gives us the participant's earnings from the asset's fundamental value. We perform the same calculation for the optimal strategy. The difference between the fundamental value earnings figure for the optimal strategy and for the participant is our measure for the fundamental component of underperformance due to bad decision-making. We calculate the non-fundamental component as the difference between underperformance due to bad decision-making and the fundamental component.

Table 7 reports the regression results for the multiple assets experiment. The coefficient on UP_i is significantly higher in the regression with the fundamental component of bad decision-making than the non-fundamental component, indicating the fundamental component contributes more to the underperformance of the worst performers than the non-fundamental component. The *TASK* treatment also targets the fundamental component; the coefficient on the interaction term between $TASK_{j,k}$ and UP_i is significantly larger when regressed on the fundamental component than the non-fundamental component. This result implies that the *TASK* treatment

helps underperformers match the optimal strategy more closely in terms of the earnings they derive from the fundamental value of the asset, and this convergence to optimal is the primary reason behind their reduction in underperformance. In other words, the *TASK* treatment helps participants guess the price increase probabilities corresponding to each asset more accurately and more quickly than they were doing in the *NOFRICTION* treatment.

< Table 7 here >

Next, we perform the same tests for the single asset experiment. Since our benchmark strategy in the single asset experiment is a no-trade strategy, calculating the fundamental component is more straightforward than in the multiple assets experiment. Like the multiple assets experiment, we use the participant's portfolios at the end of a market to calculate the fundamental component of bad decision-making. We subtract the endowed units from the number of units of the asset in the participant's terminal portfolio to calculate their net position. We then multiply this net position by the difference between the volume-weighted average price at which they bought or sold the additional units in their terminal portfolio and the asset's payout. This difference directly gives us the fundamental component of underperformance due to bad decision-making. Additionally, like the multiple assets experiment, we calculate the non-fundamental component as the difference between underperformance due to bad decision-making and the fundamental component.

Table 8 reports the regression results for the single asset experiment. Unlike the multiple assets experiment, here, the non-fundamental component of bad decision-making explains underperformance of the worst performers more than the fundamental component, as indicated by the higher coefficient on UP_i . As discussed before, speculating on prices is a viable trading strategy in the single asset experiment but not in the multiple assets experiment. Hence, it is not entirely surprising that the non-fundamental component is more important in this experiment. Interestingly, unlike the multiple assets experiment, the *TASK* treatment targets the non-fundamental component more than the fundamental component in this experiment, as indicated by the significantly larger coefficient on the interaction term between $TASK_{j,k}$ and UP_i . In effect, the *TASK* treatment helps underperformers lose less money due to speculating on prices. Even the *SLOW* treatment helps underperformers lose less money due to price speculation, however,

this benefit is almost entirely neutralized by the increase in underperformance due to the fundamental component for underperformers. We investigate the effects of the *SLOW* treatment on the fundamental component further by breaking down the fundamental component into the buys and sells domain (see Table IA2 in Section IA1 of the Internet Appendix). We find that, in the *SLOW* treatment, underperformers have a tendency to sell high payout assets at relatively low prices. They would be better off holding these assets till the end of the market and earning the high payout.

< Table 8 here >

5. Discussion

We observe that participants are more deliberative while making trading decisions in both the *SLOW* and *TASK* treatments but this additional cognitive effort only brings tangible performance benefits to one set of actors, i.e., the most underperforming participants in our experiments, in one setting, i.e., the *TASK* treatment. What explains these results? We conjecture that the answer to this question lies in the differing nature and effects of the *SLOW* and *TASK* treatments. Both treatments increase cognitive effort, but one only provides a performance benefit for a subset of participants.

Unlike the time delay friction in the *SLOW* treatment, the cognitive task friction in the *TASK* treatment specifically asks participants their beliefs about fundamental value of the asset(s) being traded. Participants are more likely to be thinking about these fundamentals in the additional time they spend before each trade in the *TASK* treatment than in the *SLOW* treatment. This deliberation on fundamentals can reap performance benefits for participants who might not be using this information in the *NOFRICTION* treatment. We conjecture that the worst underperformers who benefit from the *TASK* treatment likely fit this bracket. These actors also increase cognitive effort in the *SLOW* treatment but do not receive any performance benefits. This is likely because in that treatment they do not necessarily think about the asset fundamentals. Like Enke et al. (2020), we find that additional cognitive effort by itself is not sufficient to improve trading performance, rather participants need to look at trading decisions in the “right way”. The *SLOW* treatment only helps participants increase cognitive effort, while the *TASK* treatment, to some degree, helps participants approach the trading decisions in the right

way by asking them about the asset fundamentals. This property makes the *TASK* treatment better than other treatments. However, even in the *TASK* treatment, only the worst performers receive performance benefits. It is possible that the worst performers are inattentive and do not use information about asset fundamentals well. Gargano and Rossi (2018) find that inattentive investors tend to perform badly in financial markets. The marginal benefit of increased deliberation on fundamentals in the *TASK* treatment is highest for these inattentive investors.

We attempt to characterize the right way of approaching trading decisions in our experiments and explain how our *TASK* treatment can help inattentive investors approach trades in the right way. In both experiments, it is important to consider information holistically before making trading decisions. Neglecting certain kinds of information and relying too much on other kinds of information can be costly. Previous studies show that inattentive investors either underweight or neglect certain pieces of information. In particular, these investors are likely to absorb and rely on information that is salient and neglect information that is nonsalient (Hirshleifer and Teoh, 2003). In both experiments, market prices are more salient than information about the asset's fundamentals, i.e., the price increase probabilities in the multiple assets experiment and the trader's private information (clue) and the payout distribution in the single asset experiment. This is because prices are dynamic, changing after every trading period in the multiple assets experiment and after every trade in the single asset experiment. These dynamic prices are continuously plotted on a graph on the trader's screen during the market. New price movements represented by newly created points on the graph are more eye-catching than the static price increase probabilities, clues, and payout distributions that are displayed in one corner of the screen (see Sections IA3 and IA4 in the Internet Appendix for screenshots of the participant's trading screens in the multiple assets experiment and the single asset experiment, respectively).

Salience of prices can cause inattentive investors to overweight recent price movements in their trading decisions. Previous literature in behavioral finance finds that investors display "recency bias" or "extrapolation bias", i.e., they overweight recent price movements in their expectations about future returns (Greenwood and Shleifer, 2014; Barberis et al., 2015). Simultaneously, information participants could be neglecting or underweighting the core value formation process for the asset(s) since information about asset fundamentals is relatively nonsalient. In the multiple assets experiment, the assets derive their value from their price increase probability. A trader who buys an asset based on a recent price rise expecting the price

to rise further can make losses if the asset has a historical trend of falling prices that the trader has ignored. In the single asset experiment, the asset derives its value from the terminal payout. By focusing mainly on recent price movements and ignoring the terminal payout, a trader might be tempted to buy (sell) a low (high) payout asset that has experienced recent price increases (decreases) in the expectation of a further price increase (decrease). This trader can end up making losses if they cannot liquidate (buyback) the asset at a higher (lower) price before the end of the market.

In both experiments, in the *TASK* treatment, our question about the trader's belief about the fundamental value of the asset can help direct their attention to the core value formation process of the asset. In the multiple assets experiment, we ask participants about the assets that they think correspond to the extreme price increase probabilities. To be able to answer this question correctly, a participant must look at the full history of price movements for all assets and not just the recent price movements. In doing so, the participant forms beliefs about the fundamental values of all assets that are less affected by recency bias. In the single asset experiment, we ask participants about the payout they believe to be correct. To be able to answer this question correctly, a participant needs to use their private information, derive the likelihood of the remaining payouts occurring, and use the public information in prices (not just recent prices). In doing so, the participant performs the mental calculus necessary to generate a belief about the fundamental value of the asset. In both cases, this belief about the fundamental value is likely to stay in the participant's memory when they are finalizing a trade. In this manner, the question in the *TASK* treatment can help draw the attention of inattentive investors to nonsalient information about asset fundamentals and potentially increase their use of such information in their trading decisions.

Our results seem to provide some evidence for this explanation. In the multiple assets experiment, only the *TASK* treatment helps the worst performers (most inattentive investors) more closely match the optimal strategy in terms of the accuracy and speed with which they guess the price increase probabilities for all assets (see Section 4.4). This is likely because these investors are able to form beliefs free from recency bias due to the question in the *TASK* treatment. Another indicator of this reduction in recency bias is the reduction in momentum-driven trades in the *TASK* treatment (see Section 4.1). The interpretation of results for the single asset experiment is slightly more complex. In the single asset experiment, both the *SLOW* and

TASK treatments help the worst performers lose less money due to speculating on prices (see Section 4.4). This result and a simultaneous reduction in momentum-driven trades implies that both these treatments are effective in helping inattentive investors avoid recency bias, i.e., buying (selling) assets after recent price increases (decreases) with the expectation of further price increases (decreases). However, since the *SLOW* treatment does not direct the attention of inattentive investors to the core value formation process, they fall into another trap in this treatment. These investors end up selling assets with a high payout at relatively low prices before the market ends (see Section 4.4). If these investors had the fundamental value in mind, they would realize that they would be better off holding these assets until the end of the market. This result shows that the *TASK* treatment's ability to direct investor attention to information about asset fundamentals is an important contributor to its success. In sum, the cognitive task friction in the *TASK* treatment is the only friction that offers any performance improvements to investors since it is the only one that stops inattentive investors from relying too much on recent price movements and also directs their attention to information about the asset's fundamentals.

We recognize that an exact replication of our cognitive task friction in real-world markets might not necessarily work as well. In our experiments, there is an inherent certainty about the fundamental value of an asset. Participants might have heterogeneous beliefs about the fundamental value; however, the actual fundamental value of an asset is rigid and pre-determined. This does not apply in real-world markets. Not only can the fundamental value of an asset be uncertain, it could also change over time. Consequently, cognitive tasks meant to induce deliberation on fundamentals cannot be as straightforward as our experiments. However, even in real-world financial markets, the same investor tendencies of overweighting salient information and recent price movements apply. We conjecture that some cognitive task type friction designed to help investors properly use this information can help improve performance. Such a friction could help investors, especially relatively unsophisticated investors, improve their decision-making by helping them take a step back, assess information globally, and properly consider important but nonsalient information.

References

- Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17, 223–249.
- Barber, B.M., Lee, Y.-T., Liu, Y.-J., Odean, T., 2009. Just how much do individual investors lose by trading? *Review of Financial Studies* 22, 609–632.
- Barber, B.M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55, 773–806.
- Barber, B.M., Odean, T., 2002. Online investors: Do the slow die first? *Review of Financial Studies* 15, 455–487.
- Barber, B.M., Odean, T., 2013. The behavior of individual investors. In: Constantinides G.M., Harris M. and Stulz R.M. (eds.) *Handbook of the Economics of Finance*. Elsevier, pp. 1533–1570.
- Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2015. X-CAPM: An extrapolative capital asset pricing model. *Journal of Financial Economics* 115, 1–24.
- Barlevy, G., 2003. Credit market frictions and the allocation of resources over the business cycle. *Journal of Monetary Economics* 50, 1795–1818.
- Baumeister, R.F., Bratslavsky, E., Muraven, M., Tice, D.M., 1998. Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology* 74, 1252–1265.
- Bogan, V., 2008. Stock market participation and the internet. *Journal of Financial and Quantitative Analysis* 43, 191–211.
- Choi, J.J., Laibson, D., Metrick, A., 2002. How does the internet affect trading? Evidence from investor behavior in 401(k) plans. *Journal of Financial Economics* 64, 397–421.
- Corgnet, B., Desantis, M., Porter, D., 2018. What makes a good trader? On the role of intuition and reflection on trader performance. *Journal of Finance* 73, 1113–1137.
- Enke, B., Gneezy, U., Hall, B., Martin, D., Nelidov, V., Offerman, T., Ven, J.v.d., 2020. Cognitive biases: Mistakes or missing stakes? Unpublished working paper.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10, 171–178.
- Gargano, A., Rossi, A.G., 2018. Does it pay to pay attention? *Review of Financial Studies* 31, 4595–4649.
- Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. *Review of Financial Studies* 27, 714–746.
- Hirshleifer, D., Teoh, S.H., 2003. Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics* 36, 337–386.
- Kahneman, D., 2011. *Thinking, Fast and Slow*. Farrar, Straus and Giroux.
- Kalda, A., Loos, B., Previtro, A., Hackethal, A., 2020. Smart(phone) investing? Unpublished working paper.
- Kocher, M.G., Lucks, K.E., Schindler, D., 2018. Unleashing animal spirits: Self-control and overpricing in experimental asset markets. *Review of Financial Studies*, 2149–2178.
- Laudenbach, C., Loos, B., Pirschel, J., Wohlfart, J., 2020. The trading response of individual investors to local bankruptcies. Unpublished working paper.
- Lo, A.W., Mamaysky, H., Wang, J., 2004. Asset prices and trading volume under fixed transactions costs. *Journal of Political Economy* 112, 1054–1090.

- Marguc, J., Förster, J., Van Kleef, G.A., 2011. Stepping back to see the big picture: When obstacles elicit global processing. *Journal of Personality and Social Psychology* 101, 883–901.
- Palan, S., 2015. GIMS—Software for asset market experiments. *Journal of Behavioral and Experimental Finance* 5, 1–14.
- Palomino, F., Renneboog, L., Zhang, C., 2009. Information salience, investor sentiment, and stock returns: The case of British soccer betting. *Journal of Corporate Finance* 15, 368–387.
- Plott, C.R., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. *Econometrica* 56, 1085–1118.
- Reher, M., Sokolinski, S., 2020. Does Fintech democratize investing? Unpublished working paper.
- Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. *Journal of Finance* 52, 35–55.
- Weber, M., Camerer, C.F., 1998. The disposition effect in securities trading: An experimental analysis. *Journal of Economic Behavior & Organization* 33, 167–184.

Table 1
Experimental design summary

This table presents details about our asset market designs and treatments. Panel A displays summary information about the asset market designs used in the multiple assets experiment and the single asset experiment. Panel B displays summary information about the implementation of our treatments in both experiments.

Detail	Multiple assets experiment	Single asset experiment
Panel A: Asset markets		
Original design	Weber and Camerer (1998)	Market 9 in Plott and Sunder (1988)
Trading mechanism	Trades at displayed market price	Continuous double auction
Price process	Exogenous	Endogenous
Number of experimental sessions	8	4
Number of participants	95	47
Number of markets per session	4	12
Trading periods per market	8	1
Number of assets	4	1
Endowment per asset	4	4
Cash endowment	\$1,000	\$1,000
Panel B: Treatments		
<i>HIGHCOST</i> treatment	Four times higher transaction cost	Five times higher transaction cost
<i>SLOW</i> treatment	Additional 1 min. per trading period	20-second waiting period post order
<i>TASK</i> treatment	Question related to asset value	Question related to asset value

Table 2
Trading behavior across treatments

This table reports descriptive statistics for various trading behavior related metrics across treatments in both experiments. The statistics reported are means for subject-market observations. Base statistics are only reported for the *NOFRICTION* treatment. Statistics for all other treatments are reported relative to the *NOFRICTION* treatment statistic. Buy (sell) trades occurring immediately after a price rise (fall) are classified as momentum trades. In the multiple assets experiment, time between orders (in seconds) in a market is calculated as the time of the last order divided by the number of orders. To calculate the adjusted time between orders (in seconds), we first set the time between orders for all orders that occur within 20 seconds of the previous order to 21 seconds and then calculate the average time between orders for the market. In the single asset experiment, the same calculation is performed in every trading period. Period-level values are averaged at the market-level. To minimize reader confusion, we report statistics for time between orders only where applicable. *t*-statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Experiment	<i>NOFRICTION</i> treatment	<i>HIGHCOST</i> treatment	<i>SLOW</i> treatment	<i>TASK</i> treatment
Number of trades	Multiple assets	14.09	-1.44 (-1.51)	+0.89 (0.95)	-3.44*** (-3.96)
	Single asset	7.90	-1.08 (-1.44)	-3.93*** (-6.49)	-3.55*** (-5.40)
Momentum trades	Multiple assets	7.99	-0.96 (-1.39)	+1.04 (1.47)	-1.36** (-2.17)
	Single asset	2.77	-0.34 (-0.91)	-1.28*** (-4.52)	-1.11*** (-3.53)
Time between orders (seconds)	Multiple assets	18.98	+1.07 (0.93)	+6.84*** (4.73)	NA
	Single asset	26.14	+5.72** (2.10)	+10.65*** (4.82)	NA
Adjusted time between orders (seconds)	Multiple assets	NA	NA	NA	NA
	Single asset	31.42	NA	+5.37*** (2.60)	NA
Time between orders (seconds, excluding first market round or trading period)	Multiple assets	20.09	NA	NA	+9.36*** (6.87)
	Single asset	25.13	NA	NA	+8.16*** (3.26)

Table 3**Treatment effects on earnings in the multiple assets experiment**

This table reports regression results testing the effects in the multiple assets experiment of our treatments on participant earnings for all participants and specifically for underperformers. The unit of observation is a subject i in market j of session k . $Earnings_{i,j,k}$ is the participant's earnings (in laboratory \$) in the market. We calculate this earnings figure relative to the optimal strategy by deducting the optimal earnings from the participant's earnings. $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance in the entire experiment. Regressions reported in this table control for market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable = $Earnings_{i,j,k}$	
	(1)	(2)
$HIGHCOST_{j,k}$	-650.91*** (-8.05)	-608.36*** (-7.82)
$SLOW_{j,k}$	-260.61*** (-3.22)	-261.55*** (-3.43)
$TASK_{j,k}$	115.86 (1.43)	-1.30 (-0.02)
UP_i		-788.62*** (-7.02)
$HIGHCOST_{j,k} \times UP_i$		-164.70 (-0.99)
$SLOW_{j,k} \times UP_i$		-5.60 (-0.04)
$TASK_{j,k} \times UP_i$		454.36*** (2.74)
R^2	26.80%	52.25%
Fixed effects	Sequence	Sequence
Observations	380	380

Table 4**Treatment effects on earnings in the single asset experiment**

This table reports regression results testing the effects in the single asset experiment of our treatments on participant earnings for all participants and specifically for underperformers. The unit of observation is a subject i in market j of session k . $Earnings_{i,j,k}$ is the participant's earnings (in laboratory \$) in the market. $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance or the top-quartile for underperformance in their experimental session. Control variables include two indicator variables for participants receiving a clue that the payout is not \$50 and \$490, and two indicator variables for the payout being \$50 and \$490. Regressions reported in this table also control for session and market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable = $Earnings_{i,j,k}$	
	(1)	(2)
$HIGHCOST_{j,k}$	-134.14*** (-2.62)	-94.90* (-1.75)
$SLOW_{j,k}$	16.62 (0.32)	10.06 (0.18)
$TASK_{j,k}$	17.62 (0.35)	-28.32 (-0.52)
UP_i		-381.08*** (-5.03)
$HIGHCOST_{j,k} \times UP_i$		-154.31 (-1.44)
$SLOW_{j,k} \times UP_i$		25.88 (0.24)
$TASK_{j,k} \times UP_i$		180.15* (1.69)
R^2	70.25%	75.11%
Controls	Signal, Payout	Signal, Payout
Fixed effects	Session, Sequence	Session, Sequence
Observations	564	564

Table 5

Treatment effects on overtrading and bad decision-making in the multiple assets experiment

This table reports regression results testing the effects in the multiple assets experiment of our treatments on underperformance due to overtrading and bad decision-making for underperformers. The unit of observation is a subject i in market j of session k . $Overtrade_{i,j,k}$ is the participant's underperformance (in laboratory \$) due to overtrading. $Overtrade_{i,j,k}$ equals the transaction costs paid by the participant in excess of the optimal transaction costs. $BadDM_{i,j,k}$ is the participant's underperformance (in laboratory \$) due to bad decision-making. $BadDM_{i,j,k}$ is the participant's underperformance (in laboratory \$) due to bad decision-making. $BadDM_{i,j,k}$ is the difference between the gross earnings (with transaction costs added back) from the optimal strategy and the participant's gross earnings. $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance or the top-quartile for underperformance in the entire experiment. Regression results reported in this table control for market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable =	
	$Overtrade_{i,j,k}$	$BadDM_{i,j,k}$
$HIGHCOST_{j,k}$	177.72*** (15.24)	430.64*** (5.52)
$SLOW_{j,k}$	2.70 (0.24)	258.86*** (3.39)
$TASK_{j,k}$	-21.61* (-1.85)	20.32 (0.26)
UP_i	1.18 (0.07)	787.44*** (7.00)
$HIGHCOST_{j,k} \times UP_i$	21.18 (0.85)	143.52 (0.86)
$SLOW_{j,k} \times UP_i$	6.68 (0.29)	-12.29 (-0.08)
$TASK_{j,k} \times UP_i$	16.11 (0.65)	-470.47*** (-2.83)
R^2	60.40%	46.95%
Fixed effects	Sequence	Sequence
Observations	380	380

Table 6

Treatment effects on overtrading and bad decision-making in the single asset experiment

This table reports regression results testing the effects in the single asset experiment of our treatments on underperformance due to overtrading and bad decision-making for underperformers. The unit of observation is a subject i in market j of session k . $Overtrade_{i,j,k}$ is the participant's underperformance (in laboratory \$) due to overtrading. $Overtrade_{i,j,k}$ equals the transaction costs paid by the participant in the market. $BadDM_{i,j,k}$ is the participant's underperformance (in laboratory \$) due to bad decision-making. $BadDM_{i,j,k}$ is the difference between the earnings from a no trade strategy and the participant's gross earnings (with transaction costs added back). $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance or the top-quartile for underperformance in their experimental session. Control variables include two indicator variables for participants receiving a clue that the payout is not \$50 and \$490, and two indicator variables for the payout being \$50 and \$490. Regression results reported in this table control for session and market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable =	
	$Overtrade_{i,j,k}$	$BadDM_{i,j,k}$
$HIGHCOST_{j,k}$	119.54*** (12.23)	24.64 (-0.45)
$SLOW_{j,k}$	-14.70*** (-1.48)	4.65 (0.08)
$TASK_{j,k}$	-14.69 (-1.50)	43.00 (0.79)
UP_i	17.22 (1.26)	363.86*** (4.80)
$HIGHCOST_{j,k} \times UP_i$	58.61*** (3.05)	95.70 (0.90)
$SLOW_{j,k} \times UP_i$	-6.83 (-0.36)	-19.04 (-0.18)
$TASK_{j,k} \times UP_i$	-9.66 (-0.50)	-170.49 (-1.60)
R^2	48.75%	21.10%
Controls	Signal, Payout	Signal, Payout
Fixed effects	Session, Sequence	Session, Sequence
Observations	564	564

Table 7

Treatment effects on components of bad decision-making in the multiple assets experiment

This table reports regression results testing the effects in the multiple assets experiment of our treatments on components of underperformance due to bad decision-making for underperformers. The unit of observation is a subject i in market j of session k . $Fundamental_{i,j,k}$ is the fundamental component of underperformance due to bad decision-making (in laboratory \$). $Fundamental_{i,j,k}$ is calculated as the difference between the terminal portfolio earnings for all assets from the optimal strategy and the participant's terminal portfolio earnings. To calculate terminal portfolio earnings for each asset, we first subtract the endowed units from the number of units of the asset in the participant's terminal portfolio. We then multiply this term by the difference between the terminal price of the asset and the volume-weighted average price at which the participant bought or sold the additional units. $Non\ fundamental_{i,j,k}$ is the non-fundamental component of underperformance due to bad decision-making (in laboratory \$). $Non\ fundamental_{i,j,k}$ is the difference between total underperformance due to bad decision-making and the fundamental component of underperformance due to bad decision-making. $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance or the top-quartile for underperformance in the entire experiment. Regression results reported in this table control for market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable =	
	$Fundamental_{i,j,k}$	$Non\ fundamental_{i,j,k}$
$HIGHCOST_{j,k}$	403.76*** (5.80)	26.88 (0.72)
$SLOW_{j,k}$	226.39*** (3.32)	32.47 (0.88)
$TASK_{j,k}$	221.77*** (3.19)	-201.46*** (-5.36)
UP_i	662.59*** (6.59)	124.86** (2.30)
$HIGHCOST_{j,k} \times UP_i$	256.88* (1.73)	-113.36 (-1.42)
$SLOW_{j,k} \times UP_i$	13.03 (0.10)	-25.31 (-0.34)
$TASK_{j,k} \times UP_i$	-395.90*** (-2.67)	-74.57 (-0.93)
R^2	43.59%	29.53%
Fixed effects	Sequence	Sequence
Observations	380	380

Table 8

Treatment effects on components of bad decision-making in the single asset experiment

This table reports regression results testing the effects in the single asset experiment of our treatments on components of underperformance due to bad decision-making for underperformers. The unit of observation is a subject i in market j of session k . $Fundamental_{i,j,k}$ is the fundamental component of underperformance due to bad decision-making (in laboratory \$). To calculate $Fundamental_{i,j,k}$, we first subtract the endowed units from the number of units of the asset in the participant’s terminal portfolio. We then multiply this term by the difference between the volume-weighted average price at which the participant bought or sold the additional units and the asset payout. $Non\ fundamental_{i,j,k}$ is the non-fundamental component of underperformance due to bad decision-making (in laboratory \$). $Non\ fundamental_{i,j,k}$ is the difference between total underperformance due to bad decision-making and the fundamental component of underperformance due to bad decision-making. $HIGHCOST_{j,k}$ is an indicator variable that equals one if the market has the high transaction cost treatment. $SLOW_{j,k}$ is an indicator variable that equals one if the market has the time delay treatment. $TASK_{j,k}$ is an indicator variable that equals one if the market has the cognitive effort task treatment. UP_i is an indicator variable that equals one if the participant is in the bottom-quartile for performance or the top-quartile for underperformance in their experimental session. Control variables include two indicator variables for participants receiving a clue that the payout is not \$50 and \$490, and two indicator variables for the payout being \$50 and \$490. Regression results reported in this table control for session and market sequence fixed effects. t -statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

Variable	Dependent variable =	
	$Fundamental_{i,j,k}$	$Non\ fundamental_{i,j,k}$
$HIGHCOST_{j,k}$	3.48 (0.07)	-28.12 (-1.24)
$SLOW_{j,k}$	-29.16 (-0.58)	33.80 (1.47)
$TASK_{j,k}$	5.89 (0.12)	37.12 (1.64)
UP_i	112.63 (1.62)	251.23*** (7.93)
$HIGHCOST_{j,k} \times UP_i$	11.53 (0.12)	84.17 (1.88)
$SLOW_{j,k} \times UP_i$	157.56 (1.61)	-176.60*** (-3.96)
$TASK_{j,k} \times UP_i$	1.28 (0.01)	-171.77*** (-3.85)
R^2	13.07%	26.43%
Controls	Signal, Payout	Signal, Payout
Fixed effects	Session, Sequence	Session, Sequence
Observations	564	564

Panel A: Multiple assets experiment

The screenshot shows a dialog box with a title bar labeled "Question". The main text asks: "Which asset number do you think corresponds to the following price rise probabilities?". Below this text, there are two rows of input fields. The first row has the probability "0.65" followed by a light blue rectangular input box. The second row has the probability "0.35" followed by another light blue rectangular input box. In the bottom right corner of the dialog box, there is a grey button labeled "OK".

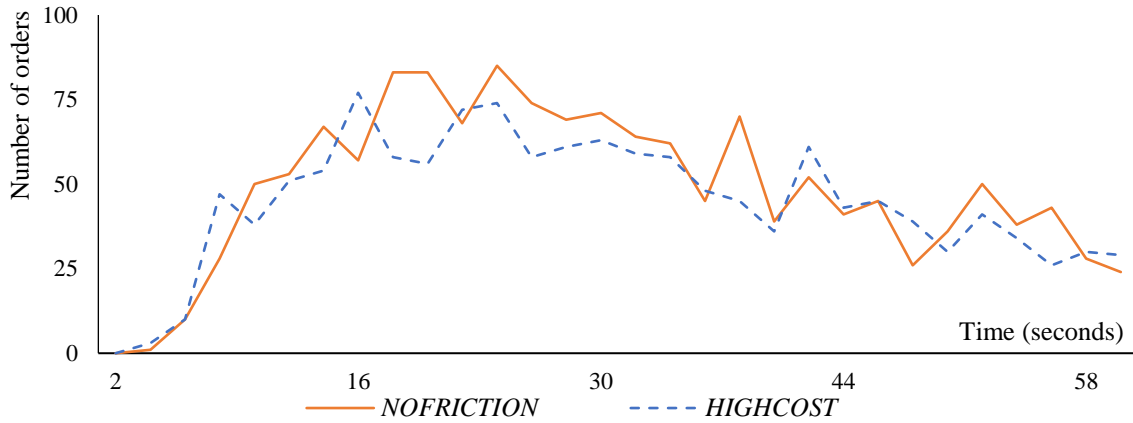
Panel B: Single asset experiment

The screenshot shows a dialog box with a title bar labeled "Question". The main text asks: "What do you think the true payoff of the asset will be?". Below this text, there are three radio button options: "\$50.00", "\$240.00", and "\$490.00". In the bottom right corner of the dialog box, there is a grey button labeled "OK".

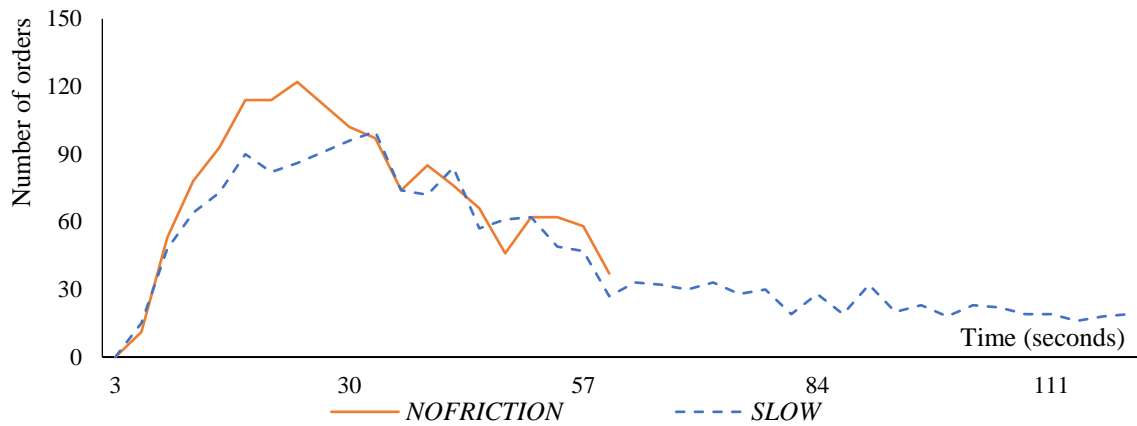
Figure 1. Screenshots of the cognitive effort task questions in the *TASK* treatment.

This figure shows the incentivized cognitive effort tasks that participants performed before placing an order in the multiple asset and single asset experiments. Panel A shows the question asked in the multiple assets experiment. Participants were required to answer this question before they were allowed to place any orders in a given trading period. The participants only needed to answer this question once, even if they placed multiple orders in the period. Panel B shows the question asked in the single asset experiment. Participants were required to answer this question before they were allowed to place an order. If participants placed multiple orders in a period, they needed to answer this question multiple times, before placing each order.

Panel A: *NOFRICTION* and *HIGHCOST* treatments



Panel B: *NOFRICTION* and *SLOW* treatments



Panel C: *NOFRICTION* and *TASK* treatments






Figure 2. Order submission activity through time in the multiple assets experiment.

This figure plots the number of orders submitted by participants in the multiple assets experiment across time. Panel A plots the number of orders in the *NOFRICTION* and *HIGHCOST* treatments in two-second intervals. Panel B plots the number of first orders in the *NOFRICTION* and *SLOW* treatments in three-second intervals. Panel C plots the number of orders (excluding the first trading period) in the *NOFRICTION* and *TASK* treatments in two-second intervals.

Panel A: NOFRICTION and HIGHCOST treatments



Panel B: NOFRICTION and SLOW treatments



Panel C: NOFRICTION and TASK treatments

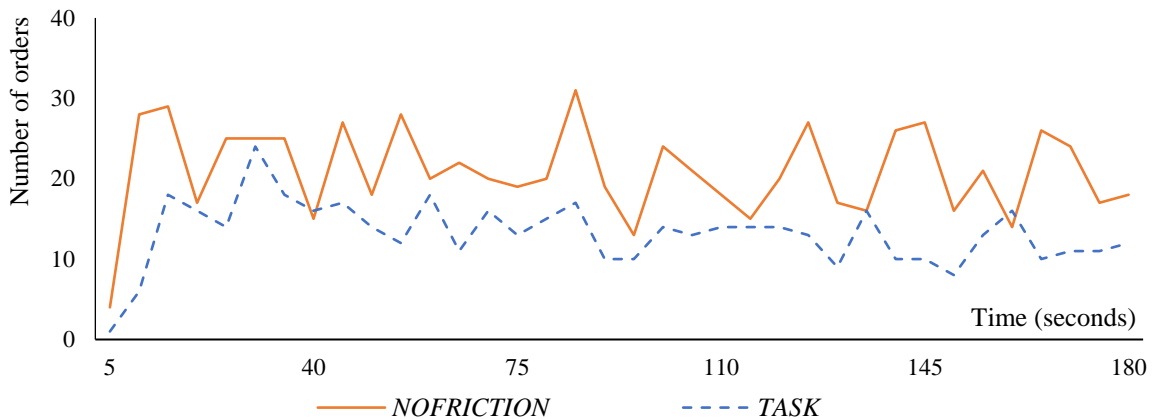


Figure 3. Order submission activity through time in the single asset experiment.

This figure plots the number of orders submitted by participants per market in the single asset experiment across time. All panels plot the number of orders in five-second intervals. Panel A plots the number of orders in the *NOFRICTION* and *HIGHCOST* treatments. Panel B plots the number of orders in the *NOFRICTION* and *SLOW* treatments. Panel C plots the number of orders (excluding the first market round) in the *NOFRICTION* and *TASK* treatments.