Investor biases in financial markets

Anirudh Dhawan

Ph.D. supervisor: Prof. Talis Putnins

A thesis submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Finance Discipline Group, UTS Business School

University of Technology Sydney

19th August 2021

Certificate of original authorship

I, Anirudh Dhawan, declare that this thesis, titled "Investor biases in financial markets", is submitted in fulfilment of the requirements for the award of the degree of Doctor of Philosophy in the UTS Business School at the University of Technology Sydney.

This thesis is wholly my own work unless otherwise referenced or acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis.

This document has not been submitted for qualifications at any other academic institution.

This research is supported by the Australian Government Research Training Program.

Signature: Production Note:
Signature removed prior to publication.

Date: 19th August 2021

To my Amma and Appu

Acknowledgements

I want to thank my Ph.D. supervisor, Talis Putnins, for being a remarkable research guide and mentor for me. I am particularly grateful to Talis for inculcating in me an understanding and appreciation of high quality research. I appreciate his patience throughout my Ph.D. journey and thank him for his interest in mentoring me. I am also grateful to my Honors supervisor, Maria Kim, for introducing me to the world of academic research.

I am thankful to Christina Sklibosios and David Michayluk for all the institutional support they offered during my time at the University of Technology Sydney (UTS). I am also thankful to the UTS Finance Discipline Group staff, including Andrea Greer, Duncan Ford, Mala Kapahi, and Stephanie Ough, for all their administrative support. I would also like to thank the faculty members and staff who manage the UTS Behavioral Laboratory for their assistance in conducting the laboratory experiments that are part of this thesis.

I gratefully acknowledge the funding provided by the RoZetta Institute, the Finance Discipline Group, the Behavioral Lab, and the Graduate Research School. I am also grateful to Binance, Kaiko, Securities Industry Research Centre of Asia-Pacific, Thomson Reuters for providing access to the data used in this thesis.

This thesis has benefitted greatly from the suggestions offered by Adrian Lee, Benjamin Loos, Christine Parlour, David Easley, David Yermack, Elvira Sojli, Gerhard Hambusch, Lee Smales, Marco Navone, Peter Bossaerts, Petko Kalev, Sean Foley, Susan Thomas, Terry Walter, and Vinay Patel.

I would like to thank my fellow Ph.D. students and, in particular, fellow members of the "Noisy corner" group—Chung Mai, Huong Nguyen, Linh Do, Mai Luong, Man Nguyen, and Prateek Daniels—for making my time at UTS so enjoyable. I also thank my friends Aaditya Jain, Atiqur Rasel, Hansini De Fonseka, and Joanita Wibowo for all their support.

I am grateful to Cyrus and Kainaz Balsara for their hospitality during my stay in Sydney. I am especially thankful to my aunt and uncle, Dipoli and Vinay Malhotra, for all their support during my time in Australia. Lastly, I would like to express my sincere gratitude to my parents, Sanjana and Vimal Dhawan, for the love, encouragement, and patience they have displayed in my academic journey and beyond.

Preface

Chapters 2–4 in this thesis have each been developed into an individual co-authored working paper. The working paper version of Chapter 2 is currently at a revise-and-resubmit stage at the *Review of Finance*. Two of these working papers have been presented or have been selected for presentation at various academic conferences. The list of working papers and conference presentations is as follows:

1. Dhawan, A., Putnins, T., 2021. "A new wolf in town? Pump-and-dump manipulation in cryptocurrency markets" (Chapter 2)
    - UWA Blockchain, Cryptocurrency, and Fintech Conference 2019
    - Australasian Finance and Banking Conference (AFBC) 2019
    - IGIDR-IIM Udaipur Market Microstructure Workshop 2020
    - Financial Intermediation Research Society (FIRS) Conference 2020 (canceled due to the COVID-19 pandemic)

2. Dhawan, A., Loos, B., Navone, M., Putnins, T., 2021. "Getting burned by frictionless financial markets" (Chapter 3)
    - Financial Research Network (FIRN) Ph.D. Symposium 2020
    - Economic Science Association (ESA) Meeting 2021
    - Peking University International Young Finance Scholars' Conference 2021
    - Academy of Behavioral Finance & Economics Annual Meeting 2021 (scheduled)
    - Financial Management Association (FMA) Annual Meeting 2021 (scheduled)
    - Melbourne Asset Pricing Meeting 2021 (scheduled)

3. Dhawan, A., Putnins, T., 2021. "Attention to information releases" (Chapter 4)

Table of contents

List of tables

List of figures

List of abbreviations

AMM	Automated Market Maker
API	Application Programming Interface
ASX	Australian Securities Exchange
AUD	Australian Dollar
BPS	Big Pump Signal
BTC	Bitcoin
CPT	Cumulative Prospect Theory
EDGAR	Electronic Data Gathering, Analysis, and Retrieval System
GMT	Greenwich Mean Time
ICO	Initial Coin Offering
NFT	Non-Fungible Token
OLS	Ordinary Least Squares
ONYSAG	Office of the New York State Attorney General
PA	PumpAnalysis
PDF	Probability Density Function
PEAD	Post-Earnings Announcement Drift
S&P	Standard and Poor's
SIRCA	Securities Industry Research Centre of Asia-Pacific
SVI	Google Search Volume Index
US	United States of America
USD	US Dollar
WE	WalletExplorer

Abstract

Behavioral biases are a key determinant of investors' wealth outcomes in financial markets. However, there is still much to learn about the nature of these biases. This thesis studies different types of financial markets to understand the individual- and aggregate-level effects of specific biases and how these biases are influenced by the environment.

The first study examines how biases drive people to participate in pump-and-dump manipulation games in cryptocurrency markets. Cryptocurrency manipulators publicly announce the intended manipulations and invite others to join them. In a simple framework, we show that these pumps are inherently negative-sum games for non-manipulators, as manipulators have an advantage. So, why do people participate in these schemes? Rational agents do not participate unless they have a skill or speed advantage. However, overconfident agents and gamblers participate, even without any advantage. We find strong empirical support for both mechanisms. Pumps generate extreme price and volume distortions, and cause large wealth transfers between participants.

The second study asks whether investor loss-making tendencies are influenced by the trading environment, particularly the trend towards frictionless access to markets. We hypothesize that adding back certain "trading frictions" in markets can make investors think harder and mitigate losses arising from impulsive trading. Using laboratory experiments, we examine how investor performance is impacted by various frictions: transaction costs, time delays in placing orders, and tasks requiring cognitive effort. High costs and time delays have no effect or harm performance, whereas cognitive tasks benefit participants who are most prone to underperforming. We conclude that frictions can yield performance benefits if they help inattentive investors consider nonsalient fundamental information they might otherwise neglect.

The third study investigates the aggregate effects of attention to fundamental information in stock markets. To isolate attention to information, we propose a new measure that involves classifying investors as inattentive when they do not cancel or update their pending orders after material overnight news events. We find strong underreactions and price drifts after material corporate announcements for stocks that receive less attention. These drifts cannot be explained

by other mispricing effects, such as mispricing after earnings announcements and for stocks favored by retail investors.

In summary, this thesis contributes by investigating how biases drive participation in loss-making schemes in financial markets, how certain elements in the market architecture can reduce investor losses by directing their attention to fundamental information, and the aggregate effects of inattention to fundamental information in stock markets.

Chapter 1: Introduction

"The investor's chief problem – and even his worst enemy – is likely to be himself."

− Benjamin Graham

Homo economicus, i.e., the human species that populates rational models in finance and economics, does not lose money easily. *Homo economicus* individuals do not gamble or purchase lottery tickets or hold on to losing stocks indefinitely because of a reluctance to realize losses. *Homo sapiens*, i.e., the human species that populates planet Earth, on the other hand, is not as disciplined as their theoretical cousin. They see value in gambling, purchasing lottery tickets, and delaying loss realizations. Casino owners and lottery ticket sellers in rational expectations models are penurious, whereas casino owners and lottery ticket sellers on Earth are prospering.

In addition to gambling, lotteries, and aversion to losses, *Homo sapiens* differ from *Homo economicus* individuals in various other respects. For example, *Homo economicus* individuals believe they are no better in terms of ability than an average person, pay full attention to all types of relevant information, and accurately weight all pieces of information in their decision-making. In contrast, *Homo sapiens* individuals overestimate their own abilities, pay more attention to information that easily catches their eye, and overweight recent events over historical ones, even if both types of events are equally important for a decision. In combination, these tendencies, among others, lead *Homo sapiens* to make biased and suboptimal decisions in certain contexts, such as financial markets.

One of the primary goals of behavioral finance and behavioral economics is to bridge the gap between *Homo economicus* and *Homo sapiens*. Behavioral finance aims to model and empirically examine the behavior of agents as it occurs in real-world financial settings, with an emphasis on deviations from optimal, rational behavior, and investigates how this behavior in turn impacts the financial ecosystems in which these agents exist. This thesis is another step in that direction.

Broadly, this thesis studies investor decision-making and its effects on financial markets. The substantive research contributions of this thesis are presented in Chapters 2–4. Each of these chapters presents a distinct research study. Chapter 2 demonstrates how certain behavioral biases, namely overconfidence and gambling preferences, propel agents to participate in speculative, negative expectation pump-and-dump manipulation games in cryptocurrency markets. Chapter 3 examines whether structural elements such as trading frictions can sometimes cause investors in financial markets to think harder before trading, thereby helping them reduce behavioral biases and improve the quality of their trading decisions. Finally, Chapter 4 investigates how investors' inability to be fully attentive to new information in markets affects stock prices.

The remainder of this chapter selectively reviews, at a broad level, particular branches of the behavioral finance literature to which this thesis contributes.

1.1 Investors and their cognitive biases

After years of being considered as a non-mainstream line of research in finance, behavioral research that studies how individuals' cognitive constraints and predilections influence their financial decisions and outcomes has gained prominence in recent times. Since its integration into the mainstream finance literature, research in this domain has made considerable advances in identifying the nuances of psychological biases exhibited by financial decision-makers. Researchers have found that these biases seep into all domains of finance that involve individual decision-making. For example, both investors and company managers can exhibit overconfidence; investor overconfidence results in excessive trading activity, whereas managerial overconfidence results in excessive takeover activity.[1] In financial markets, individuals generally favor stocks that are in the same industry as their profession, stocks that are located in their local region, and stocks that attract their attention, such as those that have recently hit upper price limits.[2] In terms of the determinants of biases, research in both the finance and psychology fields has found that such behavioral biases tend to occur more frequently when people rely on intuitive cognitive processing, i.e., processing that is fast-paced and automatic, rather than deliberative cognitive processing, i.e., processing that is slow-paced and analytical (Kahneman, 2011; Kocher, Lucks, and Schindler, 2018). In terms of the real effects of biases, they can cause investors to lose money in financial markets, either through overtrading and overpaying transaction costs or through trading losses caused by inferior trading decisions (Barber and Odean, 2000; Barber et al., 2009; Barber and Odean, 2013).

[1] See Barberis and Thaler (2002) for a thorough review of the behavioral finance literature, spanning across financial markets, investment management, and corporate finance.
[2] See Barber and Odean (2013) for an extensive review of the literature on individual investor biases and decision-making in financial markets.

Although investors in financial markets exhibit numerous biases, not all are relevant to this thesis. In this section, we introduce the reader to the psychological tendencies exhibited by investors that are the most pertinent to this thesis: overconfidence, gambling preferences, and limited attention.

Barber and Odean (2000) are the first to document the impact of overconfidence on trading decisions. They find that individual investors who exhibit the worst trading performance also trade the most frequently and, thus, pay the highest aggregate transaction costs. They conjecture that these investors are overconfident in that they overestimate the precision and quality of their information set. These overconfident traders are driven to trade excessively since they believe their information justifies a trade, even though the information is actually too imprecise to warrant a trade. This form of overconfidence is known as miscalibration. Deaves, Lüders, and Luo (2008) confirm, in a laboratory setting, the conjecture proposed in Barber and Odean (2000) that miscalibration induces excessive trading. In addition, they test whether two other forms of overconfidence, namely the better-than-average effect and the illusion of control, also affect trading frequency. The better-than-average effect refers to the tendency to believe that one's ability is better than the average person's ability, whereas the illusion of control refers to the mistaken belief that one is in control of external events. Deaves et al. (2008) find that, in addition to miscalibration, the better-than-average effect also leads to excessive trading activity (see also Dorn and Huberman, 2005; Glaser and Weber, 2007).

Some investors enjoy gambling and use stock markets as alternatives to traditional gambling venues (Kumar, Nguyen, and Putnins, 2021). Gamblers display preferences that are consistent with Cumulative Prospect Theory and are attracted to positively skewed, lottery-like payoff structures, i.e., distributions wherein there is a low probability of attaining a high positive payoff and a high probability of attaining a low negative payoff (Tversky and Kahneman, 1992; Barberis, 2012). Another related concept is sensation-seeking behavior. Sensation-seekers derive pleasure from novel, risky, and enticing experiences, such as driving at high speeds or riding roller coasters (Zuckerman, 1994). Because of the risks involved, sensation-seekers are also attracted to gambling (Raylu and Oei, 2002). In financial markets, both these actors (gamblers and sensation-seekers) tend to use trading as a source of entertainment and pleasure rather than as a method of managing investments, risks, and liquidity. In terms of trading behavior, overconfident actors, gamblers, and sensation-seekers trade more frequently than others (Dorn and Sengmueller, 2009; Grinblatt and

Keloharju, 2009). In addition, they are attracted to stocks that have a positively skewed, lottery-like payoff structure, i.e., stocks with low prices, high idiosyncratic volatility, and high idiosyncratic skewness (Kumar, 2009).

The stock selection process followed by investors, particularly retail investors, is influenced by cognitive constraints and limited attention. Limited attention refers to investors' inability to fully incorporate all relevant pieces of information into their trading decisions, either because their cognitive resources can only accommodate limited information or because, due to cognitive constraints, they adopt a hasty information search that fails to yield all relevant information. Barber and Odean (2008) explain how limited attention influences the stock selection process for individual investors. They state that retail investors face a significant search problem when selecting stocks to purchase, as the full population of stocks is large. They posit that investors simplify this search problem by first narrowing the selection to stocks that easily attract their attention, i.e., attention-grabbing stocks, instead of considering each stock in the population individually. These attention-grabbing stocks include stocks that are in the news, stocks that have recently hit their upper price ceilings, and stocks with names starting with the beginning letters of the alphabet, among others (Seasholes and Wu, 2007; Barber and Odean, 2008; Itzkowitz, Itzkowitz, and Rothbort, 2016). In sum, retail investors tend to favor attention-grabbing stocks over other stocks. Barber and Odean (2008) also explain that retail investors do not face a similar problem when selling since they mostly sell stocks they already own and do not short sell stocks.

Chapter 2 of this thesis examines how two of the three aforementioned investor biases, namely overconfidence and gambling preferences, affect investors' decision to participate in a unique type of manipulation game that occurs in cryptocurrency markets. In cryptocurrency markets, market manipulators invite market participants to join them in pump-and-dump manipulation schemes on various coins. Manipulators inform these participants about the exchange and time of the intended manipulation but not about the coin being manipulated. They only reveal the coin at the time of the manipulation. In a simple theoretical framework, we find that rational actors recognize that these games are unprofitable in expectation for non-manipulators who do not know *ex-ante* their speed of entering and exiting pumps, as manipulators can build a position beforehand and unload it on the non-manipulators. Unless rational actors know *ex-ante* that they are faster than other non-manipulators, they do not participate in these schemes. In contrast, overconfident actors and gamblers participate in pumps. Overconfident actors participate because they *believe* they are faster than others at exiting pumps (better-than-average effect), whereas gamblers join because they enjoy the positively skewed payoff structure that they can create from participating in a series of pumps. We document that overconfidence and gambling preferences can also explain participation in pumps empirically. Pump participation increases after non-manipulators become more overconfident and after more gambling money flows into the cryptocurrency ecosystem.

1.2 Investor biases and the trading environment

An interesting fact documented in the behavioral finance literature is that investor biases do not occur in a vacuum and are susceptible to the trading environment. One of the most robust behavioral biases in finance, the disposition effect, or the tendency for investors to sell winning stocks early and hold on to losing stocks for too long, is sensitive to the display of information and the nature of orders in the market. Both Frydman and Rangel (2014) and Frydman and Wang (2020) find that the disposition effect dissipates when the purchase price is made less salient. Additionally, Fischbacher, Hoffmann, and Schudy (2017) demonstrate in a laboratory setting that the disposition effect diminishes in markets where stop-loss and take-gain orders automatically execute once the pre-specified loss or gain limit is reached, but not in markets where participants are merely reminded to sell when the loss or gain limit is reached. Moreover, the effects of biases, e.g., risk-taking behavior, excessive trading activity, etc., depend on the trading environment. For example, Bazley, Cronqvist, and Mormann (2021) observe that investors take less risk when past price information is displayed in red. Similarly, Kalda et al. (2021) report that investors take more risk when trading on their smartphones. Finally, Borsboom et al. (2021) find that investors trade more frequently when their trading screens display price charts with shorter time frames.

Chapter 3 examines whether certain structural elements in markets, particularly trading frictions, affect trading behavior, including biases and loss-making tendencies. We conjecture that some trading frictions can induce deliberation before trades and thus help investors reduce the behavioral biases caused by an overreliance on intuitive thinking. The rationale for this conjecture originates from the psychology literature, which reports that some obstacles encountered before tasks help improve performance on these tasks because they enable people to process information globally, i.e., step back and see the "big picture." In this study, we select three trading frictions—high transaction costs, forced time delays, and cognitive effort tasks—and test their effects on investor decision-making quality and trading performance in two laboratory experiments. Of the three frictions, only one, i.e., cognitive tasks, helps improve the performance and decision quality for a select group of actors, i.e., participants who exhibit the worst trading performance. We conjecture that the worst performers are inattentive and underweight information about asset fundamentals that is relatively nonsalient in our experiments. We conclude that trading frictions can help improve the performance and decision quality of a trader if they help direct their attention to information that they would neglect in the absence of the frictions.

1.3 Aggregate effects of biased investor trading

Biases matter for individual-level outcomes, but do they also matter at the aggregate stock- or market-level? Yes. The most notable example is the post-earnings announcement drift, which refers to the lengthy price drifts following earnings announcements caused by investors consistently underreacting to these announcements (Ball and Brown, 1968; Bernard and Thomas, 1990; Ball and Brown, 2019). These drifts last longer when investors are inattentive, suggesting that investor inattention contributes to the initial underreaction to earnings announcements. For example, DellaVigna and Pollet (2009) find that post-announcement drifts are stronger for earnings announcements made on Friday, when investors are likely to be the most inattentive in the entire week. In addition, Hirshleifer, Lim, and Teoh (2009) show that price drifts are longer when more earnings announcements are released in a day, suggesting that limited attention explains underreaction since investors underreact more when they are more attention-constrained.

Chapter 4 examines how investor attention to corporate information releases affects price reactions to market announcements. Given the importance of attention to fundamental information at the individual level documented in Chapter 3, we also examine the effects of attention to information at the aggregate level. We use a novel measure for attention to information to determine a stock's baseline attention level. We develop a new attention measure because the existing measures confound attention to information and attention to returns, making it difficult to isolate the effects of attention to information on price reactions to announcements. Our measure is designed to disentangle attention to information from attention to returns. In essence, our measure classifies investors as inattentive if they fail to cancel or update their pending orders after overnight, material corporate disclosures or overnight material news media releases. We find that stocks receiving low investor attention experience stronger drifts after material announcements than other stocks, particularly after positive announcements. We verify that these attention effects exist independent of other mispricing effects documented in the literature.

Chapter 2 also examines some aggregate effects of investor biases. As discussed previously, this chapter shows that behavioral biases such as overconfidence and gambling preferences cause investors to participate in pump-and-dump manipulation schemes in cryptocurrency markets. These schemes cause sharp temporary price distortions for the manipulated coins and significantly increase their trading activity and volatility. These price and volume effects are directly proportional to the participation rates in these schemes, suggesting that certain behavioral biases can, at times, contribute to price and volume distortions in markets.

1.4 Why should we care about biased investor trading?

The research findings presented in this thesis are timely, pertinent, and topical for various reasons. First, the rapid uptake of online trading has been accompanied by the mass arrival of novice individual traders in stock markets (Bogan, 2008). For example, Robinhood, an American retail broker, welcomed 1.5 million first-time investors to their trading platform in the first four months of 2020.[3] Novice traders are more prone to the behavioral biases we study (see, for example, Da Costa et al., 2013). They are also more likely to be inattentive to fundamental information, at least as long as they remain inexperienced, as the process of acquiring, interpreting, and utilizing such information in real-world markets is complex. The results of all three studies in this thesis are more likely to apply to such traders.

Second, financial markets are moving toward a nearly frictionless paradigm. Online trading today is cheap, as it is commission-free; quick, as the onboarding of new clients onto trading platforms is fast and the platforms facilitate swift order input and execution; and easy, as trading platforms are designed to ensure that users navigate through as few hurdles as possible to input new orders. In fact, financial news media outlets such as Bloomberg and Wall Street Journal have criticized Robinhood for making trading “too easy” by providing a highly streamlined trading platform that contains design features that encourage investors to trade more frequently and keep returning to the platform (Wursthorn and Choi, 2020; Egkolfopoulou, Massa, and Melin, 2021). In this respect, the findings we report in Chapter 3 are quite relevant. We find that the move to frictionless markets might have the unintended consequence of amplifying trading losses for the

[3] Source: Robinhood.

most underperforming investors. According to our results, for these investors, some trading frictions, especially those that can direct their attention to the information they neglect, can be quite beneficial. Our results are relevant for brokers designing frictionless trading platforms and investors considering trading on such platforms.

Finally, both stock and cryptocurrency markets have recently become playing fields for speculative bubble games, such as the pump-and-dump games we study. In one prominent example, some users on a popular internet forum called Reddit urged other users on the platform to join them in inflating the stock price of GameStop, a video game retailer, despite the lack of any material change in the stock's fundamentals. The other users complied and helped push the stock price from $17.25 at the start of 2021 to a maximum of $347.51 on January 27, 2021, and then back down to $53.5 on February 4, 2021, a meteoric gain of more than 1,900%, followed by a precipitous drop of approximately 85% within a month. Even though there was no obvious manipulation angle in this episode, its speculative nature resembles the cryptocurrency pump-and-dump games studied in Chapter 2. Given this resemblance, the findings in this thesis can be extended to such speculative games as well. Biases such as overconfidence and gambling preferences are probable factors that attract investors to such speculative games in markets.[4] According to our findings on cryptocurrency pump-and-dump games, such games are likely to persist unless they are forcibly stopped by external entities, such as regulators or exchanges.

1.5 Thesis outline

This thesis comprises three studies on the following topics:

i. Pump-and-dump manipulation games in cryptocurrency markets (Chapter 2)
ii. Trading frictions and investor behavior and performance (Chapter 3)
iii. Attention to information releases and price reactions to announcements (Chapter 4)

Chapter 5 summarizes the findings of this thesis and suggests avenues for future research.

[4] Preliminary academic research on the subject reveals that, as conjectured, the GameStop episode was dominated by retail investors who had a history of investing in speculative assets, such as lottery-like assets, that are usually favored by investors with gambling preferences (Hasso et al., 2021).

Chapter 2: Pump-and-dump manipulation games in cryptocurrency markets

2.1. Introduction

Traditional pump-and-dump schemes typically fall into one of two categories: *information-based* manipulation or *trade-based* manipulation. Information-based manipulation involves spreading false information about the value of the security in the hope that traders will believe the false information (Vila, 1989; Van Bommel, 2003). Information-based manipulation theories require uncertainty about the fair value of a security and information asymmetry as the underpinnings of the manipulation. If there is little or no uncertainty about the fair value or no information asymmetry, then uninformed traders will not act on the rumors or false information circulated by the manipulators. Trade-based manipulation involves manipulating the price of a security by buying and then selling, or vice versa. Allen and Gorton (1992) and Jarrow (1992) respectively show that both asymmetry in liquidity-motivated trading and price momentum can drive trade-based manipulation.

We examine a new form of pump-and-dump manipulation in cryptocurrency markets. We begin by showing that this form of manipulation is widespread and accounts for considerable cryptocurrency trading. Combining hand-collected pump-and-dump data with data obtained from a pump-and-dump manipulation aggregation website, we identify as many as 355 cases of pump-and-dump manipulation within a period of seven months on two cryptocurrency exchanges. Up to 23 million individuals are involved in these manipulations. We estimate that the 355 pumps in our sample are associated with approximately \$350 million of trading on the manipulation days and that manipulators extract profits of approximately \$6 million from other participants. In all, 197 distinct cryptocurrencies or "coins" are manipulated, which implies that approximately 15% of all coins in our sample of exchanges are targeted by manipulators at least once in the seven-month period. There are, on average, two pumps per day. This rate of manipulation is considerably higher than pump-and-dump manipulation in stock markets in recent decades.[5]

[5] Market manipulation in modern equity markets is significantly lower. For example, Aggarwal and Wu (2006) only identify 142 pump-and-dump cases on US stock exchanges in an 11-year period, which implies that pump-and-dump rates in cryptocurrency markets are in the order of 40 times higher than those in equity markets.

Interestingly, while pump-and-dump manipulation of cryptocurrencies is similar to that of stocks in some regards, it is completely different in others. The most important difference is that, unlike in typical stock market manipulations, in cryptocurrency pump-and-dumps, manipulators typically make no pretense of having private information or claiming that a coin is undervalued.[6] Instead, pump group administrators (manipulators) publicly declare that they are pumping a given coin (releasing a "pump signal") and call on others to join.[7] Others then rush to buy the coin, presumably hoping to sell before the collapse of the pumped coin's price.[8] Economically, this means that manipulation schemes orchestrated by cryptocurrency pump groups do not exploit the classic mechanisms of information asymmetry and uncertainty about the fundamental value of an asset to "fool" market participants into buying a security.[9] Nor do they exploit asymmetry in price impact, which underpins trade-based market manipulation.[10]

This novel feature of cryptocurrency pump-and-dumps raises two interesting questions. First, why do individuals participate in these pumps? Second, how do manipulators profit if they are neither fooling participants into believing misinformation about the coin nor exploiting asymmetry in price impacts? We present a simple theoretical framework to address both these questions. We show that rational individuals lacking a speed or skill advantage do not participate in pump-and-dump manipulations because, for participants other than the manipulators (who can buy ahead of the pump signal), pumps provide negative expected returns. The intuition is simple. Pumps are a zero-sum game involving redistribution of wealth between players. Given that manipulators use their advantage to extract profit and given the existence of trading costs, pumps become a negative-

---

[6] In conventional pump-and-dump schemes, manipulators try to convince investors to buy the stock by spreading positive news about the stock through e-mails, phone calls, and newsletters and by claiming that the stock has the potential for large gains. Prices rise as investors who are convinced by the manipulators' promotional campaign buy the stock. After the manipulators' promotional campaign is over, these investors sell their holdings, leading to a fall in the stock price (Leuz et al., 2017).

[7] While this transparency of the manipulators' intentions is a feature of the manipulations that we examine (those orchestrated by organized pump groups), it is possible that traditional forms of pump-and-dump manipulation that rely on information asymmetry or asymmetry in price impacts also exist in cryptocurrency markets outside of these organized pump groups.

[8] In our sample, 92% pumps generate a minimum 1% return and significant trading volume within minutes of the pump signal.

[9] Although manipulators trigger the pump-and-dump episodes through an information release, the information that is released is not false information about the value of the security as is typical of information-based manipulation. Rather, the information reveals the intended manipulation. Manipulators do not exploit information asymmetry or uncertainty about the fair value, in contrast to typical cases of information-based manipulation. The data support these arguments: we find that the probability of a pump-and-dump manipulation is not higher when there is more uncertainty about a coin's value as proxied by the lagged volatility of the coin (Table 2.A1 in Appendix 2.E).

[10] See Section 2.B.1 in Appendix 2.B.

sum game for participants other than the manipulator. While the fastest or most skilled participants can profit at the expense of slower or less skilled participants, non-manipulators, in aggregate, lose money, presenting a puzzle of how these pumps can sustain participation.

We use the theoretical framework to illustrate two types of individuals who willingly take part in pumps, despite the unattractive returns. The first type is overconfident individuals, who overestimate their ability to sell at a price close to the peak. From their perspective, pumps (incorrectly) are a profitable game *ex-ante*. The second type is individuals who use pumps as a form of gambling, attracted by the possibility of large gains and the right-skewed payoff distribution that pumps can generate under certain conditions.

We find strong empirical support for both these explanations. We estimate regressions of pump participation on measures of overconfidence and gambling. The overconfidence proxy that we use is based on the tendency for past success to increase overconfidence through the self-attribution bias. The gambling proxy measures the level of gambling in cryptocurrency-based gambling services (such as SatoshiDICE), normalized by the level of cryptocurrency activity. The results show that overconfidence and gambling are statistically significant and economically meaningful determinants of the amount of participation in cryptocurrency pumps. The magnitudes of both mechanisms are similar. Therefore, we conclude that cryptocurrency pump-and-dumps are, to a large extent, a type of trading game that attract overconfident traders and people looking for gambles.[11]

We study various other properties of cryptocurrency pump-and-dumps both empirically and theoretically. We show that while pump-and-dump episodes tend to be rather short-lived, they generate extreme price distortions and abnormal volumes. The average cryptocurrency pump-and-dump manipulation is associated with a price rise of 65% within minutes. These average returns are around four standard deviations of the daily cryptocurrency returns. Therefore, even after considering the high volatility of cryptocurrencies, pump-and-dump episodes generate extreme price distortions. On average, it takes about eight minutes for a pumped coin to reach its peak price, after which the dumping phase commences and the price collapses. The trading volume on

[11] Despite functioning as a type of trading game, cryptocurrency pump-and-dumps are nevertheless a form of market manipulation as they involve a deliberate intention on the part of the manipulators to influence the price of a traded security. Therefore, these schemes not only expose unwitting and naïve investors to exploitation by manipulators but are also likely to involve similar costs as other forms of manipulation, e.g., a loss of confidence in the integrity of markets.

manipulation days is approximately 13.5 times the usual daily volume. We also find spillover effects into other markets that trade the manipulated coins but are not directly targeted by the manipulators; both volume and volatility increase significantly on these other markets during the pumps.

We also examine the characteristics of coins that are most likely to be targeted by manipulators, and we test how manipulation affects these coins. We find that manipulators target relatively illiquid coins. This observation is consistent with our framework, which predicts that individuals will be more attracted to pumps of illiquid coins. However, the most illiquid coins have a lower likelihood of manipulation, which is probably because a minimum level of liquidity is required to make the pump feasible and worthwhile on the part of manipulators, who need to build positions in the coin before sending a pump signal.

Finally, we find that although pumps create extreme price distortions *during* the pump, prices revert to their pre-pump levels within a day or two (often within an hour) following the manipulation. Therefore, pumps do not appear to have a permanent impact on the value of a targeted coin. The price distortions created by pumps are larger in the case of less liquid coins and when more individuals participate in the pump. Over time, pumps tend to speed up and take less time to reach the peak price from the time the pump signal is sent.

This chapter contributes to the literature on behavioral biases in financial markets by showing how biases such as overconfidence and gambling preferences can drive agents to participate in speculative bubble games in markets. If we strip out the manipulation angle, at their core, cryptocurrency pump-and-dumps are a type of bubble game. In such games, agents disregard a security's fundamentals and purchase it in the belief that they can sell to others purchasing after them at higher prices.[12] Such games have gained popularity in markets recently, with the

[12] Interestingly, some instruments in cryptocurrency markets are prone to speculative behavior by design. Consider the example of non-fungible tokens (NFTs) and smart contracts that rely on "bonding curves" in their primary market. The primary market prices of such instruments increase with an increase in their supply. Such instruments can witness price speculation since they naturally advantage early investors as subsequent buyers are compelled to purchase the instrument at higher prices in the primary market. Take, for example, Unisocks (SOCKS), a token that uses bonding curves. The price of SOCKS rose from $69,000 to $160,000 between February 27, 2021, and March 4, 2021, only to drop to $76,000 a day later. As an instrument, the only fundamental right that the SOCKS token grants a holder is that it can be redeemed for an NFT and a pair of socks. Another innovation in the cryptocurrency ecosystem that can favor manipulators is automated market making on decentralized exchanges. Essentially, these systems can facilitate "rug pulls," wherein manipulators create, publicize, and drive up prices of fake tokens, while simultaneously getting market participants to fill the liquidity pools held by automated market makers (AMMs). Manipulators then rely on the liquidity provided by the AMMs to liquidate their positions, thus obviating the need to rely on the liquidity provided by other pump participants or traditional market makers. Unlike traditional market makers, AMMs cannot suspend

GameStop episode being an exemplar. In this episode, users on an online forum called Reddit urged others to purchase the stock of GameStop, a video game retail company, even though there was no meaningful change in the stock's fundamentals. Other Reddit users complied, and their collective trading increased GameStop's stock price by approximately 1,900% within a month. Even celebrities are alleged to have either knowingly or inadvertently triggered speculative games in markets. For example, in his "Money Stuff" newsletter, Matt Levine discusses how, for some market participants, Elon Musk's tweets act as co-ordination mechanisms for pump-and-dump games (essentially like Telegram pump signals), with specific words in a tweet revealing the coin or stock to pump (Levine, 2021). An example of a celebrity knowingly fueling price speculation is that of social media influencer David Portnoy. In a tweet made on May 17, 2021, Portnoy announced his support for a token called SafeMoon, despite acknowledging that "it could be a Ponzi scheme." He added that "if it is a Ponzi, get in on the ground floor." Aside from these new, related examples, even cryptocurrency pump-and-dump games have regained popularity in 2021 after a brief post-2018 hiatus (see Section 2.3.2).

This chapter also contributes to the market manipulation literature by characterizing a new form of manipulation that differs from typical cases of pump-and-dump manipulation in stock markets.[13] Aggarwal and Wu (2006) provide a thorough analysis of stock market pump-and-dumps. The cryptocurrency pump-and-dumps that we examine exploit mechanisms that are quite different from those in typical stock market manipulations. Moreover, an analysis of cryptocurrency pump-and-dumps also leads to certain data advantages. For example, we can observe the exact start times for the pumps (when the pump group releases the "pump signal" to the participants). Our sample is also not contaminated by a prosecution bias, as we obtain the data on manipulations irrespective of whether these manipulations are prosecuted.[14] The cryptocurrency markets that we analyze are electronic limit-order book markets like most of the world's equity and derivatives exchanges. Thus, our findings on market manipulation can also prove useful for equity and derivatives markets.

---

liquidity provision or adjust their quotes in response to the threat of manipulation unless such a condition is pre-specified in the code.

[13] For reviews of the market manipulation literature, see Fox, Glosten, and Rautherberg (2018) and Putnins (2019).

[14] In contrast, the stock market manipulation prosecution cases that an empiricist can work with are a non-random "tip of the iceberg." For example, Comerton-Forde and Putniņš (2014) show that the prosecuted cases of manipulation reflect a tiny and non-random fraction (0.3%) of all manipulation.

Four contemporaneous papers analyze pump-and-dump manipulation in cryptocurrency markets.[15] Li, Shin, and Wang (2021) characterize the impact of pump-and-dump schemes on cryptocurrency markets. Xu and Livshits (2019) and Kamps and Kleinberg (2018) develop approaches to predict cryptocurrency pumps *ex-ante* and *ex-post*, respectively. Finally, Hamrick et al. (2020) examine the factors that affect the price jumps in cryptocurrency pumps. While these four studies and this chapter have some findings in common, our study has a different focus. For example, like this chapter, Hamrick et al. (2020) and Li et al. (2021) find that coins with lower market capitalization have larger price jumps during pumps. Similar to the findings reported in this chapter, Li et al. (2021) reveal that these schemes are widespread, generate large but short-lived price distortions, and result in substantial wealth transfers between market participants. A novel feature of their analysis lies in testing the causal effects of these schemes on liquidity. Surprisingly, they find that these schemes decrease the liquidity of cryptocurrency markets.

In contrast, the focus of this chapter is on modeling and testing the mechanisms that enable these schemes to exist in the first place and analyzing how they sustain participation. Theoretically, we show that two behavioral factors—overconfidence and gambling preferences—can explain pump participation. We find empirical support for both of these behavioral factors. We also provide further characterization of these schemes beyond what is presented in the other studies, showing their impact on the manipulated market as well as spillovers to other markets and identifying the characteristics of coins most susceptible to such manipulation.

The rest of this chapter is structured as follows. Section 2.2 provides an illustration of a cryptocurrency pump-and-dump manipulation. Section 2.3 introduces our data and sample. Section 2.4 investigates participation in cryptocurrency pumps theoretically and empirically. Section 2.5 analyzes the characteristics, determinants, and effects of pumps. Section 2.6 discusses the welfare effects of cryptocurrency pumps. Section 2.7 concludes this chapter.

[15] This chapter is also related to two recent studies of other forms of cryptocurrency market manipulation. In both cases, the (alleged) perpetrators of market manipulation are cryptocurrency exchanges or parties associated with them. Griffin and Shams (2018) show that parties related to the cryptocurrency exchange Bitfinex used Tether (a USD-pegged cryptocurrency) to inflate the price of Bitcoin. Gandal et al. (2018) examine Bitcoin trading in 2013 and argue that fraudulent transactions conducted by the Mt. Gox exchange caused the Bitcoin price to rise by more than 500% in two months.

2.2. Illustration of a cryptocurrency pump-and-dump manipulation

To explain the mechanics of a pump, we illustrate a pump conducted by the Big Pump Signal (BPS) group, one of the largest pump groups by membership (around 63,000 members on Telegram at the time of writing).[16] This pump was on the Binance exchange and involved the cryptocurrency ChatCoin (CHAT). Figure 2.1 displays communication from the BPS administrators to the group members about the CHAT pump. First, the administrators announce the exchange, date, and time at which the pump will occur, but not the actual coin that will be pumped (top left-hand-side message in Figure 2.1). This pre-announcement of the pump prepares the participants, allowing them to transfer funds to the nominated exchange and be online and wait for the pump signal at the pre-specified time.

Then comes the "pump signal," which is just the name of the coin being pumped (bottom left-hand-side message in Figure 2.1). In this illustration, the coin being pumped is "CHAT," as indicated by the red dashed line around the text "CHAT" (the pump signal is displayed in such a format to prevent machine readability). BPS sent the pump signal for the CHAT pump at 20:00:23 GMT on June 10, 2018.

Figure 2.2 Panel A shows the price of CHAT in 15-second intervals starting 15 minutes prior to the release of the pump signal and ending 90 minutes after the pump signal. Figure 2.2 Panel B shows the trading volume of CHAT in five-minute intervals during the same period. There is very little trading activity in the 15 minutes leading up to the pump signal and a sudden jump in the price and trading volume immediately after the pump signal at 20:00. Although the price fluctuates considerably during and after the pump, it peaks at around 17 seconds after the pump signal is sent (Panel C zooms in to the five minutes around the pump signal). The peak price is approximately 55% above the pre-pump price level. Following the peak, the price and volume of CHAT remain elevated until about one hour after the pump signal, at which time the price falls back down to around the pre-pump price level.

In addition to the extreme price movements, the pump is associated with a significant spike in trading volume. In the half-hour before the pump signal is sent, the trading volume in CHAT is $17,313.47 (or 2.55 BTC), which is likely to be pre-pump position-building by the manipulators (the group administrators).[17] In the half-hour after the pump signal is released, the trading volume

[16] Appendix 2.A provides more details about pump groups and their operations.
[17] BTC denotes the unit of measurement for Bitcoin. Most cryptocurrencies are traded in BTC pairs.

in CHAT jumps to $2.69 million, which is 3.6 times the median daily dollar volume for CHAT between December 2017 and July 2018.

2.3. Data and summary statistics

2.3.1. Data description

We identify pump-and-dump manipulation cases in cryptocurrency exchanges by using chat history data from Telegram pump-and-dump channels. We hand-collect pump data for the cryptocurrency exchange Binance and add to this set the pump data for the cryptocurrency exchange Yobit from the database compiled by PumpAnalysis (PA).[18] In our sample, we only include cases in which the Telegram pump-and-dump administrator pre-specifies a date, time, and exchange for a pump and releases a "pump signal" containing the name of the pumped coin. Therefore, our sample only includes events in which the pump group administrator makes it explicit to its members that the intention is to pump the given coin. For each pump, we record the coin being pumped, the exchange, and the time at which the pump signal is sent. We also extract ancillary information, such as the number of pumps that a group has conducted prior to the pump, the number of groups participating in the pump, and the total number of members in the pump group(s) participating in the pump.

We couple the pump information with data on all trades on the Binance and Yobit exchanges (sourced from the official Binance API and the data provider Kaiko, respectively). We restrict our focus to Binance and Yobit because of the limited availability of reliable trade data. Our sample consists of 355 pumps (64 on Binance and 291 on Yobit). The earliest pump in our sample occurred on December 29, 2017, while the last pump was on June 22, 2018. Based on the number of pump group members, up to 23.3 million participants were involved in the pump-and-dumps during our sample period.

[18] Although the website hosting the dataset (PumpAnalysis.com) is no longer active, we downloaded a copy of the data before it went offline. We conducted a data audit to verify the quality of the dataset and did not find any material inaccuracies. In this audit, we verified the pump-related information in the database against the actual information found in the chat history of the pump-and-dump group for a random sample of pumps. Our copy of the PA dataset is available upon request.

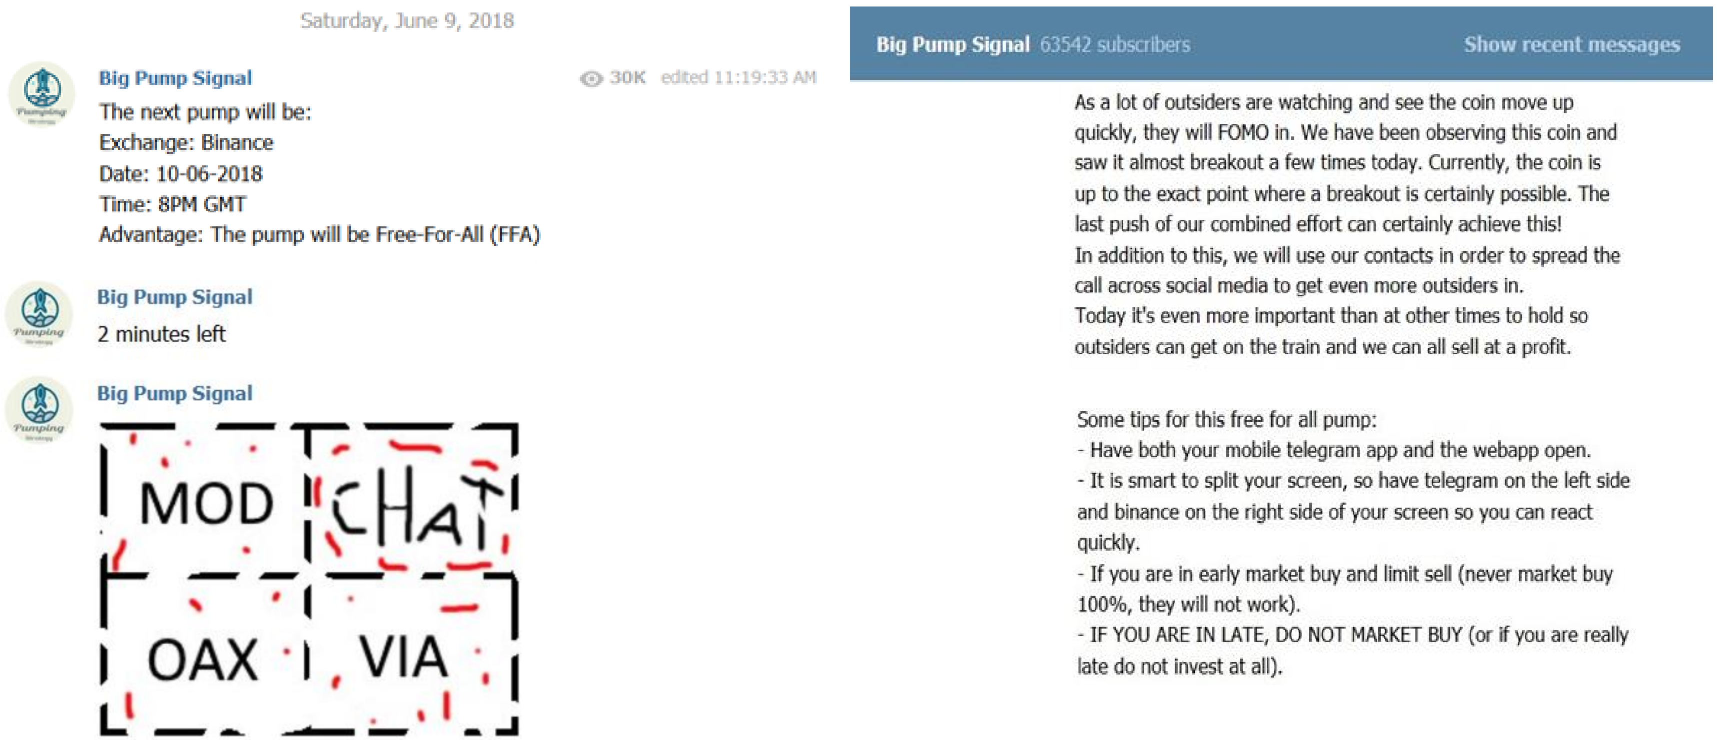

Figure 2.1. Sample messages sent on the 'Big Pump Signal' Telegram group.
This figure shows messages sent by the administrators of the Telegram group 'Big Pump Signal' to its members. In these messages, the administrators announce the date, time, and exchange for a pump in advance of the actual pump (top left-hand-side message). Then, at the commencement of the pump, the group's administrators send the pump signal by releasing the name of the coin being pumped (bottom left-hand-side message). In this illustration, the coin being pumped is 'CHAT,' as indicated by the red dashed line around the text 'CHAT.' The administrators also provide tips to pump participants on how to profit from pumps (right-hand-side message). The group's administrators advise members to enter the pump as soon as possible and dump their holdings on outsiders who are likely to buy the coin based on rapid upward price movements.

Panel A: Price movement for ChatCoin before, during, and after the 'Big Pump Signal' pump

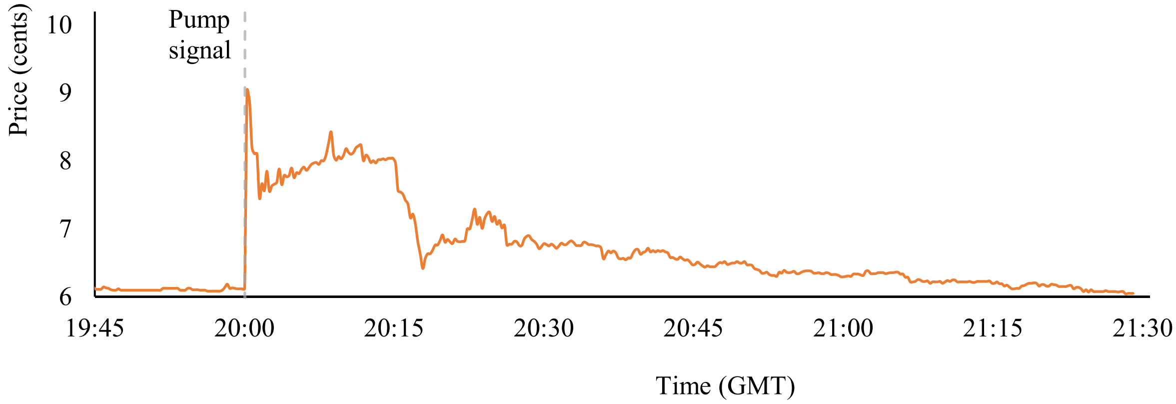

Panel B: Trading volume for ChatCoin before, during, and after the 'Big Pump Signal' pump

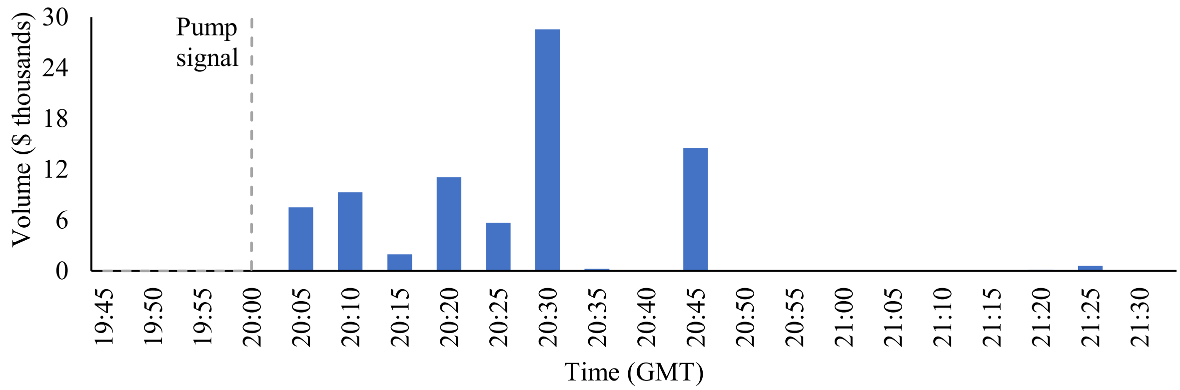

Panel C: Magnified price movement graph for ChatCoin during the 'Big Pump Signal' pump

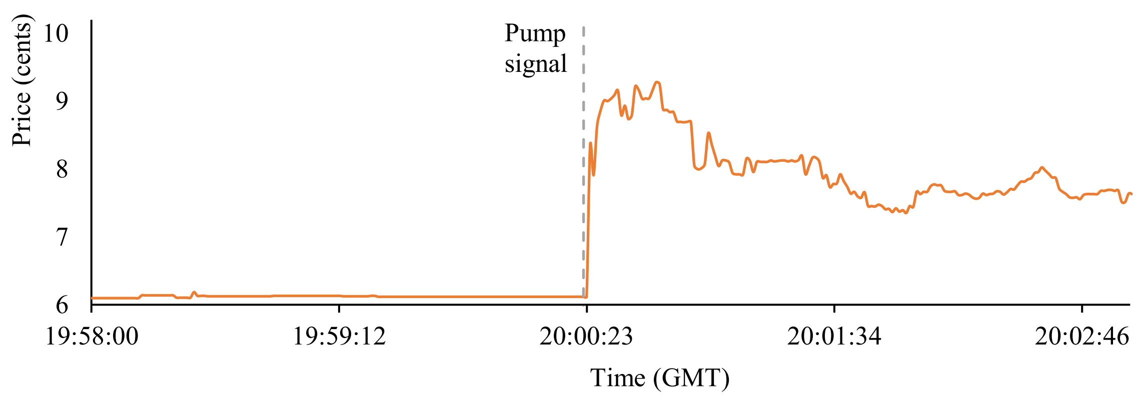

Figure 2.2. Price and volume for ChatCoin during the 'Big Pump Signal' pump.
The pump signal for the ChatCoin pump was sent by the administrators of the 'Big Pump Signal' group at 20:00:23 GMT on June 10, 2018. Panel A shows the prices (in cents) for ChatCoin in 15-second intervals between 19:45 GMT and 21:30 GMT on June 10, 2018, i.e., from 15 minutes prior to the pump signal to 90 minutes after the signal. Panel B shows the trading volume ($ thousands) for ChatCoin in five-minute intervals. Panel C shows the prices (in cents) for ChatCoin in a narrower window, between 19:58 GMT and 20:03 GMT on June 10, 2018, i.e., from two minutes before the pump signal to three minutes after the signal.

We obtain daily market capitalization and price data for all cryptocurrencies from coinmarketcap.com, a cryptocurrency data aggregation website, and daily exchange-level price and volume data for all cryptocurrencies from cryptocompare.com, another cryptocurrency data aggregation website. We also obtain Google Search Volume Index data from the Google Trends website. Lastly, we obtain daily Bitcoin gambling volumes from WalletExplorer.com, a Bitcoin blockchain explorer website.

2.3.2. Summary statistics

Table 2.1 reports the descriptive statistics for the trading activity on the two exchanges (Yobit and Binance) and the prevalence of pump-and-dump manipulation. In our seven-month sample, 1,307 cryptocurrencies ("coins") are traded on the two exchanges, with a combined volume of around $20 billion. The 355 instances of pump-and-dump manipulation that we identify and for which we have all necessary data occur in 197 distinct coins. Therefore, approximately 15% of all coins (197/1,307) experience at least one pump-and-dump manipulation during the seven-month period, with an average of 1.8 pumps per coin and 2.67 pumps per pump-day.[19] These manipulation rates (being a lower bound, as they include only instances on which we have the necessary data) suggest that cryptocurrency pump-and-dump manipulation is widespread and frequent.

The volumes traded during pump-and-dump episodes are economically meaningful, with around $350 million traded during the 355 pumps in our sample. We estimate that manipulators buy approximately $24.38 million of coins in the two hours leading up to the pumps, resulting in a conservatively estimated aggregate profit to manipulators of approximately $6.04 million.[20] This figure reflects the estimated wealth transfer from the pump participants (pump group followers) to the manipulators (pump group administrators) during our sample period. As a return, the manipulators earn around 24.77% within minutes or hours.

[19] A "pump-day" is a day on which there is at least one instance of pump-and-dump manipulation.
[20] Manipulator profits are estimated from the difference in the volume-weighted average price during the two hours preceding the pump signal and the volume-weighted average price during the pump (from start to peak). Factoring in transaction costs slightly reduces these profit estimates. For example, assuming that manipulators pay the highest trading fees on both the exchanges (0.1% of transaction value on Binance and 0.2% on Yobit; the actual fees may be lower, depending on factors such as trading volume), the profit estimate decreases only marginally to $5.98 million. If we further assume that manipulators pay an effective spread of 1% on their trades (which is 18 times larger than the average effective spread for Bitcoin trades on Gdax, see Dyhrberg, Foley, and Svec, 2018), the profit estimate decreases to $5.44 million.

Table 2.1
Aggregate trading and manipulation on cryptocurrency exchanges

This table reports descriptive statistics for the trading activity and prevalence of pump-and-dump manipulation on two cryptocurrency exchanges (Binance and Yobit) between December 2017 and June 2018. Panel A statistics are calculated for all coins listed on both exchanges during the sample period. Panel B statistics are calculated only for pumped coins. A "pump-day" refers to a day on which there is at least one manipulation. Pre-pump volume is the trading volume in a manipulated coin in the two hours preceding the release of the pump signal. Manipulators' profit is calculated as the difference between the volume-weighted average price during the pump (from start to peak) and the volume-weighted average price in the two hours preceding the release of the pump signal, multiplied by the pre-pump volume.

Characteristic	Statistic
Panel A: Market activity	
Total coins	1,307
Total trading volume ($ million)	19,788.12
Panel B: Manipulation activity	
Total number of pumps	355
Number of pumped coins	197
Average pumps per pumped coin	1.80
Number of pump-days	133
Average pumps per pump-day	2.67
Total pump-day volume ($ million)	350.77
Total pre-pump volume ($ million)	24.38
Manipulators' total profit ($ million)	6.04
Manipulators' profit (% of pre-pump volume)	24.77%

We estimate the prevalence of cryptocurrency pump-and-dump schemes beyond the December 2017 to June 2018 window, which constitutes our main sample. For this purpose, we plot the Google search activity for cryptocurrency pump-and-dump-related search terms and intraday price reversals from June 2017 to July 2021 (see Figure 2.3).[21] Google searches are an indicator of public interest in cryptocurrency pumps, whereas reversals are an indicator of pump-and-dump activity and intensity, as pumps can cause sharp intraday price spikes and reversals. Although both pump-and-dump interest and activity decrease after the end of our sample period in June 2018, there is an increase in both search activity and reversals in 2021, with the search activity reaching close to its previous 2018 peak. The increase is markedly lower for the reversals than for the search activity. However, this does not necessarily imply a lower level of pump-and-dump activity. In addition to pump-and-dump activity, reversals are also a proxy for the peak returns of pumps, as pumps with higher peak returns witness higher reversals. The 2021 period could have witnessed similar levels

[21] Intraday price reversals are calculated for all coins listed on Yobit and Binance.

of pump-and-dump activity as the 2018 period, but with lower peaks. Higher liquidity in cryptocurrency markets can explain these lower peaks according to our theoretical framework (see Section 2.5.2). In addition, the pump-and-dump activity might have moved from Yobit and Binance to other cryptocurrency exchanges. In this case, too, we would observe relatively lower intraday reversals for coins listed on these exchanges.

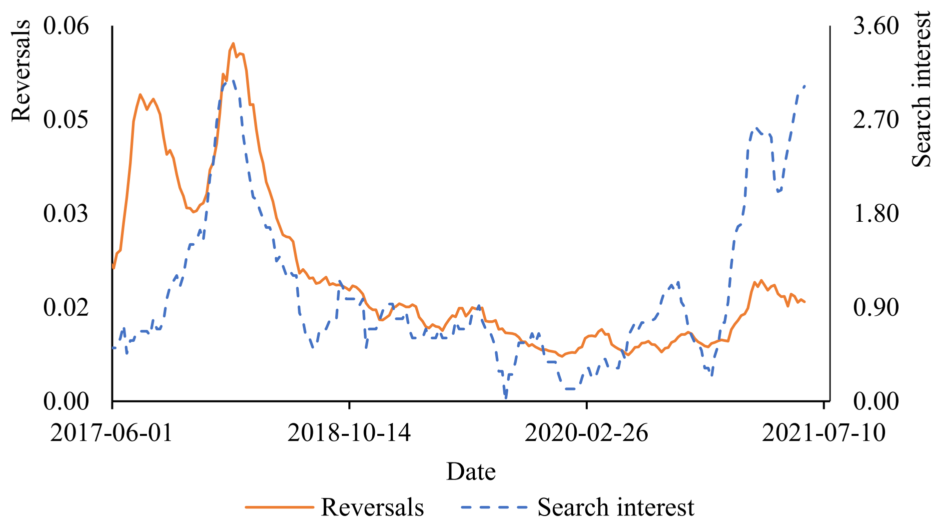

Figure 2.3. Pump-and-dump interest and potential pump activity through time.
This figure plots search interest in cryptocurrency pump-and-dumps and a potential pump activity indicator (intraday reversals) between June 2017 and July 2021. We calculate intraday reversals for all coins listed on Binance and Yobit. Reversals are calculated as the minimum of the percentage return from the day's open to the day's high and the return from the day's high to the close. Search interest is the Google search activity for search terms related to cryptocurrency pump-and-dumps.

2.4. Who participates in pumps?

In this section and the following section, we study cryptocurrency pump-and-dump manipulations theoretically by using a simple framework and empirically by using the data described in Section 2.3.1. We begin by formally expounding on the puzzle of participation in cryptocurrency pumps and then proceed to show how behavioral mechanisms such as overconfidence and gambling preferences help resolve this puzzle.

2.4.1. Framework foundation and notation

We model cryptocurrency pump-and-dumps as a four-period, simultaneous-move trading game. There are three types of agents: manipulators (pump group administrators), traders who can choose to participate in the pump following the release of a pump signal, and a simple market maker that facilitates trades.

The game starts in Period 0 when a manipulator or a group of manipulators decides to pump a particular coin. The price of the coin at the time is P_0.

In Period 1, the manipulators take a long position of M units ($M > 1$) in the coin and send a message to their pump group members stating there will be a pump in Period 2 (without releasing the name of the coin, as per the previous example of a typical pump). The M units can be thought of as M manipulators, each buying one unit of the coin or one manipulator buying M units of the coin. We assume that all trades (for manipulators and other participants) have linear price impacts. The pricing function is $P_t = P_{t-1} + \beta x_t$, where P_t is the price in period t, β is a price impact parameter between zero and one, and x_t is the net volume of buys (buys minus sells) received by the market at time t. Section 2.B.1 in Appendix 2.B details the motivation behind selecting this pricing function.

In Period 2, the manipulators send the pump signal to the $N > 1$ members of their pump group, notifying them of the coin being pumped. These N members simultaneously decide whether to participate in the pump (buy one unit of the coin) or not to participate (no trade). We restrict the trade sizes to one unit to keep the framework simple and to focus on the participation decision. Players who decide to participate in the pump race to submit their unit-volume market buy orders to the market. Matching engines in financial markets, including cryptocurrency exchanges, typically process incoming orders sequentially by placing them in a queue. Therefore, small random latencies in order submission determine the queue position or the sequence in which the participants' orders are executed by the market. These random latencies include the time taken to receive and interpret the pump signal, to make a decision, to enter the order, and for the order to be transmitted to the market. The individual participants buy at prices $\{(P_1 + 1\beta), (P_1 + 2\beta), \dots\}$ depending on their random latency, which determines their queue position.

Conditional on participation, in Period 3, players exit the pump (along with the manipulators) by simultaneously submitting unit-volume market sell orders. These orders are executed in the same way as the entry orders: random latencies determine the queue positions and execution prices. For simplicity, we assume the exit queue position is independent of the entry queue position. As per our framework, prices first rise slightly in Period 1 after the manipulators establish their initial positions and rise even higher in Period 2 after other participants join the pump. Finally, prices fall back to their base level (P_0) after all participants (manipulators and non-manipulators) exit the pump in Period 3. Thus, prices are a function of M, N, and β. Section 2.B.1 in Appendix 2.B explains the price dynamics of pumps in detail.

2.4.2. Rational agents

What types of individuals choose to participate in cryptocurrency pump-and-dumps? We start by showing that it is unclear why individuals choose to participate in pumps, given that the expected return to pump participants in aggregate is negative. We then show that a systematic skill or speed advantage can explain the participation of some individuals, shifting the puzzle to less skilled or slower participants.

Let us start with the simple case of homogenous participants with correct beliefs and consider the prices at which a participant expects to buy and sell the coin. Recall that the manipulators buy M units before sending the pump signal, driving the price to $P_1 = P_0 + M\beta$ just before the pump signal is released. If all N pump group members choose to participate, they buy the coin at prices $\{(P_0 + \beta(M+1)), (P_0 + \beta(M+2)), \ldots, (P_0 + \beta(M+N))\}$ depending on their random latency. Thus, their "entry prices" (P_{entry}) are uniformly distributed from $P_0 + \beta(M+1)$ to $P_0 + \beta(M+N)$. Similarly, the "exit prices" at which they sell the coin in Period 3 (P_{exit}) are uniformly distributed from P_0 (once all positions are liquidated, there are zero net cumulative buys) to $P_0 + \beta(M+N-1)$, which is the price received by the first seller following the peak.[22] Therefore, an individual i's expected profit is:

$$\mathbb{E}[\pi_i] = \mathbb{E}\left[P_{exit} - P_{entry}\right] = -\frac{\beta(M+2)}{2}. \tag{1}$$

Both β and M are strictly positive; therefore, the expected return, $\mathbb{E}[\pi_i]$, is negative. The expected loss has two components: the round-trip trade cost (β) and the expected loss to the manipulator, which is half the manipulators' initial price impact ($\beta M/2$).

In contrast, manipulators have an advantage over other participants, as they can buy the coin ahead of the pump signal, effectively gaining a more advantageous entry price than their followers. Therefore, pumps can have positive expected profits for manipulators as long as there is a sufficiently large number of participants (N') in the pump to cover their transaction costs:

[22] Technically, the entry and exit prices follow discrete uniform distributions. However, given that there are typically many participants in pumps (N in the hundreds or thousands), the continuous distribution is a reasonable approximation that allows us to obtain a tractable solution.

$$\mathbb{E}[\pi_m] = \frac{\beta M}{2}(N' - 2M), \tag{2}$$

which is positive if $N' > 2M$.

Therefore, while it is clear that manipulation can be profitable for the manipulators, other pump participants should expect to lose money in aggregate; collectively, they lose the sum of their transaction costs ($\beta N'$) and the gross profits of manipulators ($\beta MN'/2$). Intuitively, a rational individual with correct beliefs recognizes that a cryptocurrency pump-and-dump is a zero-sum game across all participants, including the manipulators, in the absence of trading costs. Positive trading costs and the wealth transfers to the manipulators make it a negative-sum game for participants other than the manipulators. Therefore, risk-averse or risk-neutral rational individuals with correct beliefs and no advantage over other participants choose not to participate in these pump-and-dump games.

Result 1*: Rational individuals with correct beliefs do not participate in cryptocurrency pump-and-dumps.*

However, what if the participants are not homogenous; instead, some are faster or more skilled than others? Some participants might be able to react faster than others systematically, or they may have lower latencies in getting orders filled at the exchange, or they simply might be better at "picking the peak" and sensing when the pump is about to enter the dump phase. In effect, they would be able to buy at a lower price soon after the release of the pump signal and sell at a higher price soon after the pump peaks.

To incorporate heterogeneity in speed or skill into the framework, we allow for fast and slow players by tilting the exit price distribution of fast (slow) players toward higher (lower) prices (tilting the entry prices as well would merely strengthen the effect). We introduce a speed parameter, S_i, which is the slope of the transformed exit price probability density function (pdf) for player i. When there is an equal proportion of fast and slow participants, for fast players, $S_i = s^* > 0$, whereas for slow players, $S_i = -s^* < 0$. The tilted pdfs are illustrated in Figure 2.4. The extent of the tilt (magnitude of s^*) determines the magnitude of the difference in speed or skill. In any given pump, players continue to face uncertainty around their exit prices, creating the possibility of profits or losses from any given pump. However, the tilted distributions of exit prices favor the fast players, allowing them to sell at higher prices on average, such that they earn higher payoffs from pumps on average. Section 2.B.2 in Appendix 2.B provides the exact specifications of the tilted pdfs and the expected profits for fast and slow players.

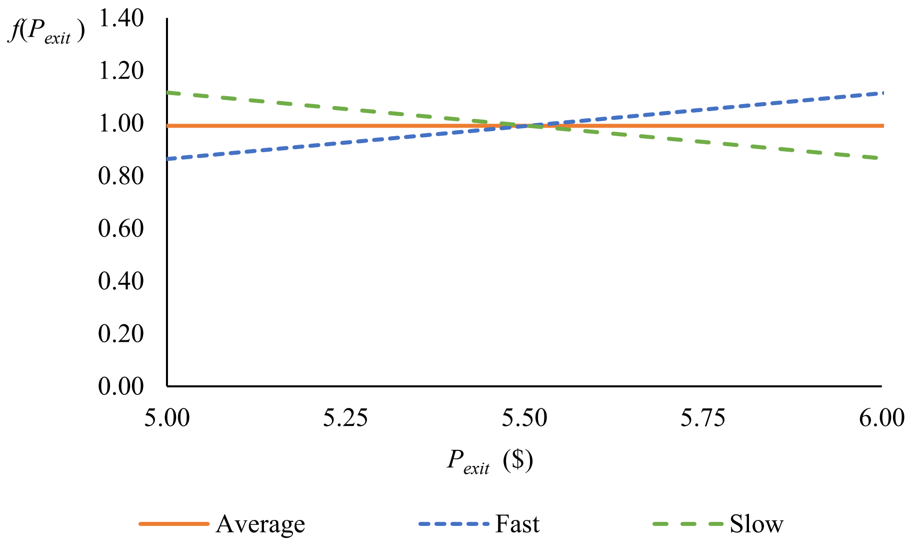

Figure 2.4. Exit price distributions for fast and slow individuals.
This graph plots the probability density functions (pdf) of exit prices (the prices at which a pump participant expects to close their long position) for fast and slow participants. In this illustration, we use the following parameter values: initial price $P_0 = \$5$, number of manipulators $M = 2$, number of participants $N' = 100$, and price impact $\beta = 0.01$. Assuming an equal proportion of fast and slow agents, the exit price distribution for the average participant is uniform. Fast (slow) agents have a high (low) probability of receiving high exit prices and a low (high) probability of receiving low exit prices. The speed parameter, S, is the slope of the transformed pdf. S takes the value $s^*(s^* > 0)$ for fast agents and the value $-s^*$ for slow agents. In this illustration, $s^* = 0.25$.

There are two interpretations of these results. First, it is possible that participants know their type (fast or slow). In this case, it is rational for sufficiently fast participants (with sufficiently low risk aversion) to participate in pumps. They might not profit from each pump, but under the conditions given above, they face positive expected returns. This shifts the puzzle of why people participate in pumps to the slow participants. In fact, participation becomes more puzzling because, in the presence of fast participants, the slower individuals face even worse expected losses from participating in pumps.

The second interpretation is that participants do not know their speed or skill relative to others. In this case, the puzzle of participation that we describe for the homogenous case remains. For example, with an equal proportion of fast and slow participants (implying symmetric tilts of the exit price distributions), a participant's expected exit price distribution if they do not know their type is uniform, as in the case of homogenous participants, resulting in negative expected returns.

Intuitively, if a participant does not know their speed or skill relative to others, they can expect a random share of the aggregate outcomes of participants, which are losses equal to the aggregate trading costs plus the manipulators' gross profits. Therefore, while the heterogeneity in speed potentially explains why some individuals willingly participate in pumps, it remains a puzzle as to why less skilled participants or participants who do not know their skill level participate in pumps.

2.4.3. Overconfident agents and gamblers

2.4.3.1. Theoretical predictions

To resolve the puzzle of participation, we propose two behavioral explanations. We show that overconfidence and/or a preference for skewness (gambling preferences) can explain the remainder of the participation puzzle.

First, we consider overconfidence, which can give individuals the *perception* of having an advantage over others. A large literature on psychology and behavioral finance shows that most people (including financial market participants) assess their own abilities to be higher than those of the average person (Barber and Odean, 2000; Gervais and Odean, 2001; Alicke and Govorun, 2005; Deaves et al., 2008). This is known as the better-than-average effect. In cryptocurrency pump-and-dumps, an overconfident individual who believes they are more skilled than the average player can expect to enter and exit pumps faster than the average participant and, therefore, obtain more favorable prices, much like the fast individuals described above.

To model this overconfidence effect, we return to the baseline of homogenous individuals and add a bias to the *perceived* distribution of exit prices for overconfident individuals. This bias can be interpreted as individuals believing they are better than average in "picking the peak" of the pump and exiting at a higher price than the average participant. Let the overconfidence parameter, ε_i, be the slope of the *perceived* pdf of exit prices. The overconfidence bias, ε_i, tilts the distribution to increase the perceived probability of exiting the pump at a high price. Individuals exceeding a minimum overconfidence threshold of ε^{min} expect to earn profits from cryptocurrency pump-and-dumps. The *perceived* pdf of exit prices and the minimum overconfidence threshold condition are presented in Section 2.B.3 of Appendix 2.B.

Thus, we arrive at the first potential explanation for why individuals participate in cryptocurrency pump-and-dumps.

Result 2*: Sufficiently overconfident individuals participate in cryptocurrency pump-and-dumps.*

Figure 2.5 illustrates how the minimum overconfidence threshold varies with the other parameters. For the illustration, we set values for two of the three parameters and plot ε^{min} against the third parameter. Panel A shows that ε^{min} is decreasing in β, implying that pumps of less liquid coins (higher price impact parameter, β) tend to have more participation from overconfident individuals due to the lower minimum overconfidence required to induce participation. This effect occurs because less liquid coins tend to have a higher dispersion of exit prices (higher pre-pump to peak return). Consequently, a smaller bias is required in the perceived probability of being able to exit the pump near the peak to make the pump attractive to an overconfident individual. This effect also explains why pumps tend to occur in relatively illiquid coins. Panel B shows that ε^{min} is increasing in M, implying that pumps with more manipulator participation tend to have less participation from overconfident players. This effect occurs because manipulators impose a cost (losses equivalent to manipulator gains) on other participants; therefore, a greater perceived ability to exit near the peak price (higher overconfidence) is required to make the pump attractive. Finally, Panel C shows that ε^{min} is decreasing in N', implying that pumps with more participants are more attractive to overconfident players. This effect arises because a larger number of participants leads to a higher peak price and more dispersion in exit prices; thus, only a small bias in the perceived likelihood of exiting the pump near the peak is required to make a pump attractive to an overconfident individual.

Panel A: Participation region for different levels of overconfidence (ε^{min}) and price impact (β)

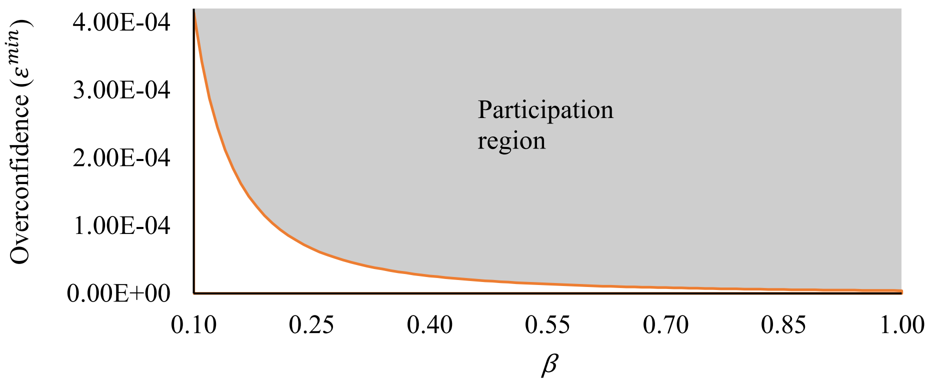

Panel B: Participation region for different levels of overconfidence (ε^{min}) and manipulators' pre-pump long position (M)

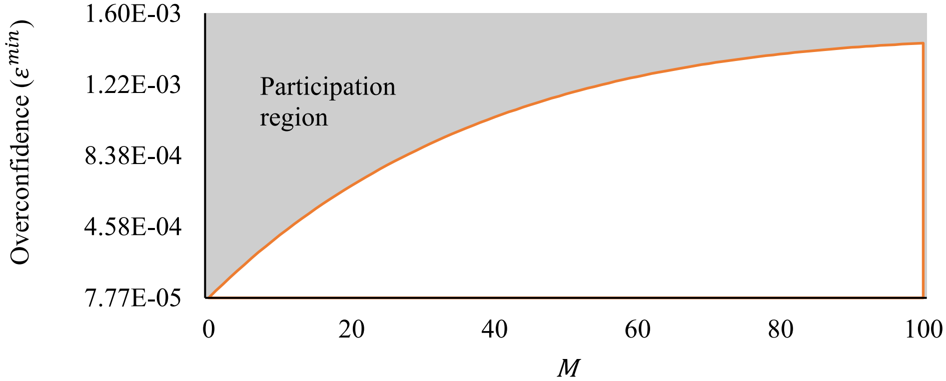

Panel C: Participation region for different levels of overconfidence (ε^{min}) and number of participants (N')

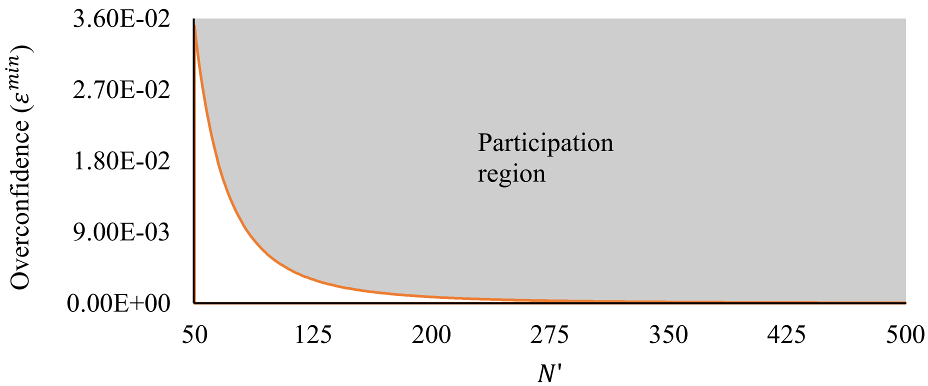

Figure 2.5. Pump-and-dump participation thresholds as a function of overconfidence and other parameters. This figure plots the minimum overconfidence level (ε^{min}) above which individuals choose to participate in pump-and-dump manipulation games. Panel A plots this overconfidence threshold for different values of the price impact parameter (β). Panel B plots the threshold for different values of the manipulators' pre-pump long position (M). Panel C plots the threshold for different values of the number of non-manipulators (N'). In all three plots, the shaded area is the region in which individuals participate in the pump. To plot the graphs, for illustrative purposes, we use the following parameter values: price impact parameter $\beta = 0.1$, number of manipulators $M = 10$, and number of non-manipulator participants $N' = 250$.

Next, we consider whether gambling preferences explain why individuals participate in pump-and-dumps. The previous literature shows that individuals prefer "lottery-like" assets with positively skewed payoffs (Barberis and Huang, 2008; Kumar, 2009). In our framework, a single pump does not have a positively skewed payoff distribution; the gains and losses are approximately symmetrical. However, even a single bet on red or black at the roulette table does not have a right-skewed payoff. To explain the attraction of non-skewed games, such as bets in roulette, Barberis (2012) shows that gamblers view these games not in isolation but as a series of bets that collectively constitute a game. When a gambler intends to play a game repeatedly and stop if their losses exceed a "walk away" threshold, a game that has symmetric payoffs as a one-off gamble becomes right-skewed for a series of bets. Other studies also suggest that gamblers favor repeated games over single games (Dickerson, 1984; Grinblatt and Keloharju, 2009).

Applying the mechanism proposed by Barberis (2012) to our setting, suppose a gambler starts with $\$a$ ($a > 0$) and considers whether to participate in a series of pumps until they either deplete their wealth to $\$b$ ($b < a$), incurring a loss of $a - b$, or accumulate a wealth of $\$c$ ($c > a$), making a gain of $c - a$. This strategy reduces to a binary gamble. Following Barberis (2012), we assume that gamblers exhibit preferences consistent with the Cumulative Prospect Theory (CPT) developed in Tversky and Kahneman (1992). Section 2.B.4 in Appendix 2.B provides details about the binary gamble, the gambler's maximization problem, and their preferences. In essence, the gambler determines the optimal values of b and c that lead to the highest expected payoff.

We solve the gambler's optimization problem numerically by determining the optimal b and c values for a given set of parameter values.[23] We find that the optimal value of b is \$0 and that of c is \$245. This implies that a gambler starting off with $a = \$6$ participates in pumps either until they go bankrupt or until they accumulate a wealth of \$245. This strategy yields a positive expected value for the gambler with CPT preferences and, therefore, provides a second possible reason for why individuals participate in cryptocurrency pump-and-dumps.

Result 3*: Individuals with Cumulative Prospect Theory preferences participate in cryptocurrency pump-and-dumps as a form of gambling.*

[23] We set $P_0 = \$5$, $\beta = 0.01$, $M = 10$, $N' = 1{,}000$, and $a = \$6$. Refer to Section B.4 in Appendix 2.B for details about the CPT parameters and their values.

In Figure 2.6, we illustrate how the attractiveness of cryptocurrency pump-and-dumps as a form of gambling varies with different parameters. We vary the parameters one at a time, plotting the gamblers' expected value of participating in a series of pumps. The figure shows that gamblers are more attracted to cryptocurrency pump-and-dumps (higher expected value) when the pumped coins are relatively illiquid (high β), there are relatively few manipulators (low M), and there are many participants (high N'). These relations are similar to those observed for overconfident individuals.

We document two additional participation-related results. First, based on the illustrations in Figures 2.5 and 2.6, we know that more pre-pump trading by the manipulators (M) makes pumps relatively less attractive to individuals of all types—rational agents expect larger expected losses to manipulators, higher levels of overconfidence are required to induce participation when M is higher, and the expected value of participating in a series of pumps decreases in M.

Result 4*: Pumps with more manipulator participation are less attractive to non-manipulators and, therefore, have less non-manipulator participation.*

Next, we model the long-run dynamics of pump-and-dumps by considering repeat iterations of the pump-and-dump game. Section 2.B.5 in Appendix 2.B details the repeated game dynamics of pumps. The primary empirical implication of this extension is that we expect to see increases in the volumes traded in pump-and-dumps through time when the participant inflow rate exceeds the outflow rate. This scenario can be expected when: (i) there is increasing interest in cryptocurrencies in general, (ii) overconfidence increases, for example, due to the self-attribution of success in past pumps, and (iii) there is a market-wide increase in the propensity to gamble, with cryptocurrency pump-and-dumps providing one outlet for gamblers.

Result 5*: Participation in pump-and-dumps through time increases when the general level of interest in cryptocurrencies increases, when the returns of past pumps are higher, and when there is an increase in market-wide gambling activity.*

Panel A: Gamblers' expected value from pumps vs. the price impact parameter (β)

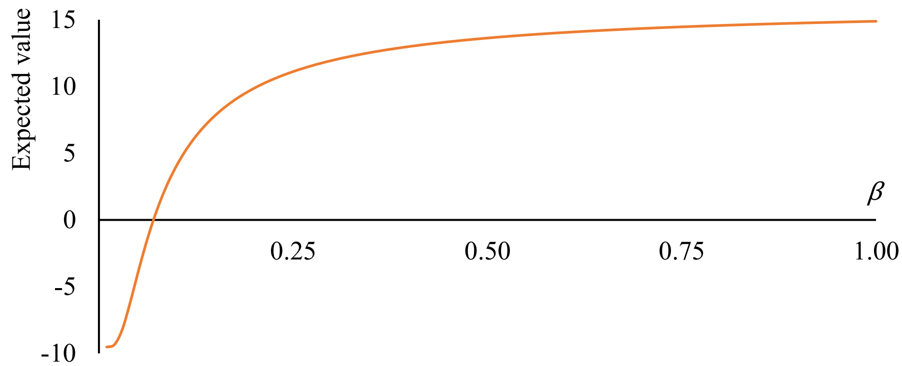

Panel B: Gamblers' expected value from pumps vs. the manipulators' pre-pump long position (M)

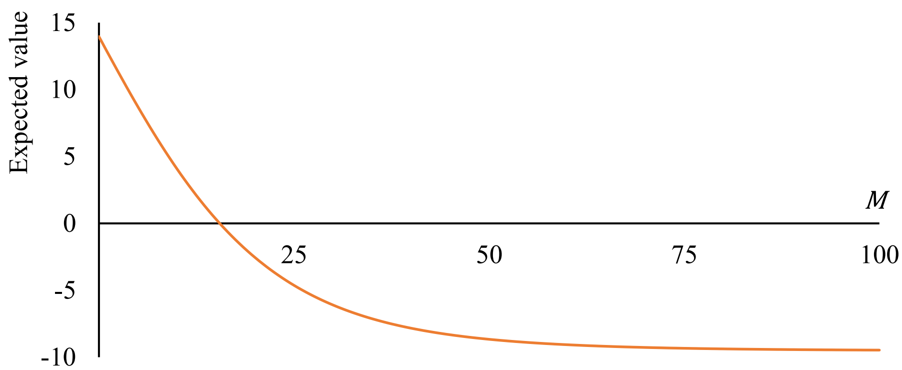

Panel C: Gamblers' expected value from pumps vs. the number of participants (N')

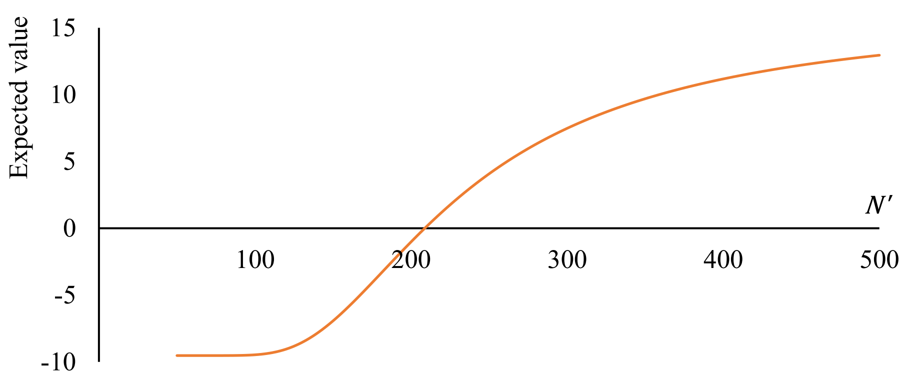

Figure 2.6. Gamblers' expected value from pumps for different parameter values.
The figure plots the expected value from participating in pumps for an individual with Cumulative Prospect Theory preferences (vertical axis). The horizontal axis is the price impact parameter (β) in Panel A, the manipulators' pre-pump long position (M) in Panel B, and the number of non-manipulators (N') in Panel C. For illustrative purposes, we set the price impact parameter $\beta = 0.1$, number of manipulators $M = 10$, number of non-manipulator participants $N' = 250$, and initial wealth $a = \$6$. We set the Cumulative Prospect Theory parameters as follows: value adjustment parameter $\alpha = 0.95$, probability weighting parameter $\delta = 0.5$, and loss aversion parameter $\lambda = 1.5$. We set the profit threshold, c, and the loss threshold, b, to their optimal values: $c = \$245$ and $b = \$0$.

A further implication is that, among pump participants, the mix of overconfident individuals and gamblers will vary over time. For example, all else equal, an increase in past returns on pumps is expected to increase the proportion of overconfident individuals among pump participants. Similarly, an increase in the market-wide propensity to gamble, all else equal, is expected to increase the proportion of gamblers.

2.4.3.2. Empirical results

We now test the predictions in Results 2 and 3. Overconfidence and gambling preferences are individual characteristics that are best measured at an individual level. However, our data do not allow us to identify individuals and measure their individual characteristics. Instead, we use a proxy for overconfidence measured at the pump level and an aggregate market-wide proxy for gambling activity.

Our main proxy for overconfidence exploits the past success or failure of the pump group; it is the average start-to-peak return earned in the two most recent pumps conducted by the Telegram pump groups orchestrating pump j of coin i on day t ($Overconfidence_{j,i,t}$).[24] Statman, Thorley, and Vorkink (2006) find that trading volumes in stock markets are positively correlated with previous returns and attribute this effect to positive returns inducing investor overconfidence. Because of the self-attribution bias, investors tend to attribute positive outcomes (e.g., earning high returns) to their own skill, thereby fueling overconfidence in their abilities. Given this logic, members of pump groups that have had recent success in conducting pumps with high returns are likely to be more overconfident about their prospects in subsequent pumps than others. Our hypothesis that overconfidence contributes to pump participation implies that we should expect to see higher participation in pumps conducted by groups likely to have more overconfident members. One potential issue is that our proxy might be capturing group speed or skill instead of overconfidence. To address this concern, we re-run our tests after adjusting our proxy to remove the effects of past group performance and report the results in Table 2.A2 in Appendix 2.E. The overconfidence results are robust to the use of this adjusted proxy.[25] Our results are also robust to

[24] The results are robust to the use of the three most recent pumps instead of two.
[25] Although the gambling proxy is not statistically significant in the combined test, the coefficient (0.24) is almost as large as the coefficient (0.28) in the individual test (see Table 2.2).

the use of an alternative, less granular proxy for overconfidence, the past five-day cryptocurrency market return.

Our main proxy for gambling among individuals who are part of the cryptocurrency ecosystem is the revenue of gambling services that accept gambles in Bitcoin. For example, SatoshiDICE is a relatively well-known cryptocurrency gambling site in which participants wager a certain amount of Bitcoin and receive a payoff determined by a random number generator. If individuals with gambling preferences use pump-and-dumps as yet another venue for gambling, pump participation rates are likely to be positively correlated with the general demand for gambling within cryptocurrencies.[26] Thus, our proxy for gambling activity ($Gambling_t$) is the daily log revenue of known Bitcoin gambling services identified in the Wallet Explorer (WE) database.[27] We remove the effects of general cryptocurrency market activity from this gambling measure by regressing it on the contemporaneous value and the three lagged values of the average daily return and trading volume for all cryptocurrencies and using the residuals from this regression.

For each of the 355 pump-and-dump manipulations in our sample, we measure the log total trading volume during pump j in coin i on day t from the release of the pump signal to three hours after the pump's peak price is reached. We regress this participation measure ($Participation_{j,i,t}$) on the $Overconfidence_{j,i,t}$ and $Gambling_t$ proxies. We control for lagged volatility, which captures differences in uncertainty about the coin value. We also control for the number of Telegram groups participating in the pump as a proxy for the number of manipulators because we expect fewer individuals to participate in pumps with more manipulators. Finally, we control for the number of members in Telegram groups participating in the pumps, the liquidity of the coin (log average daily trading volume), and the differences across the two exchanges by including exchange fixed effects (a $Yobit_i$ indicator variable). While testing the overconfidence channel in Model 1, we also include week fixed effects to absorb other time-series factors.

[26] The notion that individuals use cryptocurrencies as an alternative venue for gambling stems from the previously documented evidence of individuals substituting between traditional gambling methods (such as purchasing lottery tickets) and gambling in equity markets (Dorn, Dorn, and Sengmueller, 2014; Gao and Lin, 2015).

[27] WE identifies the Bitcoin wallets of many different entities and reports all blockchain transactions associated with these wallets. WE has a separate section listing Bitcoin wallets of various gambling websites like SatoshiDICE. We use the transaction activity for these wallets to measure the overall gambling activity in cryptocurrencies. Our proxy includes combined gambling from 43 gambling websites. The use of actual transactions recorded on the blockchain enhances reliability since self-reported data from gambling sites may be fabricated.

The results in Table 2.2 Models 1 and 2 show that both overconfidence and gambling have statistically significant positive associations with the level of participation in pumps. The positive associations are consistent with the notion that both overconfidence and gambling contribute to participation in pump-and-dump manipulations, in line with the mechanisms illustrated in the theoretical framework (Results 2 and 3). The coefficient for $Overconfidence_{j,i,t}$ is statistically significant at the 1% level, while that for $Gambling_t$ is statistically significant at the 5% level. The regressions in Table 2.2 (Model 3) show evidence that pump participation is negatively related to the number of manipulators, which is also consistent with the theory (Result 4). More manipulators imply greater aggregate losses for non-manipulators, thereby discouraging participation. Model 4 shows that uncertainty about the fundamental value of a coin (lagged volatility) does not explain pump participation, further supporting the notion that cryptocurrency pumps are not a form of information-based manipulation.

Model 5 includes all regressors. Both overconfidence and gambling retain their positive signs and statistical significance. We use the overconfidence and gambling coefficients from this regression to compare the economic impact of these two mechanisms on pump participation. When overconfidence increases from its 25th percentile to its 75th percentile (an increase of 87% in past pump returns), pump participation increases by 37.97%. Similarly, when gambling increases from its 25th percentile to its 75th percentile (an increase of $45,321 in the daily Bitcoin gambling volume), pump participation increases by 24.33%.

As an alternative way to interpret the magnitudes, when the overconfidence proxy increases by 10% from its mean of 69.16%, pump participation increases by 2.62%. Similarly, when the Bitcoin gambling volume increases by 10% from its mean of $97,528.81, pump participation increases by 3.39%. The magnitudes of both mechanisms appear to be economically meaningful and are robust to controlling for other drivers. We, therefore, conclude that both overconfidence and gambling are reasons why people participate in pumps.

We also consider an alternative proxy for overconfidence motivated by the literature on self-attribution bias, the past five-day cryptocurrency market return. We do not rely on this measure in our main tests, as it is less granular than our main proxy and can capture other confounding effects, such as the general level of interest or sentiment toward cryptocurrencies. Nevertheless, the time-series relation between this alternative overconfidence proxy and the level of pump activity is consistent with the overconfidence mechanism (unreported).

Table 2.2
Determinants of pump participation

This table reports regression results testing the determinants of participation in pump-and-dump manipulations. The dependent variable, $Participation_{j,i,t}$, is the log total trading volume during pump j in coin i on day t, measured from the release of the pump signal to three hours after the pump's peak price is reached. $Overconfidence_{j,i,t}$ is the average percentage return in the two most recent pumps organized by the Telegram groups participating in the pump. $Gambling_t$ is the log daily dollar revenue of Bitcoin gambling services after removing the effects of cryptocurrency market activity. $Volatility_{i,t-1}$ is the log of intraday volatility for the coin on day $t-1$ (the difference between the highest and lowest trade prices scaled by the volume-weighted average price). For missing values of lagged volatility, we use the average volatility for the coin. $Manipulators_{j,i,t}$ is the log number of Telegram groups participating in the pump. $Members_{j,i,t}$ is the log number of members in these Telegram groups. $Yobit_i$ is an indicator variable that equals one if the coin is traded on the Yobit exchange. $Liquidity_i$ is the log average daily dollar trading volume of the coin. The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018. t-statistics are in parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Participation_{j,i,t}$				
Variable	(1)	(2)	(3)	(4)	(5)
$Overconfidence_{j,i,t}$	0.57***				0.37***
	(4.89)				(2.93)
$Gambling_t$		0.28**			0.35**
		(2.04)			(2.33)
$Manipulators_{j,i,t}$			–0.20**		–0.23**
			(–2.38)		(–2.59)
$Members_{j,i,t}$			0.57***		0.51***
			(6.08)		(5.26)
$Volatility_{i,t-1}$				0.03	0.08
				(0.68)	(1.49)
$Liquidity_i$	0.15*	0.20***	0.20***	0.21***	0.09
	(1.78)	(2.65)	(2.76)	(2.86)	(1.18)
$Yobit_i$	–6.35***	–4.77***		–4.68***	
	(–7.84)	(–6.89)		(–6.98)	
Week fixed effects	Yes	No	No	No	No
R^2	87.58%	77.72%	21.51%	77.45%	23.80%
Observations	264	355	291	355	211

Figure 2.7 plots daily Bitcoin prices alongside the time series of the number of pumps per day (Panel A) and the average pump returns (Panel B). Both the number of pumps per day and pump returns follow patterns similar to those of the lagged Bitcoin price. This trend is consistent with the notion that overconfidence leads to higher participation in pumps and higher pump returns. Similarly, Figure 2.7 Panel C shows that the time series of our main proxy for participant overconfidence ($Overconfidence_t$) is positively related to daily pump participation (volume traded during pumps).

Panel A: Number of pumps and Bitcoin price through time

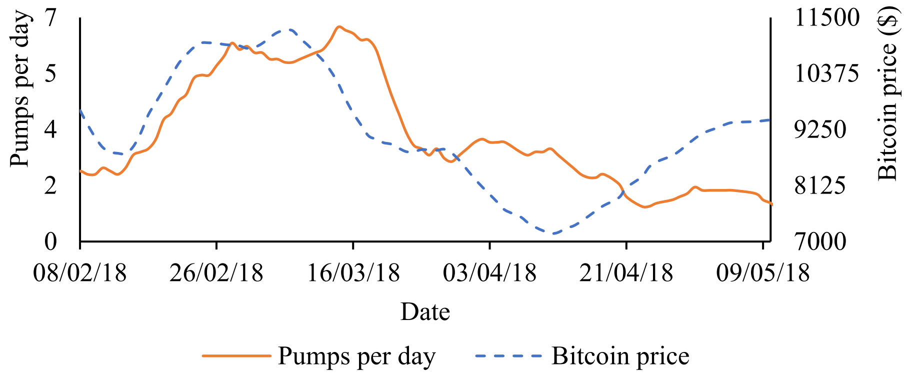

Panel B: Pump returns and Bitcoin price through time

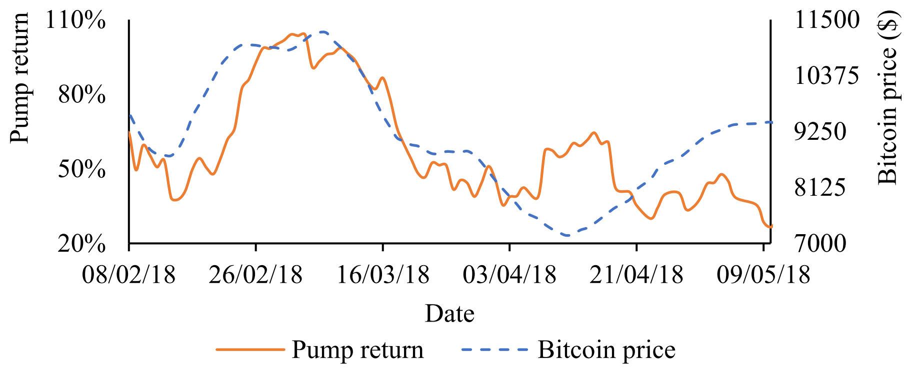

Panel C: Past pump return and pump volume through time

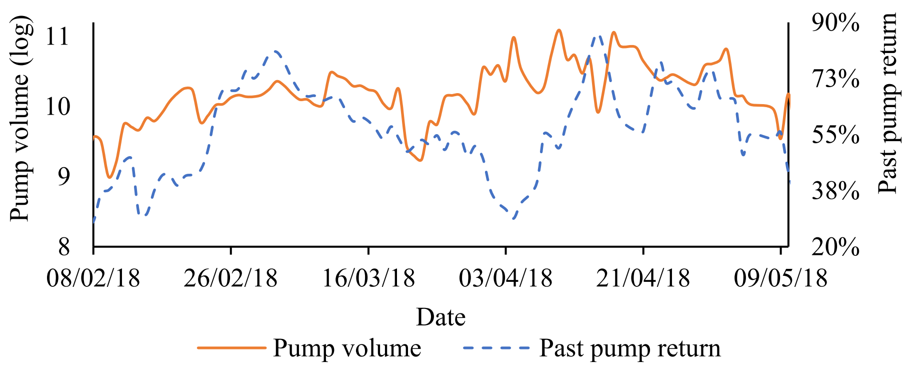

Figure 2.7. Pump dynamics through time.
Panel A plots the daily number of pumps and the volume-weighted average Bitcoin price. Panel B plots the average return on pumps (from start to peak) on a given day and the volume-weighted average Bitcoin price. Panel C plots the average return for the two most recent pumps conducted by the Telegram groups responsible for the pumps on a given day (our proxy for overconfidence) alongside the total trading volume during the pumps on that given day (our proxy for pump participation). All variables are smoothed using a ten-day moving average. The sample includes 355 pumps on the Binance and Yobit exchanges between December 2017 and June 2018.

We also implement a more direct test of the time-series drivers of participation in pumps. Recall that Result 5 predicts an increase in pump participation over time when there is an increase in the general level of interest in cryptocurrencies, higher past returns on pumps, and increasing market-wide gambling activity. To test these predictions, we estimate time-series regressions in which the dependent variable is daily pump participation (calculated as the daily average of the pump participation measure, $Participation_{j,i,t}$) and the independent variables are daily averages of the overconfidence and gambling proxies used in our previous tests. We also add a time-series proxy for the level of interest in cryptocurrencies; this proxy is the log of the average Google Search Volume Index for cryptocurrency-related search terms on day t ($SearchInt_t$).[28]

The results support the three predicted drivers of pump participation dynamics.[29] All three factors are individually positive and significant drivers of the time-series variation in cryptocurrency pump participation. The results show that much of the time-series variation in cryptocurrency pumps is attributable to variation in overconfidence, gambling, and general interest in cryptocurrencies.

2.5. Characteristics, determinants, and effects of pump-and-dumps

In this section, we document some additional predictions about cryptocurrency pump-and-dumps made by our theoretical framework and test these predictions empirically using our data. We also report certain empirical facts independent of our framework.

2.5.1. Pump-and-dump characteristics

Table 2.3 reports the characteristics of pump-and-dump manipulations. On average, pumps take around eight minutes to reach their peak price from the time the pump signal is sent (median time of 1.54 minutes) and generate an average return of 65.47% within this short period. For comparison, during our sample period, the highest daily return earned by Bitcoin is 22.72%, and that earned by the S&P 500 index is 2.72%. The return earned by the average pump in eight minutes is around three (24) times higher than the highest return earned by Bitcoin (S&P 500) in an entire day. The average pump return is also approximately four standard deviations higher than

[28] The search terms we include are "altcoin," "Bitcoin," "crypto," "cryptocurrency," and "ICO."

[29] Given these tests are similar to those reported in Table 2.2 but with less statistical power and generally similar conclusions, we report the results in Appendix 2.E (Table 2.A3).

the pumped coin's average daily return. This result suggests that pumps have a substantial effect on the prices of pumped coins, even after considering the high volatility of cryptocurrencies.

The impact of pumps is also evident in volume. The traded volume of pumped coins during manipulation days is, on average, approximately 13.5 times the average daily volume for the pumped coin. As pumps account for around 40% of the total pump-day volume, the trading volume generated by the average pump in eight minutes from start to peak is approximately five times the average daily trading volume for the pumped coin.

Finally, manipulators earn around 49% on an average pump.[30] We use the same method as in Section 2.3.2 to calculate this percentage profit figure. Profits are estimated as the difference between the volume-weighted average price during the pump (from start to peak) and that in the two hours preceding the pump signal. We calculate the percentage profit in each pump and take the average across pumps.[31] This profit figure corroborates our previous observation that cryptocurrency pumps provide manipulators with high returns within a short period.

Next, we examine the price and volume dynamics around pumps. In the theoretical framework (as illustrated in Figure 2.A1 in Appendix 2.B), prices rise before the release of a pump signal as the manipulators build their positions. Prices rise sharply once a pump signal is released as non-manipulators join the pump. Finally, prices reach a turning point, after which they fall back to pre-pump levels.

Figure 2.8 Panel A provides an empirical analog of these predicted price dynamics, showing the cumulative returns from 15 minutes before to 45 minutes after the release of a pump signal. Prices rise by approximately 10% in the 15 minutes preceding the pump signal, most likely due to the price impact of manipulators building their initial positions ($M\beta$ in the theoretical framework). Prices rapidly rise by a further 40% following the release of the pump signal because of the price impact of non-manipulators joining the pump ($N'\beta$ in the theoretical framework).[32] After reaching the peak, prices fall at a slower rate until they reach approximately the initial (pre-pump) level (P_0

---

[30] Manipulators here are pump group administrators who are aware of the coin being pumped before the pump signal and can thus build a long position before the signal and unload their positions on other participants during the pump.

[31] In contrast, the profit number in Table 2.1 was a percentage calculated from the aggregate earnings and the aggregate position of manipulators.

[32] The peak in Figure 2.7 (at around 40%) is lower than the average peak return (around 65% in Table 2.2) because in Figure 2.7, the individual pumps are aligned based on the pump signal ($t = 0$) but are not aligned in their peaks. Some pumps peak earlier than others, which is why Figure 2.7 does not reflect the average of the peak returns.

in the theoretical framework). The price trajectory in Figure 2.8 is very similar to that in the theoretical framework.

Table 2.3
Characteristics of pump-and-dump manipulations

This table reports statistics describing the characteristics of the sample of 355 pump-and-dump manipulations. Pump duration, return, and volume statistics are calculated from the start of a pump (the release of the pump signal) to its peak. A "pump-day" refers to a day on which there is at least one manipulation. Manipulators' percentage profit from a pump is calculated as the percentage difference between the volume-weighted average price during the pump (from start to peak) and the volume-weighted average price in the two hours preceding the release of the pump signal. Manipulators' dollar profit is calculated as their percentage profit multiplied by the pre-pump volume (volume in the two hours preceding the release of the pump signal). The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018.

Variable	Mean	Standard deviation	Median
Pump duration (minutes)	8.07	21.27	1.54
Pump return (%)	65.47%	84.44%	34.75%
Pump return (number of standard deviations of the coin's daily returns)	4.01	5.27	2.01
Pump volume (% of pump-day volume)	38.77%	24.22%	39.22%
Pump-day volume (% of the coin's average daily volume)	1,351%	1,978%	628%
Manipulators' profit (%)	49.02%	47.72%	39.36%
Manipulators' profit ($ thousands)	16.77	85.94	0.17

Panel B in Figure 2.8 plots the cumulative volume during the same one-hour window around the pump signal, expressing the cumulative volume as a percentage of the total volume in this one-hour window. Approximately 15% of the total trading volume during the pump occurs in the 15 minutes leading up to the release of the pump signal. This pre-pump trading activity is likely a result of the manipulators building their initial positions in the pumped coin (M in the theoretical framework). The highest trading rate occurs immediately following the release of the pump signal, as non-manipulators race to buy the coin being pumped. The rate of trading during the phase in which the price returns to its pre-pump level is more subdued.

Panel A: Return during pumps

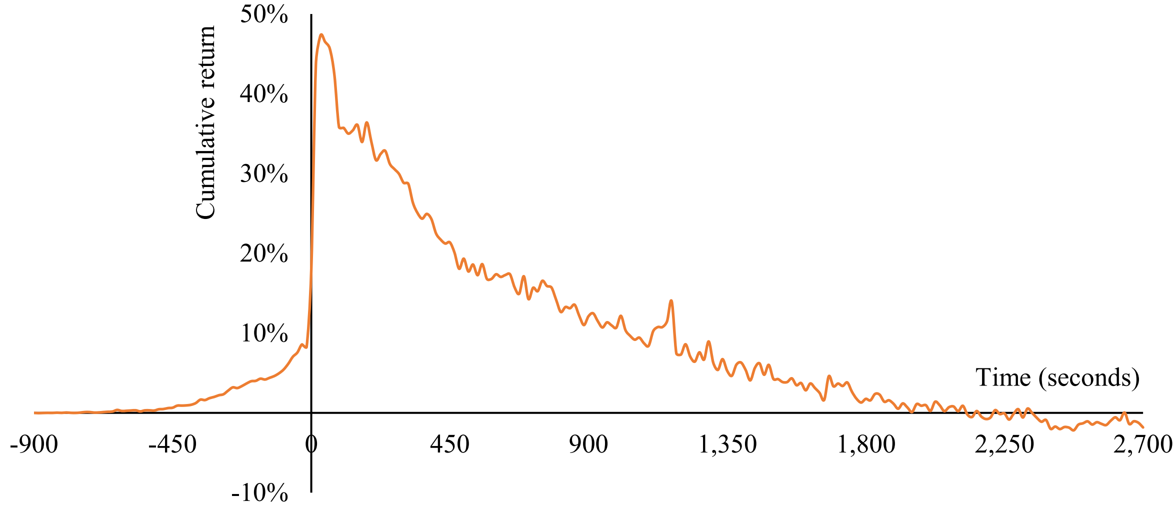

Panel B: Volume during pumps

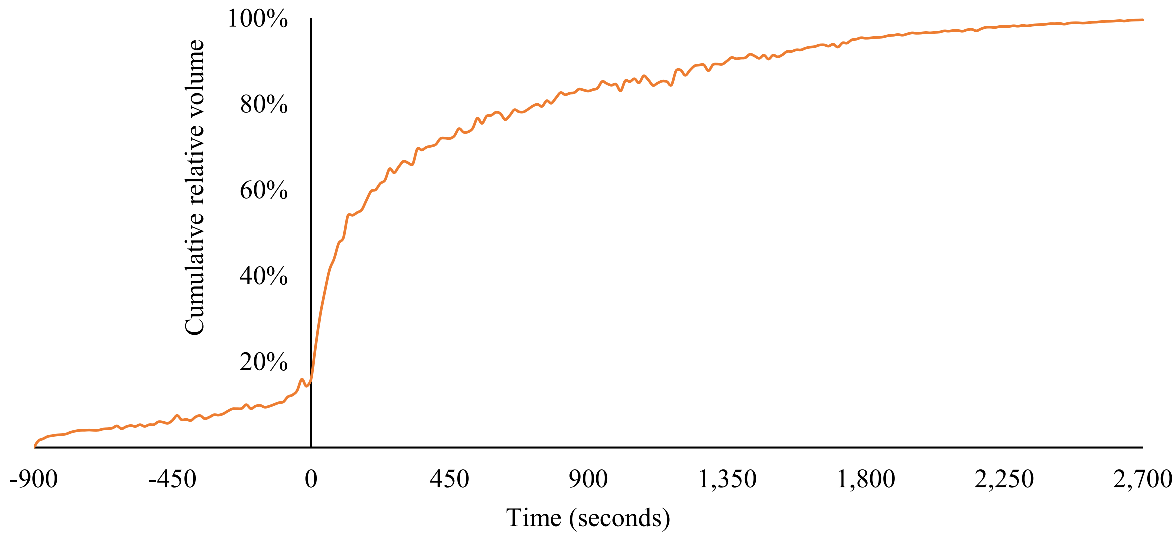

Figure 2.8. Return and volume dynamics during pump-and-dump manipulations.
The figure plots average cumulative returns (Panel A) and average cumulative volumes (Panel B) before, during, and after a pump signal ($t = 0$). The cumulative returns and volumes are measured in 15-second intervals from 15 minutes (900 seconds) before the pump signal to 45 minutes (2,700 seconds) after the pump signal. Cumulative volume is measured as a percentage of the total trading volume from 15 minutes before the pump signal to 45 minutes after the pump signal. The sample includes 355 pumps on the Binance and Yobit exchanges between December 2017 and June 2018.

2.5.2. Determinants of pump outcomes

A basic yet important prediction that follows directly from the expressions for the peak price and manipulator profits is that, *ceteris paribus*, pumps have higher peak prices, larger pre-pump-to-peak returns, and higher profits for the manipulators when there are more participants (high N') or the coin is less liquid (high β).[33]

Result 6*: Pumps with more participants and pumps in less liquid coins have higher peak prices and earn manipulators higher profits.*

Interestingly, Results 4 and 6 imply considerations of the size of the position the manipulators take initially. The manipulators face a trade-off; they naturally want to take a large position to earn large dollar profits but cannot take too large a position otherwise, there would be few or no non-manipulator participants and the pump would fail.

We test the predictions in Result 6 by regressing measures of pump outcomes on determinants at the pump level. Table 2.4 reports the results.

First, we test the determinants of manipulator profits, approximated by the difference between the volume-weighted average price in the two hours preceding a pump signal and that during the pump (from start to peak) multiplied by the trading volume in the two hours leading up to the pump. The theory predicts that manipulator profits are higher when more non-manipulators participate in the pump (higher N') and for pumps in less liquid coins (higher β). The former prediction is supported by the results in Table 2.4, which show a positive relation between manipulator profits and the level of pump participation. The estimates suggest that a 1% increase in pump participation is associated with a 0.96% increase in manipulator profits. Recall that manipulator profits come at the expense of non-manipulator pump participants. Intuitively, with more participants, manipulators can extract greater payoffs from pumps. The results are robust to the use of an alternative measure of pump participation, the log number of members in all Telegram groups participating in the pump (Table 2.A4 in Appendix 2.E). We find weak evidence, at best, in support of the second prediction regarding manipulators earning more profits for pumps in less

[33] The expression for peak price is $P_2 = P_1 + N'\beta$, while the expression for manipulators' profits is $\mathbb{E}[\pi_m] = \frac{\beta M}{2}(N' - 2M)$.

liquid coins. Although our liquidity measure has the expected direction (negative), it is not statistically significant.

Manipulators face a tradeoff in determining their optimal level of pre-pump trading; they want to trade a large volume to make a large dollar profit, but the more they trade, the lower the participation of non-manipulators. In the second regression, shown in Table 2.4, we find that manipulators tend to take larger positions before releasing the pump signal when they anticipate a higher rate of participation by non-manipulators and when the coin is more liquid since they can establish a larger position for a given level of price impact. The estimates suggest that a 1% increase in pump participation is associated with a 0.44% increase in the manipulators' pre-pump inventory position, and a 1% increase in liquidity increases the size of pre-pump positions by 0.46%.

The theory predicts that pumps have higher peak prices and thus higher pre-pump-to-peak returns when there is more participation in the pump (higher N') and less liquidity (higher β). Both these predictions are supported by the third regression in Table 2.4, which shows a positive (negative) relation between the pre-pump-to-peak returns and the level of pump participation (liquidity). The estimates suggest that a 1% increase in pump participation is associated with a 0.24% higher pump return.

Finally, we expect that, through time, pumps will become faster (the time from the pump signal to the pump peak will decrease) as slow individuals learn that they are at a disadvantage and cease to participate in pumps. The attrition of relatively slow individuals leaves a higher concentration of relatively fast individuals, thereby reducing the pump duration. This conjecture is supported by the fourth regression in Table 2.4. The pump duration is negatively related to our proxy for participant experience, which is the log number of pumps conducted in the past by the Telegram groups participating in the present pump. The estimates suggest that a 1% increase in participant experience reduces pump duration by 0.33%. This magnitude implies that pumps conducted by groups that have conducted three pumps previously are, on average, 16.5% faster than those conducted by groups that have conducted only two pumps previously.

Table 2.4
Determinants of pump outcomes

This table reports regression results testing how pump participation, participant experience, and liquidity affect manipulators' profit, pre-pump volume, pump return, and pump duration. The unit of observation is a pump j in coin i on day t. $ManipProfit_{j,i,t}$ is the manipulators' log dollar profit from the pump, calculated as the difference between the volume-weighted average price during the pump (from start to peak) and the volume-weighted average price in the two hours preceding the release of the pump signal multiplied by the trading volume in the two hours leading up to the pump. $PrePumpVol_{j,i,t}$ is the log dollar trading volume in the two hours preceding the release of the pump signal. $Return_{j,i}$ is the percentage return from the time of the pump signal to the peak of the pump. $Duration_{j,i,t}$ is the number of seconds from the release of the pump signal to the peak of the pump. $Participation_{j,i,t}$ is the log total dollar trading volume from the start of the pump to three hours after the pump's peak price is reached. $Experience_{j,i,t}$ is a proxy for participant experience in the pump, calculated as the log average number of pumps conducted before pump j by Telegram groups participating in pump j. $Liquidity_i$ is the log average daily dollar trading volume of the coin. The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018. t-statistics are in parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
Variable	$ManipProfit_{j,i,t}$	$PrePumpVol_{j,i,t}$	$Return_{j,i,t}$	$Duration_{j,i,t}$
$Participation_{j,i,t}$	0.96***	0.44***	0.24***	–0.14*
	(14.63)	(6.69)	(5.93)	(–1.49)
$Experience_{j,i,t}$	0.21**	0.17*	0.06	–0.33***
	(2.38)	(1.69)	(0.94)	(–3.14)
$Liquidity_i$	–0.08	0.46***	–0.23***	0.05
	(–1.62)	(8.65)	(–6.86)	(0.71)
R^2	84.47%	86.51%	27.38%	6.88%
Observations	181	174	189	189

2.5.3. Which coins are more likely to be pumped?

Overconfident individuals and individuals who treat pumps as a form of gambling are more likely to participate in pumps of relatively illiquid coins (see Figures 2.5 and 2.6), for which there are more extreme returns and greater dispersion in payoffs. However, coins require at least some minimum level of liquidity to be feasible; otherwise, manipulators cannot even establish their initial position in the coin.

***Result** 7: Non-manipulators are more attracted to pumps in relatively illiquid coins.*

The characteristics of coins that are more likely to be targeted by a pump-and-dump group depend on the manipulators' preferences, given that manipulators choose which coins to pump. Result 6 suggests that, all else being equal, a manipulator's profits increase with the illiquidity of the coin (β), as larger price impacts lead to higher returns, and the number of pump participants

(N'). The manipulators' profits do not depend on the type of participant or their motivation (e.g., overconfident individuals vs. gamblers) but depend on the aggregate number of participants that the manipulators can attract. Therefore, we can expect manipulators to target relatively illiquid coins and coins that are attractive to a large number of non-manipulators. Considering the preferences of non-manipulators, both overconfident individuals and gamblers are likely to prefer coins that are less liquid and have larger price impacts (see Result 7). Therefore, we expect that liquidity, whether measured directly or proxied by coin market capitalization, is the primary driver of which coins are more likely to be manipulated.

To test these predictions, we estimate logistic regressions of the probability that a coin is subject to at least one pump-and-dump manipulation during our sample period. The results in Table 2.5 indicate that pumps are more likely to occur in smaller coins (coins with lower market capitalization). This result is highly statistically significant, and the magnitude suggests that, as market capitalization is doubled, the odds of being pumped reduce by 14.79%.[34]

Although the theory predicts that both manipulators and non-manipulators prefer illiquid coins, for a pump to be feasible, there must be at least a minimum level of liquidity; otherwise, manipulators cannot establish a sufficient initial position to warrant the pumping of the coin. To allow for this potential non-linearity, we include, as regressors, dummy variables for the coin's liquidity quartile (quartiles of average daily number of trades and average daily dollar trading volume). We find that the likelihood of a coin being pumped is not monotonically related to liquidity. The second lowest liquidity quartile is the most likely to be targeted by manipulators, followed by the third lowest quartile. The highly liquid and illiquid coins are less likely to be manipulated. The coefficients of $TradesQ2_i$ and $TradesQ3_i$ indicate that coins in the second and third quartiles by the number of trades have 263.28% and 118.15%, respectively, higher odds of being pumped than those in the fourth (highest) quartile, while the lowest quartile coins have 42.31% lower odds of being pumped than the highest quartile coins.

The finding that the coins in the middle of the liquidity spectrum are most likely to be manipulated mirrors results from stock markets, where the ideal target for a manipulator is a stock that is sufficiently illiquid to be successfully manipulated but sufficiently large and liquid for the

[34] $MarketCap_i$ has a coefficient of -0.16. This coefficient implies an odds ratio of 0.8521 ($e^{-0.16} = 0.8521$). This odds ratio means that if $MarketCap_i$ increases by one unit, then the odds of the coin being pumped reduce by 14.79%. Since $MarketCap_i$ is the log to the base two of market capitalization, it increases by one unit when market capitalization is doubled.

manipulation to be worthwhile (Comerton-Forde and Putniņš, 2014). The results presented in Table 2.5 are robust to the use of an alternative measure of the likelihood that a coin is targeted by manipulators: the number of pumps conducted in the coin (Table 2.A5).

Table 2.5
Characteristics of pumped coins

This table reports the results of logistic regressions examining the characteristics of coins that are more likely to be pumped. The unit of observation is a coin i. The dependent variable, $Pumped_i$, is an indicator variable that equals one if the coin is pumped at least once during the sample period. $MarketCap_i$ is the log of the coin's average market capitalization in dollars. $TradesQ1_i$, $TradesQ2_i$, and $TradesQ3_i$ are indicator variables that equal one if the coin is in the first (lowest), second, or third quartile by average daily number of trades. $VolumeQ1_i$, $VolumeQ2_i$, and $VolumeQ3_i$ are indicator variables that equal one if the coin is in the first (lowest), second, or third quartile by average daily dollar trading volume. The sample includes all coins on the Binance and Yobit exchanges from December 2017 to June 2018. Chi-square statistics are in the parentheses below the coefficient estimates. Standard errors are clustered by coin. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Pumped_i$		
Variable	(1)	(2)	(3)
$Intercept$	2.05***	–2.29***	–2.35***
	(4.00)	(–11.95)	(–12.07)
$MarketCap_i$	–0.16***		
	(–7.30)		
$TradesQ1_i$		–0.55*	
		(–1.77)	
$TradesQ2_i$		1.29***	
		(5.65)	
$TradesQ3_i$		0.78***	
		(3.25)	
$VolumeQ1_i$			–1.10***
			(–2.92)
$VolumeQ2_i$			1.44***
			(6.24)
$VolumeQ3_i$			0.92***
			(3.85)
Pseudo-R^2	7.55%	5.36%	7.81%
Observations	609	1,307	1,307

2.5.4. Impact of pump-and-dumps on market characteristics

Finally, we analyze how pumps affect market characteristics, both in the market that is pumped and the other markets that trade the pumped coin.[35] We measure trading volumes, returns, and volatility for each coin-day (for coins that are pumped at least once during the sample period) for

[35] As illustrated in Section 2.2, pump group administrators typically specify not only a coin but also a cryptocurrency exchange on which the pump will occur, which we refer to as the "pumped market." Other exchanges that also trade the same coin may be impacted differently and therefore we analyze them separately.

the pumped market and the other markets that trade the pumped coin. We then regress these trading characteristics on an indicator for whether the coin was the target of a pump-and-dump manipulation that day ($PumpDay_{i,t}$). We control for the coin's market capitalization, market fixed effects (a dummy variable for the Yobit exchange), coin fixed effects, and time fixed effects.

The results in Table 2.6 Panel A show the effects of pumps on the pumped market. Traded volume is significantly increased during pumps on the pumped exchange, even after controlling for the various other coin and time effects. The coefficient for $Volume_{i,t}$ is statistically significant at the 1% level in all specifications. The coefficient suggests that the trading volume is about 30 to 70 times higher when a coin is pumped.[36]

Interestingly, the coin's return measured from before the pump (day $t-1$) to after the pump (day $t+2$, with day t being the day of the pump) is not significantly affected by the pump despite the earlier results showing that at intraday horizons, pumps generate large returns in the order of 65%. Thus, these regression results confirm that, consistent with the theory, the prices of pumped coins return to their pre-pump levels following the conclusion of a pump, with no permanent effects on coin valuations.

Lastly, the results show that volatility, measured from the high/low price range for the coin on day t, is considerably higher on the pumped market when a pump-and-dump manipulation occurs. The coefficient for $Volatility_{i,t}$ is statistically significant for all specifications and indicates that intraday volatility is between six and eight times higher when a pump occurs than it is at other times, after controlling for other factors.[37]

Although the impact of pumps on volatility is very large *conditional* on a pump occurrence, overall, pumps are not the primary driver of the extremely high levels of volatility in cryptocurrencies. In our sample, the identified pumps occur on less than 5% of the coin-days. Given that the identified cases provide a lower bound on the prevalence of pump-and-dump manipulations in cryptocurrencies, we assume, for illustration, that pumps occur approximately 5% of the time. If cryptocurrency volatility in the absence of pumps is σ and pumps increase this volatility to 7σ conditional on a pump occurrence (as estimated in the regressions), then the

[36] With the natural log of trading volume as the dependent variable, the estimated effect of a pump is to increase volume by $e^{3.38} = 29.4$ to $e^{4.27} = 71.5$ times its usual level.

[37] With the natural log of high/low range-based volatility as the dependent variable, the estimated effect of a pump is to increase volatility by $e^{1.86} = 6.4$ to $e^{2.04} = 7.7$ times its usual level, all else equal.

overall, unconditional effect of pumps is to increase volatility by 30% of its natural level.[38] While a 30% increase is a meaningful contribution to the overall volatility, it is only a small fraction of the "excess" volatility in cryptocurrencies as compared to other asset classes. For example, Liu and Tsyvinski (2021) show that Bitcoin returns are five to ten times more volatile than stocks, depending on the return frequency.

Table 2.6 Panel B tests for spillover effects on other markets that trade the pumped coin but are not directly targeted by the pump group administrators. The results show that the trading volume for the pumped coin tends to increase in other exchanges that trade the coin during a pump. The coefficient for $PumpDay_{i,t}$ is positive in all specifications and statistically significant in Model 2 when we include coin fixed effects. Volatility of pumped coins is also significantly elevated on other exchanges during pump-and-dumps, although not as much as the increase in the pumped market. These results are consistent with arbitrageurs somewhat, but not fully, aligning the prices of coins across markets; some of the price distortion of a manipulated coin in a manipulated market is reflected in the prices of the coin in other markets.

2.6. Welfare implications and regulation

Cryptocurrency pump-and-dump schemes affect welfare in three main ways.

First, pumps cause wealth transfers. In aggregate, wealth is transferred from the least sophisticated players (e.g., slow players, gamblers, and overconfident players) to manipulators and more sophisticated players (e.g., fast players). The losses incurred by gamblers do not necessarily decrease their welfare because gamblers receive utility from the act of gambling itself (Conlisk, 1993). Their losses are gains for manipulators and sophisticated players; therefore, wealth transfers from gamblers might increase aggregate welfare by the amount of utility obtained from gambling. There is, however, some scope for harm to gamblers given the lack of regulatory oversight. In other settings, gambling tends to be regulated to reduce the risk of gamblers being excessively exploited by gambling service providers. No such controls exist in cryptocurrency pump-and-dump games.

[38] Given these assumptions and estimates, the volatility in the presence of pumps is $0.95\sigma + 0.05(7\sigma) = 1.30\sigma$, i.e., 1.3 times its natural level, σ.

Table 2.6
Effects of pump-and-dump manipulations

This table reports regression results testing how volume, return, and volatility are impacted by pump-and-dump manipulations. The unit of observation is a coin-day, i, t. $Volume_{i,t}$ is the log dollar trading volume. $Return_{i,t}$ is the percentage return for the coin from day $t-1$ to day $t+2$. $Volatility_{i,t}$ is the log of intraday volatility for the coin on day t (the difference between the highest and lowest trade prices scaled by the volume-weighted average price). The independent variable of interest is $PumpDay_{i,t}$, which equals one if there is a pump for the given coin on the given day and zero otherwise. $Yobit_i$ is an indicator variable that equals one if the coin is listed on the Yobit exchange. $MarketCap_i$ is the log of the coin's average dollar market capitalization. Panel A displays the results for the exchange that is the target of the manipulation. Panel B displays the results for all exchanges on which the manipulated coin is listed except the target exchange. The sample includes coins targeted in 355 manipulations on two exchanges (Binance and Yobit) from December 2017 to June 2018. t-statistics are reported in parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =								
	$Volume_{i,t}$			$Return_{i,t}$			$Volatility_{i,t}$		
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Pump exchange									
$PumpDay_{i,t}$	3.38***	3.50***	4.27***	–0.02	0.01	0.03	1.86***	1.87***	2.04***
	(31.12)	(35.57)	(30.35)	(–0.60)	(0.48)	(0.97)	(37.60)	(36.97)	(34.89)
$Yobit_i$	–9.21***		–6.69***	0.01		–0.01	0.00		–0.08***
	(–164.76)		(–79.80)	(0.57)		(–1.06)	(0.07)		(–2.58)
$MarketCap_i$	0.05***		0.35***	–0.01***		–0.01***	–0.03***		–0.02***
	(5.24)		(29.63)	(–4.39)		(–4.63)	(–3.72)		(–3.35)
Coin fixed effects	No	Yes	No	No	Yes	No	No	Yes	No
Day fixed effects	No	No	Yes	No	No	Yes	No	No	Yes
R^2	82.23%	4.13%	69.34%	0.39%	0.00%	0.23%	2.32%	2.48%	2.63%
Observations	14,542	24,816	14,542	14,542	24,814	14,542	13,295	20,826	13,295
Panel B: Other exchanges									
$PumpDay_{i,t}$	0.09	0.25**	0.25	–0.09***	–0.00	–0.04**	0.30***	0.26***	0.38***
	(0.50)	(2.38)	(1.44)	(–4.95)	(–0.02)	(–2.30)	(4.30)	(3.89)	(5.45)
$MarketCap_i$	1.00***		1.00***	–0.02***		–0.02***	–0.03***		–0.03***
	(125.29)		(135.72)	(–5.86)		(–5.83)	(–11.05)		(–10.60)
Coin fixed effects	No	Yes	No	No	Yes	No	No	Yes	No
Day fixed effects	No	No	Yes	No	No	Yes	No	No	Yes
R^2	55.76%	0.00%	59.52%	0.14%	0.00%	0.14%	1.18%	0.09%	1.19%
Observations	11,995	13,240	11,995	14,452	19,276	14,452	11,887	13,094	11,887

In contrast, the transfers from slow and overconfident players to manipulators and more sophisticated players are likely to decrease aggregate welfare if the less sophisticated players tend to be less wealthy and, therefore, have a higher marginal utility of wealth than the more sophisticated players. In other financial markets, transfers from less sophisticated to more sophisticated participants are accompanied by an offsetting social benefit, i.e., information production and the provision of price discovery by sophisticated investors. No such offsetting benefit is present in cryptocurrency pump-and-dump manipulations.

Second, cryptocurrency pump-and-dumps, like other forms of market manipulation, cause price distortions that harm price accuracy and informativeness. The price distortions could, in theory, degrade the efficiency of resource allocation. However, because the price distortions from pumps are short-lived and tend to disappear within two days of the pump, it is difficult to imagine that cryptocurrency pump-and-dump schemes have any material effects on resource allocation.

Third, widespread manipulation damages the integrity of cryptocurrency markets and investor confidence with respect to tokens and tokenization. Financial institutions that are concerned about their reputation might choose not to associate themselves with cryptocurrency markets or tokenization more broadly. Additionally, regulators cite manipulation as a reason to restrict the expansion of cryptocurrency markets, tokenization, and cryptocurrency-linked products. For example, the US Securities and Exchange Commission has rejected multiple applications for Bitcoin Exchange-Traded Funds because of concerns about market manipulation in the Bitcoin market. Lack of confidence/trust in cryptocurrency markets can impede the development and adoption of cryptocurrencies and the tokenization of financial securities. These effects potentially have large negative welfare consequences.

While it is virtually impossible to quantify these welfare consequences, the effects of a stunted adoption of these new technologies could be an order of magnitude larger than the welfare implications of wealth transfers from participating in cryptocurrency pump-and-dump schemes. Therefore, widespread manipulation in cryptocurrency markets matters for reasons beyond the immediate and direct wealth transfers they create.

The cryptocurrency ecosystem has many potential benefits. For example, initial coin offerings (ICOs) are a promising approach to raising capital—see Cong, Li, and Wang (2018, 2020) and Li and Mann (2020) for models that show the benefits of using digital tokens. Howell, Niessner, and

Yermack (2019) find that platforms that list their tokens on exchanges post-ICO tend to be more successful, suggesting that a secondary market contributes to the success of token-based platforms.

A contribution of this chapter is that it increases our understanding of the downsides of the cryptocurrency ecosystem so that the sector can be efficiently regulated rather than being shut down entirely. Our results suggest that, without appropriate regulation and enforcement, widespread manipulation will prevail in cryptocurrency secondary markets. Widespread manipulation can reduce confidence in cryptocurrencies and tokens and impede their growth and adoption. Currently, a lack of regulation and enforcement, as well as weak or absent oversight from exchanges, allows this form of manipulation to persist and flourish.[39] If left unchecked, this manipulation can lead to a loss of confidence in these markets and in the technology of tokenization more broadly. Regulators must, however, be careful when opting for this path since overregulation might also cause similar harms of hindering growth and innovation.

2.7. Conclusion

Cryptocurrencies have given rise to a new form of pump-and-dump manipulation, which is similar in some respects to traditional pump-and-dump manipulation of stocks but completely different in other respects.

Like pump-and-dump manipulation of stocks, cryptocurrency pumps generate large price distortions (average price movements around 65%), generate abnormal trading volumes (13.5 times the average volume), and earn the manipulators millions of dollars. Similar to stock manipulation, manipulators target fairly illiquid coins, although they avoid coins with such little liquidity that manipulation would be infeasible or not sufficiently profitable. Although the targeted coins experience extreme returns while being manipulated, their prices subsequently revert to their pre-manipulation levels, with no apparent long-lasting effects on their valuations.

However, in contrast to pump-and-dump manipulation in stock markets, cryptocurrency pump-and-dumps do not rely on information asymmetry and uncertainty about the value of the manipulated security—manipulators openly declare their intentions to manipulate particular coins. These manipulations also do not rely on asymmetry in price impacts, as in standard trade-based manipulation of stocks. Rather, our evidence suggests that cryptocurrency pump-and-dumps are

[39] Appendix 2.C provides a background for the current regulatory and market surveillance environment in cryptocurrency markets.

akin to a gambling game in which players compete to buy a pumped coin ahead of others and sell near the peak, before the price collapses. Rational individuals with correct beliefs do not participate in cryptocurrency pumps (unless they have an advantage over other participants), as the pumps constitute a negative-sum game. However, individuals with gambling preferences participate under certain conditions, as do overconfident individuals, who overestimate their ability to sell near the peak price. We find empirical support for both of these explanations. Our findings about the determinants of participation in cryptocurrency pumps are likely to apply to other similar speculative games in markets, such as the GameStop episode.

Appendix 2.A. Pump-and-dump groups

Big Pump Signal (BPS) appropriately illustrates how pump-and-dump groups operate. Most groups communicate with their members on one of two instant messaging platforms, Telegram and Discord; BPS uses both. A unique feature of these messaging platforms is that they allow users to form public groups in which only the administrator can broadcast messages to the group members. These groups are public insofar as any Discord or Telegram user can find and join them.

Most messages sent by group administrators fit into one of the following four categories. The first category is messages that convey information about pumps, such as details about the date/time/exchange of upcoming pumps, the coin being pumped, or the pump results. The second category is messages that provide guidance on how pumps work and how members can profit from pumps. The right-hand-side message in Figure 2.1 provides an example in which the group's administrators advise members to enter the pump as soon as possible and dump their holdings on outsiders, such as technical traders, who might buy the coin based on the upward price movements. The third category is success stories of people profiting from pumps. It is possible that these success stories are intended to encourage participation by creating a perception of large potential profits. Lastly, the fourth category is messages that advertise the opportunity to become a paid member of the group. Paid members (as distinct from the individuals who follow the group messages at no cost) are provided pump-related information (e.g., the coin name) in advance of the general broadcast to all group members, allowing them to take positions in coins ahead of the official pump signal.

Unlike conventional pump-and-dump manipulation, wherein manipulators typically try to mislead market participants into thinking the manipulated security is fundamentally undervalued (Leuz et al., 2017), in cryptocurrency pumps such as the one illustrated above, the manipulators' goal is to convince people to participate in the pump. Cryptocurrency manipulators typically do not seek to trick people into believing that a coin is mispriced on the basis of fundamentals—they explicitly communicate to the pump group members that a coin is being pumped, as opposed to presenting an excellent investment opportunity. We randomly sample 70 pumps and manually categorize each pump signal sent to the pump group members. In none of the 70 cases is there any suggestion that the target coin is undervalued.

Although manipulators are explicit in telling their group members about the intention to pump, they may also hope to attract others, such as technical traders, to buy the manipulated coin following the initial sharp increase in price. For example, the message in Figure 2.1 urges pump group members to hold for a long period so that outsiders have a chance to enter and “we can all profit.” Such messages can appear in isolation as general “how to pump better” messages or accompany the pump signal.

In the sample of the 70 pump signals that we manually categorize, we find that 27 signals (39%) urge participants to hold the coin longer to attract outsiders to the pump. This feature of the schemes adds to the uncertainty about how many traders will be drawn to participate in the pump and, consequently, at what point the pump will peak.

Appendix 2.B. Framework details

2.B.1 Price dynamics

We assume that market orders have linear price impacts, consistent with microstructure models of market making (Kyle, 1985). A further motivation for linear price impacts is that they rule out simple trade-based manipulation strategies that could otherwise make unlimited profits by simply buying and selling (Huberman and Stanzl, 2004). The previously discussed anecdotal evidence suggests that cryptocurrency pumps do not seek to exploit non-linearity or asymmetry in price impacts. We fail to reject the null hypothesis that price impacts are symmetric in our empirical data.[40] As mentioned in the body of the chapter, prices are determined by the function, $P_t = P_{t-1} + \beta x_t$. Cumulating the net volume received up to time t, $X_t = \sum_{\tau=1}^{\tau=t} x_\tau$, we can rewrite the pricing function as $P_t = P_0 + \beta X_t$. At the end of Period 1, after the manipulators have bought M units, pushing the price up by $M\beta$, the price of the coin is $P_1 = P_0 + M\beta$.

In Period 2, if N' players choose to participate, having a combined price impact of $N'\beta$, the price at the end of Period 2 (which is the price paid by the participant whose order arrives at the market last) will be $P_2 = P_1 + N'\beta = P_0 + M\beta + N'\beta$.

In Period 3, the individual sell orders posted by both manipulators and non-manipulators are executed at prices $\{(P_2 - 1\beta), (P_2 - 2\beta), \ldots, P_0\}$ depending on their random latency and queue position. The price at the end of Period 3 (which is the price received by the participant whose order arrives at the market last) will be $P_3 = P_0$, because once the manipulators and all pump participants have liquidated their long positions, the cumulative net volume of buys, X_3, is zero.

Figure 2.A1 illustrates the timing and price dynamics in this simple framework under different parameter values. In the baseline illustration ($P_0 = \$5$, $M = 10$, $N' = 100$, $\beta = 0.2$), the price starts at $P_0 = \$5$ and rises to $P_1 = \$7$ once the manipulators buy ten units. The price rises further to $P_2 = \$27$ once the one hundred pump group participants buy the coin, with P_2 being the peak price of the pump. Finally, the price falls back down to $P_3 = \$5$ once the manipulators and participants exit the pump. With more manipulators (higher M), there is a larger run-up in Period 1 before the pump signal is sent and, consequently, a higher peak price. With more pump participants (higher N'), there is a sharper price rise in Period 2 and a higher peak. When there is

[40] We test whether price impacts are symmetric in our data by calculating the Amihud illiquidity ratio (Amihud, 2002) in 30 minute intervals and comparing the measure for price increases and decreases. Illiquidity ratios for price increases are not statistically different from the ratios for price decreases at the 5% significance level.

less liquidity and a higher price impact parameter (β), there is a larger run-up in Period 1 before the pump signal is sent and a sharper price rise after the signal as participants buy the coin.

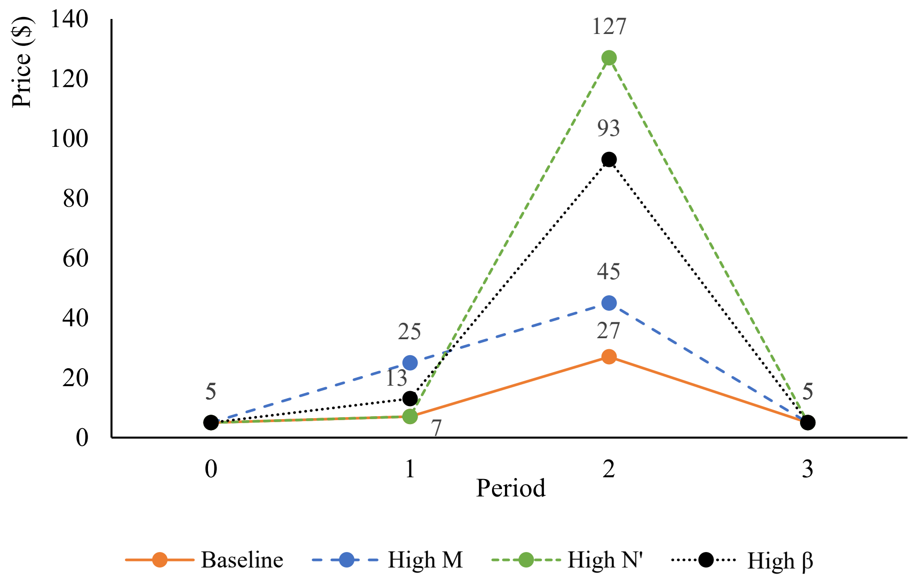

Figure 2.A1. Price dynamics in the theoretical framework.
The figure illustrates price paths during pumps based on the theoretical framework. For the baseline illustration, we set initial price $P_0 = \$5$, number of manipulators $M = 10$, number of participants $N' = 100$, and price impact parameter $\beta = 0.2$. In the high M specification, we increase the number of manipulators to $M = 100$, keeping all other parameters at their baseline levels. In the high N' specification, we increase the number of participants to $N' = 600$, keeping all other parameters at their baseline levels. In the high β specification, we increase the price impact parameter to $\beta = 0.8$, keeping all other parameters at their baseline levels.

2.B.2 Fast and slow agents

The exact specification of the tilted exit price pdfs for fast and slow agents is as follows:

$$f_i(P_{exit}) = \begin{cases} \frac{1}{\beta(M+N'-1)} + S_i P_{exit} - \frac{S_i(\beta(M+N'-1)+2P_0)}{2} & \text{if} \quad P_0 \leq P_{exit} \leq P_0 + \beta(M+N'-1) \\ 0 & \text{otherwise} \end{cases} \tag{2.B1}$$

where $S_i = s^* > 0$ for fast players and $S_i = s^* < 0$ for slow players.

A fast (slow) player earns higher (lower) expected profits than an "average" participant, i.e., $\mathbb{E}[\pi_i]$ in Eq. (1). The expected profit for fast (slow) players is higher (lower) than that for an average participant by $\frac{\beta^3 S_i}{2}(M+N-1)^3$. Manipulator profits remain unchanged; therefore, the increased expected profit of fast participants is at the expense of slow participants. Fast participants earn positive expected profits when they have a sufficiently large advantage, i.e., when $S_i > \frac{6(M+2)}{\beta^2(M+N'-1)^3}$.

2.B.3 Overconfident agents

The *perceived* pdf of exit prices for overconfident actors is written as follows:

$$f_i(P_{exit}) = \begin{cases} \frac{1}{\beta(M+N'-1)} + \varepsilon_i P_{exit} - \frac{\varepsilon_i(\beta(M+N'-1)+2P_0)}{2} & \text{if} \quad P_0 \leq P_{exit} \leq P_0 + \beta(M+N'-1) \\ 0 & \text{otherwise} \end{cases} \tag{2.B2}$$

Replacing the unbiased exit price distribution with the biased one, we find that the expected pump payoff, $\mathbb{E}[\pi_i]$, for an overconfident player i is strictly positive when

$$\varepsilon_i > \frac{6(M+2)}{\beta^2(M+N'-1)^3} \equiv \varepsilon^{min}. \tag{2.B3}$$

The condition in Eq. (2.B3) implies that cryptocurrency pump-and-dumps have positive expected payoffs for sufficiently overconfident individuals. Assuming risk-neutrality, individuals with overconfidence exceeding the threshold ε^{min} would participate in pumps.

2.B.4 Gamblers

The series of pumps following the strategy q (participate until wealth reaches either \$$b$ or \$$c$) can be reduced to a binary gamble, $\tilde{G}_q$,

$$\tilde{G}_q \sim \Big((c-a), \mathbb{P}(c-a); -(a-b), \mathbb{P}\big(-(a-b)\big)\Big) \tag{2.B4}$$

where $\mathbb{P}(c-a)$ represents the probability of achieving a $(c-a)$ gain and $\mathbb{P}\big(-(a-b)\big)$ represents the probability of suffering a loss of $(a-b)$. We can estimate the probabilities $\mathbb{P}(c-a)$ and $\mathbb{P}\big(-(a-b)\big)$ by approximating the gambler's wealth as a Brownian motion starting at \$$a$ and terminating upon reaching one of two absorbing barriers on either side of the starting point, \$$b$ and \$$c$.[41] Based on Dixit (1993), the estimates for $\mathbb{P}(c-a)$ and $\mathbb{P}\big(-(a-b)\big)$ are:

$$\mathbb{P}(c-a)=\frac{\exp\left(-\frac{2a\mu}{\sigma^2}\right)-1}{\exp\left(-\frac{2c\mu}{\sigma^2}\right)-1} \qquad \text{and} \qquad \mathbb{P}\big(-(a-b)\big)=1-\mathbb{P}(c-a) \tag{2.B5}$$

where μ and σ^2 are the mean and variance of the payoffs from a single pump: $\mu=\frac{-\beta}{2}(M+2)$ and $\sigma^2=\frac{\beta^2}{12}(M^2+2M(N'-1)+2(N'-1)^2)$.

Following Barberis (2012), a gambler chooses a strategy q (the strategy gives the values b and c that determine when the gambler stops playing) from a set of strategies Q to solve the maximization problem:

$$\max_{q\in Q} V(\tilde{G}_q) \tag{2.B6}$$

where $V(\cdot)$ is the expected value of the gamble based on the Cumulative Prospect Theory (CPT) developed by Tversky and Kahneman (1992). According to CPT, individuals overweight small probability outcomes and underweight large probability outcomes. Additionally, individuals' valuation of payoffs is concave in the region of gains and convex in the region of losses.[42] These features of CPT explain why individuals are attracted to positively skewed payoffs and serve as the basis for various theories of gambling (Barberis and Huang, 2008; Barberis, 2012). For the figures and the numerical solution, we set the CPT parameters as the benchmark estimates in Barberis (2012), $\alpha=0.95$, $\delta=0.5$, and $\lambda=1.5$.

2.B.5 Repeated games

Next, we examine the dynamics of repeated pump-and-dump games, which is crucial to understanding how cryptocurrency pumps-and-dumps evolve through time. Participation in pumps is determined by various factors that can change over time. Let N_t^{OC}, N_t^{GP}, and N_t' denote the number of sufficiently overconfident individuals (OC), the number of participants with gambling

[41] Given that the payoff from a single pump follows a trapezoid distribution rather than a normal distribution, the Brownian motion approximation assumes a sufficiently large number of small bets.
[42] The full set of equations for the value function and probability weighting function are in Appendix 2.D.

preferences (GP), and the total number of participants ($N'_t = N_t^{OC} + N_t^{GP}$), respectively, in the t^{th} iteration of the game (we also refer to these iterations as rounds).

As overconfident individuals observe the actual prices at which they buy and sell pumped coins and the profits or losses they incur, we expect them to revise their beliefs about their skill and, consequently, about their expected profits. Effectively, some of their overconfidence bias should attenuate over time as they learn about their skill and payoffs, consistent with other models of overconfidence (Gervais and Odean, 2001). In this process, as their overconfidence bias attenuates (as $\varepsilon_i \to 0$), some overconfident individuals cross the threshold ε^{min} and switch from participating in pumps to not participating. Let ρ_t^{OC} denote the number of overconfident individuals who switch to not participating after round t. Additionally, let λ^{OC} be the Poisson arrival rate of new, sufficiently overconfident individuals.

Similarly, individuals with gambling preferences stop participating in pumps when they achieve their desired gains of $\$(c-a)$ or lose $\$(a-b)$. Let ρ_t^{GP} denote the number of gambling-motivated participants who cease to participate in pumps after round t. Furthermore, let λ^{GP} be the Poisson arrival rate of new participants with gambling preferences. Given these new parameters, the dynamics for the size of the participant pool can be expressed as follows:

$$\mathbb{E}[\Delta N'_t] = \mathbb{E}[N'_t - N'_{t-1}] = (\lambda^{OC} + \lambda^{GP}) - (\rho_{t-1}^{OC} + \rho_{t-1}^{GP}). \qquad (2.B7)$$

According to the above equation, the change in the number of pump participants depends on the rate at which new overconfident individuals and individuals looking for a gamble arrive at the market and the rate at which existing participants cease to participate after learning about their ability or hitting their maximum gains or losses.

Appendix 2.C. Cryptocurrency markets and regulation

Why is such manipulation conducted so openly in cryptocurrency markets? The answer to this question is four-fold. First, cryptocurrency exchanges are underequipped to detect and prevent most forms of market manipulation, including pump-and-dump schemes. The Office of the New York State Attorney General (ONYSAG), the highest law office in New York, US, confirms this in a recent investigation into cryptocurrency exchanges.[43] Out of the ten exchanges investigated, only four have formal policies defining actions that constitute manipulative activity. Surprisingly, none of the exchanges has adequate market surveillance methods to detect manipulation, and only two exchanges are working on implementing better surveillance tools. The lack of supervision by exchanges might be an issue of incentives—by turning a blind eye to manipulation, they generate higher trading volumes from the presence of manipulation schemes, thereby increasing their revenue from fees. However, according to the literature, competition between exchanges ensures that exchanges have an incentive to curb market manipulation since they stand to lose potential income if traders wary of manipulation take their trades elsewhere (e.g., Easterbrook, 1986; Daniel and David, 1991).

Second, there is virtually no action from regulators or law enforcement agencies to counter these cryptocurrency pump-and-dump schemes. Currently, the only regulatory recognition of these schemes is a notice from the US Commodities and Futures Trading Commission advising the public to be wary of pump-and-dump manipulation on cryptocurrency exchanges and announcing a reward for whistleblowers who report manipulators.[44] This lack of attention from exchanges and regulators implies that cryptocurrencies are vulnerable to market manipulation.

The last two factors that enable cryptocurrency pump-and-dumps are the increasing participation of the public in cryptocurrency markets and strong speculative sentiment among participants. Both of these factors ensure that there is no shortage of people from whom manipulators can extract profits. Cryptocurrency markets have recently witnessed a significant rise

---

[43] The ONYSAG conducted an in-depth investigation into the working practices of ten cryptocurrency exchanges: Bitfinex, Bitflyer, Bitstamp, Bittrex, Coinbase, Gemini, Hbus, Itbit, Poloniex, and Tidex. Among the exchanges with known instances of pump-and-dump manipulation (Yobit, Cryptopia, and Binance), one exchange (Binance) was contacted by the Attorney General's office to join the investigation. Binance, however, refused to participate stating that their exchange does not allow trading from New York. ONYSAG conducted a preliminary investigation into whether Binance operated in New York and referred them to the New York Department of Financial Services for potential violation of the state's virtual currency regulations. The report can be accessed at this link (https://virtualmarkets.ag.ny.gov).

[44] This advisory statement can be found at this link (https://bit.ly/2NH3wj0).

in trading activity, with the trading volume in 2018 exceeding the total trading volume in the preceding five years. Investors entering cryptocurrency markets could increase participation in cryptocurrency pumps, as manipulators can attract a larger audience into their pump groups. Previous research shows that price movements in cryptocurrencies have a sizeable speculative component (Cheah and Fry, 2015), and speculation-driven investors exhibit a higher tendency to participate in pump-and-dump schemes (Leuz et al., 2017).

In summary, cryptocurrency pump-and-dump manipulations occur openly in cryptocurrency markets because these markets have little oversight by regulators and exchanges, high levels of investor speculation, and a growing pool of participants.

Appendix 2.D. Cumulative prospect theory

In our model, gamblers face a binary gamble, $\tilde{G}_q$. The value of this gamble to an individual is determined by Cumulative Prospect Theory (CPT) preferences, involving a value function and probability weighting (Tversky and Kahneman, 1992). As per CPT, the set of equations below give the value of a gamble to an individual $G \sim (x, P_x; y, P_y)$,

$$V(G) = w(P_x)v(x) + w\left(P_y\right)v(y)$$

$$w(P) = \frac{P^\delta}{(P^\delta + (1+P)^\delta)^{\frac{1}{\delta}}}$$

$$v(z) = \begin{cases} z^\alpha \\ -\lambda(-z)^\alpha \end{cases} \tag{2.D1}$$

where α, δ, and λ are the CPT parameters such that, $\alpha,\delta \in (0,1)$ and $\lambda > 1$. $v(z)$ is the value function that gives the perceived value of a payoff for an individual. The parameter α determines the degree of concavity over gains and convexity over losses, with higher values of α implying lower concavity over gains and convexity over losses. The parameter λ determines the degree of loss aversion; an individual with a higher λ is more loss averse. $w(P)$ is the probability weighting function that overestimates small probabilities and underestimates large probabilities. The parameter δ determines the degree of overestimation and underestimation. Like α, a lower value of δ implies a greater bias: higher overestimation of small probabilities and higher underestimation of large probabilities.

Appendix 2.E. Additional regression results

Table 2.A1
Tests of information-based manipulation theories

This table reports the results of logistic regressions testing whether theories of information-based manipulation explain cryptocurrency pump-and-dump schemes. These theories predict manipulation is more likely to happen when there is higher uncertainty about the fundamental value. The unit of observation is a coin-day, i, t. The dependent variable, $PumpDay_{i,t}$, is an indicator variable that equals one if there is a pump for the given coin on the given day and zero otherwise. The independent variable of interest is $Volatility_{i,t-1}$, which is the log of intraday volatility for the coin on day $t - 1$ (the difference between the highest and lowest trade prices scaled by the volume-weighted average price). For missing values of lagged volatility, we use the average volatility for the coin. $Liquidity_{i,t}$ is the log dollar trading volume for the coin on the day. $MarketCap_i$ is the log of the coin's average dollar market capitalization. $Yobit_i$ is an indicator variable that equals one if the coin is listed on the Yobit exchange. The sample consists of all coins listed on Binance and Yobit between December 2017 and June 2018. Chi-square statistics are in the parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $PumpDay_{i,t}$		
Variable	(1)	(2)	(3)
$Volatility_{i,t-1}$	–1.03***	–1.96***	0.26
	(–1.43)	(–4.71)	(0.42)
$Volume_{i,t}$	0.46***	0.98***	0.49***
	(14.91)	(16.55)	(13.75)
$MarketCap_i$	–0.40***		–0.42***
	(–9.60)		(–9.05)
$Yobit_i$	16.99***		10.05***
	(143.16)		(64.17)
Coin fixed-effects	No	Yes	No
Day fixed-effects	No	No	Yes
Pseudo-R^2	0.26%	1.05%	0.57%
Observations	117,644	174,542	117,644

Table 2.A2

Determinants of pump participation (adjusted overconfidence proxy)

This table reports regression results testing the determinants of participation in pump-and-dump manipulations. The dependent variable, $Participation_{j,i,t}$ is the log total trading volume during pump j in coin i on day t, measured from the release of the pump signal to three hours after the pump's peak price is reached. $Overconfidence_{j,i,t}$ is the adjusted overconfidence proxy. The original overconfidence proxy is the average percentage return in the two most recent pumps organized by the Telegram groups participating in the pump. We adjust this proxy by first regressing the contemporaneous group-level pump returns against one lagged value and using the residuals from this regression to calculate this proxy. $Gambling_t$ is the log daily dollar revenue of Bitcoin gambling services after removing the effects of cryptocurrency market activity. $Volatility_{i,t-1}$ is the log of intraday volatility for the coin on day $t-1$ (the difference between the highest and lowest trade prices scaled by the volume-weighted average price). For missing values of lagged volatility, we use the average volatility for the coin. $Manipulators_{j,i,t}$ is the log number of Telegram groups participating in the pump. $Members_{j,i,t}$ is the log number of members in these Telegram groups. $Yobit_i$ is an indicator variable that equals one if the coin is traded on the Yobit exchange. $Liquidity_i$ is the log average daily dollar trading volume of the coin. The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018. t-statistics are in parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Participation_{j,i,t}$	
Variable	(1)	(2)
$Overconfidence_{j,i,t}$	0.78**	0.56***
	(4.59)	(3.43)
$Gambling_t$		0.24
		(1.40)
$Manipulators_{j,i,t}$		–0.28***
		(–3.06)
$Members_{j,i,t}$		0.57***
		(5.07)
$Volatility_{i,t-1}$		0.10
		(0.88)
$Liquidity_i$	0.11	0.07
	(1.21)	(0.72)
$Yobit_i$	–6.67***	
	(–7.44)	
Week fixed effects	Yes	No
R^2	88.43%	25.18%
Observations	226	178

Table 2.A3
Dynamics of participation in pumps

This table reports regression results testing the predictors of participation in pump-and-dump manipulations in the time series. The dependent variable, $Participation_t$ is the log average total trading volume during pumps conducted on day t, measured from the release of the pump signal to three hours after the pump's peak price is reached. $Overconfidence_t$ is the average percentage return in the two most recent pumps organized by the Telegram groups participating in pumps conducted on day t. $Gambling_t$ is the log daily dollar revenue of Bitcoin gambling services after removing the effects of cryptocurrency market activity. $SearchInt_t$ is the log of the Google Search Volume Index for cryptocurrency related search terms on day t. $Yobit_t$ controls for exchange effects and is calculated as the number of pumps on the Yobit exchange on day t divided by the total number of pumps on the day t. $Liquidity_t$ is the log of the equal-weighted average daily dollar trading volume for all cryptocurrencies. The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018. t-statistics are in parentheses below the coefficient estimates. We include week fixed effects in Model 1 to absorb the longer horizon variation in pump returns. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Participation_t$		
Variable	(1)	(2)	(3)
$Overconfidence_t$	0.72**		
	(2.41)		
$Gambling_t$		0.32**	
		(2.18)	
$SearchInt_t$			1.25***
			(3.99)
$Liquidity_t$	0.69	1.42***	0.17
	(1.24)	(5.11)	(0.49)
$Yobit_t$	–8.12***	–6.43***	–6.79***
	(–20.19)	(–29.78)	(–33.86)
R^2	95.15%	89.25%	89.89%
Observations	119	133	132

Table 2.A4
Group membership and pump outcomes

This table reports regression results testing how pump group membership (a proxy for pump participation) affects manipulators' profit, pre-pump volume, pump return, and pump duration. The unit of observation is a pump j in coin i on day t. $ManipProfit_{j,i,t}$ is the log of the manipulators' dollar profit from the pump, calculated as the difference between the volume-weighted average price during the pump (from start to peak) and the volume-weighted average price in the two hours preceding the release of the pump signal multiplied by the trading volume in the two hours leading up to the pump. $PrePumpVol_{j,i,t}$ is the log dollar trading volume in the two hours preceding the release of the pump signal. $Return_{j,i}$ is the percentage return from the time of the pump signal to the peak of the pump. $Duration_{j,i,t}$ is the number of seconds from the release of the pump signal to the peak of the pump. $Members_{j,i}$ is a proxy for pump participation, calculated as the log of the number of members in the Telegram groups participating in the pump. $Liquidity_i$ is the log average daily dollar trading volume of the coin. The sample consists of 355 manipulations on two exchanges (Binance and Yobit) between December 2017 and June 2018. t-statistics are in parentheses below the coefficient estimates. Standard errors are clustered by coin and day. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
Variable	$ManipProfit_{j,i,t}$	$PrePump_{j,i,t}$	$Return_{j,i,t}$	$Duration_{j,i,t}$
$Members_{j,i,t}$	0.68***	0.32**	0.15***	–0.26***
	(5.29)	(2.34)	(3.57)	(–3.12)
$Liquidity_i$	–0.02	0.08	0.01	0.09
	(–0.19)	(0.71)	(0.28)	(1.10)
R^2	11.34%	3.89%	3.60%	3.75%
Observations	277	247	291	291

Table 2.A5
Determinants of pump frequency

This table reports regression results examining the coin characteristics that affect pump frequency. The unit of observation is a coin i. The dependent variable, $Pumps_i$, is the number of pumps for the coin. $MarketCap_i$ is the log of the coin's average dollar market capitalization. $TradesQ1_i$, $TradesQ2_i$, and $TradesQ3_i$ are indicator variables that equal one if the coin is in the first (lowest), second, or third quartile by average daily number of trades. $VolumeQ1_i$, $VolumeQ2_i$, and $VolumeQ3_i$ are indicator variables that equal one if the coin is in the first (lowest), second, or third quartile by average daily dollar trading volume. The sample includes all coins on the Binance and Yobit exchanges from December 2017 to June 2018. t-statistics are in parentheses below the coefficient estimates. Standard errors are clustered by coin. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Pumps_i$		
Variable	(1)	(2)	(3)
$Intercept$	1.17***	0.20***	0.19***
	(6.36)	(4.44)	(4.37)
$MarketCap_i$	–0.05***		
	(–5.45)		
$TradesQ1_i$		–0.11**	
		(–2.12)	
$TradesQ2_i$		0.31***	
		(4.10)	
$TradesQ3_i$		0.09	
		(1.55)	
$VolumeQ1_i$			–0.15***
			(–3.12)
$VolumeQ2_i$			0.34***
			(4.52)
$VolumeQ3_i$			0.13**
			(2.01)
R^2	4.46%	3.46%	4.63%
Observations	609	1,307	1,307

Chapter 3: Trading frictions and investor behavior and performance

3.1. 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 or mobile phones, or by simply giving voice commands to virtual assistants, such as Alexa, making markets highly accessible and virtually frictionless.[45]

Recently, media outlets such as Bloomberg and Wall Street Journal have criticized trading platforms such as Robinhood for making trading "too easy" and including design features in their trading platforms that have the potential to encourage investors to trade frequently and steer them toward certain securities (Wursthorn and Choi, 2020; Egkolfopoulou et al., 2021). These design features include removing frictions, such as trading costs, high initial capital requirements, and the number of screens that a user needs to navigate to input an order, as well as prominently displaying information about the top performing stocks and cryptocurrencies on the trading screen.

While reduced frictions benefit some market participants, we conjecture that they might harm individual investors by increasing their impulsive, heuristic-driven trading and exacerbating their tendency to underperform the market portfolio (Barber and Odean, 2000). Studies of individual investors primarily attribute their underperformance to biases in decision-making (Barber and Odean, 2013). The psychology literature contends that individuals making decisions rely on two types of cognitive processing: intuitive processing, which is fast-paced and automatic, and analytical processing, which is slow-paced and deliberative. Intuitive cognitive processes are more prone to systematic errors, biases, and reliance on heuristics (Kahneman, 2011). Lack of

[45] Sources: Charles Schwab and TD Ameritrade (https://bit.ly/3oeQhGn).

deliberative thinking leads to rash decisions in financial markets (Kocher et al., 2018). Therefore, trading frictions that increase deliberative thinking may 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 each 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.

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 in 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 the time between orders as a measure of cognitive effort in trading decisions. Participants who spend more time between orders are likely to be spending this time thinking about the next order. We control 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.[46] We adjust for the mechanical delays caused by reading the question in the $TASK$ treatment by removing all observations in which participants encounter the question for the first time.[47] We find that, in both experiments, participants spend more time between orders in the $SLOW$ and $TASK$ treatments as compared to the frictionless market. Participants spend between 17% and 47% more time per order in the $SLOW$ and $TASK$ treatments, after adjusting for

[46] This adjustment only applies to the single asset experiment.

[47] 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.

mechanical effects. The additional time spent per order is between 6% and 17% of the duration of a trading period in the multiple assets experiment and between 3% and 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, whereas 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 benefits at least one set of participants. In the $TASK$ treatment, participants in the bottom-quartile for performance underperform by approximately 57% (40%) less than the $NOFRICTION$ treatment in our multiple assets (single asset) experiment.[48]

Our third key finding is that the benefits of the cognitive task primarily result from 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, the performance improvement in the $TASK$ treatment is almost entirely explained by better decision-making rather than less overtrading. We further decompose the losses caused by bad decision-making into a fundamental and a non-fundamental component. In the multiple assets experiment, the $TASK$ treatment helps underperformers more closely match the Bayesian optimal strategy, i.e., helps them lose less money because of the fundamental component of bad decision-making. In the single asset experiment, the $TASK$ treatment helps underperformers lose less money due to price speculation, i.e., non-fundamental component, without affecting the losses caused by 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 reduce the amount the worst performers lose due to unsuccessful price speculation 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).

[48] In the single asset experiment, underperformers could be risk minimizers who trade-off performance for low risk exposure. The $TASK$ treatment can leave them worse-off by not allowing them to swiftly transfer risk. We can, however, rule out this possibility since, in our data, we do not find that underperformers seek to actively maintain low risk portfolios (see Section 3.4.2).

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.[49] Previous research shows that investors have limited attention, which 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 et al., 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. Investors also tend to expect recent price movements to extrapolate in the future (Greenwood and Shleifer, 2014; Barberis et al., 2015). The salience of prices in our experiments can exacerbate this tendency.

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 et al., 2020). To answer the question in the $TASK$ treatment, participants need to use 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. In the $SLOW$ treatment, although participants are thinking harder, this tendency is not corrected fully. This is possibly because the participants are not thinking in the "right way."

Evidence for the attention-inducing property of the $TASK$ treatment lies in the fact that in both the $SLOW$ and $TASK$ treatments, participants increase deliberation before trading, but only in the $TASK$ treatment can some participants 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 (in particular, information about asset fundamentals). The mechanism through which both treatments affect 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 make inaccurate estimates

[49] Assuming a linear-log relationship between performance and attention, the most inattentive investors are the worst performers. Gargano and Rossi (2018) document such a relationship for portfolio returns and investor attention.

of the assets' fundamental values, thus leading them to deviate from the optimal strategy. The $TASK$ treatment (but not the $SLOW$ treatment) helps them reduce this tendency by increasing deliberation on fundamentals, as this treatment involves asking participants about their beliefs about fundamental values.[50] 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 the asset at a low price when it 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 asset fundamentals can help them make better trading decisions.

Our findings are important, as the barriers to joining 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.[51] The number of individual investors using online trading to access financial markets is ever-increasing. Retail trading today is almost entirely conducted online. Of the brokers we mentioned in the beginning, 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 clearly explained to participants. In our setting, there are a maximum of three information sources (market prices, the order book, and private information in the form of a clue), among which only two sources (prices and the order book) are dynamic. In real markets, investors

[50] To correctly answer the question asked in the $TASK$ treatment, participants need to consider the entire price history of each asset, instead of only focusing on recent prices.
[51] Source: Robinhood (https://bit.ly/2FTe8dw).

have to sift through numerous information sources that 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 on individual investor behavior in financial markets. Barber and Odean (2002) find that investors' performance degrades 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, though only for 401(k) accounts (retirement savings accounts). They document no difference in performance across the two environments. Both of these studies compare trading in settings that implicitly have different degrees of frictions, although other factors also change, such as the framing of information about the market and the available information (Kalda et al., 2021). 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 forms of 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 can overcome the self-selection effects that might influence the results of these field studies. Our renewed analysis of frictions is crucial, as the observations reported in Barber and Odean (2002) might not hold today because 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; today, it offers commission-free trading.[52]

3.2. Hypotheses development

3.2.1. Retail investor underperformance

The behavioral finance literature finds that retail investors consistently underperform the market portfolio 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 other tendencies. Indeed, investors display biases such as overconfidence and other examples of limits to cognition in the experimental asset markets we study, specifically the Plott and Sunder (1988) market (Biais et al., 2005; Pouget,

[52] Sources: CNNMoney (https://cnn.it/31xJUUT) and Charles Schwab.

2007). Barber et al. (2009) provide the most comprehensive breakdown of underperformance sources 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 losses individuals make by trading.

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 it is higher. 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. The second strategy that traders can use 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 parameters might underperform systematically. In our experiments, we select two asset market designs that collectively capture all these sources of underperformance.

3.2.2. Psychological effects of obstacles

According to the psychology literature, 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).[53] Additional cognitive effort has been shown to mitigate biases caused by over-reliance 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 the decision. However, this reduction in biases is not guaranteed; it is possible that additional time spent on a task or even

[53] Intuitive and deliberative processing are not mutually exclusive processes that are invoked sequentially. Rather, these processes are invoked simultaneously. Recent studies indicate the presence of a control process that regulates various cognitive processes (including intuitive and deliberative processes) and arrives at a response that best fits the context (Venkatraman and Wittenbraker, 2020).

additional cognitive effort still yields the same biased decision (Nursimulu and Bossaerts, 2014; Enke et al., 2020; Venkatraman and Wittenbraker, 2020).

Further evidence from the psychology literature shows that obstacles that people encounter while making decisions can sometimes 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 required 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 appropriately 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 on phone calls than when trading online (Barber and Odean, 2002). Trading on the phone is more expensive, slower, and more effortful than online trading. Investors appear to fare better in a trading environment riddled with frictions than in a relatively frictionless trading environment. Along similar lines, Heimer and Imas (2021) find that a trading constraint, reduced access to leverage, improves investor performance and reduces trading biases by making it more difficult to avoid the psychological cost of realizing losses.

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 decisions, perhaps also helping them 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 they only trade when their trades are expected to be sufficiently profitable 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.[54] If a friction such as high transaction

[54] Thus far, we mainly build our hypotheses around the cognitive effort exerted while making trading decisions, and do not discuss specific biases such as overconfidence. However, traders employing higher cognitive effort are generally less likely to display such biases. For example, Hoppe and Kusterer (2011) find that traders with high cognitive ability, i.e., those that can engage deliberative processing more easily, are less overconfident.

costs induces more cognitive effort, it can reduce investors' overconfidence and cause them to reduce their trading activity.

The second friction is time delays before orders. This friction can help investors take a step back from the fast-paced trading environment and spend more time carefully thinking about a trading decision.

Lastly, the third friction 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 this information in their trading decisions, in case they are not doing so already. Due to this property, it can be argued that our cognitive task is not a friction that is likely to naturally arise in financial markets, but rather a deliberate "nudge" intended to influence trader behavior by inducing additional cognitive effort prior to a trade (Thaler and Sunstein, 2008). However, certain real-world frictions have a similar effect of inducing cognitive effort prior to trades. For example, while trading telephonically, an investor might discuss their order with the broker before finalizing it. This discussion can induce additional cognitive effort, particularly if the broker presents new information or informs the investor about 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. Such questions are frictions inasmuch as they impede swift order submission. These kinds of questions before orders exist on online trading platforms as well. For example, Figure 3.1 shows the questions asked before orders on the trading platform provided by SelfWealth, an Australian retail broker. These questions ask investors about their motivation and rationale to buy or sell the stock and their beliefs about the stock's future price trajectory. Answering these questions might require additional cognitive effort, as in our cognitive task.[55] Hence, our cognitive task can be considered a friction since it is quite similar in nature to real-world frictions that have a cognitive-effort-inducing effect. However, given the potential nudging effects associated with this friction, it can

[55] One might argue that these questions can easily be ignored. However, so can the questions in our cognitive task treatment. Questions such as these (including the telephonic broker and SelfWealth examples and our cognitive task) are likely to affect cognitive effort at the margins, i.e., for investors who exert little cognitive effort prior to encountering the question.

be used as a starting point to design direct nudges targeted at retail investors if we do find that it can help improve trader performance.[56]

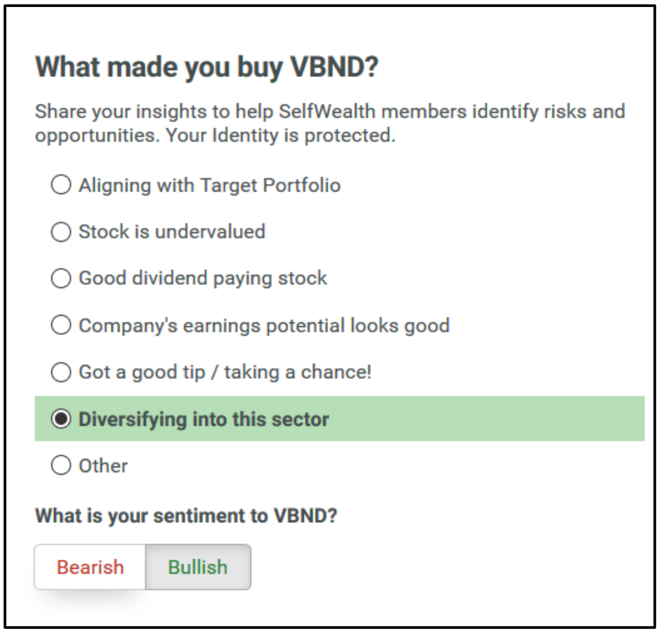

Figure 3.1. Screenshot of questions asked by SelfWealth before order submission.
This figure shows questions that investors trading on Australian broker SelfWealth's platform are asked before placing an order (market or limit). These questions appear on the order confirmation screen. Investors view these questions along with their order details and must press a 'Confirm order' button on this screen before their order is finalized.

Corgnet, Desantis, and Porter (2018) find that traders with high cognitive ability, i.e., traders who can better engage deliberative processing, perform better in financial markets. Individuals with high cognitive ability are less prone to behavioral biases (Oechssler, Roider, and Schmitz, 2009; Toplak, West, and Stanovich, 2011). Hence, the marginal benefit of any additional cognitive effort and reduction in biases caused by our frictions is likely to be the 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.

[56] One must, however, be careful when going down this path given that nudges might at times backfire. Osman et al. (2020) discuss the failures associated with nudges and offer guidance for practitioners to avoid such failures and backfiring effects.

3.3. Experiment design

3.3.1. Multiple assets experiment

Our multiple assets experiment uses a modified version of the market proposed by 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 in cash (laboratory currency) and four units of each of the four assets at the start of the market.

Asset prices are not determined endogenously through participant trading. Rather, the prices of all assets start at $180 and follow a stochastic process. The prices of all assets move by $15 after each 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 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 that 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 informed about the price increase probability corresponding to each asset. They need to infer this information from each asset's price path, which is displayed graphically on their screens during all trading periods. In addition to the eight trading periods, we add six non-trading periods before Period 1.[57] These non-trading periods help in participants' 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 prevailing market prices after the price movement in the final period, and 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.

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

3.3.2. Single asset experiment

Our single asset experiment uses an asset market design that is a modified version of Market 9 proposed by 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 in cash (laboratory currency) and four units of the asset. The asset has a cash payout at the end of each market. Each market has an independent payout. 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%, that of the $240 payout is 45%, and that of the $490 payout is 20%.[58] 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 of the traders are told one of the two 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 are 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 this market setting, it can be argued that trading frictions do not matter if the market is in equilibrium. Assuming rational expectations, in equilibrium, information is fully impounded into prices; there are either no trades (in the presence of transaction costs) or all trades occur at the equilibrium price (true asset value), and there is no heterogeneity in trader performance (Milgrom and Stokey, 1982; Biais and Pouget, 2000). However, modern replications of Plott and Sunder (1988) markets, e.g., Corgnet et al. (2019), find that, in contrast to the rational expectations theory, information is not fully revealed in these markets, suggesting that either the market is not in equilibrium or the agents do not have rational expectations. Additionally, various studies (including ours) document that agents continue to trade at out-of-equilibrium prices and face heterogeneous performance outcomes (Biais et al., 2005; Corgnet et al., 2018). This evidence suggests that frictions are likely to matter in this market setting.

[58] We follow Corgnet et al. (2018) in the modification of the payout probabilities.

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

3.3.3. Treatments

Both our experiments have four treatments: $NOFRICTION$, $HIGHCOST$, $SLOW$, and $TASK$. 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 that in the $NOFRICTION$ treatment in the single asset experiment is 2% of the transaction value.[59]

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 to 10% of the transaction value in the single asset experiment.

The $SLOW$ treatment contains the time delay friction. Here, 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.

Finally, 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).[60] 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 3.2 for the exact 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 the 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 3.2 for the exact question). The 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.

---

[59] We implement different types of transaction costs in both experiments to study both fixed and variable costs.
[60] 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.

Panel A: Multiple assets experiment

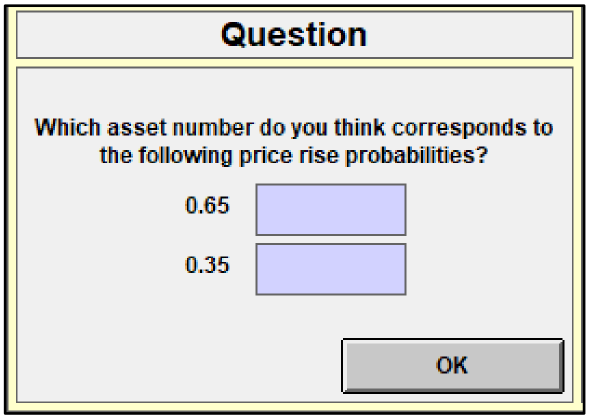

Panel B: Single asset experiment

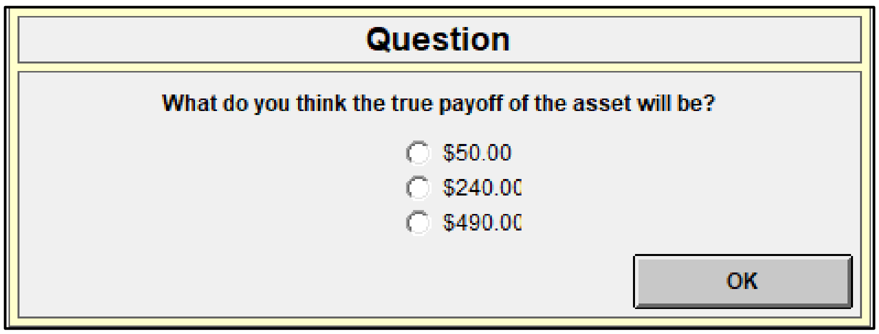

Figure 3.2. Screenshots of the cognitive effort task questions in the *TASK* treatment.
This figure shows the incentivized cognitive effort tasks that participants performed before placing an order in the multiple assets and single asset experiments. Panel A shows the question asked in the multiple assets experiment. Participants were required to answer this question before they were allowed to place 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 each order.

3.3.4. Additional experiment-related details

We conduct both experiments at 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. The 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 appropriate number of sessions required to generate a sufficiently large sample size for our statistical tests.[61] In total, 95 participants participate in the multiple assets experiment and 47 participants participate in the single asset experiment.[62] After each session, participants are ranked on the basis of their total earnings in the session. They receive a cash reward between AUD 25 and AUD 60 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 of which 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. Each participant participates in four markets. We generate 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 all, each experimental session has three markets for each treatment. We generate 564 subject-market observations in the single asset experiment, 141 observations per treatment.

Table 3.1 Panel A summarizes the structure of the asset markets, and Table 3.1 Panel B summarizes the treatments. Participant instructions and screenshots of participant trading screens for the multiple assets experiment (single asset experiment) are reproduced in Appendix 3.C (Appendix 3.D).

3.4. Experimental results

3.4.1. Cognitive effort

We start by examining whether our frictions help our participants think harder before making trading decisions. For this purpose, we examine the trade frequency and the time between orders across treatments. The time in a market or trading period is limited, and during this time, the participants are not interrupted or asked to complete any other tasks (except the cognitive task in

[61] We run more experimental sessions for the multiple assets experiment since this experiment generates fewer observations per session.
[62] Both experiments have independent sets of participants, i.e., no participant participated in both experiments.

the $TASK$ treatment), implying that their entire focus is on trading. Consequently, participants placing fewer trades and spending more time between orders are likely to be thinking more about each order than other participants.

Table 3.1
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 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		
$HIGHCOST$ treatment	Four times higher transaction cost	Five times higher transaction cost
$SLOW$ treatment	Additional 1 min. per trading period	20-second waiting period post order
$TASK$ treatment	Question related to asset value	Question related to asset value

3.4.1.1. Trade frequency

We first examine the 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 et al. (1998), Kocher et al. (2018) highlight that in the context of a financial market, trader passiveness can result in either 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 on deliberative cognitive processing. In a trading context, if passive traders associate participation in financial markets with frequent trading, their automatic systems might push them to trade frequently. Their trades would be prone to biases due to a lack of deliberation. In contrast, 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 3.2 reports the number of trades across treatments in both experiments. This table displays the mean values for subject-market observations. In the multiple assets experiment, participants make an average of 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 the 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 and 3.55 trades in the $SLOW$ and $TASK$ treatments, respectively.

The observed reduction in trading activity 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 participants 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 (Table 3.2).[63] 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 in 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 induce traders to exert additional cognitive effort and engage additional deliberative cognitive processing before trades.

3.4.1.2. Time between orders

Next, we examine the treatment effects on the time between orders. Although our trade frequency results indicate that our $SLOW$ and $TASK$ treatments induce participants to think harder before trades, these results could be misleading if the trades are clustered together in time and/or mainly occur early in the market or trading period. In this case, we cannot conclusively state that a reduced number of trades implies increased cognitive effort per trade. Hence, it is useful to

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

examine the time between orders as well, to determine exactly how much time the participants spend 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 the time between orders. We control for this mechanical effect by using an adjusted time between orders metric for the $NOFRICTION$ treatment. 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 $NOFRICTION$ treatment 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 for the waiting period and one second for order submission).[64]

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, we exclude the first trading period (market) in the $TASK$ treatment in the multiple assets (single asset) experiment. In both cases, we exclude the first instance in which participants encounter the question. Reading and processing 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 or process the question again.

Table 3.2 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% ($t = 4.73$) in the $SLOW$ treatment. The increased time spent per order is slightly more than 10% of the total extra time (1 additional minute) that the participants

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

are given per trading period in the $SLOW$ treatment. Panel B in Figure 3.3 plots the number of orders across time in the $NOFRICTION$ and $SLOW$ treatments. Although the ordering activity in both treatments is initially similar, in the $SLOW$ treatment, it dissipates gradually. Interestingly, participants place orders even after the normal trading period (1 minute). This indicates that they utilize the extra time given, which allows them to think harder about their orders.

The time between orders increases to 29.45 seconds in the $TASK$ treatment, an increase of approximately 9.36 seconds or 46.6% ($t = 6.87$). The graph of the number of orders across time in the $NOFRICTION$ and $TASK$ treatments (Panel C in Figure 3.3) clearly indicates that participants spend more time thinking about orders in the $TASK$ treatment. Ordering activity in the $TASK$ treatment increases at a much slower rate than in 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 also responding to the cognitive task), and placing orders only in the second half of the trading period. In contrast, in the $NOFRICTION$ treatment, participants jump straight to placing orders, with ordering activity peaking in the first half of the trading period itself. Participants clearly spend less time on processing the new information and deciding their trades in the $NOFRICTION$ treatment than in the $TASK$ treatment.

The absolute value of the additional time between orders might not seem high; the increases are only between 7 and 10 seconds. However, even 7 seconds is valuable in this experiment, as it is around 12% of the total time available to participants in a period to make a trading decision. Hence, although 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; an increase of 5.37 seconds or 17.09% ($t = 2.60$). Panel B in Figure 3.4 plots the number of orders in the $NOFRICTION$ and $SLOW$ treatments in the single asset experiment across time. Like the multiple assets experiment, ordering activity in both treatments is initially similar, but in the $SLOW$ treatment, it gradually decreases. In the $TASK$ treatment, the average time between orders increases by 32.64% ($t = 3.26$) from its value of 25.13 seconds in the $NOFRICTION$ treatment. This result is visualized in Panel C in Figure 3.4, which plots the number of orders in the

$NOFRICTION$ and $TASK$ treatments across time. Ordering activity in the $TASK$ treatment is lower than that in 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 and 10 seconds). However, these times are between 17% and 33% higher than the time spent on an individual order in the $NOFRICTION$ treatment. The increases are also around 3–5% of the total market duration. Hence, the extra time spent thinking about an order is meaningfully large even in this experiment.

Table 3.2
Trading behavior

This table reports descriptive statistics for various trading behavior related metrics 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 trading period is calculated as the time of the last order divided by the number of orders. Period-level values are averaged at the market level. In the single asset experiment, the same calculation is performed directly at the market level. 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. 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	$NOFRICTION$	$HIGHCOST$	$SLOW$	$TASK$
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)	NA	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)

Panel A: *NOFRICTION* and *HIGHCOST* treatments

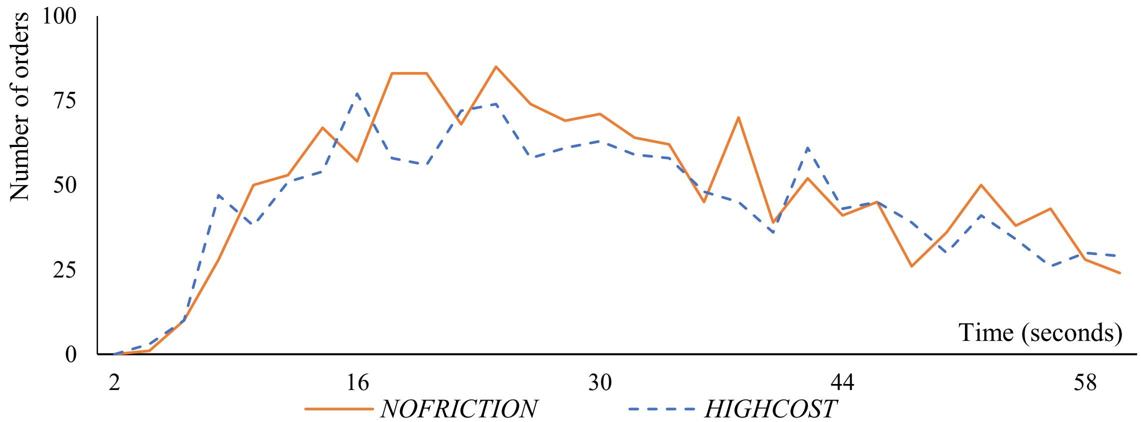

Panel B: *NOFRICTION* and *SLOW* treatments

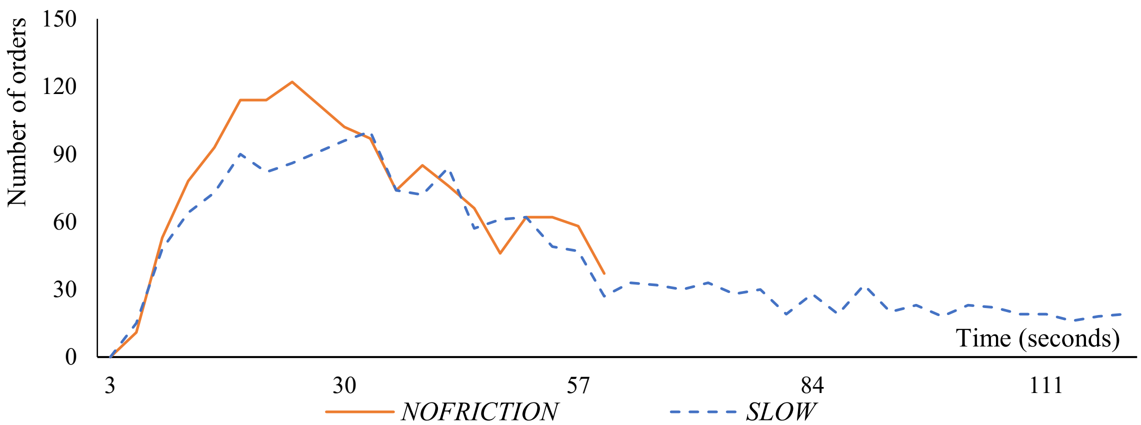

Panel C: *NOFRICTION* and *TASK* treatments

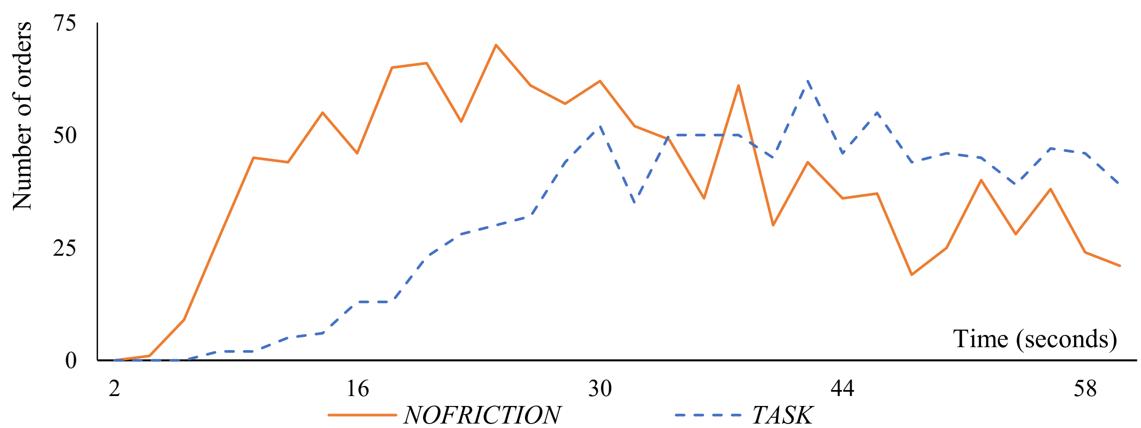

Figure 3.3. Order submission activity through time in the multiple assets experiment.
This figure plots the number of orders submitted by participants per period in the multiple assets experiment through time. Panel A plots the number of orders in the $NOFRICTION$ and $HIGHCOST$ treatments in two-second 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

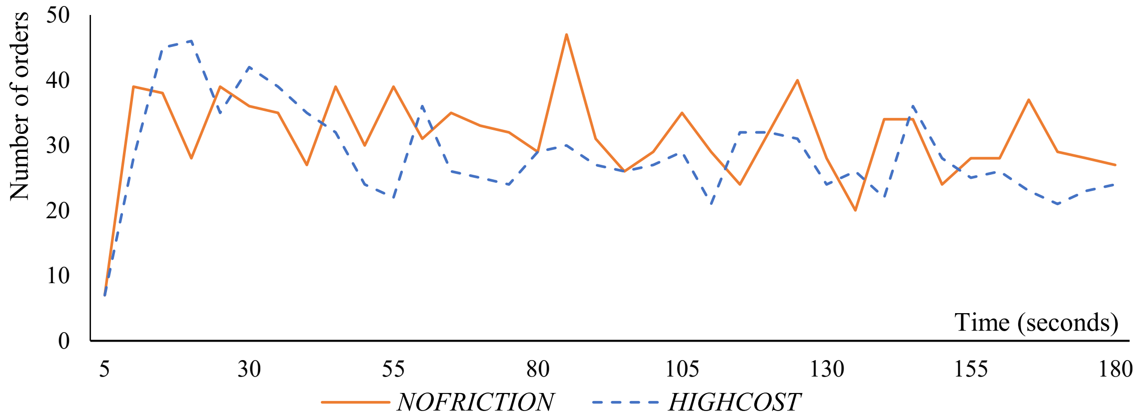

Panel B: *NOFRICTION* and *SLOW* treatments

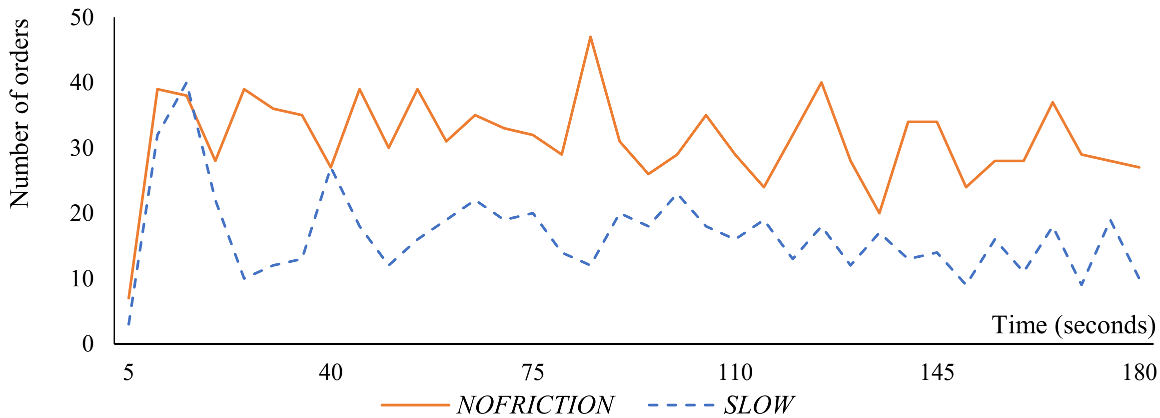

Panel C: *NOFRICTION* and *TASK* treatments

Figure 3.4. 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.

The preponderance of evidence suggests that participants consider their trading decisions more deliberatively in both the $SLOW$ and $TASK$ treatments but not in the $HIGHCOST$ treatment. Participants spend between 17% and 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 in the single asset experiment.

3.4.2. Trading performance

Now, we examine the treatment effects on investor performance. Numerous previous studies have documented that individual investors underperform the market.[65] This is true for our experiments too. In the multiple assets experiment, the average investor underperforms the Bayesian optimal strategy by 7.86%, or \$388.28, in terms of the final earnings in a market.[66] The most underperforming participants (bottom-quartile for performance) underperform the Bayesian optimal strategy by approximately \$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, this number is approximately \$64.22, or 3.54% of the participants' total portfolio value in a market. In addition, the most underperforming participants earn an average \$432.8 or 24% less than their peers in the $NOFRICTION$ treatment.

Participants might trade for risk reasons, i.e., participants who prefer not to hold risky assets transfer these assets to participants with a higher risk appetite, rather than to maximize profit. In this sense, the underperformers we identify could be risk minimizers who trade-off performance for low risk exposure. However, we can rule out this possibility, as we find that underperformers do not actively seek to maintain lower risk portfolios than the other participants. The standard deviation of earnings for underperformers does not differ from that of the other participants in a statistically significant manner (difference = –\$8.91, $t = -0.14$) when all treatments are combined. Additionally, if we look at the $NOFRICTION$ treatment individually, we find that underperformers have a higher standard deviation of earnings than the other participants

[65] See Barber and Odean (2013) for a comprehensive review on individual investor trading behavior.
[66] Appendix 3.B provides details about the Bayesian optimal strategy.

(difference = \$173.3, $t = 2.36$), indicating that, if anything, underperformers take more risk than other participants.

We begin by examining the overall effects of our treatments on participant performance in the multiple assets experiment. Here, we use participants' earnings as the measure of performance, after subtracting from it the earnings earned under the Bayesian optimal strategy.[67] Earnings from a market are calculated as the sum of the cash balance and the total asset portfolio value at the end of the market.

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.[68] We include fixed effects for market sequence.[69]

Table 3.3 Model 1 reports the performance effects of our treatments. Contrary to expectations, the $HIGHCOST$ treatment reduces participant performance compared to the $NOFRICTION$ treatment. The $HIGHCOST$ treatment reduces participant earnings by a magnitude of \$650.91, or by 16.5% of the mean earnings level in the $NOFRICTION$ treatment. Participant earnings also reduce in the $SLOW$ treatment; earnings drop by \$260.61, or 6.61% of the mean earnings level in the $NOFRICTION$ treatment. Although the $TASK$ treatment has a positive coefficient, it is not statistically significant, implying that the average participant's performance is not significantly different in the $TASK$ treatment from that in 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 explore possible individual-level heterogeneity in the treatment effects on performance. Our frictions aim to induce additional cognitive effort; the marginal benefit of this added cognitive effort is likely to be the highest for participants who exert the least cognitive effort, i.e., the worst performing participants. We divide participants into quartiles based on their total performance across all treatments, and test whether our treatments benefit the worst performing participants

[67] Performing this subtraction merely removes some of the variance in earnings caused by different asset value realizations.
[68] All our regressions (in the current and subsequent sections) are OLS regressions, and all our tests use subject-market observations.
[69] 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.

more than others.[70,71] 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 in the previous set of tests, only adding UP_i and interaction terms between the treatment variables and UP_i as additional regressors.

Table 3.3
Earnings in the multiple assets experiment

This table reports regression results testing how the treatments generally affect earnings of all participants and specifically affect earnings of underperformers in the multiple assets experiment. 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***	–608.36***
	(–8.05)	(–7.82)
$SLOW_{j,k}$	–260.61***	–261.55***
	(–3.22)	(–3.43)
$TASK_{j,k}$	115.86	–1.30
	(1.43)	(–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

[70] For the multiple assets experiment, we generate underperformance quartiles at the experiment level, i.e., we compare the underperforming participants' performance with all other participants in the experiment. In this experiment, we are able to compare underperformance across experimental sessions since the fundamentals of the game remain the same across sessions and we compare participants' underperformance relative to the optimal strategy.

[71] Our results for both experiments are robust to dividing participants into quartiles based on their underperformance in only the $NOFRICTION$ treatment.

Table 3.3 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 than in 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 the other participants, from \$938 in the $NOFRICTION$ treatment to \$254 in the $TASK$ treatment.

Table 3.4
Earnings in the single asset experiment

This table reports regression results testing how the treatments generally affect earnings of all participants and specifically affect earnings of underperformers in the single asset experiment. 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 or not \$490, and two indicator variables for the payout being \$50 or \$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.

	Dependent variable = $Earnings_{i,j,k}$	
Variable	(1)	(2)
$HIGHCOST_{j,k}$	–134.14***	–94.90*
	(–2.62)	(–1.75)
$SLOW_{j,k}$	16.62	10.06
	(0.32)	(0.18)
$TASK_{j,k}$	17.62	–28.32
	(0.35)	(–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

We now examine the effects of our treatments on participant performance in the single asset experiment. Here, we use participants' earnings from a market as a performance measure. 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 participant performance on the set of treatment indicators and include experimental session fixed effects and market sequence fixed effects.[72,73] We control for the participant's clue by including the variables $ClueNot50_{i,j,k}$ (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 the variables $Payout50_{j,k}$ (equals one when the asset payout is \$50) and $Payout490_{j,k}$ (equals one when the asset payout is \$490). For brevity, we do not report the coefficients for our control variables.[74]

Table 3.4 Model 1 reports the performance effects of our treatments. The only treatment that has a statistically significant effect on earnings is $HIGHCOST$. As in the multiple assets experiment, the $HIGHCOST$ treatment lowers the subject-market earnings compared to the $NOFRICTION$ treatment; the subject-market earnings are \$134.14 lower in the $HIGHCOST$ friction, or approximately 8% less than the mean earnings level in the $NOFRICTION$ treatment. In contrast, the average subject-market earnings in the $SLOW$ and $TASK$ treatments do not differ from those in the $NOFRICTION$ treatment in a statistically significant manner. The $HIGHCOST$ treatment reduces participant earnings on average, whereas the $SLOW$ and $TASK$ treatments have no statistically significant effect on earnings when pooling all participants into a single group.

Table 3.4 Model 2 reports the results for the underperformance quartiles.[75] The $TASK$ treatment benefits underperformers more than other participants, as indicated by the positive and

[72] 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.

[73] 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 Table 3.A1 in Appendix 3.A. All our results are robust to controlling for non-linear learning effects.

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

[75] For the single asset experiment, we generate underperformance quartiles at the experimental session level, i.e., we only compare the underperforming participants' performance with other participants in their session. We do so because participant performance in this experiment is sensitive to the asset payout distributions and price paths, which are

statistically significant coefficient of the interaction term between $TASK_{j,k}$ and UP_i.[76] 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 from \$432.8 in the $NOFRICTION$ treatment to \$193.3.

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

3.4.3. Overtrading and bad decision-making

In this subsection, 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 caused by overpaying trading costs and is calculated as the additional transaction costs that 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. Our measure for bad decision-making is calculated as the difference between the benchmark strategy's gross earnings, i.e., final earnings with transaction costs added back, and the participant's gross earnings.

---

different in each session. Additionally, in this experiment, each market has the potential to be a fundamentally different game since markets can either involve price speculation or fundamental trading, or both.

[76] It can be argued that as this market is effectively a zero-sum game (absent trading costs), this performance benefit for underperformers is a mechanical effect caused by all participants trading less in the $TASK$ treatment and thus less money being transferred from underperformers to other participants. However, even in the $SLOW$ treatment, all participants trade less by a similar magnitude to the $TASK$ treatment (see Section 3.4.1). Despite similar levels of reduced trading, there is no real performance benefit for underperformers in the $SLOW$ treatment. Consequently, it is likely that this performance benefit is not mechanical and is driven by reasons other than a reduction in trading.

We start by examining the treatment effects on sources of underperformance for the worst performers in the multiple assets experiment. As in the previous subsection, we define the indicator variable UP_i to represent underperformers. In these tests, we use the same set of independent variables as in the previous subsection, i.e., treatment dummies, the underperformer dummy, and 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.

Table 3.5
Overtrading and bad decision-making in the multiple assets experiment

This table reports regression results testing the treatment effects on underperformance due to overtrading and bad decision-making for underperformers in the multiple assets experiment. 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 difference between the gross earnings (with transaction costs added back) in 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***	430.64***
	(15.24)	(5.52)
$SLOW_{j,k}$	2.70	258.86***
	(0.24)	(3.39)
$TASK_{j,k}$	–21.61*	20.32
	(–1.85)	(0.26)
UP_i	1.18	787.44***
	(0.07)	(7.00)
$HIGHCOST_{j,k} \times UP_i$	21.18	143.52
	(0.85)	(0.86)
$SLOW_{j,k} \times UP_i$	6.68	–12.29
	(0.29)	(–0.08)
$TASK_{j,k} \times UP_i$	16.11	–470.47***
	(0.65)	(–2.83)
R^2	60.40%	46.95%
Fixed effects	Sequence	Sequence
Observations	380	380

Table 3.5 reports the results for the multiple assets experiment. The incremental reduction in underperformance for underperformers in the $TASK$ treatment observed 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 underperformers witness a significantly higher reduction in losses caused by bad decision-making than 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 is the primary driver of their improved performance.

Next, we perform the same set of tests for the single asset experiment, except that here, we use a different benchmark strategy. In this experiment, we measure participant underperformance relative to a no-trade strategy, i.e., we compare participant performance to a hypothetical scenario wherein they do not trade at all.

Table 3.6 reports the regression results for the single asset experiment. We confirm that, as in 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 less due to overtrading and \$170 less due to bad decision-making in the $TASK$ treatment as compared to the $NOFRICTION$ treatment. Although these values are not statistically significant, they are still economically meaningful, as we observe a statistically significant relation in the underperformance tests in Section 3.4.2. In the current set of tests, we are only interested in investigating which component of underperformance explains the performance improvement observed by the underperformers in the $TASK$ treatment.

The results suggest that the $TASK$ treatment helps underperformers improve their decision-making quality more than it helps others. This improvement in decision-making, and not less overtrading, explains the reduction in underperformance experienced by underperformers in the $TASK$ treatment.

Table 3.6
Overtrading and bad decision-making in the single asset experiment

This table reports regression results testing the treatment effects on underperformance due to overtrading and bad decision-making for underperformers in the single asset experiment. 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 in 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 or not \$490, and two indicator variables for the payout being \$50 or \$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.

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

3.4.4. Components of bad decision-making

In this subsection, we analyze the treatment effects on the components of underperformance due to bad decision-making. Here, we 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 the 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 assets sold by the optimal strategy). Additionally, even if the participant can 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 asset 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 caused by buying high and selling low during the market are included in the non-fundamental component. This component is more relevant in the single asset experiment than in the multiple assets experiment. This is because, unlike in 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 the bad decision-making components in the multiple assets experiment. We calculate the fundamental component of bad decision-making by using the participants' portfolios at the end of a market. We first subtract the endowed units from the participant's terminal holdings. This gives us the participant's net change in position for each asset. We then multiply this net change by the difference between the terminal price of the asset and the volume-weighted average price at which they bought the additional units or sold the existing units. 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 that 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 3.7 reports the regression results for the multiple assets experiment. The coefficient of UP_i is significantly higher in the regression with the fundamental component of bad decision-making than in that with the non-fundamental component, indicating that the fundamental component contributes more to the underperformance of the worst performers than the non-fundamental component. The $TASK$ treatment targets the fundamental component; the coefficient of the interaction term between $TASK_{j,k}$ and UP_i is significantly larger when regressed on the fundamental component than when regressed on 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 the reduction in their underperformance. In other words, the $TASK$ treatment helps participants guess the price increase probabilities corresponding to each asset more accurately and more quickly than in the $NOFRICTION$ treatment.

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. As in 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 participant's terminal holdings to arrive at their net change in position. We then multiply this net change by the difference between the volume-weighted average price at which they bought the additional or sold the existing units and the asset's payout.[77] This difference directly yields 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.

[77] This step is slightly different in the single asset experiment as compared to the multiple assets experiment. This is because, in the multiple assets experiment, we first compute the *earnings* due to the fundamental component for both the participant and the optimal strategy, and then take the difference between the two as the participant's underperformance, whereas here we directly compute the underperformance due to the fundamental component for the participant.

Table 3.7
Components of bad decision-making in the multiple assets experiment

This table reports regression results testing the treatment effects on components of underperformance due to bad decision-making for underperformers in the multiple assets experiment. 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 in 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 the additional units or sold the existing units. $Non\text{-}fundamental_{i,j,k}$ is the non-fundamental component of underperformance due to bad decision-making (in laboratory \$). We calculate the non-fundamental component by deducting the fundamental component from total 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.

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

Table 3.8
Components of bad decision-making in the single asset experiment

This table reports regression results testing the treatment effects on components of underperformance due to bad decision-making for underperformers in the single asset experiment. 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 the additional units or sold the existing units and the asset payout. $Non\text{-}fundamental_{i,j,k}$ is the non-fundamental component of underperformance due to bad decision-making (in laboratory \$). We calculate the non-fundamental component by deducting the fundamental component from total 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 or not \$490, and two indicator variables for the payout being \$50 or \$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.

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

Table 3.8 reports the regression results for the single asset experiment. Unlike in the multiple assets experiment, here, the non-fundamental component of bad decision-making explains the underperformance of the worst performers more than the fundamental component, as indicated by the higher coefficient of 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 in 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 of the interaction term between $TASK_{j,k}$ and UP_i. In effect, the $TASK$ treatment helps underperformers lose less money because of price speculation. 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 further investigate the effects of the $SLOW$ treatment on the fundamental component by breaking down the fundamental component into the buys and sells domains (see Table 3.A2 in Appendix 3.A). We find that, in the $SLOW$ treatment, underperformers tend 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.

3.5. Discussion

In both the $SLOW$ and $TASK$ treatments, participants are more deliberative when making trading decisions; however, 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 only one yields 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 about their beliefs regarding the 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 provide performance benefits for participants who might be underutilizing this information in the $NOFRICTION$ treatment. We conjecture that the worst underperformers who benefit from the $TASK$ treatment likely fit this bracket. These actors increase cognitive effort in the $SLOW$ treatment as well but do not receive any performance benefits. This is likely because, in this 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 the other treatments. However, even in the $TASK$ treatment, only the worst performers receive performance benefits. The worst performers might be inattentive and might 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 the increased deliberation on fundamentals in the $TASK$ treatment is the 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 types of information and relying too much on other types 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). This tendency is exacerbated when investors spend relatively less time in making a decision and rely more on intuitive processing (Liao et al., 2020). 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 large 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 Appendices 3.C and 3.D 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. According to the behavioral finance literature, 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, inattentive participants could be neglecting or underweighting the core value formation process for the asset(s), as 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 on the basis of 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 regarding 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. By 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 finalize their trade soon after answering the question. In this manner, the question in the $TASK$ treatment can help draw the attention of inattentive investors toward 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 3.4.4). This is likely because these investors can form beliefs that are less affected by recency bias because of the question asked 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 3.4.1).

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 because of price speculation (Section 3.4.4). This result and a simultaneous reduction in momentum-driven trades imply that both these treatments are effective in helping inattentive investors avoid recency bias, i.e., speculating on prices expecting recent price movements to extrapolate in the future. However, since the $SLOW$ treatment does not direct the attention of inattentive investors to the core value formation process, they fall into another trap. These investors end up selling assets with a high payout at relatively low prices in this treatment (Section 3.4.4). 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 toward 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 improvement to investors since it is the only one that stops inattentive investors from relying too much on recent price movements and 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. 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 can also change over time. Consequently, cognitive tasks meant to induce deliberation on fundamentals cannot be as straightforward as our experiments. Some examples of real-world interventions resembling the cognitive tasks could include asking investors about their beliefs about the fair valuation of the security and the valuation model they used to compute this valuation or asking investors whether they have seen a recent news item about the security (assuming one exists) before they can place an order. These interventions could have some positive effect since, even in real-world financial markets, the same investor tendencies of overweighting salient information and recent price movements apply. We speculate that some frictions that resemble our cognitive task can help investors appropriately use this information and improve their performance.

Appendix 3.A. Bayesian optimal strategy in the multiple assets experiment

The stochastic price process in the multiple assets experiment allows us to construct a Bayesian optimal strategy that can be used as a benchmark for participant performance. We construct this strategy from the frame of reference of a participant, implying that this strategy is constrained by the information available to a participant in a given period. This strategy involves performing a few calculations during each period in which there is new information in the form of a price movement. We begin by computing a belief about the probability of an asset having a certain price increase probability. In all, there are 16 belief probabilities; each asset has four belief probabilities, one for each price increase probability. All these belief probabilities are initialized to 25% since, at the start, each asset has a 25% chance of having a certain price increase probability. In each subsequent period, we generate a posterior belief by Bayesian updating the prior belief on the basis of the observed price movement. Suppose that we are updating our belief about the price increase probability for Asset 1 being 65%, and we observe a price increase in the previous period. Then, to generate our posterior belief ($P^1(0.65|increase)$), we update our prior belief ($P^1(0.65)$) as follows:

$$P^1(0.65|increase) = \frac{P(increase|0.65) \times P^1(0.65)}{P^1(increase)} \quad (3.B1)$$

where $P(increase|0.65) = 0.65$ and $P^1(increase) = P(increase|0.65) \times P^1(0.65) + P(increase|0.55) \times P^1(0.55) + P(increase|0.45) \times P^1(0.45) + P(increase|0.35) \times P^1(0.35)$.

Our posterior belief becomes our prior belief for the next period. We perform the above calculation in all periods, trading and non-trading. The succeeding calculations are only performed for trading periods. After generating the posterior belief, we estimate all possible terminal values for each asset given current period prices. Using our belief probabilities, we calculate the probability of each terminal value for each asset. This gives us the expected terminal value for each asset. We then compute all portfolio combinations that can be entered into, given the available cash and assets and considering the transaction costs. We rank all portfolio combinations based on their expected terminal value and select the highest ranking portfolio.[78] This portfolio becomes the optimal portfolio, and all computations in the next period are conditioned on this selected portfolio.

[78] In instances where two assets have the same high price in the initial periods, we select the highest ranking portfolio combination that allocates funds equally among these high priced assets.

In most instances, this strategy entails selling all low price increase probability assets and investing all funds in the highest price increase probability asset as soon as there is sufficient information available to reliably guess the highest probability asset.

Appendix 3.B. Multiple assets experiment instructions and participant screen

1. Welcome to this experiment. You all will earn some money, which will be paid to you in cash at the end of the experiment. Different people will earn different amounts. The amount you earn will depend on the decisions you and other people make during this experiment.

2. Please do not interact with other participants. In case you have any questions, please raise your hand and you will be assisted. Please do not move from your seat or use your mobile phone until the experiment finishes.

3. Your earnings from a market will be added to your total experiment earnings. At the end of the experiment, participants will be ranked based on their experiment earnings and paid an amount between \$25 and \$60.

Asset value

4. You will be trading 4 assets (Asset 1 to 4). A market will last for 8 trading periods (Periods 1 to 8). You will start off each market with a trading account of \$1,000 and 4 units of each of the 4 assets.

5. You can use your cash to purchase more units of any asset or sell units of assets you own for cash. Your cash and asset balance can never be negative. **The time remaining in a period is shown on the top-right corner of your screen.**

6. **Prices of all 4 assets move by \$15 at the end of every period.** Each asset has a **different probability** of experiencing a price increase in a given period. This probability is called the price increase probability. **Price increase probability** is **65% for 1 asset, 55% for 1 asset, 45% for 1 asset and 35% for 1 asset.** Since there are only 4 assets, no 2 assets have the same price increase probability. Also, since prices move in each period, the **probability of a price decrease is 1 minus the probability of a price increase**. For example, the asset which has a 65% price increase probability also has a 35% price decrease probability. Each asset has 1 price increase probability that stays constant during a market and doesn't change from period to period within the market.

7. You can trade at the current price in a period. You have to pay a transaction cost for each trade. This transaction cost is applied on each trade and doesn't depend on the number of units in the trade. Suppose you have \$500 and 2 units of Asset 2 in a period. Assume you buy 1 unit of

Asset 4 at the price of \$240 and sell 2 units of Asset 2 at the price of \$80. The transaction cost is \$5.

8. Your cash balance will decrease by \$240 due to your Asset 4 purchase. It will increase by $(\$80 \times 2) = \$160$ due to your Asset 2 sale. Your cash balance will also decrease by $(\$5 \times 2) = \$10$ to pay for the transaction cost of both the trades.

9. At the end of the period, your cash balance be $\$500 + \$160 - \$240 - \$10 = \$410$. Your Asset 2 balance will be $2 - 2 = 0$ units and your Asset 4 balance will be $0 + 1 = 1$ unit.

10. **At the end of a market, you will receive what your assets are worth, which will be determined by the asset's market price at that time.** The amount you receive from your assets will be added to your cash balance to determine your earnings from the market.

11. Assume you have 5 units of the Asset 1, 6 units of Asset 3, 1 unit of Asset 4 and \$350 at the end of a market. Say Asset 1 is worth \$150, Asset 3 is worth \$210 and Asset 4 is worth \$135 after the final period of the market. At the end of the market, your earnings from the market will be $\$350 + (5 \times \$150) + (6 \times \$210) + (1 \times \$135) = \$2{,}495$. These earnings will be added to your experiment earnings.

12. Experiment earnings account balances of all participants will be ranked at the end of the experiment. Your rank will determine how much cash you earn from the experiment.

Trading

13. During each period, you can see the price of all assets in the current period and all previous periods. This price history is visible in the charts on the left-side of the screen. You can see that prices of all assets start at \$180. You can also see price movements for each asset in 6

Asset ownership

Asset	Current price	Units you own	Total value
1	165.00	4	660.00
2	255.00	4	1020.00
3	135.00	4	540.00
4	135.00	4	540.00

periods before Period 1. You can also see all possible price increase probabilities in the "Useful information" box.

14. You can buy or sell the asset at the current period's price. You can see the current period price along with the number of units you hold of each asset and the total value of your asset holdings in the "Asset ownership" box.

15. To buy a given asset, go to the "Trade submission" box, select the asset you wish to buy, enter the number of units you wish to buy in the "Volume" box and press the "Buy" button. To sell a given asset, select the asset you wish to sell, enter the number of units you wish to sell in the "Volume" box and press the "Sell" button. Try putting in some trades now.

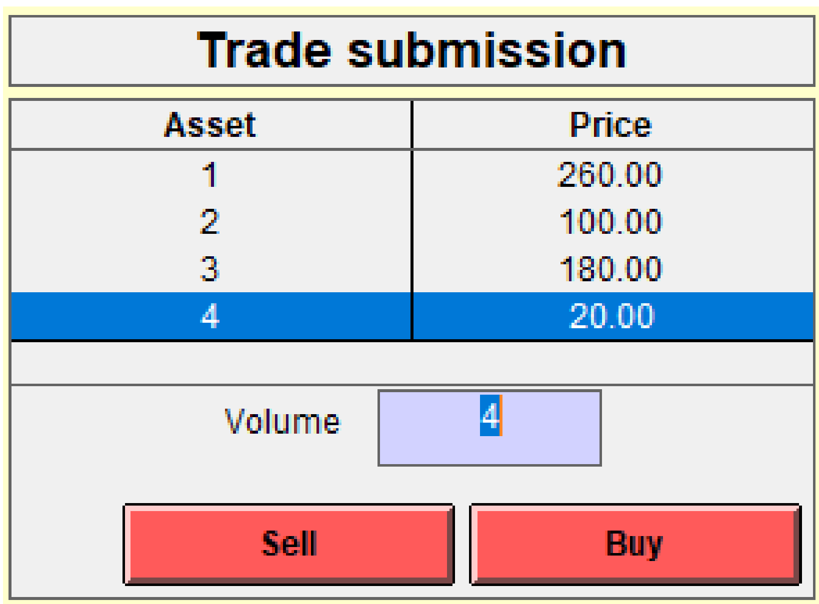

16. It is not compulsory to trade in every period and your trades will only be executed at the end of the period. Until then, you can see your trades in the current period in the "Current period trades" box. You can change your trades before the period finishes. To do so, you need to first cancel your existing trades. To cancel a trade, select the trade from the "Current period trades" box and then press the "Cancel trade" button just below the trade. After cancelling your existing trade, resubmit the new trade just like you submitted the previous trade. The "Current

period trades” also displays the total trade value (price multiplied by quantity) before and after transaction costs. Try cancelling and resubmitting a trade now.

Current period trades

Asset	Trans. type	Volume	Trade value	Trans. cost	Tot. cost/income
3	Buy	3	405.00	5.00	410.00
2	Sell	4	1020.00	5.00	1015.00
4	Buy	1	135.00	5.00	140.00

Cancel trade

17. Your cash balance and net spend during the period are visible in the “Current position” box. If you have no trades for the period or if you first put in a trade then later cancel it, your cash and asset balances won’t change.

Current position

Cash:	$1000.00
Net spend this period:	$-465.00

18. Your net spend during the period is the amount you are paying for your buy trades minus the amount you are receiving from your sell trades. This amount can be negative if your sell trades are worth more than your buy trades. If your net spend is negative, your cash balance will increase after the period by your net spend amount. If your net spend is positive, your cash balance will decrease by the net spend amount at the end of the period.

Practice market

19. Before the start of the experiment, there will be a practice market lasting for 2 periods. This practice market is meant to give you some idea about how to use the trading software.

20. We encourage you to put in some trades in the practice market to get a feel for the trading software. Your earnings from this market will not be added to your experiment earnings account.

Please feel free to ask any questions.

Multiple assets experiment participant screen

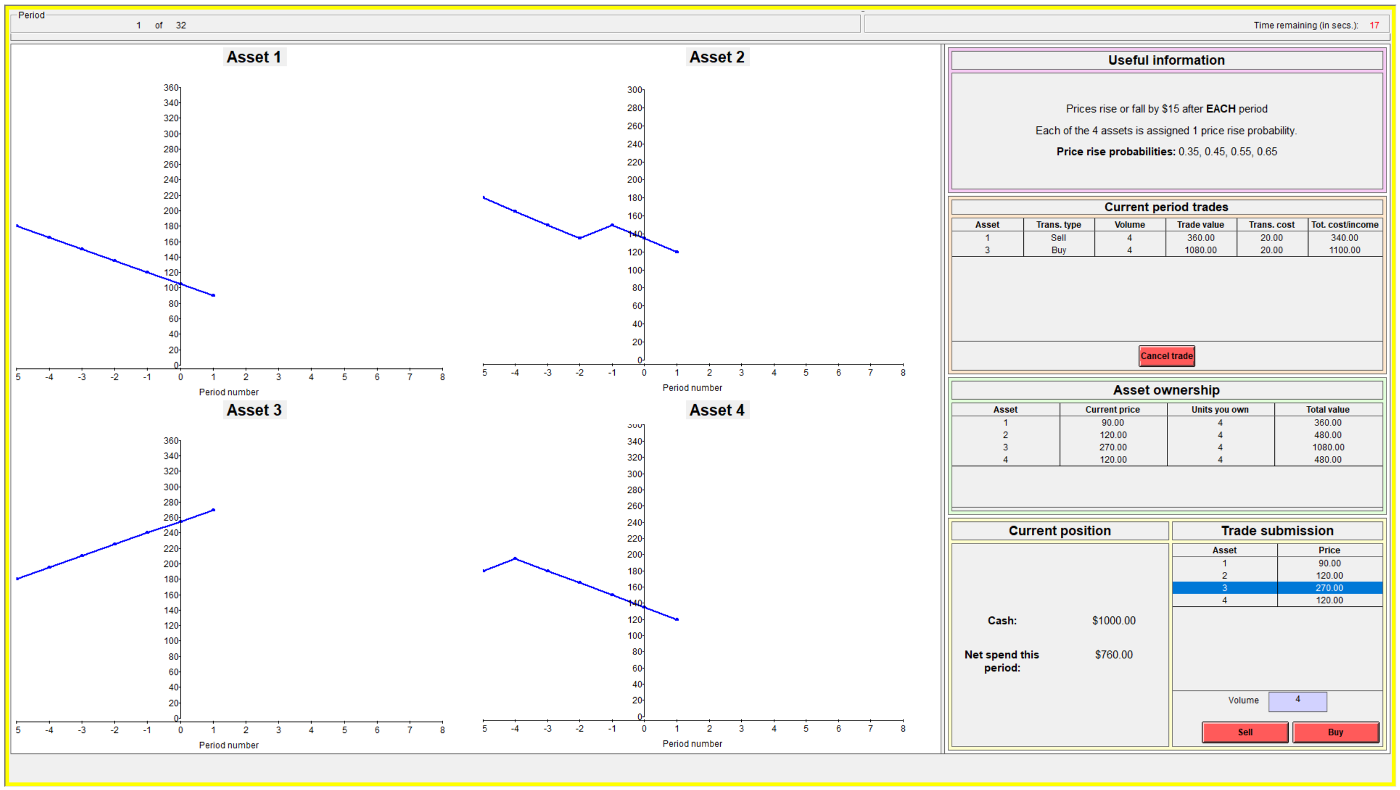

Appendix 3.C. Single asset experiment instructions and participant screen

1. Welcome to this experiment. You all will earn some money, which will be paid to you in cash at the end of the experiment. Different people will earn different amounts. The amount you earn will depend on the decisions you and other people make during this experiment.

2. Please do not interact with other participants. In case you have any questions, please raise your hand and you will be assisted. Please do not move from your seat or use your mobile phone until the experiment finishes.

3. You will be trading in a series of asset markets, with each market trading a different asset. Your earnings from each market will be added to your total experiment earnings. At the end of the experiment, participants will be ranked based on their experiment earnings and paid an amount between $25 and $60.

Trading

4. You can trade with other participants by submitting orders to buy or sell. Orders have 3 important elements – the price, the quantity (or volume), and the transaction cost. You can buy or sell at any time, but your cash and asset balance can never be negative.

5. If you only want to buy or sell for a particular price, you can place a buy or sell order in the "Order submission" box. In this box, first input the price you want to buy or sell at, then input

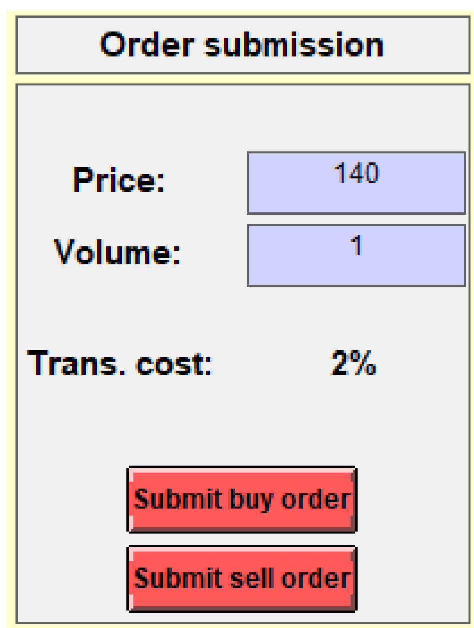

the number of units you want to buy or sell (the volume or quantity) and click “Submit buy order” or “Submit sell order”.

6. After you press “Submit buy order” or “Submit sell order”, you will get a summary of the order you are submitting in the “Order details” box. In this box you can see the total value of your order (i.e., price multiplied by volume) before and after transaction costs.

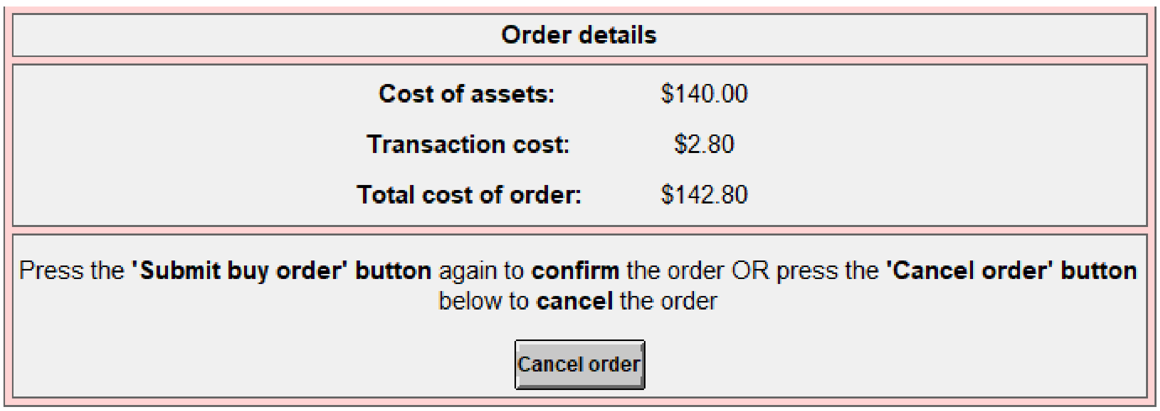

7. You can confirm the order by clicking the “Submit buy order” or “Submit sell order” button again. You can also cancel your order at this stage by clicking “Cancel order” in the “Order details” box.

8. An order submitted in this manner will not execute immediately. It will appear in the “Buy orders” or “Sell orders” box and will be visible to all participants until it is executed. The buy

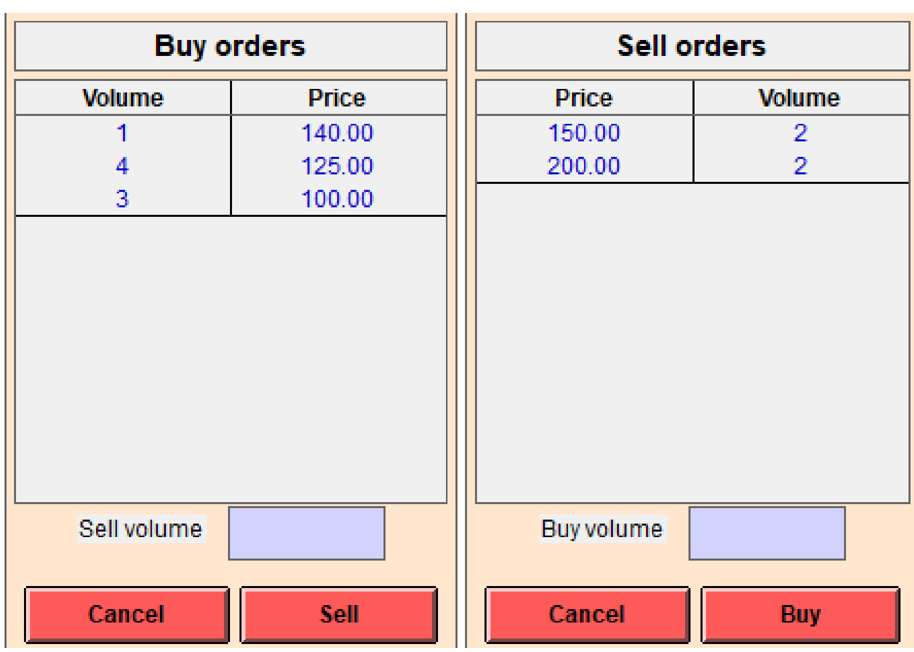

orders will appear in the “Buy orders” box and sell orders will appear in the “Sell orders” box. Try to submit a order with your desired price and quantity now.

9. You can also cancel any of your orders before they are executed by selecting your buy or sell order from the “Buy orders” or “Sell orders” box and pressing the “Cancel” button. Your own orders are in blue and others’ orders are in black. Try cancelling one of your orders now.

10. The pending orders you see in the “Buy orders” and “Sell orders” boxes only execute when someone decides to buy or sell immediately against a pending order. Buying or selling immediately is another way you can trade. If you choose to buy immediately, you won’t be able to decide the price you buy at. You will have to buy at the price of the cheapest pending sell order. Similarly, if you choose to sell immediately, you won’t be able to decide the price you sell at. You will have to sell at the price of the most expensive buy order.

11. To sell immediately, go to the “Buy orders” box, select the best buy order (the first one), input the number of units you want to sell in the box next to “Sell volume” and click the “Sell” button. To buy immediately, go to the “Sell orders” box, select the best sell order (the first one), input the number of units you want to buy in the box next to “Buy volume” and click the “Buy” button.

12. There are a few things to note when it comes to immediate buys and sells. You need to select the best order to immediately buy or sell. The first order in the box is always the best order. Additionally, if you are immediately buying or selling, the quantity you can buy or sell needs to be less than or equal to the quantity of the best order you are selecting. Lastly, you can only buy or sell from others. This means that you can only select other people’s orders (not your own) to buy or sell immediately. Try immediately buying or selling against someone else’s order now.

13. Although you don’t get a summary when you are immediately buying or selling like you do when you are submitting new orders, the details for immediate buys and sells are always visible directly below the box you need to use to immediately buy or sell. You can use these details to check the transaction cost you would need to pay if you immediately buy or sell. These details use the quantity of the best available order. If you want to check the details for a quantity lower

than that of the best available order, input your quantity next to "Custom qty." and press the "Check details" button.

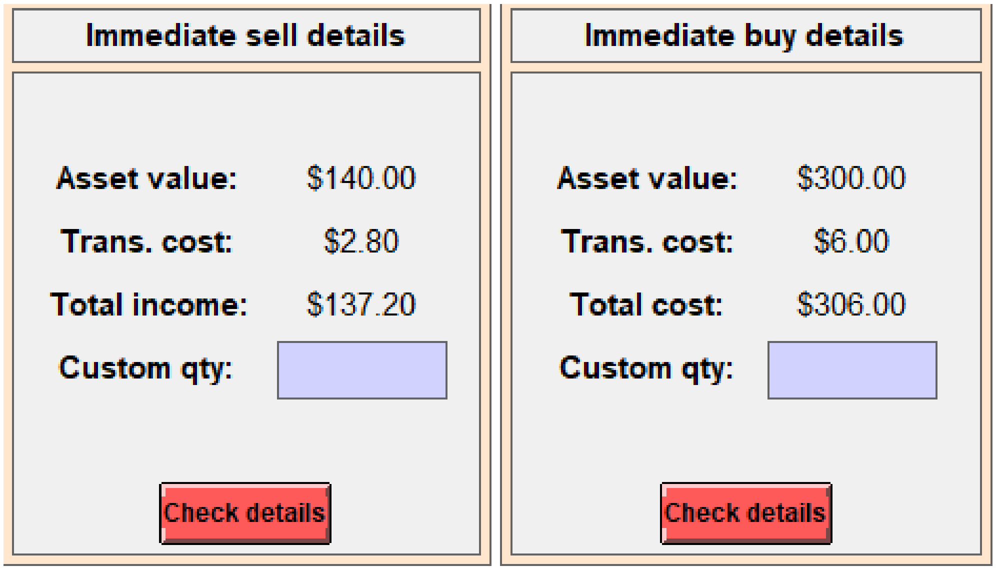

14. Remember that if you only want to buy or sell at a particular price and are happy to wait until order is executed, you need to submit an order. If you don't want to wait and are happy to pay the best price you can get in the market, you need to select someone else's order and buy or sell immediately.

15. Your cash balance will increase as soon as any of your pending sell orders execute or you immediately sell. It will decrease as soon as any of your pending buy orders execute or you immediately buy. Similarly, your asset balance will decrease as soon as any of your pending sell orders execute or you immediately sell and increase as soon as any of your pending buy orders execute or you immediately buy.

16. Each trade has a transaction cost that both the buyer and the seller have to pay. This transaction cost is deducted from your cash balance as soon as your pending order executes or you buy or sell immediately.

17. The transaction cost is a percentage of your total order value. Suppose you are selling 5 units of the asset at a price of \$100 and the transaction cost is 2%. Now, the transaction cost you pay will be $2\% \times \$100 \times 5 = \$10$. The cash you receive from this sale will be $(5 \times \$100) - \$10 = \$490$.

18. You can see the transaction cost and the total cost/income for your order either in the "Order details" box before order submission and in the "Immediate buy details" and "Immediate sell details" boxes.

19. Your cash and asset balances are listed in the "Current position" box. This box also displays the cash and assets that are available to you for submitting new orders. Cash and assets tied up in pending orders are not available for new orders or for buying/selling immediately.

Asset value

20. You start each market with a trading account of $1,000 and 4 units of the asset. You can use the cash to purchase additional units of the asset from others or your can sell the units you own to others for cash. **Each market lasts for 3 minutes. The time remaining in a market is shown on the top-right corner of your screen.**

21. At the end of EACH market, EACH unit of the asset you hold will automatically payout a certain amount. The payout in each market can be different. **The payout can either be $50, $240 or $490. The chance of the payout being $50 is 35%. The chance of the payout being $240 is 45%. The chance of the payout being $490 is 20%.**

22. Let us consider an example. Assume you start a market with $1,000 and 4 units of the asset. Suppose you buy 5 more units of the asset in this market at a price of $100. Assuming the transaction cost is 2%, this would you cost you $(5 \times \$100) + (5 \times \$100 \times 2\%) = \$510$. This buy trade would reduce your cash balance to $\$1{,}000 - \$510 = \$490$ and increase your asset balance to $4 + 5 = 9$ units.

23. Now suppose the payout in this market is $240. After this market is over, each of your 9 units of the asset will pay out $240. This payout will increase your cash balance by $\$240 \times 9 = \$2{,}160$. Your cash balance after the payout will be $\$490 + \$2{,}160 = \$2{,}650$. This $2,650 amount will be your earnings from this market. These earnings will be added to your experiment earnings account.

24. Your trading account will be reset to $1,000 and 4 units in the next market. Experiment earnings account balances of all participants will be ranked at the end of the final market. Your rank will determine how much cash you receive at the end of the experiment.

Clue

25. In each market, all of you will receive a clue about the asset payout in the market. Different people might receive different clues. As explained before, the asset will pay out 1 of 3 values at the end of the market. Your clue will be 1 value that the asset DOES NOT pay out. This clue can be seen in the "Useful information" box. Keep an eye out for your clue in EVERY market.

26. Suppose you get a clue that the asset does not pay out $50 in a market. This means that the asset can only pay out $240 or $490. Keep this information in mind since the payout is the only cash you will receive if you are holding the asset at the end of the market. The clue you get is chosen at random.

Practice rounds

27. Prior to the beginning of the experiment, there will be 2 practice markets. These markets are meant to give you some idea about how to use the trading software. Although you can trade in these markets (we encourage you to do so) and the asset will have a payout at the end of the market as well, your earnings will not be added to your experiment earnings account.

Please feel free to ask any questions.

Single asset experiment participant

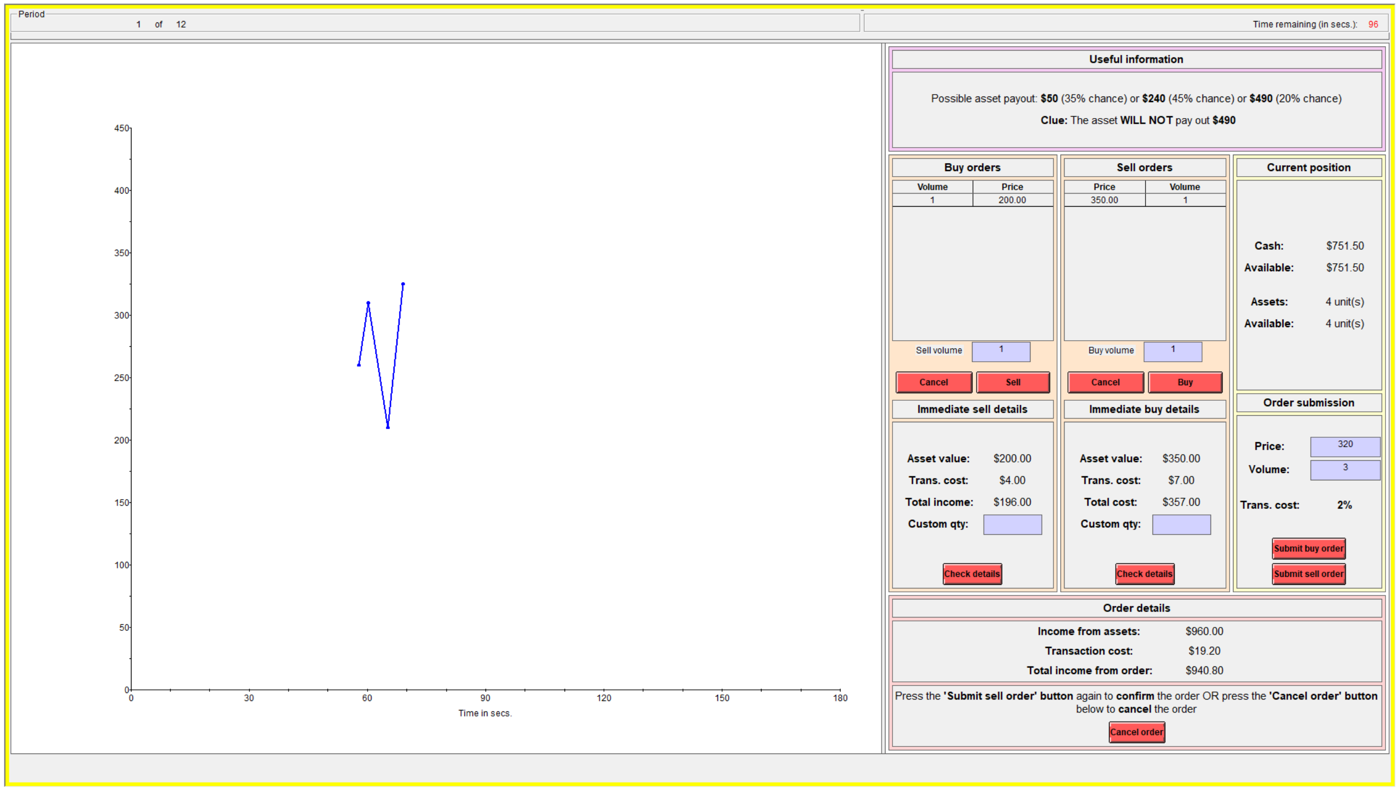

Appendix 3.D. Additional regression results

Table 3.A1
Earnings in the single asset experiment with non-linear learning controls

This table reports regression results testing how the treatments affect earnings of all participants and specifically underperformers in the single asset experiment while controlling for non-linear learning effects. 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 or not $490, and two indicator variables for the payout being $50 or $490. We also control for non-linear learning effects using the market sequence, squared market sequence, and cubed market sequence. Regressions reported in this table also control for session fixed effects. t-statistics are in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable = $Earnings_{i,j,k}$	
Variable	(1)	(2)
$HIGHCOST_{j,k}$	-134.12**	-117.29**
	(-2.57)	(-2.04)
$SLOW_{j,k}$	14.29	6.98
	(0.25)	(0.11)
$TASK_{j,k}$	14.28	-35.58
	(0.24)	(-0.55)
UP_i		-309.25***
		(-3.96)
$HIGHCOST_{j,k} \times UP_i$		-65.43
		(-0.59)
$SLOW_{j,k} \times UP_i$		29.52
		(0.27)
$TASK_{j,k} \times UP_i$		196.50*
		(1.78)
R^2	70.24%	72.87%
Controls	Signal, Payout, Sequence, Sequence^2, Sequence^3	Signal, Payout, Sequence, Sequence^2, Sequence^3
Fixed effects	Session	Session
Observations	564	564

Table 3.A2
Fundamental components of bad decision-making in the single asset experiment

This table reports regression results testing the treatment effects on fundamental components of underperformance due to bad decision-making for underperformers in single asset experiment. The unit of observation is a subject i in market j of session k. $Buys_{i,j,k}$ is the fundamental component of underperformance due to bad decision-making (in laboratory $) in the buys domain. To calculate $Buys_{i,j,k}$, we first subtract the endowed units from the number of units of the asset in the participant's terminal portfolio. If the number of units in the terminal portfolio is less than the endowed units, this term is set as zero. We then multiply this term by the difference between the volume-weighted average price at which the participant bought the additional units and the asset payout. $Sells_{i,j,k}$ is the fundamental component of underperformance due to bad decision-making (in laboratory $) in the sells domain. To calculate $Sells_{i,j,k}$, we first subtract the endowed units from the number of units of the asset in the participant's terminal portfolio. If the number of units in the terminal portfolio is greater than the endowed units, this term is set as zero. We then multiply this term by the difference between the volume-weighted average price at which the participant sold the existing units and the asset payout. $HIGHCOST_{j,k}$ 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 or not $490, and two indicator variables for the payout being $50 or $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.

	Dependent variable =	
Variable	$Buys_{i,j,k}$	$Sells_{i,j,k}$
$HIGHCOST_{j,k}$	-28.40	31.88
	(-0.97)	(1.07)
$SLOW_{j,k}$	-39.18	10.02
	(-1.32)	(0.33)
$TASK_{j,k}$	-5.36	11.25
	(-0.18)	(0.38)
UP_i	50.18	62.45
	(1.23)	(1.50)
$HIGHCOST_{j,k} \times UP_i$	8.00	3.53
	(0.14)	(0.06)
$SLOW_{j,k} \times UP_i$	43.22	114.34*
	(0.75)	(1.95)
$TASK_{j,k} \times UP_i$	5.83	-4.54
	(0.10)	(-0.08)
R^2	40.05%	37.84%
Controls	Signal, Payout	Signal, Payout
Fixed effects	Session, Sequence	Session, Sequence
Observations	564	564

Chapter 4: Attention to information releases and price reactions to announcements

4.1. Introduction

It is a well-documented fact that investors are not fully attentive to new information in financial markets.[79] Inattention can cause them to react to information disclosures with a lag. For example, on days when investors pay less attention to markets, prices underreact to earnings announcements, leading to prolonged return drifts following the announcement (DellaVigna and Pollet, 2009). Although the literature documents strong attention effects after earnings announcements, it is less clear whether similar effects exist for other types of disclosures as well. Additionally, the literature does not identify *what* attracts the investors' attention. This is important since attention to different items could have differing effects. While attention to information can help improve market efficiency by speeding up the incorporation of new information into prices, attention to returns or trading volumes might have the opposite effect of contributing to market inefficiency if it fuels price phenomena such as momentum.

Capturing cross-sectional variation in attention to information has proven to be an empirical challenge. The existing literature proposes various proxies that approach the problem of measuring investor attention from different angles. For example, many studies use stock-level proxies, such as search activity for a stock on Google, US Securities and Exchange Commission's EDGAR platform and Bloomberg Terminal, and Wikipedia page views to measure information acquisition after announcements (Da, Engelberg, and Gao, 2011; Drake, Roulstone, and Thornock, 2015; Ben-Rephael, Da, and Israelsen, 2017; Focke, Ruenzi, and Ungeheuer, 2020). Other studies use attention-grabbing events, such as media coverage and extreme returns, as exogenous shocks to attention (Barber and Odean, 2008; Peress, 2016). Finally, some studies use trading activity as a proxy for investor attention (Gervais, Kaniel, and Mingelgrin, 2001; Hou, Xiong, and Peng, 2010). Many of these measures either confound attention to information releases with attention to returns or are noisy measures of total investor attention, implying that they are not fit for purpose when

[79] See Barber, Lin, and Odean (2019) and Gabaix (2019) for thorough literature reviews on the role of inattention in finance.

measuring attention exclusively to information releases.[80] The distinction between attention to information and attention to returns or volumes is crucial given the evidence documented in Barber and Odean (2008) of extreme returns and volumes attracting significant investor attention. Separating these two aspects of attention is important since they could have differing effects on market efficiency in terms of information being impounded into prices. Intuitively, since attention to information directly captures the attention investors pay to an information release, it is more likely to matter than attention to returns with respect to the incorporation of information into prices. As this chapter studies price reactions to announcements, we primarily focus on the former.

We propose a novel measure for attention that focuses specifically on information releases. The intuition behind this measure is that investors who are attentive to releases of material information should rationally cancel or update their pending orders (if any) after the information release as they update their private valuations of the company. Because of its material nature, this new information alters the fair valuation of the stock, thus necessitating an update to an investor's pending order, which would now be based on a stale valuation. Investors who miss these information releases would not update their pending orders. Consequently, investors who cancel or update their orders after such information releases can be classified as attentive, whereas those who do not can be classified as inattentive. Our attention measure is the average proportion of overnight order cancellations and amendments for a stock on days with material disclosures or material news events in the overnight period, i.e., after trading hours.[81] We construct our measure in overnight periods since, during these periods, new information can be released by companies and investors can cancel or amend their pending orders, but there are no changes to prices, as there are no trades. This allows us to disentangle attention to information releases from attention to returns.

---

[80] The existing attention measures confound attention to returns with attention to information for various reasons. Historical data for Google search volumes are available at a daily frequency at best. On trading days, extracting attention to information releases from total attention is not feasible using daily data. Media coverage as a proxy for attention also faces a similar problem, as media coverage is endogenously related to returns, with stocks that earn extremely high returns receiving disproportionately high media attention. The proxies that are better at segregating attention to returns from attention to information releases (Wikipedia page views, EDGAR searches, and Bloomberg views) are noisy measures for attention to information. Wikipedia page views are a noisy proxy for information acquisition since new announcements rarely appear on company Wikipedia pages and investors do not necessarily need to view these pages to access the latest information. Additionally, EDGAR searches and Bloomberg views are measures of institutional attention rather than retail attention, and hence do not capture total investor attention well.

[81] Our results are robust to the use of only overnight periods with price sensitive disclosures to construct our measure.

Our measure offers at least three additional advantages as a proxy for attention to information over the most popular attention proxy, the Google Search Volume Index (SVI).[82] First, investors do not need to search for stocks on Google to obtain information. Brokers are increasingly providing integrated trading platforms wherein participants can view company announcements and news updates in real-time along with other market-related information for a stock. Out of the 16 retail online brokers registered on the Australian Securities Exchange (ASX), at least 12 brokers provide company announcements on a real-time basis, and at least 14 brokers provide the latest stock- and market-related news on their platforms.[83] Second, individuals can search for stocks on Google for other, non-information-related reasons. For example, Focke et al. (2020) find that advertising increases Google search activity for a firm's stock in the short run. The authors also find that Google searches have a weak correlation with trading activity, implying that Google search activity might not always translate into trades. Hence, in addition to returns, Google searches can also be confounded by other factors, such as advertising. Third, Google search activity does not discriminate between investors paying attention to company information releases or attention to prices or volumes. Thus, it is uninformative about *what* catches the investors' attention.[84]

We use our attention measure to test how stock-level attention affects price reactions to material (also known as "price sensitive") announcements for stocks listed on ASX. We find that the market generally underreacts initially to new information, with return drifts occurring at horizons ranging from the first day after the announcement to the subsequent 19 days. These drifts increase with inattention and are stronger for low attention stocks than for other stocks. Drifts are also stronger for announcements that contain positive news than for those that contain negative news. For positive announcements, drifts persist for approximately 40 days after the announcement. For negative announcements, drifts dissipate after the first-day post announcement.

---

[82] SVI is a relative score for the search frequency of a given search term. The relative score is obtained by scaling the number of Google searches for a search term in a given day by its time-series average in the selected period. The score ranges from zero to one hundred, with a value of zero indicating no Google searches and a value of one hundred indicating the maximum number of Google searches in the selected period. After its introduction in Da et al. (2011), SVI has been used as an investor attention proxy in various studies, such as Drake, Roulstone, and Thornock (2012); Ben-Rephael et al. (2017); Hirshleifer and Sheng (2019); Chen et al. (2020); Hansen (2021); and Liu, Peng, and Tang (2020), among others.

[83] The full list of brokers registered on ASX is available at this link (https://bit.ly/3i04g2l).

[84] A further limitation is that Google search volume is not available in many small stocks or at a high frequency.

To analyze the economic significance of these underreaction patterns that are amplified by inattention, we construct a long-short trading strategy that aims to exploit the post-announcement drifts. The strategy purchases low attention stocks after a positive announcement and short sells low attention stocks after a negative announcement. To ensure the strategy uses information available *ex-ante*, we only use the stock-level attention distribution available at the time of the trading decision. We classify news as positive or negative on the basis of the observed return at the time. We compare the performance of this attention-conditioned strategy to that of a baseline news trading strategy that purchases stocks after positive announcements and sells short stocks after negative announcements to isolate the benefits of measuring attention.

We find that while the attention-conditioned trading strategy fails to outperform the baseline strategy at the one-day horizon, it does reliably outperform at the ten-day and 20-day horizons.[85] For example, the attention-conditioned strategy outperforms the baseline strategy by around 20% at the ten-day horizon, earning a statistically significant daily excess return of 0.46%. At the 20-day horizon, the attention-conditioned strategy outperforms the baseline strategy by approximately 95% and earns a 0.41% statistically significant daily alpha. The substantially improved performance of our strategy at longer horizons adds credence to the claim that inattention has an economically significant effect of slowing down the incorporation of information into prices.

We investigate how the price drift patterns and the role of attention vary with announcement- and stock-level factors. First, we examine whether the drifts only occur for certain types of announcements. We perform our main tests for scheduled and unscheduled announcements separately.[86] Due to their unexpected nature, unscheduled announcements could be mispriced initially, as investors take time to process the information. We find that low attention drifts occur after both scheduled and unscheduled announcements, suggesting that inattention is a pervasive phenomenon, not restricted merely to unscheduled information.

Second, we examine earnings and non-earnings announcements separately. Given the existing evidence on the post-earnings announcement drift (PEAD), it could be that earnings releases drive the results. However, we do not find this to be the case. Stronger drifts for low attention stocks are

---

[85] It is important to note here that, in the 20-day case, the strategies do not trade in the entire 20-day period. Rather, the first ten-day period is spent observing the announcement return to classify announcements as positive or negative, and the strategies only trade in the next ten-day period.

[86] Scheduled announcements include announcements that are periodic in nature, e.g., earnings reports, while unscheduled announcements are not periodic, e.g., mergers and acquisitions.

prevalent after non-earnings announcements as well. In fact, non-earnings announcements drive the stronger first-day drifts after negative announcements for low attention stocks, suggesting that, if anything, drifts are more prevalent after non-earnings announcements than after earnings announcements.

Finally, given the previous evidence of higher mispricing of stocks favored by retail investors (Bali et al., 2019), we also examine whether our results are driven by the mispricing effects of these stocks rather than low attention. We classify lottery-like stocks, i.e., stocks with low prices, high volatility, and high skewness, and small capitalization stocks as stocks favored by retail investors. We find that the effects of attention on post-announcement drifts exist independent of the mispricing effects of these retail-friendly stocks. Overall, based on the evidence, attention-related inefficiencies, as captured by return drifts, are pervasive and not constrained to a particular segment of the announcements or stocks.

This chapter contributes to the literature studying the effects of investor attention on price reactions to corporate announcements. Numerous studies investigate the effects of attention on post-earnings announcement drifts in some manner (Chen, Jiang, and Zhu, 2018; Huang, Nekrasov, and Teoh, 2018; Hirshleifer and Sheng, 2019; Kottimukkalur, 2019; Li, Nekrasov, and Teoh, 2020; Liu et al., 2020; Hansen, 2021). Some studies find that these drifts are stronger when investor attention is low (Hirshleifer et al., 2009; Pantzalis and Ucar, 2014; Lawrence et al., 2016; Ben-Rephael et al., 2017; Jiang, Li, and Wang, 2021). However, as we pointed out earlier, the attention measures used by these studies confound attention to information releases with attention to returns or volumes. Disentangling the two is essential given the potential for returns and volumes to attract significant investor attention (Barber and Odean, 2008). Intuition dictates that inattention to information should matter more than inattention to returns or volumes in predicting pricing inefficiencies post announcement releases. Hence, we develop a new attention measure that specifically captures investor attention to information releases. In line with previous findings, we document that price drifts after announcements are stronger when investors pay less attention to information. We also document a novel result that drifts for stocks that receive low investor attention are stronger for announcements that contain positive news.

The Australian setting of our study offers a unique advantage over the previous studies that are set in the US market. In Australia, firms must disclose all relevant information on a real-time basis. This continuous disclosure of all, and not strategically selected, firm-specific information makes it an ideal setting to study the effects of investor attention on the incorporation of information into stock prices.

Aside from examining the effects of attention to information on post-announcement returns, this chapter makes a methodological contribution by introducing a new measure of investor attention. Researchers who are interested in specifically examining attention to information can use the measure, as it offers cleaner identification of attention to information than other measures proposed in the literature.

This chapter helps bridge the behavioral asset pricing and market microstructure literature. Activity at the limit order level has been studied extensively in the microstructure literature. Studies examine order placements, revisions, executions, cancellations, non-executions, monitoring, and other order-level activity (Handa and Schwartz, 1996; Liu, 2009; Fong and Liu, 2010; Yamamoto, 2014; Khomyn and Putnins, 2021). This order-level activity ultimately affects market efficiency through executed or forgone trades and, thus, also mediates the relation between investor attention and pricing efficiency since the only mechanism through which information is impounded into prices is through trading by placing orders. Despite this, to the best of our knowledge, research has not yet examined the link between order cancellations and price reactions to announcements. This study helps bridge that gap by using highly granular order messages data to construct the main attention proxy. These data are unique since they offer a microscopic look into investors' trading activity, and are rarely available in the US market. These data allow us to study the link between order cancellation activity and market efficiency post announcement release and, in effect, examine the behavioral relation between investor attention and asset prices using a microstructure channel.

This chapter contains five sections. Section 4.2 provides details about our data. Section 4.3 explains how we construct our attention measure. Section 4.4 presents the empirical results. Section 4.5 concludes.

4.2. Data description

Our main sample includes all announcements released outside of trading hours (overnight announcements) for all stocks listed on the Australian Securities Exchange (ASX) in 2019.[87] An advantage of conducting our study in the Australian market instead of the US market is that ASX imposes continuous disclosure obligations on all listed companies, requiring them to disclose all relevant material information (except confidential information) as they become aware of it. Thus, the Australian setting offers a unique laboratory to study how information is impounded into prices since distinct pieces of firm-specific information are released individually on a real-time basis, rather than in a combined form along with the company's financial statements. Another advantage of this setting is that the strict disclosure requirements make the Australian company announcements database a nearly complete compilation of relevant firm-specific information and render it less susceptible to strategic concealment of certain information.

The Australian equity market is among the top ten equity markets in the world by market capitalization, with ASX being among the top 20 stock exchanges.[88] ASX lists around 2,000 stocks with a total market capitalization of approximately \$2.07 trillion, or AUD 2.87 trillion, at the time of writing.

To filter out immaterial announcements, we restrict our sample to announcements classified as "price sensitive" by the ASX on the basis of their nature and information content. Further, we apply a liquidity filter by excluding all announcements for which there are no trades in the next trading day's opening auction. We collect the announcement data from the Australian company announcements database provided by the Securities Industry Research Centre of Asia-Pacific (SIRCA). From this database, we collect information about the date, time, and type of the announcement.[89]

In total, there are 11,753 announcements and 1,791 stocks in our sample. Of these, 4,110 announcements can be classified as positive news, and 7,643 can be classified as negative news based on the return in the ten trading days following the announcement release. Additionally, based

[87] Announcements released outside of trading hours offer a clean setting to test the effects of attention to information on price reactions, as investors have more time to assess the information content of the announcement. Initial price reactions to announcements made during trading hours are likely to be noisier.

[88] Source: Statista.

[89] Announcement types are based on ASX's disclosure requirements (ASX Listing Rule 3.1). There are 19 announcement types in total. The list of announcement types can be accessed at this link (https://bit.ly/3uf4vJL). A given announcement can belong to multiple types depending on its information content.

on the criteria provided in Mudalige, Kalev, and Duong (2016) and Prasad, Bakry, and Varua (2020), 7,515 announcements can be classified as scheduled, and 4,238 announcements can be classified as unscheduled announcements.[90] Lastly, 1,184 announcements are earnings announcements, while 10,569 announcements are non-earnings announcements.

We use order messages data for obtaining information on market activity. Order messages data include information about order additions, amendments, cancellations, and executions. Unlike the more conventional trades and quotes data and order book data, these data contain a unique identifier for each order that is sent to the market and, thus, enable us to trace the entire life cycle of all orders. We obtain these data from SIRCA, which in turn obtains these data on a real-time basis directly from ASX's trading platform. In addition, we obtain intraday price data from the Thomson Reuters Tick History database.

4.3. Attention measure

We use the highly granular order messages data to construct our investor attention measure, which aims to capture the cross-sectional and time-series variation in investor attention to fundamental information. The intuition behind our measure is simple. In essence, our measure categorizes investors who do not amend or cancel their orders in response to material overnight news as inattentive. An investor must pay attention to the market to realize that new material news about the stock has been released. Since the material nature of the information would cause a change in the fair valuation of the stock, this information is likely to affect an investor's pending order, which would be based on the valuation at the time of the order. Hence, cancellation or amendment of the order is warranted. Inattentive investors are likely to miss these events and, thus, are unlikely to cancel their pending orders in response to these events.

There are two primary channels through which material news can be conveyed to investors. First, material news can be disclosed by the company itself in the form of a material announcement. Given the continuous disclosure requirements imposed by ASX, this channel should cover almost all relevant firm-specific information. Second, if the news item does not necessitate a disclosure, it can reach investors through financial media outlets. This channel should cover all other relevant information such as macroeconomic news, industry news, expert commentary, analyst forecasts,

[90] Both Mudalige et al. (2016) and Prasad et al. (2020) classify announcements as scheduled or unscheduled based their type. Announcements that are periodic (non-periodic) in nature are classified as scheduled (unscheduled).

etc. In our data, the first channel can be easily identified by looking at overnight releases of material announcements. To capture the second channel, we look at the overnight return.[91] We classify all stock-days with an overnight return exceeding 0.5% as material news event days.[92]

In its most granular form, our measure can be constructed for individual announcements or in overnight intervals on each stock-day. We follow the latter method since it gives us a robust measure that captures cross-sectional variation in attention well. We start by constructing our attention measure at the stock-day level; the unit of observation is a stock i on day t ($Attention_{i,t}$). $Attention_{i,t}$ is the proportion of pending orders for stock i that are amended or canceled after trading hours on day t. We only calculate $Attention_{i,t}$ on days on which there is either an overnight material announcement for the stock or the overnight return (previous day's close to next day's open) for the stock exceeds 0.5%. To calculate this measure, we first extract the pending limit orders that remain in the limit-order book after market close on day t. We calculate the measure as the proportion of orders that are amended or canceled before the next trading day's open.[93] For our tests, we use the stock-level version of the measure, which is the average value for a given stock across all days. [94]

We restrict our focus to overnight news to try as best possible to prevent non-news-related motivations for amendments and cancellations from influencing our measure. During trading hours, in addition to news, market prices are a source of information on which investors can rely to make their trading decisions. Unless we restrict our intraday time window to after trading hours, it is difficult to know whether an investor's actions are in response to fundamental news or to price updates. Unlike Google searches and trading volumes, order amendments and cancellations allow us to restrict our intraday time window to after trading hours. For Google searches, intraday data for longer time ranges are unavailable, while there are no trades after hours. Focusing on overnight periods with material news releases offers the unique advantage of disentangling attention to fundamental information from attention to returns.

[91] Although we could directly look at firm-level and macroeconomic news coverage to capture this channel, doing so would not help us determine whether a given news item is material. Unlike announcements, there is no clean categorization of news articles into price sensitive and non-price sensitive news. In addition, it would be difficult to determine the exact impact of a macroeconomic news item on a stock without looking at the price reaction.
[92] In our sample, the mean overnight return is 0.4%.
[93] This proportion equals the number of orders that are amended or canceled divided by the total number of pending orders that remain in the order book after market close on day t.
[94] For readers' convenience, in our statistical tests, we use the inverse of this measure, $Inattention$.

4.4. Empirical tests

4.4.1. Summary statistics

We begin by discussing how our attention measure varies with certain stock characteristics. Table 4.1 reports the summary statistics for our attention measure for quartiles based on various stock characteristics. Predictably, attention level increases with market capitalization, with stocks in the top-quartile for market capitalization witnessing significantly higher investor attention than other stocks.

Additionally, our attention measure is inversely related to relative tick size, i.e., the tick size to price ratio. Stocks in the top-quartile for relative tick size have significantly lower attention levels than other stocks. This is likely because our attention measure relies on order cancellations. O'Hara, Saar, and Zhong (2019) find that when the relative tick size of a stock is high, high-frequency traders, who are often associated with extremely high order cancellation activity (Khomyn and Putnins, 2021), leave orders lying in the book for longer. As stocks with high relative tick sizes witness fewer order cancellations from these traders, these stocks have lower values for our attention measure.

Interestingly, the attention measure seems to have a U-shaped relation with volatility, implying that the most and least volatile stocks receive higher investor attention than stocks in the middle of the volatility spectrum.

Table 4.1
Summary statistics

This table reports descriptive statistics for our attention measure. Reported below are the median attention values for quartiles based on various stock characteristics. Relative tick size is the ratio of a stock's tick size to its price. The sample includes all stocks listed on the Australian Securities Exchange in the year 2019.

	Quartile			
Variable	1 (Bottom)	2	3	4 (Top)
Attention level (%)	5.06%	8.70%	13.79%	38.41%
Announcement releases (%)	14.81%	16.67%	12.50%	11.29%
Relative tick size (%)	39.88%	14.29%	10.00%	6.25%
Market capitalization (%)	6.67%	8.89%	11.76%	36.22%
Liquidity (%)	7.14%	8.22%	11.29%	37.14%
Volatility (%)	13.64%	10.00%	10.25%	23.94%
Skewness (%)	12.00%	8.22%	13.04%	25.93%

Next, we split the stocks in our sample into groups based on their attention level and examine the characteristics of each group (see Table 4.2). We classify stocks with a value below 20% as low attention stocks, those with a value between 20% and 50% as medium attention stocks, and those with a value above 50% as high attention stocks.[95] We adopt these thresholds since they classify roughly 5% stocks as high attention stocks, around 15% stocks as medium attention stocks, and approximately 80% stocks as low attention stocks. This distribution seems reasonable given that the broadest market indices only cover around 15%–25% of the stocks listed on ASX.[96] We retain this attention classification in all the subsequent tables and figures.

Unsurprisingly, the low attention group is the largest in terms of the number of announcements, containing approximately 72% of the announcements in our sample. Stocks that receive high investor attention are significantly higher in size and liquidity than other stocks. Interestingly, these stocks are also substantially more volatile than other stocks.

Table 4.2
Attention group characteristics

This table reports descriptive statistics for stock groups formed based on attention levels. Reported below are the number of stocks and announcements in each group and the median values for various stock characteristics. Liquidity is the average daily dollar trading volume. Volatility is the standard deviation of daily returns. Relative tick size is the ratio of a stock's tick size to its price. AUD is an acronym for the Australian dollar. The sample includes all stocks listed on the Australian Securities Exchange in the year 2019.

	Attention level		
Variable	Low	Medium	High
Number of stocks	1,534	326	132
Number of announcements	15,347	2,910	1,113
Market capitalization (AUD million)	45.50	625.94	3,474.91
Liquidity (AUD million)	0.04	1.11	10.37
Volatility (%)	9.37%	21.90%	197.03%
Relative tick size (%)	2.42%	0.37%	0.16%

4.4.2. Price reactions to announcements

In this section, we examine how investor attention, particularly low investor attention, impacts price reactions to announcements. For the main analysis in this section, we analyze price reactions in two post-announcement time windows: the first day, examining the intraday process of price

[95] Our main results are robust to setting the maximum threshold for low attention stocks as 25% instead of 20%.
[96] These indices are the S&P/ASX 300 index and the All Ordinaries index, which consist of 300 and 500 stocks, respectively.

adjustment, and the first 20 days, examining the medium-term process, including how information is gradually incorporated into the price. Throughout the chapter, we calculate cumulative abnormal returns using the method in Hirshleifer and Sheng (2019).

Figure 4.1 plots price reactions split by attention level. We plot announcements conveying positive and negative news separately, using the return in the ten trading days following the announcement to classify announcements as positive or negative news. Panel A plots the cumulative abnormal return in five-minute intervals on the trading day following overnight announcements. This graph shows sharp price reactions at market open and some preliminary evidence that, in some cases, this initial price reaction is an underreaction and is followed by a price drift. Although the drift might not be visually obvious, negative announcements for low attention stocks have a −1.4% abnormal return from 11:00 to market close.[97]

Figure 4.1 Panel B plots the cumulative abnormal return in one-and-a-half-hour intervals for 20 trading days after the announcement. Here, the drifts are more visually evident and exist for all stocks except high attention stocks. Drifts are more pronounced for positive announcements and low attention stocks. Prices seem to level off approximately 11 days after the announcement. We observe similar price patterns even when we plot price reactions for the extreme announcement return quintiles. These plots are provided in Figure 4.A1 in Appendix 4.B.

Next, we isolate announcements with the most extreme returns. Figure 4.2 plots the same graphs as Figure 4.1, except only for the top- and bottom-quintiles for announcement return.[98] Figure 4.2 Panel A plots the intraday price paths in the next trading day for announcements in the extreme return quintiles. As in Figure 4.1 Panel A, there is an initial underreaction and a subsequent drift for low attention stocks even though this might not be visually obvious for positive announcements (top return quintile) because of the high initial return. For low attention stocks, the abnormal return from 11:00 to market close is 1.89% and −1.02% for positive and negative announcements, respectively. In contrast, during the same period, the weighted-average abnormal returns for medium and high attention stocks combined are 1.56% and −0.81% for positive and negative announcements, respectively.

[97] To open the market every day, ASX conducts opening auctions between 10:00 and 10:10. Stocks are sorted alphabetically into groups based on the first letter of their ASX code. Groups are assigned time slots, and stocks in the same group open together during the assigned time slot. The market closes with a closing auction conducted between 16:10 and 16:11.

[98] Here, announcement return is the same return as the one used to classify announcements as positive or negative, i.e., return in the ten trading days following the announcement.

Panel A: First day

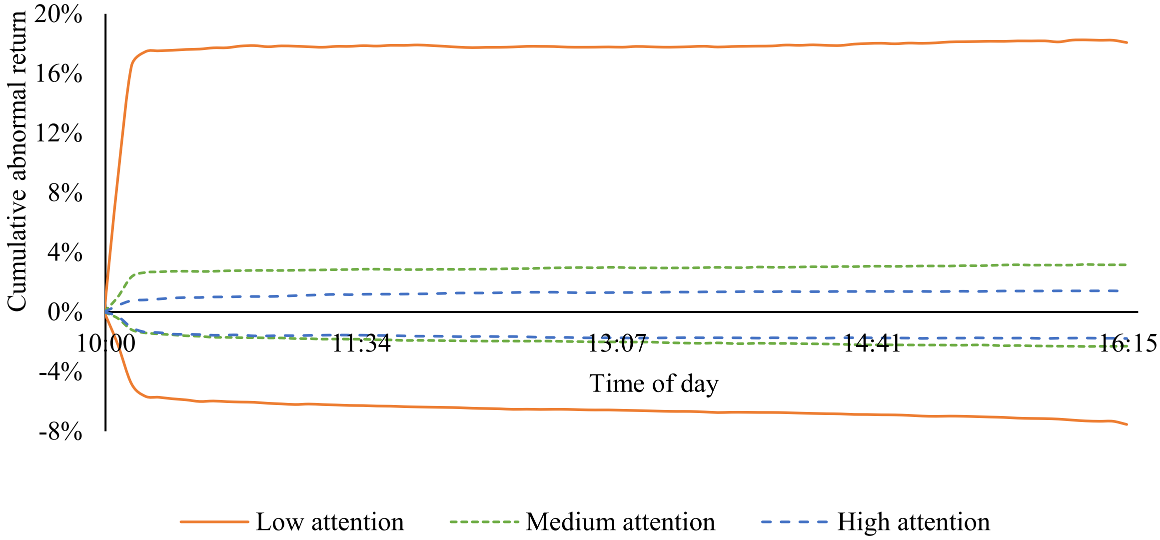

Panel B: First 20 days

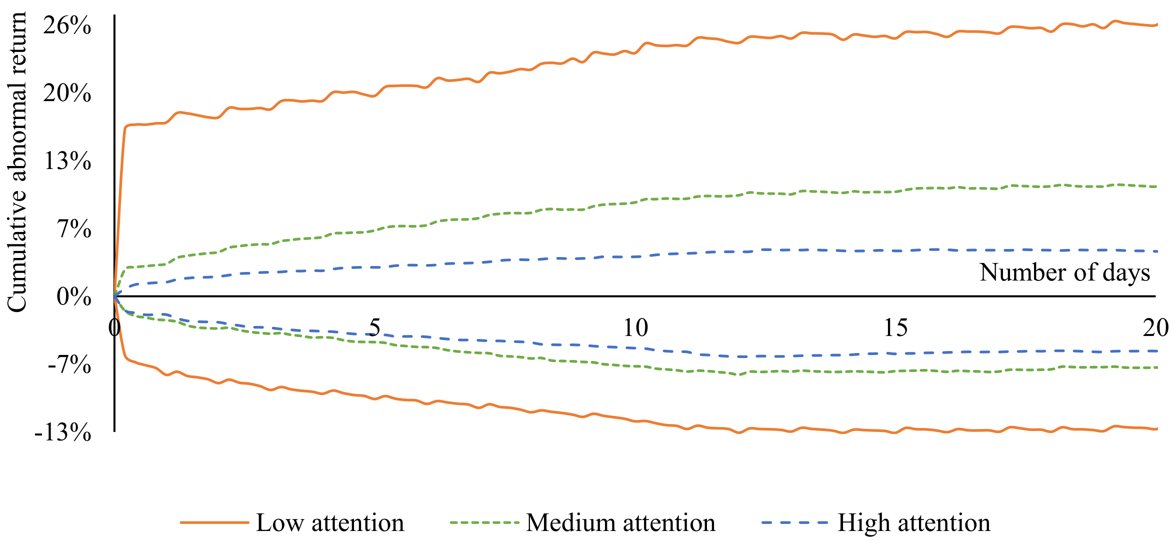

Figure 4.1. Price reactions split by attention level.
This figure plots price reactions to announcements split into buckets based on the stock's attention level. Panel A plots cumulative abnormal returns in five-minute intervals in the next trading day after the announcement. Panel B plots cumulative abnormal returns in one-and-a-half-hour intervals across the next 20 trading days after the announcement. The figure plots positive and negative announcements separately. Announcements are classified as positive or negative based on the return in the ten trading days after the announcement. Announcements are split into attention buckets based on our attention measure. The attention measure values for low attention announcements are below 20%, the values for medium attention announcements are between 20% and 50%, while the values for high attention announcements are higher than 50%.

Panel A: First day

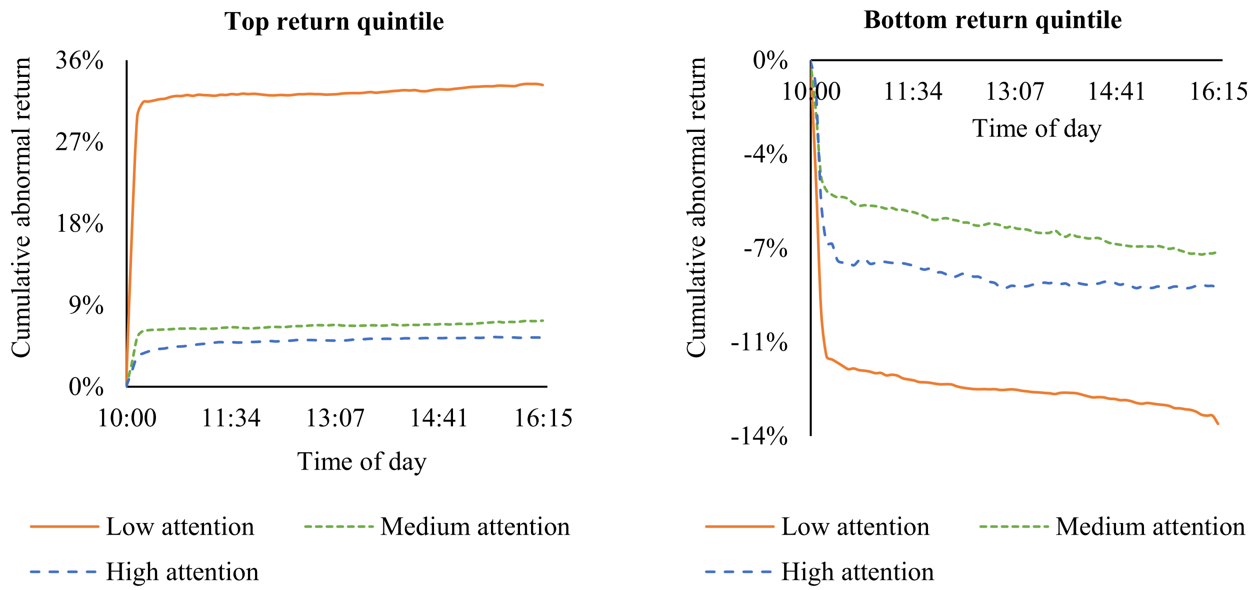

Panel B: First 20 days

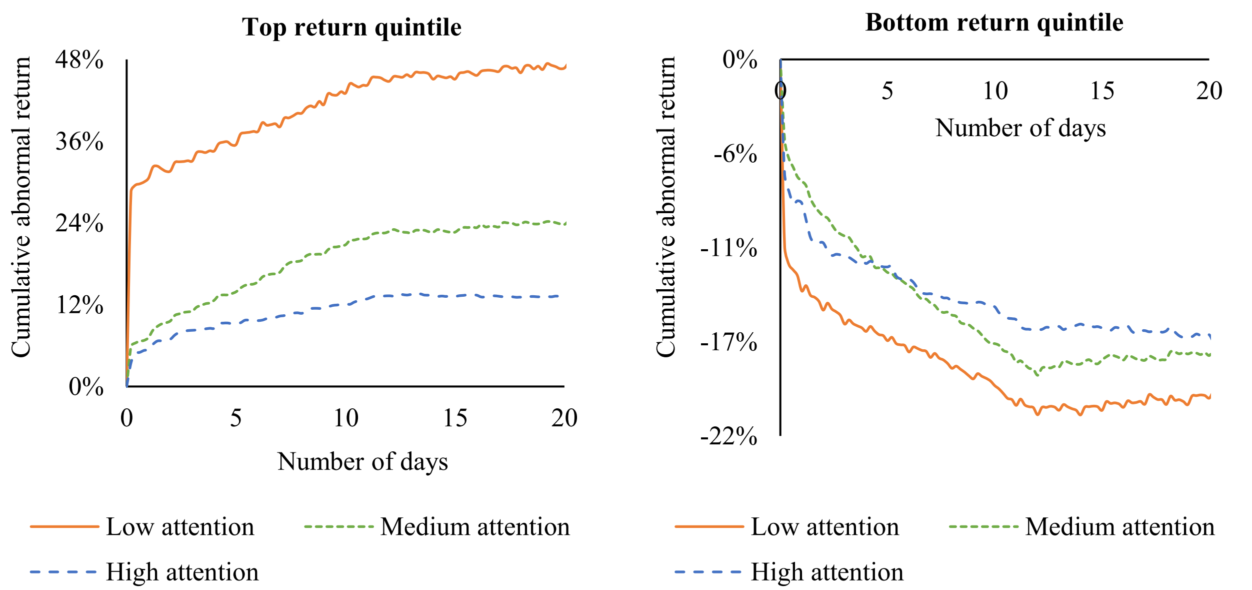

Figure 4.2. Price reactions to announcements with extreme returns.
This figure plots price reactions to announcements with extreme returns. Panel A plots cumulative abnormal returns in five-minute intervals in the next trading day after the announcement. Panel B plots cumulative abnormal returns in one-and-a-half-hour intervals across the next 20 trading days after the announcement. Price reactions are only plotted for announcements in either the top-quintile (left) or the bottom-quintile (right) for announcement return. This return is calculated from the time of the announcement to ten trading days later. Announcements in each of the extreme quintiles are split into attention buckets based on our attention measure. The attention measure values for low attention announcements are below 20%, the values for medium attention announcements are between 20% and 50%, while the values for high attention announcements are higher than 50%.

Figure 4.2 Panel B plots the price paths in the 20 trading days after the announcement. Like Figure 4.1 Panel B, these plots indicate that positive announcements for low attention stocks are the most mispriced, with these announcements earning an abnormal return of around 17% from one day after the announcement to 20 days after the announcement. The efficiency gap between low and high attention stocks is much more obvious visually when we plot the post-announcement drift (excluding the initial return) for the extreme return quintiles separately for low and high attention stocks. Figure 4.A2 in Appendix 4.B displays these plots.

To test whether the relation between the post-announcement drifts and a stock's attention level as illustrated in the previous figures is statistically significant, we regress the post-announcement return (excluding the initial return) of each announcement j for stock i on the inverse of our attention measure ($Inattention_{j,i}$) and control variables. We break down the 20-day window from the figures into two parts: the first day and the subsequent 19 days.[99] We define the announcement return as the return from the time of announcement release to ten trading days after the announcement. We use the announcement return to categorize announcements as positive or negative.[100] Following Hirshleifer and Sheng (2019), we control for the stock's trading volume and volatility, the market return, and the number of announcements on the day of the announcement. We also include day of the week and month fixed effects.

Table 4.3 reports the results, with positive news in Panel A and negative news in Panel B. For positive news, according to Models 1–4 (with and without controls and fixed effects), $Inattention_{j,i}$ does not significantly affect returns on the first day after the announcement. However, in the subsequent 19 days (Models 5–8), $Inattention_{j,i}$ has a statistically significant positive relation with the post-announcement abnormal return. This result implies that the post-announcement drift for positive announcements increases with investor inattention. For negative news, the effect of inattention on post-announcement returns is confined to the first day after the announcement. As in the case of positive announcements, the post-announcement drift increases with inattention for negative announcements.

---

[99] The first-day analysis only includes the post-announcement drift. It excludes the opening return, which we deem the initial reaction to the announcement.

[100] We use this method to classify announcements as positive or negative for all our tests.

Table 4.3
Investor attention and post-announcement drifts

This table reports regression results testing how investor attention affects the post-announcement return. The unit of observation is an announcement j for stock i. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $Inattention_{j,i}$, is the inverse of our attention measure. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for positive announcements. Panel B displays the results for negative announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =							
	First day				Next 19 days			
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Positive announcements								
$Inattention_{j,i}$	–0.34	–0.36	–0.61	–0.56	12.50***	12.59***	12.61***	12.60***
	(–0.55)	(–0.58)	(–0.92)	(–0.84)	(6.47)	(6.49)	(6.40)	(6.37)
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Time fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.01%	0.56%	0.49%	1.15%	1.00%	1.93%	2.56%	3.63%
Observations	4,108	4,108	3,886	3,886	4,108	4,108	4,108	4,108
Panel B: Negative announcements								
$Inattention_{j,i}$	–2.78***	–2.82***	–2.89***	–2.95***	–1.14	–0.53	0.79	1.25
	(–4.56)	(–4.61)	(–4.53)	(–4.59)	(–0.72)	(–0.33)	(0.49)	(0.78)
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Time fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.48%	0.96%	0.82%	1.41%	0.00%	0.92%	2.06%	2.91%
Observations	4,281	4,281	4,026	4,026	4,281	4,281	4,281	4,281

In the next set of tests, we replace $Inattention_{j,i}$ with a dummy variable ($LowAttn_{j,i}$) that equals one if stock i is a low attention stock based on the definition in Section 4.4.1.[101] Table 4.4 Panel A displays the results for positive announcements. As expected, low attention stocks have stronger post-announcement drifts in the 20-day horizon but not on the first day after the announcement, earning a 4.36% higher abnormal return than other stocks in the 20 days after positive announcements.

[101] Based on this definition, we classify all stocks with a value below 20% for our attention measure as low attention stocks.

Table 4.4 Panel B displays the results for negative announcements. Like the previous tests, low attention stocks have stronger drifts only on the first day after negative announcements. Low attention stocks earn a 1.11% lower abnormal return on the first day after a negative announcement. As an additional robustness check, we run the tests in Tables 4.3 and 4.4 after constructing our attention measure using only overnight periods with material announcements, excluding overnight periods with sharp opening price movements.[102] Table 4.A1 in Appendix 4.A reports these results. Our results are robust to the use of only overnight periods with material announcements to construct our measure.

Table 4.4
Low attention stocks

This table reports regression results testing whether the post-announcement drift is stronger for low attention stocks. The unit of observation is an announcement j for stock i. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for positive announcements. Panel B displays the results for negative announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =							
	First day				Next 19 days			
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Positive announcements								
$LowAttn_{j,i}$	0.06	0.02	–0.05	–0.05	4.13***	4.20***	4.32***	4.36***
	(0.21)	(0.07)	(–0.17)	(–0.18)	(5.17)	(5.24)	(5.36)	(5.39)
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Time fixed effects	No	Yes	No	No	No	Yes	No	Yes
R^2	0.00%	0.55%	0.47%	1.13%	0.65%	1.58%	2.27%	3.36%
Observations	4,108	4,108	3,886	3,886	4,108	4,108	4,108	4,108
Panel B: Negative announcements								
$LowAttn_{j,i}$	–1.06***	–1.05***	–1.10***	–1.11***	–0.81	–0.59	–0.34	–0.17
	(–4.54)	(–4.50)	(–4.52)	(–4.51)	(–1.34)	(–0.98)	(–0.56)	(–0.28)
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Time fixed effects	No	Yes	No	Yes	No	Yes	No	Yes
R^2	0.48%	0.94%	0.82%	1.39%	0.04%	0.94%	2.06%	2.90%
Observations	4,281	4,281	4,026	4,026	4,281	4,281	4,281	4,281

[102] We perform this robustness check since opening price movements could be a noisy proxy for material non-disclosure news if they are caused by reasons other than material news releases, such as market manipulation.

We examine the drifts more closely by performing the main tests in Tables 4.3 and 4.4 in the following time windows after the announcement: day one to day ten, day 11 to day 20, day 21 to day 40, and day 41 to day 60. Table 4.5 Panel A displays the results for positive announcements. The effects of attention on drifts are the strongest in the first ten-day period after positive announcements. However, the effects are still economically significant until 40 days after the announcement. Interestingly, after 40 days, low attention stocks experience a reversal, suggesting that some part of the initial reaction and the subsequent drift might have been an overreaction.

Table 4.5 Panel B displays the results for negative announcements. Negative announcements for low attention stocks do not experience stronger drifts than other stocks at any time horizon. They do, however, experience stronger reversals from 20 days to 40 days after the announcement. Overall, this section indicates that announcements generally witness lengthy post-announcement price drifts; these drifts are stronger when investor attention is low and when announcements disclose positive news.

4.4.3. Trading strategy

One way to test whether low attention stocks truly witness stronger drifts than other stocks and to quantify whether the effect is economically meaningful is to construct a trading strategy that exploits the drift caused by investor inattention and measure the returns of the strategy. If the market reaction to announcements for low attention stocks is efficient, such a trading strategy will, on average, not generate a profit. The basic idea is to buy stocks following good news and sell stocks following bad news to profit from the price drift. To isolate the role of inattention, we compare the returns to the strategy implemented conditional on low attention to returns from a similar strategy but not conditioning on attention. To avoid look-ahead bias, we ensure that the strategy only uses attention-related information available at the time of the announcement.

We start by getting the past distributions of our attention measure for a given day t, by using an expanding window that starts on the first day of our sample period (day one) and ends on day $t-1$. Each observation in this distribution is the average value of our daily attention measure for a given stock i until day $t-1$. We allow for a minimum number of observations to build up initially by setting the condition $t \geq 15$.[103]

[103] Our results are robust to starting the window at a lower $t = 10$ and a higher $t = 30$.

Table 4.5
Attention effects across different time horizons

This table reports regression results testing the attention effects on the post-announcement drift at different time horizons. The unit of observation is an announcement j for stock i. $CAR[t, t+1]$ is the cumulative abnormal return from t days after the announcement until $t+1$ days after the announcement. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. $Inattention_{j,i}$ is the inverse of our attention measure. $LowAttn_{j,i}$ is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for positive announcements. Panel B displays the results for negative announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =							
	$CAR[1,10]$		$CAR[11,20]$		$CAR[21,40]$		$CAR[41,60]$	
Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Positive announcements								
$Inattention_{j,i}$	8.11***		3.26***		3.55***		–41.28***	
	(5.86)		(2.93)		(2.20)		(–3.69)	
$LowAttn_{j,i}$		2.73***		1.18***		0.53		–10.22**
		(4.81)		(2.60)		(0.81)		(–2.23)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	3.24%	2.98%	1.33%	1.29%	1.93%	1.83%	15.46%	15.28%
Observations	4,108	4,108	4,108	4,108	4,108	4,108	4,108	4,108
Panel B: Negative announcements								
$Inattention_{j,i}$	–0.12		1.40		7.53***		1.74	
	(–0.11)		(1.11)		(3.73)		(0.89)	
$LowAttn_{j,i}$		–0.21		0.01		2.81***		0.78
		(–0.53)		(0.03)		(3.67)		(1.06)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	3.66%	3.67%	1.48%	1.46%	1.37%	1.36%	1.85%	1.85%
Observations	4,281	4,281	4,281	4,281	4,281	4,281	4,281	4,281

We use our attention distribution to construct our trading strategy at three different time horizons. Our strategy purchases low attention stocks (attention value below 20%) after a positive announcement and sells short low attention stocks after a negative announcement. We compare the returns for this strategy against a baseline strategy that purchases stocks after positive announcements, irrespective of the stocks' attention level, and sells short stocks after negative announcements. Since we document stronger drifts for positive announcements than for negative announcements, we can expect the long leg of our strategy to perform better than the short leg, which might not even outperform the market portfolio. However, following the literature, we still construct long-short strategies to mitigate concerns about market risk.

We implement our strategies at three time horizons: a one-day horizon, a ten-day horizon, and a 20-day horizon. These time horizons differ in the return period used to classify announcements as positive or negative. Under the one-day horizon, our strategies use the price reaction at the next trading day's open to classify announcements as positive or negative and establish initial positions after the opening auction. These positions are closed out at market close on the same trading day. Under the ten-day horizon, our strategies establish initial positions after the opening auction two trading days after the announcement and use the return until that point in time to classify announcements as positive or negative. Finally, under the 20-day horizon, our strategies establish initial positions after the opening auction ten trading days after the announcement and use the return until that point in time to classify announcements as positive or negative. The final criterion for categorizing announcements as positive or negative resembles that described in Section 4.4.2.

Table 4.6 reports the excess returns for value-weighted portfolios constructed based on the above strategies.[104] To capture the excess return for each strategy, we regress the strategy returns on the factors in the Fama-French three-factor model (Fama and French, 1992). At the one-day horizon, not only does our main attention-based strategy underperform the baseline strategy, it also does not earn a statistically significant alpha. This is likely because the opening price reaction is too noisy to accurately categorize announcements as positive or negative, particularly for low attention stocks.

In line with this reasoning, the performance of our main strategy improves substantially at the ten-day and 20-day horizons, earning a higher alpha than the baseline strategy at both time

[104] Our results are robust to the use of equal-weighted portfolios. The equal-weighted results are reported in Table 4.A2 in Appendix 4.A.

horizons. The attention-based strategy earns a statistically significant daily alpha of 0.46% and 0.41%, outperforming the baseline strategy by 21% and 95% at the ten-day and 20-day horizons, respectively.[105] Interestingly, the low attention drifts are exploitable even if investors trade ten days after the announcement, despite the abnormal returns seemingly leveling off at this point, according to the graphs in Section 4.4.2. Overall, these results suggest that the drifts for low attention stocks are not exploitable on the first day after the announcement but are exploitable ten days and 20 days after the announcement.

Table 4.6
Trading strategies exploiting the drift

This table reports results for a trading strategy exploiting the post-announcement drift. Portfolios are constructed using the stock's attention level and the announcement's observed return. The attention-based strategy purchases a low attention stock (attention measure below 20%) when the announcement released has a positive return and sells short a low attention stock when the announcement released has a negative return. The baseline strategy purchases a stock when the announcement released has a positive return, and sells short a stock when the announcement released has a negative return. Results are reported for three time horizons. Under the one-day horizon, strategies establish initial positions at market open on the trading day following the announcement, after having observed the opening price reaction to the announcement, and close-off positions at market close on the same day. Under the ten-day horizon, strategies establish initial positions at market open two trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. Under the 20-day horizon, strategies establish initial positions at market open ten trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. All portfolios are value-weighted. $MarketRtn_t$, HML_t, and SMB_t are the market, value, and size factors, respectively, from the Fama-French three-factor model. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =					
	Attention-based strategy			Baseline strategy		
Variable	One day	Ten days	20 days	One day	Ten days	20 days
$Intercept$	–0.32	0.46***	0.41***	0.65**	0.38***	0.21***
	(–0.65)	(4.76)	(3.48)	(2.53)	(4.94)	(3.26)
$MarketRtn_t$	–66.23	29.62**	34.57**	30.54	–7.17	–2.84
	(–0.98)	(2.25)	(2.16)	(0.88)	(–0.69)	(–0.33)
HML_t	34.47	–0.60	–6.21	50.18***	–6.79*	–9.91***
	(1.52)	(–0.13)	(–1.12)	(4.21)	(–1.90)	(–3.32)
SMB_t	–56.86	–3.60	–8.40	–24.87	–14.64**	–18.91***
	(–1.36)	(–0.43)	(–0.82)	(–1.13)	(–2.21)	(–3.42)
R^2	17.59%	2.57%	3.33%	41.56%	2.12%	5.15%
Observations	225	237	229	232	237	229

[105] Table 4.A3 in Appendix 4.A reports the results of the long leg and the short leg of the attention-based strategy separately. As conjectured, the contribution of the long leg is substantially higher than that of the short leg. However, at the ten-day horizon, even the short leg earns a statistically significant alpha.

4.4.4. Heterogeneity

In this section, we explore whether the role of attention in post-announcement price reactions varies with announcement- and stock-level factors. In doing so, we examine whether particular announcement- or stock-specific characteristics drive our results and whether post-announcement price reactions differ for announcements or stocks with certain characteristics.

4.4.4.1. Scheduled and unscheduled announcements

We start by examining whether the existence of post-announcement drifts for low attention stocks depends on whether the announcement is scheduled or unscheduled. Due to the periodic nature of scheduled announcements, price reactions for such announcements could, on average, be more efficient than those for unscheduled announcements for at least two reasons. First, investors might endogenously allocate more attention to scheduled announcements. This might happen if companies tend to release these announcements at certain times of the year; investors might intently watch the market at these times in anticipation of an announcement. This can help investors price scheduled announcements more accurately than unscheduled announcements, which, because of their unpredictable time of arrival, might be released when investors are not watching the market. Second, due to the recurring nature of scheduled announcements, investors might, over time, become more adept at swiftly gauging the true price impact of the information content in these announcements. This can cause them to efficiently price scheduled announcements faster than unscheduled announcements. Alternatively, if companies tend to systematically release more complex information in scheduled announcements than in unscheduled announcements, we can expect investors to price scheduled announcements less efficiently than unscheduled ones. However, if the drifts are caused by low investor attention, then all announcements, irrespective of whether they are scheduled or unscheduled, for low attention stocks would be priced less efficiently than announcements for other stocks.

We conduct the tests in Table 4.4 for scheduled and unscheduled announcements separately. The results are reported in Table 4.7. The results for both scheduled and unscheduled announcements resemble those in Section 4.4.2. For negative announcements, low attention stocks have stronger drifts only on the first day after the announcement, while for positive announcements, low attention stocks have stronger drifts following the first day. These results suggest that post-announcement underreactions are not fully driven by either scheduled or unscheduled announcements. Hence, the periodicity of the announcement does not explain the post-announcement drifts.

Table 4.7
Scheduled and unscheduled announcements

This table reports regression results testing whether attention affects post-announcement drifts for scheduled and unscheduled announcements differently. The unit of observation is an announcement j for stock i. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for scheduled announcements. Panel B displays the results for unscheduled announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	Positive announcements		Negative announcements	
Variable	First day	Next 19 days	First day	Next 19 days
Panel A: Scheduled announcements				
$LowAttn_{j,i}$	–0.30	4.59***	–0.93***	0.14
	(–0.87)	(4.60)	(–3.05)	(0.19)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.65%	4.64%	1.57%	3.59%
Observations	2,638	2,776	2,726	2,889
Panel B: Unscheduled announcements				
$LowAttn_{j,i}$	0.59	3.88***	–1.51***	–0.63
	(1.36)	(2.72)	(–3.68)	(–0.56)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	2.41%	2.58%	2.54%	3.63%
Observations	1,248	1,332	1,300	1,392

4.4.4.2. Earnings and non-earnings announcements

We separately examine post-announcement price reactions for a specific type of announcement that has been extensively studied in the literature: earnings announcements. The phenomenon of lengthy drifts following earnings announcements, in particular when investor attention is low, is well documented in the literature (Ball and Brown, 1968; Bernard and Thomas, 1990; DellaVigna and Pollet, 2009; Hirshleifer et al., 2009; Ball and Brown, 2019). It would be interesting to examine whether, in our data, the low attention drifts observed are primarily driven by earnings announcements, or whether non-earnings announcements for low attention stocks also witness post-announcement drifts.

Table 4.8
Earnings and non-earnings announcements

This table reports regression results testing whether attention affects post-announcement drifts for earnings and non-earnings announcements differently. The unit of observation is an announcement j for stock i. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for earnings announcements. Panel B displays the results for non-earnings announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	Positive announcements		Negative announcements	
Variable	First day	Next 19 days	First day	Next 19 days
Panel A: Earnings announcements				
$LowAttn_{j,i}$	0.15	6.35***	–0.79	1.46
	(0.24)	(3.83)	(–1.12)	(1.01)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	3.54%	13.82%	4.39%	5.04%
Observations	481	513	386	415
Panel B: Non-earnings announcements				
$LowAttn_{j,i}$	0.01	3.68***	–1.20***	–0.37
	(0.02)	(4.03)	(–4.46)	(–0.55)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.21%	3.05%	1.50%	2.97%
Observations	3,405	3,595	3,640	3,866

We conduct the tests in Table 4.4 for earnings and non-earnings announcements separately. The results are reported in Table 4.8. Table 4.8 Panel A reports the results for earnings announcements. Although earnings announcements for low attention stocks do not experience stronger drifts on the first day after the announcement, positive earnings announcements for low attention stocks do experience stronger drifts than those for other stocks in the subsequent 19 days. Table 4.8 Panel A reports the results for non-earnings announcements. The results for non-earnings announcements match those for the entire sample; negative announcements for low attention stocks witness stronger drifts than other stocks on the first day following the announcement, and positive announcements experience stronger drifts in the subsequent 19 days. This suggests that the first-day low attention drifts for negative announcements discussed in Section 4.4.2 are primarily driven by non-earnings announcements. However, in their entirety, low attention drifts are not fully explained by either earnings or non-earnings announcements individually.

4.4.4.3. Lottery-like and small capitalization stocks

Next, we test whether a stock's characteristics, other than its attention level, confound our results. Specifically, we examine whether the mispricing effects we document can be explained by the degree of retail investor interest in a stock. Previous literature finds that stocks favored by retail investors witness higher mispricing than other stocks (Bali et al., 2019). In this regard, we test whether the observed drifts are caused by retail interest rather than attention. If our results are driven by retail interest, the low attention drifts will either vanish once we control for the level of retail investor interest in a stock or will primarily be concentrated in stocks favored by retail investors. Given the previous evidence of retail investor preference for lottery-like stocks, i.e., stocks that have a low price, high idiosyncratic volatility, and skewness, and small capitalization stocks, we classify these stocks as stocks favored by retail investors (Kumar, 2009; Bali et al., 2019; Luo et al., 2020).[106] We label a stock as lottery-like if it has a below-median price, above-median idiosyncratic volatility, and above-median idiosyncratic skewness. Idiosyncratic volatility and skewness are the second and third moments, respectively, of the residuals obtained by regressing the daily stock returns on the factors in the Fama-French three-factor model. Since our sample period is 2019, we use stock prices in 2018 for *ex-ante* identification of stocks with lottery-like payoffs. In addition, we classify stocks in the bottom-tercile for market capitalization as small capitalization stocks.

We first examine lottery-like stocks and non-lottery-like stocks separately. Table 4.9 reports the results. The first-day low attention drift for negative announcements and the next 19 day drift for positive announcements only occur in non-lottery-like stocks and cannot be observed in lottery-like stocks. We conduct similar tests for small capitalization stocks. Table 4.10 reports the results.

[106] Our results are robust to the use of the Bali, Cakici, and Whitelaw (2011) proxy for lottery-like nature, i.e., the maximum daily return for the stock in the previous month. The results for this proxy are reported in Table 4.A4 in Appendix 4.A. Even using this proxy, lottery-like stocks do not fully explain the mispricing effects of attention.

Small capitalization stocks that witness low investor attention do not have stronger post-announcement drifts than other small capitalization stocks. Instead, the low attention drifts we observe previously primarily occur in other stocks. These results show that the mispricing effects we observe are not entirely concentrated in stocks favored by retail investors and that attention affects post-announcement mispricing independent of retail interest.

Overall, in this section, we find that inattention affects announcements universally. Inattention leads to stronger drifts after both scheduled and unscheduled releases and both earnings and non-earnings announcements. These effects are also not explained by other mispricing effects observed in the literature such as mispricing of stocks favored by retail investors. Low investor attention appears to be a key driver of post-announcement drifts.

Table 4.9
Lottery-like stocks

This table reports regression results testing whether post-announcement drifts are driven by lottery-like stocks. The unit of observation is an announcement j for stock i. Lottery-like stocks have low prices, high idiosyncratic volatility, and high idiosyncratic skewness. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for lottery-like stocks. Panel B displays the results for small capitalization stocks. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	Positive announcements		Negative announcements	
Variable	First day	Next 19 days	First day	Next 19 days
Panel A: Lottery-like stocks				
$LowAttn_{j,i}$	0.79	–6.51	0.99	1.60
	(0.35)	(–0.83)	(0.35)	(0.26)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	3.14%	4.04%	2.08%	2.87%
Observations	1,101	1,177	1,250	1,339
Panel B: Other stocks				
$LowAttn_{j,i}$	0.14	2.89***	–1.11***	–0.72
	(0.48)	(4.19)	(–5.58)	(–1.24)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.29%	3.00%	2.11%	4.39%
Observations	2,733	2,874	2,716	2,880

Table 4.10
Small capitalization stocks

This table reports regression results testing whether post-announcement drifts are driven by small capitalization stocks. The unit of observation is an announcement j for stock i. Small capitalization stocks are in the bottom-tercile for market capitalization. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for small capitalization stocks. Panel B displays the results for other stocks. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	Positive announcements		Negative announcements	
Variable	First day	Next 19 days	First day	Next 19 days
Panel A: Small capitalization stocks				
$LowAttn_{j,i}$	8.04	1.90	–12.47	6.67
	(0.99)	(0.06)	(–0.99)	(0.40)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	5.67%	3.94%	3.91%	3.04%
Observations	485	522	693	741
Panel B: Other stocks				
$LowAttn_{j,i}$	0.10	4.01***	–1.08***	–0.40
	(0.36)	(5.44)	(–5.51)	(–0.73)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.14%	4.93%	1.95%	5.00%
Observations	3,339	3,520	3,297	3,503

4.5. Conclusion

In this chapter, we introduce a novel measure of investor attention that classifies investors as inattentive if they do not update their pending limit orders in the stock market upon the release of material information. Relative to other measures of investor attention, our measure offers the advantage of isolating attention to information from attention to returns or volumes.

We use the new measure to examine how cross-sectional variation in attention affects price reactions to material corporate disclosures. We document stronger post-announcement underreactions and drifts for stocks that receive low investor attention. Underreactions after positive announcements are significantly stronger than those after negative announcements. A trading strategy that exploits drifts in low attention stocks earns a daily alpha of 0.46% from the first day to the tenth day after the announcement and 0.41% from the tenth day to the 20th day after the announcement.

Inattention affects how the market impounds scheduled and unscheduled information as well as earnings and non-earnings announcements. These effects cannot be explained by the general mispricing of stocks favored by retail investors. Collectively, these findings suggest that inattention is a pervasive driver of inefficiency in how markets impound material stock-specific information and, therefore, an impediment to efficient price formation.

Appendix 4.A. Additional regression results

Table 4.A1
Attention measured using only material announcements

This table reports regression results for our main tests after computing our attention measure using only overnight periods with material announcements. The unit of observation is an announcement j for stock i. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. $Inattention_{j,i}$ is the inverse of the modified attention measure. $LowAttn_{j,i}$ is an indicator variable that equals one if the stock's value for the modified attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. Panel A displays the results for positive announcements. Panel B displays the results for negative announcements. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	First day		Next 19 days	
Variable	(1)	(2)	(3)	(4)
Panel A: Positive announcements				
$Inattention_{j,i}$	–0.89		12.29***	
	(–1.29)		(5.99)	
$LowAttn_{j,i}$		–0.20		3.67***
		(–0.75)		(4.55)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.17%	1.14%	3.53%	3.17%
Observations	3,882	4,108	4,104	4,104
Panel B: Negative announcements				
$Inattention_{j,i}$	–2.71***		1.97	
	(–4.08)		(1.18)	
$LowAttn_{j,i}$		–0.94***		0.36
		(–3.87)		(0.59)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.29%	1.25%	2.93%	2.90%
Observations	4,021	4,021	4,276	4,276

Table 4.A2

Trading strategies exploiting the drift (equal-weighted results)

This table reports results for a trading strategy exploiting the post-announcement drift. Portfolios are constructed using the stock's attention level and the announcement's observed return. The attention-based strategy purchases a low attention stock (attention measure below 20%) when the announcement released has a positive return, and sells short a low attention stock when the announcement released has a negative return. The baseline strategy purchases a stock when the announcement released has a positive return, and sells short a stock when the announcement released has a negative return. Results are reported for three time horizons. Under the one-day horizon, strategies establish initial positions at the market open on the trading day following the announcement, after having observed the opening price reaction to the announcement, and close-off positions at market close on the same day. Under the ten-day horizon, strategies establish initial positions at the market open two trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. Under the 20-day horizon, strategies establish initial positions at the market open ten trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. All portfolios are equal-weighted. $MarketRtn_t$, HML_t, and SMB_t are the market, value, and size factors, respectively, from the Fama-French three-factor model. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =					
	Attention-based strategy			Baseline strategy		
Variable	One day	Ten days	20 days	One day	Ten days	20 days
$Intercept$	–0.55	0.46***	0.44***	–0.18	0.37***	0.36***
	(–1.08)	(4.16)	(3.73)	(–0.54)	(4.68)	(5.72)
$MarketRtn_t$	8.04	30.49**	23.71	36.64	–12.42	1.94
	(0.11)	(2.05)	(1.50)	(0.81)	(–1.16)	(0.23)
HML_t	28.65	0.19	–8.05	36.07**	–6.43	–5.58*
	(1.22)	(0.04)	(–1.48)	(2.31)	(–1.73)	(–1.91)
SMB_t	–69.08	–5.24	–10.26	–49.83*	–14.24**	–9.28*
	(–1.59)	(–0.55)	(–1.02)	(–1.73)	(–2.07)	(–1.72)
R^2	16.54%	2.72%	2.80%	41.56%	2.10%	1.82%
Observations	225	237	229	232	237	229

Table 4.A3
Sources of profitability for the attention-based strategy

This table reports results decomposing the sources of profitability for our attention-based trading strategy. Portfolios are constructed using the stock's attention level and the announcement's observed return. The long leg of the strategy involves purchasing a low attention stock (attention measure below 20%) when the announcement released has a positive return. The short leg of the strategy involves and selling short a low attention stock when the announcement released has a negative return. Under the one-day horizon, strategies establish initial positions at the market open on the trading day following the announcement, after having observed the opening price reaction to the announcement, and close-off positions at market close on the same day. Under the ten-day horizon, strategies establish initial positions at the market open two trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. Under the 20-day horizon, strategies establish initial positions at the market open ten trading days after the announcement, after having observed the price reaction to the announcement until that point in time, and close-off positions at market close ten trading days later. All portfolios are value-weighted. $MarketRtn_t$, HML_t, and SMB_t are the market, value, and size factors, respectively, from the Fama-French three-factor model. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =					
	Long leg			Short leg		
Variable	One day	Ten days	20 days	One day	Ten days	20 days
$Intercept$	–0.54*	0.29***	0.35***	0.22	0.18**	0.06
	(–1.88)	(3.43)	(3.65)	(0.56)	(2.34)	(0.78)
$MarketRtn_t$	–46.04	14.64	19.81	–20.18	14.98	14.76
	(–1.15)	(1.29)	(1.53)	(–0.37)	(1.47)	(1.35)
HML_t	–3.29	–3.77	–8.78*	37.76**	3.17	2.57
	(–0.25)	(–0.96)	(–1.97)	(2.07)	(0.90)	(0.68)
SMB_t	–6.34	–9.71	–14.46*	–50.52	6.10	6.06
	(–0.26)	(–1.34)	(–1.76)	(–1.50)	(0.94)	(0.87)
R^2	0.60%	1.96%	3.52%	24.24%	1.09%	1.00%
Observations	225	237	229	232	237	229

Table 4.A4
Lottery-like stocks (alternate identification)

This table reports regression results for post-announcement drift tests conducted using an alternate identification strategy for lottery-like stocks. The unit of observation is an announcement j for stock i. Under this identification strategy, stocks are classified as lottery-like stocks if they have an above-median maximum daily return in the previous month. The first-day return is the percentage cumulative abnormal return in the next trading day after the announcement, excluding the return earned from the time of the announcement to the next trading day's open. The return in the next 19 days is the percentage cumulative abnormal return in the 19 trading days following the next trading day. Announcements are split into positive or negative announcements based on the return in the ten trading days after the announcement. The main independent variable, $LowAttn_{j,i}$, is an indicator variable that equals one if the stock's value for our attention measure is below 20%. Control variables include the stock's average trading volume and volatility, the market return, and the number of announcements on the day. Time fixed effects include day of the week fixed effects and month fixed effects. The sample includes all overnight price sensitive announcements for all stocks listed on the Australian Securities Exchange in the year 2019. t-statistics are reported in parentheses below the coefficient estimates. Significance at the 10%, 5%, and 1% levels is indicated by *, **, and ***, respectively.

	Dependent variable =			
	Positive announcements		Negative announcements	
Variable	First day	Next 19 days	First day	Next 19 days
Panel A: Lottery-like stocks				
$LowAttn_{j,i}$	–0.05	0.88	–1.04*	–1.15
	(–0.10)	(0.50)	(–1.91)	(–0.87)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.32%	2.82%	1.42%	2.48%
Observations	1,906	2,011	2,207	2,352
Panel B: Other stocks				
$LowAttn_{j,i}$	0.77*	3.12***	–0.84***	–0.22
	(1.82)	(3.59)	(–3.43)	(–0.34)
Controls	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
R^2	1.69%	3.90%	2.27%	6.79%
Observations	1,589	1,678	1,576	1,664

Appendix 4.B. Additional figures

Panel A: First day

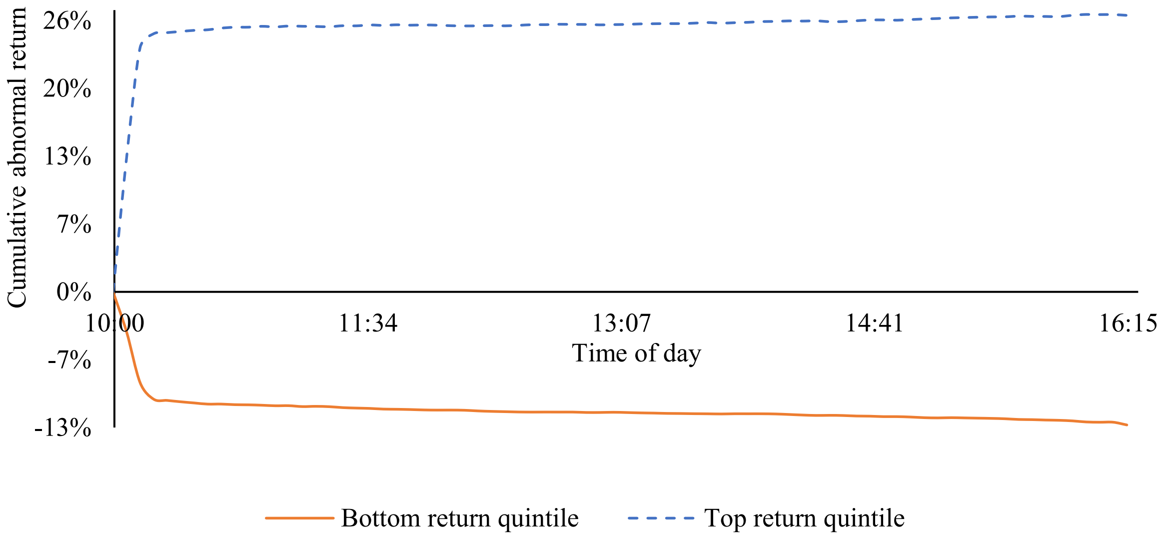

Panel B: First 20 days

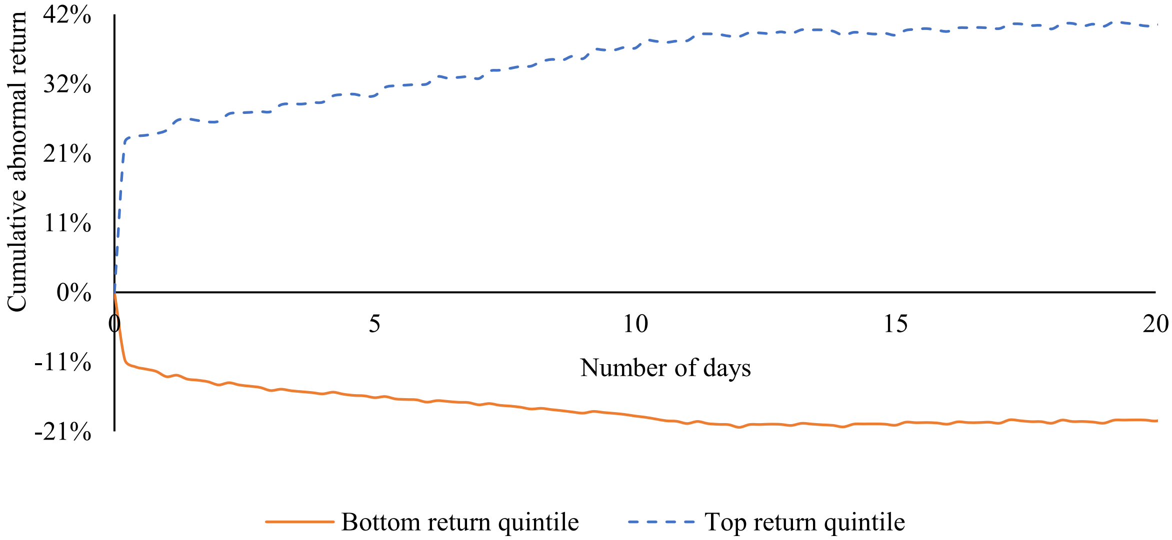

Figure 4.A1. Price reactions split by announcement return.
This figure plots price reactions split into buckets based on the announcement return. Panel A plots cumulative abnormal returns in five-minute intervals in the next trading day after the announcement. Panel B plots cumulative abnormal returns in one-and-a-half-hour intervals across the next 20 trading days after the announcement. Announcements are split into buckets based on the return from the time of the announcement to ten trading days later. High (low) return announcements are in the top-quintile (bottom-quintile).

Panel A: First day

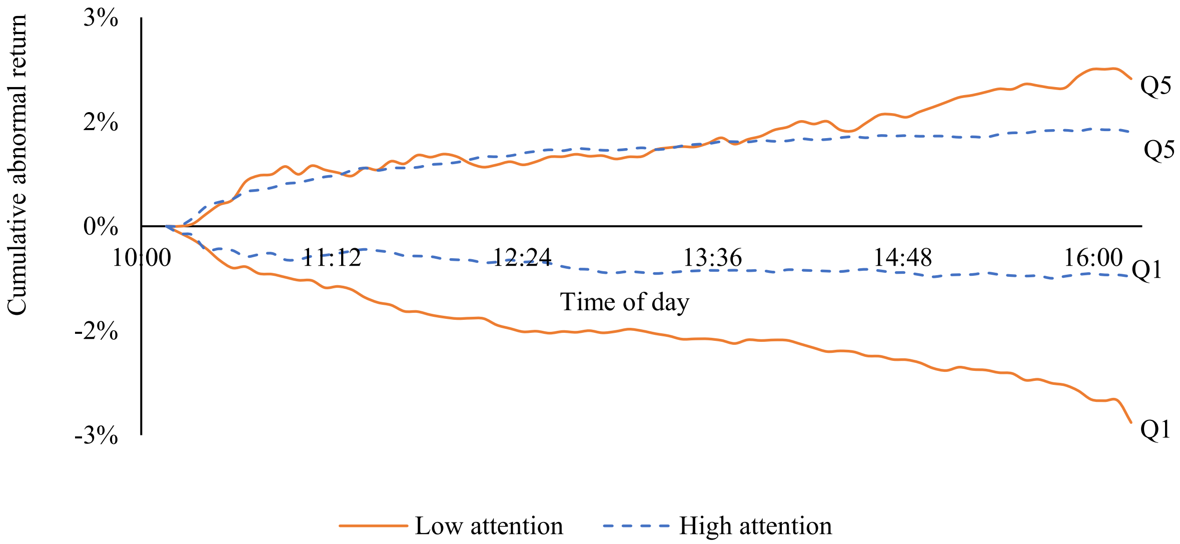

Panel B: First 20 days

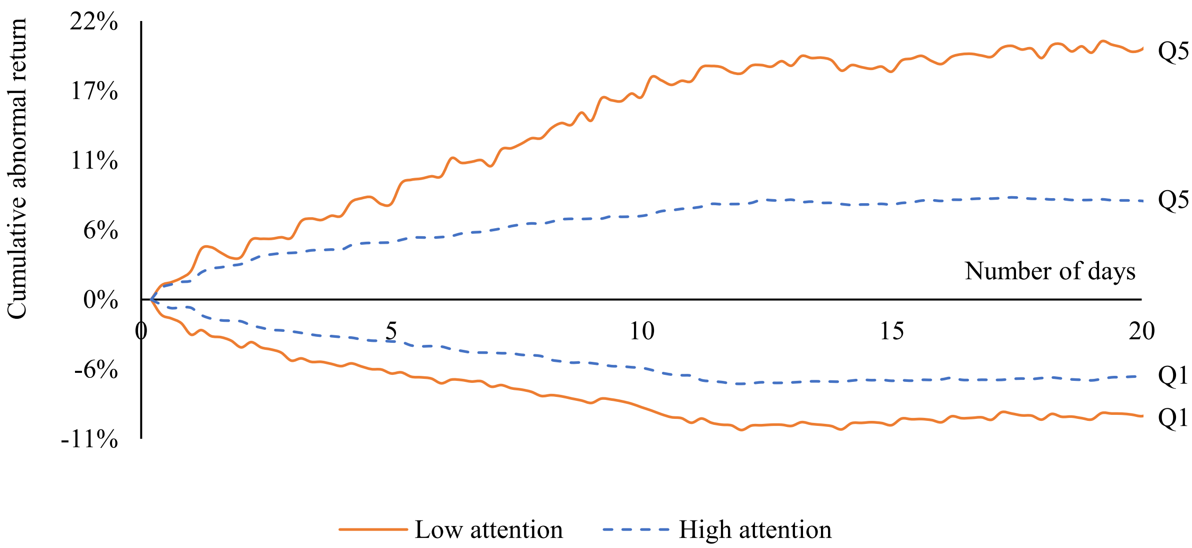

Figure 4.A2. Post-announcement drifts for extreme announcements.
This figure plots the post-announcement drifts for announcements with extreme returns split by the stock's attention level. Panel A plots cumulative abnormal returns in five-minute intervals in the next trading day after the announcement. Panel B plots cumulative abnormal returns in one-and-a-half-hour intervals across the next 20 trading days after the announcement. Both plots exclude the initial price reaction to the announcement, i.e., the return from the time of the announcement to the next trading day's open. Announcements are split into attention buckets based on our attention measure. The attention measure values for low attention announcements are below 20%, while the values for high attention announcements are higher than 50%. The 'Q1' ('Q5') marker indicates that the drift plot is for announcements in the bottom (top) announcement return quintile within the attention bucket. Announcement return is calculated from the time of the announcement to ten trading days later.

Chapter 5: Conclusion

This chapter outlines the four key questions that this thesis addresses and the conclusions that it draws for each question.

i. Why do investors participate in pump-and-dump manipulation games?
ii. How do cryptocurrency pump-and-dumps affect market activity?
iii. How do trading frictions affect investor decision-making?
iv. How does attention to information affect reactions to announcements?

This chapter also discusses avenues for future research.

5.1 Why do investors participate in pump-and-dump manipulation games?

Cryptocurrency markets witness unique pump-and-dump manipulation games that are yet to be observed in any other financial market. In these games, manipulators first gather a following on chat groups on social media platforms, such as Telegram, and openly declare their intentions to conduct pump-and-dump manipulations. At each instance, manipulators inform their followers about the time and venue (exchange) of the intended manipulation well in advance, but do not reveal the coin. They only reveal the coin at the time of the manipulation. We examine these manipulation schemes in detail in Chapter 2 of this thesis.

Using a simple theoretical framework, we demonstrate that rational actors recognize that, without a speed advantage, they face negative expected profits from participation in these schemes since manipulators can build a long position before the pump and unload their position on other participants at higher prices. Hence, these actors do not participate in cryptocurrency pumps unless they know *ex-ante* that they are faster than other non-manipulators. We find that overconfident actors, i.e., those who *believe* they are faster than others, and gamblers, i.e., those who favor the positively skewed payoff structures that a series of pumps can create, participate in these schemes. We confirm this theoretical result empirically as well. Using a sample of 355 pump-and-dump manipulations, we find that participation in these pumps increases after a group's participants become more overconfident and after more gambling money flows into cryptocurrency markets, suggesting that these behavioral mechanisms explain participation in pump-and-dump schemes. Therefore, we conclude that behavioral biases such as overconfidence and gambling preferences drive investors to participate in cryptocurrency pump-and-dump schemes.

5.2 How do cryptocurrency pump-and-dumps affect market activity?

In Chapter 2, we also examine the aggregate coin-level effects of pump-and-dump manipulations. We document significant short-term price, volume, and volatility distortions in the manipulated coins as a result of pump-and-dumps. Pumps generate price spikes of around 65%, on average. However, these spikes are short-lived and reverse within two days of the pump. Pumps also have a substantial effect on trading activity for manipulated coins. Trading volumes are approximately 30–70 times higher on days on which there is a pump. Lastly, we find that pumps have similar effects on volatility too. We find that intraday volatility is between six and eight times higher on days with a pump than on other days. Despite these sharp volatility spikes, cryptocurrency pumps are not the primary driver of the high levels of volatility observed in cryptocurrency markets.

5.3 How do trading frictions affect investor decision-making?

In Chapter 3, using laboratory experiments, we examine how trading frictions, specifically high transaction costs, forced time delays, and cognitive effort tasks, affect investor decision-making quality and performance. Do these frictions induce investors to take a step back to reconsider their trading decisions, thus helping improve their decision quality and performance? Alternatively, are these frictions mere nuisances for investors to ignore on their way to implementing a well-thought-out trading decision, or, worse, do these frictions prevent well-considered trades from ever seeing the light of day? We find that transaction costs and time delays are either ineffectual or counter-productive in improving decision quality. Cognitive tasks, on the other hand, help reduce the underperformance of the worst performing participants by approximately 40%–57%. This performance benefit can almost entirely be attributed to better quality trading decisions than to less overtrading.

We attribute the superiority of the cognitive task friction to its nature. Unlike transaction costs and time delays, in our experiments, cognitive tasks direct the attention of traders to information about asset fundamentals by asking them about their beliefs about the asset's true value. We argue that the most underperforming participants ignore or underweight this information in the frictionless setting, as it is relatively nonsalient. By directing their attention to this information, cognitive tasks help reduce this tendency. Hence, we conclude that trading frictions can be helpful when they induce inattentive investors to consider the information they otherwise neglect.

5.4 How does attention to information affect reactions to announcements?

Chapter 4 of this thesis examines the aggregate effects of investor attention to information about stock fundamentals. For this purpose, we construct a new measure that exclusively captures attention to information releases. This measure relies on the intuition that inattentive investors miss material overnight news events and, thus, do not cancel or update their pending orders after such events. We calculate this measure at the stock level as the proportion of overnight order cancellations and amendments for the stock on days with material overnight news releases. The main motivation to construct this new measure arises from the issue that most existing measures confound attention to information with attention to returns, whereas this new measure allows us to disentangle attention to information from attention to returns.

We use this measure to examine price reactions to material announcements for Australian stocks in 2019. We document strong underreactions that are more pronounced for low attention stocks and positive announcements. The underreactions for low attention stocks are economically meaningful, as a long-short trading strategy that exploits the drifts earns a daily alpha of around 0.46% until ten days after the announcement and around 0.41% from the tenth day to the 20th day after the announcement. These drifts cannot be explained by announcement periodicity, earnings announcements, or the general mispricing of stocks favored by retail investors.

5.5 Future research directions

Speculative bubble games that tend to attract retail investors have recently spread across financial markets. Examples include speculating on "meme stocks," such as GameStop, trading on Elon Musk's tweets, etc. This thesis studies one such game, i.e., pump-and-dump manipulation games in cryptocurrency markets. Future research could examine the effects of other speculative games on markets and test whether the behavioral mechanisms behind pump-and-dump games identified in this thesis explain participation in other speculative games as well.

Future research can also explore whether there is any place for such games in financial markets. Since they constitute a form of gambling, their existence can be justified as a source for people to extract utility from gambling. However, as is the case with other forms of gambling, these games must be regulated closely. Research should examine how best to regulate these games to minimize their harms, in particular, the potential for such games to become a means for sophisticated agents to exploit naïve agents and their potential to foster market inefficiency.

Future research can also test whether real-world nudges designed to resemble the cognitive task frictions in Chapter 3 of this thesis can help reduce investor biases and improve their trading performance. A controlled field experiment in collaboration with a retail broker, for example, would be a valuable way to shed light on this issue.

Lastly, research can attempt to identify other interventions that can help reduce biases and improve performance. Interventions could include steps to impart financial education or other trading frictions that have similar effects to the cognitive task friction, among others.

References

Aggarwal, R., Wu, G., 2006. Stock market manipulations. Journal of Business 79, 1915–1953.

Alicke, M.D., Govorun, O., 2005. The better-than-average effect. In: Alicke M.D., Dunning D.A. and Krueger J.I. (eds.) The Self in Social Judgment. Taylor & Francis Group, New York, pp. 85–108.

Allen, F., Gorton, G., 1992. Stock price manipulation, market microstructure and asymmetric information. European Economic Review 36, 624–630.

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31–56.

Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. Journal of Financial Economics 17, 223–249.

Bali, T.G., Cakici, N., Whitelaw, R.F., 2011. Maxing out: Stocks as lotteries and the cross-section of expected returns. Journal of Financial Economics 99, 427–446.

Bali, T.G., Hirshleifer, D.A., Peng, L., Tang, Y., 2019. Attention, social interaction, and investor attraction to lottery stocks. Unpublished working paper.

Ball, R., Brown, P., 1968. An empirical evaluation of accounting income numbers. Journal of Accounting Research 6, 159–178.

Ball, R., Brown, P., 2019. Ball and Brown (1968) after fifty years. Pacific-Basin Finance Journal 53, 410–431.

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., Lin, S., Odean, T., 2019. Mediating investor attention. Unpublished working paper.

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., 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. Review of Financial Studies 21, 785–818.

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., 2012. A model of casino gambling. Management Science 58, 35–51.

Barberis, N., Greenwood, R., Jin, L., Shleifer, A., 2015. X-CAPM: An extrapolative capital asset pricing model. Journal of Financial Economics 115, 1–24.

Barberis, N., Huang, M., 2008. Stocks as lotteries: The implications of probability weighting for security prices. American Economic Review 98, 2066–2100.

Barberis, N., Thaler, R., 2002. A survey of behavioral finance. Unpublished working paper.

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.

Bazley, W.J., Cronqvist, H., Mormann, M.M., 2021. Visual finance: The pervasive effects of red on investor behavior. Unpublished working paper.

Ben-Rephael, A., Da, Z., Israelsen, R.D., 2017. It depends on where you search: Institutional investor attention and underreaction to news. Review of Financial Studies 30, 3009–3047.

Bernard, V.L., Thomas, J.K., 1990. Evidence that stock prices do not fully reflect the implications of current earnings for future earnings. Journal of Accounting and Economics 13, 305–340.

Biais, B., Hilton, D., Mazurier, K., Pouget, S., 2005. Judgemental overconfidence, self-monitoring, and trading performance in an experimental financial market. Review of Economic Studies 72, 287–312.

Biais, B., Pouget, S., 2000. Microstructure, incentives, and the discovery of equilibrium in experimental financial markets. Unpublished working paper.

Bogan, V., 2008. Stock market participation and the internet. Journal of Financial and Quantitative Analysis 43, 191–211.

Borsboom, C., Janssen, D.-J., Strucks, M., Zeisberger, S., 2021. History matters: How short-term price charts hurt investment performance. Unpublished working paper.

Cheah, E.-T., Fry, J., 2015. Speculative bubbles in bitcoin markets? An empirical investigation into the fundamental value of bitcoin. Economics Letters 130, 32–36.

Chen, J., Tang, G., Yao, J., Zhou, G., 2020. Investor attention and stock returns. Unpublished working paper.

Chen, L.H., Jiang, G.J., Zhu, K.X., 2018. Total attention: The effect of macroeconomic news on market reaction to earnings news. Journal of Banking & Finance 97, 142–156.

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.

Comerton-Forde, C., Putniņš, T.J., 2014. Stock price manipulation: Prevalence and determinants. Review of Finance 18, 23–66.

Cong, L.W., Li, Y., Wang, N., 2018. Tokenomics: Dynamic adoption and valuation. Unpublished working paper.

Cong, L.W., Li, Y., Wang, N., 2020. Token-based platform finance. Unpublished working paper.

Conlisk, J., 1993. The utility of gambling. Journal of Risk and Uncertainty 6, 255–275.

Corgnet, B., Deck, C., Desantis, M., Hampton, K., Kimbrough, E.O., 2019. Reconsidering rational expectations and the aggregation of diverse information in laboratory security markets. Unpublished working paper.

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.

Da Costa, N., Goulart, M., Cupertino, C., Macedo, J., Da Silva, S., 2013. The disposition effect and investor experience. Journal of Banking & Finance 37, 1669–1675.

Da, Z., Engelberg, J., Gao, P., 2011. In search of attention. Journal of Finance 66, 1461–1499.

Daniel, R.F., David, J.R., 1991. Should the law prohibit "manipulation" in financial markets? Harvard Law Review 105, 503–2116.

Deaves, R., Lüders, E., Luo, G.Y., 2008. An experimental test of the impact of overconfidence and gender on trading activity. Review of Finance 13, 555–575.

DellaVigna, S., Pollet, J.M., 2009. Investor inattention and friday earnings announcements. Journal of Finance 64, 709–749.

Dickerson, M., 1984. Compulsive Gamblers. Longman Group, London.

Dixit, A., 1993. The Art of Smooth Pasting. Harwood Academic Publishers, Chur, Switzerland.

Dorn, A.J., Dorn, D., Sengmueller, P., 2014. Trading as gambling. Management Science 61, 2376–2393.

Dorn, D., Huberman, G., 2005. Talk and action: What individual investors say and what they do. Review of Finance 9, 437–481.

Dorn, D., Sengmueller, P., 2009. Trading as entertainment? Management Science 55, 591–603.

Drake, M.S., Roulstone, D.T., Thornock, J.R., 2012. Investor information demand: Evidence from Google searches around earnings announcements. Journal of Accounting Research 50, 1001–1040.

Drake, M.S., Roulstone, D.T., Thornock, J.R., 2015. The determinants and consequences of information acquisition via EDGAR. Contemporary Accounting Research 32, 1128–1161.

Dyhrberg, A.H., Foley, S., Svec, J., 2018. How investible is Bitcoin? Analyzing the liquidity and transaction costs of Bitcoin markets. Economics Letters 171, 140–143.

Easterbrook, F.H., 1986. Monopoly, manipulation, and the regulation of futures markets. Journal of Business 59, 103–127.

Egkolfopoulou, M., Massa, A., Melin, A., 2021. How Robinhood made trading too easy - and maybe even too hard to resist. In: Bloomberg

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.

Fama, E.F., French, K.R., 1992. The cross-section of expected stock returns. Journal of Finance 47, 427–465.

Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. Experimental Economics 10, 171–178.

Fischbacher, U., Hoffmann, G., Schudy, S., 2017. The causal effect of stop-loss and take-gain orders on the disposition effect. Review of Financial Studies 30, 2110–2129.

Focke, F., Ruenzi, S., Ungeheuer, M., 2020. Advertising, attention, and financial markets. Review of Financial Studies 33, 4676–4720.

Fong, K.Y.L., Liu, W.M., 2010. Limit order revisions. Journal of Banking & Finance 34, 1873–1885.

Fox, M.B., Glosten, L.R., Rautherberg, G.V., 2018. Stock market manipulation and its regulation. Yale Journal on Regulation 35, 67–126.

Frydman, C., Rangel, A., 2014. Debiasing the disposition effect by reducing the saliency of information about a stock's purchase price. Journal of Economic Behavior & Organization 107, 541–552.

Frydman, C., Wang, B., 2020. The impact of salience on investor behavior: Evidence from a natural experiment. Journal of Finance 75, 229–276.

Gabaix, X., 2019. Chapter 4 - Behavioral inattention. In: Bernheim B.D., DellaVigna S. and Laibson D. (eds.) Handbook of Behavioral Economics: Applications and Foundations 1. Elsevier, pp. 261–343.

Gandal, N., Hamrick, J.T., Moore, T., Oberman, T., 2018. Price manipulation in the bitcoin ecosystem. Journal of Monetary Economics 95, 86–96.

Gao, X., Lin, T.-C., 2015. Do individual investors treat trading as a fun and exciting gambling activity? Evidence from repeated natural experiments. Review of Financial Studies 28, 2128–2166.

Gargano, A., Rossi, A.G., 2018. Does it pay to pay attention? Review of Financial Studies 31, 4595–4649.

Gervais, S., Kaniel, R., Mingelgrin, D.H., 2001. The high-volume return premium. Journal of Finance 56, 877–919.

Gervais, S., Odean, T., 2001. Learning to be overconfident. Review of Financial Studies 14, 1–27.

Glaser, M., Weber, M., 2007. Overconfidence and trading volume. Geneva Risk and Insurance Review 32, 1–36.

Greenwood, R., Shleifer, A., 2014. Expectations of returns and expected returns. Review of Financial Studies 27, 714–746.

Griffin, J.M., Shams, A., 2018. Is bitcoin really un-tethered? Unpublished working paper.

Grinblatt, M., Keloharju, M., 2009. Sensation seeking, overconfidence, and trading activity. Journal of Finance 64, 549–578.

Hamrick, J.T., Rouhi, F., Mukherjee, A., Feder, A., Gandal, N., Moore, T., Vasek, M., 2020. An examination of the cryptocurrency pump and dump ecosystem. Unpublished working paper.

Handa, P., Schwartz, R.A., 1996. Limit order trading. Journal of Finance 51, 1835–1861.

Hansen, J.H., 2021. Attentive retail investors and a buying pressure after earnings news. Unpublished working paper.

Hasso, T., Müller, D., Pelster, M., Warkulat, S., 2021. Who participated in the GameStop frenzy? Evidence from brokerage accounts. Finance Research Letters

Heimer, R.Z., Imas, A., 2021. Biased by choice: How financial constraints can reduce financial mistakes. Review of Financial Studies

Hirshleifer, D., Lim, S.S., Teoh, S.H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. Journal of Finance 64, 2289–2325.

Hirshleifer, D., Teoh, S.H., 2003. Limited attention, information disclosure, and financial reporting. Journal of Accounting and Economics 36, 337–386.

Hirshleifer, D.A., Sheng, J., 2019. The attention trigger effect: Macro news and efficient processing of micro news. Unpublished working paper.

Hoppe, E.I., Kusterer, D.J., 2011. Behavioral biases and cognitive reflection. Economics Letters 110, 97–100.

Hou, K., Xiong, W., Peng, L., 2010. A tale of two anomalies: The implications of investor attention for price and earnings momentum. Unpublished working paper.

Howell, S., Niessner, M., Yermack, D., 2019. Initial coin offerings: Financing growth with cryptocurrency token sales. Unpublished working paper.

Huang, X., Nekrasov, A., Teoh, S.H., 2018. Headline salience, managerial opportunism, and over- and underreactions to earnings. Accounting Review 93, 231–255.

Huberman, G., Stanzl, W., 2004. Price manipulation and quasi-arbitrage. Econometrica 72, 1247–1275.

Itzkowitz, J., Itzkowitz, J., Rothbort, S., 2016. ABCs of trading: Behavioral biases affect stock turnover and value. Review of Finance 20, 663–692.

Jarrow, R., 1992. Market manipulation, bubbles, corners, and short squeezes. Journal of Financial and Quantitative Analysis 27, 311–336.

Jiang, H., Li, S.Z., Wang, H., 2021. Pervasive underreaction: Evidence from high-frequency data. Journal of Financial Economics 141, 573–599.

Kahneman, D., 2011. Thinking, Fast and Slow. Farrar, Straus and Giroux.

Kalda, A., Loos, B., Previtero, A., Hackethal, A., 2021. Smart(phone) investing? A within investor-time analysis of new technologies and trading behavior. Unpublished working paper.

Kamps, J., Kleinberg, B., 2018. To the moon: Defining and detecting cryptocurrency pump-and-dumps. Crime Science 7, 1–18.

Khomyn, M., Putnins, T.J., 2021. Algos gone wild: What drives the extreme order cancellation rates in modern markets? Journal of Banking & Finance 129

Kocher, M.G., Lucks, K.E., Schindler, D., 2018. Unleashing animal spirits: Self-control and overpricing in experimental asset markets. Review of Financial Studies 32, 2149–2178.

Kottimukkalur, B., 2019. Attention to market information and underreaction to earnings on market moving days. Journal of Financial and Quantitative Analysis 54, 2493–2516.

Kumar, A., 2009. Who gambles in the stock market? Journal of Finance 64, 1889–1933.

Kumar, A., Nguyen, T.H., Putnins, T., 2021. Only gamble in town: Stock market gambling around the world and market efficiency. Unpublished working paper.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53, 1315–1335.

Laudenbach, C., Loos, B., Pirschel, J., Wohlfart, J., 2020. The trading response of individual investors to local bankruptcies. Unpublished working paper.

Lawrence, A., Ryans, J., Sun, E., Laptev, N., 2016. Yahoo finance search and earnings announcements. Unpublished working paper.

Leuz, C., Meyer, S., Muhn, M., Soltes, E., Hackethal, A., 2017. Who falls prey to the wolf of wall street? Investor participation in market manipulation. Unpublished working paper.

Levine, M., 2021. Fraud is no fun without friends. In: Bloomberg

Li, J., Mann, W., 2020. Digital tokens and platform building. Unpublished working paper.

Li, T., Shin, D., Wang, B., 2021. Cryptocurrency pump-and-dump schemes. Unpublished working paper.

Li, Y., Nekrasov, A., Teoh, S.H., 2020. Opportunity knocks but once: delayed disclosure of financial items in earnings announcements and neglect of earnings news. Review of Accounting Studies 25, 159–200.

Liao, L., Wang, Z., Xiang, J., Yan, H., Yang, J., 2020. User interface and firsthand experience in retail investing. Review of Financial Studies

Liu, H., Peng, L., Tang, Y., 2020. Retail attention, institutional attention. Unpublished working paper.

Liu, W.M., 2009. Monitoring and limit order submission risks. Journal of Financial Markets 12, 107–141.

Liu, Y., Tsyvinski, A., 2021. Risks and returns of cryptocurrency. Review of Financial Studies 34, 2689–2727.

Lo, A.W., Mamaysky, H., Wang, J., 2004. Asset prices and trading volume under fixed transactions costs. Journal of Political Economy 112, 1054–1090.

Luo, C., Ravina, E., Sammon, M., Viceira, L.M., 2020. Retail investors' contrarian behavior around news and the momentum effect. Unpublished working paper.

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.

Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. Journal of Economic Theory 26, 17–27.

Mudalige, P., Kalev, P.S., Duong, H.N., 2016. Individual and institutional trading volume around firm-specific announcements. International Journal of Managerial Finance 12, 422–444.

Nursimulu, A.D., Bossaerts, P., 2014. Risk and reward preferences under time pressure. Review of Finance 18, 999–1022.

O'Hara, M., Saar, G., Zhong, Z., 2019. Relative tick size and the trading environment. Review of Asset Pricing Studies 9, 47–90.

Oechssler, J., Roider, A., Schmitz, P.W., 2009. Cognitive abilities and behavioral biases. Journal of Economic Behavior & Organization 72, 147–152.

Osman, M., McLachlan, S., Fenton, N., Neil, M., Löfstedt, R., Meder, B., 2020. Learning from behavioural changes that fail. Trends in Cognitive Sciences 24, 969–980.

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.

Pantzalis, C., Ucar, E., 2014. Religious holidays, investor distraction, and earnings announcement effects. Journal of Banking & Finance 47, 102–117.

Peress, J., 2016. Media coverage and investors' attention to earnings announcements. Unpublished working paper.

Plott, C.R., Sunder, S., 1988. Rational expectations and the aggregation of diverse information in laboratory security markets. Econometrica 56, 1085–1118.

Pouget, S., 2007. Financial market design and bounded rationality: An experiment. Journal of Financial Markets 10, 287–317.

Prasad, M., Bakry, W., Varua, M.E., 2020. Examination of information release on return volatility: A market and sectoral analysis. Heliyon 6

Putnins, T., 2019. An overview of market manipulation. In: Alexander C. and Cumming D. (eds.) The Handbook of Fraud, Misconduct and Manipulation in Financial Markets. John Wiley & Sons, West Sussex, UK.

Raylu, N., Oei, T.P.S., 2002. Pathological gambling: A comprehensive review. Clinical Psychology Review 22, 1009–1066.

Reher, M., Sokolinski, S., 2020. Does Fintech democratize investing? Unpublished working paper.

Seasholes, M.S., Wu, G., 2007. Predictable behavior, profits, and attention. Journal of Empirical Finance 14, 590–610.

Shleifer, A., Vishny, R.W., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.

Statman, M., Thorley, S., Vorkink, K., 2006. Investor overconfidence and trading volume. Review of Financial Studies 19, 1531–1565.

Thaler, R.H., Sunstein, C.R., 2008. Nudge: Improving Decisions About Health, Wealth, and Happiness. Yale University Press, New Haven.

Toplak, M.E., West, R.F., Stanovich, K.E., 2011. The cognitive reflection test as a predictor of performance on heuristics-and-biases tasks. Memory & Cognition 39, 1275–1289.

Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and Uncertainty 5, 297–323.

Van Bommel, J., 2003. Rumors. Journal of Finance 58, 1499–1520.

Venkatraman, V., Wittenbraker, J., 2020. Disrupting dual systems: A dynamic decision-making framework for human behavior. Unpublished working paper.

Vila, J.L., 1989. Simple games of market manipulation. Economics Letters 29, 21–26.

Weber, M., Camerer, C.F., 1998. The disposition effect in securities trading: An experimental analysis. Journal of Economic Behavior & Organization 33, 167–184.

Wursthorn, M., Choi, E., 2020. Does Robinhood make it too easy to trade? From free stocks to confetti. In: Wall Street Journal

Xu, J., Livshits, B., 2019. The anatomy of a cryptocurrency pump-and-dump scheme. Proceedings of the 28th USENIX Conference on Security Symposium, 1609–1625.

Yamamoto, R., 2014. An empirical analysis of non-execution and picking-off risks on the Tokyo Stock Exchange. Journal of Empirical Finance 29, 369–383.

Zuckerman, M., 1994. Behavioral Expressions and Biosocial Bases of Sensation Seeking. Cambridge University Press, New York, NY, US.