

Essays in Financial Economics

Josep Gisbert Rodriguez

TESI DOCTORAL UPF / Year 2022

THESIS SUPERVISOR

Dr. Gaël Le Mens and Dr. José-Luis Peydró

Department of Economics and Business



Embrace Reality and Deal with It.

Ray Dalio

Per la meua mare Lourdes i pel meu fill Jan.

Acknowledgements

This dissertation would not have been possible without the support of many generous and brilliant people. First of all, I thank my main advisor José Luis Peydró for his excellent support and guidance. He saw promise in my work when I was new to the field, gave me great feedback, and provided brilliant and close support during my job market. I would also like to thank Gaël Le Mens, my co-advisor, who has given me his best form of support during these years. I also thank Robin Hogarth, who has been a pillar of my success in the PhD. Robin has given me critical support at several moments and has become an example of the type of scholar and human being that I would like to become.

I thank Dmitry Kuvshinov, Björn Richter, and Victoria Vanasco. Despite none of them being official thesis co-advisor, they have shown interest in my research and have committed to helping me improve it. I thank Dmitry for his advice on the job market and for his support in helping me polish the job market paper. I thank Björn for supporting me in the job market, for his guidance while I was producing the job market paper, and for his later support in helping me polish it. Finally, I thank Victoria for writing me a recommendation letter for the job market and giving me insightful feedback on my academic presentations.

I thank my coauthors, Albert Banal-Estañol, Daniel Navarro-Martínez, and Jordi Quoidbach, plus the non-coauthor Professor Gert Cornelissen, for their support guidance, and patience during my first steps in the program. I especially thank Daniel and Albert for their supportive role throughout my doctoral studies. They have been people who I always felt I could count on and have been extremely generous when I have needed his help. I also thank Christian Brownlees, Miguel Espinosa, Jaume García, Libertad González, and Albrecht Glitz for their excellent advice during the PhD and their generosity with their time and kindness.

I thank my wife, Mònica, for all her incredible support. I feel truly fortunate to share my life with her. I also thank my son Jan, who during this first year of life has been an incredible source of happiness and blessings for the whole family. I also thank the rest of my close family, Lourdes, Damián, Damià, Maria, Carla, and Marta, for their support and for reminding me of what matters most in life. I also thank Úrsula, Anton, Margarita, Agustí, Aram, Vinicius, Antonio Méndez, Cristian, Jaume, Josep, Rafel, and the rest of my close friends and relatives that supported me during

these last seven years and made my life so much better.

I thank my friends and closest colleagues Miriam Artiles, Sebastian Diz, Milena Djourelova, Sebastian Ellingsen, Andrea Fabiani, Mario Giarda, Angelo Gutierrez, Rahil Hosseini, Juan Imbet, Adil Ismailov, Ilja Kantorovich, Karolis Liaudinskas, Helena Palumbo, Guillem Pros Rius, David Puig, Xinghua Wang, Sebastián Willis, Shohei Yamamoto, and Lang Zhou. They have not only been an immense source of happiness but have also actively helped with their insights and time at several moments in the PhD. They certainly are the best part of the PhD.

I thank all the talented people I met during these last years in the PhD program, Ana Costa, Federica Daniele, Estève Giraud, Fernando Fernández, Martí Guasch, Lukas Hoesch, Adam Lee, Sergio Pirla, Ana Rodríguez, Damian Romero, Niko Schoell, Paul Soto, André Souza, Thomas Woiczzyk, and Dijana Zejcirovic, as well as to the faculty who supported me, José Apesteguía, Mircea Epure, Xavier Freixas, Javier Gomez, Gábor Lugosi, Gergely Neu, Andrea Polo, and Mikhail Spektor. I am particularly thankful to Bernd Theilen and Carolina Manzano faculty at Universitat Rovira i Virgili, who supported me along with my academic career, to Luis Díaz Marcos faculty at CUNEF, who helped me in the Master's application, and to Vicente Bermejo and Anna Bayona faculty at ESADE Business School who counted on me to be their teaching assistant during the PhD.

I thank the administrative team at UPF and the Barcelona SE, especially Marta Araque, Laura Agustí, Mariona Novoa, and Sara Fernandez, for helping me navigate the bureaucratic sea of paperwork and providing their best advice.

Last but not least, I thank all the students at UPF, BSM, Barcelona SE, and ESADE that I had the opportunity to be a TA for. They have sparked a passion for teaching, and watching them grow is a constant source of joy.

Josep Gisbert Rodriguez, 2022

Abstract

This dissertation studies financial technology adoption, learning from financial information, and inter-temporal decisions. It contains three chapters. Chapter one examines the impact of bank branch closures on fintech adoption. I find that branch closures lead to a persistent increase in fintech lending. Chapter two investigates the effect of feedback on learning from financial information. The results suggest that investors are more accurate in their learning when investors are in a setting in which they have access to financial information because of endogenous choice. Chapter three studies the effect of mindfulness on inter-temporal decisions. Results show that mindfulness does not affect inter-temporal decisions.

Resum

Aquesta tesi estudia l'adopció de tecnologia financera, l'aprenentatge sobre informació financera, i les decisions intertemporals. Conté tres capítols. El primer capítol examina l'impacte del tancament de sucursals bancàries en l'adopció de tecnologia en el sector financer. Trobo que el tancament de sucursals comporta un augment persistent dels préstecs a empreses "fintech". El segon capítol investiga l'efecte de la retroalimentació d'informació en l'aprenentatge d'informació financera. Els resultats suggereixen que els inversors són més precisos en el seu aprenentatge quan es troben en un entorn en què tenen accés a la informació financera a causa d'una elecció endògena. El tercer capítol estudia l'efecte del "mindfulness" en les decisions intertemporals. Els resultats mostren que el "mindfulness" no afecta les decisions intertemporals.

Contents

List of figures	xv
List of tables	xix
1 FINTECH, BANK BRANCH CLOSINGS, AND MORTGAGE MARKETS	5
1.1 Introduction	5
1.2 Data	11
1.3 Fintech Classification	13
1.4 Empirical Strategy	15
1.5 Results	20
1.6 Conclusion	34
2 HOW SELECTIVE ACCESS TO FINANCIAL INFORMATION AFFECTS HOW INVESTORS LEARN	39
2.1 Introduction	40
2.2 Literature Review	45
2.2.1 Information acquisition	45
2.2.2 Information processing	47
2.2.3 Information acquisition and processing effects	49
2.3 Experimental Design	50
2.4 Empirical Findings	55
2.4.1 Main Results	55
2.4.2 Alternative Explanations	65

2.4.3	Learning in the last trial	67
2.4.4	Choices	68
2.5	Discussion and Implications	72
2.6	Conclusion	81
3	SHORTSIGHTED MEDITATORS? THE EFFECTS OF MIND- FULNESS ON INTER-TEMPORAL CHOICE	83
3.1	Introduction	84
3.2	Inter-temporal Decisions	89
3.3	Experiment 1	91
3.3.1	Method	91
3.3.2	Results and discussion	92
3.4	Experiment 2	99
3.4.1	Method	99
3.4.2	Results and discussion	100
3.5	Experiment 3	102
3.5.1	Method	102
3.5.2	Results and discussion	105
3.6	Experiment 4	116
3.6.1	Method	116
3.6.2	Results and discussion	117
3.7	Conclusion	124
A	APPENDIX - FINTECH, BANK BRANCH CLOSINGS, AND MORT- GAGE MARKETS	125
A.1	Spatial Variation	125
A.2	Summary statistics	136
B	APPENDIX - HOW SELECTIVE ACCESS TO FINANCIAL IN- FORMATION AFFECTS HOW INVESTORS LEARN	137
B.1	Participant Instructions Full feedback Condition	137

B.2	Participant Instructions Selective feedback Condition	140
B.3	Quiz for the Full feedback Condition	143
B.4	Quiz for the Selective feedback Condition	147
B.5	Objective Bayesian Posterior Beliefs	151
B.6	Measures of Financial Literacy	152
B.7	Payoff Sequences	154
B.8	Type of Stock Sequences	158
C	INADVERTENTLY SUPER-FARSIGHTED MEDITATORS: THE EFFECTS OF MINDFULNESS ON INTER-TEMPORAL CHOICE	163
C.1	Experiment 1	163
	C.1.1 Tasks	163
C.2	Experiment 2	167
	C.2.1 Tasks	167
C.3	Experiment 3	168
	C.3.1 Common Instructions	168
	C.3.2 Specific Mindfulness Condition Instructions	169
	C.3.3 Specific Mind-wandering Condition Instructions	169
	C.3.4 Intertemporal choice task items	171
C.4	Experiment 3	172
	C.4.1 Common information sheet & Task booklet	172

List of Figures

1.1	Exposed and Control Tracts Selection	8
1.2	Fintech Classification Method	14
1.3	Number of Fintech Mortgages in Sample Tracts	16
1.4	Exposure to Consolidation and Bank Branch Closings	22
1.5	Exposure to Consolidation and Bank Branch Levels	23
2.1	Payoffs and probabilities	50
2.2	Prior subjective estimate elicitation	52
2.3	Example of full feedback or partial feedback condition trial in the case the participant chooses the stock	53
2.4	Example of full feedback or partial feedback condition trial if a participant chose the bond	54
2.5	Average subjective estimates for the probability that the stock is the good one, as a function of the Bayesian probability . . .	56
2.6	Average subjective estimates for the probability that the stock is the good one, as a function of the Bayesian probability calculated as if participants received full feedback about the payoffs of the stock in each trial	57
2.7	Proportion of stock choices per trial for the selective feed- back and full feedback condition by type of stock and domain	72
3.1	Proportion of distant choices per condition	93
3.2	Proportion of distant choices per condition in the items in- volving decisions in the present	94
3.3	Proportion of distant choices per condition in the items in- volving decisions between alternatives in the future	96

3.4	Proportion of time consistent choices per condition	99
3.5	Proportion of money saved per condition in tasks 1 and 2 . .	101
3.6	Proportion of participants that chose the savings alternative per condition in task 3	101
3.7	Proportion of present biased participants per condition . . .	106
3.8	Average individual δ and β per condition	108
3.9	Average number of distant choices selected per condition . .	113
3.10	Histogram of δ per condition	114
3.11	Histogram of β per condition	115
3.12	Proportion of distant choices per wave	117
3.13	Proportion of distant choices per condition in the items in- volving decisions in the present	119
3.14	Proportion of distant choices per condition in the items in- volving decisions between alternatives in the future	121
3.15	Proportion of time consistent choices per wave	122
A.1	Spatial variation in fintech mortgage originations market share 2008 and 2009	126
A.2	Spatial variation in fintech mortgage originations market share 2010 and 2011	127
A.3	Spatial variation in fintech mortgage originations market share 2012 and 2013	128
A.4	Spatial variation in fintech mortgage originations market share 2014 and 2015	129
A.5	Spatial variation in fintech mortgage originations market share	130
A.6	Geographic distribution of sample states	131
A.7	Geographic distribution of bank branches in 2009	132
A.8	Geographic distribution of bank branches per 100,000 peo- ple in 2009	132
A.9	Geographic distribution of US population 2000 US Census .	133
A.10	Geographic distribution of branch closings since 2009 up until 2016	134

A.11 Geographic distribution of branch closings per 100,000 people since 2009 up until 2016	134
A.12 Geographic distribution of the share of branch closings with respect to 2009 bank branch levels in 2016	135

List of Tables

1.1	Sample Fintech Lenders	15
1.2	Sample Mergers	17
1.3	Mergers Summary Statistics	18
1.4	Tract Summary Statistics: Exposed vs. Control	21
1.5	First Stage Estimates	24
1.6	Reduced Form Estimates	25
1.7	Reduced Form Estimates	26
1.8	Second Stage Estimates - Mortgage Originations	27
1.9	Second Stage Estimates - Mortgage Originations Volume	28
1.10	Second Stage Dynamic Estimates - Mortgage Originations	29
1.11	Second Stage Dynamic Estimates - Mortgage Originations Volume	30
1.12	Second Stage Splits: Demographic	32
1.13	Representativeness of the Merger Sample	35
1.14	Complier Characteristics	36
2.1	Experimental Design	51
2.2	Differences in Probability Estimation Errors in the Selective feedback and Full feedback Conditions	59
2.3	[Differences in Probability Estimation Errors against a Full Feedback Bayesian benchmark in the Selective feedback and Full feedback Conditions	60
2.4	Mean Probability Estimation Error Updates per Condition and their Difference after Observing a High or Low Payoff	61

2.5	Mean Subjective Probability Updates per Condition and their Difference after Observing a High or Low Payoff	63
2.6	Mean Subjective Probability Updates in the Full feedback Condition when Stock Chosen or Sampled and their Difference after Observing a High or Low Payoff	65
2.7	Mean Subjective Probability Updates and Mean Probability Estimation Error Updates in the First Trial of the Selective and Full feedback Conditions by Asset Selection	66
2.8	Differences in Probability Estimation Errors in the Selective feedback and Full feedback Conditions in Trial 6	69
2.9	Differences in Probability Estimation Errors against a Full Feedback Bayesian benchmark in the Selective feedback and Full feedback Conditions in Trial 6	70
3.1	Distant Choices and Time-consistent Choices in the Mindfulness, Control, and Mind-wandering Conditions	95
3.2	Distant Choices Involving the Present Moment and not Involving the Present Moment in the Mindfulness, Control, and Mind-wandering Conditions	97
3.3	Amount of Savings in the Mindfulness, Control, and Mind-wandering Conditions	103
3.4	Amount of Savings in the Mindfulness, Control, and Mind-wandering Conditions	104
3.5	Present Biased Participants in the Mindfulness, Control, and Mind-wandering Conditions	107
3.6	Present Bias Parameter of Participants in the Mindfulness, Control, and Mind-wandering Conditions	109
3.7	Time-consistent Discounting Parameter of Participants in the Mindfulness, Control, and Mind-wandering Conditions	111
3.8	Average Number of Distant Choices in the Mindfulness, Control, and Mind-wandering Conditions	112
3.9	Distant and Time-consistent Choices in the Pre-treatment and Post-treatment waves	118
3.10	Distant Choices Involving the Present Moment and not Involving the Present Moment in the Pre-treatment and Post-treatment waves	120

3.11 Mindfulness Trait and Meditation Experience Effects on Pre
and Post-treatment Number of Distant Choices 123

A.1 Tract Summary Statistics: Exposed vs. All other 136

Introduction

This doctoral dissertation studies financial technology adoption, learning from financial information, and inter-temporal choices. It contains three chapters. Chapter one, using empirical methods, examines the impact of bank branch closings on lending by technologically intensive financial companies. Chapter two, using experimental methods, investigates the effect of feedback on learning from financial information. Chapter three, using experimental methods, studies the effect of mindfulness on inter-temporal decisions.

This dissertation contains several empirical and methodological contributions. In the first chapter, I assemble a novel data-set and show that although bank branch closings do not affect the overall supply of mortgages in the market, as previously demonstrated by Nguyen (2019). This fact is obscuring an important change. I show that closings increase fintech mortgage supply and reduce bank supply, thus significantly changing the lender mix in the mortgage market. More precisely, I show that closings cause a compositional change in the lenders of the mortgage market. The change occurs between banks, which reduce their mortgage supply, to fintech, which increases it. This change in lenders leads to an additional compositional change in the borrower mix. Since fintech lenders target a different set of borrowers than banks, the resulting public that obtains credit also changes. Crucially I then show that the effects of closings lead to bank credit rationing for information-intensive borrowers, such as minorities or poor individuals, thus hampering financial inclusion of these traditionally vulnerable groups of the population.

In the second chapter, we create an experimental design that, for the first time, follows a belief-based approach that allows us to measure the two sources of error when learning from financial information in a bandit setting (the cognitive error caused by incorrect processing of information, and the sampling error, caused by using a smaller sample of information). This design also allows to precisely measure the size of the probability

updates and test whether this probability update affects choice behavior accordingly at each point of the dynamic learning process. This set of features provides a new tool to analyze human learning. In this chapter, we also have relevant empirical results. We show an opposite empirical result to that of Kuhnen (2015) whose work shows that investors in a full feedback environment, who face outcomes only in the loss domain, make on average higher probability errors than people facing the same environment in the gain domain. In this chapter, we show an opposite result, reference point losses—perceived losses relative to the guaranteed payment provided by the riskless alternative—in the selective feedback environment in the gain domain, and not explicit losses, trigger superior adaptive learning by participants relative to those of people in the same domain observing the same outcome. We point to the post-decision regret literature to explain this finding. Finally, in this chapter, we provide new insights that can explain why access to foregone feedback has a great influence on choice but not on maximization. Here we show that the better processing of information in selective feedback environments can help overcome the loss of information, and thus if the more accurate beliefs translate into more optimal choices, they can help maximization.

The third chapter is the first study that uses an intensive and very popular onsite eight-week course in mindfulness, the mindfulness-based stress reduction program (MBSR), designed by Jon Kabat-Zinn as a treatment before making inter-temporal decisions¹. This compares to the most robust mindfulness manipulation up to now, which consisted of only a 50-minute prerecorded video on a mindfulness workshop focused on mindful eating (Hendrickson and Rasmussen (2013, 2017)). Our study also measures the trait level mindfulness of participants and collects a measure of the experience in the mindfulness practice, which is not common in other studies. Our study uses three different tasks to measure inter-temporal choice behaviors. This compares to only one task to measure the effects of mindfulness on an inter-temporal choice which was the best available in this type of study so far and which provided less varied insights about the relationship between mindfulness and inter-temporal choice (Hendrickson and Rasmussen (2013, 2017)). Our study is also novel in the following three ways: we measure inter-temporal choice behavior both in a laboratory setting and outside, combine manipulation and training-based studies, and present a mixture of purely hypothetical and potentially real tasks. Finally, our study not only studies the effects of a mindfulness-based intervention but also tests a particular mechanism that could explain the

¹in one of the studies

effects of mindfulness. We test whether the potential effect of mindfulness on inter-temporal decisions is derived from choices that involve decisions between the present and a distant moment in time or is also present in choices between two delayed moments in time.

The main result and contribution in each chapter can be summarized as follows: Chapter 1 studies whether bank branch closures affect fintech mortgage lending in the U.S. using data for 1999–2016 period. In this chapter, I use a quasi-experimental research design based on Home Mortgage Disclosure Act data. Home Mortgage Disclosure Act data is the most comprehensive source of publicly available information on the U.S. mortgage market. Studying the effect of branch closures on local mortgage markets poses an empirical challenge. Banks tend to close branches in areas where current or expected profitability is low. Since profitability is also related to tract characteristics that affect mortgage credit, a simple study comparing areas in which a branch closes and those in which there is no closing will produce a biased estimate of the effects of branch closures. To overcome this problem, I use an instrument based on within-county, tract-level variation in exposure to post-merger branch consolidation. The instrument exploits closures due to the merger of two large national banks that operated branches in close geographical proximity. I find that branch closures lead to a persistent increase in fintech lending. Fintech mortgages grow by a total of 8% relative to non-closure tracts in the nine years that follow a closing, while bank mortgage lending falls by 44%, off an annual baseline of 340 mortgages.

Chapter 2 compares learning in two common settings in financial markets. One in which investors can observe the outcome of an investment alternative only if they invest in it, and another one in which they always can observe the outcome—even if they do not invest in it. We provide empirical evidence that investors' beliefs are, on average, 5% closer to the objective Bayesian beliefs given the observed information when investors are in a setting in which they have access to the financial information because of endogenous choice. Then we are able to describe the mechanism that explains our findings. We show that the endogenous creation of the sample of information triggers different cognitive processes. These alternative processes cause better information processing and are of enough magnitude to help overcome the effect of sampling errors.

Chapter 3 studies the effects of mindfulness on inter-temporal decision-making. We run four studies, three in the lab and one in the field. The studies in the lab use a between-subjects design with three conditions. In

the mindfulness condition, participants listen to a 15-minute audio mindfulness breathing exercise. In the mind-wandering condition, participants listened to 15-minute audio that repeatedly instructed participants to think of whatever came to mind. In the control condition, participants did not listen to any audio. After all, participants made either 42 or 4 choices between receiving smaller cash amounts earlier and larger cash amounts later or, in the third lab study, responded to hypothetical but realistic scenarios in which inter-temporal decisions needed to be made. In the field, experiment participants completed an eight-week mindfulness training course from the largest provider of onsite mindfulness courses in Spain. We then collected participants' selections into four choices between receiving smaller cash amounts earlier and larger cash amounts later. Overall, we show that mindfulness does not affect inter-temporal decisions.

Chapter 1

FINTECH, BANK BRANCH CLOSINGS, AND MORTGAGE MARKETS

This paper studies whether bank branch closures affect fintech mortgage lending in the U.S. using data for the 1999–2016 period. To generate plausibly exogenous variation in the incidence of closings, I use an instrument based on within-county, tract-level variation in exposure to post-merger branch consolidation. I find that branch closures lead to a persistent increase in fintech lending. Fintech mortgages grow by a total of 8% relative to non-closure tracts in the nine years that follow a closing, while bank mortgage lending falls by 44%, off an annual baseline of 340 mortgages. Fintech mortgage growth is driven by wealthier areas and areas with relatively smaller populations of women, seniors, and minorities.

1.1 Introduction

Since the onset of the Great Recession in 2007, the U.S. consumer lending market has gone through two major disruptions. First, a persistent wave of bank branch closings, which represents a 12% overall decline in bank branches from the pre-recession peak¹. Second, the rise and consolidation in the market of a new type of lender: fintech companies —characterized

¹Closings amount to a total of 10,631 for the 2008-2020 period. Source: Federal Deposit Insurance Corporation.

by their intensive use of technology in the provision of financial services. Fintech's share of the mortgage market has expanded remarkably in the last two decades, from a zero market share in the mid-2000s to nowadays representing more than 12% of the overall US mortgage market².

Given the significance of these two disruptions for credit markets, the literature has scrutinized their effects and has unveiled important consequences. For instance, bank branch closures have been linked to an increase in the interest paid by borrowers and firms (Fuster et al., 2019; Bonfim et al., 2020), a reduction of local credit supply (Nguyen, 2019), the creation of "banking deserts"³ (Morgan et al., 2016), and to severe contractions of credit for low income and minority groups (Nguyen, 2019). While fintech has been linked to improvements in the efficiency of credit markets (Fuster et al., 2019), discrimination against minorities in the interest rate charged Bartlett et al. (2021), or regulatory arbitrage (basing its growth on the lower regulation that fintech has relative to banks), which could lead to financial instability. In this paper, I study whether these two disruptions, which have been viewed in isolation, are related. To my knowledge, this is the first study about the effect of bank branch closings on fintech adoption and its consequences for credit markets. I study their effect in the U.S. mortgage market, a \$11.05 trillion market that is of vital interest for the lives of most American citizens, and that is key for the strength, profitability, and stability of the overall U.S. financial system as the 2007–2009 financial crisis portrayed.

To study the effect of branch closures on local mortgage fintech adoption, I use a quasi-experimental research design based on Home Mortgage Disclosure Act (HMDA) data. HMDA data is the most comprehensive source of publicly available information on the U.S. mortgage market⁴. Studying the effect of branch closures on local mortgage markets poses an empirical challenge. Banks tend to close branches in areas where current or expected profitability is low, and since profitability is also related to tract characteristics that also affect mortgage credit, a simple study comparing areas in which a branch closes and those in which there is no closing will produce a biased estimate of the effects of branch closures⁵. To overcome this problem, I use an instrument based on within-county, tract-level variation in exposure to post-merger branch consolidation following Nguyen (2019).

²According to the 2015 Home Mortgage Disclosure Act data.

³Areas with no bank branch within 10 miles.

⁴It captures 90 percent of lending activity measured by loan volume.

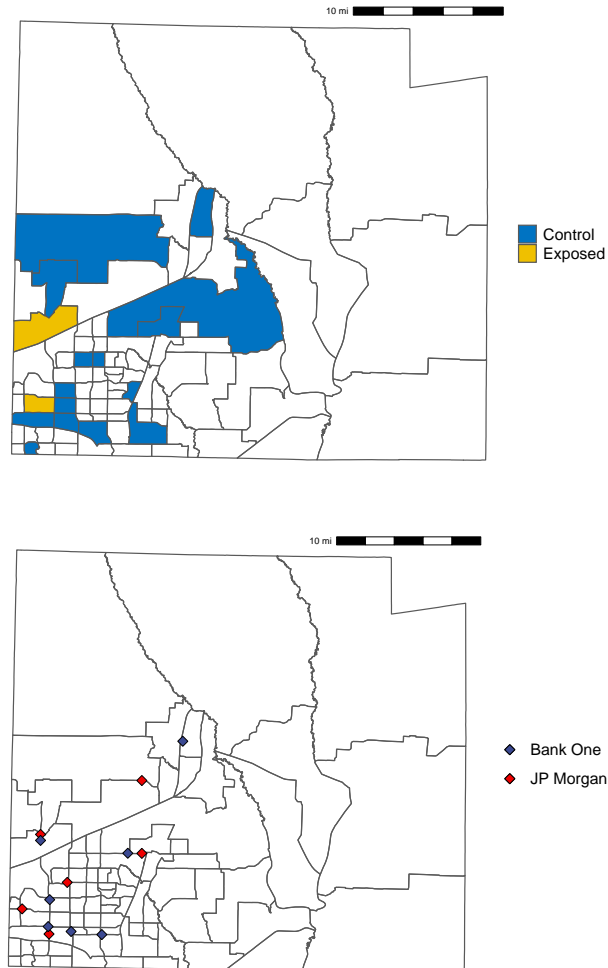
⁵U.S. Census tracts are small, relatively permanent statistical subdivisions of a county.

The instrument exploits closures due to the merger of two large, national banks that operated branches in close geographical proximity. Figure 1.1 illustrates this identification strategy for a sample merger and county. This framework compares the pre-merger and post-merger level of lending in exposed tracts (those that had branches from both merging banks prior to the merger) relative to a set of control tracts located in the same county and with branches belonging to at least two large non-merging banks. The identifying assumption is that the decision to merge is exogenous to tract characteristics that also determine credit. To warrant this, I only include in my sample mergers between very large banks. That is, banks that pre-merger had at least \$50 billion in assets, which puts them roughly at the 1% asset size distribution of U.S. banks. The business size of exposed tracts represents such a minimal share of the participating bank's profits, that the plausibility of the decision to merge being linked to these area characteristics is extremely low. The ultimate goal of this empirical framework is to compare tracts that, ex-ante, were equally likely to have been exposed to a large bank merger. To evaluate the effects of closures on fintech lenders, I then classify all institutions that report HMDA data into three types: banks, shadow banks, or fintech using established methods in the literature (Buchak et al., 2018; Fuster et al., 2019; Jagtiani et al., 2019).

As both banks and fintech lenders are competitors in the mortgage market, the first question that I explore is whether banks are fostering fintech's growth by deserting certain areas through branch closures. To test this, I estimate the causal effect of branch closures on fintech mortgage supply. I then proceed to test whether the decision of banks to downsize their branch network affected the overall supply of credit offered by banks in affected areas. To test this, I estimate the causal effect of branch closures on bank mortgage supply. Finally, I test whether the effects of branch closures are especially severe for certain population groups.

My main results are as follows: first, I show that branch closings persistently increase local mortgage fintech adoption. The cumulative increase in fintech mortgages relative to control tracts over the nine years following a closing is 27 mortgages, or 8% of a baseline of 340 mortgages in the pre-merger year (I use bank mortgages as a baseline since there are no fintech mortgages in any sample tracts in all pre-merger years). All in all, these effects suggest that the use of technology by fintech lenders may allow them to better cater to the needs of the "unbranched" borrowers. Second, I show that branch closings persistently reduce the mortgage supply offered by banks. The cumulative decrease in bank mortgages relative to control tracts over the nine years following a closing is 152 mortgages,

Figure 1.1: Exposed and Control Tracts Selection



Notes: The map on the top shows control tracts (blue) and exposed tracts (yellow) in Collin County, Texas for the sample merger between JP Morgan and Bank One, that was approved in 2004. The map on the bottom shows the network of the merger banks in the year before merger approval.

or 44% of the annual baseline. This finding highlights the critical role of branches, which, even in the fintech era, seem to help reduce information asymmetries that represent frictions for the correct functioning of the mortgage market.

Finally, I provide evidence that suggests that the effects of branch closings differ significantly across population groups. After closings, poor, high minority, low share of female, and older populated areas suffer a

steeper reduction of bank mortgages than the average area in the study. However, wealthy, low minority, low female, and younger populated areas are driving the increase of fintech mortgages. These findings suggest that the change in the lender mix has significant consequences for the type of borrower that can get a mortgage post-closing and, even more importantly, for those that no longer can get one. These findings also further raise concerns about the role of financial technologies, such as algorithm-based screening and big data use for loan approval decisions in the financial sector, and its repercussions for the financial inclusion of vulnerable groups of the population.

This paper contributes to three strands of the banking literature. First, I provide novel evidence on the effects of branch closings in credit markets. As the paper's main contribution, I show that previous tests that documented that there was no overall effect of closings on the mortgage supply, but only a reduction of credit for small business loans (Nguyen, 2019) and in pricing to firms (Bonfim et al., 2020) were obscuring a crucial change. This study shows that closings cause a significant change in the lender mix that provides the mortgage supply. More precisely, I show that closings cause a compositional change in the lenders of the mortgage market. This change occurs between banks, who reduce their mortgage supply, to fintech, which increases it. Moreover, my findings suggest that this change in lenders leads to an additional compositional change in the borrower mix. Since fintech lenders target a different set of borrowers than banks, the resulting public that obtains credit also changes. These findings are evidence of the critical role of bank branches. Even in this fintech era characterized by the intensive use of technology and new channels to reach customers, bank branches still fulfill an important role in mortgage markets. In this paper, I show that the branch role of facilitating bank-branch-borrower relationships and the transmission of soft information is important for firm credit outcomes and household ones despite new technological advancements in the mortgage market.

Second, I contribute to identifying the elements that explain the rapid growth of fintech in credit markets. An optimistic explanation might be that new technology and improved methods could be the primary source of fintech growth. Precisely these new technologies could be the factor that allows fintech lenders to produce better services or to lend more cheaply to borrowers. Supporting this view, Fuster et al. (2019) find that fintech lenders reduce frictions in the mortgage origination process, such as capacity constraints, slow processing times, and lower than optimal refinancing. Additionally, Buchak et al. (2018) show that fintech uses new

technology to provide credit and attribute part of their growth to this use. However, a more critical or pessimistic explanation is that fintech lenders are engaging in regulatory arbitrage. Fintech lenders may be benefiting from their lower regulation relative to banks to capture part of their market share in credit markets. For instance, Buchak et al. (2018) shows that, indeed, fintech lenders are filling the gap left by banks, but that they have done so in segments where regulatory burden has risen substantially for banks and relying almost exclusively on explicit and implicit government guarantees. de Roure et al. (2021) show that stricter capital requirements fostered credit reallocation from banks to peer-to-peer fintech lending providers in the German consumer credit market after 2010 while Irani et al. (2020) show a similar effect for the U.S. corporate loan market between banks and non-banks. In this paper, I provide evidence in support of the former view. I make comparisons between areas with the same regulation and show that after a branch closing, fintech lenders, and not other banks, are the ones that capture the “deserted” borrowers. These findings suggest that fintech lenders provide better products or offer cheaper mortgages and do not base their growth exclusively on regulatory arbitrage.

Third, my analysis also contributes to the literature on the effects of fintech and bank branch closings on financial inclusion⁶. Generally, the literature on the effects of closings casts a negative light on its impact on financial inclusion. Nguyen (2019) shows that bank branch closings disproportionately reduce access to credit to information-intensive borrowers, such as minorities and low-income individuals. Morgan et al. (2016) show that, especially in low-income areas, credit can be rationed after branch closings due to the creation of banking deserts⁷. In contrast, some of the literature on fintech paints a positive picture about its effects on financial inclusion. Bartlett et al. (2021) show that fintech lenders reduce disparities in the interest rate charged and do not discriminate in mortgage application rejection for Latinx/African-Americans compared to risk-equivalent borrowers. However, there is also a more pessimistic view about the effects of fintech on financial inclusion. Work by Barocas and Selbst (2016) portrays the negative effects of algorithm decision-making for the financial inclusion of large segments of the U.S. population. Buchak et al. (2018) show that fintech lenders are less likely to serve less creditworthy FHA

⁶Financial inclusion implies that individuals and businesses have access to useful and affordable financial products and services that meet their needs delivered in a responsible and sustainable way. Source: The World Bank.

⁷Banking desert is defined as a relatively homogeneous area or neighborhood containing about 4,000 people with no branches within ten miles of its center.

borrowers and higher unemployment geographies. Bartlett et al. (2021) show that fintech providers discriminate in the interest rate charged to Latinx/African-Americans. Finally, Fuster et al. (2021) recently tested the role of machine learning models in financial inclusion. They found that Black and Hispanic borrowers are disproportionately less likely to gain from the introduction of machine learning. In this paper, I show that the effects of closings indeed lead to bank credit rationing for information-intensive borrowers, such as minorities or poor individuals. Additionally, I show that, although fintech partly fills the gap left by banks, it does so, targeting richer, non-minority, younger, and male individuals, leaving other groups suffering the brunt of the decrease in credit.

The rest of the paper is organized as follows. Section 1.2 details data sources. Section 1.3 explains the fintech classification methodology. Section 1.4 discusses the details of the empirical strategy used to identify the causal effect of interest. Section 1.5 analyzes the effect of branch closings on consumer lending markets. Section 1.6 concludes.

1.2 Data

All sources of data of the paper are at the census tract level. Census tracts are defined by the U.S. Census Bureau as small, relatively permanent statistical subdivisions of a county designed to contain about 4,000 inhabitants⁸, therefore, their size varies depending on their population density. After each census, the borders of some tracts are slightly updated⁹. In this paper, I use 2000 census borders¹⁰.

To analyze the impact of branch closings on local lending, I obtain local mortgage lending data from the Home Mortgage Disclosure Act (HMDA) datasets published by the Federal Financial Institutions Examination Council (FFIEC). I use data for the 1999-2016 period¹¹. HMDA data is the most comprehensive source of publicly available information on the U.S. mortgage market. The HMDA was enacted by Congress in 1975 and was im-

⁸With a minimum of 1,200 inhabitants and a maximum of 8,000.

⁹Census tracts are split, merged or untouched, depending on population change, and small boundary corrections are sometimes allowed as well.

¹⁰For variables reported using other U.S. census borders (1990 or 2010 census) I use a set of relationship files provided by the U.S. census that show how the different census geographies relate to each other and allow to merge geographic entities over time.

¹¹The rest of the data sources are for the same 1999-2016 period and at the census tract level unless I specify that it is not.

plemented by the Federal Reserve Board. The Board requires lending institutions to report public loan data using a remarkably stable reporting criteria¹². HMDA data are at the loan application level, and include information about the census tract in which the borrower is located¹³, the amount of the application, whether the mortgage has been approved or denied, reason for denial (if denied), the name of the chartering agency of the institution, the purpose of the mortgage (i.e., home purchase / improvement / refinancing) and applicant characteristics such as gender, race or income. Crucially, HMDA data is based on the borrowers' location and not on that of the lender. That HMDA data is based on the borrowers' location allows me to estimate the impact of a branch closing on mortgage supply to borrowers in the same tract. I keep only mortgages classified by HMDA regulation as conventional. Therefore, I drop mortgages originated by the Federal Housing Administration, the Veterans Administration, and the Farm Service or Rural Housing Service. I then aggregate the remaining mortgages to create a yearly census tract-level measure of mortgage originations. Finally, I winsorize this measure at the 1 percent level.

To construct the exposure instrument, I first obtain the annual listing of all bank branches belonging to institutions insured by the Federal Deposit Insurance Corporation (FDIC). The listing is provided by the FDIC Summary of Deposits (SOD). The SOD is the annual survey of branch office deposits as of June 30 for all FDIC-insured institutions, including insured U.S. branches of foreign banks. All institutions with branch offices are required to submit the survey. Only institutions with only a main office are exempt. Apart from the branch deposits information, the SOD contains

¹²According to the 2021 reporting criteria published by the FFIEC under HMDA banks, savings associations, or credit unions that: have at least \$48 million in assets, have a home or branch office located in a metropolitan statistical area, originated at least one home purchase loan or refinancing of a home purchase loan, are federally insured or federally regulated or are insured, guaranteed or supplemented by a Federal agency or intended for sale to the Federal National Mortgage Association or the Federal Home Loan Mortgage Corporation, and meet or exceed either the closed-end mortgage loan-volume threshold or the open-end line of credit loan-volume threshold (effective January 1, 2018 through December 31, 2021, an institution that originated at least 25 closed-end mortgage loans, or originated at least 500 open-end lines of credit or exceeds the loan volume threshold) in each of the two preceding calendar years. For-profit mortgage-lending institutions other than banks, savings associations, or credit unions are subject to HMDA regulation if the institution had a home or branch office in a metropolitan statistical area and meets or exceeds either the previously mentioned closed-end mortgage or credit loan-volume threshold in each of the two preceding calendar years.

¹³Not based in the location of the lending financial institution.

data of the branch's address, the GPS coordinates of the location of the branch (only from 2008 onwards), and information related to the institution that owns the branch. I use GIS software to locate bank branches for which the GPS coordinates are available in the SOD. For branches whose GPS coordinates are not available, I use a combination of Google Maps Geocoding API (to find a branch's GPS coordinates using its address information) and GIS software to locate its census tract. The percentage of non-located branches in the complete SOD data for the 1999-2016 period is 0.7%. The percentage of unmapped observations in 1999 is 1.7% and declines to 0.03% in 2016.

To complete the construction of the exposure instrument, I obtain data on bank branch closings and merger activity from FDIC's API. To locate the branch closings, I use the same method as per the bank branch's location. The percentage of non-located branch closings for the period of interest is 2.4%. The percentage of unmapped observations in 1999 is 5.9% and declines to 0.07% in 2016. To obtain bank mergers approved by regulators in the period of interest, I downloaded from FDIC's API all mergers with effective dates of inclusion in the 1999-2020 period¹⁴. I obtain federal approval dates by searching for the corresponding order of approval documents released by the FED, a press release by participants in the merger, or other regulators' press release notes. I gather merger announcement dates by searching the announcement news or press releases in FACTIVE¹⁵ database. Information about merger participants' asset size is obtained from FDIC's statistics on depository institutions' fourth-quarter financial data report.

Finally, I gather tract-level demographic characteristics from the 2000 U.S. census. The rest of the data sources are for the 1999-2016 period.

1.3 Fintech Classification

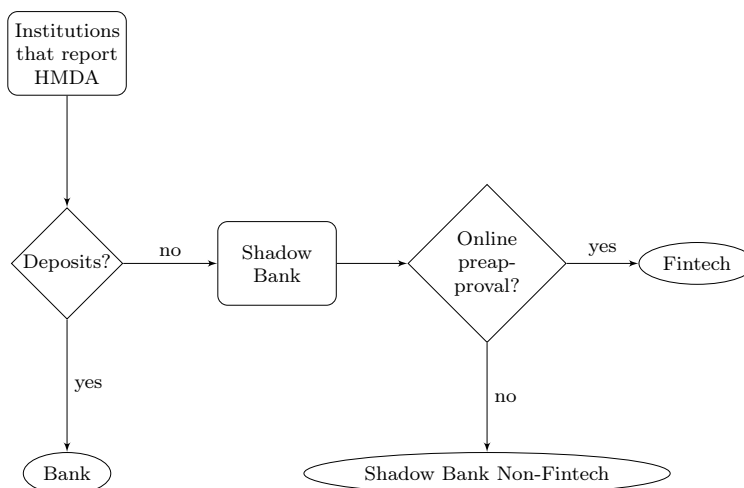
In this paper, I classify all institutions that report HMDA data in the 1999-2016 period in three types: banks, shadow banks, or fintech. To gather the list of institutions, I first collect the annual HMDA Reporter Panels (RP) for the period of interest. The RP includes information that identifies each institution (i.e., name, location...) and a variable that codes each type of

¹⁴Since approval dates are always later than the effective dates, I look for the period 1999-2020

¹⁵Global news database of more than 33,000 sources owned by Dow Jones & Company.

lending institution¹⁶. Second, following Buchak et al. (2018) institution classification methodology, I classify as banks all depository institutions, and as shadow banks the rest of institutions¹⁷.

Figure 1.2: Fintech Classification Method



Finally, to distinguish between shadow banks and fintech, I follow Jagtiani et al. (2019) classification, which is a mix of the fintech classifications by Buchak et al. (2018) and Fuster et al. (2019) plus two more recent lenders. I consider a shadow bank as fintech if it is classified as such by any of the three papers. Buchak et al. (2018) and Fuster et al. (2019) consider an institution as fintech if it allows for mortgage preapproval or full approval without the borrower having to communicate directly with a loan officer or a broker. Additionally, as in Jagtiani et al. (2019), I classify as fintech two institutions that started reporting HMDA data in 2016 based on their growing volume and media recognition as the best online mortgage providers. These institutions were not classified as fintech by Buchak et al. (2018) and Fuster et al. (2019) because these focused on earlier years and larger lenders. The list of included institutions is the following: AmeriSave Mortgage, Better Mortgage, CashCall Inc., Everett Financial,

¹⁶As either a: Bank Saving Institution, Credit Union, Mortgage Banking Subsidiary (MBS) of Commercial Bank, MBS of Bank Holding Company or Service Corporation, Department of Housing and Urban Development, Private Mortgage Institution Corporation, or Affiliate

¹⁷Therefore Banks, Saving Institutions, Credit Unions and Mortgage Banking Subsidiaries of Commercial Banks or Bank Holding Companies or Services Corporations are classified in the paper as banks.

Guaranteed Rate, loanDepot, Movement Mortgage, SoFi, and Quicken. An illustration of this classification method is displayed in figure 1.2, and a summary of fintech listings in the three cited papers is shown in table 1.1.

Table 1.1: Sample Fintech Lenders

Lender	Fintech Starting Year		
	Buchak et al. (2018)	Fuster et al. (2019)	Jagtiani et al. (2019)
AmeriSave Mortgage	2008		2008
Better Mortgage	not included		2016
CashCall Inc.	2008		2008
Everett Financial		2016	2016
Guaranteed Rate	2008	2010	2010
loanDepot		2016	2016
Movement Mortgage	2013	2014	2014
Sofi	not included	not included	2016
Quicken	2000	2010	2010

Notes: This table shows three different listings of sample fintech providers and their starting year as fintech institution type. This paper follows Jagtiani et al. (2019) fintech classification

Source: Buschak et al. (2018), Fuster et al. (2019), & Jagtiani et al.(2019).

1.4 Empirical Strategy

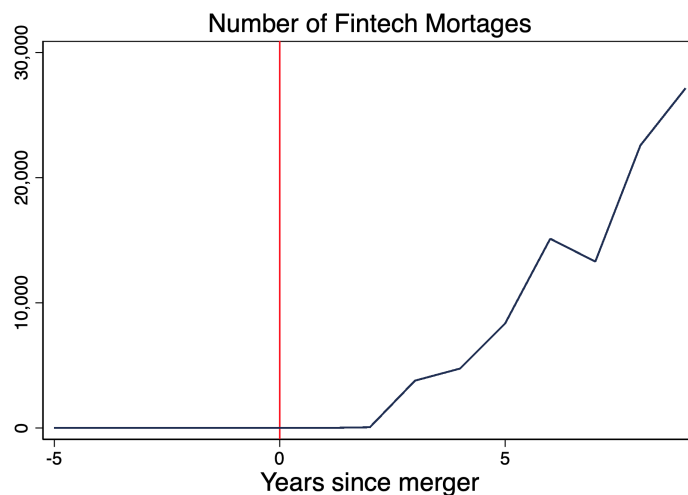
Empirical challenge The relationship of interest is the effect of branch closings on the mortgage supply of the different lender types. Estimating this relationship poses an empirical challenge: factors related to branch closings that are also associated with local economic characteristics that can correlate with mortgage lending and fintech adoption. For example, branch closings tend to occur in bank low profitability areas. These areas generally have experienced credit demand shocks, which are also related to the local level of lending.

Instrument As a solution for this empirical challenge, I use as an instrument for branch closings the exposure to post-bank-merger consolidation, following Nguyen (2019) who pioneered this approach. Specifically, I exploit the plausibly exogenous variation in the incidence of branch closings

that follows a bank merger. Bank mergers tend to be followed by a period when the resulting merged institution engages in branch closings. These closings tend to be focused on areas in which the two preceding branch networks overlap. This both (i) increases the probability of branch closings in those areas and (ii) decreases the likelihood that the decision to close a branch is based on local economic characteristics.

The key identifying assumption is that tract-level exposure to bank mergers is as good as randomly assigned. Or stated differently, that the decision of any sample pairs of banks to merge is not more likely in an exposed than in a control tract. This assumption will not hold if the decision to merge is made because of specific tract characteristics of tracts in which branch networks overlap. For instance, if particular tract economic factors motivate the decision to merge that are also associated with lending.

Figure 1.3: Number of Fintech Mortgages in Sample Tracts



Notes: This figure plots the number of fintech mortgages in sample tracts for every year relative to the merger approval. $\tau = 0$ is the year the merger was approved by federal regulators. Only years for which the panel is balanced are displayed.

The plausibility of the decision to merge being exogenous is necessary for the internal validity of the exposure instrument. However, the concern that post-merger choices regarding the specific branches to close are endogenous does not represent a threat to the internal validity of the instrument. If the decision of any sample pairs of banks to merge is not more likely in exposed tracts relative to control tracts, the internal validity of the instrument holds. However, if the post-merger election of branch

closings is indeed related to specific tract characteristics, this will threaten the identification strategy’s external validity. I will further discuss external validity later on in section 1.5.

To further address this challenge to identification (e.g., that the decision to merge is related to economic tract factors), I select only mergers between large banks. The final sample is formed by banks that at least have \$50 billion in assets in the year before the merger approval. This size situates them roughly in the top 1 percent of the size distribution of US banks. I keep only mergers between large banks because they are unlikely to be motivated by tract-level economic conditions. They tend to be driven by other factors such as increased market power, the generation of complementary business, or expansion into new markets. Although cost savings derived from consolidation may also be considered. It is doubtful that the merger decision is based on tract-level economic considerations.

Table 1.2: Sample Mergers

Buyer	Target	Year approved
Bank of America	Fleet National Bank	2004
JPMorgan Chase Bank	Bank One	2004
Wachovia Bank	SouthTrust Bank	2004
Regions Bank	AmSouth Bank	2006
Bank of America	LaSalle Bank	2007

Notes: This table shows the 5 mergers included in the sample and the year they were approved by federal regulators.

Moreover, the business size that exposed tracts represent is relatively very small. In a similar study by Nguyen (2019) with a lower bank pre-merger assets threshold for inclusion of \$10 billion, the author estimated that the median percentage of the buyer (target) banks’ deposits held in exposed tracts before the merger is only 1.4 percent (3.5 percent). Therefore, it is improbable that potential gains or savings in those tracts motivate the decision to merge.

Another potential threat to the validity of the identified effects of this study is that of reverse causality. If mergers, and thus post-merger closings, occur in areas in which fintech adoption is already higher, that could be biasing the estimated effects of branch closings. To bring evidence against this possibility, in figure 1.3 I plot the annual number of fintech mortgages for all census tracts in my sample. The plot shows that fintech mortgages only significantly increase for all tracts of all mergers in my sample after the second post-merger year. Therefore the threat of reverse

causation does not seem to represent a problem in this study since fintech mortgage levels in all tracts for all pre-merger and merger approval years are zero and could not drive the decision to merge.

I list sample mergers in Table 1.2. Included mergers are mergers that were approved during the 2000s (but before the financial crisis)¹⁸, in which the two participating institutions had two or more branches, had \$50 billion in pre-merger assets and had overlapping networks in at least one census tract.

Table 1.3 shows sample merger summary statistics. Sample banks are very large. The median buyer (target) bank holds \$565 billion (\$73 billion) pre-merger assets, has 2,554 branches (723), and operates in 15 different states (9), while the median for all US banks is \$302 million, four branches, and one state of operation, respectively.

Table 1.3: Mergers Summary Statistics

Panel A: Buyer			
	Median	Min	Max
Total assets (\$bn)	565	81	1,200
Branches	2,554	594	5,723
States of operation	15	4	31
Counties of operation	401	47	698
Panel B: Target			
	Median	Min	Max
Total assets (\$bn)	73	52	257
Branches	723	138	1,563
States of operation	9	1	13
Counties of operation	134	7	188

Notes: The table displays summary statistics for the 5 buyer and 5 target banks in the merger sample. All variables are as of the year in which the intention to merge was announced.

I define exposed tracts of any merger as those in which the buyer and the target bank had branches in the year preceding the merger. Figure 1.1 shows an illustration of the census tract allocation for the 2004 merger between JP Morgan and Bank One for a particular County —Collin County, TX. Both illustrations show a map of the county with census tracts delineated. However, the bottom map shows the geographical distribution of branches of JP Morgan (red diamonds), Bank One (blue diamonds). And

¹⁸Mergers occurred in 2008 and later years were excluded from the sample

the top map shows exposed (yellow) and control (blue) tract allocation. A census tract of Collin County is defined as exposed if it has both a branch from JP Morgan and Bank One in the pre-merger year.

Column 3 in Table 1.4 shows that exposed tracts are similar to controls tracts in many relevant dimensions. However, several differences still exist. Exposed tracts have a slightly lower percentage of college-educated population, a higher percentage of population below the poverty level, a lower median family income, fewer bank branches, a lower bank branch growth, and more bank mortgage originations than control tracts. Therefore, I first control for these differences and simultaneously use a difference-in-differences (DiD) framework to compare lending outcomes in exposed and control tracts for the same county, pre-merger, and post-merger. I also allow for time-varying trends based on pre-merger tract characteristics. Specifically, to identify the local average treatment effect of bank branch closings on mortgage supply of fintech, banks, and shadow-banks, I estimate 2SLS equations of the following form:

$$Mortgage_{tcmly} = \sigma_t + (\mu_y \times \nu_c) + \mathbf{X}_t \lambda_y + \rho_{Post} (\overbrace{Post_{my} \times Closing_{tcm}}) + \eta_{tcmly} \quad (1.1)$$

$$(\overbrace{Post_{my} \times Closing_{tcm}}) = \alpha_t + (\gamma_y \times \kappa_c) + \mathbf{X}_t \beta_y + \delta_{Post} (Post_{my} \times Expose_{tcm}) + \epsilon_{tcm} \quad (1.2)$$

where $Mortgage_{tcmly}$ is mortgage lending to borrowers in tract t , county c , of merger m in year y , by lender type l ; $Post_{my}$ is a dummy equal to 1 if year t occurs after merger m is approved by federal regulators; $Closing_{tcm}$ is an indicator equal to one if a bank branch closes in tract t , county c , after merger m ; $Expose_{tcm}$ is an indicator equal to one if tract t is an exposed tract of merger m ; σ_t are tract fixed effects; $(\mu_y \times \nu_c)$ are county-by-year fixed effects; \mathbf{X}_t is a vector of pre-merger census tract characteristics whose effects are allowed to vary by year. The pre-merger tract characteristics in vector \mathbf{X}_t are population, population density, percentage of the minority population, percentage of college-educated population, median family income, percentage of population 65 years old and over, percentage of rural population, percentage of unemployed population, percentage of population below poverty level, number of bank branches in the year before the merger is approved, and average annual bank branch growth in the two years preceding merger approval. Standard errors are clustered at the county level. Finally, the coefficient of interest δ_{Post} measures the post-closing mean shift in the level of lending for lender type l .

The first-stage equation has also a difference-in-differences specification, where the excluded instrument for the potentially endogenous interaction between the post-merger indicator $Post_{my}$ and the closing indicator

$Closing_{tcm}$, is the interaction between $Post_{my}$ and the exposure to a merger indicator $Expose_{tcm}$; α_t are tract fixed effects; $(\gamma_y \times \kappa_c)$ are county-by-year fixed effects, and all other variables as previously defined.

In this DiD framework, the identifying assumption is that outcomes of exposed tracts would have similar trends to those of control tracts in the absence of exposure to a merger.

To allow for the analysis of pre-trends in the data, I estimate year-by-year DiD and present the results in event study plots. The primary specification throughout the rest of the paper is:

$$y_{tcmyl} = \sigma_t + (\mu_y \times \nu_c) + X_t \lambda_y + \delta_\tau (D_{my}^\tau \times Expose_{tcm}) + \eta_{tcmyl} \quad (1.3)$$

where y_{tcmyl} is an outcome for tract t , county c , of merger m in year y , by lender type l ; D_{my}^τ is a dummy equal to 1 if year y is τ years after merger m is approved by federal regulators; and all other variables as previously defined. The range of τ is between -8 and 12 . Standard errors are clustered at the county level. Finally, coefficient δ_τ measures the difference, conditional on controls, in outcome y between exposed and control tracts, τ years before or after a merger.

1.5 Results

First stage Figure 1.4 shows that in the five pre-merger years for which we have a balanced panel, exposed tracts are not more likely than control tracts to experience a branch closing. However, there was a sharp increase in the number of closings one year after the merger and a moderate one two years after. Then, for the rest of the post-merger years in which we have a balanced panel (9 years post-merger), differences go back to zero until year nine after merger approval (except in the fifth year where closings are lower in exposed tracts). Column 1 of table 1.5 shows the corresponding point estimates. Since the maximum of closings for each sample tract is generally one, point estimates can be interpreted as the change in relative probability of closing in exposed tracts relative to control tracts τ years since the merger was approved. I find that the probability of closing in the first year after a merger is 27 percentage points higher in exposed tracts relative to control tracts and 33 percent combining the effects of the first and second year. For the rest of the post-merger years, for which we have a full panel, I find that the relative probability change is zero or neg-

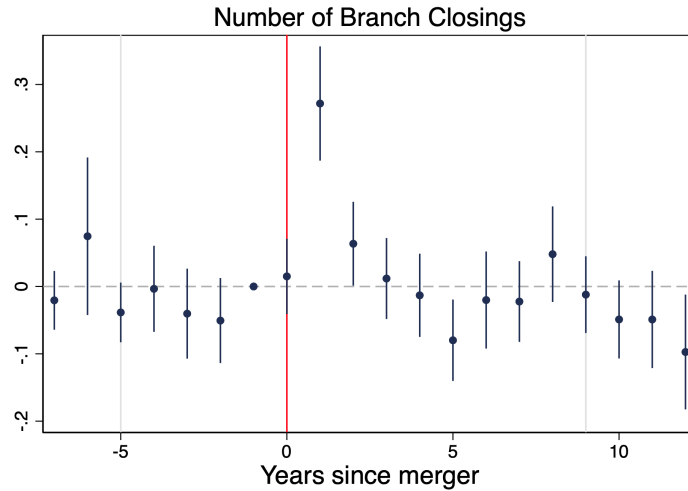
Table 1.4: Tract Summary Statistics: Exposed vs. Control

Variable	(1) Exposed	(2) Control	(3) <i>p</i> -value on difference
Population	5,752 [3,766]	5,539 [2,872]	0.519
Population density	2,498 [3,544]	6,309 [17,296]	0.290
Percent minority	28.1 [23.2]	29.3 [23.7]	0.935
Percent college educated	59.1 [19.1]	60.9 [18.6]	0.034
Percent poverty level	13.3 [11.6]	10.2 [9.1]	0.019
Percent rural population	5.8 [16.1]	3.2 [11.7]	0.024
Percent population 65 and over	16.3 [11.9]	16.0 [13.3]	0.581
Percent unemployed	5.9 [6.4]	5.0 [5.5]	0.184
Median income (000s)	56.43 [27.64]	61.1 [29.9]	0.052
Percent MSA median income	117.1 [49.6]	119.0 [55.5]	0.074
Total branches	6.9 [4.5]	4.0 [2.3]	0.000
Branch growth	0.041 [0.114]	0.076 [0.198]	0.003
Bank mortgages	339.5 [480.2]	306.4 [315.1]	0.012
Shadow bank mortgages	119.6 [189.8]	116.7 [142.3]	0.134
Fintech mortgages	0 [0]	0 [0]	n.a.
Observations	418	1,982	

Notes: Standard deviations are in brackets. Column 3 reports the *p*-value for the difference between columns 1 and 2. Here *p*-values are obtained from a regression of tract characteristics on an indicator for being an exposed tract and county fixed effects. Population density is per square mile. Percent MSA median income is the ratio of tract median income to MSA median income. Growth rates are the average annual growth rates over the two years preceding the merger approval. All demographic variables are as of the 2000 census. Credit variables are as of the year before federal merger approval.

ative in the fifth year after the merger.

Figure 1.4: Exposure to Consolidation and Bank Branch Closings

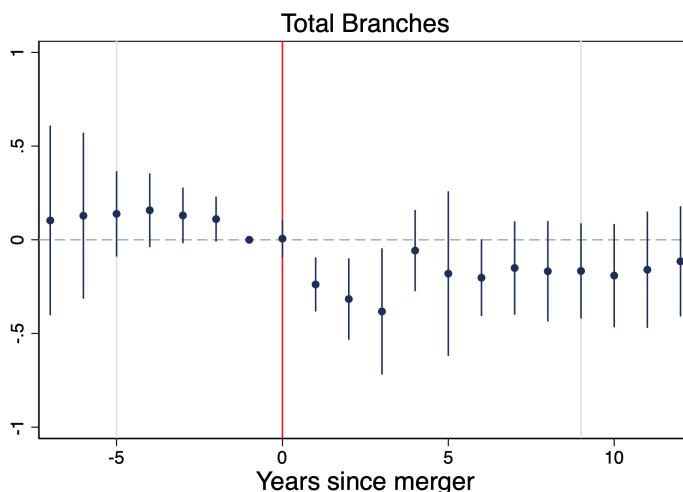


Notes: This figure plots the first-stage relationship between exposure to consolidation and the incidence of branch closings, obtained from estimating equation 1.3. The bars show 95 percent confidence intervals, $\tau = 0$ is the year the merger was approved by federal regulators, and all coefficients are normalized relative to $\tau = -1$. Robust standard errors are clustered at the county level.

To assess whether the decline in branch closings translates into a reduction in the number of branches, I follow the same method as described above to plot figure 1.5. This figure shows that there is also no evidence of pre-trends. It also shows that the number of branches in exposed tracts relative to control tracts is lower until the third year after a merger approval. This concentration of the effects of a merger in the first three years is consistent with the previous literature (Garmaise and Moskowitz, 2006; Nguyen, 2019). Column 2 of table 1.5 shows the corresponding point estimates. These show a significant decrease in the number of branches in exposed relative to control tracts between the first and third post-merger years, which ranges between 0.24 and 0.38 branches. In the years that follow, the coefficients remain negative, although the statistical significance is lost.

Intention-to-treat effects I then focus on the reduced form intention-to-treat effect of the exposure to post-merger branch consolidation on mortgage lending. In table 1.6, I estimate equation 1.3 for the number of mortgages by the three lender types and their total lending. Column 1 shows

Figure 1.5: Exposure to Consolidation and Bank Branch Levels



Notes: This figure plots the first-stage relationship between exposure to consolidation and the total number of branches, obtained from estimating equation 1.3. The bars show 95 percent confidence intervals, $\tau = 0$ is the year the merger was approved by federal regulators, and all coefficients are normalized relative to $\tau = -1$. Robust standard errors are clustered at the county level.

a significant negative reduced form effect, at the 90% level or above, for banks from the fourth year post-merger until the seventh year. The magnitude of the accumulated significant effects amounts to a reduction with respect to the bank baseline of 12% of mortgages for the overall period. The sign and significance of the effects are mirrored by those in column 1 of table 1.7 that shows lending in thousands of U.S. dollars. Here the negative reduced form effect for banks in dollar volume in the same period amounts to a reduction with respect to the bank baseline of 17% in dollar volume of bank mortgages. Column 2 and 3 of tables 1.6 and 1.7 show no significant persistent effects for shadow banks and fintech lenders. Column 4 of table 1.6 also shows no significant effects in the total number of mortgage supply. However, column 4 of table 1.7 shows that for the total dollar volume of mortgages, exposure to a merger translates into an overall significant reduction at the 90% of the total mortgages for the post-merger years fourth, sixth, and seventh, that amounts roughly to an 11% decrease with respect to the total mortgages dollar volume baseline.

Table 1.5: First Stage Estimates

	(1) Number branch closings	(2) Total branches
$\delta_{<-1}$	-0.024 (0.021)	0.124 (0.092)
δ_0	0.015 (0.028)	0.009 (0.050)
δ_1	0.271*** (0.043)	-0.236*** (0.074)
δ_2	0.063** (0.031)	-0.314*** (0.109)
δ_3	0.011 (0.030)	-0.381** (0.171)
δ_4	-0.014 (0.031)	-0.056 (0.110)
δ_5	-0.081*** (0.031)	-0.178 (0.223)
δ_6	-0.021 (0.036)	-0.200* (0.104)
δ_7	-0.023 (0.030)	-0.148 (0.126)
δ_8	0.047 (0.036)	-0.165 (0.136)
δ_9	-0.013 (0.029)	-0.165 (0.128)
$\delta_{>9}$	-0.061** (0.026)	-0.165 (0.138)
Tract FEs	Yes	Yes
County \times Year FEs	Yes	Yes
Observations	42,462	42,462
R ²	0.25	0.85
Baseline mean	0.3	6.9

Notes: This table shows estimates of equation 1.3. All coefficients are normalized relative to $\tau = -1$, and $\tau = 0$ is the year in which the merger was approved by federal regulators. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Local average treatment effects To identify the local average treatment effect of bank branch closings on mortgage lending, the main goal of this study, I estimate the 2SLS equations 1.1 and 1.2. Here, the coefficient of interest is the second stage effect of the incidence of bank branch consolidation on mortgage lending for each lender type. Columns 1 and 3 of table

Table 1.6: Reduced Form Estimates

	Reduced Form			
	(1) Bank Mortgages	(2) Shadow bank Mortgages	(3) Fintech Mortgages	(4) Total Mortgages
$\delta_{<-1}$	-2.426 (6.942)	-1.406 (2.445)	-0.333 (0.209)	-4.164 (8.995)
δ_0	4.281 (4.903)	-0.900 (2.489)	-0.0728 (0.151)	3.308 (6.932)
δ_1	4.727 (7.779)	5.004 (5.225)	-0.0385 (0.148)	9.693 (12.48)
δ_2	2.276 (9.347)	4.172 (5.720)	-0.0261 (0.154)	6.421 (14.15)
δ_3	-4.428 (7.143)	-2.024 (3.097)	0.917 (0.610)	-5.535 (10.01)
δ_4	-10.60* (6.049)	-3.244 (3.234)	0.602 (0.429)	-13.24 (9.001)
δ_5	-11.00** (5.547)	-2.216 (3.929)	1.494 (0.909)	-11.72 (9.002)
δ_6	-9.160* (5.212)	-3.232 (3.119)	0.240 (0.369)	-12.15 (7.906)
δ_7	-9.140* (4.881)	-2.666 (3.072)	0.604* (0.329)	-11.20 (7.384)
δ_8	-7.482 (4.522)	-1.379 (3.137)	0.665 (0.578)	-8.196 (6.823)
δ_9	-5.170 (4.943)	1.071 (3.524)	0.565 (0.495)	-3.534 (7.905)
$\delta_{>9}$	-2.111 (8.936)	0.0906 (3.514)	-0.259 (0.854)	-1.762 (11.07)
Tract FEs	Yes	Yes	Yes	Yes
County \times Year FEs	Yes	Yes	Yes	Yes
Observations	42,462	42,462	42,462	42,462
R ²	0.91	0.85	0.72	0.90
Baseline mean	339.5	119.6	0	459.1

Notes: This table shows estimates of equation 1.3. All coefficients are normalized relative to $\tau = -1$, and $\tau = 0$ is the year in which the merger was approved by federal regulators. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

1.8 show that the average closing is associated with a significant reduction at the 90% level of 16 bank mortgages and a significant increase of 3 fintech mortgages. Over the nine years following a closure, this amounts to a total decrease of 151 bank mortgages and a 27 increase in the number of

Table 1.7: Reduced Form Estimates

	Reduced Form			
	(1) Bank Mortgages \$(000s)	(2) Shadow bank Mortgages \$(000s)	(3) Fintech Mortgages \$(000s)	(4) Total Mortgages \$(000s)
$\delta_{<-1}$	-1,207.6 (1,682.4)	-202.6 (511.8)	-76.0 (57.5)	-1,486.2 (2,106.4)
δ_0	1,336.8 (1,253.2)	1.1 (456.3)	-7.1 (37.2)	1,330.8 (1,530.4)
δ_1	934.1 (2,237.4)	1,025.6 (1,031.5)	1.4 (36.3)	1,961.1 (3,112.6)
δ_2	1,821.3 (2,851.0)	1,154.5 (1,399.9)	5.8 (37.9)	2,981.6 (3,988.4)
δ_3	-921.0 (2,301.6)	-140.7 (713.7)	254.1 (168.9)	-807.5 (2,951.4)
δ_4	-2,943.7** (1,474.8)	-746.6 (570.5)	144.2 (120.4)	-3,546.1* (1,955.0)
δ_5	-3,169.0** (1,526.0)	-483.5 (793.9)	410.0 (252.4)	-3,242.4 (2,089.5)
δ_6	-2,314.6* (1,342.8)	-837.0 (518.0)	53.1 (110.8)	-3,098.5* (1,651.7)
δ_7	-2,372.4** (1,161.2)	-220.5 (529.8)	140.8 (87.5)	-2,452.1* (1,478.3)
δ_8	-1,281.2 (1,196.0)	-43.9 (734.8)	69.1 (131.6)	-1,256.0 (1,647.2)
δ_9	-1,304.6 (1,210.7)	243.3 (881.9)	76.1 (118.8)	-985.2 (1,807.8)
$\delta_{>9}$	-846.2 (2,299.6)	420.2 (1,019.9)	3.5 (238.9)	-422.5 (2,896.5)
Tract FEs	Yes	Yes	Yes	Yes
County \times Year FEs	Yes	Yes	Yes	Yes
Observations	42,462	42,462	42,462	42,462
R ²	0.88	0.80	0.67	0.87
Baseline mean	63,222.6	20,015.9	0.0	83,238.5

Notes: This table shows estimates of equation 1.3. All coefficients are normalized relative to $\tau = -1$, and $\tau = 0$ is the year in which the merger was approved by federal regulators. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

fintech mortgages. Column 3 of table 1.9 shows that the fintech increase in the number of mortgages also translates into a significant increase in the dollar volume of mortgages of \$640,500. This increase amounts to a total

of \$5.7 million over the nine-year period. Compared to the annual bank mortgages baseline, the variation in the number of mortgages corresponds to a 44% decrease for banks and an 8% increase for fintech. The increase in the volume of fintech mortgages corresponds to a 9% increase over the same benchmark. Columns 2 and 4 of table 1.8 show that closings have no significant impact on shadow bank mortgages and the total amount of mortgages.

To evaluate the effects of bank branch closings on mortgage lending for each lender type over time, in table 1.10, I estimate a more flexible version of the 2SLS equation 1.1:

$$Mortgage_{tcmly} = \sigma_t + (\mu_y \times \nu_c) + \mathbf{X}_i \lambda_y + \delta_\tau (\widehat{D_{my}^\tau} \times Closing_{tcm}) + \eta_{tcmly} \quad (1.4)$$

Table 1.8: Second Stage Estimates - Mortgage Originations

	2SLS			
	(1) Bank Mortgages	(2) Shadow bank Mortgages	(3) Fintech Mortgages	(4) Total Mortgages
Post × Closing	-16.87* (8.727)	2.721 (5.284)	3.001*** (0.810)	-11.15 (12.88)
Tract FEs	Yes	Yes	Yes	Yes
County × Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	1,538	1,538	1,538	1,538
Observations	42,462	42,462	42,462	42,462
R ²	0.42	0.31	0.35	0.38
Baseline mean	339.5	119.6	0	459.1

Notes: This table shows estimates of equation 1.1. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses.

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Columns 1 and 3 respectively show that the average closing is associated with a significant reduction in bank mortgages and a significant increase in fintech mortgages. Specifically, after a closing, bank mortgages experience a significant accumulated decrease of 235 mortgages at the 90% level or above from the fourth year post-merger up until the eighth year. This reduction in bank lending represents a decrease of 69% with respect to the annual baseline mean for banks. Fintech mortgages experience a significant increase of 10 mortgages at the 90% level or above, after a closing in the third and fifth-year post-merger combined. These variations in

Table 1.9: Second Stage Estimates - Mortgage Originations Volume

	2SLS			
	(1) Bank Mortgages \$(000s)	(2) Shadow bank Mortgages \$(000s)	(3) Fintech Mortgages \$(000s)	(4) Total Mortgages \$(000s)
Post \times Closing	-2,716.26 (2,487.45)	901.59 (1,236.08)	640.50*** (222.10)	-1,174.17 (3,359.50)
Tract FEs	Yes	Yes	Yes	Yes
County \times Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	1,538	1,538	1,538	1,538
Observations	42,462	42,462	42,462	42,462
R ²	0.22	0.19	0.30	0.22
Baseline mean	63,222.64	20,015.87	0.00	83,238.51

Notes: This table shows estimates of equation 1.1. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses.

Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

lending for fintech represent an increase of 3% with respect to the annual baseline mean for banks (baseline mean for fintech is 0). Columns 2 and 4 show that closings have no significant persistent impact on shadow bank mortgages and only a significant one, at the 90% level, in years fourth and sixth post-closings for the total amount of mortgage supply.

To test whether the effects of branch closings also have an impact on the \$ volume of mortgages, I estimate the same flexible version of the 2SLS equation 1.1 In table 1.11. The results closely mirror those for the number of mortgages. After the average closing, bank mortgages experience a total accumulated decrease, significant at the 90% level or above, of \$52 million between the fourth and seventh post-merger years, and a positive significant one for fintech of \$2.7 million in the third and fifth post-merger years. These variations in lending represent a decrease of 83% in bank mortgages and a 4.3% increase in fintech ones with respect to the annual baseline mean for banks.

Heterogeneity To examine how the impact of closings varies across demographic groups, in table 1.12 I separately estimate the 2SLS equation 1.1 splitting the sample by the median according to a set of relevant demographic characteristics related to the predominant gender, age, and eco-

Table 1.10: Second Stage Dynamic Estimates - Mortgage Originations

	2SLS			
	(1) Bank Mortgages	(2) Shadow bank Mortgages	(3) Fintech Mortgages	(4) Total Mortgages
$\delta_{<-1}$	-19.54 (19.99)	-8.614 (12.08)	-0.679 (1.851)	-28.83 (29.48)
δ_0	0.526 (23.14)	-5.656 (13.98)	0.549 (2.142)	-4.581 (34.14)
δ_1	1.544 (21.12)	11.82 (12.76)	0.953 (1.955)	14.31 (31.15)
δ_2	-7.782 (23.00)	8.574 (13.90)	1.079 (2.130)	1.871 (33.94)
δ_3	-29.02 (23.42)	-9.242 (14.15)	4.222* (2.169)	-34.04 (34.56)
δ_4	-48.83** (23.62)	-13.30 (14.27)	3.032 (2.187)	-59.09* (34.85)
δ_5	-52.12** (23.76)	-11.48 (14.36)	6.136*** (2.200)	-57.46 (35.05)
δ_6	-47.94** (23.80)	-14.85 (14.38)	2.231 (2.203)	-60.56* (35.11)
δ_7	-47.67** (23.80)	-12.47 (14.38)	3.526 (2.204)	-56.62 (35.12)
δ_8	-41.34* (23.73)	-8.868 (14.34)	3.363 (2.198)	-46.85 (35.02)
δ_9	-33.81 (23.68)	-1.108 (14.31)	3.067 (2.192)	-31.85 (34.93)
$\delta_{>9}$	-23.38 (21.23)	-3.537 (12.83)	2.061 (1.966)	-24.86 (31.32)
Tract FEs	Yes	Yes	Yes	Yes
County \times Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	61.49	61.49	61.49	61.49
Observations	42,462	42,462	42,462	42,462
R ²	0.41	0.30	0.34	0.38
Baseline mean	339.5	119.6	0	459.1

Notes: This table shows estimates of equation 1.4 splitting the interaction of exposure to a merger and *Post* into a set of annual interactions with leads and lags. All coefficients are normalized relative to $\tau = -1$, and $\tau = 0$ is the year in which the merger was approved by federal regulators. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

nomic characteristics of the census tract.

I start the analysis by testing whether the impact of closings varies with

Table 1.11: Second Stage Dynamic Estimates - Mortgage Originations Volume

	2SLS			
	(1) Bank Mortgages \$(000s)	(2) Shadow bank Mortgages \$(000s)	(3) Fintech Mortgages \$(000s)	(4) Total Mortgages \$(000s)
$\delta_{<-1}$	-7,527.6 (5,696.2)	-1,109.5 (2,824.7)	-171.9 (507.3)	-8,809.1 (7,691.8)
δ_0	785.4 (6,595.0)	-105.0 (3,270.4)	143.1 (587.3)	823.6 (8,905.5)
δ_1	-564.0 (6,018.8)	2,817.6 (2,984.7)	245.7 (536.0)	2,499.4 (8,127.5)
δ_2	1,721.9 (6,556.1)	3,082.0 (3,251.2)	275.1 (583.8)	5,079.1 (8,853.1)
δ_3	-7,003.2 (6,676.1)	-668.8 (3,310.6)	1,104.7* (594.5)	-6,567.3 (9,015.0)
δ_4	-13,018.7* (6,732.4)	-2,543.2 (3,338.6)	697.5 (599.5)	-14,864.4 (9,091.0)
δ_5	-14,415.8** (6,771.4)	-1,889.1 (3,357.9)	1,601.5*** (603.0)	-14,703.4 (9,143.7)
δ_6	-12,190.1* (6,782.6)	-3,134.5 (3,363.5)	461.2 (604.0)	-14,863.4 (9,158.9)
δ_7	-12,539.8* (6,784.4)	-1,118.7 (3,364.4)	787.4 (604.2)	-12,871.1 (9,161.3)
δ_8	-8,858.6 (6,764.9)	-602.6 (3,354.7)	467.8 (602.4)	-8,993.3 (9,135.0)
δ_9	-8,834.7 (6,748.0)	445.6 (3,346.3)	487.7 (600.9)	-7,901.4 (9,112.1)
$\delta_{>9}$	-7,008.1 (6,051.3)	982.3 (3,000.8)	250.2 (538.9)	-5,775.6 (8,171.4)
Tract FEs	Yes	Yes	Yes	Yes
County \times Year FEs	Yes	Yes	Yes	Yes
First-Stage F stat.	61.49	61.49	61.49	61.49
Observations	42,462	42,462	42,462	42,462
R ²	0.21	0.19	0.29	0.21
Baseline mean	63,222.6	20,015.9	0.0	83,238.5

Notes: This table shows estimates of equation 1.4. All coefficients are normalized relative to $\tau = -1$, and $\tau = 0$ is the year in which the merger was approved by federal regulators. The baseline mean is calculated for exposed tracts in $\tau = -1$. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

tract economic and racial characteristics. This analysis is consistent with findings that show that low-income and minority buyers are primarily re-

liant on relationship-intensive lending¹⁹. For banks, column 1 shows that lower median income tracts are driving the effects of branch closings on branch credit and experiencing a steeper reduction in bank mortgages after the average closing, compared to their baseline LATE. Crucially, column 4 shows that this steeper reduction in bank credit leads to an overall reduction of the mortgage supply in these poorer tracts. This result suggests that the loss of credit relationships caused by branch closures is especially severe for specific population segments. Bank borrowers, such as low-income individuals, who significantly benefit from a personal relationship with branch officers and from the exchange of soft information that the presence of a branch facilitates, suffer a steeper reduction in credit.

I then test the effects on fintech. Column 3 shows that the above-median income tracts drive the positive effect on fintech mortgages after a branch closing. These results contrast with those of banks. Here wealthier tracts are the only ones who receive the positive increase in fintech originations. And, all in all, paint a negative picture of the effects of branch closings for mortgage markets in disfavored areas, as they suffer a reduction in bank credit and do not benefit from the positive effect in fintech credit.

I then proceed to test whether the effects of branch closings differ by the percentage of the minority population, and I find similar results. In column 1, I show that both the effect of branch closings depresses more bank mortgages in below-median white percentage tracts and that these below-median white percentage tracts are the ones that are driving the effect of branch closures. Moreover, column 4 shows that this bank mortgage supply reduction translates into an overall reduction for these higher minority tracts. In turn, column 3 shows that tracts with a higher white percentage of the population also drive the positive effect in fintech credit and have a higher increase in fintech supply than the baseline. These results suggest similar implications as the ones for income. The loss of a relationship seems to carry a higher cost for minorities that both suffer a higher reduction of bank mortgages, do not benefit from the increase in fintech originations, and end up being credit rationed compared to areas not exposed to a closing.

¹⁹Nguyen (2019) shows that post-bank-branch-closings, the decline in credit that follows is especially severe in tracts with lower median income and higher-fraction of minority households. Butcher and Muñoz (2017) show that credit histories of minority and low-income borrowers tend to be thinner. Bond and Townsend (1996) show that borrowers that live in low-income and minority neighborhoods rely more heavily on informal sources of credit.

Table 1.12: Second Stage Splits: Demographic

	Mortgage Originations				Observations (5)
	Banks (1)	Shadow banks (2)	Fintechs (3)	Total (4)	
Baseline	-16.87*	2.721	3.001***	-11.15	42,462
Median income					
Above median	-2.625 (10.46)	-0.799 (6.859)	2.464** (1.111)	-0.961 (15.93)	20,466
Below median	-27.07*** (8.811)	-0.670 (4.800)	0.0448 (0.565)	-27.69** (12.42)	21,492
White percentage					
Above median	-5.423 (12.08)	-0.611 (7.206)	4.783*** (1.164)	-1.250 (17.66)	21,312
Below median	-20.00** (8.450)	-9.109* (5.241)	-1.090 (0.733)	-30.20** (12.76)	20,646
Percentage female population					
Above median	17.81 (10.98)	10.15* (5.879)	0.0254 (0.945)	27.99* (15.43)	21,258
Below median	-26.62*** (9.319)	1.157 (5.893)	3.245*** (0.881)	-22.22 (14.01)	20,430
Percentage population 65 and over					
Above median	1.414 (9.911)	4.846 (5.050)	0.488 (0.761)	6.748 (13.72)	21,600
Below median	-33.52*** (10.61)	-2.337 (7.093)	2.198** (1.071)	-33.66** (16.35)	20,124

Notes: This table shows estimates of equation 1.1. Baseline controls for the corresponding dependent variable in bold are omitted in each corresponding panel. Robust standard errors are clustered at the county level and are in parentheses. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

I continue the analysis measuring whether the effects differ by the percentage of the female population. This analysis is consistent with findings that show that there is a gender gap in fintech lending²⁰. In column 1, I show that both the effect of branch closings depresses more bank mortgages in tracts with a lower percentage of the female population and that precisely these tracts are the ones that are driving the effect of branch closures. Column 3 shows that tracts with a lower female share of the population are driving the increase of fintech mortgages after a closing. These two results suggest that males seem to be the ones driving both the decrease in bank borrowing and the increase in fintech borrowing. After a bank closure, males swiftly switch to new technology-intensive provider types. Women seem to be more conservative in approaching new lender types and either go to another branch of the surviving institution or switch to a non-tech-intensive competitor.

Finally, I test whether the effects of closings differ for the senior population. I test this subgroup of the population since people over 60 use less technology-intensive channels to access banking services²¹ and this could influence the adoption of fintech that follows a closing. In column 1, I show that in tracts with a lower percentage of people 65 and over, the reduction in mortgages caused by a branch closing is much steeper and is driving the overall effect of closings. Moreover, this reduction in bank supply leads to credit rationing in these areas. This evidence suggests that the younger segments of the population are the ones that are switching lender types after a closing. Column 3 shows that tracts with a less senior population also drive the increase in fintech mortgages that follow a closing. These two findings suggest that the senior segment of the population when facing a closing does not substitute the lender type for a more technology-intensive one. These findings could be supportive evidence of the preferences of these groups for the less technology-intensive channels through which banking services have been traditionally provided.

External validity Is the local average treatment effect (LATE) identified by exposure to post-merger branch consolidation representative of the more general effect of bank branch closings in all settings? To explore this, I construct table 1.13 and compare columns 2 and 3 to column 1. Merger sample tracts are more similar to tracts with bank branch closings than

²⁰Chen et al. (2021) show that while 29% of men use fintech products and services, only 21% of women do.

²¹Dodini et al. (2016) show that only 18% of people over the age of 60 use mobile banking.

all branched tracts in several dimensions. For instance, compared to all branched tracts, both sample, and closing tracts, have larger population, higher percentage of college-educated population, lower share of population below the poverty level, lower share of rural population, higher share of population 65 years old and over, lower unemployment rate, higher median family income, more bank branches, and larger bank and shadow bank mortgage markets. Only for population density, the percentage of the minority population, average branch growth, and in fintech originations there is no significant difference between all branched tracts and tracts with closings (and there is a significant difference compared to sample merger tracts) or the sign of its difference varies compared to the one between sample tracts and tracts with closings.

To further scrutinize the LATE identified, I construct table 1.14. In this table, I compare complier tract characteristics to those of sample tracts — complier tracts are those that closed a branch if and only because they were an exposed tract. Although exposure to a merger is assumed to be exogenous to tract characteristics (and this assumption is sufficient for the internal validity of the merger instrument), the posterior decision to close a branch and which particular branch to close need not be exogenous for the internal validity of the instrument to hold. I study this posterior selection because it affects the interpretation of the LATE. With heterogeneous treatment effects, the LATE identified by an instrument is the average treatment effect on the compliers. That is, the effect of closing a branch in an exposed tract that has closed a branch only because it has been affected by a merger —if it had not been affected, it would not have closed the branch. Table 1.14 shows that complier tracts are remarkably similar to sample tracts. However, a few differences remain; complier tracts tend to be less densely populated, have a lower share of rural population, and have a higher number of bank branches. This last point suggests that post-merger closings of bank branches are more focused on over-branched tracts. This focus on over-branched tracts, in turn, suggests that the estimated effects underestimate the impact of an average branch closing in the United States.

1.6 Conclusion

The main contribution of this paper is to show that bank branch closings significantly change the lender mix of the mortgage market. Although closings do not affect the overall supply of mortgages of the market, as

Table 1.13: Representativeness of the Merger Sample

Variable	(1) All branched tracts	(2) Tracts with closings	(3) Merger sample
Population	4,680 [2,195]	4,929 [2,448]	5,634 [3,074]
Population density	4,026 [10,242]	3,484 [7,222]	5,913 [16,556]
Percent minority	24.8 [25.6]	24.5 [23.7]	29.7 [24.0]
Percent college educated	51.7 [18.2]	54.5 [18.4]	60.9 [19.0]
Percent poverty level	12.0 [10.0]	11.7 [10.3]	10.8 [9.9]
Percent rural population	24.1 [38.0]	14.2 [29.6]	3.8 [12.9]
Percent population 65 and over	13.9 [7.4]	14.6 [8.4]	15.2 [11.9]
Percent unemployed	5.7 [4.8]	5.6 [5.0]	5.2 [5.7]
Median income (000s)	52.70 [23.02]	54.84 [24.12]	60.81 [30.09]
Percent MSA median income	104.2 [38.1]	107.2 [40.8]	119.4 [55.4]
Total branches	2.30 [2.08]	3.58 [3.04]	4.28 [2.96]
Branch growth	0.022 [0.147]	0.021 [0.173]	0.055 [0.175]
Bank mortgages	258.6 [260.8]	296.3 [317.2]	356.0 [400.1]
Shadow bank mortgages	77.8 [103.3]	89.1 [123.6]	114.9 [147.0]
Fintech mortgages	0 [0]	0 [0]	0 [0]
Observations	37,537	8,704	2,192

Notes: Standard deviations are in brackets. demographic variables are as of the 2000 census; all other variables are from 2003. Columns 1 and 2 are all tracts in the United States that were branched and had a closing, respectively, over the 2003-2008 period.

Nguyen (2019) exposed, this fact is obscuring an important change. I show that closings both increase fintech supply and reduce bank supply, thus significantly changing the lender mix in the mortgage market. More specifically, I show that the average closing leads to an 8% increase in the number of fintech mortgages and to a 44% decrease in the number of bank mortgages in the nine-year period that follows the closing. These find-

Table 1.14: Complier Characteristics

Variable	(1) Proportion of compliers above the sample median (percent)	(2) Ratio: Compliers to sample
Population	60.6	1.21
Population density	23.8	0.48
Percent minority	47.5	0.95
Percent college educated	53.9	1.08
Percent poverty level	57.1	1.14
Percent rural population	30.9	0.62
Percent population 65 and over	54.8	1.10
Percent unemployed	47.9	0.96
Median income (000s)	47.9	0.96
Percent MSA median income	51.9	1.04
Total branches	84.8	1.70
Branch growth	50.3	1.01
Bank mortgages	60.7	1.21
Shadow bank mortgages	58.3	1.17
Fintech mortgages	0	0

Notes: This table shows how complier tracts compare to the median tract in the sample. Complier characteristics are calculated using the methodology outlined in Angrist and Pischke (2009). Column 1 shows the proportion of compliers who lie above the median tract in the sample; column 2 calculates the ratio of compliers to sample by dividing each entry in column 1 by 0.50. Demographic variables are as of 2000 census; total branches and branch growth are as of the year preceding each merger.

ings are consistent with the view that fintech is improving the products and services offered to the market and not only benefiting from regulatory arbitrage. Since first, fintech lenders and not bank competitors are filling the gap left by the closure of a branch, and second, I compare constant regulation areas.

But, does this change in the lender mix matter? I provide evidence that it does. I show that closings change who gets credit and who no longer does. More specifically, I show that the effects of closings vary across population groups and that crucially, this leads to credit rationing for specific segments of the population. Poorer and higher minority areas both experience a more severe depression of bank credit and do not benefit from the increase in fintech credit after a branch closing. In contrast, richer and lower minority areas are driving the rise in fintech credit. These findings are consistent with the literature on the importance of lending relationships in bank branch lending as bank-branch-customer relationships

are especially important for information-intensive borrowers such as low-income and minority groups. They are also supporting evidence for the view that fintech lenders are increasing the financial exclusion risk of specific segments of the population by using new machine learning methods and other new technologies.

The last takeaway of this paper is that branch closings are significantly changing mortgage markets. Even in the fintech era, where the financial industry makes intensive use of technology, bank branches are still vital for mortgage markets. This paper suggests that branches allow vulnerable to financial exclusion segments of the population to obtain credit by potentially reducing information asymmetries. The findings in this paper have important implications for the financial industry. It shows that most banks' branch network reduction tactic is fostering the adoption of a new competitor type. This competitor is less regulated and employs a different set of tools that may jeopardize the future dominance of banks in mortgage markets.

The current wave of branch closings is far from being over. Branch closings in the U.S. are currently accelerating and are expected to continue in the following years. Simultaneously, the role of bank branches is being redefined. Branches are being transformed by reducing the emphasis on day-to-day operations and emphasizing the tailored service and commercial focus. Additionally, the financial industry is experiencing other vital disruptions. Large corporations with a significant advantage in data accumulation and data processing technology known as bigtech are entering the financial services industry. In this paper, I show that research on the interplay between incumbents, newcomers in the industry, and the role of bank branches will still be crucial in the coming years. The final equilibrium resulting from the interplay between these forces shapes essential outcomes for the financial industry and the lives of most American citizens.

Chapter 2

HOW SELECTIVE ACCESS TO FINANCIAL INFORMATION AFFECTS HOW INVESTORS LEARN

Joint with Gaël Le Mens¹

In this study, we compare learning in two common settings in financial markets. One in which investors can observe the outcome of an investment alternative only if they invest in it, and another one in which they always can observe the outcome—even if they do not invest in it. We provide empirical evidence that investors' beliefs are, on average, 5% closer to the objective Bayesian beliefs given the observed information when investors are in a setting in which they have access to the financial information because of endogenous choice. Then we are able to describe the mechanism that explains our findings. We show that the endogenous creation of the sample of information triggers different cognitive processes. These alternative processes cause better information processing and are of enough magnitude to help overcome the effect of sampling errors.

¹Data collection for Chapter 2 was funded by ERC Consolidator Grant #772268 from the European Commission to Gaël Le Mens.

2.1 Introduction

Investment environments, by their nature, directly affect the way information is sampled, and thus, experienced by investors. There are many important investment environments in which it is impossible to learn about the outcome of an investment unless the investment is made. As a result, the sample of information that investors use to evaluate investment alternatives is endogenously created by their own choices. They are, what the financial decision-making literature calls, selective feedback environments.

Examples of this type of environment are financial markets traded over the counter, such as most private equity, structured financial instruments, natural resources exploration, or foreign direct investment alternatives for which the risk/reward data is not available in an actionable time-frame and investors can only learn about the risk/reward trade-off if they invest in the alternatives. Other examples of this type of feedback can be found in other managerial and financial domains, for instance, a potential entrepreneur can only be certain about the performance of her potential start-up if she decides to pursue the venture, an employer can only be certain about the performance of a new employee in her company if the employee is hired, or a CEO can only be certain about the outcome of an investment project if that project is carried forward in her company.

However, not all relevant financial environments have this characteristic. There are other environments in which investors can easily learn about the outcomes of an investment alternative—even if they do not choose it. They are what the literature calls full feedback environments. For instance, in the stock market, an investor can always learn about the past or present prices of traded companies; in recruitment decisions, a manager generally has easy access to information about the performance of an employee already under her supervision or in corporate financial decisions CEOs can, in most cases, learn about the return of a realized investment alternatives.

Experimental evidence in finance and economics suggests that full feedback and selective feedback environments differ in, at least, two important ways that affect learning. First, in the way information is acquired by the investors. Second, in how that information is processed. That information acquisition and information processing components of financial decision making are key components to analyze was suggested by Rangel et al. (2008). The relevance of this decomposition is underpinned in the neuroscientific findings by O'Doherty et al. (2004) and Behrens et al. (2007) who show that the neural pathways underlying the two cognitive processes can

be dissociated.

The third group of effects that we expect to make learning outcomes differ in the two environments is the result of the interaction between information acquisition and information processing. For instance, the different information acquisition strategies mentioned above will make to differ the samples obtained in the two environments. These different samples may affect the sample size or sample proportion, and this has been shown to lead to different information processing outcomes (i.e., Griffin and Tversky, 1992).

To investigate whether learning is different when people face an environment in which they can only learn if they choose an investment alternative relative to when they face an alternative in which they can learn about an alternative regardless of their choice, we recruited adult participants from the U.S. through Amazon's Mechanical Turk online platform, who were proposed to participate in a study that required the completion of a financial decision making task similar to Kuhnen (2015). In our study, however, participants were randomly allocated (50%/50%) to one of two conditions at the beginning of the experiment. Participants could either face the full feedback condition or the partial feedback condition, which was not present in Kuhnen (2015) and is our main intervention. Crucially, the two conditions differed by the way the information regarding the stock payoffs could be accessed. In the selective feedback condition, participants accessed the information about the payoffs only if they chose the stock in that trial. In the full feedback, condition participants accessed the information regarding the payoff of the stock regardless of their choice. More precisely, the steps participants followed in each trial were the following. In either condition, participants had to choose in each trial between a bond or a stock. Those who chose the stock observed the dividend paid by the stock after making their asset choice and then were asked to provide an estimate of the probability that the stock was paying from the good distribution. However, if they chose the bond, only participants in the full feedback condition always observed the dividend paid by the stock and needed to provide the probability estimate. Those in the selective feedback condition neither observed the outcome nor had to state their probability estimate. In either condition, two types of payoff domain—gain or loss—were possible. Subjects were paid based on their investment payoffs and a fixed participation installment.

To deepen our understanding of these two environments, first, we test and measure whether there are any systematic differences in information

processing between the two conditions. Here, we provide experimental evidence that participants facing a selective feedback environment, are on average, 5% more accurate in their beliefs compared to the objective Bayesian posterior about the quality of the stock than participants who face the same risky alternative but receive information about it regardless of their choice. This is the first measurement in the literature, of the effect in the belief formation process, of being in a selective versus a full feedback environment.

Next, we measure the effect of information processing and the different information acquisition behaviors resulting from the two learning environments. Note that in the full feedback condition, there is no possible sampling error, but that in the selective feedback condition there is². As a consequence of the two learning environments, we find that the samples of information that investors use in the two environments are systematically different. People in the selective feedback condition gather smaller samples. These smaller samples of information lead to a sampling error that, on average, adds 5% error compared to the fully informed objective Bayesian beliefs³.

Finally, we measure the combined effect of the two potential sources of error in learning. We find that the better information processing in the selective feedback environment and the increasing sampling error in the same condition lead to an overall null effect on learning outcomes comparing the two environments. Crucially this is the result of a dynamic process that we reveal and measure. Unveiling this process, measuring the size of its distinct components, and explaining the gap between choice behavior and decision outcomes is the main contribution of this empirical work.

Our study is the first one that takes a double approach to analyze the effects of selective feedback environments in investors' learning in the finance field. Here, we quantify both the effects resulting from differences in information processing and the differences arising from information acquisition, and we are able to describe the process that explains our results. Moreover, there is no previous experimental study in finance analyzing the effects of these two learning environments that focus on analyzing investors' beliefs.

²In the full feedback condition participants have available the full sample of realized outcomes at each point in time, while participants in the selective feedback condition can miss some of that information if they do not choose the stock in any trial

³The Bayesian beliefs calculated using all potential information available to investors whether this information has been accessed because the investor has chosen the alternative or not

Analyzing the beliefs and not focusing only on the value function and choices is important since prior experimental evidence shows that the effects of learning environments can change both the beliefs and the value function. Thanks to this study, we can better understand why previous studies analyzing the effects of foregone outcomes and the effects of active learning had contradictory results. There are counteracting forces, both in information processing and information acquisition, with an effect on investors learning that counterbalance each other. Here we reveal that previously unaccounted effects in information processing also have a role in explaining the differences in learning outcomes and maximization.

Our study can also deal with two possible criticisms of studies that focus on eliciting the beliefs of people. First, that probability errors found in our study are not linked to investors learning. And second, that the subjective beliefs of participants in our study are possibly non-meaningful quantities to study. Kuhnen (2015), in her very similar experimental setting, show that the probability errors reflected by the beliefs of participants are related to the learning capacity outside of the experiment according to two different measures of learning. Then shows that people who were participating in the experimental task acted based on the subjective beliefs stated during the experiment. They are significantly more likely to choose the stock if they believe that the probability of it paying from the good distribution is higher.

Additionally, we have evidence that the experimental task used in this study correlates with real-life investment decisions. Häusler et al. (2018) using fMRI data from an experimental design similar to our study, shows that activity in the anterior insula during the assessment of risky vs. safe choices in an investing task is associated with self-reported real-life active stock trading. Moreover, the authors show that this association remains intact even when they control for financial constraints, education, the understanding of financial matters, and cognitive abilities. Finally, Häusler et al. (2018) using measures of preferences and beliefs about risk-taking show that both measures mediate the association between brain activation in the anterior insula and real-life active stock trading.

The work presented here contributes to the experimental literature on learning in financial markets that have been growing in recent years. For instance, Kluger and Wyatt (2004) show that there is heterogeneity across traders respect their ability to learn according to Bayes's rule. Bruguier et al. (2010) found that skill in predicting price changes in markets with insiders correlates with scores on two tests that assess the human capacity

to discern malicious or benevolent intent and not on the ability to solve complex mathematical problems. Kogan (2008) and Carlin et al. (2013) analyze both the effects of overconfidence on learning and complexity on trading and found that strategic considerations influence the two.

Payzan-LeNestour and Bossaerts (2014) find that investors that face an environment with investment alternatives that change randomly and with payoffs that are observable only if investors invest in them learn overwhelmingly in a Bayesian way as neoclassical finance assumes. On the contrary, investors stop learning in a Bayesian way and learn in a bounded rational way when not nudged into paying attention to contingency shifts. Asparouhova et al. (2015) show that under asymmetric reasoning, prices do not reflect all types of reasoning. Investors unable to produce correct probability computations prefer to hold portfolios with unambiguous returns and do not directly influence asset prices.

Kuhnen (2015) finds that being in the negative domain leads individuals to form overly pessimistic beliefs about available investment options. Kuhnen et al. (2017) show that prior portfolio choices influence investors' expectations of asset values and future choices. This is the result of people updating more from information consistent with their prior choices, and this leads to sticky portfolios over time. Banerjee et al. (2017) show that more negative financial outcome experienced histories tended to produce poorer cognitive performance. Payzan-LeNestour (2018) show that people facing tail risk overwhelmingly behaved like Bayesian learners and that this is the best strategy to survive when facing this type of risk. Hartzmark et al. (2019) show that people overreact to signals about goods that they own, but that learning is close to Bayesian for non-owned goods.

Studying differences between investment environments is becoming increasingly important. New investment environments are frequently created. With the advent of fintech companies, proptech companies, cryptocurrencies, or peer-to-peer financial platforms, the market architects that build these ecosystems need to be aware of the effects of the choices that they make when they design them. Moreover, most financial markets can be either fit in one of the two investment environments analyzed in this study. Here we show that the way information is provided to the participants on these markets has a systematic effect on the belief formation process of investors and significantly affects their financial decisions.

The remainder of the article is organized as follows. Section I describes the relevant literature. Section II describes the experimental design. Section III analyzes the results. Section IV concludes.

2.2 Literature Review

2.2.1 Information acquisition

Exploration-exploitation dilemma Differences in information acquisition strategies are very important to determine the learning outcomes between the two environments. Looking at differences in this aspect between the two environments, we have previous evidence that people, on average, acquire more information in full feedback environments relative to selective feedback ones (Erev and Haruvy, 2015). This results from investors in the selective feedback environment facing the “exploration-exploitation dilemma” —whereas those in a full feedback environment do not face it. This dilemma refers to the fact that investors in a selective feedback environment to learn about a risky investment alternative —or equivalently in order to explore it—need to choose the alternative. In contrast, in a full feedback environment, that information is available regardless of investors’ choices.

This means that, on certain occasions, to learn about investment alternatives, investors may have to choose alternatives with a lower subjective expected value. That is, investors may have to forego the alternative with the highest subjective expected value—generally known as the exploitation option—in order to explore the rest of the alternatives. A key consequence of this is that when learning in environments without access to foregone payoffs, people will tend to use a smaller sample of information to inform their decisions. More precisely, since gathering information can be costly (Selten and Chmura, 2008; Shafir et al., 2008), investors tend to reduce the size of the samples collected. Moreover, these smaller samples, compared to those of investors in a full feedback environment, will, in most cases, be less representative of the actual distribution of the outcomes of the investment alternatives Fiedler (2000).

Use of small samples That investors use unrepresentative samples matters for the quality of their financial decisions. Using these smaller, less representative samples will lead them to make “Sampling error”⁴ even if they process information perfectly. Another relevant phenomenon related to information acquisition that arises in a selective feedback environment but not in a full feedback one is that of “Adaptive sampling”. “Adaptive

⁴Sampling error is the fruit of using unrepresentative samples of information to make inferences about the payoff distributions of the alternatives

sampling” are the terms that we use to refer to the assumption that the adaptive learning models have built-in. That is that decision-makers increase the probability of choosing an alternative after a high outcome and decrease the probability of choosing the alternative after a low outcome⁵. If people evaluate the alternatives following an adaptive behavior. Adaptive sampling can lead investors to collect unrepresentative samples about the alternatives that they are learning from.

Adaptive sampling Adaptive sampling has been shown to make risk-neutral decision-makers that use an optimal policy of learning, to behave as risk-averse participants in the gain domain or risk-seeking in the loss domain, or to produce biased impressions of people or social groups⁶. “Adaptive sampling” has empirical support in the empirical finance literature. Karlsson et al. (2009) create and test a model which links information acquisition decisions to the hedonic utility of information. The authors provide evidence that individuals monitor and attend to information more actively given preliminary good news but “put their heads in the sand” by avoiding additional information given prior adverse news. They refer to this behavior with the name the “ostrich effect”. On another dimension, we know that the use of different samples in the two environments has significant economic implications. Empirical evidence in finance and economics shows that the samples of information that investors directly experience influence key investment outcomes more than other available information that investors’ do not directly experience⁷. Experimental data also underpins this statement (Cohn et al., 2015). Crucially, these empirical studies also show that the different information acquisition attitudes can lead to sub-optimal investment behavior.

⁵Adaptive learning agents base their behavior on Thorndike (1927) law of effect, which states that any behavior is followed by favorable consequences is likely to be repeated, and any behavior that is followed by negative consequences is unlikely to be repeated

⁶See: March (1996); Denrell (2005, 2007); Le Mens et al. (2011); Le Mens and Denrell (2011)

⁷See: Froot (2001); Choi et al. (2009); Malmendier and Nagel (2011); Barberis and Xiong (2012); Ingersoll and Jin (2013); Imas (2016); Dittmar and Duchin (2016); Necker and Ziegelmeyer (2016); Guiso et al. (2018); Shigeoka (2019); Liu and Zuo (2019)

2.2.2 Information processing

Access to foregone outcomes Camerer and Hua Ho (1999) show that foregone payoffs in economic games are weighted less in investors' future choices than payoffs obtained as a result of a direct choice of an alternative. Similar findings are found by Ashby and Rakow (2016) using eye-tracking technology to evaluate decisions from experience⁸ in an experimental setting. Their data suggest that vigilance to outcomes decreases as more consecutive choices are made, is greater for obtained than for foregone outcomes, and when options deliver only gains as opposed to losses or a mixture of gains and losses. Furthermore, the authors find that this variation in attentional allocation plays a central role in the apparent inconsistency in choice, with increased attention to foregone outcomes predicting switches to that option on the next choice.

Another difference between the two environments that can arise when given access to foregone outcomes is the result of the phenomena of "selective attention"⁹. Selective attention refers to the proven fact that, in certain situations, investors choose to either avoid paying attention to information or avoid internalizing information concerning the foregone outcomes they are given. This behavior is more likely when the feedback reveals that they have made a mistake than when the feedback reveals that they have made a good choice. This reduction in the attention given to foregone outcomes can have a deep impact on investors' behavior since we have ample evidence that attention influences investment decisions¹⁰. Theoretical work in finance has also focused on the importance of attention. For example, Hirshleifer and Teoh (2003) model firms' choices between alternative means of presenting information and the effects of different presentations on market prices when investors have limited attention and processing power and are able to rationalize certain behaviors.

A third difference between the two learning environments in information processing that can arise in environments with access to foregone outcomes is that of underweighting of small probabilities. Previous research shows that providing people with information about foregone out-

⁸Decisions by experience here means decisions in which participants in the experimental task do not have prior knowledge about the outcome distribution of the investment alternatives they are facing

⁹See: Ehrlich et al. (1957); Frey and Stahlberg (1986); Witte (1996); Caplin and Eliaz (2003); Kőszegi (2003, 2010); Thornton (2008); Oster et al. (2013); Golman et al. (2017)

¹⁰See: Barber et al. (2005); Barber and Odean (2007); Dellavigna and Pollet (2009); Da et al. (2011); Hartzmark (2014); Stango and Zinman (2014); Sicherman et al. (2015)

comes in repeated decisions by experience in a laboratory is linked to underweighting of small probabilities by investors (Grosskopf et al., 2006). This underweighting increases the appeal of rare attractive events and thus makes people less risk-averse. That is why foregone payoffs are considered one of the main mechanisms used to encourage participation in casino gambling and state lotteries.

However, interestingly, the literature on the effects of the availability of foregone outcomes reveals that these differences have an important impact on the investment decisions made by investors in the laboratory but not a profound impact on other financial outcomes of interest such as the maximization of rewards (Grosskopf et al., 2006). Whether having access to foregone payoffs helps people achieve maximization in experimental tasks is not clear. Depending on the environment, this information can either facilitate, impair or have no significant effect on maximization. Furthermore, even though this effect on the maximization of information about foregone payoffs is not apparent, it has been shown to have profound effects on individual choice behavior. This contradiction between no clear effects on maximization and profound effects on choice behavior has been traditionally explained using an information acquisition perspective. According to this view, foregone outcomes have two opposite sign effects on maximization. Foregone outcomes on one side have a positive effect because they increase information about the alternatives that the decision-maker is facing. This information can help avoid "getting stuck" in a sub-optimal alternative. On the other side, foregone outcomes information can lead to counterproductive switching. That is, foregone outcome information can attract people to choose a sub-optimal alternative, for example, one with higher variance but lower expected value.

Active learning Another difference in information processing that we can expect between the two environments is the result of "active learning" studied in the psychological learning literature. This literature on "active learning" shows that active control, that is, the opportunity to control the information experienced while learning, improves memory for studied premises as well as transitive inferences involving items that are never experienced. A characteristic pattern that can emerge is that self-directed learning can lead to similar levels of performance with less training even within the same learning environment¹¹. Related to this idea, in the com-

¹¹See: Markant and Gureckis (2010); Gureckis and Markant (2012); Markant et al. (2016); Markant (2018)

puter science field, the development of efficient “active learning” algorithms that can select their own training data is an emerging research topic in machine learning. Unlike traditional learning models that involve passively fed training data, this work has explored algorithms that gather their own training data and can be more efficient in certain environments (see: Settles, 2009 or Sutton and Barto, 2018). Another relevant fact related to how information is processed by investors is the evidence that shows that access to foregone outcomes is naturally linked to the experience of post-decision regret. Coricelli et al. (2007) using neuropsychological, and neuroimaging data studied the fundamental role of the orbitofrontal cortex in mediating the experience of regret. The patterns in the obtained data reflect learning based on cumulative emotional experience. This suggests that affective consequences can induce specific mechanisms of cognitive control of the choice processes, involving reinforcement or avoidance of the experienced behavior.

In the literature, we also find contradictions regarding the effects of the availability of foregone outcomes. For instance, the literature on the effects of “active learning” has no consensus about this type of learning in the accuracy of learners. Whereas Markant and Gureckis (2010); Gureckis and Markant (2012); Markant et al. (2016); Markant (2018) link “active learning” to positive outcomes in learning accuracy, Waggoner et al. (2009) show that active learners learn with the same accuracy as passive learners. Additional evidence in this direction can be found in Keehner et al. (2008). Here the authors found that learners who could actively manipulate a novel 3D object on a computer were no more accurate in learning its shape than passive learners who saw the same screen displays but were unable to manipulate them.

2.2.3 Information acquisition and processing effects

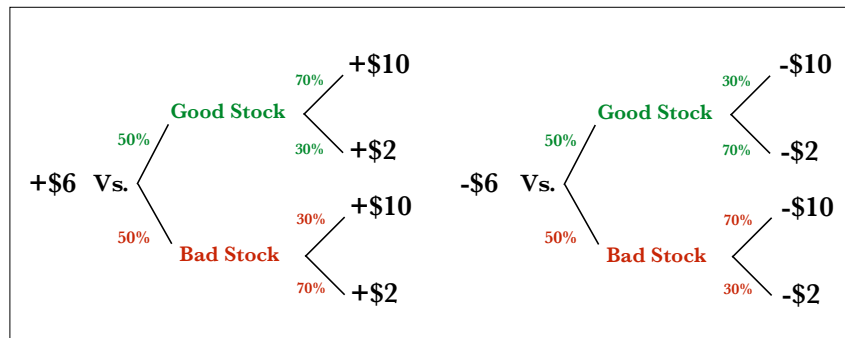
The stated differences in information acquisition can also have an impact on the information processing component. We can find interactions between information processing and information acquisition effects caused by the two learning environments. For instance, Griffin and Tversky (1992) show that people focus on the strength of the available evidence (defined by sample proportion) with insufficient regard for its weight (defined by sample size) and that this leads to substantial violations of Bayes rule. The authors suggest that this behavior can explain both underconfidence and overconfidence in investors’ judgments.

2.3 Experimental Design

One hundred nineteen participants, 65 males, 53 females, were recruited through Amazon’s Mechanical Turk online platform to participate in the experiment. Participants had to complete a financial decision-making task similar to Kuhnen (2015). All participants were presented with information regarding two options: a riskless option, called bond, and a risky one, called stock. Participants faced two different conditions: the full feedback condition and the partial feedback condition. Participants were randomly allocated (50%/50%) to one of the two conditions at the beginning of the experiment. In either condition, the Bond paid +\$6 and the Stock either +\$10 or +\$2.

The Stock could be of 2 types: good or bad. Whether the stock was good or bad was decided randomly (50%/50%) at the beginning of each block. The good stock paid +\$10 with 70% probability and +\$2 with 30% probability; the bad stock paid +\$10 with 30% probability and +\$2 with 70% probability. In Figure 1, this information is summarized in a diagram.

Figure 2.1: Payoffs and probabilities



Notes: This figure describes for the gain domain (left tree) and loss domain (right tree) the: payoffs for the three types of asset, the probabilities of participants facing the good or bad stock, and the probabilities of getting the high or low payoff for each type of stock.

The type of stock that participants faced and its payoffs were generated before the experiment according to the probability distributions mentioned in the previous paragraph. We yoked one participant of each of the two conditions condition to one of the generated sequences. We did this to reduce the variability in the stimuli that participants faced. In total, fifty-nine different sequences were used in the experiment.

The task was divided into 10 learning blocks of 6 trials each, the first

Table 2.1: Experimental Design

Type of Stock	Block Number	Trials	Domain	Condition
Good/Bad	Block 1	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain	Full feedback
Good/Bad	Block 2	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss	Full feedback
Good/Bad	Block 3	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain	Full feedback
Good/Bad
Good/Bad
Good/Bad
Good/Bad	Block 9	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss	Full feedback
Good/Bad	Block 10	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss	Full feedback

Type of Stock	Block Number	Trials	Domain	Condition
Good/Bad	Block 1	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss	Selective Feedback
Good/Bad	Block 2	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Loss	Selective Feedback
Good/Bad	Block 3	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain	Selective Feedback
Good/Bad
Good/Bad
Good/Bad
Good/Bad	Block 9	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain	Selective Feedback
Good/Bad	Block 10	Trial 1; Trial 2; Trial 3; Trial 4; Trial 5; Trial 6	Gain	Selective Feedback

Notes: Each participant had to go through 60 trials. Those trials were split into 10 learning blocks of 6 trials each. In each trial, participants had to choose between a Stock or a Bond. The Stock could be of 2 types: good or bad. Whether the stock was good or bad was decided randomly (50%/50%) at the beginning of each block. The good stock paid the high payoff with 70% probability. The bad stock paid the high payoff with 30% probability. Whether the Stock was good or bad was decided at the beginning of each learning block (with 50%/50% probabilities). In the task, there were 2 conditions: full feedback and selective feedback. Participants faced 10 learning blocks of the same condition, the first 5 blocks in the gain domain and the 5 next blocks in the loss domain. Find above an example of a sequence of full feedback (top table) and selective feedback (bottom table) blocks that a participant may have faced.

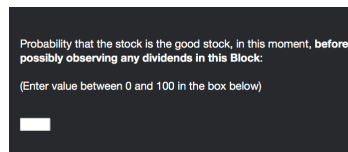
5 blocks in the gain domain and the other 5 blocks in the loss domain. Therefore, each participant had to make 60 choices. In Table I, we show a summary of the experimental design and a sequence of loss and gain learning blocks a participant may have faced during the task.

At the beginning of each block, before participants in either condition had possibly observed any payoffs of the stock, we asked them first to estimate the probability the stock they were facing was the good one. Figure 2 shows the screen that participants faced when facing this question.

Then, in each trial of a block, we first asked participants to choose between the stock and the bond. Participants who chose the stock, independently of the condition they were assigned, first observed the payoff of the stock; second, their accumulated payoffs so far for the whole experiment; and third, they were then asked to estimate the probability that the stock they were facing was the good one. Figure 3 shows the timeline of a typical trial in the full feedback and partial feedback conditions in the case participants chose the stock.

Figure 2.2: Prior subjective estimate elicitation

Subjective estimate before the first choice



Probability that the stock is the good stock, in this moment, before possibly observing any dividends in this Block:
(Enter value between 0 and 100 in the box below)

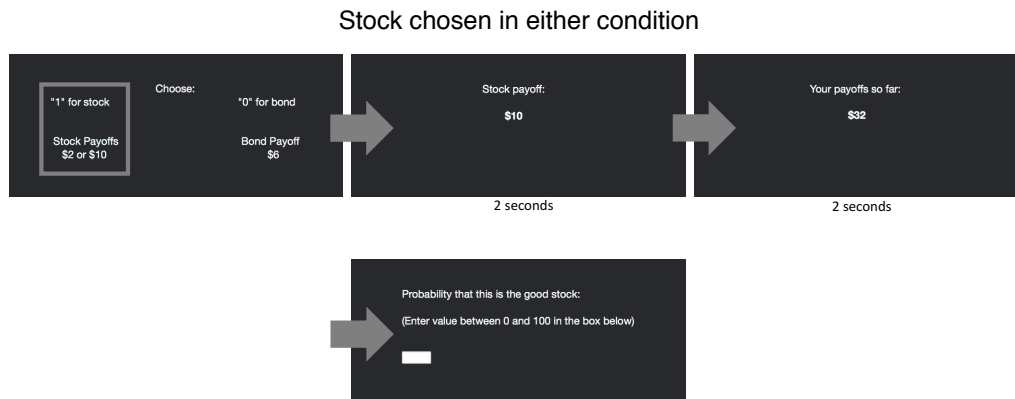
Notes: This figure describes the payoffs for the three types of asset, the probabilities of participants facing the good or bad stock, and the probabilities of getting the high or low payoff for each type of stock.

If participants chose the bond, the steps they had to follow were different in each condition. In the full feedback condition, they had to follow the same three steps as in the case they chose the stock. However, participants in the partial feedback condition participants only saw their accumulated payoffs. If participants chose the bond and were assigned to the partial feedback condition, they did not observe the payoff of the stock for that period and did not have to estimate the probability of the stock being the good one.

Not observing the payoff of the stock is a crucial difference compared to participants assigned in the full feedback condition and the ones in Kuhnen (2015), where all participants were assigned to a full feedback condition and thus observed the payoff of the stock even if they chose the bond. Figure 4 shows the timeline of a typical trial in the full feedback (top timeline) and partial feedback (bottom timeline) conditions for the case that participants chose the bond.

Each participant received a fixed payment of \$5 for participating in the experiment and a bonus corresponding to one-tenth of the accumulated

Figure 2.3: Example of full feedback or partial feedback condition trial in the case the participant chooses the stock



Notes: In either condition, the participant first must choose between the stock and the bond. Regardless of the asset choice, then observes the payoff of the stock for that trial. After this, the participant is reminded of the accumulated payoffs for the whole task. Finally, participants are asked to provide an estimate for the probability that the stock is paying from the good dividend distribution and their confidence in this estimate.

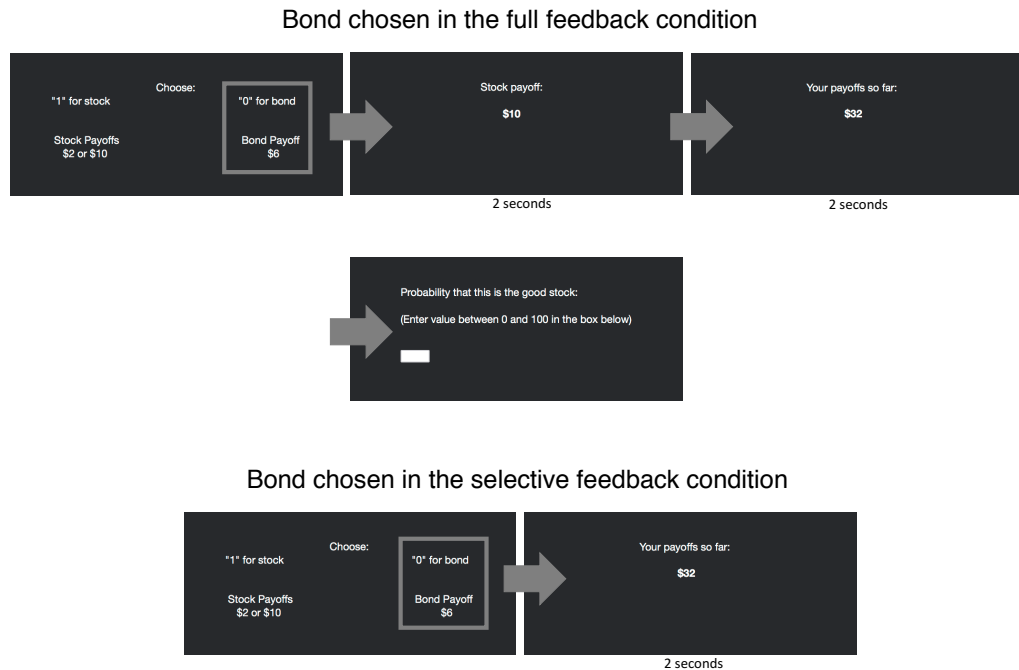
payoffs in the whole task.

For this experimental design, the value of the objective Bayesian posterior is easy to calculate. After observing, t high payoffs in n trials in which the payoff of the stock has been observed so far, the Bayesian posterior that the stock is the good one is given by: $\frac{1}{1 + \frac{1-p}{p} * (\frac{q}{1-q})^{n-2t}}$ where $p = 50\%$ is the prior that the stock is the good one (before any payoffs are observed in that learning block) and $q = 70\%$ is the probability that a good stock pays the high payoff in each trial. In the Appendix, we include all the possible posteriors given all the possible combinations of $\{n, t\}$ in the experiment.

The Bayesian posterior is our benchmark for measuring how close the participants expressed probability estimates are to the objective beliefs. In the full feedback condition, we calculated the Bayesian posterior in each trial; since participants observed the payoff of the stock in all periods. In the partial feedback condition, we calculated the Bayesian posterior only for the trials in which they chose the stock and thus could observe its payoff.

We collected measures of financial literacy and risk preferences for all participants. These two measures can be found in the Appendix, but we

Figure 2.4: Example of full feedback or partial feedback condition trial if a participant chose the bond



Notes: In either type of condition, the participant must first choose between the stock and the bond. Then, participants in the full feedback condition observe the payoff of the stock for that trial, are reminded of the accumulated payoffs for the whole task, and finally, are asked to provide an estimate for the probability that the stock is paying from the good dividend distribution. Participants in the partial feedback condition after choosing the security are only reminded of the accumulated payoffs for the whole task.

will also provide a short description here. To assess risk preferences, after each participant completed the financial decision-making task, we asked them to allocate \$10,000 into 2 different investment options: a risk-free option, in the form of a savings account, and a risky option, in the form of the stock market. Their choice provided a proxy for risk preferences outside the main task of the experiment. To assess financial literacy, we asked participants to solve a financial problem in which they had to estimate the expected amount of money that their previous choice in the risk preference task would have granted them given the financial conditions of the two types of possible investment, either a risk-free asset or a risky one. This question allowed us to check the knowledge of three concepts: probabilities, the difference between net and gross returns, and the difference between stocks and saving accounts. The answer to this question allowed

us to give a financial literacy score from 0 to 3 depending on the number of these three concepts they understood.

2.4 Empirical Findings

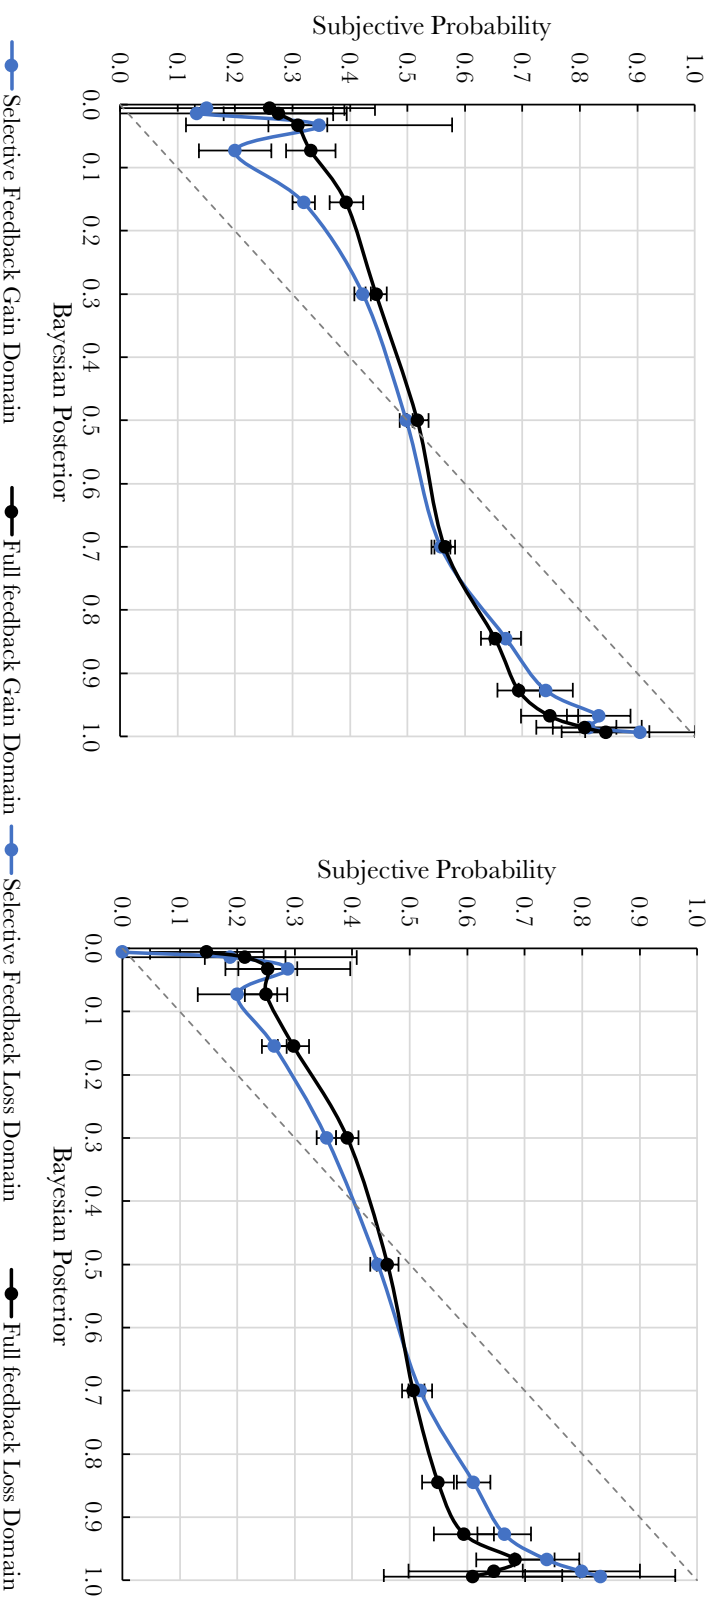
2.4.1 Main Results

We find that participants' beliefs regarding the likelihood that the stock pays from the good distribution are different in the selective feedback condition compared to those of the participants in the full feedback condition. Those differences can be summarized in three results, the main one, and two complementary ones. First and main, the subjective beliefs of participants in the selective feedback condition are closer to the objective Bayesian posterior than those of the participants in the full feedback condition. Second, reference point losses in a selective feedback environment, and not only explicit losses, are sufficient to trigger superior adaptive learning by investors. Third, even if we use as a benchmark a Bayesian posterior that is updated in each trial regardless of the participant's choice, the subjective beliefs of participants in the selective feedback condition are as close to the objective Bayesian posterior as those of the participants in the full feedback condition.

In figure 5, we can observe the first and main results. In this figure, in the x-axis, we represent all the Bayesian objective posteriors that participants faced in the experiment. The number of Bayesian objective posteriors that participants faced during the experiment are limited since there is a finite number of outcome historical paths that participants could observe. We list all of them in the Appendix. The y-axis represents the average of the subjective estimates of the probability of the stock being the good one that participants stated after observing the outcome histories that yield each of the Bayesian posteriors on the x-axis. Points outside the 45° line indicate deviations from the objective Bayesian posterior.


Analyzing figure 5, we observe that participants, in both conditions, deviate significantly from what a Bayesian learner would estimate. If we look at the two graphs, the gain domain (on the left) and the loss domain (on the right), there are very clear deviations from the 45° line in most of

Figure 2.5: Average subjective estimates for the probability that the stock is the good one, as a function of the Bayesian probability



Notes: If participants were Bayesian, all the average subjective estimates would line up in the 45° line. Observations outside that line indicate deviations from the objective Bayesian posteriors. All the Bayesian posteriors that were possible in the experiment can be found in the appendix, together with the combinations of high and low payoffs that lead to them. The average subjective estimates for each Bayesian posterior are plotted in black for participants in the full feedback condition and blue for participants in the selective feedback condition. The left panel presents belief data from the Gain domain, while the right panel presents belief data from the Loss domain.

Figure 2.6: Average subjective estimates for the probability that the stock is the good one, as a function of the Bayesian probability calculated as if participants received full feedback about the payoffs of the stock in each trial



Notes: Posteriors in the figure have been updated after each trial as if participants have observed all the payoffs of the stock up to the trial the posterior is calculated. If participants were Bayesian and had observed all the payoffs of the stock up to the trial, the posterior was calculated, all the average subjective estimates would line up in the 45° line. Observations outside that line indicate deviations from the objective Bayesian posteriors with full feedback. All the Bayesian posteriors that were possible in the experiment can be found in the appendix, together with the combinations of high and low payoffs that lead to them. The average subjective estimates for each Bayesian posterior are plotted in black for participants in the full feedback condition and in blue for participants in the selective feedback condition. The average Bayesian posterior of participants in the Selective feedback condition is plotted in blue and is dashed.

the points for which we have an observation. Crucially, these deviations vary significantly depending on the condition that participants were assigned. Specifically, subjective posteriors in the selective feedback condition tend to be closer to the objective beliefs than those of the participants in the full feedback condition. We can quantify the differences between the two conditions by looking at the results in Table II (gain domain in Panel A, loss domain in Panel B). According to the column (1) regressions on the table, the probability errors of people in the selective feedback condition are on average 6.32% lower in the gain domain and 4.90% lower in the loss domain than those of people in the full feedback condition ($p < 0.01$). If we look at column (3) and column (5) of Table II, we see that the difference between the two conditions is higher for objective probabilities $< 50\%$ in the gain domain and for objective probabilities $\geq 50\%$ in the loss domain. Specifically, in the gain domain, the average participant in the selective feedback condition was 7.02% closer to the objective beliefs than the average participant in the full feedback condition evaluating objective probabilities $< 50\%$ ($p < 0.01$). In the loss domain, the average participant in the selective feedback condition was 6.44% closer than the average participant in the full feedback condition to the objective beliefs only evaluating objective probabilities $\geq 50\%$ ($p < 0.01$). These six results are robust to the inclusion of trial fixed effects.

A relevant point to assess is in which particular scenarios do participants in the selective feedback condition make less probability estimation error than people in the full feedback condition. The analysis in Table IV presents information that can help us to better understand this, and we use this information as a foundation to explain the second result of this study. In Table IV, we find the average change from trial to trial in the probability estimation error after observing either a high payoff (column (1)) or low payoff (column (1)) in a trial for the selective feedback and full feedback conditions, and the gain (Panel A) and loss (Panel B) domains. In column (2), we observe that after participants have observed a low realization of the stock, those in the selective feedback condition are closer to the objective Bayesian beliefs in both the gain domain and loss domain by respectively 5.87% ($p < 0.01$) and 3.29% ($p < 0.1$). Note that the better performance in the gain domain nearly doubles that of the loss domain and that the better performance in the gain domain is significant at the 1% level of confidence and in the loss domain is only marginally significant. So, up to this point, our analysis suggests that the better performance in the selective feedback, compared to the full feedback, has the largest magnitude after a low realization of the stock rather than a high one. But what

Table 2.2: Differences in Probability Estimation Errors in the Selective feedback and Full feedback Conditions

Dependent Variable	<i>Absolute Probability Error_{it}</i>					
	Panel A: Gain Domain					
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	-0.06*** (-3.83)	-0.06*** (-3.82)	-0.07*** (-3.02)	-0.07*** (-3.11)	-0.05*** (-2.90)	-0.05*** (-2.92)
Constant	0.20*** (15.52)	0.20*** (15.51)	0.23*** (11.95)	0.23*** (11.89)	0.17*** (13.36)	0.17*** (13.44)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.039	0.048	0.045	0.064	0.030	0.045
Observations	3541	3541	1404	1404	2137	2137
	Panel B: Loss Domain					
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	-0.05** (-2.47)	-0.05** (-2.47)	-0.03 (-1.50)	-0.03 (-1.57)	-0.06** (-2.30)	-0.06** (-2.30)
Constant	0.21*** (13.29)	0.21*** (13.28)	0.19*** (12.40)	0.19*** (12.34)	0.22*** (10.00)	0.22*** (10.01)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.019	0.030	0.011	0.024	0.026	0.049
Observations	3539	3539	1461	1461	2078	2078

Notes: This table shows that the probability estimation errors are lower in the selective feedback condition relative to the full feedback condition in both the Gain and the Loss domain. The dependent variable in the regression models below, *Absolute Probability Error_{it}*, is the absolute value of the difference between the subjective posterior belief that the stock is the good one that participant i expressed in trial t and the corresponding Objective Bayesian Posterior which is the Bayesian posterior probability that the stock is good, given the information seen by the participant up to trial t in the learning block. The independent variable included is the *Selective feedback trial_i* indicator variable, which is equal to one if participant i is in the selective feedback condition and zero if she is in the full feedback condition. Trial fixed effects are included in the second, fourth, and sixth specifications in each panel. Standard errors are robust to heteroskedasticity and are clustered by subject. t -statistics are in parentheses. ***, ** indicate significance at the 1% and 5% level, respectively.

Table 2.3: [Differences in Probability Estimation Errors against a Full Feedback Bayesian benchmark in the Selective feedback and Full feedback Conditions

Dependent Variable	<i>Absolute Probability Error Full Feedback_{it}</i>					
Panel A: Gain Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	0.011 (0.60)	0.011 (0.60)	-0.03 (-1.25)	-0.03 (-1.43)	0.04** (2.02)	0.04** (2.02)
Constant	0.20*** (15.52)	0.20*** (15.51)	0.23*** (11.95)	0.24*** (11.90)	0.17*** (13.36)	0.17*** (13.46)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.001	0.025	0.007	0.045	0.014	0.049
Observations	3541	3541	1404	1404	2137	2137
Panel B: Loss Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	0.00 (0.16)	0.00 (0.16)	0.01 (0.43)	0.01 (0.29)	0.00 (-0.09)	0.00 (-0.10)
Constant	0.21*** (13.29)	0.21*** (13.28)	0.19*** (12.40)	0.19*** (12.39)	0.22*** (10.00)	0.22*** (10.02)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.000	0.020	0.001	0.029	0.000	0.032
Observations	3539	3539	1461	1461	2078	2078

Notes: This table shows that the probability estimation errors, using a fully informed Bayesian benchmark, are not significantly different in the selective feedback condition and full feedback conditions in both the Gain and the Loss domain. The dependent variable in the regression models below, *Absolute Probability Error Full Feedback_{it}*, is the absolute value of the difference between the subjective posterior belief that the stock is the good one that participant i expressed in trial t and the corresponding Objective Bayesian Posterior updated with full feedback, which is the Bayesian posterior probability that the stock is good, given the information of all payoffs of the stock up to trial t in the learning block. The independent variable included is the *Selective feedback trial_i* indicator variable, which is equal to one if participant i is in the selective feedback condition and zero if she is in the full feedback condition. Trial fixed effects are included in the second, fourth, and sixth specifications in each panel. Standard errors are robust to heteroskedasticity and are clustered by subject. t -statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and level, 10% respectively.

Table 2.4: Mean Probability Estimation Error Updates per Condition and their Difference after Observing a High or Low Payoff

Dependent Variable	<i>Absolute Probability Error_{it}</i>					
Panel A: Gain Domain						
	High Payoff in Trial t	Low Payoff in Trial t	High Payoff in Trial t Subjective estimate t - 1 < 50%	High Payoff in Trial t Subjective estimate t - 1 ≥ 50%	Low Payoff in Trial t Subjective estimate t - 1 < 50%	Low Payoff in Trial t Subjective estimate t - 1 ≥ 50%
Selective feedback condition	0.16	0.15	0.20	0.15	0.16	0.14
Full feedback condition	0.19	0.21	0.22	0.18	0.17	0.23
Difference	-0.03*	-0.06***	-0.02	-0.03*	-0.01	-0.08***
Panel B: Loss Domain						
Selective feedback condition	0.20	0.15	0.29	0.17	0.17	0.14
Full feedback condition	0.23	0.18	0.30	0.20	0.18	0.18
Difference	-0.04	-0.03*	0.00	-0.03*	-0.01	-0.04*

Notes: This table shows the average change from trial to trial in the probability estimation error that a participant produced after observing either a high or low payoff in a trial for the selective feedback and full feedback conditions, gain and loss domains, and for high versus low subjective priors. ***, * indicate significance at the 1% and 10% level, respectively.

happens if we study the effect of the size of the realization, but we differentiate between high and low subjective priors? Columns (3) to (6) add this level of analysis. In column (6), we can see that high subjective prior trials are the main drivers of the effect of low realizations on the subjective probability error. More precisely, after observing a low realization of the stock, people in the gain domain and the selective feedback are 8.14% ($p < 0.01$) closer to the Bayesian objective posteriors than people in the full feedback condition. So, the analysis of the scenarios indicates that the beliefs of participants in the selective feedback condition are closer to those of participants in the full feedback condition, particularly after a low realization of the stock in the gain domain and, especially when this low realization is preceded by a high subjective prior. Therefore, this allows us to state that reference point losses in a selective feedback environment, not only explicit losses, are sufficient to trigger superior adaptive learning by investors.

Additionally, it is interesting to note that, as shown in column (2) of Table V, our second result can be explained by a larger reaction to a low payoff realization that produces a more aggressive probability updating behavior in the selective feedback relative to the full feedback condition. More precisely, in the gain domain (Panel A), people in the selective feedback condition update their subjective probability after observing a low realization of the stock 2.65% ($p < 0.05$) more than people in the full feedback condition.

In Figure 6, we observe the third result. In the x-axis of this figure, we represent all the Bayesian objective posteriors that a Bayesian learner would have produced if she had received information about all payoffs of the stock up to the corresponding trial. Note that this is a crucial difference for the posteriors of participants in the selective feedback condition, but not for those of the participants in the full feedback condition, which do not change with respect to those in Figure 5. Why does the plot of the selective feedback condition change in Figure 6? Because, when we calculate the posterior for participants in a selective feedback condition, we do not only update it in the trials in which participants have chosen the stock and thus observed the payoff of the stock, but we update it in every period regardless of the asset choice. This is possible thanks to how the experiment was programmed, which allowed us to know whether a participant would have observed a high or low payoff if she had chosen the stock, even if she ended up choosing the bond, and thus, not observing the payoff. In Figure 6, the y-axis represents the average of the subjective estimates of the probability of the stock being the good one that participants

Table 2.5: Mean Subjective Probability Updates per Condition and their Difference after Observing a High or Low Payoff

Dependent Variable	<i>Probability Estimate_{t+1} - Probability Estimate_t</i>					
	Panel A: Gain Domain					
	High Payoff in Trial t + 1	Low Payoff in Trial t + 1	High Payoff in Trial t + 1 Subjective estimate < 50%	High Payoff in Trial t + 1 Subjective estimate ≥ 50%	Low Payoff in Trial t + 1 Subjective estimate < 50%	Low Payoff in Trial t + 1 Subjective estimate ≥ 50%
Selective feedback condition	7.27%	-8.09%	10.64%	6.63%	-5.05%	-8.97%
Full feedback condition	5.77%	-5.44%	11.40%	3.85%	-1.42%	-7.72%
Difference	1.51%	-2.65%**	-0.77%	2.78%**	-3.64%	-1.25%
	Panel B: Loss Domain					
Selective feedback condition	6.34%	-9.65%	9.43%	5.41%	-2.73%	-12.58%
Full feedback condition	4.31%	-6.37%	7.21%	2.62%	-1.88%	-10.13%
Difference	2.03%	-3.28%***	2.22%	2.79%	-0.84%	-2.45%

Notes: This table shows the average update in the subjective posterior that a participant produced in the selective feedback or full feedback condition after observing a high or low payoff in a trial. In the third and fourth columns, we restrict the observations to those in which participants observed a low payoff and their subjective probabilities in the previous trial were lower than 50% or equal or higher than 50%. *** indicates significance at the 1% level.

stated after having observed, or not –for participants in the selective feedback condition that did not choose the stock in any period– the outcome histories that yield each of the Bayesian posteriors on the x-axis.

Analyzing Figure 6, we observe a clear pattern. Subjective beliefs of participants in both conditions deviate significantly from the Bayesian posterior calculated as if participants in both conditions received full feedback. This is not surprising for the full feedback condition plot since this is the same plot as we observed in Figure 5, but we do indeed get new information for the selective feedback condition. Contrary to Figure 5, we observe that participants' beliefs in the selective feedback condition do not significantly deviate from those in the full feedback one. And, as reported in column (1) of Table III, this is true for both the gain domain (1.06%, $p > 0.1$) and the loss domain (0.30%, $p > 0.1$) (gain domain in Panel A, loss domain in Panel B).

Apart from the subjective beliefs, in Figure 6, we have another plot. This plot is the objective Bayesian posterior that participants in the selective feedback condition could calculate if they were Bayesian learners, as a function of the Bayesian posterior, calculated as if participants had observed the payoff of the stock in all trials up to the moment the Bayesian posterior is calculated. The first posterior is the one that participants would produce if they were Bayesian learners, given the information they observed about the payoffs of the stock; the second one is the one that they would have produced if they were Bayesian learners and had observed all the payoff of the stock so far in the experiment. The difference between the 45° line and this plot is the average Sampling error that participants made for each point in which we have an observation. The sampling error can only exist in the selective feedback condition since in the full feedback condition; the objective Bayesian benchmark is always updated regardless of the asset choice made by participants. Here, we observe that the Sampling error exists across the whole range of objective probabilities and that it is higher in the extremes for very high and very low probabilities. Crucially, columns (1) of Table III shows that the sampling error is responsible for completely erasing the better processing of information in the selective feedback condition that was reflected in the gain and loss domains in column (1) in Table II, and this allows us to quantify its size of about 5% in each condition.

2.4.2 Alternative Explanations

In figure 5 and Table II, we have shown that participants in the selective feedback condition are close to the objective Bayesian beliefs in the two domains than participants in the full feedback condition. We suggest that this outcome is the product of a different learning process between the two learning environments. Moreover, we suggest that this result is caused by already known effects of access to foregone outcomes.

Table 2.6: Mean Subjective Probability Updates in the Full feedback Condition when Stock Chosen or Sampled and their Difference after Observing a High or Low Payoff

	Dependent Variable	
	<i>Probability Estimate_{t+1} - Probability Estimate_t</i>	<i>Absolute Probability Error_{it}</i>
Panel A: Gain Domain		
Stock selected	0.52%	20.35%
Stock sampled	-0.23%	19.17%
Difference	0.75%	1.19%
Panel B: Loss Domain		
Stock selected	-0.97%	22.34%
Stock sampled	-1.13%	18.55%
Difference	0.16%	3.79%***

Notes: This table shows the average update in the subjective posterior that a participant produced in the full feedback condition, after choosing the stock, or after passively observing the outcome of the stock, for high or low payoff trials. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively.

The first known effect of foregone outcomes that could cause the different learning outcomes is the documented difference in the weight given in a full feedback setting to selected information and passively observed one. In the context of our experiment, if participants treated these two different sources of information differently, we would observe differences in probability updating and in probability estimation errors between payoffs observed as a result of selected stocks or payoffs observed as a result of sampled stocks¹². According to the literature, we would expect that selected stocks would be given more weight in probability updating and

¹²By sampled we referred that the payoffs are observed as a result of choosing the bond

that probability errors would also be lower after a stock selection. To test this, we constructed Table VI. Analyzing the first column of the table, we observe that there are no significant differences in probability updates between selected and sampled stocks neither in the gain domain (0.75%, $p > 0.1$) nor in the loss domain (0.16%, $p > 0.1$). Analyzing the second column of the table, we observe that there are no significant differences in probability estimation errors in the gain domain (0.01, $p > 0.1$), but there are in the loss domain (0.04, $p < 0.01$). However, this significant difference is against the suggested results in the literature and cannot explain the better processing of information in the selective feedback condition. This could have been a factor if sampled stocks were the origin of higher probability estimation errors.

Table 2.7: Mean Subjective Probability Updates and Mean Probability Estimation Error Updates in the First Trial of the Selective and Full feedback Conditions by Asset Selection

	Dependent Variable	
	<i>Probability Estimate_{t+1} - Probability Estimate_t</i>	<i>Absolute Probability Error_{it}</i>
Panel A: Stock chosen		
Selective feedback condition	-2.81%	0.16
Full feedback condition	-1.65%	0.18
Difference	-1.16%	-0.03***
Panel B: Bond chosen		
Selective feedback condition	0.00%	0.07
Full feedback condition	-1.74%	0.15
Difference	1.74%	-0.09***

Notes: This table shows the average update in the subjective posterior that a participant produced in the full feedback condition, after choosing the stock, or after passively observing the outcome of the stock, for high or low payoff trials. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively.

A second effect of having access to foregone outcomes cited in the literature refers to the fact that participants in a full feedback environment will get more conflicting information than those in a selective feedback environment. That participants in the full feedback condition will observe more conflicting information about the quality of the stock is also expected in our experiment. This is because participants in the full feedback condition will observe all the payoffs of the stock and will not miss any, as most

participants in the selective feedback will. The more conflicting information that participants in the full feedback condition will receive could affect their information processing. To test this, we analyze whether there are differences between the two conditions in probability updating and probability estimation errors after the first trial. We choose to measure those two outcomes at this point because this allows us to eliminate potential confounds related to conflicting information since in our test, we will measure differences in outcomes after having made the same choices and where no conflicting information from previous trials is possible. Analyzing column 1 in Table VII, we observe that there are no significant differences in probability updates between the selective feedback and full feedback condition after the stock having been chosen in trial 1 (-1.16%, $p > 0.1$), nor after the bond is chosen in the same trial (1.74%, $p > 0.1$). However, if we compare probability errors at that stage, we observe that differences already exist at that point, and moreover, the differences match the sign of our main result. More precisely, in the second column of Table VII, we observe that the probability errors are already 3% lower in the selective feedback condition compared to the full feedback condition after the stock has been chosen ($p < 0.01$) and 9% lower after the bond has been chosen ($p < 0.01$).

The third effect of foregone outcomes that could explain differential learning is that of the underweighting of small probabilities that access to foregone outcomes produces, according to the literature. To test this, we can look at columns 3, 4, 5, and 6 of Table II. There we observe that the differential learning is not only produced when the objective Bayesian posteriors are lower than 50%, but this is also produced for Bayesian posteriors higher than 50%.

2.4.3 Learning in the last trial

Given that learning is a dynamic process that should help participants to improve results with the experience, we analyze the learning outcomes also in the sixth trial of the experimental task, which is the last trial that participants faced in each block of decisions. We can quantify the differences between the two conditions by looking at the results in Table VIII (gain domain in Panel A, loss domain in Panel B). According to the column (1) regressions on the table, the probability errors of people in the selective feedback condition are on average 6.96% lower in the gain domain and 6.32% lower in the loss domain than those of people in the full

feedback condition (both $p < 0.01$). If we look at column (3) and column (5) of Table VIII, we see that the difference between the two conditions is higher for objective probabilities $< 50\%$ in the gain domain and for objective probabilities $\geq 50\%$ in the loss domain. Specifically, in the gain domain, the average participant in the selective feedback condition was 11.58% closer to the objective beliefs than the average participant in the full feedback condition evaluating objective probabilities $< 50\%$ ($p < 0.01$). In the loss domain, the average participant in the selective feedback condition was 7.87% closer than the average participant in the full feedback condition to the objective beliefs only evaluating objective probabilities $\geq 50\%$ ($p < 0.01$). These six results are robust to the inclusion of trial fixed effects.

Probability errors are another learning measure of interest. We obtained them using as a benchmark the Bayesian objective posteriors that a Bayesian learner would have produced if she had received information about all payoffs of the stock up to the corresponding trial. We can quantify the differences between the two conditions according to this measure looking at the results in Table IX (gain domain in Panel A, loss domain in Panel B). According to the column (1) regressions on the table, the probability errors against a fully informed benchmark of people in the selective feedback condition are no significantly different than those of people in the full feedback condition neither in the gain domain (0.04, $p > 0.1$) nor in the loss domain (0.01, $p > 0.1$). If we look at column (3) of Table IX, we observe that there are also no significant differences between the two conditions for objective probabilities $< 50\%$ neither in the gain domain (-0.02, $p > 0.1$), nor in the loss domain (0.04, $p > 0.1$). In column (5), we observe the results for objective probabilities $\geq 50\%$. For the gain domain, we observe that participants in the selective feedback condition produced 8.09% significantly higher error updates than participants in the full feedback condition ($p < 0.01$). However, participants in the loss domain produced no significantly different error updates between the two conditions (-0.01, $p > 0.1$). These six results are robust to the inclusion of trial fixed effects.

2.4.4 Choices

We proceed now to the analysis of other important learning outcomes. More precisely, now we will analyze differences in choice behavior dynamics in the selective and full feedback condition by the type of stock that the participants faced and by the domain of the payoffs of the decision

Table 2.8: Differences in Probability Estimation Errors in the Selective feedback and Full feedback Conditions in Trial 6

Dependent Variable	<i>Absolute Probability Error_{it}</i>					
Panel A: Gain Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	-0.07*** (-3.17)	-0.07*** (-3.17)	-0.11*** (-3.17)	-0.11*** (-3.17)	-0.05* (-1.72)	-0.05* (-1.72)
Constant	0.22*** (12.51)	0.22*** (12.51)	0.27*** (9.28)	0.27*** (9.28)	0.18*** (9.74)	0.18*** (9.74)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.034	0.034	0.072	0.072	0.016	0.016
Observations	590	590	248	248	342	342
Panel B: Loss Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	-0.06** (-2.39)	-0.06** (-2.39)	-0.04 (-1.43)	-0.04 (-1.43)	-0.08** (-2.09)	-0.08** (-2.09)
Constant	0.24*** (11.12)	0.24*** (11.12)	0.20*** (8.06)	0.20*** (8.06)	0.27*** (8.99)	0.27*** (8.99)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.024	0.024	0.017	0.017	0.028	0.028
Observations	590	590	251	251	339	339

Notes: This table shows that the probability estimation errors, using a fully informed Bayesian benchmark, are not significantly different in the selective feedback condition and full feedback conditions in both the Gain and the Loss domain in the sixth and last trial of each block. The dependent variable in the regression models below, *Absolute Probability Error Full Feedback_{it}*, is the absolute value of the difference between the subjective posterior belief that the stock is the good one that participant *i* expressed in trial *t* and the corresponding Objective Bayesian Posterior updated with full feedback, which is the Bayesian posterior probability that the stock is good, given the information of all payoffs of the stock up to trial *t* in the learning block. The independent variable included is the *Selective feedback trial_i* indicator variable, which is equal to one if participant *i* is in the selective feedback condition and zero if she is in the full feedback condition. Trial fixed effects are included in the second, fourth, and sixth specifications in each panel. Standard errors are robust to heteroskedasticity and are clustered by subject. *t*-statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and level, 10% respectively.

Table 2.9: Differences in Probability Estimation Errors against a Full Feedback Bayesian benchmark in the Selective feedback and Full feedback Conditions in Trial 6

Dependent Variable	<i>Absolute Probability Error Full Feedback_{it}</i>					
Panel A: Gain Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	0.04 (1.58)	0.04 (1.58)	-0.02 (-0.68)	-0.02 (-0.68)	0.08*** (2.81)	0.08*** (2.81)
Constant	0.22*** (12.51)	0.22*** (12.51)	0.27*** (9.28)	0.27*** (9.28)	0.18*** (9.74)	0.18*** (9.74)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.009	0.009	0.003	0.003	0.040	0.040
Observations	590	590	248	248	342	342
Panel B: Loss Domain						
	All Trials	All Trials	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors < 50%	Trials with Objective Posteriors ≥ 50%	Trials with Objective Posteriors ≥ 50%
<i>Selective Feedback Trial_i</i>	0.01 (0.45)	0.01 (0.45)	0.04 (1.26)	0.04 (1.26)	-0.01 (-0.14)	-0.01 (-0.14)
Constant	0.24*** (11.12)	0.24*** (11.12)	0.20*** (8.06)	0.20*** (8.06)	0.27*** (8.99)	0.27*** (8.99)
Trial fixed effects	No	Yes	No	Yes	No	Yes
R^2	0.001	0.001	0.011	0.011	0.001	0.001
Observations	590	590	251	251	339	339

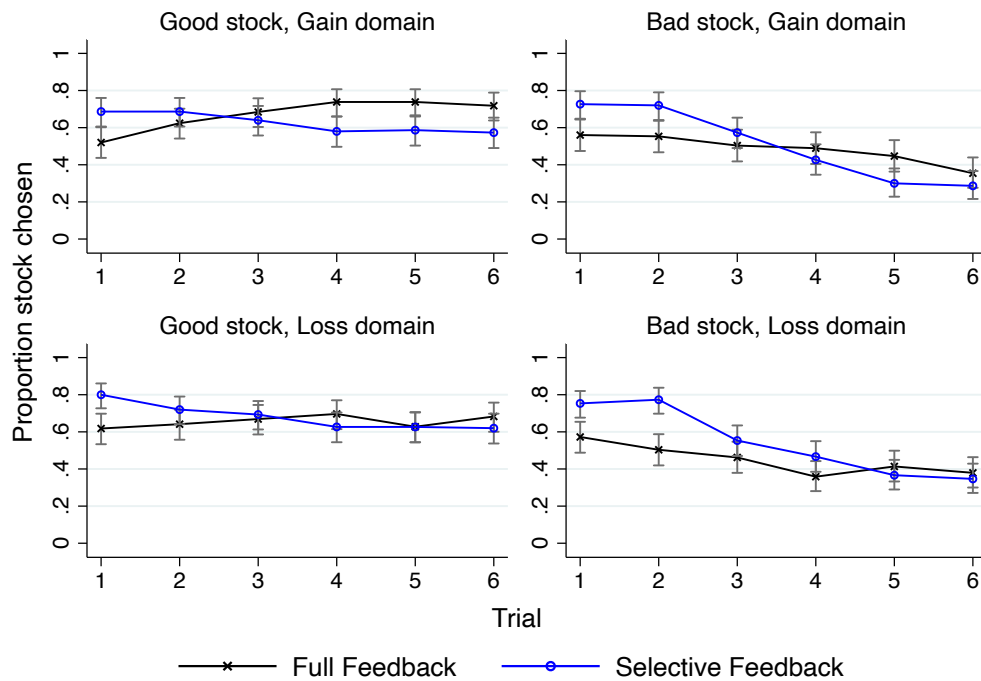
Notes: This table shows that the probability estimation errors are lower in the selective feedback condition relative to the full feedback condition in the sixth trial in both the Gain and the Loss domain. The dependent variable in the regression models below, *Absolute Probability Error Full Feedback_{it}*, is the absolute value of the difference between the subjective posterior belief that the stock is the good one that participant *i* expressed in trial *t* and the corresponding Objective Bayesian Posterior updated with full feedback, which is the Bayesian posterior probability that the stock is good, given the information of all payoffs of the stock up to trial *t* in the learning block. The independent variable included is the *Selective feedback trial_i* indicator variable, which is equal to one if participant *i* is in the selective feedback condition and zero if she is in the full feedback condition. Trial fixed effects are included in the second, fourth, and sixth specifications in each panel. Standard errors are robust to heteroskedasticity and are clustered by subject. *t*-statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and level, 10% respectively.

block. In figure 7, we have plotted the proportion of stock choices per trial for each condition. First, we observe in all four graphs that the proportion of stock choices is always higher in the first two trials in the selective feedback condition. This is the result of the exploration-exploitation trade-off. Since the expected payoffs if participants are facing the good stock of choosing that asset is higher than if they choose the bond, they have an incentive to explore the risky alternative to learn about it. Participants in the selective feedback condition, apart from the value of the payoff, get value from the information that the payoff of the stock reveals. However, participants of the full feedback condition will be able to observe the outcome of the stock regardless of their choice, so choosing the stock will not carry that information value. This higher exploration in the selective feedback condition helps participants to choose the best alternative in terms of expected payoff (the stock) when facing the good stock but hinders participants' maximization when facing the bad stock (since the asset with the higher expected value is the bond in that case).

Second, we observe that the slope of the selective feedback condition plot in the graphs in the second column (when participants faced a bad stock either in the gain or loss domain) is more inclined relative to the same plots in the graphs of the left column (when participants faced a good stock). The negativity bias that adaptive sampling predicts is found here. The negativity bias, produced by facing lower than expected payoffs when choosing the stock, leads participants in the selective feedback condition to make the same proportion of stock choices as participants in the full feedback condition when facing the bad stock. In this case, the negativity bias helps participants to choose the best alternative in terms of expected payoff (the bond). However, if we look at the left column graphs, we observe that at least for the gain domain, the absence of a high negativity bias leads participants in the selective feedback condition to make a significantly lower proportion of stock choices (the maximizing option) than participants in the full feedback condition. The effect for the loss domain would be in the same direction if participants in the full feedback condition facing the loss domain would not have incurred in higher loss aversion in that domain relative to the gain domain, as can be seen comparing the full feedback plots of the left column top and bottom graphs.

Finally, and most important of all, overall, we observe that the learning outcomes that we observed in the previous section match the ones that we find in this one. In the last trial of each block, there is no significant difference in choice behavior as a result of learning in a selective or a full feedback setting.

Figure 2.7: Proportion of stock choices per trial for the selective feedback and full feedback condition by type of stock and domain



Notes: The proportion of stock choices for the selective feedback condition are plotted in a blue line, while the proportion of stock choices for the full feedback condition are plotted in black. In light grey are plotted the 95% confidence intervals for the population proportion in either condition. The top row graphs are produced using only trials in which participants faced payoffs in the gain domain. The bottom row graphs are produced using only trials in which participants faced payoffs in the loss domain. The left column graphs are produced using only trials in which participants faced a good stock. The right column graphs are produced using only trials in which participants faced a bad stock.

2.5 Discussion and Implications

Main Results As our main finding, we uncover and measure the elements of the process that explains why and how learning differs between selective feedback and full feedback environments. This process is composed of two principal elements. The first one is the better processing of

information produced by investors in a selective feedback environment compared to those in a full feedback environment. The elicited beliefs of participants in the selective feedback environment are on average 5% closer to the beliefs that a perfect Bayesian rational learner would have given the available information. We find that investors in a selective feedback environment enter a distinct learning mode that allows them to process the available information in a better way, thus reducing the cognitive error that they make.

As per the second element, we find that the better processing of information in the selective feedback environment is powerful enough to even offset sampling errors. Sampling errors, in our experimental task, on average amount to 5% of additional error. Despite investors in the selective feedback condition having access on average to less information, the better processing of information allows them to offset this sampling error. Sampling errors in our experimental task are generated by the missed information as a result of investors in the selective feedback condition, not choosing the stock, and thus missing the information about its outcome. Therefore we should expect sampling error in a trial, not to change if the stock is sampled, or to necessarily increase if the stock is not chosen. If the latter case happens, the new sampling error will add to the potential new cognitive error and the preexisting cognitive and sampling errors. All these elements combined will produce the total error made by the investor. For these reasons, we expected the error of participants in the selective feedback condition to be higher than that of participants in the full feedback condition.

However, here we show that the better processing of information in a selective feedback environment can, in certain circumstances, eliminate this impairment to investors' learning. This is the result of a dynamic process that we just described and summarized in the following few lines. First, the better processing of information helps participants in the selective feedback condition to make less cognitive error. Second, compared to the investors in a full feedback condition, those in the selective feedback condition, as they decide not to choose the risky alternative, face a cumulative sampling error. In our particular experimental study, we find that these two opposite forces end up perfectly counterbalancing each other.

Additionally, we show that reference point losses in the gain domain, and not only explicit losses, are sufficient to trigger adaptive learning by investors in the selective feedback condition compared to those in the full feedback condition. After observing a low payoff in the gain domain, par-

ticipants in the selective feedback condition, on average, produce beliefs about the quality of the stock that are 6% closer to the objective Bayesian beliefs than those of the people of the full feedback condition. This result is mainly driven by the observations in which participants believed that the likelihood that they were facing the good stock was higher than 50%. Thus, this difference is mainly driven by data points in which participants were more optimistic about the type of stock they were facing and then received negative information—a low outcome.

Method To study whether people learn differently in a selective feedback setting, compared to a full feedback one—in which investors can learn about the outcomes of an investment alternative regardless of their choice—we choose an experimental approach. This methodological choice is based on three main advantages of this method. The first one is that using conventional archival data sources, beliefs cannot be directly observed, but they can be elicited in an experimental setting. The second one is that an experimental approach allows controlling two key aspects for studying human learning. First, we can set known objective priors from which investors should update when confronted with new information. Second, we can precisely control the information that participants can access. And, moreover, we can even control the information they would have accessed if they had chosen the investment alternative, even if the participant decided not to choose the alternative. The third advantage is that in the few available survey-based archival data sources in which investors' beliefs are elicited, it is impossible to analyze both the beliefs and the choices. Still, we can effectively do so with an experimental approach.

Contribution Our contribution, other than the main result already discussed, is sixfold. First, we focus on eliciting and analyzing investors' beliefs and not only on studying their choices. This is valuable since prior experimental studies analyzing the effects of foregone outcomes have focused on choices rather than beliefs. In those studies from the choices, then researchers fitted model parameters (i.e., Grosskopf et al., 2006 or Camerer and Hua Ho, 1999). But we know that using this approach, we can be led astray when we are forced to infer the value of parameters using observable proxies for variables previously thought to be unobservable (Nyarko and Schotter, 2002). Second, thanks to our experimental design, we can precisely measure and isolate two sources of error. The Cognitive error, caused by incorrect processing of information, and the Sampling error,

caused by using a smaller sample of information. This is the first study that can measure the two sources, following a belief approach.

Third, Kuhnen (2015) work shows that investors in a full feedback environment, who face outcomes only in the loss domain, make on average higher probability errors than people facing the same environment in the gain domain. Here we show an opposite result, reference point losses—perceived losses relative to the guaranteed payment provided by the riskless alternative—in the selective feedback environment in the gain domain, and not explicit losses, trigger superior adaptive learning by participants relative to those of people in the same domain observing the same outcome. Our findings provide new evidence of the effects of regret in learning. Having access to foregone outcomes is naturally linked to the experience of post-decision regret (Inman et al., 1997; Ritov and Baron, 1995; Taylor, 1997).

The feeling of regret has been linked to important effects in decision making and has sparked the creation of Regret theory (Loomes and Sugden, 1982) and has inspired experimental studies that have revealed that regret can have a profound influence on the decisions people make—for instance increasing the switching behavior between the different alternatives—and can promote both risk-averse as well as risk-seeking choices (Zeelenberg, 1999). Coricelli et al. (2007) using neuropsychological, and neuroimaging data studied the fundamental role of the orbitofrontal cortex in mediating the experience of regret. Their data indicates reactivation of activity within the orbitofrontal cortex and amygdala occurring during the phase of choice when the brain is anticipating possible future consequences of decisions and that this characterizes the anticipation of regret. These patterns reflect learning based on cumulative emotional experience. This suggests that affective consequences can induce specific mechanisms of cognitive control of the choice processes, involving reinforcement or avoidance of the experienced behavior. This reaction caused by post-decision regret could explain the beneficial effect of reference point losses on learning outcomes.

Our findings align with the recent literature on the adaptive role of losses. In tasks ranging from simple economic decisions to meta perception, previous studies have generally shown positive effects of losses on performance (as in, Costantini and Hoving, 1973; Denes-Raj and Epstein, 1994; Bereby-Meyer and Erev, 1998; Dawson et al., 2002). Additionally, a recent stream of literature (e.g., Yechiam and Hochman, 2014) suggests that losses may be treated as signals of attention and not only as signals

of avoidance. Our results complement previous findings showing that losses induce more controlled processing than comparable gains (Dunegan, 1993) and are associated with some of the physiological indices of attention (Yechiam and Hochman, 2013).

Fourth, when we compare the sizes of the probability updates performed by investors, in the selective feedback and full feedback condition, given the same outcome observed, we find that the same outcome in the selective feedback condition is given a larger probability update. Moreover, we are able to test whether this probability update affects choice behavior accordingly.

Fifth, traditional models to study simultaneous experienced based learning and decision, like the Bayesian Sequential Risk-Taking model or the Expectancy valence model, making use of maximum likelihood methods, or Bayesian methods that give as a result a fixed set of parameters that account for the choice behavior of participants. Our experimental design allows us to measure the dynamic process of each of the modeled features at each moment of time, thus acting as a new viewpoint to better understand participant's behavior step by step.

Finally, the evidence presented here provides new insights that can explain why access to foregone feedback has a great influence on choice but not on maximization. Using a different perspective than the traditional explanations, which take an information-gathering approach, we provide evidence that the information processing view also plays an important role in thoroughly understanding learning in selective feedback environments. This role is related to the better processing of information in selective feedback environments compared to that of full feedback environments. The better processing, we find, can help overcome the loss of information, and thus if the more accurate beliefs translate into more optimal choices, help maximization.

Implications The empirical findings provided here correlate with relevant behaviors outside the laboratory. For instance, recent empirical work in finance highlights the role of the personal experiences of investors. They show that those experiences shape their attitudes towards risky alternatives and financial decisions. We know that whether an investor has experienced the sample of information available matters for economic outcomes of interest and can lead investors to make sub-optimal investment decisions.

For instance, Malmendier and Nagel (2011) show that households that

witness bad economic times both become reluctant to invest in equities and have more pessimistic beliefs about future stock returns. More recently, Necker and Ziegelmeyer (2016), and Guiso et al. (2018) prove that individuals become more risk-averse after a financial crisis, and additionally Shigeoka (2019) shows that these effects on risk aversion triggered by personal experience are long-lasting. Effects of personal experience have also been linked to investment behavior in the insurance market or corporate financial decisions. Froot (2001) showed that after floods or earthquakes, people are more likely to buy insurance against such events even though the probability of occurrence of such events does not change; Dittmar and Duchin (2016) show that firms run by CEO's who have experienced distress have less debt, save more cash, and invest less than other firms. More recently Liu and Zuo (2019) have studied the effects of being exposed to an environment with an average risk aversion parameter different than that of your natal environment. The authors have used this setting to explain the existence of the gender gap in risk aversion, and they find that after spending time in a new environment with the majority of riskier averse children, less risk averse children change their risk preferences and adopt the risk preferences of the majority.

The effect of personal experience in such a high stakes decision as saving for retirement has also been documented. Choi et al. (2009) show that individual investors over-extrapolate from their personal experience when making savings decisions. Investors who experience particularly rewarding outcomes from 401(k) saving—a high average and/or low variance return—increase their 401(k) savings rate more than investors who have less rewarding experiences. The idea that endogenously chosen samples of information have more impact than other available information is related to the realization utility empirical findings (Barberis and Xiong, 2012; Ingersoll and Jin, 2013; Imas, 2016). According to this view, trading and its resulting realized gains and losses cause greater utility swings than paper gains and losses.

Here, we provide evidence that experiencing information that comes as a result of your own choice affects your beliefs differently than other available information. This, for instance, could help explain the presented findings by Dittmar and Duchin (2016) related to behaviors of managers that have experienced financial distress and why those decisions have impacted later on in their life as financiers. Our findings speak to the role of the effect of information that comes from endogenous choices of agents and not as a result of an exogenous event. Here, we find that the sample of information that comes as a result of endogenous choices helps partici-

pants learn more similarly to a Bayesian learner.

Our findings complement those of Hartzmark et al. (2019). In their experimental study show that people overreact to signals about goods that they own but that learning is close to Bayesian for non-owned goods. Moreover, they show that the endowment effect increases in response to positive information and disappears with negative information. Ownership, according to their findings, increases attention to recent signals about owned goods, exacerbating over-extrapolation. Here we find a different channel that affects learning. Learning is not only impacted by ownership but also by the way information is accessed. More precisely, information that comes as a result of an endogenous choice is treated differently in the belief formation process than the information available in the environment.

Our research also speaks to the theoretical literature in finance, which approaches investors' behavior from a bounded rationality approach and uses a non-classical view of the formation of beliefs of economic agents¹³. In this literature, investors are believed to learn as Bayesian learners using a possibly incorrect prior belief. Here we show how investors learn from an objective prior when facing different environments that vary in the way information is sampled, thus adding to the knowledge of how people process new information in different environments and quantifying the effect of lost information.

Another of the channels in information processing that can explain the effect of personal experiences on investment outcomes is the attention paid to the available information. Empirical evidence in the Finance field suggests that the attention investors pay to the financial information influences their financial decisions. Barber et al. (2005) and Barber and Odean (2007) find that individual investors are net buyers of attention-grabbing stocks and mutual funds. In a study of American households' investment behavior, they find that mutual funds purchases mainly occur in the top quintile of past annual returns and show that many investors consider purchasing only stocks that have first caught their attention. Da et al. (2011) using search frequency in Google, as a measure of investors' attention find that an increase in Google searches predicts higher stock prices in the next 2 weeks and an eventual price reversal within the year.

Dellavigna and Pollet (2009) compare the response to earnings announcements on Friday, when investor inattention is more likely, to the response

¹³Barberis et al. (1998); Bossaerts (2004); Brunnermeier and Parker (2005); Gabaix et al. (2006); Van Nieuwerburgh and Veldkamp (2010); Gennaioli and Shleifer (2010)

on other weekdays. They find that Friday announcements have a 15% lower immediate response and a 70% higher delayed response. These findings support explanations of post-earnings announcement drift based on underreaction to information caused by limited attention. Sichertman et al. (2015) show how aggregate and individual household trading behavior are related to investor attention. Hartzmark (2014) document that both retail traders and mutual fund managers are more likely to sell the extreme winning and extreme losing positions in their portfolio (what they call “the rank effect”). This effect is not driven by firm-specific information, holding period, or the level of returns itself, but is associated with the salience of extreme portfolio positions. Stango and Zinman (2014) find that conditional on selection into surveys, individuals who face overdraft-related questions are less likely to incur a fee in the survey month. Moreover, taking multiple overdraft surveys builds a “stock” of attention that reduces overdrafts for up to two years. Our study highlights the special role that attention has on attitude formation towards investment alternatives. Our study shows that different market environments affect the attention that investors pay to information. Moreover, we show that being in a setting with information accessed only by endogenous choice can help investors process that information better.

On the theoretical side, the literature analyzing economic games has proposed three alternative modeling approaches to explain how foregone outcomes can be processed by people. Our empirical findings provide evidence that could be incorporated into more realistic models of investor learning. Standard belief-based rational Bayesian models assume that investors consider all information available about the investment alternative whether it has been obtained as a result of their choice or as a result of others’ choices (See: Brown and Koopmans, 1951; Cournot, 1960; Fudenberg et al., 1998). However, pure reinforcement learning models completely disregard information not directly experienced by investors (See: Erev and Roth, 1998). In the middle ground between the previous two, there are models that take a mixed approach. For instance, Camerer and Hua Ho (1999) produce a hybrid model that combines reinforcement learning and belief learning and situates each as a special case. But none of these models incorporate our main finding. The better processing of information if investors learn in an environment in which they can gather information only as a result of their endogenous choice.

In the financial decision-making literature, the evidence about the weight given to foregone values is also mixed. Experimental and eye-tracking data shows that people do give weight to foregone outcomes (for a review,

see: Plonsky and Teodorescu, 2020). Other studies find that people are as sensitive to foregone outcomes as obtained outcomes (Yechiam et al., 2005; Yechiam and Rakow, 2012). While, further research finds that people are even more sensitive to foregone outcomes (Yechiam et al., 2005; Yechiam and Rakow, 2012) than experienced ones. More recently, an eye-tracking study has found evidence that people are less sensitive to foregone than to obtained outcomes in line with the findings of the literature in economic games Ashby and Rakow (2016). Here we provide new evidence, using a belief approach, that captures the effects of availability of both: foregone outcomes and outcomes that come as a result of endogenous choice. We show that the probability updates that come as a result of the combination of the two types of outcome are, on average lower than the probability updates produced only as a result of endogenous choice.

The learning by doing management literature has also analyzed experiential learning. The typical finding is that experiential learning leads to biased inferences. Levitt and March (1988) find that experiential learning is hampered by the turnover of personnel and the passage of time. Moreover, learning is further complicated by the ecological structure of the simultaneously adapting behavior of other organizations and by an endogenously changing environment. A similar view is held by Levinthal and March (1993), who suggest a certain conservatism in expectations about the outcomes of organizational learning. This conservatism of expectations about the potential benefits of organizational learning is due to the challenges and difficulties that maintaining exploration in the face of a tendency to overinvest in exploitation presents. On another level, Lejarraaga (2010); Lejarraaga and Gonzalez (2011) presented research that tackled the description-experience duality in organizational behavior. This duality refers to the fact that managers can have good descriptions of investment alternatives but also evaluate them after they have some experience with their outcomes. The authors find that although descriptive information may be objective, individuals prefer to rely on experience, which provides rougher information but is easier to interpret. The reliance on experience can even be higher when facing complex decisions, which are very common in the organizational environment. In our study, we provide a more positive view of the results of organizational learning in selective feedback environments. We find that, although the challenges of balancing the information acquisition problem are significant and affect learning, the better information processing of information can help learners that rely on experience.

Elucidating the influence of selective feedback environments in investor's

learning can shed light on the importance of environments for human learning. Our research goes in line with the reviewed findings in Erev and Roth (2014). We add to their observation that “highlight conditions under which experience does not guarantee the highest efficiency, and also suggest that small modifications of the environment can increase efficiency”. The new knowledge obtained in this research can be incorporated into new theoretical models of learning that will help us predict the future choices and behavior of investors in new environments. The new insights can also be used to inform investors of the expected errors that they could make when facing such environments and help them avoid costly mistakes. Application of these new insights can also be employed in many other domains. There are many environments in life in which people cannot learn about the outcome of risky alternatives unless they choose them.

2.6 Conclusion

In many relevant investment environments, investors only have access to the outcomes of an investment alternative if they choose it. Therefore, the sample of information they use to evaluate the alternative is endogenously created by their own choices. They are in, what the literature calls, selective feedback environments. In this study, we find evidence that investors, on average, learn better about the quality of a risky investment alternative given the information they have observed in a selective feedback environment rather than in a full feedback environment. There is a better processing of information in environments in which the only route to obtaining information about an investment alternative is by endogenous choice, compared to environments in which investors can learn about the outcomes of an investment alternative even if they do not choose the alternative.

The evidence shows that investors’ beliefs in a selective feedback environment are closer to a rational benchmark than those of investors in a full feedback environment. This is particularly the case after both receiving negative information about the quality of the investment alternative and when investors beliefs about the quality of the alternative are optimistic before receiving the negative information. Here we show that even in the gain domain, experiencing an outcome lower than the expected value of the investment is sufficient to trigger superior adaptive learning by investors in a selective feedback environment compared to those in a comparable full feedback one. The more aggressive update in the selective

feedback environment after receiving negative information seems to be explained in part, the better learning in that environment compared to a full feedback one.

Moreover, if we compare the deviations of participants' beliefs about the quality of the asset in a condition with limited feedback to the objective Bayesian beliefs a perfect Bayesian learner would have produced if she had access to all the outcomes of the investment alternative, we find that there was no significant difference in learning. That is, in our particular experimental task, the better processing of information in the condition with limited feedback completely overcomes the sampling error generated by the loss of access to financial information and makes participants in the limited feedback condition learn, as well, as participants with access to full information about the outcomes of the stock.

The evidence presented supports that investors that find themselves in a selective feedback investment environment trigger different cognitive processes that will help them be closer to the objective, rational beliefs (given the information that they have seen) and not be necessarily farther away from objective beliefs than investors with full information even if they did not have access to the full information set. We consider this empirical finding and the unveiling of the mechanism that generates it our main contribution in this empirical work. The empirical findings that we reveal add another step to allow us to understand why the effects of foregone outcomes on maximization are not clear. Although the information effects of the missed information in selective feedback environments have traditionally been considered, the better processing of information that a more active learning style sparks were previously unaccounted for.

Finally, we have shown that the measures on which we base our main findings are related to important financial decisions outside the laboratory. Our findings contribute to the growing empirical literature on the effects of personal experience and attention into determining the choices and preferences of CEOs, asset managers, and individual investors. Here we show a channel that explains why personally experienced information influences financial decisions more than other available information. Moreover, our findings reveal how the features of learning environments, and not only the information available in them, affect decision making, not only in the financial domain but in all domains in life.

Chapter 3

SHORTSIGHTED MEDITATORS? THE EFFECTS OF MINDFULNESS ON INTER-TEMPORAL CHOICE

Joint with Daniel Navarro-Martínez and Jordi Quoidbach¹

In this study, we aim to understand the effects of mindfulness on inter-temporal decision-making. To do so, we run four studies, three in the lab and one in the field. The studies in the lab use a between-subjects design with three conditions. In the mindfulness condition, participants listen to a 15-minute audio mindfulness breathing exercise. In the mind-wandering condition, participants listened to 15-minute audio that repeatedly instructed participants to think of whatever came to mind. In the control condition, participants did not listen to any audio. After that, all participants made either 42 or 4 choices between receiving smaller cash amounts earlier and larger cash amounts later, or in the third lab study, responded to hypothetical but realistic scenarios in which inter-temporal decisions needed to be made. In the field experiment, participants completed an eight-week mindfulness training course from the largest provider of on-site

¹Data collection for Chapter 3 was funded by Fundacion BBVA-EI-2019-D.Navarro from Fundación BBVA to Daniel Navarro, Fundacion Ramon Areces 2019-Navarro from Fundación Ramón Areces to Daniel Navarro, and by PID2019-105249GB-I00 from the Ministerio de Ciencia Innovación y Universidades to Daniel Navarro and Gaël Le Mens.

mindfulness courses in Spain. We then collected participants' selections into four choices between receiving smaller cash amounts earlier and larger cash amounts later. Overall, we show that mindfulness does not affect inter-temporal decisions.

3.1 Introduction

It is 9 AM on a bright sunny day in San Jose, California. The morning silence is interrupted by a high-pitched sound. English teacher Argos Gonzalez faces his class—a mix of Hispanic and black 5-year-old children—with a balanced rounded metal bowl on an outstretched palm. He has just invited a Tibetan bell. The students are already familiar with its sound. They gently close their eyes, set their back straight, and start delicately and curiously observing the physical sensations created by their breathing.

Argos Gonzales is part of Mindful Schools, a not-for-profit training organization founded in 2007 as a program for a single school in Oakland, CA. Today, it is a training organization with online and in-person courses, content, and a network of mindful educators spanning all 50 U.S. states and 100+ countries. In 2015 it had trained more than 300,000 youth worldwide.

The creation of programs and organizations like Mindful Schools is the result of the work of Jon Kabat-Zinn, who, in the last 30 years, has popularized mindfulness worldwide. According to Jon Kabat-Zinn, "Mindfulness is the awareness that arises through paying attention, on purpose, in the present moment, non-judgmentally". But mindfulness is nothing new; it is a secular philosophy and set of techniques adapted from thousands-of-years-old Buddhist meditation traditions—ones that only recently landed in mainstream Western consciousness.

The scene described above is part of the mindfulness revolution that is still occurring today. Nowadays, we find Mindfulness practitioners in all kinds of Corporations and Institutions. Companies such as Apple, Procter & Gamble, General Mills, Google, and many others offer mindfulness courses, meditation retreats, and other related resources to their employees. Legal and law enforcement organizations are also showing interest in mindfulness. For instance, recently, Harvard Law School's Program on Negotiation hosted a workshop on "Mindfulness in the Law & Alternative Dispute Resolution." Mindfulness has even grabbed the attention of Governments; in 2014, the British Parliament organized a mindfulness session

for its members, and even the U.S. Army also provides mindfulness training to its soldiers. Recently, mindfulness is entering some schools in countries like the U.S. and the U.K, where Oxford researchers have recently announced plans to launch a large-scale, seven-year, \$10 million studies on mindfulness in education.

The benefits of mindfulness in certain areas have been established. Mindfulness has been frequently applied in clinical contexts for improving a variety of conditions such as anxiety and personality disorders, substance abuse, stress, as well as chronic pain, to name just a few (Sauer et al., 2013). Several reviews aggregate the findings and, over and above, support its clinical effectiveness (Bowen et al., 2006; Burke, 2010; Chiesa et al., 2011; Creswell, 2017; Fjorback et al., 2011; Hofmann et al., 2010; Mars and Abbey, 2010; Walach et al., 2012). Results from brain imaging studies have found a positive relationship between mindfulness and neural functions associated with well-being (Farb et al., 2007; Tang et al., 2007, 2009). Other studies have found positive effects on immune parameters (Carlson et al., 2004; Davidson et al., 2003; Witek-Janusek et al., 2008).

Recently, however, mindfulness has been facing various criticisms. Buddhist commentators have criticized the movement as being presented as equivalent to Buddhist practice; however, it is possibly denatured with undesirable consequences, such as being ungrounded in the traditional reflective morality and isolated from traditional Buddhist ethics (Shonin et al., 2015). This popularization of mindfulness as a "commodity" astray from the true Buddhist practice has been criticized, being termed "Mc-Mindfulness" by some critics (Bazzano, 2014). Some media reports have even pointed out some health risks reportedly caused by this practice. According to these reports, some meditators have suffered unexpected effects of increasing fear and anxiety, panic, or "meltdowns" after practicing, which could expose bipolar vulnerability or repressed PTSD symptoms (Foster, 2016).

Looking at these developments presented in the above paragraphs, we asked ourselves: if some of our politicians, corporate managers, military personnel, and even our children practice mindfulness, should not we know more about its effects? Should not we provide some rigorous scientific research to inform all these people about the consequences of making decisions after practicing mindfulness? Should not they know if this practice biases them in any direction?

In the present study, we aim to do that. We have decided to study the effects of mindfulness in inter-temporal decision-making. This is a par-

ticularly unexplored subject of research. Surprisingly, only a handful of research papers have dealt with the topic (Hendrickson and Rasmussen, 2013; Morrison et al., 2014; Hendrickson and Rasmussen, 2017; Yao et al., 2017). But crucially, none of these papers have focused on the effects of mindfulness in decision making per se or have used formal mindfulness training as a treatment. For instance, Yao et al. (2017) analyzes the impact of a mindfulness intervention on decision impulsiveness to curb Internet gaming disorder. However, the authors use as a treatment a combination of real therapy and mindfulness meditation, which does not allow for the isolation of the effects of the mindfulness treatment. Morrison et al. (2014) uses as a treatment an acceptance-based training session to learn its effects on decision impulsiveness. Although acceptance-based therapy is rooted in the philosophy that underpins the mindfulness movement, in the acceptance-based therapy used in the study, no formal mindfulness training was given to treated participants, which we believe is a severe limitation of the study. On the other hand, in Hendrickson and Rasmussen (2013, 2017) papers, participants received a manipulation in which a mindfulness practice was used. However, the authors were not interested directly in the effects of mindfulness on inter-temporal decision-making but studied its effects to reduce impulsive choices to curb obesity. That meant that the authors only produced a unique measure of inter-temporal discounting for money using a just one task² and only one type of mindfulness induction which reflected null effects and did not provide any insights about potential mechanisms that could explain their results.

That is why there is a need for a study that uses a pure decision-making perspective to investigate the effects of mindfulness on decision-making. We believe that a study that analyzes the effects of mindfulness on inter-temporal decisions with some of the various tools that are available for behavioral scientists in the decision-making field is needed to bring light to the topic. Our study improves on existing studies in the followings six ways. It is the first study that uses an intensive and very popular onsite eight-week course in mindfulness, the mindfulness-based stress reduction program (MBSR), designed by Jon Kabat-Zinn³ as treatment⁴. This compares to the most robust mindfulness manipulation up to now, which consisted of only a 50-minute prerecorded video on a mindfulness workshop focused on mindful eating (Hendrickson and Rasmussen, 2013, 2017).

²Which was the same in the two studies

³One of the figures that have popularized mindfulness techniques in the western world.

⁴in one of the studies

Second, in our study, we measure the trait level mindfulness of participants, and we are able to correlate this trait level with particular inter-temporal outcomes. This allows us to test whether a higher baseline mindfulness level is associated with particular inter-temporal choice behaviors. Third, we collect a measure of the experience in the mindfulness practice. By measuring how long participants have practiced mindfulness, we are able to test whether the duration of its practice is associated with particular inter-temporal behaviors.

Fourth, our study uses three different tasks to measure inter-temporal choice behaviors. Our three tasks vary, from classical choices between a smaller amount of money sooner to a larger amount of money later (as it was used in the best studies up to this point) to hypothetical real-world scenarios in which particular decisions must be taken based on these scenarios (which have never been used before). These new measures allow us to test previously unknown concepts; for instance, two of the tasks allow us to test whether the effects of mindfulness are different for choices that involve exclusively delayed rewards or a mix of delayed and immediate rewards; moreover, these tasks also allow us to test whether participants display time consistency in their choices. This compares to only one task to measure the effects of mindfulness on an inter-temporal choice, which was the best available in this type of study so far and which provided less varied insights about the relationship between mindfulness and inter-temporal choice (Hendrickson and Rasmussen, 2013, 2017).

Fifth, we measure inter-temporal choice behavior both in a laboratory setting and also outside the lab, which, once again, is novel in a study of this topic and has been proven to be needed since we have evidence that some lab findings do not correlate well with field behaviors. Sixth, contrary to the previous existing studies on the topic, which have been manipulation-based, our research combines manipulation and training-based studies. This is important since training based interventions can be more useful to target individuals for which delay aversion is the main reason for steep delay discounting and thus need a treatment that allows them to change behavior and not an only particular one-off choices, and the effects of mindfulness training on inter-temporal choice have never been studied before Scholten et al. (2019). Seventh, in our study, we present a mixture of purely hypothetical and potentially real tasks⁵ this is a novel feature since already existing studies only used purely hypothetical choices

⁵Participants knew that in tasks, they could obtain for real the outcome that they chose at the proposed delay.

and had a fixed participation payment.

Finally, our study not only studies the effects of a mindfulness-based intervention but also tests a particular mechanism that could explain the effects of mindfulness. More precisely, in our research, we test whether the potential effect of mindfulness on the inter-temporal decision is derived from choices that involve decisions between the present and a distant moment in time or is also present in choices between two delayed moments in time. That is, we are able to test whether mindfulness affects inter-temporal choices by altering the present bias of participants. Previous studies on the topic have not previously tested any mechanism that could explain their findings.

But why have we chosen to study the effects of mindfulness on inter-temporal decisions? Because of the following five reasons. First, it has received very little attention, given the importance of the practice of mindfulness in our society. Mindfulness interventions are used in clinical treatments (Dimidjian and Segal, 2015), the workplace (Good et al., 2015), in schools (Sibinga et al., 2016), the military (Johnson et al., 2014), and in prisons (Samuelson et al., 2007) and a deeper understanding of its effects is important. Second, because it has been shown that impatience in inter-temporal choices have important consequences for people's lives (e.g., Mischel and Ebbesen, 1970; Mischel et al., 2011), for instance, higher impatience is associated with important health outcomes like obesity (Komlos et al., 2004), substance and alcohol use disorders (Li and Sinha, 2008), but also with important life outcomes such as a divorce rate (De Paola and Gioia, 2017) or suicide rate (Wang et al., 2014). Third, focusing on the present moment is the cornerstone on which almost all the mindfulness techniques and Buddhist teachings are based (Kabat-Zinn, 1990; Hanh, 2010). Fourth, as we have previously pointed out, the existing studies have severe limitations. Fifth and final, most practitioners are only aware of the positive benefits of mindfulness and could fail to predict unintended consequences of its practice, which is reflected in the various null and negative effects of mindfulness interventions that target undesired outcomes in people's lives (Britton, 2019).

During the design of this study, we hypothesized that mindfulness would bias choices toward the present moment or towards more immediate rewards. This hypothesis was underpinned by the fact that focusing on the present moment is a key teaching of the prominent Buddhist philosophers and top mindfulness authors. However, here we show that mindfulness does not affect inter-temporal choices. To support this claim,

we provide both laboratory evidence and field evidence.

The remainder of the article is organized as follows. Section I describes the relevant literature. Section II describes the experimental design. Section III analyzes the results. Section IV concludes.

3.2 Inter-temporal Decisions

Delayed rewards, throughout most of history, have been assumed to be discounted at a constant rate over time. However, empirical advances in economics, neuroscience, and psychology have revealed a much more complex pattern (for a review: Berns et al., 2007). Four mechanisms have been identified as determinants of inter-temporal choice.

The first one is time discounting, which has been thoroughly studied in the three mentioned disciplines during the past and current century. Time discounting refers to the degree to which a future reward is discounted. It is well known that both animals and humans discount future rewards hyperbolically Herrnstein (1961). That is, hyperboloid discount functions represent well how both animals and humans behave. A feature of these functions is that they decay at a more rapid rate in the short run than in the long run. This means that hyperbolic discounters are more impatient when making short-run tradeoffs than when making long-run tradeoffs. However, researchers have found a large degree of variation regarding the dimension of time discounting both between and within species. For instance, cotton-top tamarin monkeys display a discount factor that sharply falls to zero after a delay of about one minute (Stevens et al., 2005) which reflects a much steeper temporal discount compared to that of humans. And within the human species, people with alcohol or substance dependence display higher time discounting than healthy controls (Li and Sinha, 2008).

A potential explanation for these differences within and between species has been identified. It has been shown that the differential size and functioning of the prefrontal cortex have an important role in shaping time discounting. For instance, humans present a disproportionately large prefrontal cortex compared to animals, and some researchers speculate that this allows humans to care more and discount less delayed outcomes (Cottle and Klineberg, 1974). Evidence of the role of the prefrontal cortex in explaining the difference in time discounting within the human species comes from studies with people who experience damage in the prefrontal

cortex (Cottle and Klineberg, 1974; Damasio, 1994) or that the present a differential neurological development of the prefrontal cortex (Durstun et al., 2002). All in all, it seems that the extent to which the prefrontal cortex is developed and used can account for why some individuals can make decisions valuing more delayed rewards.

The above findings could not account for the following empirical pattern. Why the same individual sometimes can present differential time discounting? More recently, a new theory has tried to reconcile these facts in a new theory that would encompass all the previous knowledge. In McClure et al. (2004) the authors, in a correlational study using functional magnetic resonance imaging, showed that two separate systems are involved when people make choices between monetary reward options that vary by delay to delivery. More precisely, parts of the limbic system associated with the dopamine system are preferentially activated by decisions involving immediately available rewards while, in contrast, regions of the lateral prefrontal cortex and posterior parietal cortex are engaged uniformly by inter-temporal choices irrespective of delay.

More recently, however, three sometimes competing mechanisms that are implemented in the brain: representation, anticipation, and self-control, have been identified as additional key influential factors that shape inter-temporal decision-making (Berns et al., 2007). Anticipation refers to an individual's propensity to imagine and experience pleasure and pain in anticipation of a future event. Self-control refers to the tensions that people experience when they attempt to implement a far-sighted decision in the presence of immediate temptation. Representation refers to the way that the brain interprets or frames a set of choices.

Overall, the literature on inter-temporal choice does not provide a clear guide on what to expect from the effects of mindfulness on inter-temporal decisions. Mindfulness effects would vary according to how it affects a variety of factors. Our main hypothesis in this study is that mindfulness manipulations and training bias people toward the present moment or towards more immediate rewards. This hypothesis is underpinned by the fact that focusing on the present moment is a crucial teaching of mindfulness authors, is present in many mindfulness inductions, and is a central concept in the mindfulness philosophy.

3.3 Experiment 1

3.3.1 Method

We assigned 323 participants⁶ (219 females, 104 males; mean age = 21 years, range = 18–51 years⁷) to mindfulness, mind wandering, and control conditions randomizing per session. Participants were students and local residents from the Universitat Pompeu Fabra Behavioral and Experimental Sciences Laboratory participant pool who responded to an advertisement offering €8 for participation. Each participant sat in a semi-private cubicle within a laboratory. Our mindfulness and mind-wandering induction procedures drew on established methods (Kabat-Zinn, 1990; Arch and Craske, 2006; Kiken and Shook, 2011; Hafenbrack et al., 2014).

Participants listened to a 15-min audio-recorded induction created specifically for this research by a professional mindfulness-meditation instructor. Participants were led through a focused-breathing meditation exercise that instructed them to focus on the physical sensations of breath entering and leaving their body and repeatedly reminded them to focus on their experience of breathing. The content of the mind-wandering induction repeatedly instructed participants to think of whatever came to mind. This type of induction has been used as a control condition in prior mindfulness experiments (Arch and Craske, 2006; Kiken and Shook, 2011; Hafenbrack et al., 2014) because it replicates a waking, baseline mental state (Mason et al., 2007).

Participants in the control condition were not subjected to any procedure. We decided to include this condition to improve on other mindfulness studies by trying to replicate a scenario without any manipulation in which participants only had to decide the task at hand.

Then participants in the first of the two waves of the experiment⁸, in which we had a total of 129 participants (78 females, 41 males; mean age = 22 years, range = 18–51 years⁹), had to complete three manipulation check items in order to advance in the experiment. The three items are included in section A of the appendix. The first two items were designed to test

⁶In two different waves

⁷One participant stated being one year old, but it must be an erroneous input since in the subject pool there are no children

⁸In the other wave and in the rest of the studies we did not include it again since we demonstrated the effectiveness of our manipulations in this first wave

⁹One participant stated being one year old but it must be an erroneous input

to which extent participants had been focused on the present moment or on the physical sensations in their bodies. The third item was designed to test to which extent participants had been mind-wandering. To measure the three items, we used a 5-point Likert scale.

After, all participants went through some items for an unrelated study and then made four hypothetical choices, inspired by tasks in Frederick et al. (2002), between two assets, asset A and asset B. Asset A always offered a smaller amount of money €200, but gave it sooner in time (either now or in 12 weeks). Asset B always gave a higher amount of money (either €220, or €250) but gave it at a more distant moment in time (either in 4 or in 16 weeks). Crucially, in 2 of the choices, asset A offered amounts only in the present moment, while in the other two choices, both Asset A and B offered delayed monetary rewards. Thus, participants made choices between smaller, immediate rewards and larger, later rewards or between smaller, later rewards and larger, even later rewards. Additionally, by both Asset A and Asset B offering in the first two choices, and then also in the last two choices, the same monetary amounts in different moments in time, we can compare whether participants behaved in a time-consistent manner.

3.3.2 Results and discussion


We start our analysis focusing on the manipulation check items. Overall, we find that both our mindfulness and mind-wandering manipulations were effective. The first two items, that were designed to measure the level of mindfulness of participants during the task shows that, mindfulness participants reported being significantly more focused on the present moment ($mean = 3.64$) than mind-wandering ($mean = 2.61$; $t = 4.87$, $p < .00001$) or control participants ($mean = 2.98$; $t = 3.00$, $p < .01$) did. But, as expected, there was no significant difference between the control and mind-wandering conditions ($t = 1.41$, ns). Results were similar for the second item, which measure the level of attention to the physical sensations of their body. Mindfulness participants reported being significantly more focused on the physical sensations of their body ($mean = 3.73$) than mind-wandering ($mean = 1.95$; $t = 8.81$, $p < .00001$) or control participants ($mean = 2.30$; $t = 6.35$, $p < .00001$) did. Moreover, once again, there was no significant difference between the control and mind-wandering conditions ($t = 1.36$, ns).

The third item of the manipulation check was designed to measure the

level of mind-wandering experienced during the experiment. The item showed that participants in the mind-wandering condition reported being significantly more freely mind-wandering ($mean = 3.47$) than mindfulness ($mean = 2.82$; $t = 2.76, p < .01$) or control participants ($mean = 2.83$; $t = 2.47, p < .02$) did. But, as expected, there was no significant difference between the mindfulness and control conditions ($t = -.05, ns$).

Figure 1 shows the proportion of participants that chose the distant choices combining the four items in Experiment 1. In the mindfulness condition the portion of participants choosing the distant alternatives (71%) was no different than that of mind-wandering participants (73%) $\chi^2(1, N = 412) = .23, ns$; and neither significantly different than the proportion of participants choosing distant choices in the control condition (72%) $\chi^2(1, N = 436) = .06, ns$. The difference in the proportion of people selecting the distant alternative between the mind-wandering and control conditions was also not statistically significant $\chi^2(1, N = 412) = .03, ns$.

Figure 3.1: Proportion of distant choices per condition




Notes: This figure shows the proportion of distant choices combining the 4 inter-temporal decisions as a function of the condition. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

Table I shows the analysis of distant choices in a regression form. Columns (1) and (2) show that there is no significant difference between the mindfulness and both the control and mind-wandering conditions if we control for the age and gender of the participants. Regarding included covariates, we observe a marginally significant positive association between gender

and the number of distant choices in experiment one ($0.0366, p < 0.1$).

Figure 2 shows the proportion of participants that chose the distant choices combining just the two items in Experiment 1 that involved choices in which the earlier choice was in the present moment. In the mindfulness condition the portion of participants choosing the distant alternatives (69%) was no different than that of mind-wandering participants (71%) $\chi^2(1, N = 206) = .13, ns$; and neither significantly different than the proportion of people choosing distant choices in the control condition (70%) $\chi^2(1, N = 218) = .05, ns$. The difference in the proportion of people selecting the distant alternative between the mind-wandering and control conditions was also not statistically significant $\chi^2(1, N = 206) = .00, ns$.

Figure 3.2: Proportion of distant choices per condition in the items involving decisions in the present



Notes: This figure shows the proportion of distant choices combining the 2 inter-temporal decisions in which participants had to choose between monetary outcomes in the present and in the future as a function of the condition. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

Columns (1) and (2) of Table II show the analysis of distant choices combining just the two items in Experiment 1 that involved choices in which the earlier choice was in the present moment in a regression form. These columns show that there is no significant difference between the mindfulness and both the control and mind-wandering conditions if we control for the age and gender of the participants. Regarding included covariates, we observe a significant positive association between age and the

Table 3.1: Distant Choices and Time-consistent Choices in the Mindfulness, Control, and Mind-wandering Conditions


Dependent Variable	<i>Distant Choices_i</i>		<i>Proportion Consistent_i</i>	
<i>Control_i</i>	0.0387 (0.21)	0.0238 (0.13)	0.0307 (0.73)	0.0364 (0.85)
<i>Mind-wandering_i</i>	0.0691 (0.38)	0.0444 (0.24)	0.00739 (0.18)	0.00953 (0.23)
<i>Age_i</i>		0.258 (1.60)		-0.0387 (-1.11)
<i>Gender_i</i>		0.0366* (1.92)		0.00172 (0.47)
Constant	2.853*** (21.91)	1.653*** (3.12)	0.794*** (25.84)	0.819*** (7.70)
R^2	0.000454	0.0162	0.00188	0.00611
Observations	323	323	323	323

Notes: The dependent variables in the regression models below are respectively, *Distant Choices_i* which is the number of choices of asset A^{10} that participant i made in the 4 choices of Experiment 1, *Proportion Consistent_i* which is the proportion of time consistent choices per participant in Experiment 1. The independent variables included are *Control_i* indicator variable, which is equal to one if participant i was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant i was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Age_i* which is the number of years old that participant i reports, and *Gender_i* indicator variable, which is equal to one if participant i reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. t -statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

number of distant choices ($0.0256, p < 0.01$) in experiment one.

Figure 3 shows the proportion of participants that chose the distant choices combining just the two items in Experiment 1 that involved choices in which both the earlier and the later alternative were not in a future moment in time as a function of the condition. In the mindfulness condition the portion of participants choosing the distant alternatives (74%) was no different than that of mind-wandering participants (75%) $\chi^2(1, N = 206) = .05, ns$; and neither significantly different than the proportion of people choosing distant choices in the control condition (74%) $\chi^2(1, N = 218) = .00, ns$. The difference in the proportion of people selecting the distant alternative between the mind-wandering and control conditions was also not statistically significant $\chi^2(1, N = 206) = .01, ns$.

Figure 3.3: Proportion of distant choices per condition in the items involving decisions between alternatives in the future



Notes: This figure shows the proportion of distant choices combining the 2 inter-temporal decisions in which both, the earlier and the later monetary outcome were in a future moment as a function of the condition. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

Columns (3) and (4) of Table II show the analysis of distant choices combining just the two items in Experiment 1 that involved choices in which both the earlier and the later alternative were not in a future moment in time as a function of the condition. These columns show that there is no significant difference between the mindfulness, and both, the control and mind-wandering conditions, if we control for the age and gender of

Table 3.2: Distant Choices Involving the Present Moment and not Involving the Present Moment in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	<i>Distant Choices Present_i</i>		<i>Distant Choices No Present_i</i>	
<i>Control_i</i>	0.0293 (0.28)	0.0282 (0.27)	0.00942 (0.09)	-0.00444 (-0.04)
<i>Mind-wandering_i</i>	0.0413 (0.40)	0.0289 (0.28)	0.0278 (0.27)	0.0155 (0.15)
<i>Age_i</i>		0.0256*** (2.73)		0.0109 (0.98)
<i>Gender_i</i>		0.105 (1.19)		0.153* (1.66)
Constant	1.376*** (18.16)	0.657** (2.45)	1.477*** (20.49)	0.996*** (3.21)
R^2	0.000532	0.0179	0.000233	0.0106
Observations	323	323	323	323

Notes: The dependent variables in the regression models below are respectively, *Distant Choices_i* which is the number of choices of asset A^{11} that participant i made in the 4 choices of Experiment 1, *Proportion Consistent_i* which is the proportion of time consistent choices per participant in Experiment 1. The independent variables included are *Control_i* indicator variable, which is equal to one if participant i was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant i was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Age_i* which is the number of years old that participant i reports, and *Gender_i* indicator variable, which is equal to one if participant i reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. t -statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.


the participants. Included covariates show a marginally significant positive association between gender and the number of distant choices (0.153, $p < 0.1$) in experiment one.

Figure 4 shows the proportion of participants that displayed consistent time choices in Experiment 1 as a function of the condition. This proportion is obtained by calculating the number of identical asset selections between the first and second choices, which displayed the same amounts in the two choices in different moments in time, and between the third and fourth choices, which also displayed the same amounts in the two choices in different moments in time, per condition—then dividing this number by two. Then we used a Wilcoxon Rank Sum Test with continuity correction to test whether the samples were likely to derive from the same population, and we observed that in the mindfulness condition, the portion of participants choosing consistent alternatives (79%) was no different than that of mind-wandering participants (80%; $W = 5640$, *ns*); and neither significantly different than the proportion of people choosing consistent alternatives in the control condition (82%; $W = 5746$, *ns*). The difference in the proportion of people selecting consistent alternatives between the mind-wandering and control conditions was also not statistically significant ($W = 6052$, *ns*).

Columns (3) and (4) of Table I show the analysis of the proportion of time-consistent choices in a regression form. These columns show that there is no significant difference between the mindfulness and both the control and mind-wandering conditions if we control for the age and gender of the participants. Included covariates do not show significant associations between gender or age and the proportion of time-consistent choices in a regression form.

Overall, experiment 1 shows that there is no significant difference in these classical inter-temporal decision items between any of the three conditions. That is true both when we consider differences in choices involving all types of delays or when we just consider differences when one of the alternatives offers a monetary reward in the present or in the future. Moreover, there is also no difference in the time consistency of choices between the three conditions.

Figure 3.4: Proportion of time consistent choices per condition



Notes: This figure shows the proportion of time consistent choices. This proportion is obtained by calculating the number of identical asset selections between the first and second choices, which displayed the same amounts in the two choices in different moments in time, and between the third and fourth choices, which also displayed the same amounts in the two choices in different moments in time. Then dividing this number by two and plotting this proportion as a function of the condition. Large numbers indicate more proportion of time consistent choices. Error bars indicate the 95% confidence interval.

3.4 Experiment 2

3.4.1 Method

To test whether the effects of mindfulness go over and above the classical intertemporal tasks in the decision-making literature, we decided to test whether the effects hold in tasks inspired by real-world scenarios. In Experiment 2, all procedures were generally the same as in the first experiment. We assigned 140 participants (83 females, 57 males; mean age = 21 years, range = 18–37 years¹²) to mindfulness, mind-wandering, and control conditions randomizing per session. Participants in the mindfulness and mind-wandering conditions completed the same mindfulness or mind-wandering induction procedure as in Experiment 1.

After, all participants completed three hypothetical tasks in which participants had to make choices between saving money for the future or spending it now. In the first task, participants were told to imagine that they had won €10,000 in the lottery and that they should allocate an amount

¹²One participant stated being four years old but it must be an erroneous input


that sums up to the €10,000 between two different categories: leisure activities or saving in a checking account. In the second task, participants were told to imagine that they earn €2,000 net per month, and they were asked to allocate an amount that sums up to the €2,000 between these different categories: saving in a checking account, pension plan, rent and home expenses, leisure, food, and other expenses. Finally, in the third task, participants were told to consider a scenario in which they need to purchase a car, but they only have half of the required purchase price. Then they have two options; first, borrow the missing money from a bank and pay interest for it; second, they wait two years until they have saved enough, and in this way, they do not have to pay any amount. In the B section of the appendix, we include the three tasks. The tasks were not directly incentivized. Participants only received the fixed payment for participation and a potential bonus in a later experiment in the same experimental session.

3.4.2 Results and discussion

Figure 5 gives the average proportion of money saved combining task 1 and task 2. To calculate this, we added the money amounts allocated to the saving alternative in the two tasks per condition and we divided this amount by two. Mindful participants allocated an equal proportion to saving (*mean* € = .541) than did mind-wandering participants (*mean* € = .503; $W = 1155$, *ns*), or control participants (*mean* € = .573; $W = 823$, *ns*). The difference in the proportion allocated to savings between mind-wandering participants and control participants was significant ($W = 1495$, *ns*).

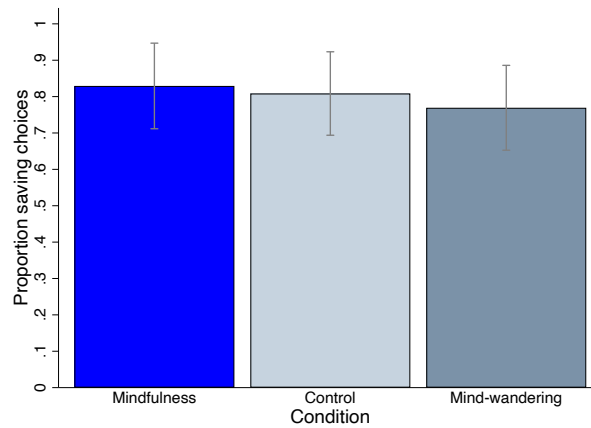
Figure 6 shows the proportion of participants that chose the savings alternative per condition in task 3. In the mindfulness condition the proportion of participants choosing the savings alternative (83%) was marginally different than that of mind-wandering participants (77%) $\chi^2(1, N = 41) = 2.83$, $p = .084$ and equal than the proportion of saving choices in the control condition (81%) $\chi^2(1, N = 41) = .86$, *ns*. The difference in the proportion of people selecting the saving alternative between the mind-wandering and control conditions was not significant $\chi^2(1, N = 47) = .85$, *ns*.

Figure 3.5: Proportion of money saved per condition in tasks 1 and 2



Notes: This figure shows the proportion of money allocated in saving alternatives as a function of the condition. Large numbers indicate more saving. Error bars indicate the 95% confidence interval.

Figure 3.6: Proportion of participants that chose the savings alternative per condition in task 3



Notes: This figure shows the proportion of saving choices per condition in task 3. Large numbers indicate more saving. Error bars indicate the 95% confidence interval.

Table III includes a regression analysis of the first two tasks. In columns (1) to (3), the regression analysis shows that controlling for the age, gender, and a measure of experience in the meditation practice of participants there is no significant difference in the amount of money allocated to the

saving alternatives between the mindfulness condition and the control or mind-wandering conditions. Moreover, Table I additionally shows that there is no significant association between previous meditation experience and the amount of money allocated to saving alternatives in any of the three experimental conditions.

Table IV shows a regression analysis of the third task. Columns (1) to (3) of the regression analysis show that controlling for the age, gender, and a measure of experience in the meditation practice of participants there is no significant difference in the probability of selecting the saving choice between the mindfulness condition and the control or mind-wandering conditions. Moreover, Table IV, once again, shows that there is no significant association between previous meditation experience and the probability of choosing the saving option in any of the three experimental conditions. However, column (3) shows that being a woman is significantly associated with a higher probability of choosing the savings choice (0.219, $p < 0.01$).

The results of Experiment 2 suggest that there are no significant differences between the mindfulness condition and the control and mind-wandering condition in the outcomes of tasks 1 and 2. However, the point estimates of task 3, a binary decision between consuming now and incurring a cost, or waiting and saving so that one does not need to borrow money, point to participants in the mindfulness condition being marginally more patient than participants in the mind-wandering condition but equally patient as participants in the control condition, however, once we control for age, gender and previous experience in the meditation practice that difference disappears.

3.5 Experiment 3

3.5.1 Method

In Experiment 3, all procedures were generally the same as in the first and second experiments. We assigned 389 participants (263 females, 126 males; mean age = 21 years, range = 18–51 years¹⁵) to mindfulness, mind wandering, and control conditions randomizing per session. Participants in the mindfulness and mind-wandering conditions completed the same mindfulness or mind-wandering induction procedure as in Experiments 1

¹⁵One participant stated being four years old but it must be an erroneous input

Table 3.3: Amount of Savings in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	<i>Amount of Savings_i</i>		
<i>Control_i</i>	423.9 (1.02)	700.3 (1.23)	775.1 (1.34)
<i>Mind-wandering_i</i>	-512.5 (-1.15)	-652.3 (-1.05)	-651.9 (-1.05)
<i>Practice Mindfulness_i</i>		91.58 (0.14)	57.30 (0.09)
<i>Control_i * Practice Mindfulness_i</i>		-778.8 (-0.88)	-762.1 (-0.86)
<i>Mind-wandering_i * Practice Mindfulness_i</i>		382.3 (0.41)	449.5 (0.50)
<i>Age_i</i>			-40.37 (-0.70)
<i>Gender_i</i>			402.6 (1.08)
Constant	8032.4*** (24.48)	7990.0*** (15.98)	8566.9*** (8.85)
R^2	0.0373	0.0507	0.0626
Observations	140	140	140

Notes: The dependent variable in the regression models below, *Amount of Savings_i* is the amount of savings that participant *i* allocated in both tasks 1 and 2 to the alternatives related to savings¹³. The independent variables included are *Control_i* indicator variable, which is equal to one if participant *i* was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant *i* was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Practice Mindfulness_i* indicator variable, which is equal to one if participant *i* reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, *Age_i* which is the number of years old that participant *i* reports, and *Gender_i* indicator variable, which is equal to one if participant *i* reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. *t*-statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

Table 3.4: Amount of Savings in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	<i>Savings Choice_i</i>		
<i>Control_i</i>	-0.0208 (-0.25)	0.0337 (0.29)	0.0574 (0.50)
<i>Mind-wandering_i</i>	-0.0600 (-0.72)	0.0710 (0.63)	0.0471 (0.45)
<i>Practice Mindfulness_i</i>		0.122 (1.05)	0.100 (0.90)
<i>Control_i * Practice Mindfulness_i</i>		-0.116 (-0.68)	-0.174 (-1.00)
<i>Mind-wandering_i * Practice Mindfulness_i</i>		-0.316* (-1.83)	-0.267 (-1.64)
<i>Age_i</i>			0.0146 (1.26)
<i>Gender_i</i>			0.219*** (3.01)
Constant	0.829*** (13.96)	0.773*** (8.46)	0.349 (1.26)
R^2	0.00392	0.0313	0.113
Observations	140	140	140

Notes: The dependent variable in the regression models below, *Amount of Savings_i* is the amount of savings that participant *i* allocated in both tasks 1 and 2 to the alternatives related to savings¹⁴. The independent variables included are *Control_i* indicator variable, which is equal to one if participant *i* was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant *i* was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Practice Mindfulness_i* indicator variable, which is equal to one if participant *i* reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, *Age_i* which is the number of years old that participant *i* reports, and *Gender_i* indicator variable, which is equal to one if participant *i* reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. *t*-statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

and 2. Once again, participants in the control condition were not subjected to any procedure.

After, all participants made 42 choices (McClure et al., 2004) between receiving smaller cash amounts (between €5 and €34) earlier (immediately, two weeks from the day of the experiment, or four weeks from the day of the experiment) and larger cash amounts (between €7 and €43) later (2, 4, or 6 weeks, respectively, from the day of the experiment). Thus, participants made choices between smaller, immediate rewards and larger, later rewards or between smaller, later rewards and larger, even later rewards. We incentivized participants to express their true preferences by randomly selecting 2 out of every 50 participants to realize one of his or her choices, paying that person his or her preferred alternative for a randomly selected choice pair. The payment was made at the chosen moment and for the chosen amount using Amazon.com gift certificates. All gift certificates were sent electronically.

This task will allow us to fit the parameters of a model that distinguishes between two types of processes that are represented in the quasi hyperbolic discounting function, $D(t) = \beta\delta^t$, for length of delay $t > 0$, and $D(0) = 1$ (Laibson, 1997; O'Donoghue and Rabin, 1999). One process (δ) reflects a time-consistent exponential discounting of rewards that is sensitive to the length of delay, t . The other process, present bias (β), discounts all future rewards when there is any delay (regardless of its length). We hypothesized that mindfulness (compared to the other conditions) would make people more patient but had no particular prediction whether it would impact more the present bias (β) or the time-consistent discounting (δ).


3.5.2 Results and discussion

To test whether the parameters differ between the three conditions, we fit the parameters in two different ways. In the first way, we fit each participant's choices to the quasi hyperbolic discounting function using maximum-likelihood estimation, constraining β and δ between 0 and 1. With this procedure, we get as many sets of parameters as participants in our study.

Then we can analyze the distribution of β parameters. In Figure 9 we observe that a significant portion of participants showed evidence of present bias by displaying a β lower than 1 (78%); β was marginally less likely to be lower than 1 among mindfulness participants (74%) than among

mind-wandering participants (77%) $\chi^2(1, N = 212) = 2.76, p = .097$ and significantly lower than the share of present biased participants of the participants in the control condition (82%) $\chi^2(1, N = 121) = 5.70, p < .05$. The difference in percentage of people with present bias between the mind-wandering and control conditions was not significant $\chi^2(1, N = 131) = 1.59, ns$.

Figure 3.7: Proportion of present biased participants per condition



Notes: This figure shows the proportion of individual present bias (β) parameters lower than 1 as a function of the condition. Large proportions indicate less patience. Error bars indicate the 95% confidence interval.

Table V shows a regression analysis of the probability of being present biased between the different conditions. Columns (1) to (3) show that after controlling for the age, gender, and a measure of experience in the meditation practice of participants, there is no significant difference in the probability of being present biased between the mindfulness condition and the control or mind-wandering conditions. Moreover, Table V also shows that there is no significant association between previous meditation experience, age, or gender and the probability of choosing the saving option in any of the three experimental conditions.

Figure 8 gives the mean values for both parameters in the three conditions. If we look at the results reflected in Figure 8 and we use Wilcoxon Rank Sum Test with continuity correction to test whether two samples are likely to derive from the same population we observe that mindful participants displayed less present bias (*mean* $\beta = .972$, *median* $\beta = .989$)


Table 3.5: Present Biased Participants in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	<i>Present Biased_i</i>		
<i>Control_i</i>	0.0813 (1.54)	0.0991 (1.34)	0.102 (1.37)
<i>Mind-wandering_i</i>	0.0382 (0.71)	0.0663 (0.89)	0.0757 (1.00)
<i>Practice Mindfulness_i</i>		0.0774 (0.97)	0.0900 (1.10)
<i>Control_i * Practice Mindfulness_i</i>		-0.0288 (-0.27)	-0.0393 (-0.37)
<i>Mind-wandering_i * Practice Mindfulness_i</i>		-0.0537 (-0.49)	-0.0617 (-0.56)
<i>Age_i</i>			-0.0114 (-1.35)
<i>Gender_i</i>			0.0198 (0.42)
Constant	0.736*** (18.27)	0.698*** (11.99)	0.918*** (4.68)
R^2	0.00617	0.0102	0.0183
Observations	389	389	389

Notes: The dependent variable in the regression models below, *Present Biased_i* is an indicator variable, which is equal to one if participant *i* has an individual β parameter lower than one, thus is what the literature calls a present biased individual, and zero if the participant has a β parameter equal to one. The independent variables included are *Control_i* indicator variable, which is equal to one if participant *i* was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant *i* was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Practice Mindfulness_i* indicator variable, which is equal to one if participant *i* reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, *Age_i* which is the number of years old that participant *i* reports, and *Gender_i* indicator variable, which is equal to one if participant *i* reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. *t*-statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

than did mind-wandering participants ($mean \beta = .95$, $median \beta = .987$; $W = 8629$, ns), discounting all nonimmediate rewards by almost 2 percentage points less than mind-wandering participants did. Mindfulness participants also displayed marginally less present bias than control participants ($mean \beta = .951$, $median \beta = .993$; $W = 8282$, ns), discounting all nonimmediate rewards by also almost 2 percentage points less than control participants did. The difference in the present bias parameter between mind-wandering participants ($mean \beta = .95$, $median \beta = .987$) and control participants was not significant ($W = 8930$, ns).

Figure 3.8: Average individual δ and β per condition



Notes: This figure shows the mean of the individual time-consistent annual exponential discount factor (δ) (left graph) and present bias (β) (right graph) as a function of the condition. Large numbers indicate more patience. Error bars indicate the 95% confidence interval.

Table VI shows a regression analysis of the present bias parameter (β) in the different conditions. Columns (1) to (3) show that after controlling for the age, gender, and a measure of experience in the meditation practice of participants, there is a significant difference in the present bias parameter of the mindfulness and the control conditions (-0.0338 , $p < 0.05$) and marginally significant between the mindfulness and mind-wandering conditions (-0.0240 , $p < 0.1$). That is, participants in the mindfulness condition, on average, discounted all nonimmediate rewards by three percentage points less than control participants did and two percentage points less than mind-wandering participants did.

When we compare conditions with decisions involving only choices between delayed rewards, mindful participants ($mean \delta = .310$, $median \delta = .154$) discounted already delayed rewards equally than mind-wandering

Table 3.6: Present Bias Parameter of Participants in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	$Beta_i$		
$Control_i$	-0.0203*	-0.0338**	-0.0338**
	(-1.88)	(-2.12)	(-2.10)
$Mind-wandering_i$	-0.0208*	-0.0241*	-0.0240*
	(-1.94)	(-1.71)	(-1.66)
$Practice\ Mindfulness_i$		-0.00773	-0.00785
		(-1.02)	(-1.04)
$Control_i * Practice\ Mindfulness_i$		0.0325	0.0321
		(1.63)	(1.57)
$Mind-wandering_i * Practice\ Mindfulness_i$		0.00671	0.00653
		(0.31)	(0.30)
Age_i			0.000231
			(0.22)
$Gender_i$			0.00259
			(0.24)
Constant	0.972***	0.975***	0.969***
	(259.78)	(220.60)	(40.30)
R^2	0.00914	0.0146	0.0148
Observations	389	389	389

Notes: The dependent variable in the regression models below, $Beta_i$ is the individual present bias β parameter fitted to the decisions made by participant i in Experiment 3. The independent variables included are $Control_i$ indicator variable, which is equal to one if participant i was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, $Mind-wandering_i$ indicator variable, which is equal to one if participant i was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, $Practice\ Mindfulness_i$ indicator variable, which is equal to one if participant i reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, Age_i which is the number of years old that participant i reports, and $Gender_i$ indicator variable, which is equal to one if participant i reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. t -statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

participants ($mean \delta = .245$, $median \delta = .104$); ($W = 9238$, ns). There were also no differences in discounting delayed rewards between mindful participants and control participants ($mean \delta = .256$, $median \delta = .126$; $W = 8757$, ns) and between mind-wandering and control participants ($W = 9045$, ns).

Table VII shows a regression analysis of the time-consistent discounting parameter (δ) in the different conditions. Columns (1) to (3) show that after controlling for the age, gender, and a measure of experience in the meditation practice of participants, there is a marginally significant difference in the δ parameter of the mindfulness and the control conditions (-0.103 , $p < 0.1$) and also marginally significant difference between the mindfulness and mind-wandering conditions (-0.111 , $p < 0.1$). That is, participants in the mindfulness condition, on average, discounted all already delayed rewards by ten percentage points less than control participants did and 11 percentage points less than mind-wandering participants did.

Figure 9 shows the average number of distant choices per condition. That is the number of choices in a more distant moment out of the 42 choices in Experiment 3. Performing t-tests to test for differences between the experimental conditions we find that mindful participants chose a marginally higher number of distant choices ($mean = 22.1$) than did mind-wandering participants ($mean = 19.8$; $t = -1.73$, $p < 0.1$). However the differences between the mindfulness participants and the control participants ($mean = 20.0$; $t = -1.57$, $p < 0.12$), was not significant. The difference in the number of distant choices between the mind-wandering participants and control participants was not significant ($t = -0.14$, ns).

Table VIII shows a regression analysis of the number of distant choices in Experiment 3. Columns (1) to (3) show that after controlling for the age, gender, and a measure of experience in the meditation practice of participants, participants in the mindfulness condition chose a significantly higher number of distant choices than participants in the control (-3.985 , $p < 0.05$) and mind-wandering conditions (-3.725 , $p < 0.05$). That is, participants in the mindfulness condition, on average, chose nearly four less distant choices than control participants did and also close to 4 less distant choices than mind-wandering participants did. Comparing these quantities to their respective means, it shows that mindfulness participants chose on average 20% less distant choices than control participants and 19% less distant choices than participants in the mind-wandering condition.

Table 3.7: Time-consistent Discounting Parameter of Participants in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	Δ_i		
$Control_i$	-0.0537 (-1.26)	-0.103* (-1.73)	-0.103* (-1.73)
$Mind-wandering_i$	-0.0644 (-1.56)	-0.109* (-1.90)	-0.111* (-1.92)
$Practice\ Mindfulness_i$		-0.0757 (-1.19)	-0.0802 (-1.26)
$Control_i * Practice\ Mindfulness_i$		0.108 (1.25)	0.107 (1.23)
$Mind-wandering_i * Practice\ Mindfulness_i$		0.0980 (1.18)	0.0984 (1.18)
Age_i			0.00523 (0.94)
$Gender_i$			0.0190 (0.54)
Constant	0.310*** (9.69)	0.346*** (7.31)	0.226* (1.83)
R^2	0.00726	0.0126	0.0155
Observations	389	389	389


Notes: The dependent variable in the regression models below, Δ_i is the individual time-consistent discounting parameter δ fitted to the decisions made by participant i in Experiment 3. The independent variables included are $Control_i$ indicator variable, which is equal to one if participant i was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, $Mind-wandering_i$ indicator variable, which is equal to one if participant i was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, $Practice\ Mindfulness_i$ indicator variable, which is equal to one if participant i reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, Age_i which is the number of years old that participant i reports, and $Gender_i$ indicator variable, which is equal to one if participant i reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. t -statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

Table 3.8: Average Number of Distant Choices in the Mindfulness, Control, and Mind-wandering Conditions

Dependent Variable	<i>Number of Distant Choices_i</i>		
<i>Control_i</i>	-2.146 (-1.57)	-4.006** (-2.15)	-3.985** (-2.13)
<i>Mind-wandering_i</i>	-2.334* (-1.74)	-3.768** (-2.03)	-3.725** (-1.99)
<i>Practice Mindfulness_i</i>		-1.796 (-0.92)	-1.800 (-0.92)
<i>Control_i * Practice Mindfulness_i</i>		4.315 (1.56)	4.195 (1.50)
<i>Mind-wandering_i * Practice Mindfulness_i</i>		3.307 (1.21)	3.248 (1.19)
<i>Age_i</i>			0.0315 (0.17)
<i>Gender_i</i>			0.646 (0.54)
Constant	22.11*** (22.63)	22.97*** (15.82)	21.87*** (5.46)
R^2	0.00925	0.0174	0.0182
Observations	389	389	389

Notes: The dependent variable in the regression models below, *Number of Distant Choices_i* is the number of distant choices in the 42 binary decisions that participant *i* chose in Experiment 3. The independent variables included are *Control_i* indicator variable, which is equal to one if participant *i* was allocated to the control condition and zero if the participant was allocated to the mindfulness or mind-wandering conditions, *Mind-wandering_i* indicator variable, which is equal to one if participant *i* was allocated to the mind-wandering condition and zero if the participant was allocated to the control or mindfulness conditions, *Practice Mindfulness_i* indicator variable, which is equal to one if participant *i* reports having any experience practicing meditation and zero if the participant reports no experience in meditation practice, *Age_i* which is the number of years old that participant *i* reports, and *Gender_i* indicator variable, which is equal to one if participant *i* reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. *t*-statistics are in parentheses. ***, ** indicate significance at the 1%, 5% and 10% level, respectively.

Figure 3.9: Average number of distant choices selected per condition

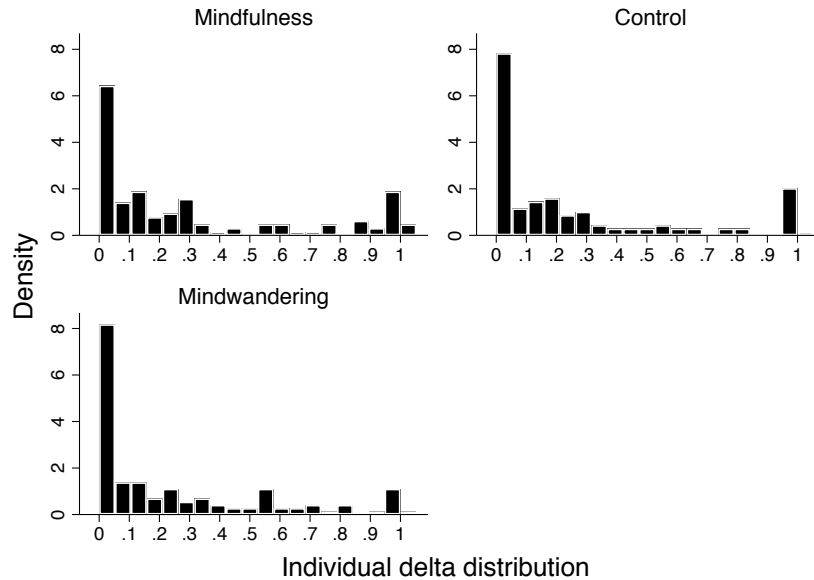


Notes: This figure shows the mean number of distant choices selected as a function of the condition. Large numbers indicate more patience. Error bars indicate the 95% confidence interval.

To further test whether the individual parameters of participants in the three conditions differ, we decided to evaluate their distributions. To do so, we produced a histogram for each of the experimental conditions and for each parameter. In Figure 10, we observe the histograms for the δ parameter. If we compare the plot of the mindfulness condition in the top left to those of the control and mind-wandering one, we observe that we tend to have a lower density for values close to zero in the mindfulness condition. Moreover, when we evaluate the values close to one, we also observe a higher density of values in that range for the mindfulness condition compared to the other ones. However, a Kolmogorov-Smirnov test suggested that the distribution of the δ parameters does not differ between the mindfulness and control conditions ($D(252) = .11$; $Z = .85$; *ns*), the mindfulness and the mind-wandering conditions ($D(258) = .11$; $Z = .84$; *ns*), and nor between the control and mind-wandering conditions ($D(268) = .07$; $Z = .09$; *ns*).

In Figure 11 we find the plots of the histograms for the β parameter. Analyzing the plots we observe that the β parameters in the mindfulness condition tend to have a higher density in values equal to one

Figure 3.10: Histogram of δ per condition

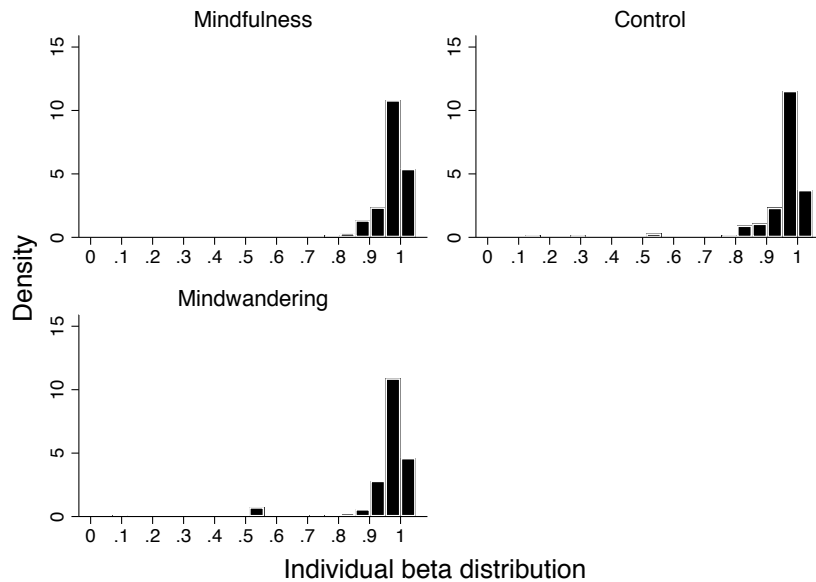


Notes: This figure shows the mean time-consistent annual exponential discount factor (δ) (left graph) and present bias (β) (right graph) as a function of the condition. Large numbers indicate more patience. Error bars indicate the 95% confidence interval.

than the plots in the other two conditions. However when we perform a Kolmogorov-Smirnov test, this suggests that the distribution of the β parameters does not differ between the mindfulness and control conditions ($D(252) = .09$; $Z = .37$; $p = .70$), the mindfulness and the mind-wandering conditions ($D(258) = .06$; $Z = .04$; $p = 0.97$), and nor between the control and mind-wandering conditions ($D(268) = .10$; $Z = .65$; $p = 0.52$).

All in all, mindfulness individual parameter estimations (compared to the mind-wandering and control conditions ones) decreased the desire to get something immediately but not sooner. However, when we test for differences in the distributions, we do not observe significant differences between the three experimental treatments. It seems thus that the difference in the present bias parameter between the mindfulness and the control and mind-wandering conditions comes from a subtle shift in the higher and lower bound of the range of the estimated betas that is not sufficient to create significant differences in the parameter distributions.

Figure 3.11: Histogram of β per condition



Notes: This figure shows the mean time-consistent annual exponential discount factor (δ) (left graph) and present bias (β) (right graph) as a function of the condition. Large numbers indicate more patience. Error bars indicate the 95% confidence interval.

To summarize, in this third experiment, we find three main findings. These findings point to an effect of mindfulness making participants in that condition more patient. First, if we compare the share of participants that experience present bias, we observe that the share of present biased participants is 3% and 8% lower in the mindfulness condition with respect to the mind-wandering and control one. Second, if we compare the average of the individually estimated betas per condition, we find that, on average, a participant in the mindfulness condition discounts all nonimmediate rewards by almost two percentage points less than mind-wandering or control participants. Third, when we analyze the average of the individual deltas per condition, we do not find significant differences. However, the point estimates of the averages suggest that participants in the mindfulness condition are more patient than participants in the other two conditions.

3.6 Experiment 4

3.6.1 Method

In experiment 4, we tested the effects of mindfulness on inter-temporal choices outside the lab. To that effect, we contacted the largest provider of the mindfulness-based stress reduction program (MBSR) in Spain, and we agreed on a partnership to do a field experiment on their premises. The MBSR program, according to the American Psychological Association, is a therapeutic intervention that involves weekly group classes and daily mindfulness exercises to practice at home over an 8-week period. MBSR taught people how to increase mindfulness through yoga and meditation and was developed by Jon Kabat-Zinn, the founder of the Stress Reduction Clinic and the Center for Mindfulness in Medicine, Health Care, and Society at the University of Massachusetts Medical School. And as of 2017, just at the University of Massachusetts, more than 24,000 people have taken this course which is the most popular standardized program on mindfulness worldwide.

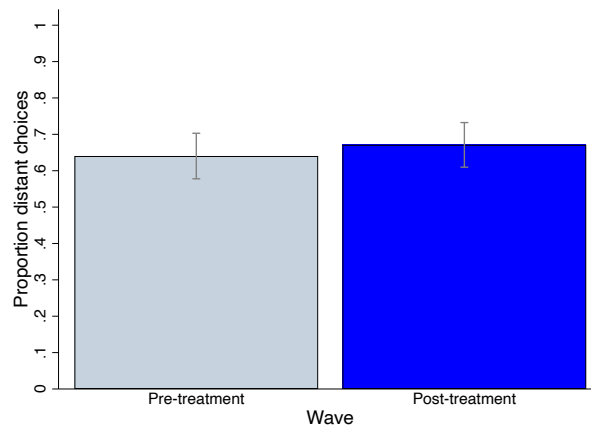
The study had a within-subject design; inter-temporal measures were taken prior to the course -before the start of the course in a guidance session- and during the last session of the course -after a guided meditation in the middle of the more than the 2 hours of the session. Participation in the study was voluntary, there was no direct monetary compensation for the participation, but a mindfulness book, written by the director of the center in which the program was taught, was offered as a gift for those that decided to participate in the two waves of the study. A total of 57 participants (37 females, 20 males; mean age = 43 years, range = 23–69 years) opted to participate in the study and completed the pre-treatment and post-treatment waves (the attrition in the second wave was 14%).

As measures of inter-temporal choice, we used the same four items of experiment 1. As in experiment 1, we also included the same unrelated tasks prior to their completion. We displayed the four inter-temporal choice items in a booklet format. In appendix D, we include a copy of the booklet used as support for our study, plus the instructions and consent form that participants completed are included in section D of the appendix.

3.6.2 Results and discussion

Figure 12 shows the proportion of participants that chose the distant choice combining the four items in Experiment 4. In the pre-treatment wave, the portion of participants selecting the distant alternatives (64%) was not significantly different than that of the participants in the post-treatment wave (67%) $\chi^2(1, N = 228) = .35, ns$.

Figure 3.12: Proportion of distant choices per wave



Notes: This figure shows the proportion of distant choices combining the 4 inter-temporal decisions as a function of the wave. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

Table IX shows the analysis of the number of distant choices in a regression form. Columns (1) and (2) show that there is no significant difference between the pre-treatment and post-treatment waves if controlling for participant fixed effects.

Results are similar if we calculate the proportion just using the two items in the study that involved choices in which the earlier choice was in the present moment. Figure 13 shows this. In the pre-treatment wave, the portion of participants choosing the distant alternatives (61%) was equal than that of the participants in the post-treatment wave (61) $\chi^2(1, N = 114) = .00, ns$.

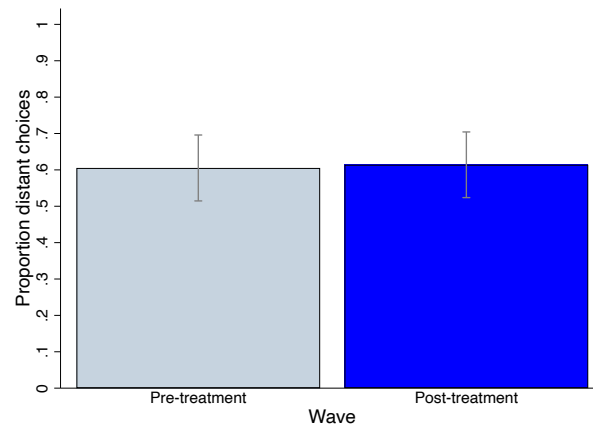
Columns (1) and (2) of Table X show the analysis of distant choices combining just the two items in Experiment 1 that involved choices in

Table 3.9: Distant and Time-consistent Choices in the Pre-treatment and Post-treatment waves

Dependent Variable	<i>Distant Choices</i> _{<i>iw</i>}		<i>Consistent Choices</i> _{<i>iw</i>}	
<i>Post-treatment</i> _{<i>iw</i>}	0.143	0.143	7.93e-18	0
	(0.85)	(0.60)	(0.00)	(0.00)
Constant	2.589***	2.589***	1.732***	1.732***
	(11.86)	(21.71)	(23.22)	(27.63)
Subject fixed effects	No	Yes	No	Yes
R^2	0.00202	0.847	2.00e-15	0.666
Observations	112	112	112	112

Notes: The dependent variables in the regression models below are respectively, *Distant Choices*_{*iw*} which are the number of choices of asset A¹⁶ that participant *i* made in the 4 choices of Experiment 4 in wave *w*, and *Consistent Choices*_{*iw*} which is the number of time consistent choices per participant in in wave *w* of Experiment 4. This number is obtained by calculating the number of identical asset selections between the first and second choices, which displayed the same amounts in the two choices in different moments in time, and between the third and fourth choices, which also displayed the same amounts in the two choices in different moments in time in each wave *w*. The independent variable included is *Post-treatment*_{*iw*} indicator variable, which is equal to one if participant *i* was deciding in the post-treatment wave *w* and zero if the participant was deciding in the pre-treatment wave *w*. Subject fixed effects are included in the second and fourth specifications. Standard errors are robust to heteroskedasticity and are clustered by participant. *t*-statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively.

Figure 3.13: Proportion of distant choices per condition in the items involving decisions in the present



Notes: This figure shows the proportion of distant choices combining the 2 inter-temporal decisions in which participants had to choose between monetary outcomes in the present and in the future as a function of the experimental wave. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

which the earlier choice was in the present moment in a regression form. These columns show that there is no significant difference between the pre-treatment and post-treatment waves if controlling for participant fixed effects.

However, the point estimates differ if we calculate the proportion just using the two items in the study that involved choices in which both the earlier and the later alternative were not in the present moment. Figure 14 shows this. Although if we look at the statistical significance, in the pre-treatment wave, the portion of participants choosing the distant alternatives (68%) was equal to that of the participants in the post-treatment wave (73%) $\chi^2(1, N = 114) = .52, ns$.

Columns (3) and (4) of Table X show a regression analysis of the distant choices combining the two items in Experiment 4 that involved choices in which both earlier alternatives were not in the present moment. These columns reveal that, once again, there is no significant difference between the pre-treatment and post-treatment waves controlling whether we control for participant fixed effects or not.


Figure 15 shows the proportion of participants that displayed time-

Table 3.10: Distant Choices Involving the Present Moment and not Involving the Present Moment in the Pre-treatment and Post-treatment waves

Dependent Variable	<i>Distant Choices Present_{iw}</i>		<i>Distant Choices No Present_{iw}</i>	
<i>Post-treatment_{iw}</i>	0.0179 (0.19)	0.0179 (0.13)	0.125 (1.22)	0.125 (0.86)
Constant	1.232*** (10.26)	1.232*** (18.53)	1.357*** (12.03)	1.357*** (18.71)
Subject fixed effects	No	Yes	No	Yes
R^2	0.000106	0.840	0.00566	0.792
Observations	112	112	112	112

Notes: The dependent variables in the regression models below are respectively, *Distant Choices Present_{iw}* which are the number of choices of asset A¹⁷ that participant *i* made in the 4 choices of Experiment 4 in wave *w*, and *Distant Choices No Present_{iw}* which is the number of time consistent choices per participant in in wave *w* of Experiment 4. This number is obtained by calculating the number of identical asset selections between the first and second choices, which displayed the same amounts in the two choices in different moments in time, and between the third and fourth choices, which also displayed the same amounts in the two choices in different moments in time in each wave *w*. The independent variable included is *Post-treatment_{iw}* indicator variable, which is equal to one if participant *i* was deciding in the post-treatment wave *w* and zero if the participant was deciding in the pre-treatment wave *w*. Subject fixed effects are included in the second and fourth specifications. Standard errors are robust to heteroskedasticity and are clustered by participant. *t*-statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively.

Figure 3.14: Proportion of distant choices per condition in the items involving decisions between alternatives in the future



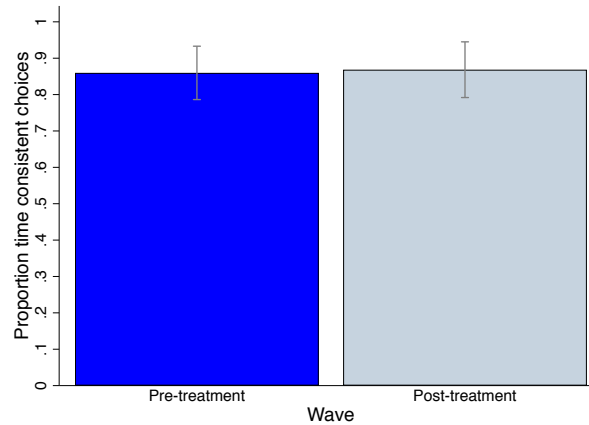
Notes: This figure shows the proportion of distant choices combining the 2 inter-temporal decisions in which both, the earlier and the later monetary outcome were in a future moment as a function of the experimental wave. Large numbers indicate more proportion of distant choices. Error bars indicate the 95% confidence interval.

consistent choices in Experiment 4 as a function of the experimental wave. This proportion is obtained with the same method as in the first experiment, but instead of doing the calculation by the condition, this time, we do it per experimental wave. We then use a Wilcoxon Rank Sum Test with continuity correction to test whether the samples are likely to derive from the same population. And we observed that in the pre-treatment wave, the portion of participants choosing consistent alternatives (86%) was no different than that of the post-treatment wave (87% $W = 1625, ns$).

Table IX shows the analysis of consistent choices in a regression form. Columns (3) and (4) show that there is no significant difference between the pre-treatment and post-treatment waves controlling for participant fixed effects.

We then explored whether mindfulness as a trait and exposure to meditation influence the inter-temporal decisions of the participants. Columns (1) and (2) of Table XI show that while there is no significant association between the analysis of distant choices and the accumulated score in the mindful attention awareness scale (MAAS), when we control for the months of meditation practice, age, and gender of participants, there is

Figure 3.15: Proportion of time consistent choices per wave



Notes: This figure shows the proportion of time consistent choices. This proportion is obtained by calculating the number of identical asset selections between the first and second choices, which displayed the same amounts in the two choices in different moments in time, and between the third and fourth choices, which also displayed the same amounts in the two choices in different moments in time. Then dividing this number by two and plotting this proportion as a function of the experimental wave. Large numbers indicate more proportion of time consistent choices. Error bars indicate the 95% confidence interval.

a significant positive association between the months of meditation practice prior to the treatment and the number of distant choices in the pre-treatment wave of Experiment 4 when we control for the same covariates and the cumulative score in MAAS. Moreover, when we look at the coefficients of the included covariates, we observe a significant positive relationship between gender and the dependent variable. That is, being a woman is positively associated with the number of distant choices in the pre-treatment wave.

Columns (3) to (5) of Table XI show that there is a significant negative association between the MAAS and the number of distant choices in the post-treatment wave if we control for the months of meditation practice, age, and gender of participants. Additionally, it shows the same positive association between the months of meditation experience and the post-treatment number of distant choices when we control for the same covariates and the cumulative score in MAAS. However this association becomes marginal when we control for the number of distant choices in the pre-treatment wave.

Combining just the 2 items in Experiment 1 that involved choices in which the earlier choice was in the present moment in a regression form. These columns show that there is no significant difference between the pre-treatment and post-treatment waves if controlling for participant fixed effects.

Table 3.11: Mindfulness Trait and Meditation Experience Effects on Pre and Post-treatment Number of Distant Choices

Dependent Variable	<i>Number of Distant Choices</i>		<i>Number of Distant Choices</i>		
	<i>Pre-treatment_i</i>		<i>Post-treatment_i</i>		
<i>MAAS_i</i>	0.00546 (0.28)	0.00108 (0.06)	-0.0254* (-1.75)	-0.0255* (-1.68)	-0.0262** (-2.10)
<i>Months Meditation Practice_i</i>	0.00318* (1.68)	0.00394** (2.04)	0.00492*** (4.55)	0.00525*** (3.75)	0.00239* (1.94)
<i>Age_i</i>		0.0135 (0.69)		-0.00438 (-0.20)	-0.0142 (-1.01)
<i>Gender_i</i>		0.929** (2.19)		0.139 (0.30)	-0.536* (-1.75)
<i>Number of Distant Choices Pre-treatment_i</i>					0.726*** (7.56)
Constant	2.254** (2.10)	1.575 (1.32)	4.075*** (5.03)	4.216*** (3.54)	3.072*** (2.80)
R^2	0.0127	0.0919	0.0584	0.0609	0.570
Observations	56	56	56	56	56

Notes: The dependent variables in the regression models below are respectively, *Number of Distant Choices Pre-treatment_i* which are the number of choices of asset A¹⁸ that participant *i* made in the 4 choices of Experiment 4 in the pre-treatment wave, and *Number of Distant Choices Post-treatment_i* which are the number of time consistent choices per participant in the post-treatment wave of Experiment 4. The independent variables include the *MAAS_i*, which is the cumulative score of participant *i* in the mindful attention awareness scale (MAAS) in the pre-treatment wave, as well as *Months of Meditation Practice_i*, which are the number of months of meditation practice stated in the pre-treatment wave, and both *Age_i* which is the number of years old that participant *i* reports, and *Gender_i* indicator variable, which is equal to one if participant *i* reports being a woman and zero if the participant reports being a man. Standard errors are robust to heteroskedasticity. *t*-statistics are in parentheses. ***, **, * indicate significance at the 1%, 5% and 10% level, respectively.

Overall, experiment 4 shows that there is no significant difference in these classical inter-temporal decision items in a field experiment that used a very popular eight-week intensive program as a base for the treatment

in the study. This is true both when we consider differences in choices involving all types of delays or when we just consider differences when one of the alternatives offers a monetary reward in the present. However, although not statistically significant, we find point estimate differences in the proportion of distant choices when we only analyze choices involving decisions in which both alternatives that participants can choose are in the future. Those differences point to participants, after receiving the mindfulness training becoming more patient. After checking time consistency, we find that there is no difference in this regard between experimental waves.

3.7 Conclusion

This study shows that mindfulness does not affect inter-temporal choices. Although focusing on the present moment is a vital teaching of the prominent Buddhist philosophers, best mindfulness authors, and even the Buddha himself, as attested by quotes such as: “When we are mindful, deeply in touch with the present moment” or “The past is gone, the future is not yet here, and if we do not go back to ourselves in the present moment, we cannot be in touch with life.” of Thich Nhat Hanh. Or like the ones of the Buddha himself: “The secret of health for both mind and body is not to mourn for the past, nor to worry about the future, but to live the present moment wisely and earnestly.” or “Do not dwell in the past, do not dream of the future, concentrate the mind on the present moment.” Here we show that mindfulness practitioners do not seem to be biased toward the present moment. Mindfulness practitioners do not prefer more immediate rewards compared to those of mind-wandering or control participants despite the protagonist role that the focus on the present moment has on the philosophy behind the modern mindfulness practice. Lots of new opportunities are ahead in this area, and lots of work is still needed to better understand the mechanisms behind this thousands-of-years-old tradition, the importance of which is growing day by day.

Appendix A

APPENDIX - FINTECH, BANK BRANCH CLOSINGS, AND MORTGAGE MARKETS

A.1 Spatial Variation

Figure A.1: Spatial variation in fintech mortgage originations market share 2008 and 2009

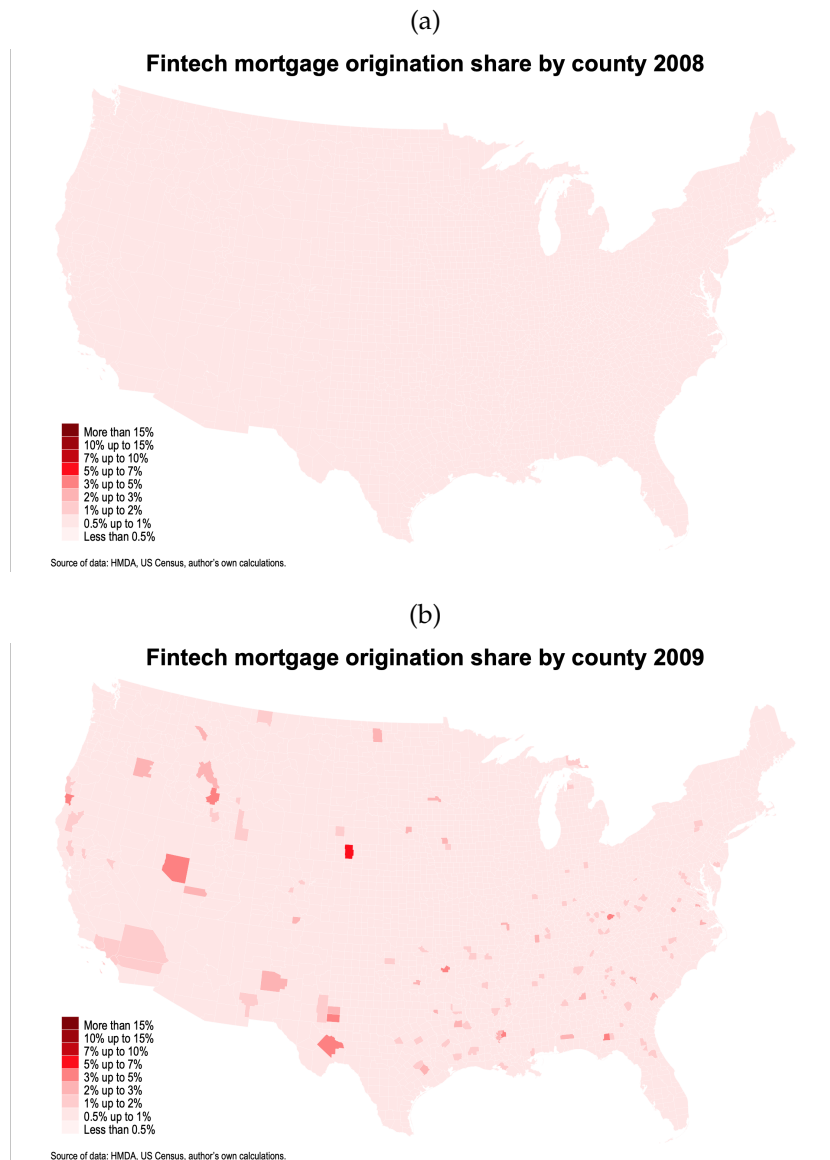



Figure A.2: Spatial variation in fintech mortgage originations market share 2010 and 2011

(a)



(b)

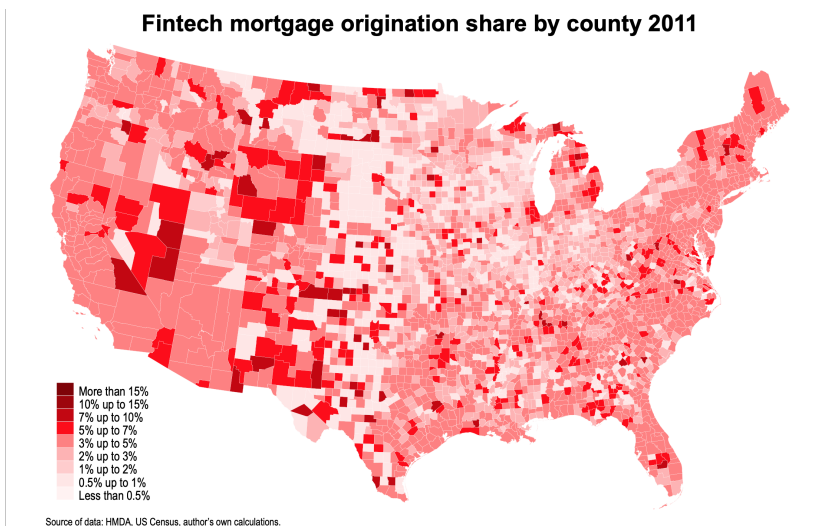


Figure A.3: Spatial variation in fintech mortgage originations market share 2012 and 2013

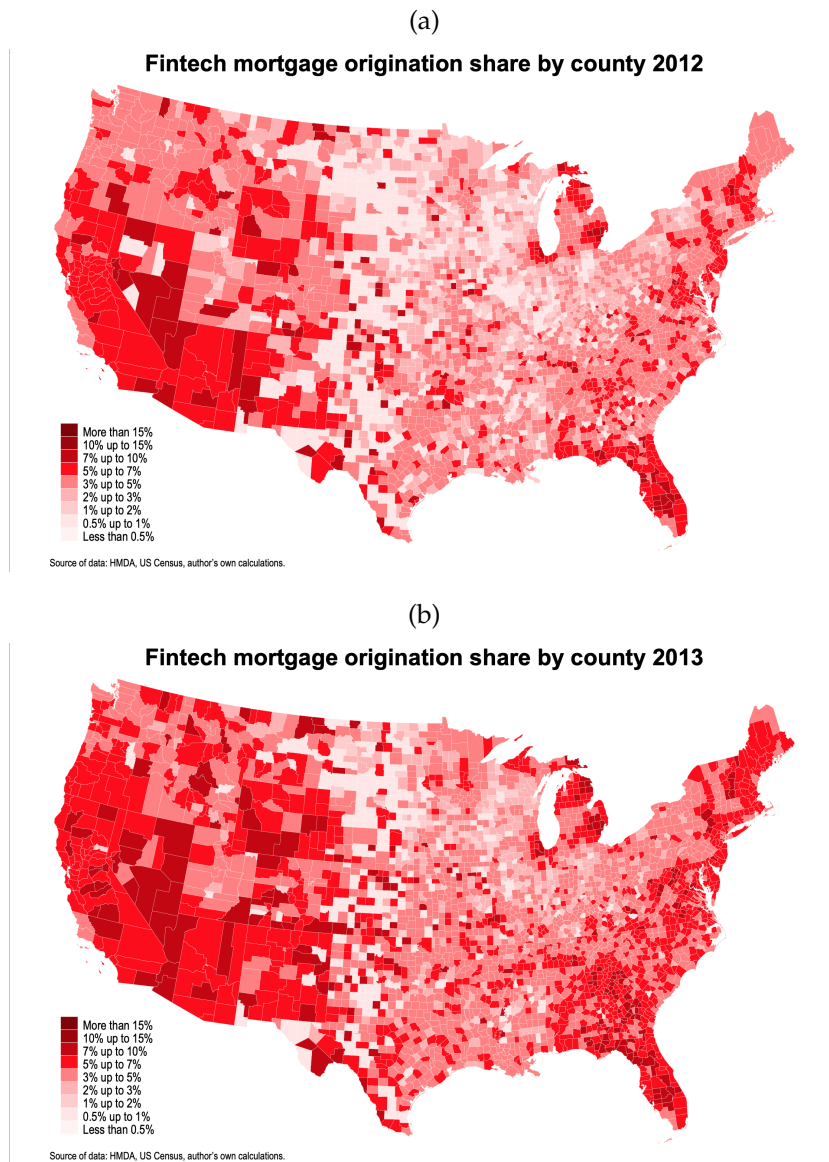



Figure A.4: Spatial variation in fintech mortgage originations market share 2014 and 2015

(a)


Fintech mortgage origination share by county 2014



Source of data: HMDA, US Census, author's own calculations.

(b)

Fintech mortgage origination share by county 2015



Source of data: HMDA, US Census, author's own calculations.

Figure A.5: Spatial variation in fintech mortgage originations market share

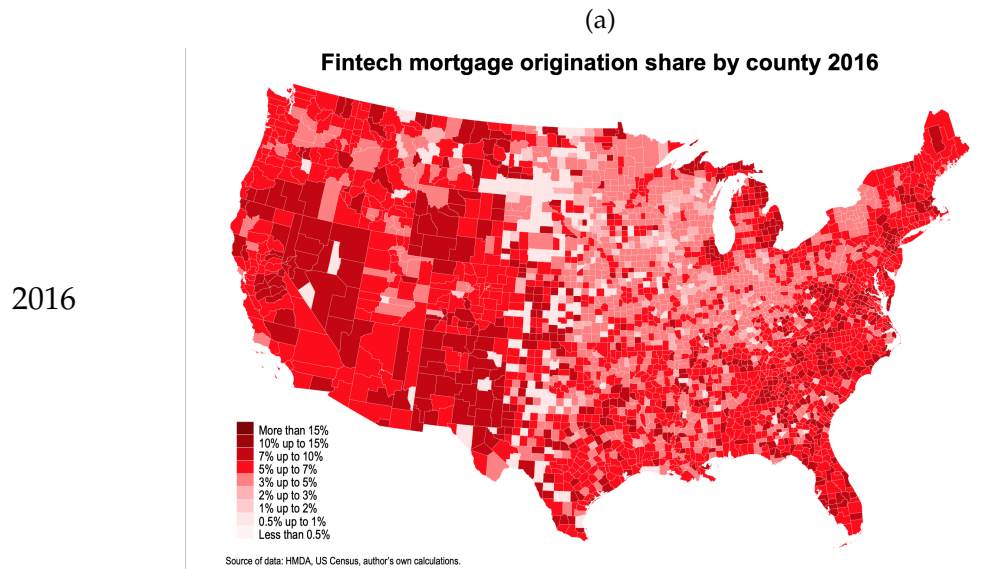


Figure A.6: Geographic distribution of sample states

(a)

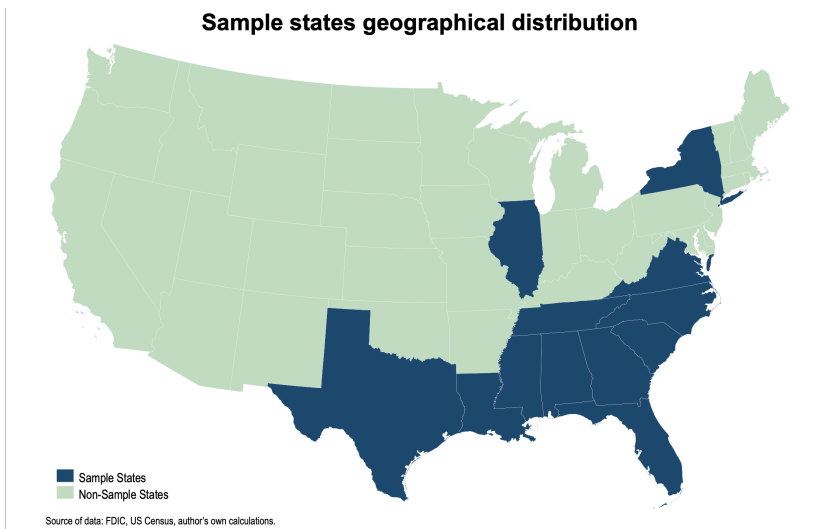


Figure A.7: Geographic distribution of bank branches in 2009

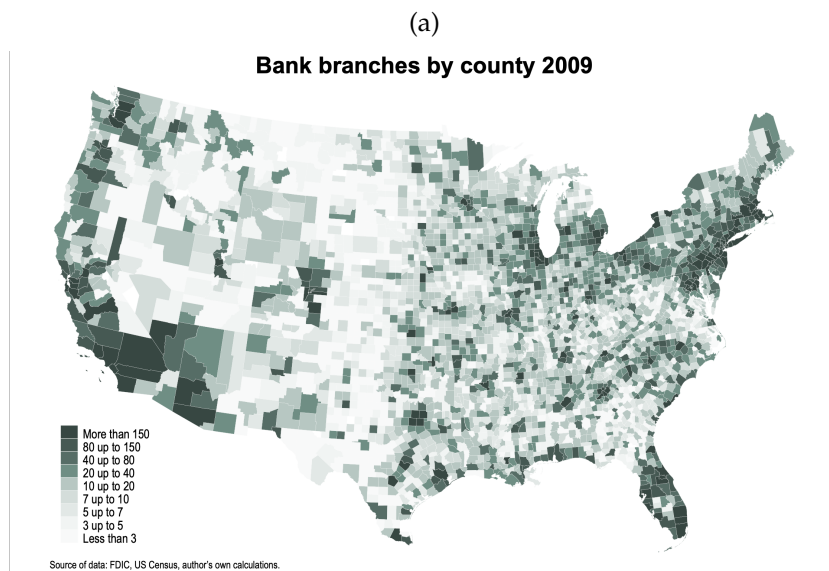


Figure A.8: Geographic distribution of bank branches per 100,000 people in 2009

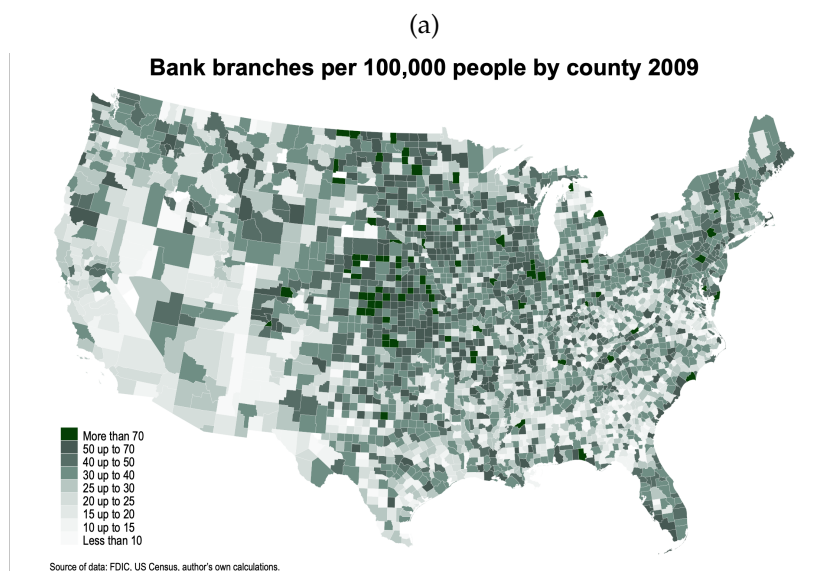


Figure A.9: Geographic distribution of US population 2000 US Census

(a)

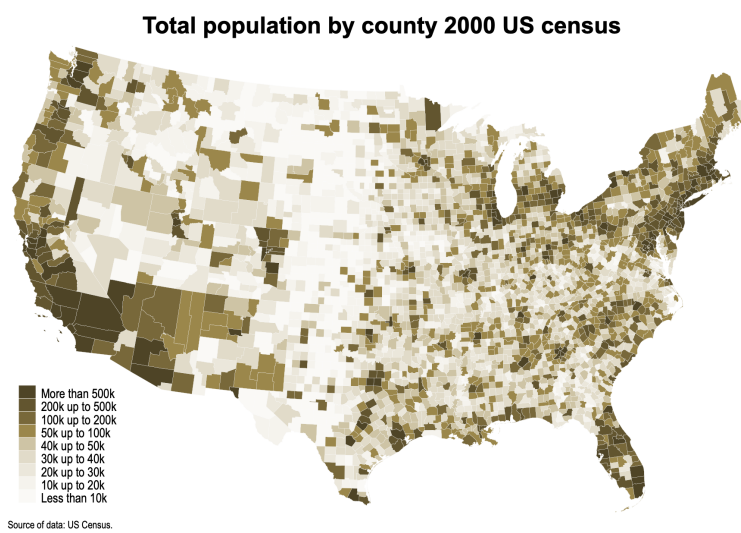


Figure A.10: Geographic distribution of branch closings since 2009 up until 2016

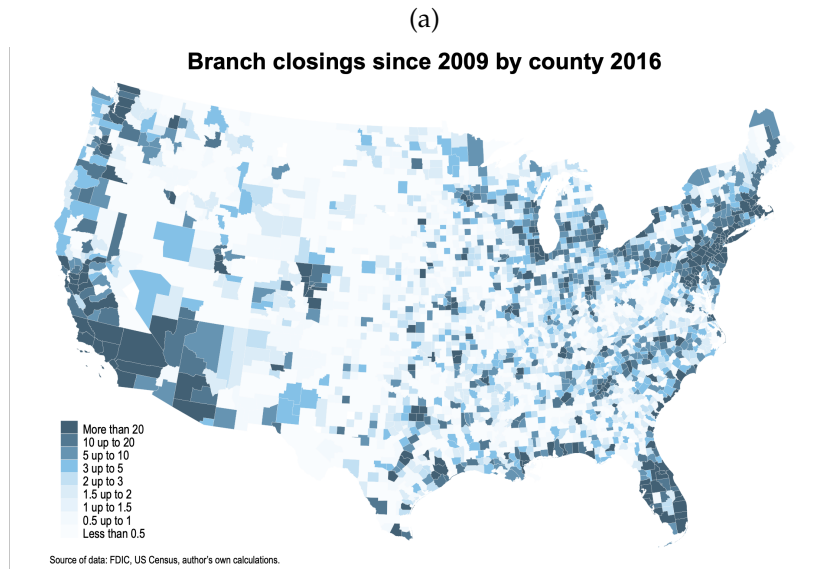


Figure A.11: Geographic distribution of branch closings per 100,000 people since 2009 up until 2016

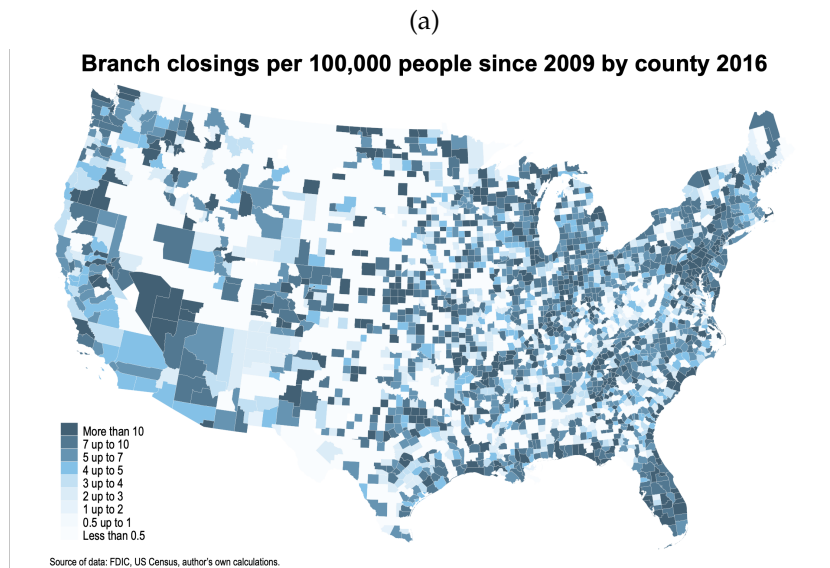
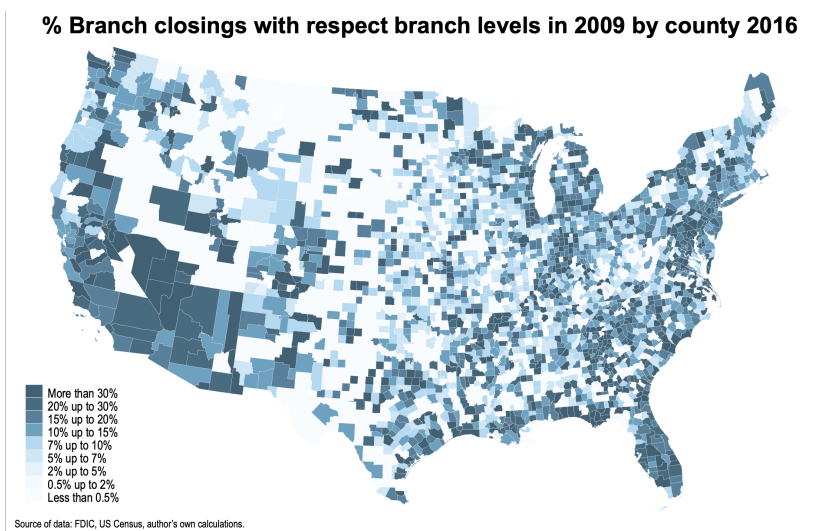


Figure A.12: Geographic distribution of the share of branch closings with respect to 2009 bank branch levels in 2016

(a)



A.2 Summary statistics

Table A.1: Tract Summary Statistics: Exposed vs. All other

Variable	(1) Exposed	(2) All other	(3) <i>p</i> -value on difference
Population	5,752 [3,766]	4,702 [2,523]	0.000
Population density	2,498 [3,544]	6,534 [16,460]	0.111
Percent minority	28.1 [23.2]	43.2 [32.0]	0.000
Percent college educated	59.1 [19.1]	52.5 [20.4]	0.000
Percent poverty level	13.3 [11.6]	14.5 [12.9]	0.025
Percent rural population	5.8 [16.1]	8.7 [23.8]	0.000
Percent population 65 and over	16.3 [11.9]	13.0 [10.1]	0.000
Percent unemployed	5.9 [6.4]	6.9 [7.0]	0.095
Median income (000s)	56.43 [27.64]	51.40 [26.63]	0.000
Percent MSA median income	117.1 [49.6]	102.1 [50.2]	0.000
Total branches	6.9 [4.5]	1.2 [1.8]	0.000
Branch growth	0.041 [0.114]	0.032 [0.187]	0.022
Bank mortgages	339.5 [480.2]	211.0 [229.8]	0.000
Shadow bank mortgages	119.6 [189.8]	85.8 [102.1]	0.000
Fintech mortgages	0 [0]	0 [0]	n.a.
Observations	418	11,737	

Notes: Standard deviations are in brackets. Column 3 reports the *p*-value for the difference between columns 1 and 2. Here *p*-values are obtained from a regression of tract characteristics on an indicator for being an exposed tract and county fixed effects. Population density is per square mile. Percent MSA median income is the ratio of tract median income to MSA median income. Growth rates are the average annual growth rates over the two years preceding the merger approval. All demographic variables are as of the 2000 census. Credit variables are as of the year before federal merger approval.

Appendix B

APPENDIX - HOW SELECTIVE ACCESS TO FINANCIAL INFORMATION AFFECTS HOW INVESTORS LEARN

B.1 Participant Instructions Full feedback Con- dition

Welcome to our financial decision making study!

In this study you will work on an investment task. In this task you will repeatedly invest in one of two securities: a risky security (i.e., a stock with risky payoffs) and a riskless security (i.e., a bond with a known payoff), and will provide estimates as to how good an investment the risky security is.

In either task, there are two types of conditions you can face: the GAIN and the LOSS conditions. In the GAIN condition, the two securities will only provide POSITIVE payoffs. In the LOSS condition, the two securities will only provide NEGATIVE payoffs.

Specific details for the GAIN condition:

In the GAIN condition, on any trial, if you choose to invest in the bond, you get a payoff of \$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either \$10 or \$2. The stock can either be good or bad, and this will determine the likelihood

of its dividend being high or low. If the stock is good then the probability of receiving the \$10 dividend is 70% and the probability of receiving the \$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being \$10 are 70%, and the odds of it being \$2 are 30%. If the stock is bad then the probability of receiving the \$10 dividend is 30% and the probability of receiving the \$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being \$10 are 30%, and the odds of it being \$2 are 70%.

Specific details for the LOSS condition:

In the LOSS condition, on any trial, if you choose to invest in the bond, you get a payoff of -\$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either -\$10 or -\$2.

The stock can either be good or bad, and this will determine the likelihood of its dividend being high or low. If the stock is good then the probability of receiving the -\$10 dividend is 30% and the probability of receiving the -\$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being -\$10 are 30%, and the odds of it being -\$2 are 70%. If the stock is bad then the probability of receiving the -\$10 dividend is 70% and the probability of receiving the -\$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being -\$10 are 70%, and the odds of it being -\$2 are 30%.

In both the GAIN and LOSS conditions:

In each condition, at the beginning of each block of six trials, you do not know which type of stock the computer selected for that block. You may be facing the good stock, or the bad stock, with equal probability.

On each trial in the block you will decide whether you want to invest in the stock for that trial and accumulate the dividend paid by the stock, or invest in the riskless security and add the known payoff to your task earnings.

You will then see the dividend paid by the stock, no matter if you chose the stock or the bond.

You will then have to tell us what you think is the probability that the stock is the good one (the answer must be a number between 0 and 100 – do not add the % sign, just type in the value).

You will also have to tell us the same thing at the beginning of each block before making any choice.

There is always an objective, correct, probability that the stock is good, which depends on the history of dividends paid by the stock already. For instance, at the beginning of each block of trials, the probability that the stock is good is exactly 50%, and there is no doubt about this value.

As you observe the dividends paid by the stock you will update your belief whether or not the stock is good. It may be that after a series of good dividends, you think the probability of the stock being good is 75%. However, how much you trust your ability to calculate this probability could vary. Sometimes you may not be too confident in the probability estimate you calculated and sometimes you may be highly confident in this estimate. For instance, at the very beginning of each block, the probability of the stock being good is 50% and you should be highly confident in this number because you are told that the computer just picked at random the type of stock you will see in the block, and nothing else has happened since then.

Throughout the task you will be told how much you have accumulated through dividends paid by the stock or bond you chose up to that point.

Your final pay for completing the investment tasks will be:

$\$5 + 1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.

Thank you!

B.2 Participant Instructions Selective feedback Condition

Welcome to our financial decision making study!

In this study you will work on an investment task. In this task you will repeatedly invest in one of two securities: a risky security (i.e., a stock with risky payoffs) and a riskless security (i.e., a bond with a known payoff), and will provide estimates as to how good an investment the risky security is.

In either task, there are two types of conditions you can face: the GAIN and the LOSS conditions. In the GAIN condition, the two securities will only provide POSITIVE payoffs. In the LOSS condition, the two securities will only provide NEGATIVE payoffs.

Specific details for the GAIN condition: In the GAIN condition, on any trial, if you choose to invest in the bond, you get a payoff of \$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either \$10 or \$2.

The stock can either be good or bad, and this will determine the likelihood of its dividend being high or low. If the stock is good then the probability of receiving the \$10 dividend is 70% and the probability of receiving the \$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being \$10 are 70%, and the odds of it being \$2 are 30%. If the stock is bad then the probability of receiving the \$10 dividend is 30% and the probability of receiving the \$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being \$10 are 30%, and the odds of it being \$2 are 70%.

Specific details for the LOSS condition:

In the LOSS condition, on any trial, if you choose to invest in the bond, you get a payoff of -\$6 for sure at the end of the trial. If you choose to invest in the stock, you will receive a dividend which can be either -\$10 or -\$2.

The stock can either be good or bad, and this will determine the likelihood of its dividend being high or low. If the stock is good then the probability

of receiving the -\$10 dividend is 30% and the probability of receiving the -\$2 dividend is 70%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is good, then on each trial the odds of the dividend being -\$10 are 30%, and the odds of it being -\$2 are 70%. If the stock is bad then the probability of receiving the -\$10 dividend is 70% and the probability of receiving the -\$2 dividend is 30%. The dividends paid by this stock are independent from trial to trial, but come from this exact distribution. In other words, once it is determined by the computer that the stock is bad, then on each trial the odds of the dividend being -\$10 are 70%, and the odds of it being -\$2 are 30%.

In both the GAIN and LOSS conditions:

In each condition, at the beginning of each block of six trials, you do not know which type of stock the computer selected for that block. You may be facing the good stock, or the bad stock, with equal probability.

On each trial in the block you will decide whether you want to invest in the stock for that trial and accumulate the dividend paid by the stock, or invest in the riskless security and add the known payoff to your task earnings.

You will only see the dividend paid by the stock if you select it.

At one or several stages in each block of six trials, you will have to tell us what you think is the probability that the stock is the good one (the answer must be a number between 0 and 100 – do not add the % sign, just type in the value).

There is always an objective, correct, probability that the stock is good, which depends on the history of dividends paid by the stock already. For instance, at the beginning of each block of trials, the probability that the stock is good is exactly 50%, and there is no doubt about this value.

As you observe the dividends paid by the stock you will update your belief whether or not the stock is good. It may be that after a series of good dividends, you think the probability of the stock being good is 75%. However, how much you trust your ability to calculate this probability could vary. Sometimes you may not be too confident in the probability estimate you calculated and sometimes you may be highly confident in this estimate. For instance, at the very beginning of each block, the probability of the stock being good is 50% and you should be highly confident in this number because you are told that the computer just picked at random the type of stock you will see in the block, and nothing else has happened since then.

Throughout the task you will be told how much you have accumulated through dividends paid by the stock or bond you chose up to that point.

Your final pay for completing the investment tasks will be:

$\$5 + 1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.

Thank you!

B.3 Quiz for the Full feedback Condition

How many conditions are there in the task? What is their name?

- o 2 conditions called HIGH and LOW
- o 2 conditions called GOOD and BAD
- o 2 conditions called LOSS and GAIN

How many securities are there in the task? Which is the risky one and which the riskless?

- o There are 2 securities. The Stock which is the risky security and the Bond which is the riskless security
- o There are 2 securities. The Stock which is the riskless security and the Bond which is the risky security
- o There are 3 securities. The Stock which is the riskless security, the Bond which is the risky security and the Option which the very risky security

In the GAIN condition if I choose a Bond I will get... / In the GAIN condition if I choose the Stock I will get...

- o ... a payoff of \$6 for sure / ...a dividend which can be \$2 or \$10
- o ... a payoff of \$2 for sure / ...a dividend which can be \$6 or \$10
- o ... a payoff of \$6 for sure / ...a dividend which can be \$3 or \$10

In the LOSS condition if I choose a Bond I will get... / In the LOSS condition if I choose the Stock I will get...

- o ... a payoff of -\$6 for sure / ...a dividend which can be -\$2 or -\$10
- o ... a payoff of -\$2 for sure / ...a dividend which can be -\$6 or -\$10
- o ... a payoff of -\$6 for sure / ...a dividend which can be -\$3 or -\$10

How many types of Stocks are there in the GAIN condition?

- o 2 types, the high Stock and the low Stock
- o 2 types, the good Stock and the bad Stock
- o 3 types, the good Stock, the bad Stock and the neutral Stock

How many types of Stocks are there in the LOSS condition?

- o 2 types, the high Stock and the low Stock
- o 2 types, the good Stock and the bad Stock
- o 3 types, the good Stock, the bad Stock and the neutral Stock

In the GAIN condition, the good Stock pays \$10 with which probability?
And \$2?

- o Pays \$10 with 70% probability and \$2 with 30%
- o Pays \$10 with 30% probability and \$2 with 70%
- o Pays \$10 with 50% probability and \$2 with 50%

In the LOSS condition, the good Stock pays -\$10 with which probability?
And -\$2?

- o Pays -\$10 with 70% probability and -\$2 with 30%
- o Pays -\$10 with 30% probability and -\$2 with 70%
- o Pays -\$10 with 50% probability and -\$2 with 50%

In the GAIN condition, the bad Stock pays \$10 with which probability?
And \$2?

- o Pays \$10 with 70% probability and \$2 with 30%
- o Pays \$10 with 30% probability and \$2 with 70%
- o Pays \$10 with 50% probability and \$2 with 50%

In the LOSS condition, the bad Stock pays -\$10 with which probability?
And -\$2?

- o Pays -\$10 with 70% probability and -\$2 with 30%
- o Pays -\$10 with 30% probability and -\$2 with 70%
- o Pays -\$10 with 50% probability and -\$2 with 50%

At the beginning of each Block... (1)

- o I know if I am facing the good or the bad Stock
- o I do not know if I am facing the good or the bad Stock
- o I will be told which type of Stock do I face

At the beginning of each Block... (2)

- o There is a 50% chance that I face the good Stock and a 50% chance that I face the bad Stock
- o There is a 70% chance that I face the good Stock and a 30% chance that I face the bad Stock
- o There is a 30% chance that I face the good Stock and a 70% chance that I face the bad Stock

At the beginning of each Block... (3)

- o I know I will be facing the same type of Stock (good or bad) for the next 6 Trials
- o I know I will be facing a different type of Stock (good or bad) in each of the next 6 Trials
- o I know I will be facing a different type of Stock (good or bad) in each of the next 10 Trials

At the beginning of each Block... (4)

- o I know that the dividends paid by the Stock are independent from Trial to Trial
- o I know that the dividends paid by the Stock are dependent from Trial to Trial
- o I know that the dividends paid by the Stock are independent from Block to Block

On each Trial of the Block you will see the dividend paid by the Stock, no matter if you chose the Stock or the Bond.

- o True
- o False

Before the first choice in each block, and after each choice in any trial we ask you: what do you think is the probability that the Stock is the...

- o Good one
- o Bad one
- o High

Is there an objective, correct probability the Stock is the good one?

- o Yes, and it depends on the history of dividends paid by the Stock already
- o No
- o Yes, and it does not depend on the history of dividends paid by the Stock already

What is the objective probability of the Stock being the good one at the beginning of each Block, when you still have not possibly seen any dividend from the Stock?

- o 50%
- o 30%
- o 70%

What is the objective probability of the Stock being the good one at the beginning of each Block, before the first choice, when you still have not possibly seen any dividend from the Stock?

- o 50% with no doubt about its value
- o 50% but could be slightly different
- o 70% or 30%

Which is the high dividend in the GAIN condition? and the low one?

- o \$2 is the high and \$10 is the low
- o \$10 is the high and \$2 is the low
- o \$6 is the high and \$2 is the low

Which is the high dividend in the LOSS condition? and the low one?

- o -\$2 is the high and -\$10 is the low
- o -\$10 is the high and -\$2 is the low
- o -\$6 is the high and -\$2 is the low

If after some trials since you began a Block you see that the Stock has always given high dividends, your estimation of the probability that the Stock is the good one should be?

- o Higher than 50%
- o Lower than 50%
- o Still 50%

If after some trials since you began a Block you see that the Stock has always given low dividends, your estimation of the probability that the Stock is the good one should be?

- o Higher than 50%
- o Lower than 50%
- o Still 50%

Your final payoff will be calculated according to the following formula:

- o $\$5 + 1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.
- o $1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.
- o $\$5 + 1/100 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.

B.4 Quiz for the Selective feedback Condition

How many conditions are there in the task? What is their name?

- o 2 conditions called HIGH and LOW
- o 2 conditions called GOOD and BAD
- o 2 conditions called LOSS and GAIN

How many securities are in the task? Which is the risky one and which the riskless?

- o There are 2 securities. The Stock which is the risky security and the Bond which is the riskless security
- o There are 2 securities. The Stock which is the riskless security and the Bond which is the risky security
- o There are 3 securities. The Stock which is the riskless security, the Bond which is the risky security and the Option which the very risky security

In the GAIN condition if I choose a Bond I will get... / In the GAIN condition if I choose the Stock I will get...

- o ... a payoff of \$6 for sure / ...a dividend which can be \$2 or \$10
- o ... a payoff of \$2 for sure / ...a dividend which can be \$6 or \$10
- o ... a payoff of \$6 for sure / ...a dividend which can be \$3 or \$10

In the LOSS condition if I choose a Bond I will get... / In the LOSS condition if I choose the Stock I will get...

- o ... a payoff of -\$6 for sure / ...a dividend which can be -\$2 or -\$10
- o ... a payoff of -\$2 for sure / ...a dividend which can be -\$6 or -\$10
- o ... a payoff of -\$6 for sure / ...a dividend which can be -\$3 or -\$10

How many types of Stocks are there in the GAIN condition?

- o 2 types, the high Stock and the low Stock
- o 2 types, the good Stock and the bad Stock
- o 3 types, the good Stock, the bad Stock and the neutral Stock

How many types of Stocks are there in the LOSS condition?

- o 2 types, the high Stock and the low Stock
- o 2 types, the good Stock and the bad Stock
- o 3 types, the good Stock, the bad Stock and the neutral Stock

In the GAIN condition, the good Stock pays \$10 with which probability?
And \$2?

- o Pays \$10 with 70% probability and \$2 with 30%
- o Pays \$10 with 30% probability and \$2 with 70%
- o Pays \$10 with 50% probability and \$2 with 50%

In the LOSS condition, the good Stock pays -\$10 with which probability?
And -\$2?

- o Pays -\$10 with 70% probability and -\$2 with 30%
- o Pays -\$10 with 30% probability and -\$2 with 70%
- o Pays -\$10 with 50% probability and -\$2 with 50%

In the GAIN condition, the bad Stock pays \$10 with which probability?
And \$2?

- o Pays \$10 with 70% probability and \$2 with 30%
- o Pays \$10 with 30% probability and \$2 with 70%
- o Pays \$10 with 50% probability and \$2 with 50%

In the LOSS condition, the bad Stock pays -\$10 with which probability?
And -\$2?

- o Pays -\$10 with 70% probability and -\$2 with 30%
- o Pays -\$10 with 30% probability and -\$2 with 70%
- o Pays -\$10 with 50% probability and -\$2 with 50%

At the beginning of each Block... (1)

- o I know if I am facing the good or the bad Stock
- o I do not know if I am facing the good or the bad Stock
- o I will be told which type of Stock do I face

At the beginning of each Block... (2)

- o There is a 50% chance that I face the good Stock and a 50% chance that I face the bad Stock
- o There is a 70% chance that I face the good Stock and a 30% chance that I face the bad Stock
- o There is a 30% chance that I face the good Stock and a 70% chance that I face the bad Stock

At the beginning of each Block... (3)

- o I know I will be facing the same type of Stock (good or bad) for the next 6 Trials
- o I know I will be facing a different type of Stock (good or bad) in each of the next 6 Trials
- o I know I will be facing a different type of Stock (good or bad) in each of the next 10 Trials

At the beginning of each Block... (4)

- o I know that the dividends paid by the Stock are independent from Trial to Trial
- o I know that the dividends paid by the Stock are dependent from Trial to Trial
- o I know that the dividends paid by the Stock are independent from Block to Block

On each Trial of the Block you will see the dividend paid by the Stock, no matter if you chose the Stock or the Bond.

- o True, I will see the dividend paid by the stock regardless of my choice.
- o False, I will only see the dividend paid by the stock if I select it.

At one or several stages during the study we ask you: what do you think is the probability that the Stock is the...

- o Good one
- o Bad one
- o High

Is there an objective, correct probability the Stock is the good one?

- o Yes, and it depends on the history of dividends paid by the Stock already
- o No
- o Yes, and it does not depend on the history of dividends paid by the Stock already

What is the objective probability of the Stock being the good one at the beginning of each Block, when you still have not possibly seen any dividend from the Stock?

- o 50%
- o 30%
- o 70%

What is the objective probability of the Stock being the good one at the beginning of each Block, before the first choice, when you still have not possibly seen any dividend from the Stock?

- o 50% with no doubt about its value
- o 50% but could be slightly different
- o 70% or 30%

Which is the high dividend in the GAIN condition? and the low one?

- o \$2 is the high and \$10 is the low
- o \$10 is the high and \$2 is the low
- o \$6 is the high and \$2 is the low

Which is the high dividend in the LOSS condition? and the low one?

- o -\$2 is the high and -\$10 is the low
- o -\$10 is the high and -\$2 is the low
- o -\$6 is the high and -\$2 is the low

If after some trials since you began a Block you see that the Stock has always given high dividends, your estimation of the probability that the Stock is the good one should be?

- o Higher than 50%
- o Lower than 50%
- o Still 50%

If after some trials since you began a Block you see that the Stock has always given low dividends, your estimation of the probability that the Stock is the good one should be?

- o Higher than 50%
- o Lower than 50%
- o Still 50%

Your final payoff will be calculated according to the following formula:

- o $\$5 + 1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.
- o $1/10 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.
- o $\$5 + 1/100 \times \text{Investment Payoffs}$, where Investment Payoffs = Dividends of securities you chose in the experiment.

B.5 Objective Bayesian Posterior Beliefs

The table below provides all possible values for the objectively correct Bayesian posterior that the stock is paying from the good dividend distribution, starting with a 50% to 50% prior, and after observing each possible dividend history path in a learning block. Every trial a new dividend (high or low) is revealed. There are six trials in each learning block. The value of the objective Bayesian posterior that the stock is paying from the good distribution can be easily calculated. Specifically, after observing t high outcomes in n trials so far, the Bayesian posterior that the stock is the good one is given by: $\frac{1}{1 + \frac{1-p}{p} * (\frac{q}{1-q})^{n-2t}}$ where $p = 50\%$ is the prior that the stock is the good one (before any payoffs are observed in that learning block) and $q = 70\%$ is the probability that a good stock pays the high payoff (rather than low) in each trial.

n Trials So Far	t High Outcomes So Far	Probability {stock is good t high outcomes in n trials}
1	0	30.00%
1	1	70.00%
2	0	15.52%
2	1	50.00%
2	2	84.48%
3	0	7.30%
3	1	30.00%
3	2	70.00%
3	3	92.70%
4	0	3.26%
4	1	15.52%
4	2	50.00%
4	3	84.48%
4	4	96.74%
5	0	1.43%
5	1	7.30%
5	2	30.00%
5	3	70.00%
5	4	92.70%
5	5	98.57%
6	0	0.62%
6	1	3.26%
6	2	15.52%
6	3	50.00%
6	4	84.48%
6	5	96.74%
6	6	99.38%

B.6 Measures of Financial Literacy

To get measures of financial literacy and risk preferences, each participant was asked the following questions after the completion of the experimental tasks: "Imagine you have saved \$10,000. You can now invest this money over the next year using two investment options: a U.S. stock index mutual fund, which tracks the performance of the U.S. stock market, and a savings account. The annual return per dollar invested in the stock index fund will be either +40% or -20%, with equal probability. In other words, it is equally likely that for each dollar you invest in the stock market, at the end of the one year investment period, you will have either gained 40 cents, or lost 20 cents. For the savings account, the known and certain rate of return for a one year investment is 5%. In other words, for each dollar you put in the savings account today, for sure you will gain 5 cents at the end of the one year investment period. We assume that whatever amount you do not invest in stocks will be invested in the savings account and will earn the risk-free rate of return. Given this information, how much of the \$10,000 will you invest in the U.S. stock index fund? Choose an answer that you would be comfortable with if this was a real-life investment decision. The answer should be a number between \$0 and \$10,000." After each participant wrote their answer to this question, they were asked the following: "Let's say that when you answered the prior question you decided to invest x dollars out of the \$10,000 amount in the U.S. stock index fund, and therefore you put $(10,000 - x)$ dollars in the savings account. Recall that over the next year the rate of return on the stock index fund will be +40% or -20%, with equal probability. For the savings account, the rate of return is 5% for sure. What is the amount of money you expect to have at the end of this one year investment period? Please choose one of the answers below. If you choose the correct answer, you will get a \$1 bonus added to your pay for this experiment. [A] $0.5 (0.4 x - 0.2 x) + 0.05 (10,000 - x)$; [B] $1.4 x + 0.8 x + 1.05 (10,000 - x)$; [C] $0.4 (10,000 - x) - 0.2 (10,000 - x) + 0.05 x$; [D] $0.5 [0.4 (10,000 - x) - 0.2 (10,000 - x)] + 0.05 x$; [E] $0.4 x - 0.2 x + 0.05 (10,000 - x)$; [F] $0.5 (1.4 x + 0.8 x) + 1.05 (10,000 - x)$; [G] $1.4 (10,000 - x) + 0.8 (10,000 - x) + 1.05 x$; [H] $0.5 [1.4 (10,000 - x) + 0.8 (10,000 - x)] + 1.05 x$." The correct answer to this question is [F]. The actual choices (if other than [F]) made by participants indicate three different types of errors that can occur when calculating the expected value of their portfolio holdings: a lack of understanding of statements regarding probabilities (answers [B], [C], [E], [G]); a lack of understanding of the difference between net and gross returns (answers [A], [C], [D], and [E]);

and confusing the stock versus risk-free asset investments (answers [C], [D], [G], and [H]). Therefore, a financial knowledge score varying between zero and three can be constructed, based on the number of different types of errors contained in the answer provided by each participant (i.e., zero errors for answer [F], one error for answers [A], [B], and [H], two errors for answers [D], [E], and [G], and three for answer [C]). Hence a financial knowledge score of three indicates a perfect answer, while a score of zero indicates that the participant's answer included all three possible types of errors.

B.7 Payoff Sequences

Block Trial	1						2						3						4						5														
	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6									
Sequence 1	10	10	2	2	2	10	2	10	2	2	10	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 2	2	2	2	10	10	10	2	2	2	2	10	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 3	-2	-2	-2	-2	-2	-2	2	10	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 4	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2
Sequence 5	10	2	2	10	10	10	2	2	2	2	10	10	2	2	10	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 6	2	10	10	10	10	10	2	10	10	10	2	2	10	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 7	2	2	2	2	2	2	10	2	10	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 8	10	2	10	2	2	2	2	10	10	2	2	10	10	2	2	10	10	2	2	2	2	10	10	2	2	2	10	10	2	2	10	10	2	2	10	10	2	2	
Sequence 9	2	10	2	10	2	2	2	10	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 10	10	10	2	10	10	2	10	10	10	2	10	10	2	10	10	10	2	10	10	10	2	10	10	2	10	10	2	10	10	2	10	10	2	10	10	2	10	10	
Sequence 11	2	10	2	2	2	2	10	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 12	10	10	10	10	10	10	2	2	2	2	10	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 13	2	2	10	2	2	2	2	10	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 14	10	10	2	2	2	10	2	2	2	2	10	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 15	10	10	10	10	10	2	10	10	10	10	2	2	10	10	2	2	10	10	2	2	2	10	10	2	2	2	10	10	2	2	10	10	2	2	10	10	2	2	
Sequence 16	2	2	10	2	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 17	10	2	10	10	10	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 18	2	2	10	10	2	2	2	10	10	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 19	2	2	2	2	10	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 20	2	2	10	2	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 21	2	2	10	10	10	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 22	10	10	10	2	2	10	10	2	2	10	10	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 23	2	2	2	10	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 24	10	2	2	10	10	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 25	2	2	10	2	2	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 26	2	10	10	10	10	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 27	10	10	2	10	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 28	2	2	10	2	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 29	10	10	2	2	2	2	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	
Sequence 30	2	10	10	10	2	10	2	2	2	2	2	2	2	2	2	10	2	2	2	2	2	10	2	2	2	2	10	2	2	2	10	2	2	10	2	2	10	2	

Block	6						7						8						9						10					
	Trial	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5
Sequence 1	-2	-10	-2	-10	-10	-10	-10	-2	-10	-2	-10	-2	-10	-2	-10	-2	-10	-2	-10	-10	-2	-10	-2	-10	-2	-2	-2	-2	-2	-2
Sequence 2	-2	-2	-2	-10	-2	10	2	2	10	10	2	-2	-10	-2	-10	-2	-10	-2	-10	-10	-2	-10	-2	-10	-2	-2	-2	-2	-2	-2
Sequence 3	10	10	10	10	10	10	-2	-10	-2	-2	-2	10	2	10	2	2	2	2	10	10	10	10	10	10	2	2	2	2	2	10
Sequence 4	2	2	2	2	10	10	-10	-10	-10	-2	-2	-2	-2	-2	-2	-2	-2	-2	10	10	10	2	10	10	2	2	10	2	10	2
Sequence 5	2	10	10	10	10	10	-10	-2	-2	-10	-10	-10	-10	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2	-2	
Sequence 6	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	
Sequence 7	-2	-2	-2	-10	-2	-2	-10	-2	-2	-2	-10	-10	-10	-10	-10	-10	-10	-2	-2	-2	-2	-2	-2	-10	-10	-10	-10	-10	-10	
Sequence 8	-2	-2	-2	-2	-2	-2	-10	-2	-10	-10	-2	-10	-2	-2	-10	-10	-10	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	-2	
Sequence 9	-10	-2	-10	-10	-2	-2	-10	-10	-10	-10	-2	-2	-2	-10	-10	-2	-10	-2	-2	-10	-10	-10	-2	-2	-2	-10	-10	-10	-10	-10
Sequence 10	-2	-2	-2	-2	-2	-2	-2	-10	-2	-2	-2	-2	-10	-10	-2	-2	-10	-2	-10	-10	-2	-2	-10	-2	-2	-2	-2	-2	-2	-2
Sequence 11	-2	-10	-2	-2	-2	-2	-2	-10	-2	-2	-2	-2	-2	-10	-2	-2	-10	-2	-10	-10	-2	-2	-10	-2	-2	-2	-2	-2	-2	-10
Sequence 12	-10	-2	-10	-10	-10	-2	-10	-10	-10	-2	-10	-2	-10	-2	-2	-10	-2	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-2	
Sequence 13	-10	-2	-2	-2	-10	-10	-10	-2	-2	-2	-2	-2	-10	-2	-2	-2	-10	-2	-2	-10	-10	-10	-2	-10	-2	-2	-2	-2	-2	-10
Sequence 14	-10	-10	-2	-10	-2	-10	-2	-10	-10	-10	-10	-2	-2	-2	-10	-10	-10	-2	-2	-2	-2	-2	-10	-2	-10	-10	-10	-10	-2	
Sequence 15	-2	-2	-2	-10	-10	-2	-2	-10	-10	-2	-10	-10	-2	-10	-2	-2	-10	-2	-10	-10	-2	-2	-10	-2	-2	-10	-10	-10	-10	-2
Sequence 16	-2	-2	-10	-2	-10	-10	-2	-2	-10	-10	-10	-2	-2	-2	-10	-10	-2	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2	
Sequence 17	-2	-10	-10	-10	-10	-10	-2	-10	-2	-10	-10	-10	-10	-10	-2	-2	-2	-2	-10	-10	-10	-10	-10	-10	-2	-2	-2	-10	-10	
Sequence 18	-2	-10	-10	-10	-2	-10	-2	-10	-2	-2	-10	-10	-10	-10	-2	-2	-10	-10	-2	-10	-10	-10	-10	-10	-2	-10	-10	-10	-10	-10
Sequence 19	-2	-10	-2	-2	-2	-10	-10	-10	-10	-10	-10	-2	-2	-10	-10	-2	-10	-2	-10	-10	-2	-2	-10	-2	-2	-2	-2	-10	-10	-10
Sequence 20	-2	-2	-10	-10	-2	-10	-10	-2	-2	-10	-10	-10	-10	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-2	-2	-10	-10	-2	-2	
Sequence 21	-2	-2	-10	-10	-2	-2	-10	-2	-10	-2	-10	-10	-2	-2	-10	-2	-2	-2	-10	-10	-10	-2	-10	-2	-10	-2	-2	-2	-10	
Sequence 22	-2	-2	-2	-10	-10	-2	-2	-10	-2	-2	-10	-2	-2	-2	-2	-2	-2	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	
Sequence 23	-10	-10	-2	-10	-2	-10	-2	-2	-2	-2	-10	-2	-10	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	
Sequence 24	-10	-2	-10	-2	-2	-10	-10	-10	-2	-2	-10	-2	-10	-10	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-2	-2	-2	-2	-10	
Sequence 25	-10	-10	-2	-10	-10	-2	-10	-2	-2	-10	-10	-2	-2	-10	-10	-10	-10	-2	-10	-10	-10	-10	-10	-2	-10	-10	-10	-10	-2	
Sequence 26	-10	-10	-2	-2	-2	-2	-10	-2	-10	-2	-10	-2	-2	-2	-10	-2	-2	-2	-10	-10	-10	-2	-10	-2	-2	-2	-10	-10	-2	
Sequence 27	-2	-10	-2	-2	-10	-2	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-2	-2	-2	-2	-2	-2	-10	-2	-2	-2	-10	-2	-2	
Sequence 28	-10	-2	-2	-2	-2	-10	-10	-2	-10	-10	-2	-10	-10	-2	-10	-2	-2	-2	-2	-2	-2	-2	-10	-2	-10	-10	-2	-10	-10	
Sequence 29	-2	-2	-2	-2	-10	-10	-2	-2	-2	-2	-2	-10	-10	-2	-10	-2	-2	-10	-10	-10	-10	-10	-10	-2	-2	-2	-2	-2	-2	
Sequence 30	-10	-10	-2	-2	-2	-10	-2	-10	-10	-10	-2	-2	-10	-10	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2	

Block Trial	1						2						3						4						5					
	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Sequence 31	2	10	2	2	2	2	2	2	2	2	2	10	2	10	10	10	10	2	2	10	10	10	2	2	2	10	10	10		
Sequence 32	10	10	10	2	2	10	2	10	2	2	2	10	10	10	10	2	2	10	10	10	10	2	2	10	10	10	10	10		
Sequence 33	10	10	10	10	10	2	10	2	2	2	10	10	2	2	2	2	2	10	2	2	2	2	10	10	10	10	10	10		
Sequence 34	2	10	10	2	10	10	10	10	10	2	2	10	10	10	10	10	10	10	10	10	10	2	2	2	2	2	2	2		
Sequence 35	2	10	2	10	2	2	10	10	10	10	2	10	2	2	10	2	2	10	10	10	10	2	2	2	2	2	2	2		
Sequence 36	10	10	2	10	10	2	2	2	2	2	2	10	2	2	10	10	2	10	2	10	10	10	2	2	2	2	2	2		
Sequence 37	10	10	2	2	10	2	2	10	10	2	2	2	2	2	10	10	10	10	2	2	10	2	2	2	2	2	2	2		
Sequence 38	2	2	2	2	2	2	2	2	2	10	10	2	2	2	10	10	10	2	2	10	2	2	2	2	2	2	2	2		
Sequence 39	2	2	10	2	10	2	2	2	2	2	2	10	10	2	2	10	10	2	2	2	10	10	2	2	10	10	2	2		
Sequence 40	2	2	2	2	2	2	2	2	2	10	10	2	10	10	10	10	10	10	10	10	10	2	2	10	10	10	10	10		
Sequence 41	2	10	2	10	10	10	2	2	2	10	10	2	2	2	10	10	2	2	2	10	10	2	2	2	2	2	2	2		
Sequence 42	2	10	10	2	10	2	2	10	10	2	2	10	10	10	10	10	10	10	2	2	10	10	2	2	2	2	2	10		
Sequence 43	10	10	2	10	10	2	10	10	10	2	2	10	10	10	10	10	10	10	2	2	10	10	2	2	2	2	2	10		
Sequence 44	2	2	10	10	2	10	10	2	2	2	2	10	2	2	10	10	2	2	2	10	10	2	2	2	2	2	2	2		
Sequence 45	2	2	2	2	10	10	2	2	2	2	2	10	10	10	10	10	10	2	2	10	10	2	2	2	2	2	2	2		
Sequence 46	10	10	10	2	10	10	10	10	10	10	10	10	10	10	10	10	2	2	10	10	10	2	2	2	2	2	2	2		
Sequence 47	2	2	2	10	10	2	2	2	2	10	10	2	2	2	10	10	10	2	2	2	10	10	2	2	2	2	2	2		
Sequence 48	2	10	2	2	2	2	2	2	2	2	2	10	10	2	2	10	2	2	2	10	2	2	2	2	2	2	2	2		
Sequence 49	2	10	10	10	2	10	2	10	2	2	10	2	10	2	10	2	10	2	2	10	2	2	10	10	2	10	10	10		
Sequence 50	2	10	10	10	10	10	2	10	2	2	10	2	10	2	10	10	2	2	10	10	10	2	2	10	2	10	10	2		
Sequence 51	10	2	10	2	2	10	10	2	10	10	2	10	2	10	10	10	2	2	10	10	10	2	2	2	2	2	2	10		
Sequence 52	10	10	10	10	10	10	10	2	2	2	2	10	10	10	10	10	2	2	10	10	10	2	2	10	10	10	10	2		
Sequence 53	10	10	2	10	2	10	2	10	10	2	2	2	10	2	2	10	10	2	2	10	10	10	2	2	10	10	10	2		
Sequence 54	2	2	2	2	10	2	10	2	10	2	2	2	2	2	10	2	2	2	2	10	10	10	2	2	10	2	2	10		
Sequence 55	2	10	2	2	2	2	10	10	10	2	10	2	2	2	10	2	2	2	2	10	10	10	2	2	10	10	10	10		
Sequence 56	2	2	10	2	10	2	2	2	2	10	10	10	10	10	10	10	2	2	2	10	10	10	2	2	10	10	10	10		
Sequence 57	2	2	2	2	2	2	2	2	2	10	10	10	10	10	2	2	10	2	2	10	10	10	2	2	10	10	10	10		
Sequence 58	-10	-10	-2	-2	-2	-2	-2	-10	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2	-10	-2	-10	-10	-2		
Sequence 59	-10	-10	-10	-10	-10	-2	-10	-10	-10	-2	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-10	-2		
Sequence 60	-2	-2	-2	-2	-2	-2	2	10	2	2	10	10	2	2	10	10	10	10	10	10	10	10	10	10	10	10	10	-2		

B.8 Type of Stock Sequences

For any trial in a block a 1 in the table represents that a participant was confronting the good stock, a 0 that the participant was confronting the bad stock. Whether the stock was good or bad was decided randomly at the beginning of each block, with each state having 50% probability.

Block	1						2						3						4						5					
Trial	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Sequence 1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 2	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0		
Sequence 3	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 4	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
Sequence 5	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0		
Sequence 6	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 7	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0		
Sequence 8	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0		
Sequence 9	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 10	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	1	1	1		
Sequence 11	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0		
Sequence 12	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 13	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0		
Sequence 14	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0		
Sequence 15	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
Sequence 16	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0		
Sequence 17	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1		
Sequence 18	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1		
Sequence 19	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 20	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 21	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1		
Sequence 22	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 23	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 24	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 25	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 26	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0		
Sequence 27	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1		
Sequence 28	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 29	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0		
Sequence 30	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		

Block	6						7						8						9						10					
Trial	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Sequence 1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 3	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	
Sequence 4	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 5	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 7	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 8	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 9	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	
Sequence 10	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 11	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 12	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	
Sequence 13	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 14	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	
Sequence 15	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 16	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 17	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 18	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 19	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 20	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 21	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 22	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 23	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 24	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 25	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 26	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 27	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 28	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	
Sequence 29	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 30	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	

Block	1						2						3						4						5					
Trial	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Sequence 31	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 32	1	1	1	1	1	1	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 33	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 34	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 35	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 36	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 37	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 38	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 39	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 40	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 41	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 42	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 43	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 44	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 45	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 46	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 47	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	
Sequence 48	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 49	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 50	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 51	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 52	1	1	1	1	1	1	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	
Sequence 53	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 54	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 55	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	
Sequence 56	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 57	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 58	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	
Sequence 59	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	
Sequence 60	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	

Block	6						7						8						9						10					
Trial	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6	1	2	3	4	5	6
Sequence 31	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 32	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 33	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 34	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 35	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0		
Sequence 36	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0		
Sequence 37	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0		
Sequence 38	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
Sequence 39	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
Sequence 40	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1		
Sequence 41	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0		
Sequence 42	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	0	1	1	1	1	1		
Sequence 43	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0		
Sequence 44	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0		
Sequence 45	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0		
Sequence 46	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 47	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0		
Sequence 48	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1		
Sequence 49	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0		
Sequence 50	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0		
Sequence 51	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 52	1	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0	1	1	1	1	1		
Sequence 53	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 54	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 55	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	0	0	0	0	0	0		
Sequence 56	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 57	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	0	0	0	0	0	0		
Sequence 58	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		
Sequence 59	0	0	0	0	0	0	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
Sequence 60	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1		

Appendix C

INADVERTENTLY SUPER-FARSIGHTED MEDITATORS: THE EFFECTS OF MINDFULNESS ON INTER-TEMPORAL CHOICE

C.1 Experiment 1

C.1.1 Tasks

Specific Instructions mindfulness condition

Ahora vas a pasar unos 15 minutos haciendo un ejercicio relacionado con tu respiración. Cuando acabes de leer estas instrucciones, inicia el archivo de audio que hay debajo (dándole al botón de "play") y utiliza tus auriculares para escucharlo. El audio te guiará en el ejercicio. Un vez inicies el audio, por favor sigue las instrucciones que éste te vaya dando lo mejor posible y no lo interrumpas hasta el final. No saltes tampoco adelante o atrás en el audio y deja que siga su curso. Cuando acabes, espera a que te demos el código que necesitas para continuar.

Por favor haz el ejercicio y espera a que te demos el código que necesitas para continuar.

Specific Instructions mind-wandering condition

Ahora vas a pasar unos 15 minutos haciendo un ejercicio relacionado con tu pensamiento. Cuando acabes de leer estas instrucciones, inicia el archivo de audio que hay debajo (dándole al botón de "play") y utiliza tus auriculares para escucharlo. El audio te guiará en el ejercicio. Un vez inicies el audio, por favor sigue las instrucciones que éste te vaya dando lo mejor posible y no lo interrumpas hasta el final. No saltes tampoco adelante o atrás en el audio y deja que siga su curso. Cuando acabes, espera a que te demos el código que necesitas para continuar.

Por favor haz el ejercicio y espera a que te demos el código que necesitas para continuar.

Decision Instructions for all conditions

Ahora tendrás que responder a una serie de decisiones que se te van a presentar. Por favor lee cuidadosamente cada una de ellas y responde lo que crees que harías en este momento si la decisión fuese de verdad. Dale al botón que hay debajo para empezar.

Manipulation check items

Cómo de acuerdo estás con las afirmaciones que aparecen debajo. Por favor responde utilizando las escalas proporcionadas:

- | | | | | | | | |
|----|---|--|---|---|---|---|---|
| | Muy poco
o nada | Mucho | | | | | |
| 1. | Durante los últimos 15 minutos, he estado mayormente absorto en el momento presente. | <table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td style="width: 20px; text-align: center;">1</td><td style="width: 20px; text-align: center;">2</td><td style="width: 20px; text-align: center;">3</td><td style="width: 20px; text-align: center;">4</td><td style="width: 20px; text-align: center;">5</td></tr></table> | 1 | 2 | 3 | 4 | 5 |
| 1 | 2 | 3 | 4 | 5 | | | |
| 2. | Durante los últimos 15 minutos, he estado centrado en las sensaciones físicas de mi cuerpo. | <table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td style="width: 20px; text-align: center;">1</td><td style="width: 20px; text-align: center;">2</td><td style="width: 20px; text-align: center;">3</td><td style="width: 20px; text-align: center;">4</td><td style="width: 20px; text-align: center;">5</td></tr></table> | 1 | 2 | 3 | 4 | 5 |
| 1 | 2 | 3 | 4 | 5 | | | |

Ahora situate en la escala que hay debajo:

- | | | | | | | | |
|----|--|--|---|---|---|---|---|
| | Absorto en
el presente | Divagando
libremente | | | | | |
| 3. | Durante los últimos 15 minutos, he estado. | <table border="1" style="display: inline-table; border-collapse: collapse;"><tr><td style="width: 20px; text-align: center;">1</td><td style="width: 20px; text-align: center;">2</td><td style="width: 20px; text-align: center;">3</td><td style="width: 20px; text-align: center;">4</td><td style="width: 20px; text-align: center;">5</td></tr></table> | 1 | 2 | 3 | 4 | 5 |
| 1 | 2 | 3 | 4 | 5 | | | |

Cuando hayas respondido, haz click en el botón que hay debajo para continuar.

Task 1

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo. Cuál de los dos escogerías?

o Activo A: Da **€200 hoy**

o Activo B: Da **€220 dentro de 4 semanas**

Cuando hayas tomado tu decisión, dale al botón que hay debajo para continuar.

Task 2

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo. Cuál de los dos escogerías?

o Activo A: Da **€200 dentro de 12 semanas**

o Activo B: Da **€220 dentro de 16 semanas**

Cuando hayas tomado tu decisión, dale al botón que hay debajo para continuar.

Task 3

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo. Cuál de los dos escogerías?

o Activo A: Da **€200 hoy**

o Activo B: Da **€250 dentro de 4 semanas**

Cuando hayas tomado tu decisión, dale al botón que hay debajo para continuar.

Task 4

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo. Cuál de los dos escogerías?

- o Activo A: Da **€200** dentro de **12 semanas**
- o Activo B: Da **€250** dentro de **16 semanas**

Cuando hayas tomado tu decisión, dale al botón que hay debajo para continuar.

C.2 Experiment 2

C.2.1 Tasks

Task 1

Imagina que inesperadamente recibes €10.000 en una lotería. Imagina que tienes las siguientes dos opciones para asignar todo el dinero.

¿Cómo querrías asignarlo?

- o Gastos en actividades de ocio
- o Ahorro en una cuenta de ahorro a plazo fijo

Task 2

Imagina que al graduarte consigues tu primer trabajo y ganas €2.000 netos al mes. Por favor, asigna el dinero que dedicarías al mes a estas categorías:

- o Cuenta de ahorro a plazo fijo
- o Plan de pensiones
- o Alquiler y gastos de vivienda
- o Ocio
- o Comida
- o Otros

Task 3

Estas considerando comprar un coche pero únicamente tienes ahorrado la mitad del dinero necesario. Tienes dos opciones: 1) pedir prestado el dinero restante al banco y poder disfrutar ya del coche, por lo que tendrías pagar un interés de €1.000 en los próximos 2 años; 2) esperar 2 años hasta que hayas ahorrado lo suficiente para poder tener el coche y no pagar ningún interés. ¿Qué opción eliges?

- o Pedir prestado el dinero
- o Ahorrar el dinero

C.3 Experiment 3

C.3.1 Common Instructions

First Part

Gracias por participar en este experimento del BES Lab (Behavioral and Experimental Sciences Laboratory, Universitat Pompeu Fabra).

En este experimento tendrás que realizar varias tareas distintas que se te irán explicando conforme vayas avanzando. La duración total debería ser de menos de 40 minutos. Al final recibirás un pago de 8 euros por participar en el experimento y seguir las instrucciones correctamente.

Si tienes alguna duda durante el experimento, pregunta por favor a la persona a cargo de la sesión experimental. Haz click en el botón que hay debajo para empezar.

Second Part

Ahora vas a tomar una serie de decisiones en las que siempre tendrás 2 opciones disponibles. Las opciones proporcionarían distintas cantidades de dinero en distintos momentos del tiempo. Escoge en cada decisión la opción que preferirías escoger en este momento si las decisiones fuesen reales.

En esta tarea puedes obtener un pago adicional al de participación. El pago se determinará de la siguiente manera. Hoy, después de la sesión, seleccionaremos aleatoriamente a una de las personas participantes y también una de las decisiones que ha tomado esa persona. La persona seleccionada recibirá como pago adicional lo que haya escogido en la decisión seleccionada en el momento del tiempo indicado. El pago se efectuará mediante un cheque regalo de Amazon.es por el importe y en el plazo elegidos, que será enviado por email al participante. Por tanto, cualquiera de las decisiones que tomes puede ser la que determine tu pago adicional. Haz click en el botón "Avanzar" para empezar con las decisiones de esta parte.

¿Qué opción preferirías?

o 27,10 euros dentro de 2 semanas

o 27,10 euros dentro de 6 semanas

¿Qué opción preferirías?

- o 0,16 euros hoy
- o 34,04 euros dentro de 6 semanas

¿Qué opción preferirías?

- o \$lm://Field/1 euros \$lm://Field/3
- o \$lm://Field/2 euros \$lm://Field/4¹

C.3.2 Specific Mindfulness Condition Instructions

Este es un test para verificar que el audio de tu ordenador y tus auriculares funcionan correctamente. Por favor, escribe abajo lo que está haciendo la persona en la pista de audio. Si no escuchas el audio correctamente levanta la mano y acudiremos a ayudarte.

Ahora vas a pasar unos 15 minutos haciendo un ejercicio relacionado con tu respiración. Cuando acabes de leer estas instrucciones, inicia el archivo de audio que hay debajo (dándole al botón de "play") y utiliza tus auriculares para escucharlo. El audio te guiará en el ejercicio. Una vez inicies el audio, por favor sigue las instrucciones que éste te vaya dando lo mejor posible y no lo interrumpas hasta el final. No saltes tampoco adelante o atrás en el audio y deja que siga su curso. Cuando acabes, pulsa el botón de continuar.

C.3.3 Specific Mind-wandering Condition Instructions

Este es un test para verificar que el audio de tu ordenador y tus auriculares funcionan correctamente. Por favor, escribe abajo lo que está haciendo la persona en la pista de audio. Si no escuchas el audio correctamente levanta la mano y acudiremos a ayudarte.

Ahora vas a pasar unos 15 minutos haciendo un ejercicio relacionado con tu respiración. Cuando acabes de leer estas instrucciones, inicia el archivo de audio que hay debajo (dándole al botón de "play") y utiliza tus auriculares para escucharlo. El audio te guiará en el ejercicio. Una vez inicies el audio, por favor sigue las instrucciones que éste te vaya dando lo mejor posible y no lo interrumpas hasta el final. No saltes tampoco adelante o

¹This is an automatic function that displays in a random order the next 40 choices. The complete list of choices will be provided later on in this annex.

atrás en el audio y deja que siga su curso. Cuando acabes, pulsa el botón de continuar.

C.3.4 Intertemporal choice task items

List of Choices

Trial	Early delay	Late delay	% of Early over Late	Early reward	Late reward
1	2 weeks	6 weeks	0%	27,10	27,10
2	today	6 weeks	n.a.	0,16	34,04
3	today	2 weeks	1%	21,53	21,75
4	today	2 weeks	3%	24,45	25,19
5	today	2 weeks	5%	21,49	22,57
6	today	2 weeks	10%	31,70	34,87
7	today	2 weeks	15%	30,45	35,02
8	today	2 weeks	25%	27,12	33,91
9	today	2 weeks	35%	18,12	24,46
10	today	2 weeks	50%	27,61	41,41
11	today	4 weeks	1%	16,75	16,92
12	today	4 weeks	3%	6,34	6,53
13	today	4 weeks	5%	19,77	20,76
14	today	4 weeks	10%	20,23	22,25
15	today	4 weeks	15%	23,11	26,58
16	today	4 weeks	25%	18,83	23,54
17	today	4 weeks	35%	17,86	24,10
18	today	4 weeks	50%	23,43	35,14
19	2 weeks	4 weeks	1%	33,20	33,53
20	2 weeks	4 weeks	3%	16,85	17,35
21	2 weeks	4 weeks	5%	10,43	10,95
22	2 weeks	4 weeks	10%	28,93	31,83
23	2 weeks	4 weeks	15%	17,48	20,10
24	2 weeks	4 weeks	25%	14,59	18,24
25	2 weeks	4 weeks	35%	27,56	37,21
26	2 weeks	4 weeks	50%	14,50	21,75
27	2 weeks	6 weeks	1%	14,94	15,09
28	2 weeks	6 weeks	3%	16,44	16,93
29	2 weeks	6 weeks	5%	23,05	24,20
30	2 weeks	6 weeks	10%	11,94	13,13
31	2 weeks	6 weeks	15%	34,29	39,44
32	2 weeks	6 weeks	25%	20,70	25,88
33	2 weeks	6 weeks	35%	11,79	15,92
34	2 weeks	6 weeks	50%	28,79	43,19
35	4 weeks	6 weeks	1%	30,49	30,80
36	4 weeks	6 weeks	3%	5,13	5,28
37	4 weeks	6 weeks	5%	25,12	26,37
38	4 weeks	6 weeks	10%	24,05	26,45
39	4 weeks	6 weeks	15%	5,60	6,44
40	4 weeks	6 weeks	25%	16,59	20,73
41	4 weeks	6 weeks	35%	20,84	28,13
42	4 weeks	6 weeks	50%	6,70	10,05

C.4 Experiment 3

C.4.1 Common information sheet & Task booklet

Universitat Pompeu Fabra



Departamento de Economía y Empresa
Universitat Pompeu Fabra
Barcelona, Ramon Trias Fargas 25-27, 08005
Teléfono: 93.542.2000
Fax: 93.542.2002

Título del Proyecto de Investigación: Estudio en el Instituto EsMindfulness

Investigadores: Josep Gisbert Rodríguez, UPF josep.gisbert@upf.edu; Dr. Natalia Karelaia, INSEAD natalia.karelaia@insead.edu; Dr. Andrés Martín Asuero, Instituto EsMindfulness andres.martin@esmindfulness.com; Dr. Daniel Navarro Martínez, UPF daniel.navarro@upf.edu; Dr. Jordi Quoidbach, UPF jordi.quoidbach@upf.edu.

Introducción y Propósito: Agradeceríamos tu participación en este proyecto de investigación, dedicado a entender los efectos del Curso MBSR (Reducción del Estrés Basada en Mindfulness) en sus participantes.

Desarrollo del Estudio: El estudio constará de dos partes. Si accedes a participar, la primera parte la contestarás hoy; la segunda al finalizar el curso. En ambas partes deberás simplemente contestar lo más sinceramente posible una serie de preguntas. No prevemos ningún tipo de riesgos ni problemas derivados de la participación en el estudio. Cada parte durará aproximadamente 10 minutos.

Confidencialidad: Toda la información que proporciones será estrictamente confidencial. Esta hoja de consentimiento es el único documento que contendrá la información con tus datos identificativos; la hoja será almacenada en un armario con cerradura en una oficina de la Universitat Pompeu Fabra. Todas las otras respuestas serán codificadas numéricamente y el acceso a ellas estará limitado a los investigadores mencionados en este documento. Los resultados de este estudio pueden ser usados en la tesis doctoral de Josep Gisbert Rodríguez.

Remuneración: Este estudio no será remunerado económicamente. Sin embargo, una vez completada la segunda parte se entregará a cada participante que la haya completado una copia del libro de Andrés Martín Asuero *Plena Mente: Mindfulness o el Arte de Estar Presente* como agradecimiento por su colaboración.

Información de Contacto: Este estudio será llevado a cabo por los investigadores presentados en el primer párrafo de este documento. Por favor contacta con Josep Gisbert Rodríguez a través de josep.gisbert@upf.edu si tienes alguna pregunta sobre el estudio.

Consentimiento: La participación en el estudio es completamente voluntaria por lo que puedes retirar tu consentimiento en cualquier momento en que lo desees y por cualquier razón. Si participas o no, o completas o no el estudio, no tendrá efecto alguno en tu trato en el resto del Curso MBSR (Reducción del Estrés Basada en Mindfulness).

Puedes preguntar cualquier duda que tengas durante el transcurso del estudio. Intentaremos responder a tus preguntas lo mejor posible, de manera que entiendas el desarrollo del estudio en todo momento. Al firmar este documento estas indicando que aceptas participar en el estudio y reconoces haber recibido una copia de este formulario de consentimiento.

Nombre del Participante: _____ Fecha de hoy: _____

Firma del Participante: _____

Email del Participante: _____



**Universitat
Pompeu Fabra**
Barcelona

Estudio en el
Instituto
EsMindfulness

Universitat Pompeu Fabra

Instituto EsMindfulness

Barcelona

Estudio en el
Instituto EsMindfulness

INSTRUCCIONES:

En las siguientes páginas tendrás que responder a una serie de decisiones que se te van a presentar (una por página). Por favor lee cuidadosamente cada una de ellas y responde lo que crees que harías en este momento si la decisión fuese de verdad. En la parte final tendrás que responder también a un cuestionario. El tiempo total debería ser inferior a 15 minutos.

Pasa a la siguiente página del folleto para empezar.

DECISIÓN 13:

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo.Cuál de los dos escogerías? (Marca con una X la casilla correspondiente)

- Activo A: Da **200€ hoy**
- Activo B: Da **220€** dentro de **4 semanas**

DECISIÓN 14:

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo.Cuál de los dos escogerías? (Marca con una X la casilla correspondiente)

- Activo A: Da **200€** dentro de **12 semanas**
- Activo B: Da **220€** dentro de **16 semanas**

DECISIÓN 15:

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo.Cuál de los dos escogerías? (Marca con una X la casilla correspondiente)

- Activo A: Da **200€ hoy**
- Activo B: Da **250€** dentro de **4 semanas**

DECISIÓN 16:

Imagina que tuvieses que elegir entre los dos activos (A y B) que se muestran a continuación, los cuales te darían distintas cantidades de dinero en distintos momentos del tiempo.Cuál de los dos escogerías? (Marca con una X la casilla correspondiente)

- Activo A: Da **200€** dentro de **12 semanas**
- Activo B: Da **250€** dentro de **16 semanas**

Estudio en el
Instituto EsMindfulness

- Para acabar, por favor proporci6nanos la siguiente informaci6n.

1) Cu6al es tu edad? (por favor esc6ribela en la casilla que hay debajo)

Edad:

2) Cu6al es tu g6nero? (Marca con una X la casilla correspondiente)

Mujer Hombre

Bibliography

- Arch, J. J. and Craske, M. G. (2006). Mechanisms of mindfulness: emotion regulation following a focused breathing induction. *Behaviour research and therapy*, 44(12):1849–1858.
- Ashby, N. J. S. and Rakow, T. (2016). Eyes on the prize? evidence of diminishing attention to experienced and foregone outcomes in repeated experiential choice. *Journal of Behavioral Decision Making*, 29(2-3):183–193.
- Asparouhova, E., Bossaerts, P., Eguia, J., and Zame, W. (2015). Asset pricing and asymmetric reasoning. *Journal of Political Economy*, 123(1):66–122.
- Banerjee, P., Das, T., and Gulati, P. (2017). The impact of past outcomes on choice in a cognitively demanding financial environment. SSRN Scholarly Paper 2849348, Social Science Research Network.
- Barber, B., Odean, T., and Zheng, L. (2005). Out of sight, out of mind: The effects of expenses on mutual fund flows. *The Journal of Business*, 78(6):2095–2120.
- Barber, B. M. and Odean, T. (2007). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The Review of Financial Studies*, 21(2):785–818.
- Barberis, N., Shleifer, A., and Vishny, R. (1998). A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343.
- Barberis, N. and Xiong, W. (2012). Realization utility. *Journal of Financial Economics*, 104(2):251–271.
- Barocas, S. and Selbst, A. D. (2016). Big data’s disparate impact. *Calif. L. Rev.*, 104:671.

- Bartlett, R., Morse, A., Stanton, R., and Wallace, N. (2021). Consumer-lending discrimination in the fintech era. *Journal of Financial Economics*.
- Bazzano, M. (2014). *After Mindfulness: New Perspectives on Psychology and Meditation*. Palgrave Macmillan UK.
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., and Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9):1214–1221.
- Bereby-Meyer, Y. and Erev, I. (1998). On learning to become a successful loser: A comparison of alternative abstractions of learning processes in the loss domain. *Journal of Mathematical Psychology*, 42(2):266–286.
- Berns, G. S., Laibson, D., and Loewenstein, G. (2007). Intertemporal choice—toward an integrative framework. *Trends Cogn Sci*, 11(11):482–488.
- Bond, P. and Townsend, R. M. (1996). Formal and informal financing in a Chicago neighborhood. *Economic Perspectives*, 20(Jul):3–27.
- Bonfim, D., Nogueira, G., and Ongena, S. (2020). Sorry, we’re closed” bank branch closures, loan pricing, and information asymmetries. *Review of Finance*, 25:1211–1259.
- Bossaerts, P. (2004). Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory. *The Review of Economic Studies*, 71(1):63–86.
- Bowen, S., Witkiewitz, K., Dillworth, T. M., Chawla, N., Simpson, T. L., Ostafin, B. D., Larimer, M. E., Blume, A. W., Parks, G. A., and Marlatt, G. A. (2006). Mindfulness meditation and substance use in an incarcerated population. *Psychology of addictive behaviors : journal of the Society of Psychologists in Addictive Behaviors*, 20(3):343–347.
- Britton, W. B. (2019). Can mindfulness be too much of a good thing? the value of a middle way. *Current Opinion in Psychology*, 28:159–165.
- Brown, G. and Koopmans, T. (1951). Activity analysis of production and allocation. In *Activity Analysis of Production and Allocation: Proceedings of a Conference*. New York, USA: Wiley.
- Bruguier, A. J., Quartz, S. R., and Bossaerts, P. (2010). Exploring the nature of “trader intuition”. *The Journal of Finance*, 65(5):1703–1723.

- Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118.
- Buchak, G., Matvos, G., Piskorski, T., and Seru, A. (2018). Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics*, 130:453–483.
- Burke, C. A. (2010). Mindfulness-based approaches with children and adolescents: A preliminary review of current research in an emergent field. *Journal of Child and Family Studies*, 19(2):133–144.
- Butcher, K. F. and Muñoz, A. P. (2017). Using credit reporting agency data to assess the link between the community reinvestment act and consumer credit outcomes. *Cityscape*, 19(2):85–108.
- Camerer, C. and Hua Ho, T. (1999). Experience-weighted attraction learning in normal form games. *Econometrica*, 67(4):827–874.
- Caplin, A. and Eliaz, K. (2003). Aids policy and psychology: A mechanism-design approach. *RAND Journal of Economics*, 34(4):631–646.
- Carlin, B. I., Kogan, S., and Lowery, R. (2013). Trading complex assets. *The Journal of Finance*, 68(5):1937–1960.
- Carlson, L. E., Specia, M., Patel, K. D., and Goodey, E. (2004). Mindfulness-based stress reduction in relation to quality of life, mood, symptoms of stress and levels of cortisol, dehydroepiandrosterone sulfate (dheas) and melatonin in breast and prostate cancer outpatients. *Psychoneuroendocrinology*, 29(4):448–474.
- Chen, S., Doerr, S., Frost, J., Gambacorta, L., and Shin, H. S. (2021). The fintech gender gap. BIS Working Papers 931, Bank for International Settlements.
- Chiesa, A., Calati, R., and Serretti, A. (2011). Does mindfulness training improve cognitive abilities? a systematic review of neuropsychological findings. *Clinical psychology review*, 31(3):449–464.
- Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Reinforcement learning and savings behavior. *The Journal of Finance*, 64(6):2515–2534.
- Cohn, A., Engelmann, J., Fehr, E., and Maréchal, M. A. (2015). Evidence for countercyclical risk aversion: An experiment with financial professionals. *American Economic Review*, 105(2):860–85.

- Coricelli, G., Dolan, R. J., and Sirigu, A. (2007). Brain, emotion and decision making: the paradigmatic example of regret. *Trends in Cognitive Sciences*, 11(6):258–265.
- Costantini, A. F. and Hoving, K. L. (1973). The effectiveness of reward and punishment contingencies on response inhibition. *Journal of Experimental Child Psychology*, 16(3):484–494.
- Cottle, T. J. and Klineberg, S. L. (1974). *The present of things future: Explorations of time in human experience*. Free Press New York.
- Cournot, A. (1960). Researches into the mathematical principles of the theory of wealth, english edition of cournot (1838), translated by nathaniel t.
- Creswell, J. D. (2017). Mindfulness interventions. *Annual Review of Psychology*, 68(1):491–516.
- Da, Z., Engelberg, J., and Gao, P. (2011). In search of attention. *The Journal of Finance*, 66(5):1461–1499.
- Damasio, A. (1994). Descartes'error: Emotion, reason, and the human brain (gp putnam's sons).
- Davidson, R. J., Kabat-Zinn, J., Schumacher, J., Rosenkranz, M., Muller, D., Santorelli, S. F., Urbanowski, F., Harrington, A., Bonus, K., and Sheridan, J. F. (2003). Alterations in brain and immune function produced by mindfulness meditation. *Psychosomatic medicine*, 65(4):564–570.
- Dawson, E., Gilovich, T., and Regan, D. T. (2002). Motivated reasoning and performance on the was on selection task. *Personality and Social Psychology Bulletin*, 28(10):1379–1387.
- De Paola, M. and Gioia, F. (2017). Does patience matter in marriage stability? some evidence from italy. *Review of Economics of the Household*, 15(2):549–577.
- de Roure, C., Pelizzon, L., and Thakor, A. V. (2021). P2P lenders versus banks: Cream skinning or bottom fishing? SAFE Working Paper Series 206, Leibniz Institute for Financial Research SAFE.
- Dellavigna, S. and Pollet, J. M. (2009). Investor inattention and friday earnings announcements. *The Journal of Finance*, 64(2):709–749.

- Denes-Raj, V. and Epstein, S. (1994). Conflict between intuitive and rational processing: When people behave against their better judgment. *Journal of Personality and Social Psychology*, 66(5):819–829.
- Denrell, J. (2005). Why most people disapprove of me: experience sampling in impression formation. *Psychological review*, 112(4):951–978.
- Denrell, J. (2007). Adaptive learning and risk taking. *Psychological review*, 114(1):177.
- Dimidjian, S. and Segal, Z. V. (2015). Prospects for a clinical science of mindfulness-based intervention. *The American psychologist*, 70(2):593–620.
- Dittmar, A. and Duchin, R. (2016). Looking in the rearview mirror: The effect of managers' professional experience on corporate financial policy. *Review of Financial Studies*, 29(3):565–602.
- Dodini, S., Lopez-Fernandini, A. A., Merry, E. A., and Thomas, L. (2016). Consumers and Mobile Financial Services 2016. Reports and Studies 1777, Board of Governors of the Federal Reserve System (U.S.).
- Dunegan, K. J. (1993). Framing, cognitive modes, and image theory: Toward an understanding of a glass half full. *Journal of Applied Psychology*, 78(3):491–503.
- Durston, S., Thomas, K. M., Yang, Y., Uluğ, A. M., Zimmerman, R. D., and Casey, B. (2002). A neural basis for the development of inhibitory control. *Developmental science*, 5(4):F9–F16.
- Ehrlich, D., Guttman, I., Schönbach, P., and Mills, J. (1957). Postdecision exposure to relevant information. *The journal of abnormal and social psychology*, 54(1):98–102.
- Erev, I. and Haruvy, E. (2015). *Learning and the Economics of Small Decisions*, pages 638–702. Princeton University Press.
- Erev, I. and Roth, A. (1998). Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *American Economic Review*, 88(4):848–881.
- Erev, I. and Roth, A. E. (2014). Maximization, learning, and economic behavior. *Proceedings of the National Academy of Sciences*, 111(3):10818.

- Farb, N. A. S., Segal, Z. V., Mayberg, H., Bean, J., McKeon, D., Fatima, Z., and Anderson, A. K. (2007). Attending to the present: mindfulness meditation reveals distinct neural modes of self-reference. *Social cognitive and affective neuroscience*, 2(4):313–322.
- Fiedler, K. (2000). Beware of samples! a cognitive-ecological sampling approach to judgment biases. *Psychological review*, 107(4):659–676.
- Fjorback, L. O., Arendt, M., Ornbøl, E., Fink, P., and Walach, H. (2011). Mindfulness-based stress reduction and mindfulness-based cognitive therapy: a systematic review of randomized controlled trials. *Acta psychiatrica Scandinavica*, 124(2):102–119.
- Foster, D. (2016). Is mindfulness making us ill. <https://www.theguardian.com/lifeandstyle/2016/jan/23/is-mindfulness-making-us-ill>, note = Accessed: 27th May 2020.
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: A critical review. *Journal of Economic Literature*, 40(2):351–401.
- Frey, D. and Stahlberg, D. (1986). Selection of information after receiving more or less reliable self-threatening information. *Personality and Social Psychology Bulletin*, 12(4):434–441.
- Froot, K. (2001). The market for catastrophe risk: a clinical examination. *Journal of Financial Economics*, 60(2-3, SI):529–571.
- Fudenberg, D., Levine, D., Drew, F., K, L., Levine, D., Levine, D., and of Technology, M. I. (1998). *The Theory of Learning in Games*. MIT Press.
- Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2021). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance*, n/a(n/a).
- Fuster, A., Plosser, M., Schnabl, P., and Vickery, J. (2019). The Role of Technology in Mortgage Lending. *The Review of Financial Studies*, 32:1854–1899.
- Gabaix, X., Laibson, D., Moloche, G., and Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review*, 96(4):1043–1068.

- Garmaise, M. J. and Moskowitz, T. J. (2006). Bank mergers and crime: The real and social effects of credit market competition. *The Journal of Finance*, 61:495–538.
- Gennaioli, N. and Shleifer, A. (2010). What comes to mind. *Quarterly Journal of Economics*, 125(4):1399–1433.
- Golman, R., Hagmann, D., and Loewenstein, G. (2017). Information avoidance. *Journal of Economic Literature*, 55(1):96–135.
- Good, D. J., Lyddy, C. J., Glomb, T. M., Bono, J. E., Brown, K. W., Duffy, M. K., Baer, R. A., Brewer, J. A., and Lazar, S. W. (2015). Contemplating mindfulness at work: An integrative review. *Journal of Management*, 42(1):114–142.
- Griffin, D. and Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology*, 24(3):411–435.
- Grosskopf, B., Erev, I., and Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, 34(2):285–302.
- Guiso, L., Sapienza, P., and Zingales, L. (2018). Time varying risk aversion. *Journal of Financial Economics*, 128(3):403–421.
- Gureckis, T. M. and Markant, D. B. (2012). Self-directed learning: A cognitive and computational perspective. *Perspectives on Psychological Science*, 7(5):464–481.
- Hafenbrack, A. C., Kinias, Z., and Barsade, S. G. (2014). Debiasing the mind through meditation: mindfulness and the sunk-cost bias. *Psychological science*, 25(2):369–376.
- Hanh, T. N. (2010). *Peace is every step: The path of mindfulness in everyday life*. Random House.
- Hartzmark, S. M. (2014). The worst, the best, ignoring all the rest: The rank effect and trading behavior. *The Review of Financial Studies*, 28(4):1024–1059.
- Hartzmark, S. M., Hirshman, S., and Imas, A. (2019). Ownership, learning, and beliefs. (3465246).

- Häusler, A. N., Kuhnen, C. M., Rudolf, S., and Weber, B. (2018). Preferences and beliefs about financial risk taking mediate the association between anterior insula activation and self-reported real-life stock trading. *Scientific Reports*, 8(1):11207.
- Hendrickson, K. L. and Rasmussen, E. B. (2013). Effects of mindful eating training on delay and probability discounting for food and money in obese and healthy-weight individuals. *Behaviour research and therapy*, 51(7):399–409.
- Hendrickson, K. L. and Rasmussen, E. B. (2017). Mindful eating reduces impulsive food choice in adolescents and adults. *Health psychology : official journal of the Division of Health Psychology, American Psychological Association*, 36(3):226–235.
- Herrnstein, R. J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *Journal of the experimental analysis of behavior*, 4(3):267–272.
- Hirshleifer, D. and Teoh, S. H. (2003). Limited attention, information disclosure, and financial reporting. *Journal of Accounting and Economics*, 36(1):337 – 386.
- Hofmann, S. G., Sawyer, A. T., Witt, A. A., and Oh, D. (2010). The effect of mindfulness-based therapy on anxiety and depression: A meta-analytic review. *Journal of consulting and clinical psychology*, 78(2):169–183.
- Imas, A. (2016). The realization effect: Risk-taking after realized versus paper losses. *American Economic Review*, 106(8):2086–2109.
- Ingersoll, J. E. and Jin, L. J. (2013). Realization utility with reference-dependent preferences. *The Review of Financial Studies*, 26(3):723–767.
- Inman, J. J., Dyer, J. S., and Jia, J. (1997). A generalized utility model of disappointment and regret effects on post-choice valuation. *Marketing Science*, 16(2):97–111.
- Irani, R. M., Iyer, R., Meisenzahl, R. R., and Peydró, J.-L. (2020). The rise of shadow banking: Evidence from capital regulation. *The Review of Financial Studies*, 34(5):2181–2235.
- Jagtiani, J., Lambie-Hanson, L., and Lambie-Hanson, T. (2019). Fintech lending and mortgage credit access.

- Johnson, D. C., Thom, N. J., Stanley, E. A., Haase, L., Simmons, A. N., Shih, P.-A. B., Thompson, W. K., Potterat, E. G., Minor, T. R., and Paulus, M. P. (2014). Modifying resilience mechanisms in at-risk individuals: a controlled study of mindfulness training in marines preparing for deployment. *The American journal of psychiatry*, 171(8):844–853.
- Kabat-Zinn, J. (1990). Full catastrophe living: The program of the stress reduction clinic at the university of massachusetts medical center.
- Karlsson, N., Loewenstein, G., and Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty*, 38(2):95–115.
- Keehner, M., Hegarty, M., Cohen, C., Khooshabeh, P., and Montello, D. R. (2008). Spatial reasoning with external visualizations: What matters is what you see, not whether you interact. *Cognitive Science*, 32(7):1099–1132.
- Kiken, L. G. and Shook, N. J. (2011). Looking up: Mindfulness increases positive judgments and reduces negativity bias. *Social Psychological and Personality Science*, 2(4):425–431.
- Kluger, B. D. and Wyatt, S. B. (2004). Are judgment errors reflected in market prices and allocations? experimental evidence based on the monty hall problem. *The Journal of Finance*, 59(3):969–997.
- Kogan, S. (2008). Distinguishing the Effect of Overconfidence from Rational Best-Response on Information Aggregation. *The Review of Financial Studies*, 22(5):1889–1914.
- Komlos, J., Smith, P. K., and Bogin, B. (2004). Obesity and the rate of time preference: is there a connection? *Journal of biosocial science*, 36(2):209–219.
- Kőszegi, B. (2003). Health anxiety and patient behavior. *Journal of health economics*, 22(6):1073–1084.
- Kőszegi, B. (2010). Utility from anticipation and personal equilibrium. *Economic Theory*, 44(3):415–444.
- Kuhnen, C. M. (2015). Asymmetric learning from financial information. *The Journal of Finance*, 70(5):2029–2062.

- Kuhnen, C. M., Rudolf, S., and Weber, B. (2017). The effect of prior choices on expectations and subsequent portfolio decisions. Working Paper 23438, National Bureau of Economic Research.
- Laibson, D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112(2):443–477.
- Le Mens, G. and Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological review*, 118(2):379–392.
- Le Mens, G., Hannan, M. T., and Pólos, L. (2011). Founding conditions, learning, and organizational life chances: Age dependence revisited. *Administrative Science Quarterly*, 56(1):95–126.
- Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making*, 23(1):100–116 0894–3257.
- Lejarraga, T. and Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, 116(2):286–295.
- Levinthal, D. A. and March, J. G. (1993). The myopia of learning. *Strategic Management Journal*, 14(S2):95–112.
- Levitt, B. and March, J. G. (1988). Organizational learning. *Annual Review of Sociology*, 14(1):319–338.
- Li, C.-s. R. and Sinha, R. (2008). Inhibitory control and emotional stress regulation: neuroimaging evidence for frontal-limbic dysfunction in psycho-stimulant addiction. *Neuroscience and biobehavioral reviews*, 32:581–597.
- Liu, E. M. and Zuo, S. X. (2019). Measuring the impact of interaction between children of a matrilineal and a patriarchal culture on gender differences in risk aversion. *Proceedings of the National Academy of Sciences*, 116(14):6713–6719.
- Loomes, G. and Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, 92(368):805–824.

- Malmendier, U. and Nagel, S. (2011). Depression babies: do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics*, 126(1):373–416.
- March, J. G. (1996). Continuity and change in theories of organizational action. *Administrative Science Quarterly*, 41(2):278–287.
- Markant, D. and Gureckis, T. (2010). Category learning through active sampling. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Cognitive Science Society.
- Markant, D. B. (2018). Active transitive inference: When learner control facilitates integrative encoding. Working paper.
- Markant, D. B., Ruggeri, A., Gureckis, T. M., and Xu, F. (2016). Enhanced memory as a common effect of active learning. *Mind, Brain, and Education*, 10(3):142–152.
- Mars, T. S. and Abbey, H. (2010). Mindfulness meditation practise as a healthcare intervention: A systematic review. *International Journal of Osteopathic Medicine*, 13(2):56–66.
- Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., and Macrae, C. N. (2007). Wandering minds: the default network and stimulus-independent thought. *Science*, 315(5810):393–395.
- McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695):503–507.
- Mischel, W., Ayduk, O., Berman, M. G., Casey, B. J., Gotlib, I. H., Jonides, J., Kross, E., Teslovich, T., Wilson, N. L., Zayas, V., and Shoda, Y. (2011). 'willpower' over the life span: decomposing self-regulation. *Social cognitive and affective neuroscience*, 6(2):252–256.
- Mischel, W. and Ebbesen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*, 16(2):329–337.
- Morgan, D. P., Pinkovskiy, M. L., and Yang, B. (2016). Banking Deserts, Branch Closings, and Soft Information. Liberty Street Economics 20160307, Federal Reserve Bank of New York.
- Morrison, K. L., Madden, G. J., Odum, A. L., Friedel, J. E., and Twohig, M. P. (2014). Altering impulsive decision making with an acceptance-based procedure. *Behavior therapy*, 45(5):630–639.

- Necker, S. and Ziegelmeyer, M. (2016). Household risk taking after the financial crisis. *The Quarterly Review of Economics and Finance*, 59(C):141–160.
- Nguyen, H.-L. Q. (2019). Are credit markets still local? evidence from bank branch closings. *American Economic Journal: Applied Economics*, 11:1–32.
- Nyarko, Y. and Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3):971–1005.
- O’Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., and Dolan, R. J. (2004). Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science*, 304(5669):452.
- O’Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.
- Oster, E., Shoulson, I., and Dorsey, E. (2013). Optimal expectations and limited medical testing: evidence from huntington disease. *American Economic Review*, 103(2):804–30.
- Payzan-LeNestour, E. (2018). Can people learn about ‘black swans’? experimental evidence. *The Review of Financial Studies*, 31(12):4815–4862.
- Payzan-LeNestour, E. and Bossaerts, P. (2014). Learning about unstable, publicly unobservable payoffs. *The Review of Financial Studies*, 28(7):1874–1913.
- Plonsky, O. and Teodorescu, K. (2020). The influence of biased exposure to forgone outcomes. *Journal of Behavioral Decision Making*, 33(3):393–407.
- Rangel, A., Camerer, C., and Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9(7):545–556.
- Ritov, I. and Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behavior and Human Decision Processes*, 64(2):119–127.
- Samuelson, M., Carmody, J., Kabat-Zinn, J., and Bratt, M. A. (2007). Mindfulness-based stress reduction in massachusetts correctional facilities. *The Prison Journal*, 87(2):254–268.
- Sauer, S., Walach, H., Schmidt, S., Hinterberger, T., Lynch, S., Büssing, A., and Kohls, N. (2013). Assessment of mindfulness: Review on state of the art. *Mindfulness*, 4(1):3–17.

- Scholten, H., Scheres, A., de Water, E., Graf, U., Granic, I., and Luijten, M. (2019). Behavioral trainings and manipulations to reduce delay discounting: A systematic review. *Psychonomic bulletin & review*, 26(6):1803–1849.
- Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2x2-games. *American Economic Review*, 98(3):938–66.
- Settles, B. (2009). Active learning literature survey. Technical report.
- Shafir, S., Reich, T., Tsur, E., Erev, I., and Lotem, A. (2008). Perceptual accuracy and conflicting effects of certainty on risk-taking behaviour. *Nature*, 453(7197):917–920.
- Shigeoka, H. (2019). Long-term consequences of growing up in a recession on risk preferences. Working Paper 26352.
- Shonin, D. E., Van Gordon, W., and Griffiths, M. (2015). Mindfulness in psychology: A breath of fresh air? *The Psychologist*, 28:28–31.
- Sibinga, E. M. S., Webb, L., Ghazarian, S. R., and Ellen, J. M. (2016). School-based mindfulness instruction: An rct. *Pediatrics*, 137(1).
- Sicherman, N., Loewenstein, G., Seppi, D. J., and Utkus, S. P. (2015). Financial Attention. *The Review of Financial Studies*, 29(4):863–897.
- Stango, V. and Zinman, J. (2014). Limited and varying consumer attention: Evidence from shocks to the salience of bank overdraft fees. *The Review of Financial Studies*, 27(4):990–1030.
- Stevens, J. R., Hallinan, E. V., and Hauser, M. D. (2005). The ecology and evolution of patience in two new world monkeys. *Biol Lett*, 1:223–226.
- Sutton, R. S. and Barto, A. G. (2018). *Reinforcement learning: An introduction*. MIT Press.
- Tang, Y.-Y., Ma, Y., Fan, Y., Feng, H., Wang, J., Feng, S., Lu, Q., Hu, B., Lin, Y., Li, J., Zhang, Y., Wang, Y., Zhou, L., and Fan, M. (2009). Central and autonomic nervous system interaction is altered by short-term meditation. *Proceedings of the National Academy of Sciences of the United States of America*, 106(22):8865–8870.
- Tang, Y.-Y., Ma, Y., Wang, J., Fan, Y., Feng, S., Lu, Q., Yu, Q., Sui, D., Rothbart, M. K., Fan, M., and Posner, M. I. (2007). Short-term meditation training improves attention and self-regulation. *Proceedings of the National Academy of Sciences*, 104(43):17152–17156.

- Taylor, K. (1997). A regret theory approach to assessing consumer satisfaction. *Marketing Letters*, 8(2):229–238.
- Thorndike, E. L. (1927). The law of effect. *The American Journal of Psychology*, 39(1/4):212–222.
- Thornton, R. L. (2008). The demand for, and impact of, learning hiv status. *American Economic Review*, 98(5):1829–1863.
- Van Nieuwerburgh, S. and Veldkamp, L. (2010). Information Acquisition and Under-Diversification. *The Review of Economic Studies*, 77(2):779–805.
- Waggoner, A. S., Smith, E. R., and Collins, E. C. (2009). Person perception by active versus passive perceivers. *Journal of Experimental Social Psychology*, 45(4):1028–1031.
- Walach, H., Ferrari, M.-L. G., Sauer, S., and Kohls, N. (2012). Mind-body practices in integrative medicine. *Religions*, 3(1):50–81.
- Wang, L., He, C. Z., Yu, Y. M., Qiu, X. H., Yang, X. X., Qiao, Z. X., Sui, H., Zhu, X. Z., and Yang, Y. J. (2014). Associations between impulsivity, aggression, and suicide in chinese college students. *BMC Public Health*, 14:551.
- Witek-Janusek, L., Albuquerque, K., Chroniak, K. R., Chroniak, C., Durazo-Arvizu, R., and Mathews, H. L. (2008). Effect of mindfulness based stress reduction on immune function, quality of life and coping in women newly diagnosed with early stage breast cancer. *Brain, behavior, and immunity*, 22(6):969–981.
- Witte, K. (1996). *Fear as motivator, fear as inhibitor: Using the extended parallel process model to explain fear appeal successes and failures*, pages 423–450. Elsevier.
- Yao, Y.-W., Chen, P.-R., Li, C.-s. R., Hare, T. A., Li, S., Zhang, J.-T., Liu, L., Ma, S.-S., and Fang, X.-Y. (2017). Combined reality therapy and mindfulness meditation decrease intertemporal decisional impulsivity in young adults with internet gaming disorder. *Computers in Human Behavior*, 68:210–216.
- Yechiam, E. and Hochman, G. (2013). Loss-aversion or loss-attention: The impact of losses on cognitive performance. *Cognitive Psychology*, 66(2):212–231.

- Yechiam, E. and Hochman, G. (2014). Loss attention in a dual-task setting. *Psychological Science*, 25(2):494–502.
- Yechiam, E. and Rakow, T. (2012). The effect of foregone outcomes on choices from experience: An individual-level modeling analysis. *Experimental Psychology*, 59(2):55–67.
- Yechiam, E., Stout, J. C., Busemeyer, J. R., Rock, S. L., and Finn, P. R. (2005). Individual differences in the response to forgone payoffs: an examination of high functioning drug abusers. *Journal of Behavioral Decision Making*, 18(2):97–110.
- Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, 12(2):93–106.